Sunday, July 20, 8:00 AM - 9:30 AM

SA01

Enterprise

Recent Developments in Adaptive Experimentation for Bandit Model and Simulation Optimization

Invited Session

Applied Probability Society

Chair: Wei You, The Hong Kong University of Science and Technology, Hong Kong, Hong Kong

1 - Cost-Awared Exploration for Structured Bandits

Wei You, The Hong Kong University of Science and Technology, Hong Kong, Hong Kong, Chao Qin

We study a general framework for cost-aware structured bandits that integrates exploration efficiency and measurement costs into a unified decision-making model. In our approach, we reinterpret the complexity of cost-aware exploration as a two-player zero-sum game: the max-player selects a sampling distribution over arms while the min-player chooses a mixed strategy over alternative hypotheses. The Nash equilibrium of this game characterizes both the attainable cost and the structure of universally efficient policies. Unlike unstructured bandit problems, the equilibrium strategies in our setting are generally non-unique, which poses significant analytical challenges. To overcome these, we develop a game-perspective algorithm whose continuous-time dynamics are analyzed using replicator dynamics and Lyapunov methods. We derive sufficient conditions under which the algorithm converges to the global optimal value without requiring the optimal solution to be unique.

2 - Learning to Ask the Right Questions: a Multi-Armed Bandits Approach

Vikas Deep, National University of Singapore, Singapore, Singapore, Achal Bassamboo, Sandeep Juneja, Assaf Zeevi

Adaptive testing, where the hardness of next question depends on the response of the candidate on prior questions, is used in a variety of settings. In this paper, we aim to develop characteristics of an optimal adaptive test that correctly classifies a candidate's ability into one of several categories or grades. The candidate's ability is considered an unknown factor which, combined with hardness of the question, determines likelihood of answering the question correctly. The learning algorithm only observes whether the candidate answers a given question correctly or not. Specifically, the adaptive test should arrive at the correct grade for a given candidate with the fewest number of questions asked, while guaranteeing that the probability of error is less than a pre-specified and small δ. We derive a lower bound on the expected number of questions asked for any sequential adaptive testing strategy, which is a solution to a min-max optimization problem. We develop geometrical insights into this optimization problem structure and its dual formulation. In addition, we propose an algorithm that matches these lower bounds for small δ. Our key conclusions characterize the optimal adaptive test as having the following property: asymptotically, under certain structural assumptions, the questions asked to a candidate are almost all at one (candidate ability-specific) difficulty level. In general, however, the questions asked are at most two (candidate ability-specific) difficulty levels. Finally, we also propose a related algorithm based on Gaussian approximation that performs well numerically and admits suitable δ-correct performance guarantees in an asymptotic regime.

3 - Identifying All E-Best Arms in Linear Bandits with Misspecification

Zhekai Li, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Tianyi Ma, Cheng Hua, Ruihao Zhu

Motivated by the need to efficiently identify multiple candidates in high trial-and-error cost tasks such as drug discovery, we propose a near-optimal adaptive algorithm to identify all ε-best arms (i.e., at most ε worse

compared to the optimum). Specifically, we introduce LinFACT, an algorithm designed to optimize the identification of all ε-best arms in linear bandits. We establish a novel information-theoretic lower bound on the sample complexity of this problem and show that LinFACT achieves instance optimality by matching this lower bound up to a logarithmic factor. A fundamental piece of our proof directly integrates the lower bound into the scaling process for upper bound derivation and defines the round where the classification terminates to determine sample complexity. Additionally, we extend the analysis to provide upper bounds for LinFACT under model misspecification and in the context of generalized linear models. Numerical results show the practical advantages of LinFACT over several baseline methods. Experiments with synthetic data and real data from drug discovery demonstrate that our algorithm can identify more promising candidates, reduce sample complexity, and improve computational efficiency, showing its ability to accelerate early-stage exploratory experiments.

4 - Optimizing Input Data Collection for Ranking and Selection

Taeho Kim, HKUST, Hong Kong, Hong Kong, Eunhye Song

This work considers a ranking and selection under input uncertainty when all solutions share a common input model. Given the Bayesian input model, we aim to propose a sequential sampling procedure to find the best solution as efficiently as possible by balancing two types of sampling decisions: input data collection to reduce the input uncertainty and simulation to learn each solution's performance. The most probable best (MPB) is adopted as an estimator of the best solution, and its asymptotic properties have been investigated. Building upon these results, we devise a sequential sampling procedure. Furthermore, an extension to a continuous-valued input parameter case is discussed by employing the kernel ridge regression. Numerical studies are provided to support our findings.

SA02

Mercury

Air Traffic Management

Invited Session

Aviation Applications

Chair: Wayne Ng, Singapore University of Technology and Design, Singapore, Singapore

1 - Optimizing Aircraft Flows and Trajectories in the Terminal Airspace with Environmental Considerations

Wayne Ng, Singapore University of Technology and Design, Singapore, Singapore

We propose an integrated approach to optimize both Aircraft Flows in the Terminal Airspace, as well as Aircraft Trajectories given the output of the flow model. We aim to minimize flight delays and fuel consumption. Our solution comprises of an efficient reformulation of our previous work that optimizes aircraft flows, using a Linear Programming Model, to additionally account for environmental considerations. We validate with real-world data from the congested terminal airspace of Changi Airport in Singapore and show a significant improvement in fuel consumption at the same level of delays. For the trajectory optimization module, we utilize an Optimal Control model and propose additional constraints for integration with the flow model.

2 - Optimizing Aircraft Trajectories in Terminal Airspace Under Convective Weather Conditions

Jetanat Datephanyawat, Singapore University of Technology and Design (SUTD), Singapore, Singapore, Nuno Ribeiro

Efficient Air Traffic Management (ATM) is critical to ensuring safety, reducing delays, minimizing fuel consumption, and lowering the environmental impact of aviation. Weather—especially convective events—is the primary disruptor of air traffic operations, often forcing aircraft to deviate from planned trajectories and complicating flow management within Terminal Maneuvering Areas (TMAs).

This research proposes an integrated modeling framework that combines real-time geostationary satellite data with predictive and optimization models to improve TMA operations under uncertain weather. The objectives are: (i) to develop probabilistic models for short-term convective weather prediction, and (ii) to optimize aircraft trajectories and TMA flow through speed control and holding strategies.

By leveraging machine learning and stochastic optimization, this work enhances the adaptability and resilience of ATM systems. The approach is particularly relevant to convective-weather-prone regions like Singapore, where preliminary findings show a 45% increase in airborne delays during such events. This study contributes novel tools for real-time, data-driven air traffic optimization and advances the integration of predictive analytics in operations research for aviation.

3 - Deep-Learning Driven Contrail Detection for the Future Sustainable Traffic Flow Management

Kam Hung Ng, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, Rong Tang

Contrails, formed in ice-supersaturated regions, are a major aviation-induced climate forcer, necessitating a deep understanding of their evolution for targeted flight rerouting to mitigate potential climate impact. Recent advancements in satellite imagery and deep learning offer promising contrail detection, tracking, and lifespan estimation tools, supported by human-labelled datasets from GOES-16 and Landsat-8 satellites and semantic segmentation algorithms. However, existing studies towards contrail detection lack systematic comparisons of model architectures, loss functions, and evaluation metrics. This work conducts a series of experiments to determine the valid enhancement tricks and losses for deep learning-based contrail detection and introduces a novel spatial-frequency synergy learning paradigm, leveraging the Fourier transform to regularize contrail detection in both learning and matching stages. By integrating frequency-domain features into existing encoder-decoder architectures, we improve contrail detection performance while maintaining computational efficiency. Meanwhile, the presented benchmarks on GOES-16 and Landsat-8 datasets provide insights for future research in contrail detection and tracking.

4 - A Two Stage Model for Predicting Air Traffic Network Flow Control Strategies

YANJUN WANG, Nanjing University of Aeronautics and Astronautics, Nanjing, China, People's Republic of, Nianxi Xie

Miles-In-Trail (MIT) is a widely used Traffic Management Initiative (TMI) to manage the imbalance between airspace demand and capacity, alleviating congestion. However, human expert-driven MIT decision-making faces challenges in complex environments, particularly in multi-airport regions. To address this, we propose a two-stage MIT issuance prediction model. In the first stage, three scenarios are explored to analyze the mechanisms behind single- and multi-airport waypoint selection in depth. This model predicts the waypoint combinations requiring MIT issuance, which then serve as input for the second stage. In the second stage, three models are compared to refine MIT predictions for each waypoint under specific constraints, optimizing flow control strategies. The two-stage model's performance is also compared to a one-stage direct prediction using the LSTM model. An empirical analysis of two major airports in the Guangdong-Hong Kong-Macau Greater Bay Area demonstrates that the two-stage model significantly enhances prediction performance. Specifically, the average improvements in $\text{Precision}{\text{macro}}$, $\text{Recall}{\text{macro}}$, and F1_{macro} for key waypoints were 6.455\%, 10.245\%, and 10.254\%, respectively. Based on these results, we identify key factors influencing MIT issuance and provide guidance for selecting MIT values, improving MIT implementation efficiency in multi-airport regions and offering a scientific foundation for air traffic flow management in complex airspace.

SA03

Minto

Bridging Artificial Intelligence, Stochastic Systems, and the Economics of Data

Invited Session

Applied Probability Society

Chair: Yunbei Xu, National University of Singapore, Singapore, Singapore

1 - Regret Minimization in Heavy-Tailed Bandits

Shubhada Agrawal, Indian Institute of Science (IISc), Bangalore, Bangalore, India

In this talk, we will consider the classical regret minimization problem in the multi-armed bandit (MAB) setting with minimal assumptions about the arm distributions. We will discuss an optimal algorithm for this problem and

the ideas involved in its design and analysis. We will also look at a novel mean estimator for heavy-tailed distributions, which may be of independent interest to researchers and practitioners.

2 - Proper Dataset Valuation by Pointwise Mutual Information

Shuran Zheng, Tsinghua University, Beijing, China, People's Republic of

Data plays a central role in the development of modern artificial intelligence, with high-quality data emerging as a key driver of model performance. This has prompted the development of various data curation methods in recent years. However, measuring the effectiveness of these data curation techniques remains a major challenge. Traditional evaluation methods, which assess a trained model's performance on specific benchmarks, risk promoting practices that merely make the data more similar to the test data. This issue exemplifies Goodhart’s law: when a measure becomes a target, it ceases to be a good measure.

To address this, we propose an information-theoretic framework for evaluating data curation methods, where dataset quality is measured by its informativeness about the true model parameters using the Blackwell ordering. We compare informativeness by the Shannon mutual information of the evaluated data and the test data, and we propose a novel method for estimating the mutual information of datasets by training Bayesian models on embedded data and computing the mutual information from the model’s parameter posteriors. Experiments on real-world data demonstrate that our mutual information-based evaluation assigns appropriately lower scores to data curation strategies that reduce dataset informativeness, while traditional test score-based evaluation methods may favor data curation strategies that overfit to the test set but compromise the training data's informativeness.

3 - Finite-TIME Information-Theoretic Bounds in Queueing Network

Yujie Liu, National University of Singapore, Singapore, Singapore, Vincent Y. F Tan, Yunbei Xu

We establish the first \emph{finite-time} information-theoretic lower bounds—and derive new policies that achieve them—for the total queue length in scheduling problems over single-hop stochastic processing networks with \emph{both} adversarial and stochastic arrivals (with arbitrary spatial correlations). Prior analyses of MaxWeight guarantee only stability and asymptotic optimality in heavy traffic; we prove that, at finite horizons, MaxWeight can incur strictly larger backlog by problem-dependent factors. Our main innovations are 1) a \emph{minimax framework} that pinpoints the precise problem parameters governing any policy's finite-time performance; 2) an information-theoretic lower bound on total queue length; 3) fundamental limitation of MaxWeight that it can never be optimal in finite time; and 4) a new scheduling rule that minimizes the \emph{full} Lyapunov drift—including its higher-order terms—thereby matching the lower bound under some specialized conditions, up to universal constants. These results reveal a fundamental limitation on "drift-only" methods and points the way toward principled, non-asymptotic optimality in queueing control.

4 - Assouad, Fano, and Le Cam with Interaction: A Unifying Lower Bound Framework and Characterization for Bandit Learnability

Yunbei Xu, National University of Singapore, Singapore, Singapore, Fan Chen, Dylan Foster, Yanjun Han, Jian Qian, Alexander Rakhlin

We develop a unifying framework for information-theoretic lower bound in statistical estimation and interactive decision making. Classical lower bound techniques -- such as Fano's method, Le Cam's method, and Assouad's lemma -- are central to the study of minimax risk in statistical estimation, yet are insufficient to provide tight lower bounds for interactive decision making algorithms that collect data interactively (e.g., algorithms for bandits and reinforcement learning). Recent work of Foster et al. (2021, 2023) provides minimax lower bounds for interactive decision making using seemingly different analysis techniques from the classical methods. These results -- which are proven using a complexity measure known as the Decision-Estimation Coefficient (DEC) -- capture difficulties unique to interactive learning, yet do not recover the tightest known lower bounds for passive estimation. We propose a unified view of these distinct methodologies through a new lower bound approach called interactive Fano method. As an application, we introduce a novel complexity measure, the Fractional Covering Number, which facilitates the new lower bounds for interactive decision making that extend the DEC methodology by incorporating the complexity of estimation. Using the fractional covering number, we (i) provide a unified characterization of learnability for any stochastic bandit problem, (ii) close the remaining gap between the upper and lower bounds in Foster et al. (2021, 2023) (up to polynomial factors) for any interactive decision making problem in which the underlying model class is convex.

SA04

Indiana

Emerging Topics in Behavioral Operations

Invited Session

Behavioral Operations Management

Chair: Eirini Spiliotopoulou, Tilburg University, Tilburg, Netherlands

1 - Predictably Irrational: Taking Advantage of Systematic Biases to Improve Human-Machine Collaboration

Varun Karamshetty, National University of Singapore, Singapore, Singapore, Christopher Chen, Nitish Jain

Modern-day firms often empower humans to improve machine recommendations by applying their private information in a discretionary manner, but humans are vulnerable to biases when applying their private information. The challenge lies in how to regulate human revisions to improve recommendation quality.

Previous work suggests category-wise regulation of revisions– permitting (or prohibiting) all revisions in categories where, on average, human revision is beneficial (harmful). We expand this literature by conceiving a revision-level regulation framework that advocates the use of each revision as per its predicted quality. We leverage the extensive evidence on human biases being systematic to predict a revision's quality using bias-related (BR) features.

We test the framework in two distinct contexts: demand forecasting and image labelling. For demand forecasting, we obtain a large proprietary dataset from a European retailer. In out-of-sample evaluation, the framework yields a 21% improvement in generating final recommendations compared to the retailer's approach of always using human revisions. Likewise, in our image labelling study, using gig-workers, the framework provides a 43% improvement. Furthermore, in line with the literature on systematic biases, we find that BR features substantially improve the accuracy of revisions' predicted quality compared to using only category-level predictors of private information (PI).

Our framework offers firms a practical way to improve human-machine collaboration. We showcase how firms can gain by adopting simple strategies to regulate the use of human revisions based on their predicted quality score.

2 - Human and Algorithm in Tandem: Insights from Automation in Sequential Decision-Making Tasks

Eirini Spiliotopoulou, Tilburg University, Tilburg, Netherlands, Seyyed Iman Moosavi, Jan Fransoo

Automation of decision-making has become increasingly prevalent in operations management due to its potential to enhance performance. While automation of a single decision-making task can improve within-task performance, it may also influence the behavior of human decision-makers responsible for other related tasks, potentially affecting overall process performance improvement. This study examines this phenomenon in a setting where a demand forecasting task directly informs an inventory planning task. Through a series of lab experiments, we compare a human forecaster-human planner baseline with scenarios where either demand forecasting or inventory planning is automated. Our results indicate that inventory planning automation does not improve overall process performance due to changes in forecasters' behavior: forecasters provide forecasts with reduced service level anchoring but increased pattern noise when sending them to an automated planner. The net effect of these behavioral shifts results in lower forecast accuracy. Conversely, demand forecasting automation improves overall process performance. Furthermore, unlike the commonly observed behavior of algorithm aversion in single decision-making tasks, planners discount forecasts from human forecasters more than those from algorithmic forecasters. This occurs even when both forecasts exhibit similar characteristics. Our results suggest that within-task performance improvements resulting from automation do not necessarily guarantee overall performance improvements.

3 - Impact of Returns Risk in a Buyback Contract: An Experimental Study

Wanshan Zhu, Renmin University, Beijing, China, People's Republic of, Ao Zhuo, Xiaobo Zhao, jinxing xie

Manufacturers may fail to fulfill their buyback commitments under buyback contracts, exposing retailers to buyback risk. We conduct experimental study to examine how the manufacturers and the retailers make decisions under buyback risk in cases of symmetric information and asymmetric information. We find that manufacturers exhibit significant mental accounting behavior, while retailers display probability-weighting

behavior in the symmetric case. In the asymmetric information case, we find that manufacturers use buyback prices to signal the buyback risk, but the retailers only partially interpret these signals. To better explain these observed behaviors, we propose a modified pooling belief model. Our findings challenge traditional signaling game predictions and reveal how buyback risk signals are processed by human decision makers. By uncovering these behavioral patterns, we offer new insights into asymmetric information game decision-making in practice.

4 - When Alarms Fall Silent: Human-Adaptive Maintenance Optimization

Melvin Drent, Tilburg University, Tilburg, Netherlands, Bas Van Oudenhoven, Collin Drent, Rob Basten, Philippe Van de Calseyde

Sensors that monitor equipment can sometimes generate false warnings. When this occurs frequently, service engineers may begin to ignore future alerts, resulting in increased downtime costs, a phenomenon known as the cry-wolf effect. We conduct an experiment to study how people respond to warnings under the assumption of full adherence. Based on these insights, we optimize the timing of warnings, anticipating that repeated false alarms may reduce adherence. A second experiment confirms that our improved timing leads to better performance.

SA05

Orchard A

Optimization, Learning, and Quantum Multi-Agent Systems

Invited Session

Quantum Computing and OR

Chair: Antonios Varvitsiotis, N/A

1 - Quantum Finance: Pricing of Classical and Quantum Assets

Bin Cheng, CQT, NUS, Singapore, Singapore, Patrick Rebentrost

Quantum computers have the potential to provide advantages for solving problems in finance. First, we briefly review the pricing of financial derivatives using quantum subroutines for Monte Carlo estimation. Second, we investigate new market scenarios (quantum markets) and asset classes that could arise from the emergence of quantum technologies for computation and communication. We define the notion of quantum assets and develop an extended definition of arbitrage. We provide a quantum version of the first fundamental theorem of asset pricing, with application to new types of financial derivatives.

2 - Learning Pure Quantum States (Almost) Without Regret

Josep Lumbreras Zarapico, Centre for Quantum Technologies, National University of Singapore, Singapore, Singapore

We initiate the study of quantum state tomography with minimal disturbance to the samples. Can we learn a precise description of a quantum state through sequential measurements of samples while at the same time making sure that the post-measurement state of the samples is only minimally perturbed? Defining regret as the cumulative disturbance of all samples, the challenge is to find a balance between the most informative sequence of measurements and measurements incurring minimal regret. Here we completely resolve this question for pure qubit states by exhibiting an optimal protocol that achieves maximal precision in the description while incurring a regret that grows only logarithmically with the number of samples.

3 - No-Regret Learning and Equilibrium Computation in Quantum Games

Ryann Sim, Singapore University of Technology and Design, Singapore, Singapore, Wayne Lin, Georgios Piliouras, Antonios Varvitsiotis

As quantum processors advance, there is potential to design large-scale decentralized systems involving interacting quantum-enabled agents. Recent research efforts have explored quantum versions of Nash and correlated equilibria as solution concepts of strategic quantum interactions, but these approaches did not directly connect to decentralized adaptive setups where agents possess limited information. We study the dynamics of quantum-enabled agents within decentralized systems that employ no-regret algorithms to update their behaviors over time. Specifically, we investigate two-player quantum zero-sum games and poly-matrix quantum zero-sum

games, showing that no-regret algorithms converge to separable quantum Nash equilibria in the time-average sense. In the case of general multi-player quantum games, our work leads to a novel solution concept, that of the separable quantum coarse correlated equilibria (QCCE), as the convergent outcome of no-regret algorithms, offering a natural solution concept for decentralized quantum systems. Finally, we show that computing QCCEs can be formulated as a semidefinite program and establish the existence of entangled (i.e., non-separable) QCCEs, which are unreachable via the current paradigm of no-regret learning.

4 - Quantum Volunteer's Dilemma

Dax Enshan Koh, Agency for Science, Technology and Research, Singapore, Singapore, Kaavya Kumar, Siong Thye Goh

The volunteer's dilemma is a well-known game in game theory that models the conflict players face when deciding whether to volunteer for a collective benefit, knowing that volunteering incurs a personal cost. In this work, we introduce a quantum variant of the classical volunteer's dilemma, generalizing it by allowing players to utilize quantum strategies. Employing the Eisert–Wilkens–Lewenstein quantization framework, we analyze a multiplayer quantum volunteer's dilemma scenario with an arbitrary number of players, where the cost of volunteering is shared equally among the volunteers. We derive analytical expressions for the players' expected payoffs and demonstrate the quantum game's advantage over the classical game. In particular, we prove that the quantum volunteer's dilemma possesses symmetric Nash equilibria with larger expected payoffs compared to the unique symmetric Nash equilibrium of the classical game, wherein players use mixed strategies. Furthermore, we show that the quantum Nash equilibria we identify are Pareto optimal. Our findings reveal distinct dynamics in volunteer's dilemma scenarios when players adhere to quantum rules, underscoring a strategic advantage of decision-making in quantum settings.

5 - Towards Quantum Intelligent Agents: on the Potential and Possibilities of Quantum-Enhanced Decision-Making

Mile Gu, Nanyang Technological University, Singapore, Singapore

The world is witnessing unprecedented automation — from self-driving vehicles and sophisticated large language models to automated trading algorithms and virtual agents. We are developing increasingly advanced AI systems to power autonomous agents operating in ever more complex environments. However, as making favourable decisions requires tracking vast and growing amounts of contextual data, the associated memory and energy costs are escalating at unsustainable rates. The search for more efficient mechanisms of reasoning has thus become imperative.

In this talk, I ask: could quantum intelligent agents — agents equipped with the ability to process information quantum mechanically — provide a pathway to mitigate these costs? I will introduce our recent research advances at the Nanyang Quantum Hub, where new discoveries offer compelling evidence that quantum-enhanced decision-making can reduce memory and energy consumption beyond what is achievable by any classical counterpart. Time permitting, I will also outline the broader implications for the future quantum ecosystem — and how such efficiency gains could accelerate the timeline toward useful quantum technologies.

SA06

Orchard B

Machine Learning and Contextual Decision-Making in Operations Management

Invited Session

Efficiency and Fairness in Supply Chain and Service Operations Management

Chair: Shaochong Lin, The University of Hong Kong

1 - Reducing Confounding Effects in Contextual Stochastic Optimization via Proxy Matrix

Yulong Huang, Zhejiang University, Hangzhou, China, People's Republic of, Shaochong Lin, Zuo-Jun Shen, Man-Chung Yue

Data-driven decision-making in real-world scenarios often faces the challenge of endogeneity between decisions and outcomes, introducing confounding effects. While existing literature typically assumes unconfoundedness,

this is often unrealistic. In practice, decision-making relies on high-dimensional, heterogeneous-type proxy features of confounders, leading to suboptimal decisions due to limited predictive power for uncertainty. We propose a novel semi-parametric decision framework to mitigate confounding effects. Our approach combines exponential family matrix completion to infer the confounders matrix from proxy features, with non-parametric prescriptive methods for decision-making based on the estimated confounders. We derive the true regret bound for data-driven decisions under confounding effects and demonstrate how our framework improves this bound. Experiments on both synthetic and real datasets validate our method's efficacy in reducing confounding effects across various proxy dimensions. We also show that our approach consistently outperforms benchmarks in practical applications.

2 - Data-Driven Promotion Optimization: Transfer Learning for Super-Quantile Treatment Effects

Jiayi Zeng, The University of Hong Kong, Hong Kong, Hong Kong, Shaochong Lin, Zuo-Jun Shen

For offline chain retailers, the firm faces the challenge of optimizing product promotions with sparse and heterogeneous data. The firm profits from two sources, revenue sharing of branch sales and wholesale to branches. This paper proposes a data-driven two-stage framework to address promotion selection and inventory decisions, accounting for the hierarchical interplay between the firm and branches. The firm first selects promoted products to maximize profit, considering both wholesale margins and shared revenue, while branches subsequently determine order quantities under demand uncertainty. To account for the heterogeneous nature of products, we focus on the heterogeneous treatment effects of promotion, which are decomposed into mean, quantile, and super-quantile terms. A transfer learning framework is proposed to overcome the data sparsity issue and enhance estimation of heterogeneous treatment effects, which leverages auxiliary datasets from related products. Probability bounds are derived for our transfer learning method, and we further provide theoretical guarantee for the regret of the promotion optimization problem. We examine our method on a real-world retailing dataset. Empirical validation shows the framework’s ability to improve promotion decisions under limited data, enabling the firm to enhance profitability across retail networks.

3 - Data-Driven Capital-Constrained Newsvendor Problems with Default Risks

Jindou LUO, City Univ of HK, Tat Chee Avenue, ON, Hong Kong, Frank Chen, Shaochong Lin

This study investigates the ordering decisions of a capital-constrained retailer who relies on external financing to optimize order quantities under working capital limitations. If the income from sales is insufficient to cover the debt, the retailer is assumed to default. The demand distribution of the focal product is unknown, and only limited historical sales data of itself and similar products are available. We apply machine learning techniques to assess covariate-based similarity between new and historical products, deriving sample-dependent weights for expected profit approximation. This approach formalizes the problem as a data-driven capital-constrained newsvendor problem, which is a non-convex optimization challenge. Proof of the asymptotic optimality of the data-driven solution is provided. Our findings reveal that, compared to the classical data-driven newsvendor problem, the capital-constrained newsvendor adopts a more risk-taking ordering strategy, transferring part of the demand risk to the lender through default mechanisms. Furthermore, we analyze how data influences the ordering decisions of capital-constrained sellers, supplementing the distribution-based capital-constrained newsvendor model with insights on leveraging data to make decisions. Our findings on retailers’ ordering decision patterns provide lenders with insights into assessing retailers’ risk behaviors. Retailers with high internal capital or valuable collateral are less risk-taking. Additionally, “poor” retailers, even those with zero collateral, can be of low-risk if they have relatively stable historical demand. Lenders should pay particular attention to retailers with extremely high sales outliers. These retailers may order much larger quantities than average historical sales, leading to a higher probability of default.

4 - Preference Learning for Dynamic Assortment and Inventory Planning

Yijie Zheng, The Polytechnic University of Hong Kong, Hong Kong, China, People's Republic of, Shaochong Lin, Youhua (Frank) Chen, Yi Yang, Zuo-Jun Max Shen

We address the problem of preference learning for joint assortment and inventory planning with stockout-based substitution using transaction data. In this context, the retailer determines the initial inventory levels for products at the beginning of the selling period, while customers arrive sequentially and choose from the available

inventory. Rather than assuming all demand information is known or imposing parametric models on the distributions of customer arrivals or choice structures, we learn these nonparametrically from historical transaction data, incorporating product feature information. In this paper, we propose a novel preference learning algorithm that effectively captures customer choices, thereby facilitating data-driven decisions for dynamic assortment and inventory management. Our method is based on a crafted kernel regression that adapts to the complex choice structures inherent in multi-product systems. By leveraging large deviation theory and partitioning estimates in statistical learning, we establish the universal consistency of our preference learning algorithm and derive its convergence rate. Building on these consistent estimates, we demonstrate the asymptotic optimality of our data-driven solution and show its regret. Furthermore, we incorporate approximation strategies to enhance practical efficiency. Extensive numerical studies using both synthetic and real data validate the promising efficacy of our data-driven framework for dynamic assortment and inventory planning.

SA07

Orchard C

FinTech Innovation, Investment, and Decision-Making

Invited Session

FinTech

Chair: Ruiting Zuo, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou

1 - Stock Price Informativeness and Inventory Management

Siguang Li, Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China, People's Republic of, Xiaohong Huang, Jian Ni

We investigate the relationship between stock price informativeness and inventory management in a novel setting with informational feedback from financial markets. Specifically, managers can learn information about customer demand from stock prices to effectively guide their inventory decisions, leading to a reduction in excess inventory holding of safety stock due to precautionary motives in fear of stock-out risks. Furthermore, we provide empirical evidence that supports the effect of managerial learning on inventory management. Our findings offer a novel perspective on how revelatory price efficiency can affect corporate decisions aside from investment.

2 - The Double-Edged Sword: Corporate Venture Capital and Startup Innovation in The Medical Device Industry

Moyan Li, Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China, People's Republic of

Medical device innovation is key to the U.S. healthcare industry. Numerous innovative medical devices originate with startup firms that rely on external funding through two primary routes, independent venture capital (IVC) and corporate venture capital (CVC). Whereas IVC primarily brings funds alone, CVC also brings industry knowledge and experience. Using difference-in-difference models and instrumental variable strategy, we study the effect of CVC on innovation outcomes of U.S. medical device startup firms. We find that CVCs foster more product approvals, but once launched, CVC-backed startups experience worse innovation product quality, indicating a double-edged sword to CVC funding. Post-hoc analyses reveal that this negative quality effect is accentuated when the CVC firm has no prior experience in the product market of the startup innovation, and when the CVC has their own recent experiences of poor quality.

3 - Individual Versus Institutional Philanthropy: Crowdfunding During Crises

Anqi Wu, Florida International University, Miami, FL, United States, Aravinda Garimella, Ramanath Subramanyam, Mehmet Ahsen

Global crises are accompanied by resource shortages, logistical and operational challenges, and shifts in demand for certain resources. During these crises, non-profit organizations have historically played a crucial role in mitigating these challenges, with crowdfunding platforms emerging as a new solution for resource mobilization. In this study, we examine how individual and institutional donors respond to the sudden shifts in demand that

accompany a global crisis. We analyze rich, granular data from multiple sources including donation data from a leading education crowdfunding platform before and after the onset of the COVID-19 pandemic and find notable contrasts between individual and institutional responses. While there was a significant increase in institutional donations during the pandemic, individual donations remained relatively stable. Adopting a contingency framework, we center donor heterogeneity in our approach, deriving further nuanced insights. Small and local institutions played a pivotal role in supporting high-need schools during the crisis. Donations from individuals in economically disadvantaged communities worst impacted by the crisis saw a decrease, but these communities prioritized immediate and time-sensitive needs. These findings make distinct contributions to the operations literature on crowdfunding. For platform designers, our findings have important implications on how to tailor their messaging to different donor groups during crises.

4 - Does Green Credit Promote Real Economic Development? Dual Empirical Evidence of Scale and Efficiency

Song Jiang, Chongqing University of Technology, Chongqing, China, People's Republic of

tba

SA08

Bras Basah A

Optimizing Sustainable Transportation

Contributed Session

Computing Society

Chair: Zhang XinYue, Northwestern Polytechnical University, Xi'an, N/A

1 - Centrality of Shortest Paths: Algorithms and Complexity Results

Dmytro Matsypura, The University of Sydney, Sydney, Australia, Johnson Phosavanh

The centrality of a node is often used to measure its importance to the structure of a network. Some centrality measures can be extended to measure the importance of a path. In this paper, we consider the problem of finding the most central shortest path. We show that the computational complexity of this problem depends on the measure of centrality used and, in the case of degree centrality, whether the network is weighted or not. We develop a polynomial algorithm for the most degree-central shortest path problem with the worst-case running time of $O(|E||V|^2\Delta(G))$, where $|V|$ is the number of vertices in the network, $|E|$ is the number of edges in the network, and $\Delta(G)$ is the maximum degree of the graph. In addition, we show that the same problem is NP-hard on a weighted graph. Furthermore, we show that the problem of finding the most betweenness-central shortest path is solvable in polynomial time, while finding the most closeness-central shortest path is NP-hard, regardless of whether the graph is weighted or not. We also develop an algorithm for finding the most betweenness-central shortest path with a running time of $O(|E|^2|V|^2)$ on both weighted and unweighted graphs. To conclude our paper, we conduct a numerical study of our algorithms on synthetic and real-world networks and compare our results to the existing literature.

2 - Can a Drone-Truck System End Heat Risks and Enhance Safety in Short-Term High-Density Events?

SHROUQ GAMAL, King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia, Saudi Arabia

Large-scale, short-term high-density events pose challenges in crowd management, thermal comfort, and emergency response, particularly under extreme weather conditions. This study introduces a Two-Phase optimization framework for a cooperative drone-truck system, pioneering spraying drones for thermal comfort and Lab-on-Drone technology for real-time environmental monitoring and intervention approaches never before integrated in mass gatherings.

The framework aligns with the three pillars of sustainability: (1) Environmental sustainability- minimizing energy consumption and carbon emissions through optimized drone-truck coordination; (2) Social sustainability- enhancing public safety and well-being by mitigating heat risks and providing rapid emergency response; and (3)

Economic sustainability- optimizing travel time and operational costs to ensure feasible large-scale implementation.

The framework operates in two phases: Phase I- Patrolling, where drones with Lab-on-Drone technology collect real-time air quality and weather data to identify high-risk zones, and Phase II- Thermal Comfort and Emergency Response, where drones and trucks coordinate targeted cooling interventions and medical support. Each phase is mathematically modeled, incorporating warm-start approaches for efficiency.

Scenario-based simulations optimize deployment strategies to reduce health risks from extreme weather and crowd congestion. The Hajj pilgrimage in Mecca, Saudi Arabia, serves as a case study, demonstrating the first-ever integration of spraying drones for thermal comfort. Future applications include the FIFA World Cup and the Asian Games. This research establishes a new standard for sustainable event management, leveraging cutting-edge drone technology for real-time, data-driven interventions.

3 - Optimization of Bus Schedule: An Integrated Approach to Mitigate the Bus Overlapping

Gajendra Nagar, Indian Institute of Technology Delhi, New Delhi, India, Hemant Kumar Suman, Lakshay Lakshay, Nomesh Bolia

For any developing country, public transport plays a crucial role in the smooth movement of commuters throughout the network. To attain this goal, an effective timetable is critical, and if it is not implemented efficiently, a bus-overlapping-like condition may emerge at various locations in the transit network. The term "bus overlapping" refers to the situation in which multiple buses are scheduled to serve a particular route simultaneously rather than at specific headways, which results in poor bus utilization, overcrowding, inconsistent headway, and passenger discomfort. Therefore, from the perspective of both transit agencies and customers, it is imperative that the timetable and bus schedule should be strategically optimized. However, this objective was challenging to accomplish in the majority of developing countries that relied on manual planning and the complexity of the process. So, to mitigate the bus overlapping and optimize the bus schedule, we developed a novel strategy based on optimization techniques. Initially, we develop an integrated optimization model for the simultaneous creation of timetables and the scheduling of buses in public transport systems. Later, to solve the model and generate the results for the large instances, we developed a heuristic algorithm. Finally, we test the developed methodology on a real network in India. The findings of this study provide valuable insight for the transit agencies that are aiming to optimize and enhance their transit network planning strategies. By implementing this, they can mitigate the bus overlapping, optimize the bus schedule, and provide reliable service.

4 - Assortment Optimization Under Cross-Nested Logit Model: Exact Approach via Cutting Plane and Branch-and-Cut

Hoang Giang Pham, Singapore Management University, Singapore, Singapore, Tien Mai

We investigate the assortment optimization problem under general linear constraints, where customer choice behavior is modeled using the cross-nested logit (CNL) framework - one of the most general and flexible choice models in the literature. In this setting, products are grouped into multiple overlapping subsets (nests), allowing each product to belong to more than one nest, thereby capturing complex substitution patterns. The objective is to determine an optimal assortment that maximizes expected revenue. To solve this problem efficiently in practice, we introduce a novel methodology based on outer approximation techniques that approximate key convex components of the objective function. This leads to the development of effective cutting-plane and branch-and-cut algorithms that enable exact solutions within reasonable computational effort.

5 - Flexible Process Planning: Polyhedral Analysis and a Branch-and-Cut Algorithm Based on Augmented or-Precedence Poset

Kaiping Luo, School of Economics and Management, Beihang University, Beijing, China, People's Republic of, Yongbo Xiao, Xiaoqiang Cai, Mabel C. Chou

Flexible Process Planning (FPP) is recognized as the most crucial problem to enable smart/intelligent manufacturing. The existing computer-aided process planning systems have, however, largely relied on meta-heuristic algorithms because of the vast complexity in deriving an optimal solution for the FPP problem. In this paper, we propose to tackle this problem from a distinct perspective, by developing a polyhedral approach based

on which nearly optimal solutions can be found in substantially faster time. We first formulate the FPP problem using least decision variables and then reformulate it in two different modes. To improve computational efficiency, we further propose a large number of strong valid inequalities by taking full advantage of the strict partial order of the OR-precedence relation between operations, and accordingly develop a branch-and-cut algorithm. Our experimental results show that our proposed algorithm performs extraordinarily well in searching for the optimal solution.

6 - Drone-Based Pickup and Delivery Problem with TIME Windows and Docking Capacity

Zhang XinYue, Northwestern Polytechnical University, Xi'an,Shannxi,China, China, People's Republic of

In this study, we introduce a novel urban delivery network that combines automated locker stations and drones to enable on-demand delivery services. These locker stations can store multiple packages and offer battery swapping services for drones; however, they have a limited docking capacity for drones at any given time. Packages for on-demand requests can be picked up from a station or another designated location and must be delivered to either a station or a non-station location within a specified time window. Drones are allowed to fly between and dock at different stations. In this work, we investigate the drone pickup and delivery problem with time windows and station docking capacity. To solve this complex optimization problem, we propose a two-phase matheuristic approach. The first phase utilizes an adaptive large neighborhood search (ALNS) to generate a diversified route pool. The second phase employs a logic-based benders decomposition (LBBD) heuristic, featuring a set partitioning master problem model and two constraint programming subproblem models. The master problem constructs delivery routes covering all tasks using routes from the route pool, while the subproblems adjust the solution to ensure it satisfies station docking capacity constraints. The algorithm has been tested on benchmark instances generated from real-world drone station coordinates in several cities, with results demonstrating its effectiveness.

7 - Screening Cut Generation for Sparse Ridge Regression

Haozhe Tan, National University of Singapore, Singapore, Singapore, Guanyi Wang

Sparse ridge regression is widely utilized in modern data analysis and machine learning. However, computing globally optimal solutions for sparse ridge regression is challenging, particularly when samples are arbitrarily given or generated under weak modeling assumptions. This paper proposes a novel cut-generating approach, Screening Cut Generation (SCG), to eliminate non-optimal solutions without sample generation assumptions. In contrast to recent safe variable screening approaches, SCG offers superior screening capability by identifying whether a specific $\{\pm 1\}$ combination of multiple features is part of an optimal solution. This identification is based on a convex relaxation solution rather than directly solving the original sparse ridge regression. The cuts generated by SCG can construct a safe outer approximation of the optimal solution set in the pre-processing step of branch-and-bound and its variants. Numerical experiments are reported to validate the theoretical findings and demonstrate the efficiency of SCG, particularly in synthetic instances with challenging parameter configurations.

SA09

Bras Basah B

Innovations and Challenges in Modern Education: Peer Learning, AI Integration, and Safety Strategies

Contributed Session

Education

Chair: Ying Yang, University of Washington Bothell, Bothell, WA, United States

1 - From Code to Confidence: How Students Built a Multi-Agent AI to Navigate Job Interviews

Graeme Warren, Johns Hopkins University, Rockville, MD, United States, Jacques Domenge, Xiangyang Li, Rahul Pandey, Miller Simmons, Minghong Xu

I will share insights from a case-based teaching initiative in a graduate course on Generative AI in which students designed and assessed a multi-agent AI tool. Course participants were asked to develop and

assess an agentic AI tool that can be used by students in an MS Analytics program to prepare for job interviews. The tool should ideally be a personalized, interactive system that helps students reflect on their experiences, clarify their value proposition, rehearse authentic storytelling, and increase their awareness of affective interview considerations. The case fosters technical fluency, design thinking, and professional self-reflection opportunities for case participants. Leveraging the experience of two life coaches, an industry professional, and several faculty members who designed and supported the case, I will discuss: a) how multi-agent AI systems can be used as a pedagogical framework for experiential learning, b) how generative AI can bridge the gap between technical expertise and professional communication, helping students articulate their value in ways that resonate with diverse interviewer personas, c) how students can apply human-AI interaction principles—like simulating behavioral cues, adapting to interviewer styles, and offering emotional intelligence coaching—within agent-based systems, d) the value of combining academic literature and practitioner insights to guide AI outputs (e.g., interview strategies, skill gaps), helping students ground their learning in research and practice, and e) how students addressed concerns like bias mitigation, data privacy, and transparency.

2 - Intersecting Dimensions: Advanced Analytical Approach to School Climate and Injury Prevention In Health Policy

Michael Khalemsky, Jerusalem Multidisciplinary College, Jerusalem, Israel, Anna Khalemsky, Eli Jaffe

Child safety in schools is a critical concern worldwide, with increasing attention on reducing injuries and ensuring safe environments. Despite this, school-related injuries remain prevalent, necessitating a comprehensive, interdisciplinary approach to understand their causes. This study aims to analyze the factors contributing to such injuries and assess the role of school climate in shaping student safety. Using advanced data analysis, the research develops targeted strategies to enhance safety in educational settings by constructing a multidimensional taxonomy.

The study analyzed data from 363 Israeli primary schools and 10,855 injury cases handled by Magen David Adom, Israel's primary emergency medical services provider. A two-level taxonomy was developed using clustering techniques to classify schools into climate "patterns", and to further identify sub-patterns based on injury characteristics.

Results revealed five school climate clusters, ranging from positive to negative. Each exhibited two distinct injury sub-clusters. In schools with a positive climate, boys often sustained head injuries during breaks, while girls experienced limb injuries from playing in corridors. In schools with a negative climate, injuries were differentiated by activity type - either from play or from falling from heights.

The findings highlight a complex relationship between school climate and injury patterns, emphasizing the need for analytical methods beyond traditional approaches. The taxonomy developed in this study offers a detailed understanding of the interplay between school environment and student safety, supporting evidence-based, gender-sensitive interventions. Integrating diverse data sources enables the design of more effective school health policies and provides insights for strategic decision-making in educational safety management.

3 - Generative AI Toolbox for Academics: Guidance Through Different Research Stages

Stefan Voß, University of Hamburg, Hamburg, Germany, Michael Palk

The integration of generative artificial intelligence (AI) into everyday life is becoming increasingly widespread, particularly in the field of conversational human-computer interaction systems. Recent advancements in chatbot technology have opened up numerous possibilities for enhancing the efficiency of academic research. Previously time-consuming tasks such as conducting literature reviews, summarizing scientific papers, extracting data, and performing quantitative analyses can potentially be handled more efficiently by AI-driven systems. These technologies have the potential to accelerate research processes, reduce the workload for researchers, and enhance access to knowledge. However, despite their advantages, these systems also present several challenges. Issues such as inconsistent responses, hallucinated information, and incorrect citations pose significant risks, particularly in academic settings where accuracy and reliability are crucial. The presence of these flaws raises concerns about academic integrity, as misinformation or incorrect attribution of sources could undermine the

credibility of research. This study contributes to the growing body of knowledge on conversational human-computer interaction systems by examining the role of generative artificial intelligence technologies in academic research. Both the benefits and limitations of these technologies, offering insights into their impact on research efficiency, reliability, and ethical considerations are explored. By critically analyzing their applications and potential risks, this study aims to provide a comprehensive understanding of how chatbots can be responsibly integrated into scholarly work during different research stages. The concepts are exemplified by means of (stochastic) vehicle routing problems.

4 - ***Impact of a Simulation Game & Providing Feedback on Enhancing Students&Rsquo; Knowledge***

Hossein Rikhtehgar Berenji, Pacific University, Forest Grove, OR, United States, Hossein Najmi

We use the revised Bloom's taxonomy framework to empirically measure students' knowledge gain in capacity and inventory management in a simulation game. Our results in the knowledge experiment show that participating in this simulation game improves students' conceptual and procedural knowledge. Our results illustrate that detailed feedback enhances students' performance relative to a control group, and that the effect is more pronounced among high-performing students. Together, these findings extend our understanding of the learning that results from participation in game-based pedagogy.

5 - Peer Feedback at Scale: Insights into Student Learning and Critical Thinking

Ying Yang, University of Washington Bothell, Bothell, WA, United States, Siyu Wang, Pradeep Chintagunta, Carine Marette

Critical thinking is a cornerstone of higher education, yet fostering this skill effectively and at scale remains a significant challenge. This study examines the role of structured peer feedback in promoting critical thinking and learning engagement using a unique dataset from a digital peer feedback platform that implements a three-stage learning model—creation, evaluation, and feedback—designed to enhance cognitive engagement through iterative peer-to-peer interaction. The dataset includes millions of observations across diverse higher education courses in North America. Students participated as both creators, producing work for evaluation, and evaluators, assessing peers' submissions using standardized rubrics. Evaluators were rated by creators on two dimensions—motivational and critical—yielding a feedback score. Analysis reveals a strong positive correlation between evaluation scores and feedback scores, suggesting that students who provide more motivational and critically constructive evaluations tend to be more effective at recognizing high-quality work. This reciprocity indicates deeper engagement with content and the peer learning process. The findings support the hypothesis that giving thoughtful feedback reinforces the evaluator's own understanding and critical thinking. Our results align with prior research on the cognitive benefits of peer assessment, including a meta-analysis in the British Journal of Educational Technology, which found that online peer assessment positively impacts higher-order thinking. This study expands on that work by providing large-scale empirical evidence from authentic classroom settings. Future analysis, including text mining of feedback content and robustness checks with additional variables, will further explore the mechanisms behind these relationships.

SA10

Hullet Room

Algorithmic Fairness for OR and AI

Invited Session

OR and Social Wellness

Chair: Zhu Wang, Shanghai Maritime University, Shanghai, N/A

1 - One-to-Many Stable Matching for Integrating Ridesharing and Public Transit On a Mobility-as-a-Service Platform

Yating Wu, Southeast University, Nanjing, China, People's Republic of, Minghui Lai, YUAN WANG

Motivated by the practices of SocialCar project, we study the integration of one-to-many peer-to-peer ridesharing and public transit for a Mobility-as-a-Service (MaaS) platform in the morning commute setting. Drivers and riders can be flexibly matched through either a first-mile or last-mile connection but not both simultaneously. However, MaaS is essentially a decentralized user-centric system, which is challenging to operate since both service providers and travelers are self-interested. We adopt the concept of two-sided stable

matching and formulate the decentralized planning problem as an integer program that typically includes an extremely large number of matches and complex stability constraints. The problem is further complicated by the possible nonexistence of a one-to-many stable matching. We propose two new algorithmic frameworks based on column-and-row-generation method to efficiently check the existence of stable matching or determine approximate stable matchings. We test our algorithms with a real-world case study and show that our algorithms can quickly determine stable matchings or approximate stable matchings with an almost negligible number of blocking matches and only minor stability deviations. Extensive computational experiments also reveal various new managerial insights into match selections, users' detours and utilities, stability deviations, public transit utilization, the effects of pricing schemes, and the significance of stability.

2 - Fair Truckload Freight Matching Algorithm for a Digital Logistics Platform

Yimiao Wang, Southeast University, Nanjing, China, People's Republic of, Minghui Lai

Order assignment represents a significant challenge for digital logistics platforms. Typically, carriers prefer bundles of orders scheduled over a planning horizon. In these scenarios, each order is unique and indivisible, and the complementarity between orders is crucial. In particular, bundles that minimize empty travel distances can enhance operational efficiency for carriers. The underlying mechanism design stresses equitable opportunities among carriers, ensuring that no carrier favors another’s assigned bundle over its own. Notably, the package assignment problem is NP-hard, and conventional price-based mechanisms that require carrier payments are often impractical. Building on recent advancements in market design theory, we propose a randomized matching mechanism that coordinates fair and near-efficient allocations without monetary transfers. Initially, we solve the relaxed optimal fractional allocation problem. Recognizing the limitations of the classical Birkhoff-von Neumann theorem in yielding integer solutions, we extend a geometry-based convex decomposition framework to derive probabilistic allocation schemes expressed in convex combination form. This innovative approach establishes a robust coordination paradigm. Finally, we conduct numerical experiments using operational data from an actual logistics network to assess both fairness and social welfare outcomes.

3 - Greedy Mechanism-Based Bi-Objective Optimization for Green Scheduling in Manufacturing Systems: Balancing Energy Consumption and Efficiency

zhu wang, Shanghai Maritime University, Shanghai, China, People's Republic of, Hongtao Hu

This study investigates sustainable production scheduling in intelligent hybrid flow manufacturing systems with crane transportation (HFMS-CT), addressing the dual challenges of mass customization and eco-friendly manufacturing. Distinct from conventional approaches that treat processing and logistics separately, we explicitly model the interdependency between machine operations and crane transportation to achieve an equitable balance between production efficiency (makespan minimization) and energy conservation. A novel bi-objective mixed-integer programming formulation is established, with the ε-constraint method employed for exact solutions in small-scale scenarios. To tackle the NP-hard nature of real-world applications, we propose a metaheuristic solution framework—Modified Multi-objective Harris Hawk Optimization (MMOHHO). It adopts greedy mechanisms by integrating Laplace crossover, tent-based chaotic mapping, elite selection, and nonlinear optimization strategy to balance exploration and exploitation capabilities. These mechanisms collectively ensure fair consideration of both competing objectives throughout the optimization process. *The experimental results reveal that the proposed algorithm outperforms other methods regarding NPS, DPO, IGD, and ES evaluation metrics.* The results particularly highlight the algorithm's exceptional capability in maintaining equitable solutions across the energy-efficiency spectrum, avoiding the common pitfall of sacrificing one objective for marginal gains in the other. A real-world case study from automotive component manufacturing further validates the approach. The findings provide practitioners with actionable strategies for implementing sustainable manufacturing paradigms that fairly reconcile ecological responsibility with operational competitiveness, supported by sensitivity analyses of crane utilization patterns and energy-price trade-offs.

4 - Proximal Autonomy V.S. Uav Delivery: a Self-Service-Enabled Uav Lorp Under Dynamic Demand and Cvar Constraints for Spatiotemporal Fairness

Zilu Zhao, Xi'an Jiaotong University, SHANXI, China, People's Republic of, Yongxi Cheng, ShiZeng Zhao

Unmanned Aerial Vehicles (UAVs) have garnered significant attention in post-disaster rescue operations, with most studies relying on Location-Routing Problem (LRP) models that uniformly address demand points regardless of their proximity to distribution centers (DCs). This work proposes a novel dynamic Location-or-Routing Model (LoRP) for post-disaster relief logistics by distinguishing between proximal and distant demand points: allowing self-service resource acquisition for nearby populations while prioritizing UAV-based delivery for medium-to-long-range demand points to enhance equity and efficiency. Conditional Value-at-Risk (CVaR) is employed to characterize dynamic demand and inventory uncertainties, while the Nash bargaining solution is integrated to capture decision-makers' preference structures. A Voronoi-enhanced NSGA-II algorithm with preference algorithm is developed, demonstrating computational expediency and effectiveness in large-scale simulations and real-world scenarios. Sensitivity analyses reveal critical trade-offs between infrastructure costs and UAV commuting distances, with star-shaped network topologies identified as optimal for pre-disaster warehouse placement.

SA11

Morrison Room

Emerging Techniques in Pricing and Choice Models

Invited Session

Revenue Management and Pricing

Chair: Zhenzhen Yan, Nanyang Technological University, Singapore, Singapore

1 - Assortment Optimization Without Prediction: An End-to-End Framework with Transaction Data

Shaoyu Wang, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Pin GAO, Andre Cire, Ningyuan Chen

Firms leverage assortment optimization to identify a set of products to offer to consumers in order to maximize the potential revenue. Traditionally, firms first use the historical transaction data to estimate a discrete choice model that predicts the choice behavior of future consumers, and subsequently solve the optimal assortment with such an estimated discrete choice model. This is typically referred to as the **predict-then-optimize** approach. In this paper, we propose an alternative approach that maps the data to the assortment decision in an **end-to-end** fashion, without the need to estimate a discrete choice model. The formulation extracts the partial orders induced from each customer in the historical transaction data. The assortment is then selected to maximize the worst-case revenue for a new customer when his/her preference is uniformly drawn from historical partial orders. We establish the NP-hardness of this problem and develop a mixed-integer programming formulation for exact solutions. Furthermore, we provide a number of approximation algorithms with theoretical guarantees. To demonstrate the external validity of the framework, we show the recommended assortment can achieve 1/2 of the optimal revenue when the data is generated by a rank-based choice model. The numerical result shows that the end-to-end framework is particularly attractive when (1) the data is scarce so that estimating model parameters can be unstable or (2) when the consumer behavior is complex and a pre determined discrete choice model is likely to be misspecified.

2 - From Small to Large: a Graph Convolutional Network Approach for Solving Assortment Optimization Problems

Guokai Li, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Pin GAO, Stefanus Jasin, Zizhuo Wang

Assortment optimization refers to selecting a subset of substitutable products to maximize the expected revenue obtained from customers whose purchase behaviors are modeled by some discrete choice model. It is a classic problem in revenue management and finds its application in multiple industries, e.g., e-commerce, advertisement, and transportation. Despite its widespread relevance, the problem is typically intractable due to its combinatorial and nonlinear nature. In this work, we investigate how to use neural networks to facilitate constrained assortment optimization. Specifically, we demonstrate how neural networks can effectively represent various choice models and constraints, and train the network to predict the optimal assortment. Surprisingly, despite being trained on a small-scale dataset (i.e., with a few products), the neural network performs remarkably

well on large-scale datasets (i.e., with lots of products), achieving a revenue ratio of approximately 90%. This imples that we can (approximately) solve large-scale constrained assortment problem by generating some small-scale instances. Our numerical experiments illustrate the efficiency and scalability of this approach for solving large-scale problems.

3 - Rkhs Choice Model

Zhi Wang, University of Texas at Austin, Austin, TX, United States, Yiqi Yang, Rui Gao, Shuang Li

Discrete choice models are essential for understanding human decision-making. Traditional approaches, rooted in Random Utility Models, rely on the assumption of stochastic rationality, limiting their ability to capture context-dependent behavior—where the presence of certain alternatives influences the perceived utility of others. To address these limitations, we propose a novel vector-valued Reproducing Kernel Hilbert Space (RKHS)-based choice model that flexibly captures context effects without imposing strong structural assumptions. Our framework supports both featureless and feature-based settings by introducing set kernels to model choice set interactions and product kernels to incorporate item-specific features. To ensure scalability, we leverage random feature approximations and Neural Tangent Kernels. We demonstrate that our models can be efficiently learned using convex optimization or stochastic gradient descent, and we derive generalization error bounds to guarantee robust performance. Empirical results on synthetic and real-world datasets show that our approach outperforms traditional discrete choice models in both predictive accuracy and generalization.

4 - Optimizing Product Pool Selection and Pricing for Size-Based Bundling Promotions

Xiaobo Li, National University of Singapore, Singapore, Singapore, Ruijiu Mao

We study size-based bundling promotions in the form of “p for s pieces.” Existing research on bundle promotions often assumes that all products are included in the promotion pool and fails to account for the impact of diverse product costs and customer preferences. To address these limitations, we develop a bundle choice model that extends the multinomial logit (MNL) model for size-based bundle settings. This model jointly optimizes pricing and the selection of products for the promotion pool, considering products with heterogeneous costs. We demonstrate that the promotion product pool selection problem is NP-hard and propose an approximation algorithm for optimal bundle promotion design with strong performance guarantees. Additionally, we present two tractable parameter estimation methods: one using utility-based data and the other using choice-based data, both of which involve either closed-form solutions or convex optimization. Numerical experiments show that the choice-based approach provides superior predictive power. Structural insights into the optimal promotion product pool are derived for special cases, including cost-interval and cost-nested structures under different conditions. Our study equips decision-makers with tools to design profitable size-based bundling promotions by optimizing product pools and pricing while incorporating customer preferences and cost variations. The proposed IEBCM framework and efficient estimation methods enable data-driven decision-making, enhancing the effectiveness of bundling strategies in multi-purchase settings.

5 - Optimizing Multiproduct Presale Pricing Strategies Using a Markovian Logit Choice Model

Yanxi Li, South China University of Technology, Guangzhou, China, People's Republic of

Presale is a common policy for online retailers, which consists of a deposit collected in advance and a final payment. Such a policy usually results in lower shipment costs due to the early knowledge of demand. However, the waiting period in presales may reduce consumers’ willingness to pay. In practice, online retailers sell substitutable products through both presale and traditional channels, where the substitutability of products and the interdependence of channels lead to a complex pricing landscape. We study the pricing problem for an online retailer who sales multiple substitutable products through presale and traditional channel. The consumer’s sequential decisions on paying the deposit and the final payment are characterized by a generalized Markovian logit choice model, which allows us to formulate a joint pricing optimization problem within the framework of marginal distribution model. We provide expressions for the optimal pricing, corresponding profits, and product market shares, and study how the retailer should adjust the waiting time in presales. In addition, we also consider the further introduction of a return policy. Finally, we conduct a case study using the data from a major online retailer in China, which verifies the precision and profitability of the model.

SA12

Moor Room

Learning in Revenue Management

Invited Session

Revenue Management and Pricing

Chair: Qinzhen Li, National University of Singapore, Singapore, Singapore

1 - Confusion Matrix Design for Downstream Decision-Making

Wei Tang, Chinese University of Hong Kong, Hong Kong, Hong Kong

We initiate the study of confusion matrix design. In this problem, an algorithm designer needs to generate a machine learning model (for a classification task from contexts to labels) which makes predictions for a population of downstream decision makers. The prediction accuracy of the machine learning model is characterized by its confusion matrix, which is a stochastic matrix where each entry encodes the probability of predicting the true label to another label. Each downstream decision maker faces a separate optimization task and will decide his binary action based on his own context, realized prediction given his context, and the confusion matrix selected by the algorithm designer. Decision makers are heterogeneous, as they may hold different contexts. Both the decision makers and the algorithm designer will obtain utilities that are determined by the actions the decision makers take, and their true labels. The goal of the algorithm designer is to design a public confusion matrix that is used for all decision makers subject to some feasibility constraints in order to maximize her net utility. We consider a general class of net utility functions, which could be a combination of both decision makers' utilities and the algorithm designer's utility. Classic outcome-independent utility and utilitarian/Nash/egalitarian social welfare are all special cases of our net utility formulation.

2 - Adaptive Bidding Policies for First-Price Auctions with Budget Constraints Under Non-Stationarity

Yige WANG, HKUST, Hong Kong, Hong Kong, Jiashuo Jiang

We study how a budget-constrained bidder should learn to adaptively bid in repeated first-price auctions to maximize her cumulative payoff. This problem arose due to an industry-wide shift from second-price auctions to first-price auctions in display advertising recently, which renders truthful bidding (i.e., always bidding one's private value) no longer optimal.
We propose a simple dual-gradient-descent-based bidding policy that maintains a dual variable for the budget constraint as the bidder consumes her budget. In analysis, we consider two settings regarding the bidder's knowledge of her private values in the future: (i) an uninformative setting where all the distributional knowledge (can be non-stationary) is entirely unknown to the bidder, and (ii) an informative setting where a prediction of the budget allocation in advance.
We characterize the performance loss (or regret) relative to an optimal policy with complete information on the stochasticity. For uninformative setting, We show that the regret is $\tilde{O}(\sqrt{T})$ plus a variation term that reflects the non-stationarity of the value distributions, and this is of optimal order. We then show that we can get rid of the variation term with the help of the prediction; specifically, the regret is $\tilde{O}(\sqrt{T})$ plus the prediction error term in the informative setting.

3 - A Behavioral Model for Multi-Armed Bandits: Theoretical Framework and Experimental Evidence

Yifan Feng, National University of Singapore, Singapore, Singapore, Jingying Ding, Ying Rong

How do people navigate the exploration-exploitation (EE) trade-off in repeated decisions under uncertainty? We study this question using multi-armed bandit problems and introduce a novel behavioral model: Quantal Choice with Adaptive Reduction of Exploration (QCARE). QCARE generalizes Thompson Sampling by incorporating a dynamic mechanism that reduces exploration as information accumulates, with the reduction rate serving as a key parameter to capture EE dynamics. We theoretically analyze how different reduction rates affect decision quality, revealing the impacts of over- and under-exploration. Empirically, QCARE is validated with behavioral data from human participants and outperforms alternative models in predictive accuracy. Our results highlight a consistent tendency toward over-exploration in human decision-making.

4 - When Worst-Case Isn't Robust: On the Limitations of Distributionally Robust Formulations in Secretary Problems

Anas Abdelhakmi, National University of Singapore, Singapore, Singapore, Andrew Lim

Secretary Problems are a fundamental class of sequential search decision-making problems where the decision-maker's objective is to select the best item from a sequence presented online. We formulate a Distributionally Robust Optimization (DRO) version of secretary problems and show that the optimal robust decision coincides exactly with the nominal decision across various secretary problem variants and under various uncertainty sets (e.g., Kullback-Leibler divergence, Wasserstein distance). In the special case where only items relative ranks are shown, the well-known 1/e rule is worst-case optimal. We show, however, that the DRO solution is not robust in that it does not reduce the sensitivity of the expected reward to worst-case deviation from the nominal model. We show that sensitivity in the expected reward comes from the exploration shocks over time and that sensitivity can be reduced by introducing hedging instruments that offset these shocks. Simple levers, such as granting an option to recall previously rejected items at a cost, can reduce sensitivity. More generally, these results offer new insights of what it means for sequential decisions to be robust and show how practitioners can achieve robustness beyond traditional worst-case strategies.

SA13

Bailey

Data-Driven Decision Support and Optimization

Contributed Session

Decision Support Systems

Chair: Alberto Costa, 127715

1 - The Divorce of Word and Deed --- a Data-Mining Approach to Identify and Evaluate Customer Requirements

Kexin Ai, Xi’an Jiao Tong University, Xi'an, China, People's Republic of, Juan Feng, Xinyu Sun

The advancement of IT has made it possible to understand customer requirements through publicly available information, such as online word-of-mouth. It is challenging, however, to apply these results to guide product design, without considering consumers’ purchase (deed). We propose a framework to understand customer requirements by incorporating one additional information source, that is, consumers’ purchase records, and examine the relationship between customers’ expressed satisfaction (words) and their actual purchase (deeds). Utilizing unique data sets in a cell phone market, we mine and characterize each product feature in a two-dimensional model, based on how important they are to customers’ expressed satisfaction as well as to the actual purchase. We find that, consumers’ words and deeds are not always consistent: while some requirements are expressed as important features in determining customer satisfaction in “words”, they do not stand out in factors that influence consumers’ purchase; there are also requirements that are “silent” in satisfaction, while customers value them when making purchase. This implies the importance of incorporating the “deed” dimension into the understanding of customer requirements. Understanding the differences in product requirements reflected from these two dimensions offers guidelines for better product design, such as which features worth investment to improve/innovate, and which features not.

2 - Integrated Last Mile Delivery and Workforce Scheduling to Meet Demand Variations In E-Commerce Retail

Saurabh Chandra, Indian Institute of Management Indore, Indore, India, Amit Vatsa

Last mile distribution in e-retail of consumer items are labor intensive operations. E-commerce retail rms frequently update their last mile distribution networks and assign deliveries to pre-defined customer routes to cater to changing demand patterns. For a given distribution network and delivery assignments for a group of customers, an important operational plan comprises of workforce scheduling to meet the delivery assignments. Seasonal variations and special events like the COVID-19 pandemic lead to sharp rise in demand for e-commerce retail. This may even be followed by uncertainty in the workforce availability. In these circumstances, e-commerce retail rms need to operate under short periods of workforce shortage followed by quick hiring to meet the demand. We present a problem of managing workforce to schedule delivery assignments across fixed customer routes in last mile e-commerce retail distribution. We present two related problems, one estimating the minimum workforce requirement to meet the given customer demand and the second aiming to maximize

customer deliveries using the available workforce. Mixed integer programming and equivalent constraint programming models are proposed for the two problems. A two-step solution framework is presented to solve real size instances of the workforce scheduling models. A case study demonstrates the applicability of the models in a real scenario.

3 - Observational Data Driven Decision Making

Huan Xu, Shanghai Jiaotong University, Shanghai, China, People's Republic of

In the era of data analytics, business decisions are increasingly shifting from being model-driven to data-driven. Data-driven decision-making brings forth many new research topics, including decision-focused machine learning, the impact of data characteristics on decisions, and the integration of learning and decision-making.

In this report, based on Alibaba’s practices, we explore business decision-making and insights derived from observational data (i.e., non-experimental data). In a project involving markdown clearance of perishable goods on a new retail fresh food platform, we proposed an aggregated semi-parametric price-quantity model to address the endogeneity of observational data and to cope with data sparsity. We also designed a multi-stage optimization algorithm to dynamically set clearance prices.

In a policy study, we leveraged a natural experiment resulting from the implementation of a waste-sorting policy in Shanghai and discovered that among Generation Z, environmental awareness is a key driving force behind participation in the secondhand market.

These practices highlight the importance of thoroughly understanding data characteristics for effective decision-making.

4 - Leveraging Generative AI for Automated Dashboard Design and Insight Generation

Priyabrata Purohit, Adobe, Bangalore, India

As organizations increasingly rely on data to drive decisions, the need for rapid, insightful, and user-friendly dashboards is more critical than ever. This poster presents a novel approach that integrates Generative AI (GenAI) into the Power BI ecosystem to streamline dashboard creation and enhance data storytelling. By leveraging natural language prompts, users can generate report layouts, select relevant visuals, and even receive automated narrative insights — all within Power BI. This reduces development time, lowers the technical barrier for non-technical users, and accelerates the path from raw data to actionable insights. The poster showcases real-world implementations, highlighting how GenAI can augment the traditional BI workflow and empower data professionals with faster prototyping, intelligent recommendations, and adaptive storytelling. We also address key considerations such as data privacy, model accuracy, and ethical use. This work sets the foundation for the future of human-AI collaboration in business intelligence.

5 - Inferno: a Framework for Knowledge Generation with Expensive Sample Evaluations

Alberto Costa, Singapore-ETH Centre, Singapore, Singapore, Riccardo Talami

Real world problems often require complex modeling and computational efforts
to be effectively addressed. Relying solely on data-driven approaches, without
integrating physics-based models, can lead to limited predictive capabilities.
Even advanced techniques such as deep learning may be impractical for decision-makers due to challenges in obtaining sufficient data and the need to justify and explain
results.

To address these challenges, we introduce INFERNO (INference Framework for Efficient
Rule-based kNOwledge generation). This framework combines
derivative-free optimization, clustering, and a novel procedure for deriving inference rules. These inference rules can be easily interpreted and used by decision-makers to identify explainable solutions. We propose a metric called the Price of eXplainability (PoX) to assess the trade-off between solution quality and explainability.

The framework was validated using a building design problem, with two test cases related to energy and thermal comfort optimization. Comparisons with a Classification Tree show that INFERNO's Inference Rules achieve a lower PoX while effectively narrowing the search space for decision-makers.

SA14

Wilberforce

Simulation and Learning

Invited Session

Simulation Society

Chair: Guangwu Liu, City University of Hong Kong, Kowloon, Hong Kong

Co-Chair: Mingbin Ben Feng, University of Waterloo, Waterloo, ON, Canada

1 - Large-Scale AI-Enabled Virtual Screening

Weiwei Fan, Tongji University, Shanghai, China, People's Republic of, Zaile Li, Jeff Hong

Screening tasks that aim to identify a small subset of top alternatives from a large pool are common in business decision-making processes. These tasks often require substantial human effort to evaluate each alternative’s performance, making them time-consuming and costly. Motivated by recent advances in large language models (LLMs), particularly their ability to generate outputs that align well with human evaluations, we consider an LLM-as-human-evaluator approach for conducting screening virtually, thereby reducing the cost burden. To achieve scalability and cost-effectiveness in virtual screening, we identify that the stochastic nature of LLM outputs and their cost structure necessitate efficient budget allocation across all alternatives. To address this, we propose using a top-m greedy selection mechanism, a simple yet effective approach that keeps evaluating the current top-m alternatives, and design the explore-first top-m greedy (EFG-m) algorithm. We prove that EFG-m is both sample-optimal and consistent in large-scale virtual screening. Surprisingly, we also uncover a bonus ranking effect, where the algorithm naturally induces an indifference-based ranking within the selected subset. To further enhance practicality, we design a suite of algorithm variants to improve screening performance and computational efficiency. Numerical experiments validate our results and demonstrate the effectiveness of our algorithms. Lastly, we conduct a case study on LLM-based virtual screening. The study shows that while LLMs alone may not provide meaningful screening and ranking results when directly queried, integrating them with our sample-optimal algorithms unlocks their potential for cost-effective, large-scale virtual screening.

2 - Efficient Nested Simulation for Value-at-Risk

qianwen zhu, City University of Hong Kong, Hong Kong, China, People's Republic of, Guangwu Liu, Xianyu Kuang

Nested simulation plays a vital role in financial risk management by providing estimates of critical risk measures such as Value-at-Risk (VaR). Traditional methods, however, estimate VaR utilizing all scenarios, neglecting the fact that estimating VaR using only scenarios near the targeted risk level would be more efficient.

To address this, we propose a nested simulation method enhanced by scenario screening for simultaneous point and interval estimation of VaR. The procedure begins by generating a large set of outer scenarios and retains only a small set of scenarios close to the targeted risk level. It then draws inner samples for surviving scenarios to estimate VaR efficiently.

Theoretical analysis guarantees the asymptotic coverage probability of the constructed confidence interval. Numerical experiments confirm the effectiveness of the method in delivering accurate point estimate and highly reliable interval estimate for VaR.

3 - Top-M Arms Identification with a Generalized Ucb Approach

PEIWEN YU, Soochow University, Suzhou, China, People's Republic of

We study the problem of identifying the top-m arms in stochastic multi-armed bandits under a fixed-budget setting. To address this problem, we propose a generalized UCB-based algorithm (GUCB) designed to efficiently

identify the m arms with the highest expected rewards. We rigorously establish that as the sample size approaches infinity, the set of top-m arms selected by our algorithm converges to the true top-m arms with probability one. Furthermore, we derive a theoretical bound on the probability of misidentification, demonstrating the algorithm's reliability in finite-budget scenarios. Numerical experiments validate the superior empirical performance of our proposed algorithm.

4 - Contextual Strongly Convex Stochastic Optimization: Optimize, Then Predict with Inexact Solutions

Heng Luo, Fudan university, Shanghai, China, People's Republic of, Nifei Lin, Jeff Hong

This paper address the real-time decision making challenge of contextual strongly convex stochastic optimization problem by proposing an "optimize, then predict" algorithmic framework. The stochastic optimization algorithms and smoothing techniques are employed to approximate the optimal solution function with respect to covariate. We highlight that the solutions derived from stochastic optimization algorithms may not necessarily be unbiased, which is neglected by previous works. Through analyses of various smoothing techniques, we demonstrate the efficiency of our algorithm and the substantial impact of bias on sample allocation and the rate of convergence. Numerical examples further validate the efficacy of the algorithm and offer practical insights into the selection of suitable smoothing techniques.

5 - Enhanced Derivative-Free Optimization Using Adaptive Correlation-Induced Finite Difference Estimators

Guo Liang, Renmin University of China, Beijing, China, People's Republic of, Guangwu Liu, Kun Zhang

Gradient-based methods are well-suited for derivative-free optimization (DFO), where finite-difference (FD) estimates are commonly used as gradient surrogates. Traditional stochastic approximation methods, such as Kiefer-Wolfowitz (KW) and simultaneous perturbation stochastic approximation (SPSA), typically utilize only two samples per iteration, resulting in imprecise gradient estimates and necessitating diminishing step sizes for convergence. In this paper, we first explore an efficient FD estimate, referred to as correlation-induced FD estimate, which is a batch-based estimate. Then, we propose an adaptive sampling strategy that dynamically determines the batch size at each iteration. By combining these two components, we develop an algorithm designed to enhance DFO in terms of both gradient estimation efficiency and sample efficiency. Furthermore, we establish the consistency of our proposed algorithm and demonstrate that, despite using a batch of samples per iteration, it achieves the same convergence rate as the KW and SPSA methods. Additionally, we propose a novel stochastic line search technique to adaptively tune the step size in practice. Finally, comprehensive numerical experiments confirm the superior empirical performance of the proposed algorithm.

SA15

Fairmont Ballroom

Enhancing Humanitarian Logistics with Network Analytics

Invited Session

Transportation Science and Logistics

Chair: Abdullah Konak, Penn State Berks, P.O. Box 7009 Tulpehocken Road, Boston, MA, 19610, United States

1 - Promoting Equity in Emergency Food Assistance: a Data-Driven Approach

Sadan Kulturel-Konak, Penn State Berks, Reading, PA, United States, Abdullah Konak

This study presents a research framework to improve fairness in emergency food assistance programs (EFAP) by addressing geographic and accessibility inequities in the U.S. The framework emphasizes the equitable allocation of resources to vulnerable populations, particularly those with limited transportation options. Integrating stakeholder engagement, spatial analysis, and optimization techniques leverages granular, location-specific data to identify systemic inefficiencies and inform decision-making. To enhance adoption and practicality, the framework includes tailored visualizations to engage diverse stakeholders, such as food banks, community organizations, volunteers, and beneficiaries. It bridges the gap between raw data collection and actionable tools, accounting for stakeholders' technical limitations. Operating at a highly localized level, the framework prioritizes accessibility for underserved communities and incorporates feedback mechanisms to enable continuous improvement. This collaborative, adaptable approach aims to transform food distribution networks into more equitable and efficient systems, contributing to long-term food security and resilience. This

presentation will include demonstrations of implementing the framework in local food bank operations in the U.S.

2 - Predictive Analytics in Food Bank Operations

Abdullah Konak, Penn State Berks, Reading, PA, United States, Steven Jiang, Lauren Davis, Sadan Kulturel-Konak

Food banks rely on donated food items to serve individuals in need. Major sources of these donations include food distribution centers, warehouses, and grocery stores. Unfortunately, food donations often come with a high degree of uncertainty. Predicting food donations and understanding the donation patterns of individual suppliers are crucial for food banks for several reasons. Firstly, it allows them to schedule their pickup operations more effectively and make better decisions regarding supplemental food drives or purchases. In this paper, we analyze the food donations received by a food bank and use predictive analytics tools to support scheduling decisions for food pickups under uncertainty. We also demonstrate a prototype dashboard designed to aid in these decisions.

3 - Coherent Photonic Quantum Computer for Transportation and Logistics Optimization: Breakthroughs from Theory to Implementation

Qi Gao, QBoson, Beijing, China, People's Republic of, Chelsea Gong

Transportation and logistics optimization problems often involve NP-hard challenges that traditional computing methods struggle to solve efficiently at scale. This presentation explores how coherent photonic quantum computer offers a promising alternative through quantum approaches. We examine potential implementations across several case studies, including AGV routing, storage allocation, subway scheduling, and three-dimensional packing problems. Unlike other quantum technologies requiring cryogenic environment, coherent photonic quantum computers can operate at room temperature with potentially lower energy consumption and greater scalability. Our presentation suggests improvements in solution quality and computational efficiency compared to conventional methods. We present our exploratory research on potential coherent photonic quantum processing with classical optimization frameworks, discuss potential business value, and outline a roadmap for this emerging field. This work highlights how coherent photonic quantum computer represents a valuable new direction for transportation optimization despite current technological limitations.

4 - Cost Sharing in Quantum Volunteer's Dilemma

SiongThye Goh, Institute of High Performance Computing, Singapore, Singapore, Dax Enshan Koh

The volunteer's dilemma is a well-known game in game theory that models the conflict players face when deciding whether to volunteer for a collective benefit, knowing that volunteering incurs a personal cost. Studying volunteer's dilemma can help to benefit multiple parties and encourage sharing of cost. In this work, we introduce a quantum variant of the classical volunteer's dilemma, generalizing it by allowing players to utilize quantum strategies. Employing the Eisert–Wilkens–Lewenstein quantization framework, we analyze a multiplayer quantum volunteer's dilemma scenario with an arbitrary number of players, where the cost of volunteering is shared equally among the volunteers. We derive analytical expressions for the players' expected payoffs and demonstrate the quantum game's advantage over the classical game. In particular, we prove that the quantum volunteer's dilemma possesses symmetric Nash equilibria with larger expected payoffs compared to the unique symmetric Nash equilibrium of the classical game, wherein players use mixed strategies. Furthermore, we show that the quantum Nash equilibria we identify are Pareto optimal. Our findings reveal distinct dynamics in volunteer's dilemma scenarios when players adhere to quantum rules, underscoring a strategic advantage of decision-making in quantum settings.

SA16

VIP A

Digital Transformation and Innovation in Sustainability

Contributed Session

Sustainable and Socially Responsible Operations

Chair: Ankit Sharma, Indian Institute of Management Amritsar, Amritsar, 143105, India

1 - Integration of ESG Practices In Supplier Assessment and Selection: a Case Study

Simon Yuen, CPCE, The Hong Kong Polytechnic University, Hong Kong, China, People's Republic of, Calvin Cheng, Helen S.M. Wong

This study explores Towngas's commitment to integrating Environmental, Social, and Governance (ESG) practices into its supplier assessment and selection process. As a leading gas utility provider in Hong Kong, Towngas recognizes the critical importance of ESG practices, viewing them as both ethical imperatives and essential elements of strategic operational commitment. Through a qualitative analysis of Towngas's ESG initiatives, including interviews with sustainability and procurement leaders, the study illustrates how the company prioritizes suppliers who demonstrate a commitment to ESG practices. The findings reveal that Towngas adopts a holistic approach to integrating ESG into its supply chain, emphasizing collaboration with suppliers to enhance sustainability capabilities. This approach not only strengthens operational resilience but also enhances the company's market reputation. Besides, the study reveals the challenges Towngas encounters in implementing ESG criteria, such as the need for comprehensive assessments and ongoing monitoring of supplier performance. It highlights the best practices and innovative strategies employed by Towngas, including supplier collaboration and leveraging advanced technologies for greater transparency. The originality of this study lies in its focus on a leading utility provider's practical implementation of ESG factors, showcasing how such practices can foster a responsible and resilient business model in the energy sector.

2 - Name Matters: the Impact of Supplier Disclosure on Firm Performance

Xinrui Zhan, Anhui University of Finance and Economics, Bengbu, China, People's Republic of, Zhi Cao

Purpose: Supply chain transparency (SCT) has garnered increasing scholarly attention, yet many firms remain cautious about disclosing supplier information due to uncertain financial consequences. While prior research has emphasized the general benefits of transparency, the specific financial implications of disclosing supplier names remain underexplored.

Design/methodology/approach: Using on a panel dataset of publicly listed firms in China from 2012 to 2022, we examine the relationship between supplier name disclosure and financial performance, measured by return on assets (ROA). To address endogeneity concerns, we employ propensity score matching and the Heckman two-stage model. We further investigate the moderating effects of market position, industry competition, and supplier concentration.

Findings: Our results show that supplier name disclosure is associated with a 0.5% increase in ROA, indicating a positive financial effect. However, this benefit diminishes—and may even reverse—in highly competitive industries or among firms with concentrated supplier bases, where disclosure can expose strategic vulnerabilities. These findings underscore the need to account for contextual factors when implementing transparency strategies.

Originality/value: This study provides robust empirical evidence on the financial consequences of supplier transparency, addressing a gap in prior theoretical discussions. By uncovering key moderating conditions, it reveals when disclosure enhances firm value and when it introduces strategic vulnerabilities. The findings yield actionable insights for managerial decision-making and contribute to policy debates on corporate disclosure practices. Overall, this research deepens the understanding of supply chain transparency by establishing its contingent effects on firm profitability and highlighting its nuanced strategic role.

3 - Resilient Closed-Loop Supply Chain Network Design via Machine Learning Aided Outer Approximation Algorithm

Hongzhen Li, Beijing Institute of Technology, BEIJING, China, People's Republic of, Yanzi Zhang

This study explores the location-inventory problem in a resilient closed-loop supply chain, considering interdependent disruption risks of co-located forward and reverse distribution centers. To address this problem, a

machine learning aided outer approximation algorithm is proposed. Numerical experiments are conducted to evaluate algorithm's performance and explore marginal insights.

4 - Mine OR Not to Mine? Voluntary Vehicle Modernization Program's Impact on Auto Supply Chains

Ankit Sharma, Indian Institute of Management Amritsar, Amritsar, India

Voluntary Vehicle Modernization Programs (VVMP) have gained significant attention as a policy tool to address the environmental and economic challenges posed by aging and polluting vehicles. This study aims to evaluate the effectiveness of the vehicle scrappage policy, launched in India in late 2021, and its associated claims in the policy document. The study focuses on investigation around the circular economy promises of VVMP by the introduction of Registered Vehicle Scrappage Facilities (RVSFs) to bring recycled material back into the automotive supply chain as government's push towards sustainability.
To assess the impact of VVMP, firstly, a comprehensive data analysis is conducted to explore the current state of the scrappage facilities in the country through various sources, including government reports, industry statistics and news articles. The study employs quasi-experimental approach to investigate the recycling claims made by the policy. We further deploy systems thinking and system dynamics approach and deploys a mental model and resulting structure of the problem. Subsequently, a stock flow diagram is built to simulate various scenarios to generate workable policy alternatives. The model measures the reduction in virgin metals in automotive supply chain, resulting from the scrappage of old vehicles and the subsequent adoption of newer, cleaner vehicles. Furthermore, the study examines the extent to which the vehicle scrappage policy contributes to the integration of recycled materials back into the vehicle supply chain. The analysis considers the availability, feasibility, and adoption of recycled materials within the automotive industry.

SA17

VIP B

Theories and Applications in Data-driven Learning

Contributed Session

Computing Society

Chair: KUMAR RAHUL, Indian Institute of Management Kozhikode, Kozhikode, N/A, India

1 - Solving Sparse & High-Dimensional-Output Regression via Compression

Renyuan Li, National University of Singapore, Singapore, Singapore, Zhehui Chen, Guanyi Wang

Multi-Output Regression (MOR) has been widely used in scientific data analysis for decision-making. Unlike traditional regression models, MOR aims to simultaneously predict multiple real-valued outputs given an input. However, the increasing dimensionality of the outputs poses significant challenges regarding interpretability and computational scalability for modern MOR applications. As a first step to address these challenges, this paper proposes a Sparse & High-dimensional-Output REgression (SHORE) model by incorporating additional sparsity requirements to resolve the output interpretability, and then designs a computationally efficient two-stage optimization framework capable of solving SHORE with provable accuracy via compression on outputs. Theoretically, we show that the proposed framework is computationally scalable while maintaining the same order of training loss and prediction loss before-and-after compression under arbitrary or relatively weak sample set conditions. Empirically, numerical results further validate the theoretical findings, showcasing the efficiency and accuracy of the proposed framework.

2 - Systemic-Risk Mitigation in Financial Networks: a Distributionally Robust Spatial-Var Approach

Hancheng Lu, University of Chinese Academy of Sciences, Beijing, China, People's Republic of, Jianping Li, YunFong Lim, Shuming Wang

We study a systemic-risk mitigation problem over a financial network under random shocks to the value of banks’ assets. The decision maker (government) determines the excess-reserve allocation across different banks so as to maximize the systemic stability level in terms of their expected total repayment (sum of inter-bank lending, required reserves, and post-shock bail-out levels). To address the distributional uncertainty of the random shocks over the network, we develop a predictive distributionally robust optimization (DRO) model that leverages the spatial vector auto-regressive (SVAR) prediction model. In particular, we construct a predictive

distribution using the estimated SVAR model with feature information (covariates) and develop the resultant Wasserstein ambiguity set for the DRO model, which captures both the network effect and ambiguity effects of the random shocks in a statistically coherent fashion. Computationally, we show that the resultant DRO model can be transformed into a mixed-integer conic linear program that can be solved by off-the-shelf MIP solvers when a proper norm is employed. Operationally, we analyze the optimal excess-reserve allocation, and identify a necessary condition for the optimal allocation and a sufficient condition for the uniform optimal allocation. We also perform the stress-testing analysis on the worst-case expected bail-out level and the sensitivity of the worst-case systemic stability to the variation in the liability over the network. Statistically, we derive finite-sample and asymptotic performance guarantees of our proposed SVAR-based DRO model. Finally, sufficient numerical experiments demonstrate the effectiveness of the proposed approach.

3 - Prescriptive Svm with Stochastic Constraints

Shiming Deng, Huazhong University of Science and Technology, Wuhan, China, People's Republic of, Xianghui Cheng, Chu Chu

We reformulate the prevalent SVM models to incorporate the two salient characteristics of decision problems, asymmetric costs and stochastic constraints. To exam their distinguishing effects separately, we first formulate SVM models accounting for only asymmetric costs. We develop a Prescriptive SVM model (PSVM), and an approximation model using hinge function (PSVM-A), which can be efficiently solved. We prove that both the PSVM and PSVM-A are universally consistent, i.e., the decision rules obtained using the two models converge to the true Bayes cost decision rules. With consideration of operational constraints, we construct a Constrained Prescriptive SVM (CPSVM) model with a sample-estimated resource constraint. To ensure that the estimated optimal decision can fully utilize resources with no budget waste, we propose a budget-elevation method and provide a lower bound for the elevation term to ensure that the empirical feasible region contains the true feasible region with a high probability. Moreover, we demonstrate that the adaptive elevation term converges to zero and the modified empirical feasible region converges to the true feasible region in set distance as sample sizes goes to infinity. We conduct extensive numerical studies to validate our results and provide interesting managerial insights.

4 - On Distributionally Robust Optimization for Cross-Docking Platforms Design Under Uncertainty

M. Araceli GARIN, University of the Basque Country, UPV/EHU, Leioa, Bizkaia, Spain, Laureano Escudero, Aitziber Unzueta

Distributionally Robust Optimization (DRO) is considered to deal with the different uncertainties in the Cross-dock Door Design Problem (CDDP) that consists of deciding the strip and stack doors and nominal capacity of the cross-dock infrastructure.

A two-stage DRO mixed binary quadratic model is presented for CDDP solving; the first stage decisions are related to the design of the cross-dock infrastructure; the second stage ones are related to the assignment of the commodity flows to the doors in a finite set of scenarios for the ambiguity set members.

The problem combines three highly computational difficulties, namely, NP-hard combinatorics, uncertainty in the main parameters and their probability distribution. Distributionally robust optimization is considered to deal with these uncertainties.

The goal is to minimize the total highest cost in the ambiguity set, subject to the constraint system for each of those members and the stochastic dominance risk averse functional.

As far as we know, the challenging problem that results has not been addressed before, although its application field is a very broad one.

Given the problem solving difficulty, a scenario cluster decomposition and a matheuristic are proposed for obtaining lower and upper bounds, respectively. A computational study validates the proposal; the approach overperformances the straightforward use of the state-of-the-art solvers Cplex and Gurobi.

5 - Reimagining Bibliometric Analysis - a Novel Deep Learning Framework Bridging the Pitfalls

KUMAR RAHUL, Indian Institute of Management Kozhikode, Kozhikode, India, Shovan Chowdhury, Gopalakrishnan Narayanamurthy

The high volume of research publications is causing the time and effort researchers must invest to conduct a comprehensive literature review to grow exponentially. Recent research has highlighted the need for computational literature review to augment human tasks. Bibliometric review is a type of computational and domain review offering meaningful insights from a large corpus of articles. But the techniques listed as part of the bibliometrics review have several shortcomings, including the inability to standalone identify emerging themes in literature review. To address these shortcomings, we propose a smart literature review framework named Thematic Review with Inbuilt Validation (TRIVAL). We demonstrate that TRIVAL can improve, if not substitute, bibliographic coupling, co-word analysis, and co-citation analysis, thereby contributing to the knowledge of the computational literature review. It will augment the interpretive analysis of the researcher, within each theme, to generate novel theoretical insights. We show the application of TRIVAL framework using the bi-directional encoder representation from transformer (BERT) topic, Large Language Model Meta AI (LLaMA2), and Bidirectional and autoregressive transformer (BART). BERT topic is an unsupervised machine learning algorithm for topic modeling. LLaMA2, a large language model, is used to name thematic clusters and finally BART is deployed for validation of the themes.

Sunday, July 20, 10:00 AM - 11:00 AM

Fairmont Ballroom

Mathematical Optimization in the Era of AI

Invited Session

Plenary

1 - Mathematical Optimization in the Era of AI

Yinyu Ye, Stanford University, Stanford, CA, United States

This talk aims to present several mathematical optimization problems/algorithms for AI such as the LLM training, tunning and inferencing. In particular, we describe how classic optimization models/theories can be applied to accelerate and improve the Training/Tunning/Inferencing algorithms that are popularly used in LLMs. On the other hand, we show breakthroughs in classical Optimization (LP and SDP) Solvers aided by AI-related techniques such as first-order and ADMM methods, the low-rank SDP theories, and their implementations on GPU.

Sunday, July 20, 11:15 AM - 12:45 PM

SB01

Enterprise

Stochastic Modelling and Analysis

Invited Session

Applied Probability Society

Chair: Dacheng Yao, Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing, N/A

1 - A Non-Convex Optimization Approach of Searching Algebraic Degree Phase-Type Representations for General Phase-Type Distributions

Yujie Liu, National University of Singapore, Singapore, Singapore, Dacheng Yao, Hanqin Zhang

For a continuous-time phase-type (PH) distribution, starting with its Laplace–Stieltjes transform, we obtain a necessary and sufficient condition for its minimal PH representation to have the same order as its algebraic degree. To facilitate finding this minimal representation, we transform this condition equivalently into a non-convex optimization problem, which can be effectively addressed using an alternating minimization algorithm. The algorithm convergence is also proved. Moreover, the method we develop for the continuous-time PH distributions can be used directly for the discrete-time PH distributions after establishing an equivalence between the minimal representation problems for continuous-time and discrete-time PH distributions.

2 - Multi-Agent Queueing Network: Coordination with Learning

Zhaoxuan Wei, Institute of Operations Research and Analytics, NUS, Singapore, Singapore, Andrew Lim, Hanqin Zhang

For the Jackson queueing network with J servers, we consider its decentralized control problem among $(J+1)$ agents, job-admission decision maker and J job-service providers. The admission decision maker controls the exogenous job arrivals by dynamically setting the admission fees while the service providers decide their job-processing rates. The system is operated in a decentralized way: all agents optimize their own objective subject to their local models which may mis-specify the system components controlled by other agents. We prove that this decentralized queueing network can be coordinated through an optimal state-dependent revenue-transfer mechanism in the sense of the decentralized system welfare equal to the optimal centralized system welfare. The algorithm for computing this optimal mechanism and its convergence rate are also established. Further, based on the interaction with local generative models, we propose a reinforcement learning algorithm for the case where the decision makers may not even know the performance given by the components they are controlling. In our analysis, a new approach, that integrates the decomposition of centralized system profit and the coupling of local objectives of agents, is developed. The decomposition helps us to identify the structure properties on the optimal state-dependent revenue-transfer mechanism while the coupling provides us a bridge to establish the algorithms convergence.

3 - Heterogeneous Information Provision and Bayesian Effect for Queues with Strategic Customers

Jian Cao, Nanjing University of Finance and Economics, Nanjing, China, People's Republic of, Yongjiang Guo

The costless information provided by third-party service platforms is a key ingredient in influencing the customers' behavior and the store's operation. We consider a single-server queueing system with uncertain service quality. There are two third-party service platforms with different levels of information accuracy, and each service platform sends a noisy signal of service quality to its subscribed customers. Based on signals, customers update their beliefs by using Bayes' rule and make the queueing decisions. We first characterize the customers' equilibrium joining decisions, and then investigate the system's metrics, such as throughput. We find that the customers' equilibrium joining probabilities may be multiple, but the customers' equilibrium arrival rate is always unique. Moreover, a larger potential arrival rate may not induce a higher equilibrium arrival rate of customers. Furthermore, under the information with heterogeneous accuracy, the system's profit may not become worse off when customers have access to more accurate information about the service quality. Finally, we confirm that a zero-information strategy is more beneficial for the system manager to maximize the throughput.

4 - Dynamic Inventory Management with Lost-Sales and General Ordering Cost

Peng Hu, Huazhong University of Science and Technology, Wuhan, China, People's Republic of

We study a lost-sales single-product dynamic inventory system, where the ordering cost can be a general increasing functions. It turns out that if the inventory holding costs are relatively larger than shortage costs in all

periods, quasi-convexity together with a so-called semi-convexity property can be preserved in this stochastic dynamic programming problem, which ensures the optimal dynamic ordering policy to be completely characterized: an order is placed when the initial inventory level falls below some threshold, and higher initial inventory levels lead to lower ordering quantities in each period.

SB02

Mercury

Airport Slot Allocation

Invited Session

Aviation Applications

Chair: Nuno Ribeiro, Singapore University of Technology and Design, Singapore, Singapore

1 - Airport Slots Supply and Demand Management in An Airport Network

Linzi Zhang, Nanjing University of Aeronautics and Astronautics, Nanjing, China, People's Republic of, YANJUN WANG

As the core element of civil aviation operation, the reasonable allocation of flight time resources is very important to improve the operational efficiency and service quality of the civil aviation industry. By constructing Airport Slots Demand Index (ASDI) and Airport Slots Availability Index (ASAI), this study quantitatively evaluates the supply and demand of flight time resources, develops Airport slots identification classification (ASIC) matrix, and divides airports into four categories: Insufficient stock - High demand (IH), Sufficient stock - High demand (SH), Insufficient stock - Low demand (IL), and Sufficient stock - Low demand (SL).

Through the analysis of the flight data for the whole year of 2024, it is found that 31 airports have the problem of flight demand exceeding supply at all times, and the ASDI value of Chengdu Shuangliu (CTU) Airport and Urumqi (URC) Airport is high, indicating that these airports are facing serious contradiction between supply and demand. The ASAI value of Shenzhen (SZX) Airport is the lowest, indicating that its flight time stock is extremely short. According to the classification of ASIC matrix, 9 airports are classified as IH, facing serious contradiction between supply and demand; 12 airports belong to the IL and SH categories, which have certain supply and demand problems but still have buffer space. 16 airports belong to the SL category and are fully stocked with flights at all times.This study provides a scientific basis for the optimal allocation of flight time resources.

2 - An Integrated Slot Allocation Model for Aircraft Emission Reduction

Hu Rong, Nanjing University of Aeronautics and Astronautics, Nanjing, China, People's Republic of, Yutong Zhang

This study focuses on the potential contribution of slot allocation to the reduction of pollutant emissions at schedule-coordinated airports. We develop a bi-objective optimization model for a single airport, integrating scheduling efficiency and pollutant emissions, and propose regulatory strategies that consider both environmental sustainability and operational benefits. The results indicate that slot allocation plays a positive role in reducing pollutant emissions at Shanghai Hongqiao International Airport. Further analysis reveals a marginally increasing cost relationship between emission cost reduction and total displacement. Additionally, variations in the maximum allowable displacement and unit emission prices significantly affect the emission reduction potential. In contrast, although increasing airport capacity yields certain emission reduction benefits, its marginal effect is relatively limited, and its role in slot allocation is not significant. Different slot strategies exhibit distinct fairness performances, therefore, establishing a coordinated mechanism that balances efficiency, emission costs reduction, and airline fairness remains a key direction for future optimization.

3 - Airport Capacity Declaration for Slot Allocation: A Modelling Framework

Nuno Ribeiro, Singapore University of Technology and Design, Singapore, Singapore

This work presents an original modelling framework to support airport managers in setting declared capacities for slot allocation at slot coordinated airports. The framework integrates a delay prediction model, an integer programming model of slot allocation and a multi-objective genetic algorithm (NSGA-II) to approximate the pareto-frontier of three conflicting objectives: slot rejections, slot displacement and expected flight delays. We test the performance of our approach using slot requests from the Singapore Changi Airport during the 2019 Summer season (April to October). Results confirm the suitability of the proposed methodology to efficiently construct the full pareto frontier of the problem. This paper also provides original and timely insights that can support policy makers and slot allocation stakeholders in adopting better capacity declaration strategies. The results obtained suggest that time-varying declared capacities may reduce avg. hourly delays by 3.5% in relation to a flat declared capacity. Also, combining two time-scales for the declared capacity may reduce avg. hourly delays by 6% in relation to one time-scale capacity.

SB03

Minto

Reinforcement Learning and Its Applications

Invited Session

Applied Probability Society

Chair: Yi Xiong, Shanghai University of Finance and Economics, Shanghai, N/A

1 - Reinforcement Learning for Intensity Control: An Application to Choice-Based Network Revenue Management

HUILING MENG, The Chinese University of Hong Kong, Hong Kong, China, People's Republic of, Ningyuan Chen, Xuefeng Gao

Intensity control is a class of dynamic optimization problems featuring a continuous-time horizon and a discrete state space. Many well-known problems in Operations Research including queueing control and dynamic pricing/assortment in revenue management fall into this category. To apply reinforcement learning (RL) as a computational framework to intensity control, the major obstacle is the continuous-time horizon, as RL is mostly developed for discrete-time MDPs. Naïve uniform discretization of the time horizon is typically used, without a clear understanding of the stability, convergence or computational trade-off of such discretization. In this study, we propose to use the intrinsic discretization provided by the realized jump process and adapt RL algorithms to the continuous-time intensity control problems. By focusing on the choice-based network revenue management as a case study, we theoretically analyze the framework and develop the corresponding policy evaluation (Monte Carlo and TD), policy gradient, and Actor-Critic algorithms. Our approach is computationally efficient and does not have discretization errors compared to naïve discretization when the function approximations are chosen properly. In the numerical experiments, we show the strength of our approach compared to a number of benchmarks in the literature, including choice-based deterministic linear programming and approximate dynamic programming.

2 - Reward-Directed Score-Based Diffusion Models via Q-Learning

Jiale Zha, Chinese University of Hong Kong, Hong Kong, China, People's Republic of, Xuefeng Gao, Xunyu Zhou

We propose a new reinforcement learning (RL) formulation for training continuous-time score-based diffusion models for generative AI to generate samples that maximize reward functions while keeping the generated distributions close to the unknown target data distributions. Different from most existing studies, our formulation does not involve any pretrained model for the unknown score functions of the noise-perturbed data distributions. We present an entropy-regularized continuous-time RL problem and show that the optimal stochastic policy has a Gaussian distribution with a known covariance matrix. Based on this result, we parameterize the mean of Gaussian policies and develop an actor–critic type (little) q-learning algorithm to solve the RL problem. A key ingredient in our algorithm design is to obtain noisy observations from the unknown score function via a ratio estimator. Numerically, we show the effectiveness of our approach by comparing its performance with two state-of-the-art RL methods that finetune pretrained models. Finally, we discuss extensions of our RL formulation to probability flow ODE implementation of diffusion models and to conditional diffusion models.

3 - No Algorithmic Collusion in Blindfolded Games with Thompson Sampling

Yi Xiong, Shanghai University of Finance and Economics, Shanghai, China, People's Republic of, Ningyuan Chen, Xuefeng Gao

When players are engaged in a repeated game with unknown payoff matrices, they may be completely unaware of the existence of others and use multi-armed bandit algorithms to choose the actions, which is referred to as the ``blindfolded game'' in this paper. We show that when the players use Thompson sampling, the game dynamics converges to the Nash equilibrium under a mild assumption on the payoff matrices. Therefore, algorithmic collusion doesn't arise in this case despite the fact that the players do not intentionally deploy competitive strategies. To prove the convergence result, we find that the framework developed in stochastic approximation doesn't apply, because of the sporadic and infrequent updates of the inferior actions and the lack of Lipschitz continuity. We develop a novel sample-path-wise approach to show the convergence.

SB04

Indiana

Behavioral Operations in Markets and Queues

Invited Session

Behavioral Operations Management

Chair: Rihuan Huang, The Chinese University of Hong Kong, Shenzhen, Shenzhen, N/A

1 - Product Selection with Signals in Social Commerce Market: An Experimental Study

Juan Li, Nanjing University, Nanjing, China, People's Republic of, Manqi Li, Wanshan Zhu

In social commerce markets, selecting the right product to sell is a critical decision influenced by signals indicating market popularity. Although prior research (Griffin and Tversky 1992) has shown that human decision makers exhibit biases in interpreting the strength and quantity of these signals. However, it is not obvious whether these biases will inevitably lead to product selection mistakes, by which a less profitable product is selected to sell. To address this question, we develop a behavioral model that identifies the necessary and sufficient conditions under which such mistakes occur. Our follow-up experiment confirms the predictions of the model, showing that product selection mistakes arise specifically when one product has higher signal strength but lower signal quantity compared to another, or vice versa. These findings reveal that social commerce practitioners can mitigate product selection mistakes by strategically gathering more balanced market signals, thereby overcoming the impact of behavioral biases.

2 - Designing Multidimensional All-Pay Procurement Auction with Loss-Averse Bidders

Xianghua Wu, University of New South Wales, Sydney, Australia, Shan Li, Kay Yut Chen

Procurement auctions involving multidimensional bids, such as prototypes and prices, are are illustrated in examples of competitive prototyping in government defense acquisitions. A key feature of these auctions is the all-pay component, where bidders must invest upfront to develop a prototype, regardless of whether they win the contract. If a bidder loses, the cost of the prototype is unrecoverable. Loss aversion, where individuals are more sensitive to losses than to gains of equal size, plays a significant role in this context. We use game theoretical analysis to examine the effects of bidder loss aversion, finding that it consistently reduces prototype quality in the equilibrium. The impact on the buyer’s expected utility depends on the extent of loss aversion and how the buyer values the bidder’s prototype. To address this, we explore two compensation policies—flat and proportional—often used in practice. Both policies improve social welfare without harming either party, with the flat compensation policy being more effective when loss aversion is low. We also examine the role of competition, suggesting that limiting the number of bidders can be beneficial. Our findings indicate that implementing a compensation policy, combined with restrictions on bidder participation, can increase buyer value and improve auction outcomes.

3 - Design of Wait TIME Intervals in Virtual Queues

Yiming Zhang, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Qiuping Yu, Yong-Pin Zhou

This paper examines the impact of displayed wait time intervals on customer abandonment in virtual queues. Through a field experiment with a ride-sharing platform, we analyze standard, LB-extended, and UB-extended intervals, uncovering mechanisms behind treatment effects and offering insights for firms to optimize wait time displays.

4 - To Keep Prices Consistent OR Not: Multi-Channel Retailing with Consumers' Fairness Concerns

Yumeng Li, Shanghai University of Finance and Economics, Shanghai, China, People's Republic of, Xiaomeng GUO, Guang Xiao, Wenxin Xu

In this talk, we study a stylized game-theoretic model to investigate a multi-channel retailer's optimal pricing strategy in the presence of consumers' fairness concerns. We investigate the impact of such concerns on the firm's pricing strategy, profitability, and consumer surplus. Our results can provide useful managerial insights for mult-channel retailers to adopt proper pricing strategies when facing fairness-concerned consumers (This research is supported by the Research Grants Council of Hong Kong [GRF Grant: PolyU 15505621]).

SB05

Orchard A

Advances in Operations and Decision Analytics

Invited Session

Decision Analysis Society

Chair: Xingyu Fu, University of New South Wales, Sydney, Australia

1 - *Algorithmic Transformation to Supermodularity with Provable Distance Reduction: Theory, Simulations, and Applications in Decision Analytics*

Hao Li, University of Science and Technology of China, Hefei, China, People's Republic of

This paper presents a groundbreaking algorithm designed to transform observed matrices into supermodular equivalents while minimizing the distance to their original structure, thereby enhancing decision-making in operations and analytics. Supermodularity, a mathematical property that captures synergistic interactions among variables, is pivotal for optimizing interdependent decisions in resource allocation, scheduling, and network design. However, real-world matrices often lack inherent supermodularity due to measurement noise or incomplete data, limiting their utility in prescriptive models. Addressing this challenge, our algorithm projects arbitrary matrices onto the space of supermodular matrices using a convex optimization framework that prioritizes proximity to the original data, ensuring interpretability and operational fidelity.

Central to this work is the integration of a distance metric (e.g., Frobenius norm) with supermodularity constraints, enabling systematic trade-offs between structural accuracy and complementarity-driven benefits. We demonstrate through extensive simulations that the algorithm significantly reduces computational complexity in downstream optimization tasks while preserving critical decision features. Case studies in multi-echelon inventory optimization and healthcare staff scheduling illustrate that transformed matrices achieve up to 30% improvement in cost efficiency and robustness compared to non-supermodular baselines.

2 - The Dynamics of Player Engagement in Online Games

Yifan Jiao, University of Macau, Macau, Macau, Christopher Tang, Jingqi Wang

This study investigates the dynamics of player engagement in online games and examines how different key factors affect its dynamics. We focus on two primary types of factors: in-game performance and promotional activities. To model these dynamics, we develop a hidden Markov model and estimate it using individual-level gameplay data. Our results reveal that extreme in-game performance—whether exceptionally good or poor—negatively impacts long-term engagement for players in all states. Promotional activities primarily influence players in a medium engagement state. Specifically, incentives that encourage gameplay reduce long-term engagement for these players, whereas promoting social play (i.e., playing with friends) has a positive effect. Price discounts, however, show no significant impact on engagement dynamics. These findings offer valuable

insights for designing targeted strategies to improve player satisfaction and retention in online gaming environments.

3 - Aiming for All: Optimal Design of Reward in Word-of-Mouth Marketing

tianzhuang xu, The Hong Kong University of Science and Technology, Hong Kong, China, People's Republic of, Boer Ouyang, Ying-Ju Chen

Word-of-mouth (WoM) has emerged as a critical marketing channel, enabling firms to acquire new customers almost exclusively through social referrals—thus bypassing cumbersome traditional campaigns. By providing rewards for successful referrals, companies incentivize their customers to exert valuable effort promoting products to friends and acquaintances. However, the need to spur active participation while maintaining a reasonable budget poses a complex optimization problem.

We model the design of optimal WoM reward program as an optimization problem and provides solutions to it. One key takeaway is that contrary to the common belief of prioritizing well-connected individuals (such as Key Opinion Leaders), the optimal strategy often involves ensuring higher referral efforts among customers with fewer social ties. Furthermore, we introduce a mechanism that overcomes the firm's lack of information about each customer's connections: By providing a menu of reward plans, the menu will induce customers to self-report their number of connections, so that firms can effectively implement the same first-best outcome.

Lastly, we benchmark this optimal design against reward programs that are widely adopted in practical words. We find that these popular approaches suffer from substantial inefficiencies because they fail to differentiate incentives based on a customer's number of connections. Our results highlight the significant gains that arise when firms tailor rewards more precisely, unlocking WoM's full potential while managing costs effectively.

4 - Innovation Exclusivity In a Supply Chain

Kanglin Chen, Southern University of Science and Technology, Shenzhen, China, People's Republic of

Manufacturers often collaborate with suppliers to develop innovative products but face a critical decision when structuring supply agreements: whether to require full, partial, or no exclusivity on the supplier's future distribution of the innovation. This study investigates how manufacturers should design exclusivity policies to balance innovation incentives and profitability. We model a supply chain consisting of a manufacturer (he) and a supplier (she), where the manufacturer provides essential resources (e.g., capital, prototypes) but relies on the supplier's technological expertise (e.g., specialized equipment, knowledge) to develop the innovation. The manufacturer determines an exclusivity period during which the supplier must supply the innovation exclusively to him before selling it elsewhere. Our analysis reveals that the optimal exclusivity policy can take three forms: (1) full exclusivity – the supplier cannot sell the innovation to others; (2) partial exclusivity – the supplier is restricted only for a limited time; (3) non-exclusivity – the supplier is free to supply the innovation to competitors. The key trade-off lies in balancing two effects: (1) profit retention – longer exclusivity increases the manufacturer's share of innovation profits; (2) innovation incentives – allowing the supplier to earn additional profits (e.g., through external sales) boosts their investment, raising the likelihood of successful innovation. Interestingly, the optimal exclusivity policy does not always tighten as the supplier's outside opportunity grows; instead, it may follow a non-monotonic relationship, depending on how external profits influence the supplier's effort. Our findings provide actionable insights for structuring R&D partnerships while aligning incentives across the supply chain.

SB06

Orchard B

New Technology in Operation Management

Invited Session

Efficiency and Fairness in Supply Chain and Service Operations Management

Chair: Xu Guan, Huazhong University of Science and Technology, Wuhan, N/A

1 - Platform Certification and Consumer Verification

Xu Guan, Huazhong University of Science and Technology, Wuhan, China, People's Republic of

The rise of online platforms has brought public concerns about the quality of products sold on platforms by third-party sellers. In response, many platforms have launched certification programs to endorse selected sellers. Consumers are increasingly gaining access to tools that enable them to independently verify product quality. This raises two key questions: How should platforms design certification programs and set commission rates when consumers have or lack the ability to verify products? Furthermore, how do the two quality assurance mechanisms—platform certification and consumer verification—interact with sellers' price signaling to shape market outcomes? We address these questions using a Bayesian persuasion framework to analyze the platform's information design, accounting for these strategic interplays. Our analysis reveals three key insights: First, when consumers lack verification tools, the platform may intentionally withhold sellers' quality information in its certification design. As the entry cost for sellers to join the platform increases, the platform adjusts the informativeness of certification and the commission rate in a nonmonotonic manner. Second, introducing consumer verification can trigger high-quality sellers to signal quality by distorting prices upward, inducing the platform to implement a fully transparent certification program to mitigate price distortions. Third, although both certification and verification enhance market transparency, they are not necessarily beneficial. Platform certification may harm consumers and society when consumers lack verification tools. Empowering consumers with verification tools can be detrimental to not only the platform and sellers but also consumers and society.

2 - Quality Verification in the Presence of Review Bias

Yaxu Guo, Huazhong University of Science and Technology, Wuhan, China, People's Republic of, Xu Guan, Yinliang Tan, Yuan Jiang, Guangrui Ma

Over the years, companies across the commercial spectrum have embraced quality verification as a badge of product excellence, utilizing it as a reliable means to convey crucial product information to consumers. In the era of social media, consumers turn to online reviews as a relevant source of product information before making purchase decisions, even though these reviews are inevitably biased. This study thus aims to investigate how firms can best manage the interplay between quality verification and reviews, and also determine the firm's optimal timing for obtaining quality verification in the presence of consumer review bias. We consider two product quality verification formats, pre-release and responsive, based on the timing of the quality verification either before or after the consumer reviews are published. Our study shows that under both quality verification formats, the presence of consumer reviews may not always necessarily benefit the firm. We find that in comparison to the scenario without consumer reviews, the firm is worse off under pre-release quality verification format when there is a greater magnitude of review bias because consumers may interpret this as an indication of poor product quality and lower their expectations. With the responsive quality verification format, the firm can also suffer lower expected profits due to the lack of quality verification in the first period. Comparing revenues between the two formats, the firm does not necessarily benefit from delaying verifying the quality of its products after observing the negative consumer reviews, and so it can profit under either format.

3 - Information Sharing and Behavior-Based Personalized Pricing In An Imbalanced Competitive Environment

Dong Mao, Huazhong University of Science and Technology, Wuhan, China, People's Republic of, Xu Guan, Ruixiao Dong

Innovative technologies enable firms to accurately identify customer segments through purchase history, allowing for differentiated prices for repeat and new consumers. These technologies further empower firms with substantial resources to customize services or products tailored to the preferences of repeat consumers, enhancing satisfaction and market competitiveness. In practice, it is common for firms unable to offer personalized products to share information about their old consumers with firms that can provide such services. To understand the insights behind this phenomenon, we develop a two-period dynamic game-theoretic model to analyze the firms' information sharing and pricing decisions. We show that firms will establish the information sharing collaboration if the firm who can provide personalized product has a relatively low or high ability to reduce the mismatch cost. Interestingly, we find that the firm's profitablity may decrease with its ability to reduce the mismatch cost.

4 - To Cross-Licensing OR Not Cross-Licensing: Implications on Investment, Pricing Decisions, and Consumer Welfare

Tang Jianxi, Huazhong University of Science and Technology, Wuhan, China, People's Republic of, Haijun Wang

This study investigates the strategic dynamics of cross-licensing agreements between asymmetric manufacturers in the presence of government subsidies. Using a Hotelling model with sequential decision-making, we analyze how technological efficiency disparities and competition intensity influence firms' technological investments, pricing strategies, and consumer welfare. The model incorporates vertical differentiation through innovation efforts, where a cost-efficient manufacturer (M1) and a less efficient competitor (M2) engage in cross-licensing to leverage complementary technologies while navigating intensified market competition. We find that cost-efficient manufacturer is willing to adopt cross-licensing when the technological efficiency disparities is not large or the initial market competition is relatively small and the less efficient manufacturer always prefers cross-licensing. Besides, we examine the impacts of government subsidy and find that, government subsidies amplify M1's innovation dominance, and harm M2's profitability at moderate subsidy levels due to exacerbated competition.

SB07

Orchard C

Financial Asset Pricing and Computational Finance

Invited Session

FinTech

Co-Chair: George Yuan, Guangzhou, 510000

1 - Deep Learning Approaches of Backward Stochastic Differential Equation and Applications in Financial Asset Pricing

Yufeng Shi, Shandong University , Jinan, China, People's Republic of

Deep learning can effectively overcome dimensional disasters in the numerical methods of Nonlinear partial differential equation and backward Stochastic differential equation, and has become an important research direction of numerical computation in recent years. We study the numerical calculation methods of high-dimensional backward doubly stochastic differential equations and high-dimensional mean field backward doubly stochastic differential equations, in which the deep neural network is introduced as the key step to achieve numerical solution. Based on the numerical method of backward Stochastic differential equation, this paper studies the option data driven g-pricing modeling method, and verifies the model performance with SPX options. At the same time, with the help of deep learning data mining capabilities, a model for mining high-frequency order book price change information was established, and empirical research was conducted using high-frequency order book data.

2 - A Generalized Portfolio Network Architecture to Improve Financial Prediction Model

Jianwu Lin, Tsinghua University and CSIAM FinTech Committee, Shenzhen, China, People's Republic of

In quantitative finance, the conventional workflow
involves predicting stock returns and subsequently utilizing these
predictions to optimize asset allocation to maximize returns and
minimize risk. This approach typifies a predict-then-optimize
problem, which, despite its widespread application, harbors
inconsistencies between the prediction and optimization stages.
Traditionally, prediction efforts focus on minimizing return
prediction error, while optimization seeks to minimize a risk-adjusted
objective function. This dichotomy often results in
sub-optimal asset allocation solutions due to the misalignment
of goals between the two stages.
In this paper, we propose a novel network architecture
that integrates the prediction and optimization processes into
a unified stage, thus harmonizing the objective of minimizing

prediction error with that of optimizing the allocation objective. By incorporating a differentiable optimization layer within the network, our model directly leverages gradients with respect to the optimization objective, enabling a single, cohesive goal that significantly enhances optimization outcomes.
To validate our approach, we employ our architecture on Qlib, an AI-oriented quantitative investment platform and benchmark, to simulate exchange environments and conduct rigorous backtesting. The results demonstrate that our model achieves substantial improvements in key financial indicators, including cumulative returns and average risk-adjusted objective function value, over traditional two-stage models. These findings underscore the efficacy of our integrated approach in addressing the inherent inconsistencies of the predict-then-optimize framework, offering a promising direction for future quantitative finance methodologies.

3 - Distributed Federated Learning and Low-Latency Edge Computing Methods for Privacy Protection in Financial Data Transaction and Computing

Zigan Wang, Tsinghua University, Beijing, China, People's Republic of

The increasing demand for privacy-preserving distributed analytics in financial systems calls for innovative approaches to secure and efficient data processing. This study proposes a novel framework that integrates federated learning with edge computing to enable secure, privacy-aware financial analytics across multiple institutions. Utilizing quarterly financial report data from commercial banks, the framework employs a Byzantine fault-tolerant architecture to facilitate robust cross-institutional data aggregation and privacy-conscious value discovery. Experimental results confirm the system’s effectiveness in distributed financial environments, supporting synchronized and periodic analytics under strict privacy constraints. By enabling secure forecasting of sensitive financial indicators, this framework holds significant potential for regulatory compliance and inter-institutional collaboration. It establishes a scalable foundation for privacy-preserving data analytics in complex, distributed financial scenarios.

4 - Optimize FinTech Solutions with Copt

Tiancheng Zhang, Cardinal Operations, Beijing, China, People's Republic of

Optimization solvers serve as the backbone of modern FinTech research and applications, driving innovation and efficiency across the industry. In this talk, we will share the current state of Cardinal Optimizer (COPT), and explore the latest performance benchmarks and developments, highlighting its advancements in matrix modeling, conic programming, and GPU-accelerated algorithms. Through real-world industrial examples, we will demonstrate how these innovations have empowered practical FinTech use cases, including portfolio optimization, index tracking, and beyond.

SB08

Bras Basah A

Healthcare Payment, Incentives, and Appointment Scheduling

Contributed Session

Healthcare Applications

Chair: Yu Lu, Tsinghua University, Beijing, N/A

1 - Blended Payment Scheme of Nhs England: Issues and Remedies

Sara Mohammadi, UCL School of Management, University College London, London, United Kingdom, Kenan Arifoglu, Bilal Gokpinar

The *blended payment scheme* introduced by NHS England aims to balance financial stability and hospital incentives by combining *fixed and variable payment components*. The fixed component is a lump sum payment made to hospitals at the start of the year, based on planned activity levels and national tariff prices, while the variable component adjusts payments at year-end to account for differences between planned and actual activity using a proportion of the tariff prices (i.e., a variable rate). We develop a *game-theoretic model* to analyze hospitals' treatment costs and activity level decisions under this scheme, considering *information asymmetry* between hospitals and the regulator. We show that, due to information asymmetry, the blended payment scheme fails to fully align hospital incentives, resulting in deviations from socially optimal treatment costs and planned activity levels. Additionally, we demonstrate that *block contracts* and *Payment by Results (PbR)* are extreme cases of the blended payment scheme, with hospitals' treatment costs and planned activity levels under the blended payment model falling between these extremes and aligning more closely with socially optimal outcomes. We propose a new *modified* blended payment scheme that combines a fixed payment, similar to the NHS's blended payment scheme, with a variable payment that adjusts hospital payments based on the average actual activity levels of other hospitals, rather than a hospital's own planned activity levels. Our modified blended payment scheme addresses information asymmetry and aligns hospital incentives with socially optimal outcomes.

2 - Generic Drug Shortages and Undifferentiated Competition: a Game-Theoretic Policy Implication

Seung-chan Choi, Rutgers University, Piscataway, NJ, United States

Shortage of affordable and quality drugs continues to be a global problem. The extant literature indicates several root causes, including commoditized generics, lack of economic supplier incentives, supply chain and inventory glitches, and government regulations. Most studies point to regulatory policy changes to lower the entry barrier for generics, to promote more competition, and to increase manufacturing capacity. However, this paper argues that drug shortages are unavoidable even with a lower entry barrier, increased competition, and increased capacity. Furthermore, some of these policy recommendations may even facilitate drug shortages.

This paper develops a Stackelberg-Cournot game theory model in which the branded drug manufacturer faces multiple generic makers. The number of diagnosed patients fixes the potential market in a therapeutic drug category. The potential market is fully covered if the drugs are available and their prices are affordable (to either the payers or patients). However, we show that the suppliers (both branded and generic) are not incentivized to supply to meet the full demand. We find that the drug shortage grows as the unit variable cost increases and the perceived quality gap between the branded and generic drugs increases. Our results show that policy changes to (1) allow product differentiation among generics (quality or feature), (2) improve the public's perception of generics' quality, and (3) lower the cost of major ingredients will mitigate drug shortages.

3 - Efficient Engagement: Opportunities for Mobile Methadone Maintenance in Geographically Underserved Areas

Anthony Bonifonte, Denison University, Granville, OH, United States, Erin Garcia

Methadone treatment is proven to be highly successful in combating opioid addiction and preventing overdoses, but many patients have geographic restrictions on access. In 2021 the Drug Enforcement Agency of the United States changed regulations to allow for easier operation of narcotic treatment programs with a mobile component, also known as mobile methadone clinics. These mobile clinics can expand access to methadone in rural and underserved communities. In this talk we estimate the demand for methadone in areas that do not currently have access and formulate a MIP to maximize the number of clients served. We solve the model to optimality within each state nationwide, and explore the effects of multi-vehicle routes. We find mobile routes can be effective in providing service to geographically isolated areas. However, additional policy changes need to be considered and efficiency improvements made to incentivize the operation of mobile methadone units for serving rural and unmet demand for methadone maintenance treatment. These findings are generalizable to planning other mobile health services with uncertain demand.

4 - Calendar Design for Assignment of Ongoing Appointments: Modeling, Analysis and Application

Yossi Luzon, AFEKA - The Tel Aviv Academic College of Engineering , Tel Aviv, Israel, Yariv Marmor

Appointment booking (AB) is widely utilized for managing elective services such as hospital surgical units, courtrooms, and other resource-intensive operations. We propose a scheduling method rooted in an appointment booking system designed to address scheduling challenges driven by demand and historical data. Specifically, we introduce the pre-generated appointment scheduling system, referred to as the appointment booking calendar (PAC), which is developed offline using demand forecasts and historical information. The PAC aims to minimize patients' sojourn time—the period from the initial request to the completion of service. Despite significant public expenses associated with long waiting times, AB strategies often neglect sojourn time, affecting patient outcomes and resource efficiency. Our approach directly addresses this gap by generating an offline schedule based on historical demand data. The proposed method also maintains quality of service (QoS), balancing effective resource utilization and service quality. Additionally, we present the chained-PAC (CPAC), a mechanism that connects multiple smaller PACs to manage dynamic demand flexibly, facilitating straightforward real-time adjustments and scalability. We illustrate our method’s effectiveness by implementing the CPAC, where smaller PAC units are sequentially combined. The CPAC method provides practicality, ease of use, and real-time adaptability, making it well-suited for operational application. We demonstrate the applicability of our approach through practical implementation, validating its efficacy in delivering user-friendly appointment scheduling solutions that enhance resource utilization and reduce patient waiting times, resulting in a robust, adaptable AB system.

5 - Location Planning, Resource Reallocation and Patient Assignment During a Pandemic Considering the Needs of Ordinary Patients

Yu Lu, Tsinghua University, Beijing, China, People's Republic of, Shaochong Lin, Zuo-Jun Shen, Junlong Zhang

During the initial phase of a pandemic outbreak, the rapid increase in the number of infected patients leads to shortages of medical resources for both pandemic-related and non-pandemic (ordinary) patients. It is crucial to efficiently utilize limited existing resources and strike a balance between controlling the pandemic and sustaining regular healthcare system operations. To tackle this challenge, we introduce and investigate the problem of optimizing the location of designated hospitals, reallocating beds within these hospitals, and assigning different types of patients to these hospitals. Designated hospitals isolate pandemic-related patients from ordinary patients to prevent cross-infection. Moreover, isolation beds can be converted into ordinary beds and vice versa. Considering the stochasticity and evolving nature of the pandemic, we formulate this problem as a multi-stage stochastic programming model, integrating a compartmental model with time-varying random parameters to enable dynamic resource allocation as the pandemic progresses. The model is then solved by a data-driven rolling horizon solution approach. We illustrate the effectiveness of our model using real data from the COVID-19 pandemic. Compared with two other approaches, our model demonstrates superior performance in controlling the spread of the pandemic while addressing the needs of both pandemic-related and ordinary patients. We also conduct a series of experiments to uncover managerial insights for policymakers to better utilize existing resources in response to pandemic outbreaks. Results indicate that admitting as many pandemic-related patients as possible during the initial phases of the outbreak can effectively flatten the pandemic peaks and alleviate strain on the healthcare system.

SB09

Bras Basah B

Supply Chain Finance, Inventory Management, and Asset Sustainability

Contributed Session

Inventory Management

Chair: Yinuo Lin, N/A

1 - Inventory Agility, Environmental Dynamism, and Supply Chain Relationship

Shih-Sian Jhang, National Sun Yat-sen University, Kaohsiung, Taiwan, Lin-Chin Wu, SHAO-CHI FANG

This paper aims to investigate inventory agility and its association with external environments and supply chain relationships. The specific research questions addressed in this study are: (1) Do firms adjust inventory toward a target level? (2) How do environmental dynamism and supply chain relationships affect inventory agility? The

data sources in this study include secondary data and hand-collected data. The secondary data include Compustat-Capital IQ and I/B/ES; both are provided by Wharton Research Data Services. The hand-collected data is related to supply chain relationship data, collected via the website of the U.S. Securities and Exchange Commission. Our sample focuses on US public firms for the years 1980-2023. Research methods include firm (year) fixed effects models, instrumental variables, two-stage least squares, etc. As an interdisciplinary research effort, this study offers a new perspective to integrate operations management, supply chain management, corporate finance, and economics.

2 - Use the Mix of Purchase Order Financing and Factoring to Support a Small Enterprise

Andy Wu, National Chung Hsing University, Taichung, Taiwan

This study examines the combined use of Purchase Order Financing (POF) and Account Receivable Financing (ARF) as a hybrid supply chain finance (SCF) strategy, particularly for SME suppliers. While ARF and its evolution into Reverse Factoring have become common in major economies and increasingly adopted in Taiwan, POF remains underdeveloped locally. Recent collaborations between major Taiwanese enterprises and banks have started implementing these tools through tri-party agreements.

Existing SCF literature largely overlooks the joint application of POF and ARF. To address this gap, we analyze real-world cases and develop a quantitative decision model to explore how suppliers can optimize pricing, payment terms (e.g., net 30), and order quantities when financing production via both tools. The model incorporates financing cost variability, cash flow timing, and fulfillment risk.

Findings reveal that pricing and credit terms are interrelated: when price concessions are limited, suppliers may extend payment periods to attract buyers. This impacts order size, financing costs, and profitability. We demonstrate that under certain conditions, combining POF and ARF reduces total financing costs and provides strategic flexibility in managing cash flow. Our contribution lies in offering a novel framework for integrating financing strategies into supply chain decision-making.

3 - Inventory Disclosure for a Capital-Constrained Retailer: Impact of Financing Mode

ying wei, Jinan University, Guangzhou, China, People's Republic of

In a dual-channel supply chain consisting of a capital-constrained store retailer and an online manufacturer, we intend to study the impact of financing mode on the store inventory disclosure policy. Customers strategically make their first channel choices based on their beliefs on the store stock availability and turn to the other once stockout occurs; or exit. Sellers in different channels then compete with their stock levels. With a game-theoretical model, we investigate the impact of store inventory disclosure on the channel selection, the channel stocks and the supply chain participants' profits, as well as the impact of financing mode on inventory disclosure policy. One salient finding that differs with the existing work is that the capital constraint of the retailer deters store sales through both the supply side and demand side, so-called *negative supply effect* and *negative demand effect*. However, a manufacturer-provided credit guarantee can resolve the supply constraint and even incentivize more consumers to visit the store, thus incurring *positive demand effect*. In addition, the manufacturer may voluntarily reduce the online inventory to mitigate the channel competition and potential losses incurred under retailer bankruptcy. Interestingly, inventory disclosure can eliminate the demand effect of capital constraint by revealing the store stock information to the consumers, which induces the B&M retailer to stock more and the manufacturer to stock less, regardless of the financing modes. Both manufacturer and store retailer benefit from inventory disclosure under the bank financing mode, whereas the store retailer may get worse under the credit guarantee mode.

4 - Automated Detection of Internal Defects in Maritime Shipping Containers Using Machine Learning

TZU-YUAN LIAO, National Cheng Kung University, Tainan, Taiwan, PEICHUN LIN, Jenhung Wang

Maritime shipping container transportation, characterized by high security, large capacity, and low cost, constitutes a crucial component in the global logistics supply chain. However, internal defects in containers may lead to cargo damage, financial losses, increased logistics costs, and reduced container lifespan. Traditional container inspection methods require substantial manual labor, are time-consuming, potentially hazardous, and

subject to subjective interpretation, resulting in inconsistent assessments. Moreover, with widespread labor shortages at ports worldwide, addressing workforce challenges through AI-enabled smart port development represents a future trend. This research employed a machine learning platform that operates without coding knowledge (Teachable Machine®) to train and test automatic identification of internal defects in maritime shipping containers. For subsequent application in robotic inspection systems or mobile devices, further exploration of model compression techniques (Pruning, Quantization) and Knowledge Distillation will be necessary to optimize the model for real-time computation. Integration with image capture equipment installed on cranes or Automated Guided Vehicles (AGVs) will enable real-time image uploading to cloud or edge computing platforms for analysis. The development of a database linking historical maintenance records, transportation routes, loading/unloading frequency, and climate information for each container will facilitate prediction models for defect frequency and maintenance costs, evolving into a Preventive Maintenance system. Following model deployment, continuous collection of misclassification cases from operational environments will inform Incremental Learning processes, allowing ongoing model optimization and accuracy maintenance. Accumulated defect detection results will enable advanced AI analysis investigating defect causation prediction, container lifespan estimation, and dynamic maintenance cycle adjustments.

5 - Robust Multi-Product Inventory Replenishment

Yinuo Lin, National University of Singapore, Singapore, Singapore, Xiaobo Li

This paper investigates a multi-period inventory management problem involving a single replenishment of multiple products in a capacity-constrained warehouse, motivated by the need to prepare for high and intensive demand over a short period. The problem is initially formulated as a stochastic program. We first consider a deterministic version with known future demand and develop an efficient solution algorithm. This algorithm is then extended to the stochastic setting using the Sample Average Approximation (SAA) method.

To further address demand uncertainty, we introduce a partial-sum uncertainty set and reformulate the problem as a robust optimization model. We show that the number of worst-case demand scenarios grows linearly with the product of the lead time and the planning horizon. This insight allows us to reduce the robust model to a tractable linear program with a polynomial number of constraints, solvable using the same algorithm developed for the deterministic case.

Additionally, we explore the optimal scheduling of multiple replenishments under deterministic demand. A case study using real-world data from RiRiShun Logistics demonstrates the practical effectiveness of the proposed robust optimization framework.

SB10

Hullet Room

OR in Agriculture

Invited Session

OR and Social Wellness

Chair: Xujin Pu

Co-Chair: Yaping Fu

1 - Integrated Harvest and Distribution Scheduling of Fresh Agricultural Products for Multiple Farms Using a Q-Learning-Based Artificial Bee Colony Algorithm with Problem Knowledge

xiaomeng ma, Jiangnan University, Wuxi, China, People's Republic of, Xujin Pu, Yaping Fu

Nowadays, a great many of farmers sell fresh agricultural products through direct online sales. In this context, the farm-to-door supply mode has emerged, playing a crucial role in reducing transportation cost and quality deterioration. This work addresses an integrated harvest and farm-to-door distribution scheduling problem involving multiple farms. First, a mixed integer programming model is formulated to minimize total operation cost and maximize customer satisfaction regarding product quality. Second, a Q-learning-based artificial bee colony algorithm with problem knowledge (Q-ABC-K) is developed in particular. The algorithm is featured with the following strategies: (i) a hybrid initialization method with two rules to generate a highquality population;

(ii) a crossover operation to prompt a collaborative search between the population and external archive at the employed bee phase; (iii) a Q-learning method to favorably select premium neighborhood structures at the onlooker bee phase; and (iv) a knowledge-based local search method to refine the nondominated solutions. Finally, a large number of comparison experiments are conducted on a set of test instances. Through observing and analyzing the experimental results, three conclusions are acquired as follows: (i) The design of Q-learning and knowledge-based local search methods plays a significant role in enhancing the performance of Q-ABC-K; (ii) QABC-K performs better than four state-of-the-art approaches in dealing with the considered problem; and (iii) Q-ABC-K has an advantage over an exact solver CPLEX in solving small-scale cases.

2 - Product selling, social utility and virtual tips in live streaming

Mingzhuo Dai, Huazhong University of Science and Technology, Wuhan, China, People's Republic of

3 - The Impact of Community Supported Agriculture Contract upon Promoting High-Value Agricultural Products in Developing Economies

Shizhen He, Xiamen University, Xiamen, China, People's Republic of, Jiguang Chen, Xujin Pu

In the ever-growing farming practice of community supported agriculture contracts in developing economies, consumers are allowed to buy claims on a farm’s future production. For example, in this type of practice called "adopting a vine", before the production the consumer pays to buy claims on a high-quality vine from the producer, and receives all the realized output yield of this vine after the producer's effort. Motivated by these farming operations, this paper formalizes the impact of the community supported agriculture contract upon promoting high-value agricultural products (such as grape with more freshness) competing with ordinary products (such as grape after a long distance transportation of about 4,000 kilometers) , we prove that the producer will prefer the community supported agriculture contract for high-value agricultural products only when the price in the external market of ordinary products is neither too low nor too high.

4 - Seeing is Not Believing: Popularity Manipulation in Live-Streaming Reward

Dongliang Guo, Northeastern University, Shenyang, China, People's Republic of, Zixuan Meng, Zhe Zhang, Zhi-Ping Fan

Influencers on live-streaming platforms leverage their popularity to attract viewers, who come to the influencers’ live-streaming rooms and voluntarily reward them with virtual gifts. Gift-rewarding has emerged as a key monetization model in the influencer economy, generating substantial revenue within the live-streaming industry. To maximize these rewards, many influencers manipulate their popularity to entice viewers into their live-streaming rooms, posing significant challenges to the platform's reputation and revenue. This paper develops a game-theoretic model to analyze an ecosystem comprising a live-streaming platform and competitive influencers with heterogenous inherent abilities. We explore their joint decisions regarding popularity manipulation and self-improvement efforts and assess how the platform's detection policy affects equilibrium outcomes. Our findings reveal different incentives for popularity manipulation among influencers, based on their differing abilities. We identify two types of popularity manipulation strategies—offensive and defensive—that impact their self-improvement efforts. Interestingly, our analysis shows that the platform does not always benefit from detecting popularity manipulation; its detection policy may backfire, reducing viewer surplus and social welfare. To address this inefficiency, we propose a popularity-specific commission policy as a complement or substitute tool to the platform's detection policy. This policy has the potential to Pareto improve social welfare, ensuring that no party is worse off while some are better off.

SB11

Morrison Room

Frontier in Revenue Management & Pricing

Invited Session

Revenue Management and Pricing

Chair: Xingyu Fu, University of New South Wales, Sydney, Australia

Co-Chair: Pin GAO, The Chinese University of Hong Kong, Shenzhen, Shenzhen

1 - Review Image Quality and Business Performance: Evidence and Theory

Chong Zhang, Tilburg University, Tilburg, Netherlands, Jiawei Chen, Yunke Mai

This paper examines the impact of review image quality on business performance. Utilizing panel datasets of Yelp, we empirically analyze the relationship between image quality, business ratings, and sales. Our findings reveal that while high-quality images significantly enhance the performance of businesses with high ratings, they may adversely affect those with low ratings. To explain these contrasting observations, we develop an analytical model that considers the varying degrees of customer sophistication in interpreting visual information. The analytical predictions replicate the contrasting empirical observations. We further validate our theory by conducting additional empirical analysis using our Yelp datasets and show that the main effects are indeed moderated by relevant measures that are potentially correlated with customer sophistication. These investigations provide a thorough analysis of the impact of review image quality and offer strategic insights for business owners to effectively manage their online presence on review platforms.

2 - Meta Dynamic Pricing with Nonparametric Empirical Bayes

Jinhui Han, Peking University, Beijing, China, People's Republic of, Ming Hu, Yunzong Xu, Xilin Zhang

We investigate a multi-product pricing problem in which a retailer or platform sells N products. Leveraging a nonparametric empirical Bayes approach, we establish that the regret bound grows at most polynomially in the logarithm of the number of products, even when the available data for each product is limited. Building on this insight, we further consider a dynamic setting and propose an online algorithm that achieves a regret bound with the same dependence on the number of products and the standard square-root dependence on the time horizon.

3 - Strategic Data Analytics Under Supplier Encroachment

Yonghui Chen, University of Liverpool, Liverpool, United Kingdom, Guangrui Ma, Qiaochu He, Ying-Ju Chen, Zuo-Jun Shen

To enhance supply chain collaboration, major retailers increasingly grant suppliers access to their data platforms, leveraging data analytics to deliver valuable demand insights. However, this access introduces the risk of channel encroachment by suppliers. This paper examines a retailer’s strategy for deploying data analytics in the presence of supplier encroachment. We analyze two approaches: a shared data platform, which delivers identical demand predictions to both the retailer and the supplier, and a separate data platform, which provides accurate predictions for the retailer while offering less precise insights to the supplier. On a shared data platform, we find that the retailer adopts perfectly accurate data analytics when the supplier’s encroachment cost is either very high or low. When the encroachment cost is moderate, the shared data platform generates informative yet distorted demand forecasts—either upward or downward—which surprisingly benefits both firms due to the ease of double marginalization. On a separate data platform, informative predictions are generated for the supplier only when its encroachment cost is either sufficiently low or high. By comparing these two platform types, we find that the shared platform becomes the retailer’s optimal choice when the encroachment cost is sufficiently high. Our findings highlight the complementary roles of strategic data analytics and channel encroachment in supply chain digital transformation. With aligned demand forecasts, retailers and suppliers can achieve “win-win” outcomes, explaining the widespread adoption of shared data platforms.

4 - Redirecting Consumer Demand: Managing Product Substitutions with Coupons

Wentao Lu, Carey Business School, Johns Hopkins University, baltimore, MD, United States, Guillermo Gallego

When consumers cannot locate their preferred product, they frequently opt for an available substitute. Identifying how demand changes between products is crucial in discrete choice modeling, revenue management, and assortment optimization. We tackle the challenge of managing demand substitution by leveraging coupons within discrete choice models. Specifically, we explore how coupons can influence consumer decisions to amplify expected revenues. We seamlessly integrate coupons into the discrete choice modeling framework and examine the associated assortment optimization problem. Our insights reveal that the benefits of employing coupons can be arbitrarily large compared to optimizing the assortment without coupons. When the choice model adheres to the multinomial logit (MNL) model, we illustrate that optimal assortments can be identified in polynomial time. Regarding the coupon pricing problem, we discover that all products exhibit the same mark-up at optimality for the MNL, with coupons serving as a price discrimination mechanism. In contrast to assortment

optimization, we ascertain that the optimal profit ratio with coupons can be at most double the optimal profits without coupons when firms have the flexibility to adjust prices.

SB12

Moor Room

Sharing Economy: Challenges and Opportunities

Invited Session

Revenue Management and Pricing

Chair: Sixing HU, National University of Singapore, Singapore, Singapore

Co-Chair: Peihua Mai, National University of Singapore, Singapore, N/A, Singapore

1 - Vehicle Relocation with Location-Based Compensation

Sixing HU, National University of Singapore, Singapore, Singapore, Bomin Bian, Zhenyu Hu, Long He

We study a vehicle pre-allocation problem for an on-demand platform that earns commissions from completed trips. To better align idle vehicles with future demand, the platform uses location-specific monetary transfers to influence drivers' relocation decisions instead of directly assigning vehicles. Drivers' choices are modeled using an endogenized multinomial logit (EMNL) model that captures competition through the probability of matching with a request. We show the uniqueness of the resulting equilibrium and interpret transfers as adjustments to reach a target expected compensation. The platform's profit maximization becomes independent of the commission rate, and the optimal compensation equalizes an adjusted markup metric across locations. The problem can be efficiently solved as an exponential cone program. Case studies using real-world data from a taxi operator in Singapore demonstrate the effectiveness of monetary control in guiding vehicle pre-allocation.

2 - Vertical Federated Graph Neural Network for Recommender System

Peihua Mai, National University of Singapore , Singapore, Singapore, Yan Pang

Conventional recommender systems are required to train the recommendation model using a centralized database. However, due to data privacy concerns, this is often impractical when multi-parties are involved in recommender system training. Federated learning appears as an excellent solution to the data isolation and privacy problem. Recently, Graph neural network (GNN) is becoming a promising approach for federated recommender systems. However, a key challenge is to conduct embedding propagation while preserving the privacy of the graph structure. Few studies have been conducted on the federated GNN-based recommender system. Our study proposes the first vertical federated GNN-based recommender system, called VerFedGNN. We design a framework to transmit: (i) the summation of neighbor embeddings using random projection, and (ii) gradients of public parameter perturbed by ternary quantization mechanism. Empirical studies show that VerFedGNN has competitive prediction accuracy with existing privacy preserving GNN frameworks while enhanced privacy protection for users' interaction information.

3 - Promoting Mobility Equity in Micromobility Sharing Services

WU HAO, National University of Singapore, Singapore, Singapore, Long He, Sheng Liu

This study examines how equity-oriented policies influence service outcomes in shared micromobility systems. We develop a stylized model in which a profit-maximizing operator makes operational decisions between two regions, first without any equity policy and then under regulatory constraints designed to promote equitable access. The analysis identifies when inequities arise, compares the impact of different policy approaches, and considers how regulators might influence operator decisions through policy design. To assess practical relevance, we apply this framework to real-world cases, embedding city-specific rules and demand patterns into a set of what-if simulations. These ongoing experiments evaluate the effectiveness and trade-offs of various policy levers under realistic operational settings. By bridging theoretical insights with empirical considerations, this work provides a structured foundation for assessing and improving equity outcomes in urban shared micromobility systems.

4 - The Strategic Secretary Problem

Yufei Zhu, National University of Singapore, Singapore, Singapore, Joel Goh

In the classical secretary problem, secretaries have no decision-making agency and are assumed to arrive in a random order. In this paper, we consider a novel variation of the classic problem. Motivated by modern online labor platforms, where job-seekers have access to much more information, we assume that secretaries are strategic agents and can choose their arrival time based on a privately revealed quality score. We analyze a game between the hiring firm and $N+1$ strategic secretaries, focusing on the structure of the equilibria. This paper addresses the following questions: Do strategic secretaries gain more utility than non-strategic ones? Does the firm's ability to select the best secretary improve or deteriorate when secretaries act strategically?

SB13

Bailey

AI and ESG in Service Operations

Invited Session

Service Science

Chair: Kejia Hu, University of Oxford Said Business School, Oxford, United Kingdom

Co-Chair: Yirui Jiang, Cranfield University, Bedford, United Kingdom

1 - A Four-Level Framework for Generative AI in Operations Management: Empirical Insights into Automation-Augmentation Dynamics and Value Creation

Jasper Shi Da Xu, Waseda University, Tokyo, Japan, Kejia Hu, Teng Wang

Generative Artificial Intelligence GAI is revolutionizing business and societal operations at an unprecedented pace. This study proposes a novel four-level framework to empirically examine the impact of GAI on operations management (OM) through controlled laboratory experiments. The framework comprises Process Automation L1, Product Personalization L2, Operational Innovation L3, and Market Creation L4. Using Epitome, an advanced LLM-driven social science experiment platform, the study will progressively introduce AI capabilities across OM scenarios such as call center operations, inventory management, and medical scheduling. This study will also embed insights from the Automation-Augmentation Paradox, which highlights the interdependence between automation and augmentation across time and space. At each level, human roles will be explicitly integrated to explore their influence on AI outcomes. For example, humans will oversee automated processes L1) to mitigate biases, refine AI-driven personalized recommendations L2, co-create innovative solutions with AI L3, and govern ethical implications in AI-driven market creation L4. Experiments will begin with manual task execution L0) and advance through automation and augmentation cycles, emphasizing the dynamic interplay between efficiency and adaptability. Key performance indicators - task efficiency, customer satisfaction, innovation rate, and decision confidence - will be measured to evaluate GAI's transformative impact. Preliminary expectations suggest that while automation enhances efficiency at lower levels, augmentation fosters creativity and strategic decision-making at higher levels. By empirically validating these dynamics, the study aims to contribute to OM theory by reframing AI adoption as a paradoxical tension requiring balanced differentiation and integration of automation and augmentation for sustainable value creation across multiple contexts.

2 - Leveraging Pre-Sales for Logistics Efficiency: when is It a Valuable Strategy?

RUIJIE ZHANG, ShanghaiTech University, Shanghai, China, People's Republic of

This paper examines the pre-sales model in which customers first pay a deposit, termed a down payment, and then wait several days or weeks to make the final payment. The retailer manages the shipment upon receiving the final payment. While the down payment offers a discount to customers, it also allows retailers to manage demand and pre-plan logistics more effectively. To highlight the associated benefits, this study focuses on the logistics cost reduction role of pre-sales for in-stock durable products in coping with bottlenecked demand surges during promotion seasons. We first propose a joint optimization model for the down payment and the product's price. In the analysis, we derive the value of introducing the pre-sales channel for the retailer and identify the conditions that make adding a pre-sales channel beneficial. Furthermore, we provide a guideline for the retailer to determine the optimal pre-sales duration. Subsequently, we investigate thoroughly the effect of

various parameters such as logistics cost and customer's waiting sensitivity on the total profit. Our results aid online retailers in understanding the conditions under which adding a pre-sales channel is beneficial. Moreover, the findings help to decide the appropriate duration of pre-sales. In the case study, we calibrate our model using real data from an e-commerce company in China and highlight the importance of considering the production cost for deriving the linear demand function. The results show that our model can predict demand well. By applying the strategies proposed in this paper, retailers can drastically improve profits.

3 - Rethinking Appointment Scheduling: Addressing Counselor Burnout and Improving Patient Outcomes

Yuchen Liang, National University of Singapore, Singapore, Singapore, Sarah Yini Gao, Kejia Hu, Chung Piaw Teo

The growing demand for mental health services has placed significant pressure on counseling clinics, where counselors must balance both task load and emotional load while delivering effective care. This study explores how these stressors influence counselors' scheduling decisions, particularly in terms of follow-up appointments, lead times, and their impact on patient no-show rates. Using data from a university counseling clinic in Singapore, we employ regression models to examine the relationship between counselor task load, emotional load, and scheduling outcomes. Our findings show that increased task and emotional load raise the likelihood of scheduling follow-up appointments but with longer lead times, leading to more treatment sessions, patient disengagement, and higher no-show rates. To address these inefficiencies, we propose a scheduling policy that balances task and emotional load, improving clinic performance while maintaining counselor well-being. A discrete-event simulation (DES) study demonstrates that this policy reduces care access times by 4.98% and patient pathway lengths by 2.92% compared to current scheduling policies. These findings have important theoretical and practical implications, highlighting the need to incorporate service providers' emotional load into service operation models. They also offer a practical solution for improving counselor welfare and patient outcomes in high-emotional labor environments.

4 - Cracking the Code of Livestream Success: Impacts of Domain Knowledge and Sales Knowledge

Demei Kong, Tsinghua University, Beijing, China, People's Republic of, Kejia Hu, Yongbo Xiao

Livestreaming commerce has garnered significant recognition due to its unique advantage in facilitating real-time and interactive engagement between streamers and online viewers. The success of individual livestreams hinges on various factors, notably the streamer's domain knowledge and sales expertise showcased during the live-shows. These elements are particularly vital for selling non-standardized products, as product-related knowledge fosters trust among viewers, while effective sales skills encourage prompt purchasing decisions. How to optimize the integration of domain knowledge and sales expertise within livestream scripts is crucial for enhancing overall throughput. This paper explores which factor is more critical in the live-selling of non-standardized products, leveraging Dual-Process Theory to conceptualize sales expertise as activating Type 1 (automatic) cognitive processing, and domain knowledge as encouraging Type 2 (deliberative) processing. By analyzing livestream videos and corresponding sales records of jade products on Douyin using Large Language Model-powered video analysis, we empirically identify several key insights. Our findings reveal: both types of knowledge effectively boost sales, but sales expertise exhibits a stronger impact; extensive sales knowledge significantly increases the return rate of jade products, whereas domain knowledge does not have this effect; sales expertise is most impactful at the beginning and end of a live-show, while domain knowledge is more effective during the middle phase; the efficacy of different types of communication is moderated by streamer credibility, viewer engagement, and product price. These results highlight a pivotal trade-off in livestream persuasion: while persuasive tactics can enhance immediate sales, they may also elevate product returns.

5 - Corporate Social Responsibility: Selling and Donating Strategies Across Markets

Xinyue Cai, Tsinghua University, Beijing, China, People's Republic of, Yongbo Xiao, Tony Cui, Chen Hu

With the growing emphasis on common prosperity and poverty alleviation, it is increasingly common for firms to demonstrate concerns on the well-being of consumers in underdeveloped regions, often referred to as the ``Bottom of the Pyramid (BoP)'' markets. By implementing markdowns and/or donations, firms are able to improve the well-being of BoP consumers through enhanced product accessibility, which in turn increases the willingness-to-pay of consumers in more developed ``primary markets'' due to the so-called ``warm-glow effect.'' In this paper, we examine the profit potential of implementing social responsibility in the presence of prosocial

consumers. More specifically, we study the impact of consumers' prosocial behavior (as reflected by the extent of warm-glow sensitivity) on a firm's optimal pricing and donation decisions. We uncover an intricate interplay between these two strategies, where donations can either complement or substitute price markdown. With a moderate warm-glow sensitivity, the firm may even charge a low price together with donations. Surprisingly, our findings reveal that primary consumers' prosociality may inadvertently hurt consumers in both markets. We also extend the model to capture consumers' fairness concerns and strategic behavior, and examine how they may influence the firm's decisions to enhance BoP product accessibility.

SB14

Wilberforce

Generative AI for Optimization

Invited Session

Simulation Society

Chair: Yun Hua, Shanghai Jiao Tong University, Shanghai, N/A

Co-Chair: Jun Luo, Shanghai Jiao Tong University, Rm 1216, Bao Zhaolong Building, Add. 1954, Huashan Rd., , Shanghai, 200030

1 - Bench4opt: Benchmarking Large Language Models for Optimization Modeling with Graph-Based Equivalence Evaluation

Ziwei Zhu, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Zhuohan Wang

In operations research (OR), formulating optimization problems for industrial applications is often labor-intensive and demands specialized domain expertise. Recent advancements in large language models (LLMs) have demonstrated significant potential to automate this task. However, evaluating the performance of LLMs in optimization modeling remains challenging due to the lack of comprehensive datasets and robust evaluation frameworks. To address this gap, we introduce Bench4Opt, a new benchmark designed to systematically assess the capability of LLMs in formulating linear programming (LP) and mixed-integer linear programming (MILP) models. Bench4Opt features a diverse dataset comprising 818 optimization modeling problems, presented in a model-data separation format with two levels of description abstraction, spanning 16 problem classes and over 40 practical domains. Additionally, Bench4Opt incorporates an innovative evaluation methodology based on a modified Weisfeiler-Lehman graph isomorphism test (WL-test) algorithm. We provide theoretical proof demonstrating that this method can accurately determine the equivalence of two optimization models. Using Bench4Opt, we evaluate the performance of several state-of-the-art LLMs, including GPT, LLaMA, DeepSeek, and Claude, revealing significant variations in their capabilities. GPT-4o and DeepSeek-V3, when directly prompted, achieve overall accuracies of 50.49% and 46.94%, respectively, outperforming all other models. Notably, all models exhibit substantial performance variability across problem classes. For example, GPT-4o achieves over 90% accuracy on the knapsack problem class but drops below 5% on the traveling salesman problem class. These results offer critical insights into the strengths and limitations of LLMs in optimization modeling, guiding future developments in automatic optimization modeling.

2 - Learning Virtual Machine Scheduling in Cloud Computing Through Language Agents

Ziwei Wang, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Junjie Sheng, Jun Luo, Xiangfeng Wang, JieHao Wu

In the field of cloud services, virtual machine scheduling is a typical Online Dynamic Multidimensional Bin Packing (ODMBP) problem, facing significant challenges due to large-scale complexity and fluctuating demands. Traditional optimization methods struggle to adapt to real-time changes, domain experts-designed heuristic approaches suffer from rigid strategies, and existing learning-based methods often lack interpretability. To address these issues, this paper proposes a hierarchical language-agent framework, MiCo, which models ODMBP as a Semi-Markov Decision Process with Options (SMDP-Option) and enables dynamic scheduling through a two-stage "Option Miner-Option Composer" architecture. Specifically: (1) Option Miner leverages large language models (LLMs) to generate context-aware strategy options, automatically constructing interpretable scheduling rules; (2) Option Composer employs an LLM-driven context-aware switching

mechanism to analyze historical VM patterns and dynamically select the optimal option from a refined option set. Experimental results on real-world enterprise datasets demonstrate that MiCo achieves a 96.9% competitive ratio in large-scale scenarios with over ten thousand VMs, maintaining high performance even under non-stationary request flows, thus validating its effectiveness in complex large-scale cloud environments.

3 - Graphthought: Graph Combinatorial Optimization with Thought Generation

JieHao Wu, East China Normal University, ShangHai, China, People's Republic of

Large language models (LLMs) have demonstrated remarkable capabilities across various domains, especially in text processing and generative tasks. Recent advancements in the reasoning capabilities of state-of-the-art LLMs, such as OpenAI-o1, have significantly broadened their applicability, particularly in complex problem-solving and logical inference. However, most existing LLMs struggle with notable limitations in handling graph combinatorial optimization (GCO) problems. To bridge this gap, we formally define the Optimal Thoughts Design (OTD) problem, including its state and action thought space. We then introduce a novel framework, **GraphThought**, designed to generate high-quality thought datasets for GCO problems. Leveraging these datasets, we fine-tune the Llama-3-8B-Instruct model to develop Llama-GT. Notably, despite its compact 8B-parameter architecture, Llama-GT matches the performance of state-of-the-art LLMs on the GraphArena benchmark. Experimental results show that our approach outperforms both proprietary and open-source models, even rivaling specialized models like o1-mini. This work sets a new state-of-the-art benchmark while challenging the prevailing notion that model scale is the primary driver of reasoning capability.

4 - Autoadmm: Learning Adaptive Penalty for Admm with Large Language Models

Yun Hua, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Shiqin Wang, Jun Luo, Xiangfeng Wang

The Alternating Direction Method of Multipliers (ADMM) has demonstrated remarkable efficiency in addressing complex optimization challenges across machine learning, statistical modeling, and signal processing domains. Nevertheless, the algorithm's convergence behavior exhibits pronounced sensitivity to the penalty parameter β (commonly treated as a hyperparameter) in the augmented Lagrangian formulation, which remains non-adaptive in standard implementations. In this paper, we will propose **AutoADMM**, a computational optimization framework that leverages large language models (LLMs) to autonomously discover adaptive control policies for the penalty parameter β in ADMM optimization processes. By framing the challenge as an algorithmic discovery task, we leverage an extended version of the LLM-based Funsearch framework to automatically generate innovative self-adaptive parameter schemes. The AutoADMM framework integrates three synergistic agents: the Pre-trained LLM agent, Creator agent, and Judger agent. The Pre-trained LLM agent serves as the core component for generating and iteratively refining self-adaptive parameter configurations through combinatorial search. The Creator agent dynamically optimizes high-performance candidate programs through bidirectional interactions with the LLM, while the Judger agent rigorously evaluates program viability through dual criteria: functional correctness verification and convergence efficiency analysis. This tripartite architecture harnesses the zero-shot/few-shot learning capabilities and program synthesis strengths of large language models to dramatically enhance exploration efficiency for self-adaptive parameter schemes compared to conventional manual design approaches. We conducted comprehensive evaluations of AutoADMM across 17 distinct tasks from the LibADMM benchmark, encompassing both sparse and low-rank optimization problems, to systematically validate its superior efficiency and robust generalization capabilities in automatically learning adaptive penalty structures for ADMM.

SB15

Fairmont Ballroom

Applications of Hybrid Quantum Optimization in Transportation and Logistics

Invited Session

Transportation Science and Logistics

Chair: Yuli Zhang, Beijing Institute of Technology, Beijing, N/A

1 - Quantum-Enhanced Logic-Based Benders Decomposition for Assembly Line Balancing Problems

Yuli Zhang, Beijing Institute of Technology, Beijing, China, People's Republic of, Chun Cao, Muyang Wen, Qi Gao

This paper studies an assembly line balancing problem with sequence-dependent setup times and heterogeneous worker skills (ALBP-SW) to minimize cycle time. We develop a quantum-enhanced logic-based Benders decomposition (LBBD) method that partitions the problem into: (1) a CPU-solved master problem (MP) for job assignments, and (2) coherent Ising machine (CIM) solable QUBO subproblems for job sequencing. The approach features three key innovations: first, MP enhancements including sequencing relaxations, warm-start cuts, and novel valid inequalities; second, strengthened Benders cuts derived from infeasible task subsets; third, a three-phase variable fixation strategy by exploiting weak-persistence properties. Experimental results demonstrate the method's effectiveness, fixing 68% of variables on average (with 100% maximum fixation) while achieving 39.9% faster convergence than classical approaches on CIM hardware. The variable fixation strategy proves particularly impactful, reducing problem complexity while maintaining solution quality.

2 - Benchmarking of Quantum and Classical Computing in Large-Scale Dynamic Portfolio Optimization Under Market Frictions

Ying Chen, National University of Singapore, Singapore, Singapore

Quantum computing is poised to transform the financial industry, yet its advantages over traditional methods have not been evidenced. As this technology rapidly evolves, benchmarking is essential to fairly evaluate and compare different computational strategies.
This study presents a challenging yet solvable problem of large-scale dynamic portfolio optimization under realistic market conditions with frictions. We frame this issue as a Quadratic Unconstrained Binary Optimization (QUBO) problem, compatible with digital computing and ready for quantum computing, to establish a reliable benchmark. By applying the latest solvers to real data, we release benchmarks that help verify true advancements in dynamic trading strategies, either quantum or digital computing, ensuring that reported improvements in portfolio optimization are based on robust, transparent, and comparable metrics. This is a joint work with Thorsten Koch, Hanqiu Peng and Hongrui Zhang. https://arxiv.org/abs/2502.05226

3 - Hybrid Quantum Benders' Decomposition for Optimal Classification Tree

Peng Sun, Tianjin University, Tianjin, China, People's Republic of

In this paper, we address the challenge of learning optimal binary classification trees with univariate splits by employing a hybrid quantum Benders' decomposition approach. The strategy of Benders' decomposition, essentially a divide-and-conquer technique, has been widely applied across various fields such as communication, networking, and machine learning. However, the master problem within Benders' decomposition remains NP-hard, which motivates the integration of quantum computing. Specifically, we reformulate the master problem as a quadratic unconstrained binary optimization (QUBO) model and solve it using a state-of-the-art quantum annealer. Our computational analysis demonstrates the feasibility and efficiency of this approach. By leveraging the strengths of both classical and quantum computing through our reformulation, our hybrid method ensures high-quality solutions for generating optimal binary classification trees.

SB16

VIP A

AI in Business and Finance

Contributed Session

Artificial Intelligence

Chair: Tamer Boyaci, ESMT Berlin, Berlin, Germany

1 - From Pixels to Profits: Ceo&Rsquo;S Digital Human and the Art to Persuade Investors

Mijia Gong, Northwestern Polytechnical University, Xi'an, China, People's Republic of, Xiaoying Jiang

Digital human (DH) has been an emerging topic of interest to scholarly work on stakeholder reactions in the era of artificial intelligence (AI). However, existing studies examining the relationship between DH's appearance anthropomorphism and stakeholder responses have come to inconsistent results. To explore this inconsistency,

we focus on CEO's DH in the context of disclosing firm financial information and draw upon transportation theory to investigate the impact of their appearance anthropomorphism on investors' responses through both emotional and cognitive processes. Moreover, we also consider the moderating effect of content concreteness/abstractness when disclosing financial information. A laboratory experimental study using Tobbi eye-tracking technology reveals the power of CEOs' DH. Specifically, investors will experience less cognitive load, exhibit more positive emotions, and thus show more willingness to invest when interacting with CEOs' DH possessing high human-like appearance. Besides, we find that content concreteness of financial information will strengthen the positive relationship between CEO's DH appearance anthropomorphism and investors' willingness to invest. These findings provide important theoretical and practical insights into the literature on DH, and the communication channel between CEOs and stakeholders.

2 - Strategic Interactions in the AI Ecosystem

Ravi Prakash Ranjan, Africa Business School - UM6P, Rabat, Rabat, Morocco, Ravi Shankar Pandey

We develop a game-theoretic framework to analyze strategic interactions in the AI ecosystem, involving a chip infrastructure provider (e.g., Nvidia), proprietary AI firms (e.g., OpenAI), open-source AI firms (e.g., Deepseek), and AI investors. We model a three-stage sequential game: **First**, the infrastructure provider chooses whether to prioritize the demand of proprietary or open-source AI firms. **Second**, observing the provider's choice, the proprietary and open-source firms simultaneously select strategies balancing: a) *scale* (hardware capabilities) versus *efficiency* (software capabilities) and b) *expansion* (user base growth) versus *monetization* (revenue per user). **Finally**, investors choose capital allocation across infrastructure, proprietary AI, open-source options, or diversification. Payoff structures incorporate network effects (θ) and investor risk aversion (λ). Using backward induction, we derive the Subgame Perfect Nash Equilibrium (SPNE), solving the Stage 2 simultaneous-move subgame via Nash Equilibrium. Our analysis identifies potential equilibrium archetypes conditional on parameters: an *Infrastructure Provider-led ecosystem* with proprietary dominance (e.g., under low θ), *Competitive coexistence* with market segmentation (moderate θ), and an *Open-source driven transformation* (high θ). The framework supports how network effects influence strategic choices and outcomes, and how high investor risk aversion can favour Infrastructure Provider, offering critical insights for stakeholders navigating evolving AI market.

3 - Not All Bots are Equal: Exploring Algorithm Aversion and Support Goals in AI Chatbot use

Chong Oh, University of Utah, Salt Lake City, UT, United States, Carina Pals, Jae Won Oh, Rishi Ramesh

This study investigates how college students respond to AI chatbots when dealing with academic stress, using *Algorithm Aversion Theory* as a guiding framework. We compare two types of support: high-contingency chatbots versus low-contingency chatbots. Participants are randomly assigned to interact with one of these and are then asked to complete a survey assessing how supported they felt whether the experience helped them reframe their thinking, and how much emotional relief they experienced.

High-contingency chatbots are designed to offer personalized, context-aware responses using natural language processing, while low-contingency bots follow simple scripts and provide more generic replies. We expect that people who are more skeptical of algorithmic decision-making, those with high algorithm aversion, will prefer low-contingency bots because they feel more predictable and less "human-like." However, we also believe that users' underlying goals will play a role: people looking for emotional support may be more open to high-contingency bots, while those looking for instrumental information may lean away from them.

We hypothesize that when the chatbot's behavior matches the user's goal—whether that's getting information or feeling emotionally heard—high-contingency bots can be just as effective as human support, and possibly more so in certain contexts. This research contributes to our understanding of how design choices in AI tools intersect with user expectations and biases. The findings can help developers create more thoughtful and user-sensitive chatbot experiences, especially for mental health and academic support services.

4 - *Beyond the Black Box: Unraveling the Role of Explainability in Human-Ai Collaboration*

Tamer Boyaci, ESMT Berlin, Berlin, Germany, Caner Canyakmaz, Francis de Véricourt

Explainable Artificial Intelligence (AI) models have been proposed to mitigate overreliance and underreliance on AI, which reduce the effectiveness of human-AI collaborative tools. Yet, empirical evidence is mixed, and the impact of explainable AI on a decision-maker (DM)'s cognitive load and fatigue is often ignored. This paper offers a theoretical perspective on these issues. We develop an analytical model that incorporates the defining features of human and machine intelligence, capturing the limited but flexible nature of human cognition with imperfect machine recommendations. Crucially, we represent how AI-based explanations influence the DM's belief in the algorithm's predictive quality. Our results indicate that explainable AI has varying effects depending on the level of explainability it provides. While low explainability levels have no impact on decision accuracy and reliance levels, they lessen the cognitive burden of the DM. In contrast, higher explainability levels enhance accuracy by improving overreliance but at the expense of increased underreliance. Further, the relative impact of explainability (c.f. a black-box system) is higher when the DM is more cognitively constrained, the decision task is sufficiently complex or when the stakes are lower. Importantly, higher explainability levels can escalate the DM's cognitive burden and hence overall processing time and fatigue, precisely when explanations are most needed, i.e. when the DM is pressed for time to complete a complex task and doubts the machine's quality. Our study elicits comprehensive effects of explainability on decision outcomes and cognitive effort, enhancing our understanding of designing effective human-AI systems in diverse decision-making environments.

SB17

VIP B

Carbon Emissions and Climate Policy

Contributed Session

Energy, Natural Resources and the Environment

Chair: Yiming Du, Tianjin university, 92 Weijin RoadNankai District，Tianjin，China, Tianjin, N/A

1 - India's Green Targets: An Assessment of Performance, Growth, and Forecasts

Siddhartha Rastogi, Indian institute of management Indore, Indore, India

World over, climate change and growth trade-off has emerged as a major debate over the past few years, particularly since the demise of the Kyoto protocol. Although India ranks low among carbon emitters on per capita basis, it is a major carbon polluter on an aggregate basis. However, India has self-imposed high commitments in international climate negotiations since 2015 Paris Climate Agreement. In this paper, we explore and forecast India’s relative performance on climate and pollution frontier. We take a historical assessment of the climate negotiations and targets. Focusing on India, we assess India’s relative rank and performance on various pollutants including carbon, greenhouse gases, and other polluting gases in terms of volume and per capita basis. Further, we evaluate and forecast India’s performance *vis-à-vis* self-declared targets over the past decade. We also explore the measures taken by India to curb emissions, such as renewable energy missions, biofuels adoption, and issues faced by India, such as technological, capital, and developmental challenges.

2 - Unlocking the Potential of Carbon Cost Pass-Through in Cross-Regional Electricity Markets: a Switch-China Model Approach to Enhancing Price Fairness and Reducing Emissions

Shanyuan Zhang, Tianjin University, Tianjin, China, People's Republic of, Hongyang Zou, Huibin Du

The low-carbon transition of the power sector is essential to achieving carbon reduction goals, yet regional disparities in energy mixes and carbon pricing pose challenges to fairness and competitiveness. This study examines the carbon cost pass-through mechanism in cross-regional electricity markets, analyzing its effects on renewable energy competitiveness, electricity prices, and national carbon emissions. Using the Switch-China-Open model, the research simulates carbon pricing strategies, power generation structures, and emission factors across regions. Addressing gaps in previous studies, this work aims to identify the optimal carbon cost pass-through level to balance the competitiveness of renewable energy and consumers' electricity cost burden. The results show that high carbon pricing is essential to achieving China's dual carbon targets, reducing emissions by 19% by 2050 compared to the BAU scenario and aligning with IPCC goals, while low carbon pricing delays the emissions peak by five years. High carbon pricing also accelerates renewable energy growth, with wind and

solar contributing 67% of power generation by 2050, as coal's share drops to 16%. Provincial carbon emission factors simulations reveal that high carbon pricing effectively lowers emission intensities and promotes balanced low-carbon development, though regional disparities persist, with high-carbon regions like Shaanxi remaining coal-dependent, while low-carbon regions like Fujian achieve successful transitions even under low-price scenarios.

3 - Carbon Emission Efficiency Measurement and Spatial Non-Equilibrium Analysis of Steel Enterprises Based on the Undesirable-Sbm Model

rengaowa wu, University of Science and Technology Beijing, beijing, China, People's Republic of

In carbon neutrality, steel enterprises with high energy consumption, high pollution, and high emission characteristics are essential for carbon emission reduction. There is an urgent need to explore ways to improve the carbon emission efficiency of steel enterprises. The study constructs an input-output model for steel enterprises from a total factor perspective of carbon emissions. The inputs are measured using capital, labor, and energy, and the outputs are measured using economic and environmental benefits. Secondly, the Undesirable-SBM model is constructed to measure the carbon emission efficiency of steel enterprises. Finally, a spatial disequilibrium analysis model is built based on the Dagum Gini coefficient theory. The study is conducted in 29 provinces in China. The research results show that the mean value of carbon emission efficiency of steel enterprises increases over time. Moreover, the provinces with higher carbon emission efficiency are mainly distributed in the eastern region. Regarding spatial unevenness, the inter-regional disparity is the most important source of regional differences in the carbon emission efficiency of Chinese steel enterprises. Optimizing the ability to allocate production factors and reducing negative externality output are important ways for steel enterprises to improve carbon emission efficiency.

4 - Deployment Strategies of Long-Duration Energy Storage Under Renewable Portfolio Standards and Carbon Cap-and-Trade Policies

Yiming Du, Tianjin University, Tianjin, China, People's Republic of, Huibin Du, Hongyang Zou

Long duration energy storage (LDES) installations, which enhance flexibility and enable renewable integration, are expected to surge in a zero-emission future. However, where and to what extent LDES should be deployed under evolving policy constraints remains unclear. Here, we investigate how policy design affects LDES deployment by focusing on two policies in China's power sector: renewable portfolio standards (RPS) and carbon cap-and-trade mechanisms. Ten policy-technology scenarios are simulated using the SWITCH-China capacity expansion model to assess their impacts on system cost, carbon emissions, storage scale, duration, and spatial allocation. Results show that under the RPS-autonomous optimization scenario, the system minimizes cost but achieves limited emission reductions. In contrast, the RPS-as-usual scenario, which reflects the continuation of existing provincial target trajectories, fails to sufficiently stimulate system flexibility, resulting in limited storage deployment and utilization. Combining RPS with carbon caps or high carbon prices significantly increases LDES investment, especially in the form of 10-hour or longer flow batteries, with deployments concentrated in resource-rich and grid-constrained regions such as Yunnan, Qinghai, and Inner Mongolia. Introducing LDES cost reductions in carbon-constrained scenarios further lowers system costs and enhances dispatch flexibility. However, this improved adaptability slightly weakens the reliance on high renewable penetration, leading to a marginal rebound in emissions. Among all scenarios, the combination of autonomous RPS, carbon caps, and reduced LDES costs achieves a Pareto-optimal trade-off between cost and emissions, demonstrating strong policy synergy between renewable integration and carbon control.

Sunday, July 20, 2:00 PM - 3:30 PM

SC01

Enterprise

Frontiers in Probabilistic Modeling for the Service Economy

Invited Session

Applied Probability Society

Chair: Zhaowei She, Lee Kong Chian School of Business, Singapore Management University, Singapore, Singapore

1 - Queueing-Informed Causal Inference for Behavioral Queueing Data

Likang Ding, Dongbei University of Finance and Economics, Dalian, China, People's Republic of, Armann Ingolfsson

Empirical work on behavioral queueing science focuses on establishing causality between system drivers and human behavior. With observational data, a reliable estimate of the causal impact requires the causal structure to include the most important factors. We argue that knowledge of queueing theory can provide important input regarding the causal structure. Specifically, we demonstrate how knowledge of queueing theory can helps determine an appropriate set of control variables and validate an instrumental variable. We argue that the econometric model should depend on whether the hypothesized behavioral mechanism is static or dynamic. We show that commonly used econometric models can be overly restrictive for particular mechanisms and suggest using a semi-parametric model. Our results are intended to provide guidance for empiricists that study queueing systems with human servers or customers, to help develop credible causal estimates to explain human behavior.

2 - Routing with Delayed Revelation of Feature Information

Xiaoquan Gao, Singapore Management University, Singapore, Singapore, Pengyi Shi

We study the routing problem for heterogeneous customer arrivals, where individual features are initially unavailable and only revealed after evaluation. This delay can lead to inefficiencies in resource allocation, particularly in systems with diverse user classes requiring personalized assessments. To address this challenge, we propose a two-stage approach. In the first stage, individuals are allocated to different programs based on arrival counts, historical class data, and facility occupancy, enabling proactive resource planning. In the second stage, individualized placement decisions are made after evaluations. We establish a performance bound relative to a full-information Markov Decision Process (MDP), provided that second-stage routing follows treatment effect rankings. Computationally, we develop a fluid model and apply constrained linear quadratic regulation (LQR) to derive an asymptotically optimal, easy-to-implement linear decision rule for the first stage. Simulations calibrated with real-world data from community corrections demonstrate that our approach outperforms existing threshold-based policies, improving staffing and resource coordination without requiring immediate judicial decisions for each individual.

3 - The "Netflix Model": a New Payment Model for Asymptomatic Disease Management

Zhaowei She, Singapore Management University, Singapore, Singapore, Huaiyang Zhong, Yueran Zhuo, Turgay Ayer, Jagpreet Chhatwal

Tackling the diagnosis and management of asymptomatic diseases is vital, as many go unnoticed in early stages, leading to severe complications later. The high costs of screenings and treatments pose additional challenges. The "Netflix Model" offers a solution: government agencies pay pharmaceutical companies an upfront fee for unlimited access to treatments, mirroring a subscription system. We offer a theoretical mathematical framework to understand the conditions under which such Netflix-style contracts can be most beneficial, presenting a potential solution to the challenges of asymptomatic disease management.

4 - Reservation Property and the Advantage of Being Small: a Queueing Analysis in Bike Sharing

Hailun Zhang, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of

In this paper, we build transient queueing models to analyze the dynamics of bike-sharing (or scooter-, moped-, and car-sharing) systems. In particular, we compare a small brand operating a small number of bikes and owning a small user base, with its larger counterpart.
By conducting a detailed analysis of the modified Bessel function of the first kind, we obtain novel structural properties on important waiting-time-related measures, and show the advantage of small-brand bikes. To enhance the robustness and applicability of our findings, we incorporate various settings of our queueing model, covering a wide range of business settings. The main result holds consistently, and it coincides with our practical observations.
We also discuss other interesting practical examples exhibiting special advantages of being small. Our analysis discovers a unique characteristic of the small brand. Namely, the low market share and bike accessibility leads to a small (and less competing) user base, which creates the consecutive-usage reservation property, hence attracts

and benefits the small pool of dedicated users. This advantage of being small could hold under a wide range of business settings. Our work provides a new strategic angle, in general, for small and medium-sized enterprises (SMEs) to be profitable under the competition against large companies.

SC02

Mercury

Airside Operations

Invited Session

Aviation Applications

Chair: Xiaorong Ren, Northwestern Polytechnical University, Xi'an, N/A

1 - Airline-Oriented Apron Assignment Problem with Towing Operations

xiaorong Ren, Northwestern Polytechnical University (NWPU), Xi’an Shaan'xi, China, People's Republic of, Chenhao Zhou

Air transport has seen significant growth and resilience, with global passenger traffic projected to reach 9.5 billion, surpassing pre-pandemic levels. Concurrently, airlines are rapidly and widely deploying narrow-body aircraft on both international and domestic routes, leading to operational challenges at hub airports designed for wide-body planes. These challenges include gate and apron capacity constraints and increased ground congestion. To enhance gate utilization, airports mandate the airlines to tow aircraft between gates and remote aprons, a process that incurs high costs and labor demands for airlines. This study introduces the apron allocation problem with towing operations (AAP-TO), which assigns multiple gates or aprons to an aircraft during its ground time by segmenting its stay into arrival, idle, and departure activities. A logic-based Benders decomposition approach is proposed to solve the AAP-TO, where the master problem assigns aircraft activities to aprons, and the subproblem checks for time window conflicts and determines precise activity timings. Numerical experiments demonstrate that the proposed approach significantly outperforms existing methods, achieving optimal solutions in 100 % of instances involving 60 and 90 aircraft activities, and 60 % of instances involving 120 activities, while superior lower bounds are obtained in the remaining instances. This framework offers a practical tool for airlines and airport operators to navigate the complexities introduced by the increasing airport traffic.

2 - Solving Air-Ground Collaborative Scheduling Problem in Airports: a Model-Based Monte Carlo Tree Search Approach

Haiyue Bi, Northwestern Polytechnical University, Xi'an, China, People's Republic of, Chenhao Zhou

Smooth coordination between aircrafts and ground resources are always expected to enhance safety and efficiency. However, the ever increasing airport traffic and incidents not only make the coordination much more difficult, but also give the operators tremendous stress. To address this challenge, this study proposes an air-ground collaborative scheduling problem that aims to improve the punctuality of aircraft ground operations. The scheduling involves a series of ground resources, including runway, taxiway and apron, and also considers the aircraft circling at approaching airspace. A complete mixed-integer linear programming scheduling model is developed to solve this problem, and a model-based Monte Carlo tree search (MCTS) algorithm is introduced to handle large-scale instances. The proposed model-based MCTS outperforms both the improved adaptive large neighborhood search and the first-come-first-serve policy, which reflects current manual scheduling practices. In the improved MCTS, a relaxed runway sequencing model is adopted in the runway simulation phase, enabling pre-consideration of key scheduling factors to guide the search direction. A partial scheduling model is also integrated to evaluate the current runway sequence within the collaborative framework, the result of which is backpropagated to the search tree as latter selection criterions, optimizing sequencing and scheduling solutions iteratively. Experiments using real-world data from Xiamen Gaoqi Airport demonstrate that the model-based MCTS algorithm achieves optimal solutions for 1-hour scenarios and near-optimal solutions for 2-hour scenarios within seconds, reducing average flight delays by 41.78%. The results also show the integration of proposed model can enhances the effectiveness of the MCTS.

3 - Vehicle Scheduling & Task Allocation Procedure for Mixed Electric and Ice Airside Fleet Using Mixed Integer Linear Programming

Wenqiang Wang, Institute of High Performance Computing, Singapore, Singapore, Muhamad Azfar Ramli

In this study, we present a methodology for assigning ground support equipment (GSE) to ground handling tasks, formulated as a vehicle scheduling problem in the context of a partially electrified fleet. The transition from internal combustion engine (ICE) GSEs to electric GSEs (e-GSEs) introduces operational challenges, mainly due to the need for battery charging and the associated vehicle downtime. Unlike ICE fleets, where GSE availability is largely continuous throughout operations, e-GSEs require recharging time, which can disrupt task scheduling and reduce overall availability. As a result, a larger fleet size is often necessary to maintain the same service level and ensure uninterrupted ground handling operations. To address these challenges, we develop an optimization model using Mixed Integer Linear Programming (MILP) to determine the minimum number of GSEs required to fulfil all scheduled tasks. For ICE GSEs, the model focuses on standard availability, whereas for e-GSEs, additional constraints are introduced to incorporate charging requirements, which include gap times between tasks, travel time to charging stations, battery capacity, charging speed, etc. We also demonstrate how different levels of fleet electrification can be introduced progressively, together with several interesting findings from the study, which we believe could support airport stakeholders in decision makings about transition to electric ground support operations.

4 - Multi-Fleet Scheduling and Optimization of Electric Specialized Vehicles at Airports: a Systematic Solution with Deep Reinforcement Learning

Hongwei Wang, IHPC, A*STAR, Singapore, Singapore

The electrification of Ground Support Equipment (GSE), such as air tugs, baggage tractors, and cargo trucks, presents emerging challenges for airport management, including energy supply-demand fluctuations and charging schedule coordination. Traditional scheduling methods often struggle to effectively balance real-time task demands, adapt to complex operational environments, and manage electrical grid loads. To address these challenges, this study proposes a deep reinforcement learning (DRL)-based scheduling framework with an enhanced maskable Proximal Policy Optimization algorithm designed for real-world airport operations involving multiple inputs and diverse optimization objectives. The model's goals are to minimize flight support delays by ensuring timely task completion, balance vehicle charging to mitigate grid impact, optimize fleet size for cost-efficiency, and reduce unnecessary movements of high-value vehicles to extend service life. Key elements of airport operations, including battery levels, charging station availability, task demands, and vehicle charging dynamics, have been comprehensively modeled within the DRL framework through clearly defined environments, agents, and actions. Flight schedules and vehicle operational requirements at Singapore Changi Airport were employed as an experimental scenario to evaluate the proposed RL method. Experimental results demonstrate that the RL model dynamically adapts to real-time requirements, effectively manages grid loads, reduces peak charging power, and improves energy utilization efficiency while ensuring high task completion reliability. Compared with traditional rule-based scheduling methods, the proposed DRL approach demonstrates stronger adaptability and enhanced optimization performance, offering a feasible and intelligent solution for operating electric specialized vehicles at airports.

SC03

Minto

Data-driven Decision-making in Operations Research

Invited Session

Applied Probability Society

Chair: Yi Chen, Hong Kong University of Science and Technology, Hong Kong, Hong Kong

1 - Persuasive Calibration

Wei Tang, Chinese University of Hong Kong, Hong Kong, Hong Kong, Yiding Feng

We introduce and study the persuasive calibration problem, where a principal aims to provide trustworthy predictions about underlying events to a downstream agent to make desired decisions. We adopt the standard calibration framework that regulates predictions to be unbiased conditional on their own value, and thus, they can reliably be interpreted at the face value by the agent. Allowing a small calibration error budget, we aim to answer the following question: what is and how to compute the optimal predictor under this calibration error budget, especially when there exists incentive misalignment between the principal and the agent? We focus on standard ℓt-norm Expected Calibration Error (ECE) metric.

<blockquote mathjax"="">We develop a general framework by viewing predictors as post-processed versions of perfectly calibrated predictors. Using this framework, we first characterize the structure of the optimal predictor. Specifically, when the principal's utility is event-independent and for L1-norm ECE, we show: (1) the optimal predictor is over-(resp. under-) confident for high (resp. low) true expected outcomes, while remaining perfectly calibrated in the middle; (2) the miscalibrated predictions exhibit a collinearity structure with the principal's utility function. On the algorithmic side, we provide a FPTAS for computing approximately optimal predictor for general principal utility and general Lt-norm ECE. Moreover, for the L1- and L-Infinity-norm ECE, we provide polynomial-time algorithms that compute the exact optimal predictor.

2 - Statistical Inference for Weighted Sample Average Approximation in Contextual Stochastic Optimization

Yanyuan Wang, The Hong Kong University of Science and Technology, Hong Kong, Hong Kong, Xiaowei Zhang

Contextual stochastic optimization provides a framework for decision-making under uncertainty incorporating observable contextual information through covariates. We analyze statistical inference for weighted sample average approximation (wSAA), a widely-used method for solving contextual stochastic optimization problems. We first establish central limit theorems for wSAA estimates of optimal values when problems can be solved exactly, characterizing how estimation uncertainty scales with covariate sample size. We then investigate practical scenarios with computational budget constraints, revealing a fundamental tradeoff between statistical accuracy and computational cost as sample sizes increase. Through central limit theorems for budget-constrained wSAA estimates, we precisely characterize this statistical-computational tradeoff. We also develop "over-optimizing" strategies for solving wSAA problems that ensure valid statistical inference. Extensive numerical experiments on both synthetic and real-world datasets validate our theoretical findings.

3 - Online Convex Optimization with Offline Data

Yi Chen, Hong Kong University of Science and Technology, Hong Kong, Hong Kong, Jize Xie, Rachel Zhang

To be filled later

4 - Stopping Rules for Sampling in Precision Medicine

Mingrui Ding, City University of Hong Kong, Kowloon, Hong Kong, Gao Siyang

Precision medicine (PM) is an approach that aims to tailor treatments based on patient profiles (patients' biometric characteristics). In PM practice, treatment performance is typically evaluated through simulation models or clinical trials. Although these two methods have differences in their sampling subjects and requirements, both are based on a sequential sampling process and require determining a stopping time for sampling to ensure that, with a prespecified confidence level, the best treatment is correctly identified for each patient profile. In this research, we propose unified stopping rules applicable to both simulation and clinical trial-based PM sampling processes. Specifically, we adapt the generalized likelihood ratio (GLR) test to determine a stage where samples collected are sufficient to make good enough treatment decisions for each profile. The stopping thresholds for the GLR statistics are calibrated using mixture martingales and a peeling method. Our stopping rules are theoretically grounded and can be integrated with different types of sampling strategies. Numerical experiments on synthetic problems and a case study demonstrate their effectiveness.

SC04

Indiana

Behavior in E-Commence and Healthcare

Invited Session

Behavioral Operations Management

Chair: Juan Li, Nanjing University, Nanjing, N/A

1 - *Heuristics and* Deep Reinforcement Learning for An Inventory Problem with An All-or-Nothing Yield Pattern and a Positive Leadtime

Wenjie Bi, Business School of Central South University, Changsha, China, People's Republic of, Yuting YAN, Frank Chen, Zhe Fu

We consider the single-item, periodic-review inventory system with stochastic demand, an all-or-nothing random yield pattern, and non-zero leadtimes, where the unmet demand is backlogged. The firm's objective is to find an optimal ordering policy to minimize its total expected discounted cost. As the optimal policy remains unknown, we first develop a new heuristic that takes advantage of the all-or-nothing yield pattern and then propose a deep reinforcement learning (DRL) algorithm based on a tailored neural network model. The proposed DRL algorithm combines the results of existing heuristics to improve performance and stabilize the training process. Moreover, the total cost to be computed is convex in the decision variable, which facilitates the search. Extensive numerical experiments show that our new solution methods outperform known heuristics and demonstrate near-optimal performance when leadtimes are relatively short and the minimal cost can be achieved numerically.

2 - The Value of AI Cheap Talk: a Field Experiment in Online B2b Procurement

Xiaojing Feng, Southwestern University of Finance and Economics, Chengdu, China, People's Republic of, Jingying Ding, Jiajia Wang

Price bargaining in B2B transactions is a core challenge for procurement managers. Research indicates that the initial price quote acts as an anchor in price bargaining, significantly influencing the final transaction price. Through a field experiment in online B2B procurement, we investigate how retailers leverage cheap talk strategies with varying confidence levels to obtain lower initial quotes, while systematically examining the role of AI technology integration in this process. Our research findings reveal that cheap talk strategies are effective in some cases, especially when AI initiates inquiries with high-confidence cheap talk postures. We explore the moderating effects of product list price and price volatility. Robustness analyses and managerial implications are further discussed.

3 - From Black to Grey: Improving Access to Antimalarial Drugs in the Presence of Deceptive Counterfeits

Jiatao Ding, Nanyang Technological University, Singapore, Singapore, Michael Freeman, Sasa Zorc

In malaria-endemic countries, the limited availability of affordable antimalarial medication has contributed to the widespread distribution of counterfeit drugs. This paper examines such markets to determine how philanthropic donors can best allocate limited funds to subsidize the purchase or sales of antimalarial drugs via private-sector distribution channels. To achieve this, we develop a game-theoretic model of the antimalarial supply chain, wherein the retailer strategically sources legitimate drugs from a donor-certified supplier, potentially counterfeit drugs from an uncertified supplier, or both. In contrast with the extant literature, we demonstrate that in the presence of counterfeits, exclusive reliance on purchase subsidies may no longer be optimal. Specifically, under donor budget constraints, offering sales subsidies that incorporate both legitimate and counterfeit drugs may be preferable, or in some cases, abstaining from subsidy provision completely. Additionally, we evaluate five pertinent non-subsidy strategies that have been employed to combat counterfeit drugs, analyzing the conditions under which these strategies can either enhance or compromise outcomes. Our study concludes with extensive numerical analysis calibrated to malaria data from Mozambique. This paper offers strategic guidance for improving outcomes in the presence of counterfeit drugs. Our results highlight the need for regulators and donors to understand specific market characteristics, such as retailers' pricing power and demand uncertainty,

when designing effective subsidy schemes and selecting appropriate technologies and policies to improve access to life-saving medicines.

4 - Economy Size vs. Compact Packaging: the Analysis of Product Scaling in Consumer Purchasing Behavior

Qing Zhang, Nanjing University of Aeronautics and Astronautics, Nanjing, China, People's Republic of

When selling products, a firm often seeks to expand its market reach by offering a diverse range of items in various sizes. In addition to standardized products, one can introduce options like compact-packaging (i.e., small-package) or economy-size (i.e., large-package) products to cater to varying consumer preferences. Within this framework, we conduct an investigation into the product line strategy and pricing mechanisms utilized by a monopolistic firm, focusing on how product size influences consumer purchasing decisions. Our findings indicate that expanding the product line to include large-package options may paradoxically reduce the total number of consumers making purchases, potentially harming the firm's bottom line. Conversely, introducing large-package products can offer benefits contingent upon the cost of standardized products. By comparing three distinct product line extension strategies, we demonstrate that product line expansion generally results in an increased retail price for standardized products. When consumers assign high value to products, the optimal strategy is to offer only standardized options. If consumer valuations are moderate, introducing small-package products is advisable; otherwise, the focus should shift to expanding offerings with large-package products. We additionally investigate the motivations behind a firm's simultaneous introduction of both small- and large-packaging products, and offer valuable insights to help the monopoly make informed decisions regarding optimal product sizing.

SC05

Orchard A

Operations Design Across Digital Frontiers

Invited Session

Decision Analysis Society

Chair: Lin Chen, University of Hong Kong, Hong Kong

Co-Chair: Abhishek Deshmane, Georgia Institute of Technology, Atlanta, GA, United States

1 - The Impact of Test Reporting on Autonomous Vehicle Training and Its Welfare Implications

Zhi Chen, National University of Singapore, Singapore, Singapore, Wenjie Xue

The success of autonomous vehicles (AVs) depends on its access to various real-world driving scenarios. To improve AVs' algorithms, firms send testing vehicles to discover unknown scenarios. While these scenarios are valuable for AVs' future improvements, they trigger failures. We study how a test reporting regulation affects firms' testing strategies.

2 - Ads for Loot: A Field Experiment on Rewarded Ads in Freemium Games

Mingqing Cai, University of Hong Kong, Hong Kong, Hong Kong, Xiao Lei, Lei Huang, Zhixi Wan

In-game advertising, a rapidly expanding sector of the digital advertising industry, is forecast to yield revenue of $109.6 billion by 2024. Distinguished from conventional advertising, in-game ads, frequently termed "rewarded ads," are optional and provide players with rewards such as additional lives or in-game items in exchange for ad viewing. Despite the substantial monetization potential of rewarded ads, game developers often exhibit reluctance to implement them due to uncertainties surrounding their impact on players. We conducted the first-ever field experiment on rewarded ads in collaboration with a leading advertising platform and a freemium game, introducing rewarded ads that offer progression acceleration rewards. Our findings reveal that while rewarded ads significantly reduce the usage of premium speedup coupons (their closest functional substitutes), they generate incremental revenue without cannibalizing in-app purchases or harming user engagement. We identify two behavioral mechanisms driving the substitution effect on speedup coupons: (1) players conserve coupons due to mental accounting shifts triggered by ad availability, and (2) players strategically avoid using rewards for short-duration quests, reserving them for high-value scenarios. We further explore player behavior to

provide plausible explanations for the insignificant impact on engagement, retention, and monetization. The negligible impact on engagement and retention is attributed to players reallocating their playing time from progression-increasing activities to non-progression activities (e.g., social interactions). Rewarded ads do not significantly affect player purchase behavior because players demonstrate stockpiling behavior for the rewarded items even in the absence of rewarded ads, which indicates little demand for them.

3 - Dali OR Dall-E? Consumer Perception of Human Versus AI Creative Processes

Abhishek Deshmane, Georgia Institute of Technology, Atlanta, GA, United States, Diwas KC

As AI-generated content proliferates in creative domains, understanding how consumers evaluate the authenticity and value of such works has become increasingly important. This study examines how attribution -- human versus AI-assisted -- shapes consumer perceptions of visual art, and whether disclosures about creative effort and uniqueness moderate these effects. Drawing on theories of authenticity, creative labor valuation, and algorithmic aversion, we theorize that AI involvement may trigger a perceived deficit in originality or intentionality, lowering perceived value. However, we also propose that transparency about the creative process can restore some of this lost value by reintroducing signals of effort and uniqueness. Using a VR-based mixed within-between subjects experiment, participants viewed artworks with systematically varied attribution and process information before expressing preferences and willingness to pay. Our findings show a consistent “AI discount” in perceived value, but also reveal that this discount is significantly reduced -- sometimes eliminated -- when the AI-assisted creation is accompanied by high-effort disclosures or uniqueness cues. These results advance our understanding of how consumers make sense of creativity in the age of AI, and suggest that attribution and process transparency can meaningfully shape public reception of algorithmically generated content.

4 - Designing Surprise Bags for Surplus Foods

Fan Zhou, Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Hansheng Jiang, Andrea Li, Joline Uichanco

Too Good To Go (TGTG) has emerged as a global leader in reducing food waste by connecting consumers with retailers’ surplus food through “surprise bags.” While TGTG provides both environmental and economic benefits, participating stores face significant operational challenges in determining the optimal number of bags to offer. Specifically, stores must decide on the number of surprise bags before knowing the actual surplus for the day. Additionally, key questions such as whether to include regular sale items in the bags and how to allocate food across bags are critical for balancing customer satisfaction and profitability. We explore these challenges and model the stores’ decision-making process by stochastic dynamic programming. Our analysis reveals that the commonly adopted practice of evenly distributing surplus food across all surprise bags may not maximize consumer satisfaction, particularly when customer utility functions are non-concave. We propose an alternative —offering up to two types of bags—that achieves maximal consumer satisfaction while still remaining simple to implement. Determining the optimal number of bags and the total food volume to allocate at each period is a non-convex two-stage stochastic optimization problem. To address both the analytical and computational complexities, we propose two approximation policies that are computationally efficient with intuitive managerial insights and provable performance guarantees. Our findings provide practical guidance for stores in balancing revenue generation, customer satisfaction, and sustainability objectives, thereby enhancing the long-term viability of TGTG’s mission to reduce food waste.

SC06

Orchard B

Financial Innovations and Market Behavior

Contributed Session

Finance

Chair: Jinjun Liu, Hong Kong Baptist University/BOCHK, Hong Kong, N/A

1 - Interest Rate Optimization for Installment Loan

SIRONG LUO, Shanghai University of Finance and Economics, Shanghai, China, People's Republic of

Consumer lending has been the driving force for economic growth. In this paper, we study the credit decision optimization problems in the consumer lending industry. Specifically, we analyze the interest rate optimization problem for installment loan. The interest rate decisions are crucial for lending product, which requires to carefully balance market acceptance, credit loss, and interest income. We estimate the effects of interest rate on both default risk, early payoff and acceptance probability by considering causal inference. We formulate the interest rate optimization problem as an integer programming model with risk constraints. Moreover, we provide an efficient converged algorithm to find the optimal interest rate. We also investigate the differences in the model estimation and optimal interest rate decision with and without causal inference, which provides valuable management insights.

2 - Invoice Tokenization and the Payables Finance In Supply Chains

Hariprasad Bellamkonda, Indian Institute of Management-Indore, Indore, India

Supply chain finance programs generally leaves lower-tier suppliers with higher financing costs. Blockchain-based invoice tokenization enables creditworthy manufacturers to transfer the tokens to tier-2 suppliers. How the invoice tokenization affects the profitability of the supply chain players is examined in the three-tier serial supply chain configuration. It is observed that the tier-2 suppliers' receivables quality is enhanced and the tier-1 supplier's payment liability is isolated while reducing the manufacturer's sourcing costs. The manufacturers with higher credit risk may derive greater value from invoice tokenization. The negative impacts of Invoice tokenization stem from its indirect influence on sourcing strategies in multi-tier supply networks. Suitability of reverse factoring for financing the deep-tier network also examined.

3 - Sword of Damocles - Effect of the Juventus Verdict on Fraudulent Earnings Management in Football

Wuwei Wang, Southwestern University of Finance and Economics, Chengdu, China, People's Republic of

This study employs triple difference and difference-in-difference approaches to investigate whether football clubs use swap deals to inflate capital gains, a form of fraudulent earnings management, and whether the Juventus verdict acts as a sword of Damocles threatening Italian clubs to reduce such earnings management. The dependent variable of interest is constructed as the log-ratio of the recorded value of the player transfer to the fair value of the transfer. I find that Italian clubs are more inclined to inflate capital gains in swap deals than their European counterparts, consistent with prior research. However, the Juventus verdict failed to threaten Italian clubs to reduce such manipulation statistically, while the verdict does on average mitigate earnings management in Italian clubs' swap deals. In a separate investigation, results from regressions with subsamples show that, the Juventus verdict may have threatening effects partially since the interaction term is significant for the sample pre-verdict but insignificant for the sample post-verdict. These findings shed light on the complexity of relation between law and economic behavior. It shows the difference between statistical evidence and legal evidence in determining fraudulent earnings management. This study also contributes to the literature by enhancing the understanding of a type of potentially fraudulent earnings management that's largely ignored in the literature, and ideas and techniques used in this paper could apply analogically to financial markets and general industries and sectors.

4 - Short-Term Anchoring Effects on Stock Returns: Empirical Evidence from the China Market

Wenlan Xie, The University of Sydney, Sydney, Australia

This study investigates the role of short-term anchoring effects in the Chinese stock market by examining how proximity to recent price peaks influences future returns and trading behavior. Using daily trading data from 2004 to 2023 for A-shares listed on the Shanghai and Shenzhen stock exchanges, we document a significant negative relationship between a stock's proximity to its 30-day high and subsequent returns, with a one standard deviation increase in proximity associated with a 9.2% decline in future 30-day returns. This effect is strongest at the 30-day horizon and demonstrates pronounced heterogeneity across market segments. Most notably, we identify a striking asymmetry across market states: the negative anchoring effect during bear markets reverses to a positive relationship during bull markets. Furthermore, the effect is more pronounced for state-owned enterprises and financially unconstrained firms. These findings, robust to multiple specifications including

Fama-French factor adjustments, reveal how investor psychology interacts with institutional features and market conditions to shape asset pricing in emerging markets characterized by high retail participation.

5 - Enhancing U.S. Treasury Yield Curve Forecasting Through Machine Learning and Forecast Combination Techniques

Jinjun Liu, Hong Kong Baptist University, Hong Kong, Hong Kong, Ming-Yen Cheng

This study presents an innovative framework for improving US Treasury yield curve forecasting, with a dual emphasis on enhancing prediction accuracy and ensuring interpretability in machine learning models, particularly during periods of financial stress. The framework integrates advanced machine learning algorithms with adaptive forecast combination techniques, emulating the decision-making process of professional forecasters to improve economic transparency.

We introduce the Enhanced Factor-Augmented Dynamic Nelson-Siegel (Enhanced FADNS) model, which leverages 10 principal components derived from 151 macroeconomic variables (adjusted for stationarity) and the market value of Treasury debt. To capture non-linear dependencies and dynamic interactions among market drivers, we incorporate Random Forest and Long-Short-Term Memory (LSTM) models. Forecasts are combined using eight techniques, including Equal Weights, Rank-Weighted Combinations, three Thick Modeling approaches, and three variants of the Aggregated Forecast Through Exponential Reweighting (AFTER) algorithm. The AFTER algorithm, which dynamically adjusts weights based on historical performance, consistently outperforms other methods and demonstrates resilience to structural breaks.

Our empirical evaluation extends to sensitivity assessments, cointegration tests, and the development of error correction models to assess the robustness of the AFTER algorithm. Rigorous robustness assessment is conducted to ensure the validity of results. First, all analyses are replicated using end-of-month yield observations to verify consistency across data sampling methodologies. Second, model performance is evaluated without the LSTM framework to ensure that improvements in predictive accuracy are not reliant on deep learning models.

Although Treasury yields are influenced by short-term market fluctuations, our findings reveal that they ultimately converge to underlying macroeconomic and supply-demand fundamentals.

SC07

Orchard C

Applications of AI and Alternative data in Finance

Invited Session

FinTech

Chair: Yunchuan Sun, Beijing Normal University, 19 Xinjiekouwai St, Beijing, 100875

Co-Chair: Junhuan Zhang, Beihang University, Beijing

1 - Gubasenti: Measuring Individual Investor Sentiment in Chinese Stock Markets Based on Big Text Data

Yunchuan Sun, Beijing Normal University, 100875, Beijing, China, People's Republic of

A distinctive feature of China’s capital market lies in the massive scale of individual investors, with their total number exceeding 223 million. These investors serve as a vital source of market vitality and a fundamental pillar for the effective operation of market mechanisms. Investigating the impacts of individual investor sentiment on the stock market and its underlying mechanism is of great theoretical and practical significance. This paper develops a systematic individual investor sentiment index, named GubaSenti, through in-depth mining and analysis of textual big data from Eastmoney Guba, the most popular online stock forum in China. The GubaSenti exhibits the following distinctive features, 1) Targeted precision: directly and accurately capturing individual investor sentiment through granular textual analysis of user-generated content; 2) Real-time updates: leveraging an automated system to enable high-frequency sentiment tracking; 3) Multi-granular coverage: aggregating

sentiment data hierarchically at the stock, industry, and market levels; and 4) Frequency flexibility: providing sentiment data at daily, weekly, and monthly frequencies. Empirical results reveal that GubaSenti exerts a significant and positive impact on stock returns, and the impact is more pronounced for stocks with greater individual investor attention and stronger limits to arbitrage. Moreover, the sentiment divergence among stocks with distinct cross-sectional characteristics significantly and positively influences the corresponding anomaly return.

2 - Profitability of Collusive Sandwich Attack in Automated Market Maker Exchanges

Junhuan Zhang, Beihang University, Beijing, China, People's Republic of

In decentralized exchanges (DEXs), sandwich attackers may collude with validators through private memory pools. We investigate the collusive behaviors between sandwich attackers and colluding validators in terms of a new game-theoretic market microstructure model of automated market maker based decentralized exchange. We also explore the impacts of this collusion on the other agents including arbitrageurs employing the two-point (TA) arbitrage strategy, liquidity providers using the balancing swap fees and impermanent loss strategy, swap traders using the moving average (MA) and zero intelligence (ZI) strategies, and ordinary validators with honest strategy. The deep reinforcement learning is used to optimize the collusive trading strategy of sandwich attackers. We validate our model using daily closing prices of USDT/BTC from January 1, 2023 to December 31, 2023. The results show that: 1) the collusive sandwich attack strategy is effective and profitable; 2) colluding validators receive larger rewards than those from only verifying transactions; 3) the collusion does not affect the rewards of the ordinary validators, but reduces the rewards of the liquidity providers; 4) rewards for the TA arbitrageurs are diminished as a result of collusion; 5) MA and ZI swap traders consistently suffer losses.

3 - Fintech and Business Income of Disadvantaged Groups

Xin Deng, Nanyang Technological University, Singapore, Singapore, Huai Zhang, Jiaxin Peng

Using a novel dataset of microbusiness owners from China for the period 2017-2019, we investigate how fintech adoption affects income inequality faced by seniors and females. Our results reveal an unexpected pattern - microbusiness owners from these two disadvantaged groups are less likely to adopt a value-enhancing fintech solution. However, once they adopt it, they experience a greater increase in their business revenue than their non-disadvantaged counterparts. This effect is more pronounced in areas with high levels of discrimination against seniors and females. Additional evidence suggests that fintech helps improve cash management, ease financing constraints, and advance management capability for disadvantaged groups. Overall, our results underscore the potential of fintech to mitigate income inequality by levelling the playground for disadvantaged groups and highlight the importance of government policies in promoting fintech adoption among disadvantaged individuals.

4 - Corporate Default Prediction Through News Insights

Hui Bu, Beihang University, Beijing, China, People's Republic of, Wenjin Tang, Jun Tu, Junbo Wang, Shouyang Wang

This study explores the informational value of media news and its predictive capability in forecasting the first-time default of corporate bond issuers—an inherently challenging task. By employing a novel credit risk dictionary that incorporates domain-specific topics, we construct a new risk-related measure based on news headlines, drawing from a comprehensive dataset encompassing almost all corporate bonds and media coverage in China from May 2015 to December 2022. Our findings demonstrate that media news provides incremental information beyond traditional financial ratios and economic variables, providing robust early-warning capabilities with a predictive horizon extending at least six months. Furthermore, the study highlights the distinct roles and informational channels of state-controlled versus market-oriented media, considering the unique media characteristics, ownership structures of firms in China, and varying information density. Our results support quality-oriented incentive hypothesis for market-oriented media and reveal “silencing effect” of state-controlled media, particularly in their reporting on central state-owned enterprises (CSOEs) and state-owned enterprises (SOEs). Additionally, our findings identify the distinct topic content emphasized by different types of media.

SC08

Bras Basah A

Optimization and Machine Learning in Emergency Medical Services

Contributed Session

Healthcare Applications

Chair: Adele Marshall, Queen's University Belfast, University Road, Belfast, BT71NN, United Kingdom

1 - Near-Optimal Dispatch Policies for Emergency Medical Services

Cheng Hua, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Tong Wang, Jingwei Zhang, Ziyan Zhou

Emergency medical services (EMS) are vital for ensuring timely and effective healthcare delivery. Ambulance dispatching in EMS directly influences patient outcomes. These systems face the challenge of managing limited resources to respond promptly to emergency calls while maintaining the capacity to address potential future incidents. We model the problem as a continuous-time stochastic system and determine the dispatch decision for each sequentially observed call to minimize the system-wide average cost. We consider the heterogeneous nature of call arrivals and general service time distributions with heterogeneous rates. To address this problem, we develop an easy-to-implement and near-optimal policy based on Lagrangian relaxation of the original problem. Using a novel proof technique, we show that our policy achieves performance within O(1/√theta) of the optimal, where theta represents the linear scaling factor for the number of ambulances and arrival rates. Our case study, based on real data from New York City, demonstrates that implementing the proposed policy can effectively reduce costs under varying levels of requests. Empirically, our policy consistently outperforms widely-applied benchmarks in both high and low traffic regimes. This study also provides insights into the effectiveness and limitations of myopic policies and underscores the value of properly adopting mutual-aid policies.

2 - Attention is What You Need: the Impact of An Information Support Tool in Addressing Trauma Center Deserts

Zidong Liu, Cambridge Judge Business School, Cambridge, United Kingdom, Feryal Erhun, Houyuan Jiang

Major trauma is a critical medical emergency requiring rapid intervention. In the UK, inclusive trauma networks (TNs) operate under a tiered model, balancing centralized resources with geographical access. However, trauma center deserts—regions with limited access to Major Trauma Centers (MTCs)—worsen health inequalities. To support pre-hospital teams in the triage process, a standardized trauma triage tool was introduced to enhance decision-making, improve awareness of patient conditions, and facilitate rapid medical decisions in emergencies. This study evaluates the impact of this information support tool on system efficiency and patient outcomes, particularly in trauma center deserts.

Analyzing nine years (2013–2021) of trauma submission data from the East of England Trauma Audit and Research Network (TARN), we assess the tool’s impact. Our findings indicate that while the triage tool does not significantly affect survival rates, it reduces the average hospital length of stay by 0.5 days. More importantly, it increases good recovery rates by 11% overall, with a notable 17% improvement for patients located more than 45 minutes from an MTC.

These results highlight the critical role of operational decision-support tools in optimizing trauma care and reducing geographical disparities. By standardizing assessments, enhancing healthcare workers’ awareness, and improving real-time triage decisions, such tools ensure timely and appropriate treatment, particularly in underserved areas.

Our study provides empirical evidence for policymakers and healthcare providers to enhance trauma system efficiency and address geographical healthcare inequities in trauma center deserts.

3 - Modelling the TIME from Emergency Call to Ambulance Discharge for out of Hospital Cardiac Arrests in Northern Ireland

Adele Marshall, Queen's University Belfast, Belfast, United Kingdom, Julia Wolfe, Emma McCorkell, Aleksandar Novakovic

An out of hospital cardiac arrest (OHCA) occurs when an individual's cardiac activity ceases when they are in any location outside of a hospital facility. The Northern Ireland Ambulance Service is the sole NHS Emergency Medical Service (EMS) that responds to anyone who experiences an OHCA in Northern Ireland. This study has for the first time collected data over a five year period on all OHCA cases in Northern Ireland including three years of data rxploring the time at which the emergency call was made, time of crew arrival, time spent on scene, right through to the ambulance delivering the patient to hospital where appropriate. This paper models the time to event data and patient flow using survival methods for the emergency response considering separately the time for the different stages to be completed and potential influencing factors.

4 - Development and Validation of An Electronic Chinese-Adapted Stroop Color-Word Test for Executive Functions Evaluation: a Pilot Study

I-Chun Lin, National Yunlin University of Science and Technology, Douliou, Taiwan

Executive functions play a critical role in promoting healthy aging, supporting autonomy, and preventing cognitive decline among older adults. The Stroop Color-Word Test (SCWT) is a widely recognized cognitive assessment tool that evaluates key executive functions, including selective attention, inhibitory control, and cognitive flexibility. This study developed and validated an electronic Chinese-adapted version of the Stroop test (eSCWT) aimed at improving accessibility and early screening for executive function impairments in Chinese-speaking populations, particularly among sub-healthy individuals. A total of 123 healthy adults participated in a test-retest design with a 7-day interval. Reliability metrics demonstrated strong internal consistency (Cronbach’s alpha = 0.84–0.94) and moderate to high test-retest reliability (intraclass correlation coefficient = 0.47–0.88). The electronic format not only automates data collection but also addresses cultural and linguistic gaps inherent in the original SCWT, while aligning with user familiarity with tablet-based applications. These features enhance the usability and scalability of the test in both clinical and community settings. Future work will expand the sample size and age range to further establish the generalizability and robustness of the eSCWT for broader executive function screening and intervention planning.

5 - Not All Labs Belong at Home: a Delphi-Based Framework and Machine Learning Roadmap for Smart Decentralization of Diagnostic Testing

Tarit Taerakul, Technology Management and Innopreneurship (Interdisciplinary Program), Graduate School, Chulalongkorn University, Rajavithi Hospital, College of Medicine, Rangsit University, Bangkok, Thailand, Krit Pongpirul

Background:
Decentralized and at-home laboratory testing offers a promising solution to improve diagnostic accessibility and reduce system burden. However, not all lab tests are suitable for home-based services. Perception-based prioritization can guide public health systems in making data-informed decisions on which tests should remain hospital-based and which could be decentralized.
Objective:
To evaluate and prioritize 210 laboratory tests based on perceived (1) testing complexity, (2) result urgency, and (3) at-home testing value
and to establish a foundation for developing a machine learning model to automate prioritization, predicting suitability for decentralized testing, enabling scalable, and evidence-based policy decisions.
Methods:
A Delphi study was conducted with 18 laboratory professionals (mean experience: 7.8 years; range: 5–20 years) across 7 clinical and technical units at Rajavithi Hospital, Thailand. Participants rated 210 tests on 5-point Likert scales for complexity, urgency, and at-home value. Mean Round 3 scores were used as final consensus. We also calculated changes (Δ) across rounds to assess convergence and identify tests with variable expert perceptions.
Results:
After data cleaning, 174 tests were analyzed. Round 3 means were: complexity = 2.33, urgency = 1.98, and at-home value = 2.92. NT-pro BNP emerged as a balanced candidate for staff-led home testing. Tests like Glucose Strip (high urgency, low complexity) are better suited for self-testing. Notable variability was seen in Procalcitonin, HIV Ag/Ab Combo, and Coombs Test, indicating evolving expert consensus.

Conclusion:
Delphi-derived perceptions offer a structured foundation for training machine learning models to support scalable, data-driven decentralization of diagnostic services.

6 - Scheduling Urgent Patients with Different Priorities in Dedicated Emergency Capacity

Oki Almas Amalia, Universitas Gadjah Mada, Yogyakarta, Indonesia, Maartje van de Vrugt, Jasper Bos, Irwan Endrayanto, Richard Boucherie

We present two extreme policies to schedule urgent and semi-urgent patients to dedicated operating room (OR) capacity: last-minute scheduling and near-online scheduling. Under *last-minute scheduling*, the scheduler assigns capacity to the waiting patients just before the start of the time-slot in which the patient will be treated. Under *near-online scheduling*, the scheduler assigns a future time-slot to the patient shortly after the patient's arrival. For both scheduling policies we develop a Markov Decision Process (MDP). As emergency capacity is fairly restricted and the number of shifts a semi-urgent patient may be delayed is small, the MDPs may be solved to obtain explicit policy tables that allow us to make the trade-off between increasing patient-friendliness through near-online scheduling, and resource efficiency under last-minute scheduling. We demonstrate the usefulness of our method by applying it to the emergency OR of our partnering hospital and compare it with some heuristics that are typically applied by the hospitals.

SC09

Bras Basah B

Reliability and System Performance Analysis

Contributed Session

Manufacturing and Service Operations Management

Chair: Elham Heydari, York University, Guelph, ON, Canada

1 - Robotics in Retail: Identifying the Threshold for Feasibility in Order Fulfillment

Sebastian Koehler, Karlsruhe Institute of Technology, Karlsruhe, Germany, René de Koster

Offering omnichannel services has become a major challenge for retailers, especially in developed countries where labor shortages and low profit margins are common. One strategy to provide additional services is to use existing store infrastructure to fulfill online orders, as stores are typically located close to customers. In this study, we explore how robotic solutions can support and improve the logistics of omnichannel services such as Buy-Online-Pick-Up-In-Store (BOPIS) and Ship-From-Store (BOSS). Therefore, we evaluate a hybrid (robot–human) picker-to-parts order fulfillment system within an in-store environment. Our research is guided by the following research question:

Which robotic capabilities are required to make hybrid in-store omnichannel retail operations beneficial compared to purely manual operations?

To do so, we use an analytical Semi-Open Queuing Network (SOQN) modeling approach. We develop an automated state-space generator to create the input for solving the steady-state probabilities of the SOQN using the Matrix Geometric Method. This allows us to model and analyze system behavior under various conditions, such as different numbers of robots and varying capability levels. We apply our model to a real-world case study from a European drugstore chain to demonstrate how the proposed hybrid system would perform in practice and compare it with a fully manual system. A key focus of our evaluation is to understand the impact of robotic integration on workload and financial outcomes.

2 - Realizing the Supply Chain Benefits Through the Affordance of Artificial Intelligence: Affordance and Adaptive Structuration Theory Integrated Perspective

Qingyun Yang, Xi'an Jiaotong University, Xi'an, China, People's Republic of

Artificial intelligence (AI) has been widely recognized for its potential in supply chain management, but its empirical benefits remain underexplored. This study examines the affordance of AI - autonomy and interactivity

- in enhancing supply chain performance (SCP). Integrating adaptive structuration theory (AST), we analyze both direct and indirect effects of AI on SCP. We propose that companies can improve SCP by adjusting social structures to better leverage the affordance of AI. Social structure results from the social results generated by the application practices of technology. Organizations can only exert the dominant influence of technology when they follow the usage cues it provides, which relies on their understanding of the technology's capabilities and affordances. Supply chain entrainment (SCE) and supply chain collaboration (SCC) elucidate how to organize social structures to improve SCP based on leveraging the affordance of AI. Our study posits that SCE and SCC serve as a bridge between the AI adoption and SCP. Using data from 249 Chinese manufacturing firms, we test these hypotheses with structural equation modeling and interaction analysis. Results show that AI impacts SCP indirectly through SCE and SCC, which positively moderate the AI-SCP relationship. This study not only contributes to the literature on AI's impact in SC management, the application of affordance and AST, but also provides practical guidance for firms on how to unlock the potential of AI to achieve SC benefits.

3 - Shop Floor Management in Manufacturing Systems: a Decision Maker&Rsquo;S Perspective

Upasana Haldar, Indian Institute of Technology Kharagpur, Kharagpur, India

The dominance of various uncertainties in the shop floor results in unpredictability of performance of the system and makes decision-making an intricate task. Industry 4.0 exploits an unprecedented growth in agility within the manufacturing systems results in ever increasing demand for real-time data-driven decision-making mechanisms, as a result smart analytics is gaining the center stage and its evolving. The current work aims to facilitate real-time control of manufacturing operations by connecting the two vital aspects of manufacturing of future i.e. digital simulation and advanced analytics. The work attempts to develop a framework for designing a prescriptive decision support system for a manufacturing shop floor management which considers production scheduling, maintenance and quality functions in conjunction and rely on the perspective of a shop floor manager's decision-making dilemma. So, a predictive control mechanism capable to predict any undesirable transition on the shop floor due to uncertainties related to the shop floor functions on a projected timeline and identify those that needs intervention is developed. Mathematical models and simulation-based assessment has been proposed to develop the prediction scheme. Intervention scheme is heuristics based and depending on the outcomes of various interventions, design of an indicative prescription is generated. Hence, the flow of the work initiates from designing prediction scheme to intervention scheme which will lead to generate the prescription. This mechanism can be an effective tool for the shop floor decision-making, enhance the system preparedness, better resource management and avoid any impending issues on the shop floor.

4 - Leveraging Data Contributors to Enhance Social Welfare Through Crowdsensing

Elham Heydari, York University, Toronto, ON, Canada, Hossein Abouee Mehrizi, Mehdi Nourinejad

Problem definition: Crowdsensing collects data from individuals participating in an activity to obtain information about a system, eliminating the need for costly data collection infrastructure. This paper seeks to answer whether increasing crowdsensing data providers is always socially beneficial in congestion-sensitive systems like queues. Methodology/results: We use a single-server queue with Poisson service demand and completion rates to examine how social welfare changes based on the ratio of crowdsensing data contributors. We assume all customers have free access to crowdsensing information, but only a subset reports their observed queue length. We analyze the impact of data contributors on customer decisions to join a queue in a two-stage decision-making process. Managerial implications: The effect of data contributors on social welfare is influenced by the system load factor. When the load factor is high, there exists an optimal ratio of data contributors. Increasing the data contribution beyond its optimal leads to frequent updates about the queue, which encourages customers to join at every opportunity, increasing wait times and reducing social welfare. Below the optimal value, the system cannot provide accurate enough information to customers, which leads to socially inefficient decisions. In contrast to high demands, social welfare increases monotonically with the ratio of data contributors in low-load factor systems as it avoids congested queues. Furthermore, the findings suggest that providing public access to crowdsensing information without mandatory data contributions maximizes social welfare. In contrast, restricting customer access to crowdsensing information diminishes social welfare and reduces system throughput.

SC10

Hullet Room

Social Factor in OR

Invited Session

OR and Social Wellness

Chair: Quanyan Zhu, New York University, New York, NY, United States

1 - The Paradox of AI Adoption in Healthcare: Operational Efficiency vs. Patient Wellbeing

Wentao Zhan, Harbin Institute of Technology, Harbin, China, People's Republic of, Xuchuan Yuan, Minghui Jiang

The widespread adoption of Artificial Intelligence (AI) in medical diagnostics has prompted extensive exploration of its potential to enhance system performance and patient utility. This study constructs a medical system consisting of two stages-detection and treatment-to investigate the impact of AI-assisted diagnostics on system performance and patient utility. AI-driven diagnostics can significantly improve detection efficiency, with both hospitals and social planners optimizing AI capability and detection time standards to maximize customer surplus. The results indicate that when diagnostic results are not mutually recognized, AI capability consistently leads to higher total customer surplus for hospitals in a fully covered market. However, surprisingly, in a partially covered market, increasing AI capability may have a detrimental effect on total customer surplus. In such cases, the overall customer surplus of the medical system initially increases with AI capability but eventually declines, possibly resulting in a lower total customer surplus for customers using AI diagnostics than for those not using AI. When diagnostic results are mutually recognized, the negative effects of AI capability not only impact hospitals using AI but also spill over to other hospitals. Specifically, hospitals not using AI tend to experience the negative effects of increased AI capability earlier. Furthermore, in a partially covered market, it is found that regardless of whether a hospital uses AI diagnostics, total customer surplus follows an inverted U-shaped pattern, revealing a paradox: while AI enhances diagnostic efficiency, it also increases the burden on the treatment stage, leading to a mismatch between the two stages.

2 - A User-Centric Assistant for Interactive Mixed-Integer Optimisation Modelling

Jyotheesh Gaddam, Deakin University, Melbourne, Australia, Vicky Mak-Hau, Lele Zhang, Diego Molla-Aliod, Bahadorreza Ofoghi

Optimisation is a powerful tool for solving decision-making problems across industries, from manufacturing and logistics to resource planning. However, defining an optimisation problem, formulating a mathematical model, and solving it can be very challenging, particularly when the problem is complex. In this work, we introduce MILO (Mixed-Integer Linear Optimisation), a smart assistant designed to guide users through the process of building Mixed-Integer Linear Programming (MILP) models for combinatorial optimisation problems, especially when users lack a background in operations research.

MILO is implemented as a conversational agent built on the AutoGen multi-agent framework. It leverages large language models (e.g., GPT-4o), advanced prompt engineering, and dynamic beam search over structured domain knowledge graphs (KGs) to interact with users in a context-aware and adaptive manner. This architecture enables MILO to ask “right” questions in a logical order, based on user responses and the evolving problem context. The system draws on a built-in understanding of various optimisation domains, such as production planning and routing, to ensure all key components of the MILP are captured. As users respond, MILO constructs a clear problem summary and prepares the foundation for a formal MILP model.

By focusing on interaction with the user, guided by KGs, MILO makes the model-building process intuitive and educational. It empowers users to define optimisation problems, avoid modelling errors, and gain insights into their scenarios. MILO is flexible and modular and can be extended to support new domains and contexts, making optimisation more accessible even for those without coding or advanced mathematical backgrounds.

3 - The Effect of Social Influence on Collective Forecasts

Steffen Keck, University of Vienna, Wien, Austria, Wenjie Tang

We explore the effect of social influence, in the form of information about other individuals' estimates, on the accuracy of collective forecasts formed by the statistical aggregation of several individual predictions. Motivated by tension in the prior literature, we attempt to better delineate the relative effects of social influence by focusing on the specific setting in which forecasts are made.

Specifically, we suggest that when the individual forecasters make independent individual predictions and are incentivized accordingly, social influence will be generally beneficial for collective accuracy. However, when forecasters know that their individual estimates will be aggregated and are paid based on collective accuracy, social influence will be detrimental. This is because in the former situation, individuals will perceive social influence as mostly informational and update their judgments depending on the perceived quality of this information. By contrast, in the latter setting, individuals will perceive social influence as mostly normative and will more uniformly adapt majority opinions.

To test this suggestion, we conducted an experimental study in which participants made forecasts for the outcomes of soccer games played during the group stage of the 2024 UEFA Euro. As predicted, we find that when forecasts are made in an independent individual setting social influence increased the accuracy of collective forecasts. By contrast, for forecasts made in an explicitly collective setting social influence had the opposite effect. Moreover, we found that this effect was driven by the different relative effects of social influence on forecast correlation and individual accuracy across the two settings

4 - Unlocking the Green Inertia of Metropolitan Area Travelers: the Hidden Motivations of Their Continuous Intentions on Low-Carbon Travel

Wenting Zhan, Sichuan University, Chengdu, China, People's Republic of, Yubo Guo, Jinchan Liu

As metropolitan areas transcend traditional administrative boundaries, they emerge as fundamental spatial units, fostering new functional relationships and interconnections. Decoupling transportation development from greenhouse gas emissions is critical to meeting the growing intercity travel demand within these regions. While prior research highlights travelers' willingness to adopt low-carbon transportation, few studies explore their continuous intention (CI) to sustain such choices, particularly in metropolitan contexts. This study investigates the underlying motivations driving travelers' CI toward low-carbon travel in metropolitan areas by integrating expectation confirmation theory, the theory of planned behavior, and mind-flow theory. Using structural equation modeling, the proposed model was empirically validated with data from 1028 travelers in the Chengdu-Chongqing Economic Circle (CCEC). Findings reveal that perceived enjoyment (PE), intrinsic motivation (IM), and perceived usefulness (PU) are the primary determinants of CI. PE and IM indirectly shape CI by influencing attitudes (ATT), while expectation confirmation (EC) and PU indirectly affect CI through satisfaction (SAT). This study elucidates these relationships and offers evidence-based policy recommendations to promote sustained low-carbon travel behaviors in metropolitan areas, thereby advancing environmentally sustainable urban mobility.

5 - Cognitive Arbitrage in Cyber Warfare: a Cross-Layer Game-Theoretic Framework for Strategic Deception

Quanyan Zhu, New York University, New York, NY, United States, Ya-Ting Yang

Cognitive arbitrage refers to the strategic exploitation of cognitive asymmetries between adversaries to gain decision advantage. In cybersecurity, this concept is particularly powerful: defenders can both leverage inherent cognitive vulnerabilities in attackers and actively create cognitive asymmetries through cyber deception. By deploying deceptive artifacts such as honeypots, fake data paths, or misinformation banners, defenders can manipulate the attacker's perception of the network, concealing the true location of valuable assets.

Achieving cognitive arbitrage requires a cross-layer approach that integrates strategic planning and tactical execution. At the strategic layer, defenders must determine the optimal timing for deploying a limited set of deception techniques, recognizing that once a deception is exposed, it loses future effectiveness. At the tactical layer, the challenge is to configure and apply each deception mechanism in a manner that maximizes its impact,

accounting for the attacker's beliefs, behavior, and situational context. These two layers are deeply interdependent, requiring a holistic co-design framework.

In this talk, we present a cross-layer Stackelberg game-theoretic framework that formalizes the joint planning and execution of cyber deception under system constraints and cognitive dynamics. The framework highlights the pivotal role of cognitive arbitrage in modern cyber warfare and enables principled reasoning about timing, belief manipulation, and adversarial response. We also introduce a Parrondo-like paradox in which individually losing tactical engagements, when orchestrated strategically, can yield a net advantage at the strategic level. This insight highlights the need for quantitative tools to analyze and design deception-based cyber defense strategies in the evolving landscape of cyber conflict.

SC11

Morrison Room

Recent Advances in Airline Revenue Management I

Invited Session

Revenue Management and Pricing

Chair: Lan Wu, National University of Singapore, Singapore, Singapore

1 - Modeling and Optimizing Flight Delay Propagation in Airline Networks Under Stochastic Block TIMES

Jingying Li, National University of Singapore, Singapore, Singapore, Wenxuan Shan, Chung Piaw Teo

Flight delays significantly disrupt airline operations, stemming from both external (primal) and internal (propagated) factors within an airline's network. Accurately modeling delay propagation is thus crucial for effective air traffic management. This study models aircraft rotations as a spatiotemporal network, representing scheduled flights as nodes and delay propagation as dynamic edges using a novel network flow framework. Assuming stochastic block times, we characterize primal delays by their first and second moments without specifying their exact distributions. Under these conditions, we develop an exact formulation to compute the maximum expected total arrival delay time for a given flight network, incorporating both primal and propagated delays. This formulation is converted into a convex conic optimization problem, further relaxed into a tractable semidefinite relaxation, providing near-optimal solutions.

Building on this model, we propose a two-stage stochastic optimization approach to minimize the maximum expected total arrival delay of scheduled flights under fleet size constraints. A branch-and-price algorithm is employed to efficiently identify optimal aircraft routes, facilitating the solution of large-scale assignment problems. Using synthetic data generated from Singapore Airlines, the proposed algorithm efficiently allocates flights across fleets of up to 24 aircraft within a 30-day planning horizon. Our results reveal a critical threshold in fleet size utilization: beyond this point, delays grow linearly; below it, they increase exponentially. This threshold offers practical insights for fleet management, indicating the minimum fleet size required to handle scheduled operations effectively, or the maximum feasible task set that can be supported by a given fleet.

2 - Feature-Based Multi-Product Pricing for Airline Luggage Ancillaries

xiao JIN, National University of Singapore, Singapore, Singapore, Junjie Ren, Chung Piaw Teo

We introduce a feature-based pricing approach for optimising ancillary product prices, focusing on luggage in low-cost carrier airlines. We employ the Marginal Exponential Model, a discrete choice model, to capture multi-product substitution effects. Our framework integrates a linear expression to address customer heterogeneity and reduce the computational complexity of segment-level optimisation. We validate the method through real-world field tests, demonstrating that it consistently outperforms off-the-shelf pricing rules even when sales data are limited. These results highlight the importance of accurately modelling substitution effects and customer heterogeneity for effective pricing, establishing our approach as a robust, data-driven alternative to conventional methods.

3 - Net Promoter Score as a Driver of Revenue Growth

Junjie Ren, National University of Singapore, Singapore, Singapore, Chung Piaw Teo

Net Promoter Score (NPS) is one of the most widely used metrics in business practices to gauge customers' willingness to recommend, and thus customer loyalty and satisfaction, but its validity as a predictor of revenue remains disputable. We provide causal evidence of NPS effect on revenue growth using data from a major airline. By utilizing flight delays as instruments, we find that an increase of one percentage point in NPS has a lagged positive effect on cumulative monthly revenue, peaking at approximately 0.5% after two months. Moreover, in a field experiment offering passengers free access to on-board Wi-Fi, we document evidence of NPS mediating the treatment effect on revenue growth.

4 - Price Genie

Aadhar Dutta, National University of Singapore, Singapore, Singapore

This paper introduces a feature-based pricing approach for optimizing ancillary product prices, with a focus on luggage in Low-Cost Carrier (LCC) airlines. We employ the Marginal Exponential Model (MEM)—a discrete choice model—to capture multi-product substitution effects, marking its first application in airline ancillary pricing. Our framework integrates a linear term to address customer heterogeneity and reduce the computational complexity of segment-level optimization. We validate the method through real-world field tests, demonstrating that it consistently outperforms off-the-shelf pricing rules even when sales data are limited. These results highlight the importance of accurately modeling substitution effects and customer heterogeneity for effective pricing, establishing our approach as a robust, data-driven alternative to conventional methods.

SC12

Moor Room

Empirical Revenue Management and Pricing

Invited Session

Revenue Management and Pricing

Chair: Freddy Lim, National University of Singapore, Singapore, Singapore

1 - Does Human-AI Co-Creation Drive Wage Premium? Evidence from Freelance Marketplaces

Jinya Hu, National University of Singapore, Singapore, Singapore, Khim Yong Goh, Chenshuo Sun

With generative artificial intelligence advancing rapidly, collaboration between human creators and AI in creative projects is on the rise. However, the effects of this human-AI co-creation (HAICC) on freelancer wages remain largely unexplored. This study investigates freelance service marketplaces to understand the wage outcomes resulting from creative freelancers' shift to HAICC. Using data from a leading freelancer platform where some creative freelancers staggered their shift to the provision of HAICC services, we employ the difference-in-differences identification method to examine the effects of the HAICC shift by creative freelancers on their wage outcomes. Our findings reveal that the transition to HAICC leads to a wage premium of $71.285, equivalent to a 32.1% increase in monthly wages for creative freelancers. Furthermore, our analysis

demonstrates that HAICC shift is associated with a significant increase in the visual complexity of creative work, as well as enhancements in novelty and emotion elicitation of the work produced. Interestingly, we do not observe a statistically significant improvement in aesthetic quality. Last, our findings show the HAICC-induced wage premium is moderated by how freelancers position their expertise in service descriptions. Emphasizing domain expertise results in higher HAICC-induced wage premium, while emphasizing AI expertise leads to diminishing returns with an inverted U-shaped moderating effect. This study provides insights for creative freelancers on how HAICC affects their wages and creative work, while also highlighting its revenue implications for freelancer platforms that facilitate HAICC-related transactions.

2 - Does Increased Platform Control over Complementors' Operations Lead to Better Vendor Product Performance? Evidence from Mandatory Semi-Managed Services on Aliexpress

Liang Zhao, Hong Kong Baptist University, Hong Kong, Hong Kong, Jicheng Zeng, Ling Xue, Xu Huang

In the evolving landscape of cross-border e-commerce, platforms are increasingly seeking a new equilibrium between control and autonomy to better navigate the complexities of this environment. Many platforms are experimenting with governance models, referred to as semi-managed services, that enhance their control over operational components while allowing vendors to retain autonomy over core business functions. However, the effects of such governance adjustments on vendor performance are not well understood. This research investigates the impact of mandatory semi-managed services on vendor product performance in cross-border e-commerce. Specifically, we examine how such services, which grant platforms greater control over certain operations, including logistics and marketing, influence product sales and consumer satisfaction. We utilized a difference-in-differences (DID) approach to analyze data from 6,000 products listed on both AliExpress and Walmart platforms, using AliExpress's mandatory adoption of a semi-managed service model as a natural experiment. Our findings indicate that while the implementation of mandatory semi-managed services led to an increase in sales, it simultaneously caused a decline in consumer satisfaction. Utilizing deep learning techniques for segment analysis and text analysis based on online reviews, we further find that this negative impact on consumer satisfaction was particularly pronounced among high-rating small vendors, who are vulnerable to increased operational risks and diminished control over their operations. This study contributes to platform governance literature by providing empirical evidence on the effects of platform intervention and offers practical insights for platform owners and policymakers.

3 - Points Plus Money: Uncovering the Impact of Mixed Currency Redemptions

Freddy Lim, National University of Singapore, Singapore, Singapore, So Yeon Chun

In recent years, loyalty programs have rapidly expanded in scale. Loyalty points effectively serve as a new currency, and consumers often use loyalty point currency in the redemption of goods and services. Despite this redemption growth, consumers can still find it difficult to attain the points required to redeem rewards. In response, some firms allow consumers to perform mixed currency (points plus money) redemptions, instead of full point redemptions.

This study examines the impact of mixed currency redemptions on consumer behaviors and firm revenue, using data from a major loyalty program. We adopt an instrumental variables control function approach to model consumers' endogenous selection to perform mixed currency redemptions. We find that mixed currency redemptions significantly increase consumers' direct point purchases with the firm compared to full point redemptions, which increases firm revenue. Interestingly, earnings from traditional accrual channels such as credit cards remain similar between consumers using mixed currency and full point redemptions. This suggests that mixed currency redemptions deliver comparable impact as full point redemptions within traditional point earning channels, but uniquely drive higher levels of direct point purchases.

Although mixed currency redemptions may encourage consumers to drain their low point balances and become inactive, we find no evidence to support this. On the contrary, consumers who drain their balances through mixed currency redemptions remain engaged in the program. We also find that a consumer segment which is often overlooked in marketing efforts contributed significantly to the revenue increase. Finally, we consider plausible mechanisms that drive these results.

4 - Unpacking the Impact of Product Variety on the Demand for Online Professional Services

Qingyu Xu, City University of Hong Kong, HONG KONG, Hong Kong, Jingyi Jiang, Hongfei Li, Weiguang Wang

Online platforms are increasingly offering professional services with credence attributes, yet the impact of product variety on demand in such contexts is unclear. Although expanding product variety often benefits search and experience goods, even a small increase in service variants can amplify information asymmetry and lead to choice deferral for professional services, whose quality is difficult to evaluate. This study investigates whether and how variety expansion affects demand under such conditions, focusing on a new feature introduced in August 2021 by one of China's largest online fitness platforms, which allowed users to adjust the difficulty level of select courses. Contrary to expectations that more options would boost uptake, our findings reveal that offering adjustable difficulty had an overall negative effect on demand for treated courses. However, this negative effect was less pronounced for courses associated with higher credibility signals in the platform's reputational system—namely, platform verification of instructor credentials, seller self-reported details, and fit-oriented user reviews. In particular, courses that featured detailed attribute information or robust user feedback emphasizing "who this difficulty suits" experienced smaller drops in demand or even a positive effect. These asymmetric treatment effects underscore that information content strategies can mitigate the information asymmetry underlying choice deferral. Our results offer important insights for the design of reputational systems in online professional service platforms, showing how transparent and user-centered information can offset the potential downsides of increased variety.

SC13

Bailey

Advancing Decision-Making through Simulation and Optimization

Invited Session

Simulation and Optimization of Healthcare Systems

Chair: Tianxiang Wang, East China Normal University, Shanghai, N/A

1 - Thompson Sampling-Based Policy for Dynamic Assortment Optimization in Live-Streaming Sales

Haidong Li, University of Chinese Academy of Sciences, Beijing, China, People's Republic of, Zishi Zhang, Ying Liu, Yijie Peng

This study explores the dynamic assortment optimization problem in the context of live-streaming sales. In response to the dynamic characteristics of live-streaming e-commerce, we propose a new choice model that extends the traditional Multinomial Logit model into the continuous time domain. This novel model enables the decoupling of parameter estimation for different products and facilitates the convenient updating of posterior parameters at any time. Finally, we introduce a myopic optimization algorithm based on Thompson Sampling. This algorithm effectively balances exploration and exploitation, captures fluctuations in the number of audiences, and demonstrates superior numerical performance.

2 - Multi-Product Inventory Optimization with Returns and Exchanges via Green Simulation

Haitao Liu, Sichuan University, Chengdu, China, People's Republic of

Motivated by the high consumer returns in E-retailing operations, we study a multi-periodic multi-item inventory control problem with item returns and exchanges. After receiving products, each customer has the right to randomly keep, return, or exchange purchased products within a designated time window. After remanufacturing and relabeling, these returned products can be restocked for resale. Unlike the stylized models, we do not impose probabilistic assumptions about demand, return, and exchange processes, and allow them to be either stationary or nonstationary. The complex interaction and correlation among demand, returns, exchanges, and inventory make inventory optimization intractable. To this end, we propose a green dynamic learning framework that experimentatize inventory alternatives, in which the main idea is reusing sample trajectories to evaluate alternatives based on low-fidelity models. A novel estimator for inventory profit combining observations given by real environments and low-fidelity models is proposed. Under mild conditions, we show that this estimator is unbiased. Also, we prove that the number of learning rounds allocated to each inventory alternative exhibits a

linear growth rate, and the asymptotic allocations converge to the optimal computing budget allocation. Numerical experiments have been conducted to analyze the impact of returns and exchanges and evaluate the finite-time performance of the proposed method.

3 - Ranking and Selection with Two-Stage Decision

Tianxiang Wang, East China Normal University, Shanghai, China, People's Republic of

Ranking & selection (R&S) is concerned with the selection of the best decision from a finite set of alternative decisions when the outcome of the decision has to be estimated using stochastic simulation. In this paper, we extend the R&S problem to a two-stage setting where after a first-stage decision has been made, some information may be observed and a second-stage decision then needs to be made based on the observed information to achieve the best outcome. We then extend two popular single-stage R&S algorithms, expected value of information (EVI) and optimal computing budget allocation (OCBA), to efficiently solve the new two-stage R&S problem. We prove the consistency of the new two-stage EVI (2S-EVI) and OCBA (2S-OCBA) algorithms. Experiment results on benchmark test problems and a two-stage multi-product assortment problem show that both algorithms outperform applying single-stage EVI and OCBA in the two-stage setting. Between 2S-EVI and 2S-OCBA, numerical results suggest that 2S-EVI tends to perform better with smaller number of decisions at first and second stage while 2S-OCBA has better performance for larger problems.

4 - Rate-Optimal Strategy for Best Policy Identification in Reinforcement Learning

Mingjie HU, Department of Management Science, Fudan University School of Management, Shanghai, China, People's Republic of, Edward Huang, Jie Xu, Kan Wu, Jian-Qiang Hu, Chun-Hung Chen

Best policy identification (BPI), which seeks to identify the policy with the highest expected cumulative reward, has recently gained considerable attention in reinforcement learning (RL). This paper formulates a BPI problem for finite-horizon, time-inhomogeneous, tabular Markov decision processes (MDPs) within a fixed-budget ranking and selection (R&S) framework, with the goal of maximizing the probability of identifying the best policy. To tackle this problem, we first derive a rate-optimal sampling ratio using large deviation theory, incorporating the dynamics and rewards of MDPs into the rate function. We then quantify the benefits of incorporating MDP structural information and show that the optimal sampling ratio balances variance, optimality gap, and visitation probabilities. Motivated by abundant covariate information in practical decision-making, we propose a contextual BPI framework with linear reward functions and extend the optimal sampling ratio to this setting. In addition, we introduce an asymptotically rate-optimal strategy and an efficient approximate strategy based on a closed-form sampling ratio. Numerical experiments demonstrate that our proposed algorithm outperforms existing R&S and RL methods. Furthermore, we illustrate its practical applicability through a complex diffusion furnace scheduling problem. These results highlight the potential of R&S techniques for solving BPI problems.

5 - Optimal Control of Patient Flow in the Emergency Department with Triage Standing Orders

Shuangjie Huang, Xiamen University, Xiamen, China, People's Republic of, Weifen Zhuang

Triage Standing Orders (TSO) in the emergency department (ED) involves triggering lab tests for patients by a nurse at triage prior to their visit with a physician. TSO has the potential to alleviate ED congestion by diverting patients; however, it may also lead to unnecessary tests resulting from triage errors. Additionally, physicians must dynamically decide which queue to serve, initial assessments or re-assessment patients with lab tests. In this study, we formulate the optimal patients sequencing problem in the ED as a Markov Decision Process (MDP) to examine the effects of TSO. Our analysis reveals that the optimal patients sequencing policy is a switching-curve policy depending on the queue lengths. Comparing system performance with and without TSO, we identify specific scenarios where implementing TSO proves beneficial. We develop efficient heuristics and evaluate their performance through numerical study, demonstrating their effectiveness in alleviating ED crowding.

SC14

Wilberforce

Simulation and Learning with Rare Events

Invited Session

Simulation Society

Chair: Anand Deo, Indian Institute of Management Bangalore, Bangalore, India

1 - Scaling Scenario-Based Chance-Constrained Optimization Under Rare Events

Anand Deo, Indian Institute of Management Bangalore, Bangalore, India, Jaeseok Choi, Anirudh Subramanyam, Constantino Lagoa

Chance-constrained optimization is a suitable modeling framework for mitigating extreme event risk in many practical settings. The scenario approach is a popular solution method for chance-constrained problems, due to its straightforward implementation and ability to preserve problem structure. However, for safety-critical applications where violating constraints is nearly unacceptable, the scenario approach becomes computationally infeasible due to the excessively large sample sizes it demands. We address this limitation with a new yet straightforward decision-scaling technique. Our method leverages large deviation principles and relies on only mild nonparametric assumptions about the underlying uncertainty distributions. The method achieves an exponential reduction in sample size requirements compared to the classical scenario approach for a wide variety of constraint structures, while also guaranteeing feasibility with respect to the uncertain constraints. Numerical experiments spanning engineering and management applications show that our decision-scaling technique significantly expands the scope of problems that can be solved both efficiently and reliably.

2 - Quantifying Asymmetry in Extreme Events with Applications to Cryptocurrency Price Movements

Bikramjit Das, Singapore University of Technology and Design, Singapore, Singapore, Xiangyu Liu

Measures of asymptotic tail dependence between two variables often tend to

be symmetric in nature. We define a measure of extreme tail association between

two variables based on the upper tail quantiles of the data, which is asymmetric and easily computable in sample data. Under mild conditions on the underlying distribution we show that the sample tail association measure defined is consistent which allows us to create confidence bounds around the measure. We also propose a test for bivariate tail asymmetry and devise bootstrapping methods to perform the test in real data sets. Data from movement in cryptocurrency prices is used to exhibit the

efficacy of the tools developed.

3 - Self-Similar Gen: Simulating Tails of Heavy-Tailed Distributions Using Generative Models

Mantu Gupta, Indian Institute of Management Bangalore, Bangalore, India

We introduce Self-Similar (SS)-Generative Model: a novel procedure for estimating rare-event probabilities in multivariate heavy-tailed settings. Our procedure leverages the inherent semi-parametric nature of multivariate heavy-tailed distributions to learn a non-parametric component using deep generative models (DGMs), such as Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), or Diffusion Models. We demonstrate that this approach enables accurate generative modeling for the extreme regions of multivariate heavy-tailed data.
We establish the statistical consistency of our estimates and provide convergence rate guarantees under mild smoothness assumptions on the true distribution. Unlike state-of-the-art methods that rely on specialized neural architectures for tail estimation or require parametric structure assumptions, our framework is flexible and can be integrated with various deep generative models.
We showcase the applicability of our methodology by developing a tractable approach to compute rare-event probabilities far beyond the range of observed data.

4 - Data-Driven Sequential Sampling for Tail Risk Mitigation

Dohyun Ahn, CUHK, Shatin, Hong Kong, Taeho Kim

Given a finite collection of stochastic alternatives, we study the problem of sequentially allocating a fixed sampling budget to identify the optimal alternative with a high probability, where the optimal alternative is defined as the one with the smallest value of extreme tail risk. We particularly consider a situation where these

alternatives generate heavy-tailed losses whose probability distributions are unknown and may not admit any specific parametric representation. In this setup, we propose data-driven sequential sampling policies that maximize the rate at which the likelihood of falsely selecting suboptimal alternatives decays to zero. We rigorously demonstrate the superiority of the proposed methods over existing approaches, which is further validated via numerical studies.

SC16

VIP A

Sustainable Supply Chain and Operations Management

Contributed Session

Sustainable and Socially Responsible Operations

Chair: Alexandra Yizhuo Dong, UCLA Anderson, Los Angeles, CA, 90024, United States

1 - Digital Breakthroughs in Drifting Green Governance: Synergistic Inhibition Effects of Financing Constraints and Information Asymmetry in a Dynamic Capability Perspective

yan du, University of Science and Technology Beijing, beijing, China, People's Republic of

Corporate greenwashing has emerged as a major environmental governance concern in the context of the global economy's shift to a green and low-carbon economy. This study employs a sample of Chinese A-share listed businesses from 2012 to 2022 to methodically investigate the governance mechanism of digital transformation on greenwashing behaviour, drawing on the dynamic capability perspective. By building ESG disclosure and performance gap assessing greenwashing indicators and using annual report text mining technology to evaluate the degree of digital transformation, the governance mechanism is investigated.According to the study's findings, corporate greenwashing may be effectively discouraged by digital transformation; for every 1% rise in digital transformation, there is a 1.9% decrease in greenwashing. The results show that the relationship between corporate greenwashing and digital transformation is mediated by knowledge asymmetry and financial limitations. According to a heterogeneity analysis, the negative relationship between corporate greenwashing behaviour and digital transformation is more noticeable in non-high-tech companies, heavily polluted industries, and areas with strong environmental regulations, strong environmental concerns, and weak environmental concerns. A theoretical basis for the creation of a digital and green synergistic governance system is provided by this paper, which clarifies the underlying logic of environmental governance enabled by digital technology. It also provides policy insights for developing nations looking to tackle the problem of greenwashing.

2 - R&D Alliances Under the AI Flywheel Effect: Equity Joint Venture vs. Contractual Partnership

Feiying Jiang, Huazhong University of Science & Technology, Wuhan, China, People's Republic of, Jun Yang

The AI flywheel effect, wherein artificial intelligence (AI) performance improves through iterative user data feedback and sustained algorithm efforts, is reshaping inter-firm collaboration in AI development. This paper examines a R&D alliance between an AI firm and a service firm for the joint development and commercialization of a domain-specific foundation model. The AI firm provides algorithmic expertise, while the service firm contributes domain-specific data and marketing resources. Under alliances initiated either by the service firm or the AI firm, we explore how the technical characteristics of the AI flywheel effect affect the choice between two alliance forms across industries: Equity Joint Venture (EJV) and Contractual Partnership (CP). We find that under alliances initiated either by the service firm or the AI firm, in an industry with high industry admission standards for the foundation model, CP is preferred. In an industry with low industry admission standards, stronger bargaining power drives the initiator toward EJV. Firms exhibit divergent alliance form preferences, leading to profit conflicts. Under the service firm-initiated alliance, higher data-driven efficiency in the AI flywheel effect increases the preference for CP, despite the associated risks of profit conflicts and the potential to constrain AI development. Surprisingly, even improvements in algorithm-driven efficiency, which seem to strengthen the AI firm’s position, may induce the service firm to choose CP, thereby further exposing the AI firm to potential profit erosion. The effects of data-driven efficiency and algorithmic-driven efficiency on the choice between two alliance forms are reversed under the AI firm-initiated alliance.

3 - Indian Start-Ups Adopting Circular Economy for Creating Sustainable Products

Bhavani Shankar Saripalli, Indian Institute of Management Indore, Indore, India, Vinaysingh Chawan

Young Indians are turning to various waste materials to create sustainable products for regular use. These efforts range from simple reimagination of the waste to physical or chemical transformation through innovative methods. The start-up ecosystem in India is encouraging, enabling these efforts to come to fruition. For example, some of the start-ups make tiles from carbon powder, repurpose cigarette butts as material for toys, produce energy from left over vegetables, use discarded flowers to make incense sticks, produce shoes from plastic bags, manufacture disposable cutlery from bagasse, and yarn produced from water hyacinths to interweave in sarees. Though in India, recycling and repurposing is not done at a large scale, these organisations are pioneering the efforts in saving the planet one garbage pile at a time.

The conscientious Gen-Z consumers in India who are witnessing the adverse impact of climate change are interested in trying out such products creating scope for expanding markets. The processes followed by these organisations pave way for future entrepreneurs to think of new ways of utilising waste materials to produce something useful, thereby changing the perspective of society towards waste. We study these organisational efforts to see the processes followed, low-cost methods of processing adopted, labour intensive ways in which low-cost workforce is utilised, etc to determine their lasting impact on innovative waste disposal. In a way these efforts are sowing seeds for zero-waste life ideology.

KEYWORDS

Sustainable products, waste management, start-up ecosystem, value from waste

4 - From Waste to Taste: Structural Estimation and Market Design in Secondary Food Markets

Yizhuo Dong, UCLA Anderson, Los Angeles, CA, United States, Auyon Siddiq, Jingwei Zhang

Commercial food waste presents a significant challenge to urban sustainability, with one-third of food wasted annually, and 31% lost at retail and consumer levels. Recently, two-sided marketplaces have emerged to connect retailers with consumers seeking surplus food at discounted prices.

In this paper, we examine how two key aspects of market design --- the degree of seller price competition and the visibility of seller ratings --- affect seller entry on a leading food waste platform. We estimate a structural model capturing both consumer demand and seller entry. The demand model incorporates Poisson arrivals with an MNL consumer choice model, incorporating spatial, temporal, and demographic factors, along with sellers' prices and ratings. For supply side, we specify a seller entry model based on a non-cooperative game. To estimate seller entry costs, we propose a discrete optimization approach that greatly improves computational efficiency over traditional simulated method of moments (SMM) and allows scaling to more covariates.

We estimate our model using a comprehensive dataset from a major secondary marketplace operating in four large U.S. cities, covering 2,700 retailers and 2.5 million hourly observations over 10 weeks in 2024. This temporal and spatial granularity allows us to observe detailed store-level operations and spatial variation in consumer demand patterns. We also incorporate Google Places data on food establishments and ACS demographics to control for neighborhood characteristics that influence demand and entry. Our results quantify how flexible pricing and rating visibility shape seller entry, offering actionable prescriptions for such platforms to boost participation and reduce food waste.

SC17

VIP B

Innovation and Strategic Decision-Making

Contributed Session

Business Analytics

Chair: JaSeung Koo, International Christian University, Mitaka, Japan

1 - Financial Analytics in Operations Management

Hariprasad Bellamkonda, Indian Institute of Management-Indore, Indore, India

In this article, we describe representative contributions in several major application areas of financial analytics in operations management to summarize the recent development, discuss the common themes, identify the current trends, and speculate the future directions. Many important contributions have been made in various application areas that are either directly or indirectly related to financial analytics, and there are important theoretical developments made by scholars in this field. The intention is not to provide a complete survey for work in this field. Instead, the focus will be on the aspect of *data integration* in operational decision-making by describing the most popular applications of financial analytics.

2 - Buffering the Shock: Business Group Affiliation and the Resilience of Exploratory Innovation in the Face of Natural Disasters

JaSeung Koo, International Christian University, Tokyo, Japan, Taiki Ito

This study examines how natural disasters influence firm search orientation in technology development. While prior research has extensively explored the environmental antecedents of exploration and exploitation, the potential impact of these environmental antecedents on decision-makers' risk preference has received limited attention. Focusing on natural disasters as an environmental factor that alters the level of managerial risk taking, we theorize how such disasters affect firm search orientation. Specifically, we argue that experiencing a natural disaster leads firm decision makers to engage in less managerial risk taking by increasing their perceived variance of outcomes, reducing their confidence in their capabilities to foresee the future, and arousing anxiety and stress.

As a result, this diminished risk preference directs search activities more towards exploitation rather than exploration, because the returns from exploration are often more uncertain and distant compared to those from exploitation. To validate our hypothesis, we utilize the Great East Japan Earthquake as a natural experiment. We analyze patent data for 985 Japanese firms in 2010 and 2013 to assess changes in exploration and exploitation activities.

The results show that firms significantly decreased exploratory search after the earthquake, supporting the argument that natural disasters lead to risk-averse behavior and a shift towards exploitation. This study contributes to the literature by highlighting the role of managerial risk-taking as a mechanism through which environmental shocks influence firm search orientation. The findings suggest that firms may benefit from strategies to mitigate the impact of natural disasters on managerial risk preferences to maintain innovation capabilities.

4 - A Comparative Study on the Competitive Advantages of Enterprises from the Perspective of Standard Essential Patents

Zidong Wang, University of Science and Technology Beijing, Beijing, China, People's Republic of, Wu Yang

In the context of a new round of global scientific and technological revolution and accelerated technological innovation, standard essential patents are playing an increasingly important role in today's era of intelligent digitalization. Based on the perspective of standard essential patents, this paper studies the formation mechanism and measurement of corporate competitive advantages. Using the "technology-economy-law" three-dimensional analysis framework, a competitive advantage measurement index system containing 15 indicators was constructed. Taking the artificial intelligence industry as an example, the Innojoy patent database was used to measure the technological competitive advantages of different countries through the entropy method and the synthetic index method. The study found that artificial intelligence standard essential patents are mainly concentrated in the fields of electrical digital data processing, digital information transmission, and image data processing. The United States has an advantage in the competition for technical standards, and China has outstanding performance in the economic and legal dimensions.

5 - Digitalization and Circular Economy: Implications for Financial Reporting

Abdullah Albizri, Montclair State University, Montclair, NJ, United States, Ahmad Jumah

In recent years, the Circular Economy (CE) has emerged as a key strategy for achieving sustainable development through cleaner production, waste reduction, and pollution mitigation. Concurrently, digital technologies play a critical role in driving sustainability by enhancing efficiency, fostering collaboration, and stimulating innovation. As businesses increasingly integrate Circular Economy principles – 10 Rs, understanding the role of digitalization and 10 Rs in financial reporting becomes essential. This research investigates the factors supporting 10 Rs and digitalization within the CE framework, with a specific focus on their implications for financial reporting, particularly in the Management Discussion and Analysis (MD&A) sections of publicly traded U.S. companies' 10-K fillings.

Our findings contribute to both academia and industry by advancing predictive analytics for materiality disclosures, improving risk assessment, and strengthening financial transparency in publicly traded companies. By leveraging AI-driven financial analysis, this research supports more informed decision-making for investors, regulators, and other corporate stakeholders. Furthermore, our framework provides organizations with a proactive approach to sustainability disclosure, ensuring alignment with evolving regulatory standards and market expectations. Ultimately, our study highlights the transformative potential of digital technologies in shaping the future of CE reporting and financial transparency.

6 - Environmental Management as Competitive Strategy: Market Reactions to Environmental Management Systems and Disclosures in Japanese Firms

JaSeung Koo, International Christian University, Tokyo, Japan, Hyuga Takeuchi

Environmental issues such as global warming and resource depletion have made corporate environmental responsibility a pivotal management concern. With the rise of ESG investment, firms increasingly adopt Environmental Management Systems (EMS) and disclose environmental information to mitigate risks and enhance economic value. This study empirically examines the impact of EMS adoption, specifically ISO 14001 certification, on corporate competitiveness in Japan's chemical industry. Using an event study methodology, it analyzes market reactions to ISO 14001 certification through cumulative abnormal returns (CAR), revealing statistically significant positive returns shortly after certification. Regression analysis further demonstrates that environmental disclosures improve corporate performance by fostering trust among investors and stakeholders.

The findings highlight that ISO 14001 certification not only signals enhanced environmental management but also strengthens brand value and competitive advantages. Environmental disclosures amplify these benefits by increasing transparency, reducing capital costs, and improving profitability. This research underscores the strategic importance of environmental management in modern business practices, offering valuable insights for corporate strategies and policy development. Future research should expand explanatory variables, diversify data sources, and include cross-industry analysis to deepen understanding of the relationship between environmental initiatives and corporate valuation.

Sunday, July 20, 4:00 PM - 5:30 PM

SD01

Enterprise

Optimization and Control in Queueing Models

Invited Session

Applied Probability Society

Chair: Ping Cao, University of Science and Technology of China, Hefei, N/A

1 - Dynamic Scheduling and Admission Control of a Multi-Class Queueing System with Customer Transfers and Abandonments

Ping Cao, University of Science and Technology of China, Hefei, China, People's Republic of, Xiaodao Wang, Zhiheng Zhong

Motivated by the fact that waiting patients may leave with being seen or his condition may deteriorate or improve in healthcare systems, we consider a service system with multiple classes of customers, in which customers may abandon the system or transition to another class while waiting, which is unfavorable as it debases the system's service level. To mitigate this service quality degradation, the system manager can intervene the customers' admission into the system when the system is too crowded, and schedule the waiting customers appropriately, with the feature of customers abandonment and transition taken into consideration. Utilizing the heavy traffic technique, we consider a related Brownian control problem, whose optimal solution is explicitly obtained. Interpreting this solution, we devise a practical and easy-to-implement admission and scheduling policy for the original system. Extensive numerical experiments demonstrate the effectiveness of our proposed policy, and highlight the importance of considering customers' transition.

2 - Patient Assignment and Prioritization for Multi-Stage Care with Reentrance

Wei Liu, University of Science and Technology of China, Hefei, China, People's Republic of, Mengshi Lu, Pengyi Shi

In this paper, we study a queueing model that incorporates patient reentrance to reflect patients' recurring requests for nurse care and their rest periods between these requests. Within this framework, we address two levels of decision-making: the priority discipline decision for each nurse and the nurse-patient assignment problem. We introduce the shortest-first and longest-first rules in the priority discipline decision problem and show the condition under which each policy excels through theoretical analysis and comprehensive simulations. For the nurse-patient assignment problem, we propose two heuristic policies. We show that the policy maximizing the immediate decrease in holding costs outperforms the alternative policy, which considers the long-term aggregate holding cost. Additionally, both proposed policies significantly surpass the benchmark policy, which does not utilize queue length information.

3 - Impact of Server Flexibility on Pooling Configuration in Stochastic Service Systems

Yanting Chen, University of Shanghai for Science and Technology, Shanghai, China, People's Republic of, Jingui Xie, Nan Yang, Zhe George Zhang, Taozeng Zhu

Resource pooling is often introduced in service systems to cope with the variability in customer demand. The primary motivation behind creating a more flexible system is to utilize resources efficiently—namely, assigning customers whose dedicated resources are fully occupied to available non-dedicated resources (referred to as off-service placement). In such a setup, the service system manager expects to serve more customers within a fixed timeframe. However, recent empirical evidence shows that, due to limited server flexibility, the service time at non-dedicated providers can be significantly longer than that at dedicated ones. In this study, we develop a two-server stochastic model to examine how server flexibility levels and other factors—such as overall workload and demand asymmetry—affect pooling configurations. We derive conditions that specify the optimal system flexibility configuration across these parameters. Our findings reveal that a partially flexible system can outperform a fully flexible one, particularly in asymmetric scenarios with low server flexibility. This advantage becomes more pronounced when the cost of server flexibility is taken into account. These insights from our two-server model offer guidance on designing more efficient flexibility in complex multi-class, multi-server service systems.

4 - On Strategic Customers in the Group-Merging Israeli Queue

Efrat Perel, Afeka Academic College of Engineering, Tel Aviv, Israel, Gabi Hanukov, Nir Perel

This paper introduces the Group-Merging Israeli Queue, an extension of the Israeli Queue model, where arriving customers, in addition to joining or creating groups, can merge two adjacent waiting customers and their respective groups into a single group for batch service. The Group-Merging Israeli Queue is formulated, and various performance measures are calculated. Furthermore, strategic customer behavior in the system is analyzed. Novel findings include the derivation of customer joining thresholds under various scenarios and the demonstration of the effect of the system parameters on customer behavior. Despite customer strategic behavior, the study reveals that the queue can remain unbounded. In addition, an economic study from the system’s perspective is presented, from which managerial decisions are derived.

SD02

Mercury

Airport Ground Operations

Invited Session

Aviation Applications

Chair: Yuan Yuan, Northwestern Polytechnical University, Xi’an, N/A

1 - A Matheuristic Approach for Optimizing Airport Baggage Transport with Drop-and-Pull Vehicles

Yuan YUAN, Northwestern Polytechnical University, Xi’an, China, People's Republic of, Xintong ZHAO, Diego CATTARUZZA, Chenhao Zhou

Airport baggage transportation is a crucial component of ground handling operations. In practice, tractors and trailers often operate in a drop-and-pull mode to shuttle baggage between the handling center and aircraft stands. The efficiency of their scheduling impacts baggage processing speed, ground handling costs and flight punctuality. However, the flexible drop-and-pull operations and the collaborative scheduling between trailers and tractors have been largely overlooked in existing research. Moreover, limited attention has been given to stand-side parking capacity, which is crucial for ensuring operational safety. This study addresses an integrated routing and scheduling problem for tractors and trailers operating under the drop-and-pull mode, considering the parking capacity at aircraft stands. Tractors are allowed to leave an aircraft stand while their trailers are being loaded or unloaded. A mixed-integer programming model is first developed, incorporating constraints such as time windows, baggage release/waiting times, multiple trips per vehicle, and simultaneous pickup and delivery. To solve this problem, an efficient two-stage matheuristic approach is proposed. The first stage employs an adaptive large neighborhood search algorithm, while the second stage applies a branch-and-check framework. Finally, instances based on real-world airport and flight data are generated. Experimental results demonstrate that the proposed mathematical model provides optimal solutions for small-scale instances using the CPLEX solver, and the proposed algorithm delivers high-quality solutions for medium and large-scale instances within reasonable time.

2 - Online Retrieval and Storage Algorithm for Airport Robotized Baggage Handling System

Chenhao Zhou, Northwestern Polytechnical University, Xi'an, China, People's Republic of

Amidst the global increase in air travel and growing concern regarding labor shortages, baggage handling systems confront substantial challenges as a crucial component of airport operations. Recently, robotic arms have been positioned alongside carousels to retrieve and store baggages directly onto baggage carts. To avoid potential demages, stacking protocol, light and small on top, is enforced. However, due to the randomness in passengers' arrivals and the weights and sizes of their pieces of baggage, challenges arise in managing robotic arms, which require the retrieval and storage decisions to be made in an online mode. To address this issue, this study proposes an online algorithm and has the algorithm tested in a high-fidelity simulatio model. Subsequently, extensive simulation experiments are conducted using 120 sets of real passenger check-in data across various departure aircraft.

3 - Optimizing System's Level Aircraft Turnaround Operations

Marcus Koh, Singapore University of Technology and Design, Singapore, Singapore

The aircraft turnaround process comprises a series of operations carried out between an aircraft’s arrival at the gate and its subsequent departure.

This research aims to develop a more comprehensive model of these operations to optimise resource planning across the airport. Given the involvement of multiple service vehicles performing diverse tasks, the Vehicle Routing Problem (VRP) provides a natural framework for modelling the process. The proposed model advances existing VRP approaches by incorporating key elements such as task precedence, vehicle reloading, and multi-service capabilities—factors often overlooked in previous studies.

4 - Evaluating Charging Infrastructure Adequacy for Electrified Airside Vehicles: a Simulation-Based Approach

Muhamad Azfar Ramli, Institute of High Performance Computing, A*STAR, Singapore, Singapore, Wenqiang Wang, Rishikeshan Rajendram, Mohasha Sampath, Joseph Chan, Hongwei Wang

In order to fulfill commitments to transport sector decarbonisation, one of the key strategies for airports is to electrify their existing airside vehicle fleet. However, this transition presents various challenges in infrastructure planning and operations management. Airport operators face significant uncertainty in understanding the adequacy of chargers required to support electrification of airside operations vehicles and the additional grid capacity needed.

In this presentation, I will share our approach to developing a data-driven modelling and simulation workflow to address the electrification planning problem for airside operations. The workflow includes:

- Modeling the broad energy requirements of the electrified airside fleet vehicles.
- Determining the requirements for a mixed fleet of internal combustion engine (ICE) and electrified vehicles (EVs) to support a specified airport's operational services.
- Running a dynamic movement and charging simulation of the various electrified vehicles and chargers
- Analysis of the detailed simulation output including spatiotemporal patterns of charging demand, charger utilisation and variations of battery state-of-charge for the various simulated vehicles

The model output allows planners to determine the patterns of electrical charging energy requirements and the required number of chargers distributed across the area of operations.

We demonstrate our methodology using a case study of Singapore's Changi Airport airside operations and showcase selected model output results based on historical flight demand

SD03

Minto

Service Management

Invited Session

Applied Probability Society

Chair: Junfei Huang, The Chinese University of Hong Kong, Shatin, Hong Kong

1 - Near-Optimal Pricing and Resource Allocation in a Large-Scale Service System

Xu Sun, Miami Herbert Business School, Miami, FL, United States, Zerui Wu, Ran LIU

We study dynamic pricing and capacity allocation in large-scale service systems where multiple service units serve customers who are both price- and delay-sensitive. Customers are segmented into classes, each characterized by class-specific service rates and demand functions shaped by posted prices and perceived delays. To jointly optimize revenue and delay performance, we propose a family of state-dependent greedy heuristics. These heuristics (i) assign dedicated service capacities to each customer class and (ii) dynamically set prices by solving a tractable one-step optimization problem at each decision epoch. Despite their simplicity, the proposed policies achieve provably strong performance: in large-market regimes where demand scales with the number of servers n, they attain a relative optimality gap of O(n^(-2/3) log n). Moreover, we show that this performance guarantee extends to any work-conserving allocation, not just the dedicated-capacity rule. This robustness highlights the practicality of our approach in settings where centralized control over service units is limited. Extensive numerical experiments validate the theoretical results and offer further managerial insights. In particular, we observe that queue-based wait-time estimates can lead to non-monotonic congestion pricing behavior—i.e., optimal prices may decrease with increasing congestion—a counterintuitive yet rational response under certain demand sensitivities.

2 - Fleet Sizing and Rebalancing in Free-Floating Bike-Sharing Systems

Hang Guo, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Jun Luo, Jiheng Zhang

The free-floating (i.e., dockless) bike-sharing system, widely adopted by operators such as Hellobike and Meituan Bike, offers greater flexibility to users. This paper addresses the joint decisions of fleet sizing and rebalancing in free-floating bike-sharing systems. We present a novel modeling approach that formulates these systems as a load-dependent closed queuing network. The model captures spatial features through the load-dependent form, reflecting the relationship between demand fluctuations and bike distributions. Given the extreme complexity of the system, we apply fluid approximation and formulate the fluid optimization problem to determine the optimal fleet size and relocation policies. We propose periodic rebalancing policies that are both practical for real-world applications and asymptotically optimal under a certain regime. These policies balance the operational need to maintain user satisfaction with practical constraints, providing a robust framework for managing the dynamic nature of free-floating bike-sharing systems. The simulation results highlight the superior performance of the proposed policies.

3 - Qed Appointments

Junfei Huang, The Chinese University of Hong Kong, Shatin, Hong Kong, Petar Momcilovic, Avishai Mandelbaum

There is a vast world of service-systems that are driven by appointments (reservations). Many of these systems are large, time-varying and stochastic, which motivates the present study. Our model captures “newsvendor tradeoffs” between underage (too few appointments, hence excessive idleness) vs. overage (too many appointments, hence excessive congestion). In large systems, this tradeoff can be sometimes ameliorated, via careful appointment-design that exploits economies-of- scale; and borrowing the terminology of many-server queueing theory, the operation is then Quality- and Efficiency-Driven (QED), that adheres to the recipe of square-root safety-appointments. More concretely, we optimize appointments to a service system with soft capacity (e.g. an amusement or a nature reserve), with penalties for over- and under- capacity utilization. We offer an asymptotically-optimal two-stage solution, first of a first-order fluid model and then its stochastic diffusion refinement.

4 - Learning and Optimization of Patient&Ndash;Physician Matching Index in Specialty Care

Debiao Li, Fuzhou University, Fuzhou, China, People's Republic of

It is challenging for a patient without medical knowledge to select a ‘capable’ physician based on nontransparent medical information. The situation becomes pervasive in specialty care. Motivated by this existing patient-physician matching problem, we propose a novel physician matching index (PMI) obtained by an analytical framework integrated with an improved multi-disease pre-diagnosing Bayesian network (BN) model. The pre-diagnosis BN structure learning is critical, as it provides a causal map among diseases and symptoms; however, it has been proven to be NP-hard. To improve the computational tractability of the BN structure learning, we propose a dynamic programming based cache calculation algorithm integrated with expert knowledge. The optimal BN structure is obtained through an improved branch-and-bound algorithm. Given patients’ symptoms and physicians’ specialty information, we apply the trained pre-diagnosis BN model to obtain PMI, which can be extended to the weighted matching index by considering patient preferences. A case study of the patient-physician matching problem in the ear, nose, and throat (ENT) department is conducted. The branch-and-bound algorithm with the proposed cache calculation algorithm learns the optimal BN structure with high pre-diagnosing accuracy and time efficiency. We disclose that the proposed PMI can rectify the misdiagnosis since the highly related diseases usually belong to one specialty. Moreover, we demonstrate the significance of the consistency between the physicians’ specialty and the patients’ disease distribution. We also highlight that the proposed PMI guides the patients in choosing physicians more appropriately under independent patient preferences.

SD04

Indiana

How Information, Incentives, and AI Shape Decisions in Health and Digital Platforms: Evidence from Field & Experiments

Invited Session

Behavioral Operations Management

Chair: Vivek kumar Choudhary, Nanyang Technological University, SIngapore, Singapore

1 - Screening Process Information Matters More than Clinical Information: a Randomized Controlled Trial of Information Sharing on Cancer Screening Decision

Xiaodong Wang, Singapore Management University, Singapore, Singapore, Yongjian Zhu, Sarah Yini Gao, Zhichao Zheng, Shao Wei Lam

This study examines which types of information effectively promote cancer screening uptake using data from a randomized controlled trial. The results underscore the importance of screening process information and suggest that messages should be tailored to specific subpopulations.

2 - A Field Experiment Study on Artificial Intelligence-Assisted Extubation Decision

Zhichao Hu, Nanyang Technological University, Singapore, Singapore, Singapore

Extubation decision-making in intensive care units (ICUs) remains a high-risk challenge, with traditional predictors often failing to account for dynamic physiological changes. This study introduces the Extubation Risk Prediction (ERP) algorithm, a deep learning framework integrating Long Short-Term Memory (LSTM), Convolutional Neural Networks (CNNs), and Temporal Convolutional Networks (TCNs) to predict extubation readiness and pneumonia risk. By leveraging multimodal data—including high-frequency ventilator waveforms and clinical records—the ERP algorithm provides real-time predictions and hourly risk forecasts. Clinical validation demonstrated its superior performance, achieving AUC-ROC values of 0.946 for extubation success and 0.992 for pneumonia risk prediction, significantly outperforming conventional metrics like the Rapid Shallow Breathing Index (RSBI, AUC=0.68). To assess its impact on clinical decision-making, we conducted a randomized field experiment with 60 ICU physicians. Physicians assisted by the ERP system improved their decision accuracy by 46%, reduced overly conservative choices, and maintained decision efficiency. Moreover, the system's explainable AI (XAI) elements influenced risk-taking behavior, encouraging greater confidence in easier cases while promoting caution in more complex ones. These findings underscore the potential of AI-driven decision support to enhance extubation outcomes, reduce ICU costs, and optimize clinical workflows.

3 - Ka-Ching: the Impact of Cash on Delivery on Gig Drivers&Rsquo; Behavior

Li Ding, Nanyang Technological University, Singapore, Singapore, Saharsh Agarwal, Vivek Choudhary

The rise of food delivery platforms has transformed the gig economy, with Cash on Delivery (COD) remaining a dominant payment method. With COD, drivers temporarily retain the entire order payment, effectively receiving a temporary interest-free loan. While prior gig economy studies have focused on explicit incentives such as bonuses and wages, our study investigates the implicit liquidity benefit that COD provides, an important aspect that has received little attention in previous literature. COD provides drivers with immediate liquidity to cover expenses such as fuel and vehicle maintenance but also introduces risks of theft and cash-handling complexities. We examine how COD influences drivers' decisions, particularly whether receiving cash impacts their intra-day shift ending decisions, the likelihood of returning to work the next day, and long-term attrition. Using data from a major Asian food delivery platform covering 8,371 drivers and about 4.1 million orders over six months, we exploit the exogenous lottery-like assignment of COD orders to analyze its effects. Our findings reveal that COD significantly shapes workforce dynamics. An additional US$10 in COD makes drivers 15.3% more likely to stop working in the next hour compared to non-COD payments, reflecting lower cash based income targets. COD also increases drivers' probability of returning to work the next day by 1.2% points, stabilizing weekly labor supply while reducing attrition and quitting rates, resulting in significant economic savings for the firm. These insights challenge platforms' aversion to COD, showing it can serve as a non-monetary tool for workforce management by enhancing driver retention.

4 - From Information Aggregators to Transaction Facilitators: An Empirical Investigation of Platform Business Model Transformation

Zhenming Shi, Nanyang Technological University, Singapore, Singapore, Vivek Choudhary, Serguei Netessine

Over the last decade, digital platforms have grown significantly, operating under two distinct models from the perspective of where the final transactions occur. Some platforms function as information aggregators, solely matching buyers with sellers, while transactions occur outside the platform. Others serve as transaction facilitators, integrating payment systems that enable users to complete transactions directly within the platform. We examine this evolution from information aggregator to transaction facilitator in the novel context of trucking platforms. Using a quasi-experimental study of a major Chinese freight logistics platform that introduced transaction capabilities, we analyze over 2.6 million transportation records through a Difference-in-Differences (DID) approach to understand how this transition affects platform stakeholders. Our results show that integrating a transaction system improved operational efficiency by reducing order acceptance times and delivery costs. Weekly driver earnings increased by ~15%, primarily because the platform's transaction system builds trust between drivers and unfamiliar shippers, leading to higher order acceptance rates and better resource utilization. However, earnings per kilometer per ton decreased by ~3.4%, suggesting drivers needed to handle larger workloads to maintain their income levels.

5 - Enhancing the Wisdom of Crowds when Experts use AI Advice

Angshuman Pal, Indiana University Kelley School of Business, Bloomington, IN, United States, Asa Palley, Ville Satopaa

In many managerial settings, good decision making depends on obtaining an accurate forecast of an uncertain variable of interest. Useful knowledge about the variable may be held by both human experts and algorithmic models, where each may have their own relative advantages in forming an accurate forecast. To make use of both sources of information, AI advice can be provided to human experts who update their prior belief based on the advice received. Since useful information may be held by different individuals, combining forecasts from multiple experts can also boost the accuracy of the aggregate forecast. However, while providing AI advice to humans can be helpful at an individual level, it can lead to two forms of bias in the aggregate forecast. First, it may induce correlations in judgment errors arising from a shared information bias. Second, humans with behavioral predispositions toward algorithms may utilize the advice improperly.

Using a stylized Bayesian model of information held by human experts and the AI, we propose a pivoting procedure using which the decision maker may recalibrate a potentially biased aggregate forecast. We develop a prescriptive method to calculate a robust pivoting factor which minimizes the maximum regret for a given range of beliefs about the information environment. Interestingly, we observe that the robust pivoting factor is close to zero when the experts are expected to be algorithm-averse; this indicates that the shared information bias and the behavioral bias cancel each other out, therefore mitigating the need for further adjustment via pivoting.

SD05

Orchard A

Optimization in Modern Service Systems

Invited Session

Decision Analysis Society

Chair: Rui Chen, The Chinese University of Hong Kong, Shenzhen

1 - The Impact of Automation on Workers when Workers are Strategic: The Case of Ride-Hailing

Zicheng Wang, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Saif Benjaafar, Xiaotang Yang

Motivated by the behavior of drivers on ride-hailing platforms (individual drivers decide whether or not to work based on the offered wage and where to locate themselves in anticipation of future fares), we examine how the introduction of autonomous vehicles impacts the strategic behavior of human drivers and driver welfare. Specifically, we consider a setting where a ride-hailing platform deploys a mixed fleet of conventional vehicles (CVs) and autonomous vehicles (AVs). The CVs are operated by human drivers who make independent decisions about whether to work for the platform and where to position themselves when they become idle. The AVs are under the control of the platform. The platform decides on the wage it pays the drivers, the size of the AV fleet, and how the AVs are positioned spatially when they are idle. The platform can also make decisions on

whether to prioritize the AVs or the CVs in assigning vehicles to customer requests. We use a fluid model to characterize the optimal decisions of the platform and contrast those with the optimal decisions in the absence of AVs. We show that, although the introduction of AVs can displace drivers and depress effective wages, there are settings where the introduction of AVs leads to higher effective wages and more drivers being hired.

2 - Embedded Queues: Offering An Ancillary Service in a Queue

Baolong Liu, ShanghaiTech University, Shanghai, China, People's Republic of, Yang Li, Rowan Wang

Restaurants and theme parks often implement queueing systems where primary service providers offer an ancillary service. In these systems, customers can be classified into two types: Type-P customers who request only the primary service, and Type-A customers who desire both services with an additional reward. All customers must wait in the primary service queue, but Type-A customers can better utilize their waiting time to enter a separate queue for the ancillary service. Although the ancillary service enhances customer experience and may increase demand, it may further congest the primary service — thus creating a trade-off between demand expansion and queueing congestion. We treat such a system as an unobservable queue to analyze customer equilibrium strategies, system throughput, provider profit, and social welfare. We evaluate the performance of a system with ancillary service (AS) against two benchmarks: a first-in-first-out (FIFO) queue and a priority (PRI) queue for the primary service. An ancillary service may outperform both FIFO and PRI systems across all metrics, especially when primary service capacity is limited. Moreover, a value-creating AS system may generate higher profits and social welfare than a PRI system which is traditionally considered an effective profit-enhancing strategy by prioritizing customers with high unit waiting costs. In some cases, introducing an ancillary service proves more effective in improving system performance than expanding primary service capacity.

3 - The Impact of Iot on Cold Chain Logistics

Manlu Ouyang, Stern School of Business, New York University, New York, NY, United States, Rui Chen, Duo Shi

The rise of the Internet of Things (IoT) provides a new operational mode for cold chain logistics, making it possible to monitor whether a product perishes during transportation from origin to destination. This paper investigates how IoT adoption impacts a firm’s key cold chain decisions regarding the reefer system: investment in refrigeration and preparation of product quantity. We build a stylized model in which a firm first invests in refrigeration—where higher investment improves system reliability and reduces the likelihood of perishment—then makes a quantity decision. Without IoT, the firm makes a one-shot quantity decision before transportation; with IoT, it can detect product status en route and, if any units perish, follow up with emergency replenishment. We characterize two basic strategy patterns: aggressive (conservative) quantity preparation paired with inactive (active) refrigeration investment. We find that IoT may serve as either a complement or substitute for refrigeration. If refrigeration is inefficient, IoT boosts its use by shifting the firm from inactive investment with aggressive quantity to active investment with conservative quantity. If refrigeration is efficient, the firm invests actively and prepares conservatively regardless of IoT, but with IoT, relies less on refrigeration as IoT serves as an added tool to mitigate perishment risk. Thus, the incremental benefit of IoT is non-monotonic in refrigeration investment efficiency. These findings reveal a nontrivial interplay between IoT and refrigeration technologies, offering guidance on whether and how IoT should be adopted to enhance cold chain logistics.

4 - Learning While Repositioning in on-Demand Vehicle Sharing Networks

Hansheng Jiang, University of Toronto, Rotman School of Management, Toronto, ON, Canada, Chunlin Sun, Zuo-Jun Shen, Shunan Jiang

We consider a network inventory problem motivated by one-way, on-demand vehicle sharing services. Due to uncertainties in both demand and returns, as well as a fixed number of rental units across an n-location network, the service provider must periodically reposition vehicles to match supply with demand spatially while minimizing costs. The optimal repositioning policy under a general n-location network is intractable without knowing the optimal value function. We introduce the best base-stock repositioning policy as a generalization of the classical inventory control policy to n dimensions, and establish its asymptotic optimality in two distinct limiting regimes under general network structures. We present reformulations to efficiently compute this best base-stock policy in an offline setting with pre-collected data.

In the online setting, we show that a natural Lipschitz-bandit approach achieves a regret guarantee of O(T^{n/(n+1)}), which suffers from the exponential dependence on n. We illustrate the challenges of learning with censored data in networked systems through a regret lower bound analysis and by demonstrating the suboptimality of alternative algorithmic approaches. Motivated by these challenges, we propose an Online Gradient Repositioning algorithm that relies solely on censored demand. Under a mild cost-structure assumption, we prove that it attains an optimal regret of O(n^{2.5} T^{0.5}), which matches the regret lower bound in T and achieves only polynomial dependence on n. The key algorithmic innovation involves proposing surrogate costs to disentangle intertemporal dependencies and leveraging dual solutions to find the gradient of policy change. Numerical experiments demonstrate the effectiveness of our proposed methods.

5 - Fleet Repositioning for Vehicle Sharing Systems: Asymptotic Optimality of the Balanced Myopic Policy

Yimin Yu, City University of Hong Kong, Kowloon, Hong Kong, Qian WANG, Yihang Yang, Junming Liu

We investigate the fleet repositioning problem aimed at dynamically optimizing vehicle distributions to maximize long-run average social welfare in a vehicle-sharing system. We model the problem as a Markov decision process under the ex ante committed decision scheme, characterizing the balanced myopic policy as optimal for the average reward setting. This policy efficiently aligns vehicle supply with trip demand and mitigates the curse of dimensionality, enhancing computational efficiency significantly. Our analysis demonstrates that although the balanced myopic policy operates with less information, potentially leading to performance losses, the maximum performance gap relative to the ex post decision scheme asymptotically converges to zero as the system size increases. This finding underscores the asymptotic optimality of the balanced myopic policy, particularly in large systems, making it a robust and effective solution for fleet repositioning. Moreover, we extend our investigation to settings with seasonal demand, confirming that a generalized balanced myopic policy remains optimal. Through comprehensive numerical experiments and a counterfactual case study of a real-world vehicle-sharing system, we quantify the operational value of our approach. This study not only validates the balanced myopic policy against more information-intensive solutions but also illuminates effective heuristic design strategies for improving the efficiency of fleet repositioning in vehicle sharing systems.

SD06

Orchard B

Innovative Stochastic Methods in Finance and Insurance

Invited Session

Finance

Chair: Zuoquan Xu, The Hong Kong Polytechnic University, Kowloon, N/A

1 - Mean-Variance Portfolio Selection with Stochastic Dominance Constraints

Jiaqin Wei, East China Normal University, Shanghai, China, People's Republic of

This paper is concerned with the mean-variance (MV) portfolio selection problem with a first- or second-order stochastic dominance (FSD or SSD) constraint in a complete market.
First, the FSD-constrained efficient payoffs are explicitly derived and the SSD-constrained efficient payoffs are fully characterized in terms of the solution of a dual problem. Second, it turns out that, if the benchmark is appropriately chosen, the SSD constraint is binding in the bad market states and hence the left-tail risk can be effectively controlled. Third, a numerical solution is subtly constructed to guarantee the convergence of a combination of an ε-algorithm and a projected gradient method (Rosen's algorithm).

2 - Optimal Dividend Payout Under Path-Dependent Constraint

Zuo Quan Xu, The Hong Kong Polytechnic University, Hong Kong, China, People's Republic of

This paper is concerned with a long standing optimal dividend payout problem in insurance subject to the so-called ratcheting constraint, that is, the dividend payout rate shall be non-decreasing over time. The surplus process is modelled by a drifted Brownian motion process and the aim is to find the optimal dividend ratcheting strategy to maximise the expectation of the total discounted dividend payouts until the ruin time. We propose a

novel partial differential equation method to solve the problem and derive an optimal dividend ratcheting strategy, and thus solve the open problem completely. Economically, we find that if the surplus volatility is above an explicit threshold, then one should pay dividends at the maximum rate, regardless the surplus level. Otherwise, by contrast, the optimal dividend ratcheting strategy relays on the surplus level and one should only ratchet up the dividend payout rate when the surplus level touches the dividend ratcheting free boundary.

3 - Linear&Ndash;Quadratic Mean-Field Game for Stochastic Systems with Partial Observation

Na Li, Shandong University of Finance and Economics, Jinan, China, People's Republic of, Min Li, Zhen Wu

This work is concerned with a class of linear–quadratic stochastic large-population problems with partial information, where the individual agent only has access to a noisy observation process related to the state. The dynamics of each agent follows a linear stochastic differential equation driven by the individual noise, and all agents are coupled together via the control average term. By studying the associated mean-field game and using the backward separation principle with a state decomposition technique, the decentralized optimal control can be obtained in the open-loop form through a forward– backward stochastic differential equation with the conditional expectation. The optimal filtering equation is also provided. Thanks to the decoupling method, the decentralized optimal control can also be further presented as the feedback of state filtering via the Riccati equation. The explicit solution of the control average limit is given, and the consistency condition system is discussed. Moreover, the related ε-Nash equilibrium property is verified. To illustrate the good performance of theoretical results, an example in finance is studied.

SD07

Orchard C

Financial Asset Allocation

Invited Session

FinTech

Chair: Ling Zhang, Guangdong University of Finance, Guangzhou, N/A, China, People's Republic of

Co-Chair: Ling Zhang, Guangdong University of Finance, Guangzhou, N/A, China, People's Republic of

1 - Early Warnings of Non-Life Insurance Companies: An Application of Explainable Distillation Learning Based on Ensemble Trees

Zhiyong Li, Southwestern University of Finance and Economics, Chengdu, China, People's Republic of, Tianhao Yi, Chen Feng

As a key type of key players, insurance companies play an important role in the stability of the financial system. However, they received limited attention regarding their own soundness compared to banks. In this research we focus on the failure of insurance companies and introduce semi supervised learning to predict the failure of insurance companies. We extract explainable decision rules from ensemble trees through distillation learning and automatically screen valuable rule subsets. The white box model is used to train the rule subsets and obtain statistical indicators such as significance level. In this paper, we compare the proposed explainable method with four white box classification models and nine ensemble tree models on a 10-class rating prediction problem using financial ratios and AMBest rating data of US non-life insurance companies from 2009 to 2021. The empirical results demonstrate that the proposed explainable method can maintain the prediction power of the distillation object while providing explainablity. This paper further applies the proposed method to a 5-class prediction problem of rating changes, and utilizes semi-supervised learning methods to further improve the prediction power. Finally, we analyze the results of the white box model based on rule subsets to identify the risk factors for insurance company rating levels and rating transitions.

2 - Information Acquisition, Portfolio Choice, and Diversification

Jiming PAN, City University of Hong Kong, Hong Kong, Hong Kong, Xiao QIAO, Moris Strub

We study the joint decision of information acquisition and portfolio choice for a mean-variance investor. The investor is concerned with risk from both the information acquisition and investment periods. Optimal policies indicate that the investor should learn about assets with lower prior squared Sharpe ratios in order to improve the prospects of a diversified portfolio, because a concentrated portfolio is associated with an increased risk of extreme returns. Our results show that diversified portfolios can arise in a model without perfect information, suggesting that diversification has wider applicability than previously recognized.

3 - Using Conditional Information to Improve Performance of Mean-Variance Portfolios

Qiong Ji, Xi'an Jiaotong-Liverpool University, Suzhou, China, People's Republic of, Xing Jin, Ming Cheng

There is large empirical evidence that stock returns are generally predictable. However, stock return predictability is largely ignored for the expected stock return forecasting in the Markowitz mean-variance portfolio selection. This article explores the return predictability within three widely-used risk targeting, return targeting, and risk-return tradeoff mean-variance frameworks with or without regularization. Various empirical tests illustrate that the return predictability with Fama-French factors as predictors drastically improves performance of risk targeting mean-variance portfolio, especially outperforming the naive 1/N strategy overwhelmingly. Surprisingly, we both theoretically and empirically demonstrate striking differences among the above three conditionally equivalent mean-variance models in exploring the return predictability. Specifically, we find that targeting a constant return limits the exploitation of return predictability and the success of exploring return predictability critically hinges on how a mean-variance investor controls portfolio risk.

SD08

Bras Basah A

Supply Chain and Product Design Optimization

Contributed Session

Manufacturing and Service Operations Management

Chair: Shiyu Wang, Nanning University, Nanning, N/A

1 - Utilization of Aggregation: a Temporal-Frequency Deep Neural Network with Mixture-of-Experts for Retail Sales Forecasting

Shihang Sun, School of Economics and Management, Dalian University of Technology, Dalian, China, People's Republic of, Xianneng Li

Although sales forecasting can be enhanced by incorporating exogenous variables, many upstream supply chain enterprises often lack access to such data. Particularly in offline channels, brands and manufacturers can only obtain time-series sales data from individual merchants as their forecasting input. Utilizing the advantages of data aggregation becomes crucial in such contexts. However, this task is not trivial. Although numerous studies have been conducted, the trade-off between overall forecasting stability and individual merchant heterogeneity has rarely been investigated. Many investigations have indicated that the sales trends of different merchants are usually distinct, but also exhibit some similarities in patterns. How to discover aggregation similarity under individual heterogeneity becomes essential in retail sales forecasting. To address the above issue, we propose a novel end-to-end deep neural network framework, which introduces the Mixture-of-Expert (MoE) mechanism and multi-task learning (MTL) as a systematic whole to learn the high-level shared information under sales aggregation. The MoE module learns three experts, named temporal,global frequency, and local frequency modelling, to capture the overall shared information from aggregated sales . The MTL module adaptively clusters and learns high-level similarity in sales among individual merchants, ultimately accurate aggregated sales forecasting for each merchant. Experimental studies are carried out under two real-world retail datasets, while the results demonstrate the superiority of our proposal over the most recent state-of-the-art baselines. The ablation experiments further validate the effectiveness of each algorithmic component. Further explanatory analysis shows that the sales similarity we found has practical management implications for marketing channel categorization.

2 - Order Release and Lead TIME Planning Through Price Coordination: a Lagrangian Approach

Gregor Blossey, University of Innsbruck, Innsbruck, Austria, Hubert Missbauer, Reha Uzsoy

Optimization models that determine order releases for a multi-period planning horizon simultaneously perform production smoothing over time and cycle time control by controlling the level of work-in-progress inventory. The crucial modelling challenge is that of anticipating the effects of a specified order release schedule on the time-dependent cycle time distribution. Using analytical or empirical approximations for this purpose, such as clearing functions, works reasonably well, but it seems extremely difficult to improve the models beyond the current state of the art. Several authors have sought to overcome these difficulties by decomposing the order release problem into two subproblems: a production smoothing subproblem, typically represented as a linear program, which determines the order releases for given time-dependent lead times, and a cycle time estimation subproblem for given releases, usually addressed through discrete-event simulation. The subproblems are solved iteratively until convergence. However, this procedure can fail to converge, and it has been shown recently that the procedure is a flawed application of mathematical decomposition. We propose a Lagrangian decomposition approach that decomposes the order release problem into order release and production scheduling subproblems. Coordination is achieved through Lagrange multipliers that act as prices for capacities and penalize long lead times. Since linear prices cannot coordinate the subproblems we use a variant of augmented Lagrangian relaxation that ensures separability. Preliminary computational results are presented and open questions are discussed. The paper aims to contribute both to the literature on order release planning models and to the integration of mathematical programming and discrete-event simulation.

3 - Tackling Decision Dependency in Contextual Stochastic Optimization

Xiangting Liu, Tsinghua University, 北京市, China, People's Republic of, Wenxuan Liu, Maoqi Liu, Zhi-Hai Zhang, Hanzhang Qin

This paper addresses contextual stochastic optimization (CSO) problems with decision-dependent uncertainty, where the distribution of random parameters depends on both observed covariates and the decision itself. Existing approaches often impose restrictive structural assumptions (e.g., linear demand-price relationships) or face computational bottlenecks in nonparametric methods like weighted sample average approximation (wSAA). While wSAA improves classic SAA by weighting samples via covariates and decisions, its integration with machine learning-based weight estimation becomes intractable due to intertwined dependencies between weights and decision variables. To resolve this, we propose a **contextual gradient descent (CGD)** method that applies the wSAA principle directly to the gradient of the loss function. This yields a "contextual gradient" that integrates covariate and decision-dependent effects into a gradient descent framework. Theoretical analysis shows that under strong convexity, CGD achieves a bounded suboptimality gap relative to the global optimum, where the loss function's convexity mitigates uncertainty from decision-dependent effects. In general settings, CGD converges to a stationary point of the expected gradient. Numerical experiments on real-world datasets demonstrate CGD's consistent superiority over existing contextual optimization methods with decision-dependent uncertainty. Key innovations include (1) circumventing structural assumptions via nonparametric gradient weighting and (2) enabling efficient computation by decoupling weight estimation from optimization. The results highlight the critical role of loss function convexity in stabilizing decision-dependent CSO solutions.

4 - Bo-LSTM Driven Warehouse Optimization: from Demand Forecasting to Scheduling

Shiyu Wang, Nanning University, Nanning, China, People's Republic of, Jun Qin, Tailin Li, Bote Liu, Hua Zhang, Yiyong Han

With the rapid development of the warehousing and logistics industry, the traditional warehousing management faces problems such as inefficiency and resource waste. In this paper, based on the constructed BO-LSTM combined prediction model, an optimization study is carried out in warehousing scenarios. A cross-domain integration framework fusing temporal prediction, spatial optimization, and operations scheduling is proposed. Experiments show that the use of Bayesian optimization ensures the stability of the algorithm, the modular design adapts to the heterogeneous warehousing scenarios, and the BiLSTM model outperforms the traditional LSTM in logistics demand prediction, with more accurate prediction. The integrated framework can significantly reduce warehousing costs, improve operational efficiency, and enhance the adaptability and scalability of the warehousing system, providing new ideas and methods for the intelligent development of the warehousing and logistics industry.

SD09

Bras Basah B

Pricing Strategies and Revenue Management

Contributed Session

Manufacturing and Service Operations Management

Chair: Hyun Seok Lee, Korea University Business School, Seoul, Korea, Republic of

1 - Multiarmed Bandit Algorithms for Dynamic Pricing of Perishable Products

Jingcheng Lei, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Shaoxuan Liu, Ruijing Wu

When making pricing decisions, retailers face not only the challenge of incomplete customer valuation information but also the need to account for the direct influence of product popularity on customer valuations. In the absence of prior knowledge, they must gradually identify the profit-maximizing price through exploratory price experiments. Given the time-sensitive nature of the information, they also need to actively detect changes in the environment and discard outdated information in a timely manner. In this paper, we propose two multiarmed bandit (MAB) algorithms, incorporating elements of microeconomic choice theory, to address dynamic pricing for retailers. These algorithms first utilize historical binary purchase feedback for online learning, and then improve the limitations of partial feedback in MAB algorithms by leveraging the monotonicity of customer purchase probabilities with respect to price and time. In addition, they actively detect environmental changes rather than passively adapt to them. In our theoretical analysis, we establish upper bounds for both algorithms in a piecewise stationary environment. Building on this, we conduct a comprehensive numerical analysis across different customer valuation distributions. Our results demonstrate the superior performance of the proposed algorithm compared to traditional passive and active detection MAB algorithms presented in the existing literature.

2 - Two-Way Substitution in Revenue Management with Multiple Flexible Products

R K Amit, IIT Madras, Chennai, India, Dhandabani Srinivasan

We address the challenge airlines face in managing overbooking and capacity control for nonstop origin-destination pairs, considering a new flexible product that facilitates the reallocation of ticket holders to free up capacity in business class. In the airline industry, revenue loss can occur due to low-value passengers in business class who book well in advance, as those passengers cannot be reallocated to accommodate higher-value requests. We propose selling BusinessLite products, offering flexibility in business class allocations while protecting the airline’s brand reputation. We utilize dynamic programming to model sequential arrivals over the selling season and linear programming to determine allocation at the departure date. We characterize the acceptance policy based on opportunity cost using multimodularity concepts. Theoretical findings indicate improved revenues and load factors for full-service carriers. For low-cost airlines, the decision to offer BusinessLite products depends on the level of cannibalization, which is in turn, influenced by demand. Empirical data from two major airlines support these findings, particularly benefiting low-cost carriers. Interviews with industry practitioners offer practical insights and guidance for airlines considering the implementation of BusinessLite products, suggesting potential revenue gains without additional costs.

3 - Mitigating Horizontal Double Marginalization in Complementary Products: Pay-Per-use and Selling Models

Zanlu Hua, School of Management, Lanzhou University, Lanzhou, Gansu, China, People's Republic of, Kangzhou Wang, Donglong Liao

Independent pricing by complementary products manufacturers leads to higher total price, lower customer demand and lower total profit, which is known as horizontal double marginalization. In practice, the pay-per-use (PPU) model, in which manufacturers charge customers based on usage and incur extra operating costs, is emerging. As some manufacturers of complementary products adopt the PPU models while others retain selling models, these manufacturers strategically adjust their pricing decisions, thereby affecting the horizontal double marginalization. Therefore, this paper investigates how the adoption of the PPU model by one manufacturer affects pricing decisions, profits, and horizontal double marginalization in a market with two strictly complementary products. The results show that, compared to the case where both manufacturers adopt selling models, the existence of the PPU model increases profits for both the PPU manufacturer and the selling

manufacturer when the operating cost is low. Moreover, the PPU model mitigates the horizontal double marginalization when the usage heterogeneity is high, or when the usage heterogeneity is medium and the operating cost is low.

4 - Try Before You Buy: How Fitting Room Usage Influences Apparel Sales and Spillover Effects

Hyun Seok Lee, Korea University Business School, Seoul, Korea, Republic of, Saeyoung Yoon

With the rise of online shopping, brick-and-mortar (B&M) retailers face increasing challenges in maintaining market share. Despite this trend, physical stores retain a key advantage - the ability for customers to experience products firsthand. In the apparel industry, fitting rooms play a crucial role in purchase decisions, yet the relationship between fitting room usage and sales remains underexplored. This study utilizes RFID data from a major Asian fast-fashion retailer to investigate how fitting room interactions impact clothing purchases, specifically focusing on spillover effects related to sizes and colors. Using a dataset comprising 152,077 observations across 28,716 SKUs, we perform regression analysis to quantify the link between fitting room trials and actual sales. Our findings confirm that increased fitting room usage positively correlates with higher sales of the product tried on (*direct effect*). Additionally, we identify a *spillover effect*, revealing that trying on alternate sizes or colors of the focal product increases the likelihood of purchasing that product, with size spillovers having a stronger influence than color spillovers. Store traffic and stock availability also significantly affect purchase behavior. This research contributes to the literature in two primary ways. First, it provides empirical evidence linking fitting room usage to SKU-level sales, offering a deeper understanding of customer decision-making. Second, it delivers actionable insights for retailers, highlighting the importance of optimizing fitting room experiences and strategically managing product displays. Given the stronger spillover effect associated with size compared to color, maintaining a diverse inventory of sizes is essential for maximizing sales opportunities.

SD10

Hullet Room

OR and Wellness

Invited Session

OR and Social Wellness

Chair: PENG WANG, Singapore University of Social Sciences, Singapore, Singapore

Co-Chair: Jing Ren, Singapore University of Social Sciences, Singapore, N/A, Singapore

1 - Resource Allocation and Capacity Planning in Intensive-Care Operations

Nicholas Yeo, Xi'An Jiaotong-Liverpool University, Suzhou, China, People's Republic of

Effective resource allocation and capacity planning are critical in intensive care unit (ICU) operations, where fluctuating patient demand, high acuity levels, and constrained resources pose significant challenges to healthcare delivery. This study applies the pipeline-queue (PQ) technique as a novel approach to optimizing ICU resource utilization, aiming to simultaneously balance patient flow dynamics and bed availability under uncertainty. The PQ framework models patient progression through successive ICU stages—such as critical care, step-down care, and eventual discharge—by incorporating stochastic elements including variable arrival rates, heterogeneous service times, and capacity constraints at both the ICU and the stand-down unit (SDU). By integrating queueing theory with operational analytics, the study aims to develop a robust predictive model that supports strategic and real-time decision-making in (i) allocating limited resources between ICU and SDU units, and (ii) minimizing operational disruptions such as delays in admissions, premature discharges, or ICU congestion. The model is evaluated using simulated and real-world hospital data to assess its effectiveness in different demand scenarios. Our findings aim to offer valuable insights for hospital administrators seeking data-driven strategies to enhance ICU efficiency while maintaining high-quality patient care under limited resources.

2 - Understanding Healthcare Course Demand in Singapore: Skillsfuture Trends and Workforce Implications

Jing Ren, Singapore University of Social Sciences, Singapore, Singapore, Shu Ling Tee

With the rapidly evolving healthcare landscape, especially the increasing demand for specialized skills in the post-pandemic era, lifelong learning has become essential for healthcare professionals to remain current and enhance the quality of patient care. The COVID-19 pandemic highlighted the need for adaptability, resilience, and updated knowledge across the healthcare sector, therefore, continuous skills development is no longer optional but a necessity. SkillsFuture Singapore (SSG) promotes lifelong learning and ensure that Singapore's workforce remains agile, competent, and competitive in the long run. It plays a crucial role in planning, developing, and supporting adult education by collaborating with higher education institutions and industry experts.

This study collected information of over 700 healthcare courses from SkillsFuture portal, aiming to analyze course popularity and explore the course designing trends. Using public data, this work leverages multiple data-driven techniques to examine key variables such as course category, duration, fees, enrolment, ratings, and learner feedback. In addition, text analytics methods are adopted to learn the course content. Prediction models are further developed to identify how these variables correlate with course popularity and learner engagements, for forecasting future trends in course demand. The learned insights can help optimize course offerings to better align with industry needs, guide policymaking, and support education providers in designing impactful programs.

3 - Optimizing Medical Escort and Transport Services for the Elderly: a Meta-Heuristic Approach

YUAN WANG, Singapore University of Social Science, Singapore, Singapore

The growing demand for elderly-friendly healthcare services has led to the emergence of medical escort and transport programs, designed to assist seniors with limited mobility or without caregiver support. These services ensure that elderly clients are accompanied to and from medical appointments—such as hospital visits, specialist outpatient clinics, or polyclinics—and are supported throughout their journey, including navigating complex medical facilities and managing appointment logistics.

This study introduces a novel formulation of the **Medical Escort and Transport Problem (METP)**, which jointly considers routing, scheduling, and escort assignment decisions to optimize operational efficiency while maintaining service quality. The METP is motivated by real-world needs in aging societies, particularly for clients who require assistance but lack reliable caregiving support. Given the combinatorial nature and practical constraints of the problem, we prove that METP is NP-hard. To tackle this complexity, we propose a multi-start meta-heuristic framework that balances solution efficiency and accuracy. Our method incorporates dynamic neighborhood search, escort-client matching heuristics, and adaptive restart strategies to explore the solution space effectively. Numerical experiments based on realistic demand scenarios demonstrate the effectiveness of the proposed approach in generating high-quality solutions within acceptable computation time.

This work contributes to the growing field of healthcare operations and elderly mobility services, providing both methodological advances and practical implications for service providers aiming to deliver scalable, compassionate transport solutions in aging urban environments.

4 - Tracking Longitudinal Health Outcomes in Social Prescribing

PENG WANG, Singapore University of Social Sciences, Singapore, Singapore, Caroline Lim

As Singapore faces an aging population and rising rates of chronic illness, social prescribing has gained traction as a complementary approach to traditional healthcare. By connecting patients to non-medical community resources—such as social groups, arts activities, and volunteer opportunities—social prescribing addresses the social and emotional determinants of health that clinical interventions often overlook. This study tracks the health and well-being outcomes of patients enrolled in a social prescribing programme over a one-year period, comparing them with patients who do not participate in such interventions. Outcomes assessed include hospital readmissions, emergency department visits, levels of loneliness, subjective well-being, and community engagement.

Adopting a longitudinal research design, the study captures changes over time to better understand the long-term impact of social prescribing on patients' health trajectories. Quantitative data from surveys will be complemented by qualitative interviews with patients, caregivers, and healthcare practitioners to provide a holistic view of how social prescribing influences individual and system-level outcomes. This approach enables a more nuanced understanding of the patient experience and the broader implications for integrated care in Singapore. Findings from this research will support evidence-based decisions on resource allocation and policy development, with the goal of strengthening community-based healthcare and improving population health outcomes.

5 - Towards Trustworthy Heart Disease Risk Prediction: Explainability Framework for Multimodal Models for Clinical Decision Support in Low-Resource Health Systems

Ira Promasanti Rachmadewi, Monash University, Tangerang, Indonesia, Taufiq Asyhari

Heart disease remains the deadliest disease globally. with especially severe impact in low- and middle-income countries (LMICs) due to limited cardiologist availability, fragmented EHR usage, lack of trust in AI-based information systems and constrained computational infrastructure. An explainability framework of an intelligent risk prediction model that fuses multimodal EHR data, including structured attributes (e.g., vitals, labs) and unstructured formats (e.g., clinical notes, ECGs) is explored to enhance diagnostic accuracy while maintaining operational efficiency in resource-limited environments. This research applies advanced machine learning and transformer-based models with optimization for low-latency deployment.

To overcome the “black-box” nature of AI and support clinical adoption, a layered explainability framework is introduced, augmented with an orchestrator powered by large language models (LLMs). This orchestrator dynamically adapts model explanations to the user’s role and expertise, translating complex inferences into domain-appropriate insights for both clinicians and technical stakeholders. Through natural language interaction, it fosters human-AI collaboration, enabling clinicians to interrogate predictions, validate evidence, and align AI outputs with clinical reasoning.

Evaluation of the framework will be conducted using real-world multimodal EHR datasets to benchmark prediction performance, explainability metrics, and user-centred utility. This research contributes a novel architecture that bridges predictive power with transparency, helping frontline healthcare professionals in LMICs make timely, informed decisions. By combining intelligent prediction with tailored transparency, this research aims to inform the design of trustworthy AI-based decision support tools that are deployable in diverse healthcare ecosystems.

SD11

Morrison Room

Recent Advances in Airline Revenue Management II

Invited Session

Revenue Management and Pricing

Chair: Lan Wu, National University of Singapore, Singapore, Singapore

1 - Leveraging Fully-Parametric and Machine Learning Techniques to Airline Sale Steering Model

wenxuan Shan, National University of Singapore, Singapore, Singapore, Xun Zhang, Weirong Wang, Chung Piaw Teo

The ever-evolving and highly competitive airline industry demands seamless integration of data analytics in decision-making to address complex challenges. With the industry’s revenue size growth averaging 7% annually, bringing the overall figure to US$889 billion in 2019, revenue management becomes a crucial area where insights from data are essential to maximize revenue. This study focuses on presenting accurate demand models for airline ticket context, considering intricacies specific to airline industry such as interdependency between fare classes, non-stationarity and seasonality of demand, and price elasticity of customer. Utilizing two constrasting approaches with different advantages, namely fully-parametric and machinelearning models, this study aims to provide accurate demand forecast under various conditions. Both models were able to capture intricate patterns in airline demands, with consistent results in forecasts.

2 - Revenue Management Under Branded Fare Structure: Mitigating the Spiral Down Effect

Lan Wu, National University of Singapore, Singapore, Singapore, Changchun Liu, Chung Piaw Teo, Xun Zhang

Airline revenue management relies on branded fares, which categorize ticket options into structured fare families with tiered service offerings. Traditional revenue management (RM) models often assume static customer choice behavior, where travelers have fixed preferences for specific fare tiers. However, in practice, customer decisions are context-dependent. A key example is the buy-down phenomenon, where passengers opt for lower-tier fares despite being willing to pay for higher ones. This behavior can trigger a spiral-down effect, where reactive price reductions unintentionally erode long-term revenue.

This paper introduces a novel pricing methodology that integrates an adjusted fare framework with the online Bayes Selector algorithm to dynamically optimize fare structures. Unlike classical RM approaches, our method accounts for changing customer responses, ensuring near-optimal revenue performance even in the absence of precise demand segmentation. Additionally, when airlines can observe customer types and distinguish between non-purchase and no-show behavior, our approach enables more robust revenue management across multiple fare brands.

Our framework consists of a two-phase process: (i) an exploratory phase, where empirical data informs fare adjustments, and (ii) an exploitation phase, where linear programming techniques dynamically refine pricing strategies. Theoretical results establish that our algorithm achieves a logarithmic regret bound, ensuring adaptability even in non-stationary demand environments with uncertain purchase probabilities. Through simulations and a real-world case study, we demonstrate that our approach significantly enhances profitability by better aligning seat allocation and pricing strategies with shifting customer demand.

3 - Data-Driven Pricing for Airline Economic Seats

Jiajie Zhang, National University of Singapore, Singapore, Singapore

We study a data-driven seat pricing problem for Scoot, an economic company in Singapore. We collect login and seat transaction data to classify customers into different segments and fit a specific multinomial logit model (MNL) for each segment, repsectively. Within each MNL, we model utility as a linear combination of 7 itinerary features and seat prices. Based on the demand model, we formulate the seat pricing model as a mixed-integer second order cone program, which can be directly solved by off-the-shelf solvers. We conduct a case study with Scoot by deploying our framework on airlines between Singapore and Sydney, compared with Scoot's current machine learning pricing method.The empirical results demonstrate the superority of our framework.

SD12

Moor Room

Urban Safety and Resilience: EV Service Operations

Invited Session

Revenue Management and Pricing

Chair: Ningwei Zhang, Lanzhou University, Lanzhou, N/A

Co-Chair: Zhiyuan Wang, Tsinghua University, Beijing, N/A

1 - Joint Routing and Inventory Model for Electric Vehicle Battery Swapping-Charging System Under Uncertain Demand

Ningwei Zhang, Lanzhou University, Lanzhou, China, People's Republic of, Yuli Zhang, Shanshan Wang, Long He

This paper addresses the joint routing and inventory management problem in an electric vehicle (EV) battery swapping-charging system (BSCS), consisting of a central battery charging station (CS) and multiple local battery swapping stations (SS). Local SSs experience random EV arrivals, where each vehicle exchanges a depleted battery for a charged one in stock. Batteries that meet the pre-specified charging time threshold are charged at the SS, while the remaining batteries are transported to the CS for centralized charging. The CS

performs the central charging of batteries and periodically redistributes them to SSs via truck. Given the uncertainty in charging times, we aim to determine the optimal initial inventory levels at both the CS and SSs, ensuring that the stock-out probabilities do not exceed pre-specified thresholds. We propose a joint routing and inventory optimization model that minimizes both battery holding and delivery costs. By analyzing the model's structure, we reformulate the original non-convex mixed-integer programming problem into a mixed-integer second-order cone programming problem and develop a constraint generation algorithm with enhanced extended polymatroid. Experimental results based on real-world data validate the efficiency of the proposed algorithm and provide valuable operational insights for managing BSCS.

2 - First-in-First-Out (Fifo) OR Highest-Energy-First-Out (Hefo)? a Tale of Two Electric Vehicle (EV) Battery Swapping Policies

Mengyi Sha, Tsinghua University, Beijing, China, People's Republic of, Wei Qi, Yuli Zhang

Battery swapping is gaining traction in cities and spurs wider EV adoption, as the business mode can alleviate range anxiety efficiently. In this paper, considering random service requests and random status of incoming batteries, we focus on the operations management of an individual battery swapping station, and propose mean-variance distributionally robust optimization (DRO) models to compare FIFO and HEFO swapping policies in different scenarios. The analytical results demonstrate that HEFO policy is superior to FIFO policy in most scenarios, except when all EVs wait for the recharged batteries to reach a target state-of-charge (SOC). We also verify the conclusions via numerical experiments using real-world data sets.

3 - Robust Multistage Location of Battery Swapping Stations

Zhiyuan Wang, Tsinghua University, Beijing, China, People's Republic of, Xian Guo, Lun Ran, Yuli Zhang

Battery swapping is undergoing a revival with the thriving electric vehicle industry and advancements in battery technology. Limited budgets and growing demands make the battery swapping industry can only develop over time. However, there are little literature that concerns the multistage location of battery swapping stations. To fill the gap, this paper aims to address the multistage location of battery swapping stations, considering decision-dependent and time-dependent swapping demands. To the best of our knowledge, this work is the first attempt to address this specific challenge. We propose a dynamic distributionally robust optimization framework that incorporates a specific ambiguity set to encapsulate the time-and-decision-dependent characteristics of the problem. Furthermore, we model the swapping station operation as a continuous Markov process and integrate it into the location model. To efficiently solve the problem, we propose a revised stochastic dual dynamic integer programming algorithm, utilizing a bias-adjusted sampling-based upper-bound estimation. Through a comprehensive numerical analysis, we validate the efficacy of our proposed model and give management insights to guide the battery swapping station planning.

4 - Mechanism of the Dual Credit Policy in Production Competition with Nev Credit Transfer

Haojie Sun, Beijing Institute of Technology, Beijing, China, People's Republic of, Xiang Zhang

The dual credit policy (DCP) regulates both new energy vehicle (NEV) manufacturers by imposing positive new energy vehicle credit requirements and internal combustion engine vehicle (ICEV) manufacturers by setting fuel economy goals. In this way, the DCP can substitute subsidies to promote NEV development and fuel economy improvements using a market-based mechanism. However, there is little study that systematicaly reveals the mechanism. We aim to uncover how the DCP works. To address this, we present a Cournot game model that incorporates both an ICEV manufacturer and a NEV manufacturer within a two-period time horizon. This model encompasses three distinct credit transfer schemes: full-transfer, zero-transfer, and credit pool. Furthermore, we examine four strategies concerning green technology investments: Without implementing, Delay implementing, Halt implementing, and Continue implementing. Our results indicate that the DCP can achieve the dual goals of promoting the achievement of fuel economy target and NEV promotion by adjusting the NEV credit ratio to a certain interval. All the optimal decisions are dependent on the NEV credit ratio. The ICEV manufacturer will choose Continue implementing strategy only when the cost-benefit ratio of technological advancements is lower than the threshold. The zero-transfer scheme exhibits superior performance compared to the full-transfer scheme in terms of achieving the dual goals of the DCP. The credit pooling scheme may generate increased profits for both NEV and ICEV manufacturers when the cost-benefit ratio of technological advancement is at a high level.

SD13

Bailey

Advancing Healthcare Systems through Innovative Simulation and Optimization

Invited Session

Simulation and Optimization of Healthcare Systems

Chair: Xiangyu Yang, Shandong University, Jinan, N/A

1 - Guided Empathetic Reasoning of Large Language Models: a Novel Deep Learning Method for Generating Empathetic Responses to Enhance Conversational Service Satisfaction

Gang Chen, Fudan University, Shanghai, China, People's Republic of

While enhancing the empathy capabilities of large language models signifies immense commercial value, identifying users' empathy needs based on their conversations with an AI agent is non-trivial due to the difficulty in gauging their subjective needs and understanding their minds beyond the conversations. In this computational design science research, we design a novel empathetic mind simulation-empowered deep learning (EMS-DL) method, which identifies users' cognitive needs through in-conversation need spreading and activation with a cognitive mind simulator, recognizes users' emotional needs through emotion contagion and activation with an emotional mind simulator, and optimizes its empathy need identification capability by drawing guidance from human agent experience with a generative contrastive learning-based optimization scheme. Empirical evaluation based on two (patient-doctor and e-commerce customer-sales agent) conversation datasets demonstrates the superiority of the empathy needs identified by EMS-DL over those by state-of-the-art alternatives, in terms of both empathy needs-empowered satisfactory response generation and empathy needs-based conversational satisfaction prediction. User studies demonstrate the practical effectiveness of EMS-DL in enhancing users' conversational satisfaction by identifying and addressing their empathy needs. Explanatory analyses render insights into users' need spreading paths and emotion contagion patterns, as well as how the spread needs and contagious emotions enhance empathy need identification.

2 - Adaptive Simulation Optimization and Its Application

Zhenxia Cheng, Shanghai Jiao Tong University, Shanghai, China, People's Republic of

This study focuses on the constrained selection of the best alternative, considering the selection of the alternative that has the best primary mean performance and satisfies finitely many secondary mean constraints. In order to achieve the statistical probability guarantee of correct selection (PCS), most existing procedures adopt the sampling rule that allocates samples equally among alternatives in contention. Unlike the prevalent procedures, this work proposes constrained selection of the best procedures that utilize adaptive sampling rules and, meanwhile, can sustain the statistical guarantee by using a new theoretical framework. In addition, a specific sampling rule is designed for the constrained procedure by solving an optimization problem with the aim of minimizing the asymptotic total sample size. Finally, numerical experiments are conducted to showcase the efficacy of our proposed procedures compared to representative ones.

3 - Patient-Centered Doctor Recommender System of Online Health Communities: a Multidimensional Sequence Alignment Approach

Xue Meng, Shandong University, Jinan, Shandong, China, China, People's Republic of, Jianghua Zhang, Xuemei Fu

This study develops a patient-centered Doctor Recommender System (DRS), facilitating patients' doctor choice performance by calculating doctors' "biological distance" based on Multidimensional Sequence Alignment (MSA) method. The distance is measured by doctor-patient interaction attributes, which are acquired from a combination of Levenshtein Distance and opinion mining technology. The XGBoost is used to identify the importance of these attributes. The precision of patient-centered DRS is 72.60 %, and the recommended doctors are significantly different from the doctors chosen by patients lacking medical knowledge. It not only furnishes practical guidance for OHCs in increasing patients' consulting satisfaction, but also illuminates the important properties in improving doctors' popularity.

4 - A Penalty-Based Exterior-Interior Point Stochastic Approximation Method for Stochastically Constrained Simulation Optimization

Xiangyu Yang, Shandong University, Jinan, China, People's Republic of, Jiaqiao Hu, Jianghua Zhang

We propose a penalty-based method for solving stochastically constrained simulation optimization problems with differentiable structure. We introduce a novel class of penalty functions and reformulate the original constrained problem as a sequence of unconstrained subproblems, which are approximately solved using a multi-timescale stochastic gradient descent framework. The method combines features of both exterior and interior penalty approaches: it acts like an exterior method to guide an infeasible solution toward the feasible region and transitions to an interior method once feasibility is achieved. In contrast to existing methods, our approach operates fully online, eliminates the need for explicit feasibility checks, and requires no auxiliary optimization subroutines during the iterative process. Under suitable assumptions, we establish almost sure convergence of the algorithm and derive an explicit convergence rate for the final iterate via a bias–variance decomposition of the stochastic gradient estimators. Numerical experiments on both synthetic and real-world datasets are performed to illustrate the practical effectiveness and robustness of the proposed method.

5 - Optimization Strategies for Online Advertising Portfolios of Individual Sellers Under Budget Constraints: a Conversion Rate-Driven Decision Framework

Bo Song, School of Management, Xiamen University, Xiamen, China, People's Republic of, Miao Li, Zhaowei Miao, Yongquan Lan

This paper investigates online advertising optimization strategies for individual retailers under strict budget constraints, emphasizing the allocation of resources among content, sales, traffic ads, and their synergistic combinations. By integrating key performance metrics—including conversion rates, click-through rates, and bounce rates—the research constructs profit-maximization models for each advertising type and their hybrid forms. This paper applies numerical optimization techniques to derive decision thresholds and prioritize budget allocation under varying market conditions. Key findings demonstrate that: 1) Traffic ads yield the highest returns when initial content conversion rates are low, as they compensate for weak organic engagement by amplifying visibility; 2) In moderate conversion scenarios, combining content and sales ads achieves optimal efficiency, leveraging enhanced product trust and improved informational appeal; 3) At high conversion levels, prioritizing content ads preserves brand equity and minimizes margin erosion from excessive discounts. The paper establishes decision intervals based on conversion rate thresholds to guide retailers at distinct operational stages and advocates adaptive strategies that align with evolving consumer behaviors and platform dynamics. The research further highlights the diminishing marginal returns of traffic ads under saturated exposure and quantifies trade-offs between short-term sales boosts and long-term brand sustainability. Practical recommendations include dynamic budget reallocation protocols and scenario-specific hybrid campaigns, offering retailers actionable frameworks to maximize ROI in competitive digital marketplaces.

SD14

Wilberforce

Simulation and AI

Invited Session

Simulation Society

Chair: Guangxin Jiang, Harbin Institute of Technology, Harbin

1 - Ranking and Selection with Unknown Sampling Variances

Jianzhong Du, University of Science and Technology of China, Hefei, China, People's Republic of, Siyang Gao, Ilya Ryzhov

Expected improvement (EI) is a common ranking and selection (R&S) method for selecting the optimal system design from a finite set of alternatives. Ryzhov (2016) observed that, under normal sampling distributions with known variances, the limiting budget allocation achieved by EI was closely related to the theoretical optimum. However, when the variances are unknown, the behavior of EI was quite different, giving rise to the question of whether the optimal allocation in this setting was totally distinct from the known-variance case. This research

solves that problem with a new analysis that can distinguish between known and unknown variance, unlike previously existing theoretical frameworks. We derive a new optimal budget allocation for this setting, and confirm that the limiting behavior of EI has a similar relationship to this allocation as in the known-variance case.

2 - Parallel Simulation of Large-Scale Queueing Networks

Tan Wang, University of Science and Technology of China, Hefei, China, People's Republic of, Jeff Hong, Yingda Song

With the accelerating pace of digital transformation and the deep integration of artificial intelligence technologies, the scale of queueing networks continues to expand. However, traditional discrete-event simulation encounters two major challenges when applied to large-scale queueing networks. First, the sheer volume of events results in a substantial increase in computational cost. Second, the strong temporal dependencies among events hinder the effective application of conventional parallel acceleration strategies. To address these challenges, we propose two simulation schemes based on Euler approximation, namely the backward and forward schemes. We demonstrate that the two schemes provide stochastic upper and lower bounds for the system state, while ensuring that the approximation error remains bounded over the simulation horizon. With the recommended choice of time step, we further demonstrate that our approximation schemes exhibit diminishing asymptotic relative error as the system scale increases, while significantly reducing computational complexity compared to traditional discrete-event simulation. Moreover, we design parallel algorithms that further enhance simulation efficiency, achieving speedups up to tens of thousands times. Overall, this study provides both theoretical insights and methodological foundations for intelligent decision-making in large-scale queueing systems.

3 - Large Language Model Agents for TIME Series Forecasting

jin yang, Harbin Institute of Technology, Harbin, China, People's Republic of, Guangxin Jiang

Recent advances in large language models (LLMs) have demonstrated significant potential for time series forecasting, primarily due to their powerful feature extraction capabilities. However, research exploiting the reasoning, planning, and information integration abilities of LLM-driven agents for time series forecasting is still in its early stages. In this paper, we propose a novel multi-agent collaborative framework leveraging these capabilities to automate time series forecasting tasks. Our framework consists of specialized agents dedicated to problem definition, multi-source information integration, transformation of complex raw data into structured formats suitable for modeling, and automated forecasting model design and optimization. To facilitate collaboration among these specialized agents, we introduce a central agent responsible for coherent global management and effective coordination. Specifically, an information-sharing mechanism is introduced to dynamically integrate insights independently discovered by individual agents, thereby enhancing overall modeling performance. Experimental evaluations on diverse time series forecasting tasks demonstrate that our proposed framework achieves promising results in terms of effectiveness and feasibility.

4 - Supply Chain Decision-Making Based on Multi-Agent Ai Simulation

Haojie Wang, Harbin Institute of Technology, Harbin, China, People's Republic of, Guangxin Jiang

With the systematization and integration of production, supply chain networks are increasingly showing highly complex characteristics. The entities within the network are highly heterogeneous, the relationships between entities are intricate, and the interactions between entities and system constitute a complex dynamic network. At the same time, the decision-making and interaction behaviors between enterprises in the real world are diverse and difficult to predict. Due to the subjective participation of decision makers, the management behaviors such as orders, production, and inventory between supply chain nodes often make it difficult for traditional supply chain simulation models to flexibly describe the interactions between enterprises and the response to the external environment. With the support of prompt engineering and retrieval-augmented generation technology, the large language model not only has the ability to autonomously understand, perceive, plan, remember, and use tools, but also can highly simulate human behavior in a specific complex environment. This paper will use the large language model to empower heterogeneous agents in the multi-agent supply chain network model and construct the Multi-Agent AI supply chain network model with learning and adaptive capabilities. The model can finely simulate the decision logic of decision makers in the supply chain, provide assistance for the modeling of the

collaboration and interaction process between upstream and downstream enterprises, and thus simulate the dynamic decision-making behavior in complex supply chain systems.

5 - Gradient-Based Optimization for Dynamic Sample Selection in Active Learning

You Yuanyuan, Harbin Institute of Technology, Harbin, China, People's Republic of, Chen Ying

Deep active learning seeks to improve model performance by selecting representative samples from unlabeled data pools with limited annotation budgets. Recent studies have shown that uncertainty measures based on training dynamics can identify informative samples, yet current methods face two main challenges. First, during the initial phases of active learning, model performance instability undermines the reliability of uncertainty-based sampling strategies, leading to suboptimal annotation decisions. Second, existing training dynamics-based methods fail to assess training dynamics in each iteration. To address these issues, we propose a gradient-driven dynamic framework featuring two key innovations: a) We developed a gradient projection sample selection module, which evaluates the projected values of the entropy gradient for unlabeled data along historical optimization directions. This module reduces reliance on error-prone early predictions by prioritizing samples aligned with stable parameter update directions. b) We designed a gradient-diversity-based adaptive learning rate module to accelerate convergence during active learning training. Our theoretical analysis offers three guarantees: First, selecting and annotating unlabeled data with high projected gradients along cumulative optimization directions decreases the upper bound of test loss. Second, our method reduces gradient variance of selected samples compared to entropy-based uncertainty measures, ensuring monotonic convergence under non-convex objectives. Finally, we derive that the gradient diversity-driven learning rate adaptation minimizes the bias of training dynamics estimation in initial phases, thereby improving the accuracy of uncertainty quantification. Experiments on CIFAR-10 and CIFAR-100 demonstrate superior performance over seven state-of-the-art baselines with identical labeling budgets, and ablation studies validate the necessity of each component.

SD15

Fairmont Ballroom

Sustainable Operations in Urban Logistics

Invited Session

Transportation Science and Logistics

Chair: Lei Zhao, Tsinghua University, Beijing

1 - A Neural Combinatorial Optimization Model for Facility Location Problems

Feng Guo, Sichuan University, Chengdu, China, People's Republic of, Zhaoxia Guo, Bowen Jia, Tao Huang

The facility location problem, a classic challenge in combinatorial optimization, involves determining optimal placements for facilities such as delivery outlets and charging stations. This study examines facility location problems in environments constrained by geographic barriers. Inspired by the success of the AlphaGo Zero algorithm in the game of Go, a novel neural combinatorial optimization model is developed to address these problems. Extensive experiments are conducted to evaluate the model’s effectiveness across two distinct types of facility location problems. The results indicate that the proposed model outperforms Gurobi and two representative meta-heuristics. In particular, compared with Gurobi, the proposed model reduces computation time by up to 23 times, with a relative performance difference of less than 1.3%.

2 - From Resisting to Feeding Algorithms: Adaptation Strategies and Improved Wellbeing of Couriers

Junjie Xiong, Tsinghua University, Beijing, China, People's Republic of, Chen Wang, Lei Zhao, Jan Fransoo

To cope with the algorithmic control of digital platforms, couriers have developed various strategic behaviors. This study adopts a qualitative research method based on grounded theory, summarizing couriers’ adaptive process from resisting algorithms to "feeding" algorithms and examining its impact on their well-being. Behaviors of resisting algorithms involves bypassing platform rules (e.g., fake GPS usage) to regain autonomy,

while behaviors of "feeding" algorithms refers to couriers providing specific behavioral data (e.g., maintaining a fixed daily task volume) to reduce task allocation uncertainty. The study also finds that couriers accelerate adaptation through peer support networks and similar mechanisms, prioritizing income while still attending to well-being. Furthermore, the research connects the adaptation process (from resistance to "feeding") with the person-environment fit theory in technology acceptance. A questionnaire is designed to explore how environmental factors and algorithmic features influence couriers' adaptation and well-being, offering insights for more well-being-oriented algorithm design.

3 - Courier-Centered Order Dispatching for on-Demand Delivery Platforms

Zhiqin Zhang, Singapore Management University, Singapore, Singapore, Si Zhang, Hoong Lau

The optimization of dispatch algorithms in on-demand (e.g., food) delivery platforms traditionally prioritizes system-wide efficiency, typically by minimizing total travel distances for couriers. However, such global optimization can inadvertently create workload imbalances, causing perceptions of unfairness among couriers, negatively affecting morale and operational sustainability. Motivated by the challenge on improving platform worker welfare, we investigate the joint optimization of delivery efficiency and workload fairness in food delivery courier systems.

We define this problem as a multi-objective optimization task, seeking to minimize both total travel distance across the couriers and individual workload imbalance. To tackle this computationally intensive real-time decision-making challenge, we propose a Graph Neural Network (GNN)-based deep learning approach. It involves two interrelated neural models: a bipartite GNN that handles courier-order assignment, and another spatial GNN capturing geographical relationships among order nodes and courier locations. We introduce a data-driven preprocessing step to predict courier acceptance of orders based on historical behavior, embedding courier preferences into the dispatch decision-making process. Training the proposed model via a GFlowNet-based (GFN) framework enables efficient generation of diverse, high-quality dispatch solutions under various fairness-efficiency trade-offs. Our proposed methodology is also scalable as it draws upon the GFN approach where the inference step to generate solutions is cheap computationally.

Our initial results, derived from the mathematical formulation of the problem, confirm the potential of achieving a more equitable workload distribution among couriers while maintaining system-wide efficiency. Experimentation with our proposed learning-based approach demonstrates its ability to generate diverse, high-quality dispatch solutions in real time.

4 - Service Network Design in Intra-City Express Delivery

Lei Zhao, Tsinghua University, Beijing, China, People's Republic of, Tirui Cao, Yixiao Huang, Alexandre Jacquillat, Alexandria Schmid

Intra-city express delivery is an emerging service that enables customers to send and receive packages in the same city within a short delivery window. Its time- and cost-efficiency relies on a dedicated intra-city service network, where packages are transported between and consolidated at a number of satellite stations located in the city. A key challenge in designing and operating such service networks lies in the limited docking capacity to unload/load packages at satellite stations, which often leads to congestion and delays. Besides increasing docking capacity at certain satellite stations, one may also reroute a subset of intra-city packages using the spare capacity of the inter-city service network (independently designed for inter-city express delivery). This mitigates congestion at the satellite stations and shortens the delivery time for the remaining packages, yet introduces extra detour and sorting time for the rerouted packages. We show in a theoretical model that, when these trade-offs are properly balanced, leveraging the inter-city service network can shorten average delivery time and reduce total operational cost. Further, we formulate the intra-city express delivery service network design problem as an integer programming model and explore two solution algorithms: a column generation-based algorithm and a multi-arc generation-based algorithm. Based on real world instances from a major express delivery carrier in China, our preliminary numerical experiments demonstrate the scalability of the proposed algorithms and validate the value of leveraging the inter-city service network to improve service quality and reduce operational cost in intra-city express delivery.

SD16

VIP A

Innovations in Socially Responsible Operations

Invited Session

Sustainable and Socially Responsible Operations

Chair: Xiaoquan Gao, Singapore Management University, Singapore, N/A, Singapore

1 - Regulating Adaptive Medical Artificial Intelligence: Can Less Oversight Lead to Greater Compliance?

Xin Fang, Singapore Management University, Singapore, Singapore, Jiayi Lai, Liang (Leon) Xu, Tinglong Dai

As of June 2024, the U.S. Food and Drug Administration (FDA) has approved 950 medical artificial intelligence (AI) devices. The current regulatory framework freezes AI algorithms after approval, requiring new submissions for updates to ensure compliance with Good Machine Learning Practices (GMLP). This approach imposes a significant administrative burden, while hindering the ability of AI algorithms to learn from new data. To address these challenges, the FDA has explored a novel pathway known as Predetermined Change Control Plans (PCCP), allowing developers to outline future changes during initial submissions and exempting approved changes from regulatory review. Yet, the impact of this exemption on GMLP compliance remains uncertain. In this paper, we model the strategic interaction between a developer and a regulator in a two-stage game with asymmetric information. The developer may choose to follow or deviate from GMLP in developing and retraining the AI algorithm, whereas the regulator reviews the marketing-clearance application for approval. Our analysis shows that, contrary to intuition, less review can actually lead to greater compliance. This scenario arises, even without considering the administrative burden saved, when (1) auditing capability is moderate and (2) the potential for efficiency improvements through retraining is substantial. Conversely, reclearance is valuable when regulatory review effectively detects noncompliance or when efficacy improvements from retraining are unlikely. We also show adaptive algorithms offer advantages over frozen algorithms in terms of not only improved device efficiency but also greater compliance. Interestingly, these advantages are particularly salient when regulatory oversight has limited ability to detect noncompliance.

2 - Incentives In Outcome-Based Regulation for Lung Transplantation

Taewoo Lee, Seoul National University, Seoul, Korea, Republic of

The past twenty years have seen the enactment of regulations on solid organ transplantation outcomes in the United States. These regulations are designed to identify transplant programs with below-average outcomes and incentivize improvements in post-transplant outcomes. However, the severe penalties associated with these regulations have led some programs to engage in adverse patient selection, rejecting critically ill patients who require immediate transplantation. In this paper, we present a game-theoretic model to design an incentive and reimbursement system that encourages transplant programs to accept high-risk patients, promoting fairness in patient selection while ensuring sustainable and socially beneficial operations for individual programs.

3 - Flexible Versus Dedicated Technology Choice and Capacity Investment in Commodity Processing with Input Substitution and Its Application in Battery Recycling

Zuyi Liu, Lee Kong Chian School of Business, Singapore Management University, Singapore, Singapore, Onur Boyabatli, Çerağ Pince

With the continued growth of electric vehicle (EV) sales, tens of millions of EV batteries are expected to reach end-of-life in the coming decade. This paper examines the flexible versus dedicated technology choice and capacity investment decisions of a battery recycling firm recovering useful materials from various types of used EV batteries (e.g., lithium-ion, solid-state, lead-acid batteries, and others). The firm processes two commodity inputs (types of used batteries) to produce three outputs in fixed proportions: a common commodity output from both inputs and a unique commodity output from each input. All inputs and outputs can be transacted through spot markets; the firm also has a contracted demand for the common output. We develop a two-stage stochastic model: In the first stage, the firm makes decisions on the flexible versus dedicated technology choice and the capacity investment with the chosen technology while facing uncertainty in spot prices for both inputs and

outputs. In the second stage, uncertainties are realized and the firm chooses the optimal processing volumes for each input (which then determines the production volume for each output). We characterize the firm's optimal decisions and examine how uncertainty in various spot prices affect these decisions as well as firm's profitability.

SD17

VIP B

Scenario Planning and Predictive Modeling for Strategic Decisions

Contributed Session

Decision Support Systems

Chair: Michael Bendersky, Holon institute of technology, Holon, Israel

1 - Narrative Logic and Helpfulness of Online Reviews: An Analysis Based on Topic Coherence and Sentiment Coherence

Kexin Tan, Xi'an Jiaotong University, Xi'an, China, People's Republic of, Xinyu Sun

Online consumer reviews serve as a critical manifestation of electronic word-of-mouth (eWOM), exerting substantial influence on consumer decision-making processes, with more informative reviews demonstrating disproportionately greater impact. While extant research has predominantly focused on the linguistic content of reviews as determinants of perceived helpfulness, the role of narrative logic in review composition has received limited scholarly attention. Using Aspect-Based Sentiment Analysis (ABSA), this research analyzes the narrative coherence of over 400,000 reviews on Yelp. The empirical findings reveal that both topic coherence and sentiment coherence exert positive effects on review helpfulness. Further analysis demonstrates significant contextual heterogeneity in these effects: (1) The impacts of both coherence types are amplified in negative reviews, longer reviews, and those associated with chain restaurants; (2) Sentiment coherence manifests stronger influence for reviews authored by professional users; (3) Topic coherence demonstrates higher impact for reviews of low-rated restaurants. The findings extend the literature on online review text analysis by introducing narrative logic as a novel analytical dimension, and provide theoretical underpinnings for improved review recommendation and ranking algorithms.

2 - Artillery Firing Shift with Multiple Registration Targets

Michael Bendersky, Holon institute of technology, Holon, Israel

Firing Shift is the shifting of artillery fire from one target ('registration') to another by the application of corrections determined from the adjustment of the initial firing data for the first target to the second ('application'). In this work a shift method based on simultaneous firing at multiple registration targets is introduced. A method is presented for using the observed data to explicitly solve for major environmental parameters which affect ballistic trajectories and subsequently update them for the application mission. The distribution of the remaining errors is determined and compared with that of other correcting algorithms. Numerical examples show the potential of the proposed method.

Monday, July 21, 8:00 AM - 9:30 AM

MA01

Enterprise

Operations in Platform Economy

Invited Session

Applied Probability Society

Chair: Hailun Zhang, The Chinese University of Hong Kong, Shenzhen, Shenzhen, N/A

1 - Near-Optimal Control and Design of Critically Loaded Service Systems when Service and Patience TIMES are Dependent

Lun Yu, Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Ohad Perry

Customer behavior has significant impacts on the operations and performance of service systems. For example, even a small amount of abandonment, which is due to customer impatience, has dramatic impacts on key performance measures, and in turn, on the optimal staffing and control of such systems. However, in contrast to the prevalent models in the literature, the time that customers are willing to wait in queue for service often depends on their service requirement, and the impacts of this dependence can be sufficiently substantial so as to render those existing models futile.

In this paper, we consider large service systems that handle customers with dependent service and patience times, when those systems are staffed according to the Square-Root Staffing (SRS) rule in order to operate in the quality-and-efficiency (QED) asymptotic regime. We first develop bounds for the steady-state queue, delay and customer loss and show that the value of a single dependence scalar determines whether SRS guarantees that the system operates in the QED regime. We then propose design-and-control mechanisms that guarantee that SRS is near-optimal, in the sense that the QED regime is achieved. We further show that our near-optimality results are robust to the dependence structure, requiring no knowledge of the exact service-patience distribution, which is hard to estimate in practice. Finally, we provide simple methods for computing the optimal control parameters. The effectiveness of our analytical results and of our proposed solutions to the dependence-caused congestion are supported by extensive simulation studies.

2 - Dynamic Control of An Outsourcing Service System with Reentrant Customers

Zhenghua Long, Nanjing University, Nanjing, China, People's Republic of, Jin Fang, Wei You, Jiheng Zhang

This study addresses a control problem for a service system comprising both in-house and outsourcing service stations. Customers who have had their service requests served but not resolved by an outsourcing server may return to the in-house station for further assistance. The in-house servers are able to serve both incoming and reentrant customers, while the outsourcing servers can only serve incoming customers. To asymptotically minimize the long-term average delay cost for the entire system, we design a control policy that combines an admission control for the outsourcing service station and a dynamic priority control for the in-house service station based on fluid analysis. The admission control (which has the same effect as a threshold control but without the need to know the threshold) may intentionally idle some of the outsourcing servers, while the dynamic priority control results in shared service resources in the in-house service station. We demonstrate the effectiveness of our policy through extensive simulation experiments.

3 - Multi-Sided Matching: a Cost Balancing Algorithm

Hailun Zhang, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of

In various matching scenarios such as ridesharing, game platform matching, and beyond, a fundamental dilemma emerges: the delicate balance between swiftly securing a match with minimal waiting costs and the potential benefits of delaying decisions for a more optimal pairing. This tradeoff pervades numerous domains, underlining the essence of efficient matching mechanisms.We present a versatile framework for modeling N-sided matching with general cost structures. Drawing insights from the realm of balancing algorithms in inventory management literature Levi et al. (2007), we introduce a straightforward yet effective cost-balancing algorithm, with a competitive ratio of 2+ \epsilon. We establish a lower bound of $\frac{\sqrt{5}+1}{2}$ for such problems. Through numerical studies in game matching and on-demand delivery, we validate the algorithm's efficiency in practical business applications.

4 - Discrete Choice Models with Piecewise Linear Utility: Modeling, Estimation and Pricing

Chenxu Ke, Shanghai University of Finance and Economics, Shanghai, China, People's Republic of

This paper incorporates a piecewise linear structure into the utility-price relationship of the classic multinomial logit (MNL) model, and studies the associated operations problems such as estimation and pricing. The derived model is referred to as the piecewise MNL model and provides greater modeling flexibility by allowing

consumers to exhibit asymmetric price sensitivities around unknown inflection points. We study the model identification for the piecewise MNL and further propose a maximum likelihood estimator (MLE) for its calibration by real data. Due to the presence of inflection points, the log-likelihood function is non-differentiable, which poses major challenges to both numerical and statistical analyses. We propose a novel profile-based numerical optimization procedure which locates the MLE efficiently and further establish statistical guarantees for the MLE based on the empirical process theory. We fully solve price optimization under the piecewise MNL and show that the optimal pricing policy can be quite different from the standard MNL. In particular, the "equal-(adjusted-)markup" policy is no longer optimal and the optimal price may not even be unique. For price competition, we show that multiple equilibria may exist depending on the magnitude of the inflection point and then fully characterize the conditions for the existence of each equilibrium. Our extensive numerical experiments on synthetic and real data suggest that the piecewise MNL can improve model fitting and prediction accuracy compared with popular choice models in the literature, and ignoring varying price sensitivities could lead to sub-optimal solutions or even substantial losses for firms.

MA02

Mercury

Airline Operations Management

Invited Session

Aviation Applications

Chair: Siqi Guo, Tongji University, Shanghai, N/A

1 - Flight Recovery with Pre-Cancellation

Yi Su, Tongji University, Shanghai, China, People's Republic of

Airlines often adopt a wait-and-see strategy to disruptions, resulting in canceling flights at the last moment. This not only incurs extra compensation costs but also significantly affects passengers' travel experiences. To mitigate these losses, we introduce the concept of flight pre-cancellation, defined as canceling flights one to several days before departure. To make pre-cancellation decisions with regards to the stochastic future weather condition, we develop a two-stage stochastic model aimed at minimizing the overall recovery cost. To solve this model, we design a Lagrangian Dual Decomposition (LDD) approach, which efficiently decomposes the model into scenario-independent sub-models. These sub-models are then solved by a column generation framework. Additionally, we propose a Dual-based Variable Evaluation Strategy (DVS) to accelerate the solving process of LDD. We evaluate the effectiveness and efficiency of our model and algorithms using real operational data from three airlines, tested under real typhoon scenarios. The computational results show that LDD can obtain optimal LP solutions and near optimal IP solutions. Compared to the baseline column generation algorithm, the solution times for LDD and LDD-DVS are reduced by 41% and 46%, respectively. Additionally, tests conducted on real typhoon scenarios demonstrate that, by incorporating pre-cancellation decisions, it achieves an average cost saving of 17% compared to solutions only considering real-time cancellation decisions.

2 - Tactical Fleet Assignment and Crew Pairing Problem with Crew Flight TIME Allocation

Fan Xiao, Shanghai University, Shanghai, China, People's Republic of

Driven by the challenges of crew resource shortages and mismatches between crew and aircraft resources in the aviation industry, we address a tactical fleet assignment and crew pairing problem with crew flight time allocation (TFACPP).

We conduct a comparative analysis of two widely used decomposition approaches, column generation and Benders decomposition, to solve this large-scale combinatorial optimization problem. In the column generation approach, two types of subproblems are designed to generate fleet assignment schedules and crew pairings. Specifically, the fleet assignment problem is solved using a time-space network model directly, while crew pairings are generated through a label-setting shortest path algorithm. Additionally, in the Benders decomposition approach, the complete fleet assignment problem serves as the master problem, and the crew pairing problem acts as the subproblem. Both the master problem and subproblems are a tackled using the column generation algorithm. A computational comparison between the two approaches is then conducted using

industrial cases from a medium-sized airline in China. The results demonstrate the effectiveness and high quality of the solutions.
Moreover, we develop a quantitative method based on shadow prices (dual variables) to evaluate the scarcity of crew and aircraft resources as well as their matching degree, which provides valuable managerial insights for crew training and aircraft acquisition.

3 - Integrated Airline Crew Pairing and Rostering Problem

Siqi Guo, Tongji University, Shanghai, China, People's Republic of

In this paper, we study a nested multi-stage set covering problem, in which the results of the previous stage become the entities that comprise the next stage decisions. We propose a nested set covering model, which can be viewed as a traditional set covering problem with DDOIs. Different from the standard DDOIs, this new type also increases the dimension of the dual problem, and we herein name it as Lift-DDOIs. Furthermore, we demonstrate that a large number of DDOIs/Lift-DDOIs could exacerbate the problem degeneracy. To resolve this issue, we derive strengthened reduced costs that can prescreen and prune degenerate variables, and a pruning-and-pricing mechanism that utilizes both standard and strengthened reduced costs to price out promising variables. An efficient nested column-and-row generation algorithm is developed to take advantage of both dual stabilization and degeneracy alleviation. We also demonstrate the above findings and methods can be adapted to the nested set packing problems. A comprehensive computational study on the airline crew scheduling problem shows that the proposed Lift-DDOIs and the column-and-row generation algorithm with pruning-and-pricing mechanism can reduce the solution time by an average of 73.52% compared to a standard column generation approach for traditional set covering problems.

4 - Machine Learning: Copy OR Learn? - An Example from Aircraft Recovery

Yuqing Liao, Tongji University, Shanghai, China, People's Republic of

The emergence of artificial intelligence and machine learning has introduced promising avenues for solving large-scale combinatorial optimization problems. In most cases, training machine learning algorithms requires extensive real or simulated data, and the model’s performance is generally reliable when new instances closely resemble those in the training set. From this perspective, machine learning can be seen as "copying" outcomes from seen scenarios. However, limited data availability in many real-world applications—particularly in airline operations—makes it impractical or costly to simulate all possible disruption scenarios. This study investigates whether machine learning can infer effective decisions for novel, significantly different scenarios using only partial training data—essentially, whether it demonstrates true "learning" capability beyond simple replication. We evaluate the model’s generalization performance in the context of flight recovery, focusing on airline networks and disruption events (e.g., mechanical failures, weather-related delays) that differ significantly from the training data in terms of time, location, and structural characteristics. Our findings aim to assess the robustness and adaptability of machine learning-based approaches in airline disruption management, and more broadly, to provide insight into the potential of machine learning to support decision-making in operations research under data-scarce and dynamically evolving environments.

MA03

Minto

Innovation in Online Marketplaces

Invited Session

Auctions and Market Design

Chair: Yiding Feng, Hong Kong University of Science and Technology, Hong Kong, Hong Kong

Co-Chair: Yiding Feng, Hong Kong University of Science and Technology, Hong Kong, Hong Kong

1 - When Learning Hurts: the Cost of Algorithmic Instability on Platforms

Yilun Chen, CUHK Shenzhen, Shenzhen, China, People's Republic of, Jiaqi Lu

Learning algorithms increasingly shape how digital platforms allocate demand across supply-side agents—such as sellers, freelancers, or tutors—using bandit-style exploration to optimize long-run performance. However, this

exploration introduces variance in agent exposure, especially when agent quality is hard to distinguish. Such fluctuations are often undesirable: they impose planning costs on agents—idle time, wasted resources, or missed opportunities—that may ultimately reduce platform efficiency and revenue.

Motivated by this, we introduce a multi-objective bandit learning framework that seeks to minimize (1) classical regret from allocating demand to suboptimal agents (arms), and (2) variance in the realized demand (i.e., number of pulls) each agent receives. We show that popular regret-optimal algorithms (UCB, Thompson Sampling and EXP3) perform poorly on the second objective by inducing high demand fluctuation. We then establish a fundamental trade-off between the two objectives by charaterizing a Pareto Frontier. We propose a simple class of algorithms that achieves Pareto optimality.

2 - Two-Sided Pricing and Learning with Buyer Choice: a Generalized Sgd Method

Meichun Lin, Singapore Management University, Singapore, Singapore, Woonghee Huh

Motivated by online used-car platforms, we study pricing decisions for purchasing and selling multiple products in a two-sided market. While supply for each product is independently acquired from individual sellers, demand is substitutable and follows a choice model. With uncertainty from both supply and demand, a platform sequentially adjusts purchase and selling prices to maximize profit while satisfying inventory constraints. Moreover, the platform does not know in advance the supply and demand as functions of prices. The pricing decisions should account for learning of the supply and demand functions in addition to the two sides of uncertainty and potential stockouts. When the supply and demand functions are known, we find a class of fixed-price policies that achieve provably good performance. Our finding suggests that we can target a fixed-price policy rather than the state-dependent optimal policy when the supply and demand functions are unknown. We propose a novel generalization of the classical SGD method for learning the target prices on each side. Our algorithm with the generalized SGD is asymptotically optimal when the planning horizon is large.

3 - Personalized Policy Learning through Discrete Experimentation: Theory and Empirical Evidence

Ruohan Zhan, The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong, Zhiqi Zhang, Zhiyu Zeng, Dennis Zhang

Randomized Controlled Trials (RCTs), or A/B testing, have become the gold standard for optimizing various operational policies on online platforms. However, RCTs on these platforms typically cover a limited number of discrete treatment levels, while the platforms increasingly face complex operational challenges involving optimizing continuous variables, such as pricing and incentive programs. The current industry practice involves discretizing these continuous decision variables into several treatment levels and selecting the optimal discrete treatment level. This approach, however, often leads to suboptimal decisions as it cannot accurately extrapolate performance for untested treatment levels and fails to account for heterogeneity in treatment effects across user characteristics. This study addresses these limitations by developing a theoretically solid and empirically verified framework to learn personalized continuous policies based on high-dimensional user characteristics, using observations from an RCT with only a discrete set of treatment levels. Specifically, we introduce a deep learning for policy targeting (DLPT) framework that includes both personalized policy value estimation and personalized policy learning. We prove that our policy value estimators are asymptotically unbiased and consistent, and the learned policy achieves a root-n regret bound. We empirically validate our methods in collaboration with a leading social media platform to optimize incentive levels for content creation. Results demonstrate that our DLPT framework significantly outperforms existing benchmarks, achieving substantial improvements in both evaluating the value of policies for each user group and identifying the optimal personalized policy.

4 - Information Disclosure Makes Simple Mechanisms Competitive

Yingkai Li, National University of Singapore, Singapore, Singapore, Yang Cai, Jinzhao Wu

In classical mechanism design, the prevailing assumption is that the information structure about agents' types is exogenous. This assumption introduces complexity, especially with multi-dimensional agent types, leading to mechanisms that, while optimal, may appear complex and unnatural. Furthermore, Hart and Nisan (2019) show that the gap between the performance of any simple mechanism and the optimal solution could be potentially unbounded. We challenge this conventional view by showing that simple mechanisms can be highly competitive if the information structure is endogenous and can be influenced by the designer.

We study a multi-dimensional generalization of a single-dimensional model proposed by Bergemann and Pesendorfer (2007), where the designer can shape the information structure via information disclosure. Specifically, we consider a fundamental multi-dimensional mechanism design problem, where a seller is selling m items to a single unit-demand buyer to maximize her revenue. The buyer's values can be arbitrarily correlated across the items. Our main result shows that, following an appropriately chosen information disclosure scheme, item pricing, i.e., set a take-it-or-leave-it price on each item is highly competitive and guarantees to attain at least 50.1% of the optimal revenue. To our knowledge, this is the first result demonstrating the (approximate) optimality of simple mechanisms in this extensively studied multi-dimensional setting, without making any assumptions about the buyer's value distribution. We believe our result not only demonstrates the power of information disclosure in enhancing the performance of simple mechanisms but also suggests a new framework for reevaluating their efficacy in multi-dimensional settings.

5 - Information Design of the Monty Hall Problem

Zhenyu Hu, National University of Singapore, Singapore, Singapore, Shouqiang Wang, Yangge Xiao

We consider the following variant of the Monty Hall Problem. There are n doors, behind each of which hides an a priori random reward, independently distributed according to a common prior. A clue box is attached to each door, which, if inspected, yields a clue about the rewards behind possibly all doors. The game host designs what clues to be put in each box and announces at the beginning of the game what each clue reveals about the rewards. Given the design, the player decides sequentially which clue box to inspect by paying an inspection fee and when to stop inspecting the boxes. Upon stopping, the player can then utilize the clues collected so far to open a door whose attached box has already been inspected and take the reward behind it. The player wishes to maximize the reward obtained from the opened door, while the host maximizes the fees collected from inspections. Under specific designs, the player’s problem can be reduced to several classic sequential information acquisition problems, such as the Pandora’s box problem or the Wald’s sequential sampling problem. We study the optimal designs of the clue boxes when the host faces different informational constraints. These informational constraints arise from the potential applications in different settings such as product-page design on online platforms, information design in delegated search, and reward design in video games.

Print Close

MA04

Indiana

Trending Topics in BOM

Invited Session

Behavioral Operations Management

Chair: Xiaomeng GUO, Hong Kong Polytechnic University, Kowloon, Hong Kong

1 - Haggling with a Bot: Human vs. LLM Negotiation in Supply Chain Contracts

Rihuan Huang, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Yang Chen

Buying and selling by autonomous AI agents are no longer visions of the future. Large companies like Walmart employ AI agents in contract negotiations. However, AI-human negotiation in the supply chain context is not well understood. In this paper, we experimentally investigate the negotiation dynamics and outcomes of supply chain contract terms between human and large language model (LLM) agents. Specifically, we consider a two-tier supply chain comprising a supplier and a retailer who purchases units from the supplier at a per-unit wholesale price and fulfills random demand. We design an experiment in which participants act as retailers and negotiate the wholesale price and quantity with an LLM-powered supplier. Our experiment follows a 2×2 design, varying two dimensions: 1) which party, retailer or supplier, bears the cost of unsold inventory, and 2) whether the retailer can engage in free-form conversations during negotiations. The first dimension builds on experimental results of human-human negotiation in literature which serves as our benchmark, while the second dimension examines how human-AI negotiation can be facilitated. Our results show that, compared to human-

human negotiations, suppliers represented by LLM agents achieved similar bargaining outcomes when they bear the risk of unsold inventory. In contrast, when retailers take on this risk, the LLM agents' performance depends on the negotiation format. When free-form conversations are enabled, LLM agents perform similarly to humans. However, when free-form conversations are disabled, LLM agents significantly outperform humans by increasing supplier profit through higher wholesale prices, while the channel profit remains unchanged.

2 - Lenient OR Stringent: Managing Returnless Refunds on Retail Platforms

Xiaoli Zhang, Hong Kong Polytechnic University, Kowloon, China, People's Republic of, Xiaomeng GUO, Shi Duo

Many retail platforms have adopted a returnless refund policy, allowing consumers to receive refunds without returning products. By eliminating returns, this policy can enhance consumer satisfaction, reduce retailer costs, and mitigate waste from unnecessary shipments and landfills. However, its generosity also risks exploitation by opportunistic consumers who may engage in dishonest behavior. This paper examines the optimal returnless refund policy and its implications for retailers, platforms, consumers, and social welfare. We develop a game-theoretical model in which a retailer sells its product via a platform, facing both ethical consumers (request refunds only for low-value products) and opportunistic consumers (always request refunds). The request type is unobservable to the retailer and platform, but it can send an imperfect signal through detection. We consider two scenarios where the policymaker for returnless refunds is either the retailer or the platform. The policy pertains to the level of strictness—whether to accept each refund request given the corresponding signal received. Our analysis yields several interesting findings: First, the platform and the retailer do not necessarily align in their preferences for policy strictness. Second, fewer opportunistic consumers or higher detection accuracy may be detrimental to both the retailer and the platform, leading to adjustments in policy strictness towards either direction. Lastly, we find that offering returnless refund can always benefit consumers and sometimes serve as an effective mechanism to reduce total social waste. Moreover, it matters that who should be the decision maker for policy strictness level, yet the answer to this is nontrivial.

3 - Impact of Social Preferences on Demand Information Sharing and CSR Investments in a Triadic Supply Chain

Junlin Chen, Central University of Finance and Economics, Beijing, China, People's Republic of

We study a triadic supply chain with a retailer and two complementary suppliers. The retailer is close to the market and can privately acquire the market segmentation information, i.e. the proportion of prosocial consumers in the market, while the suppliers only know its distribution. The retailer reports the demand information of prosocial customers to the suppliers publicly. Then, suppliers decide independently their own CSR investments to expand prosocial demand in the market under the positive externality of CSR investments. We consider supply chain members' social preferences, including the retailer's lying cost for the possible manipulation of information shared with suppliers, the suppliers' trust preferences towards the retailer's information sharing, and the suppliers' peer-induced fairness concerns caused by the possible free-riding given CSR investments' positive externality characteristics. The results show that: (1) Under full rationality, suppliers' CSR investments are independent of the retailer's report, and the equilibrium is uninformative. However, the existence of social preferences can facilitate the effective sharing of prosocial demand information and enhance suppliers' CSR investments, leading to a win-win-win condition for supply chain members and prosocial customers. (2) The critical factors in achieving this condition are the matching between the retailer's lying cost and the suppliers' trust levels, and the alignment between suppliers' trust levels towards the retailer's report. (3) Extended analysis shows that both supplier cooperation and peer-induced fairness under non-cooperation, although driven by different approaches, can ultimately lead to an increase in individual profits and improve the overall profit of the supply chain.

MA05

Orchard A

Information, Incentives, and Optimal Decisions in Complex Systems

Invited Session

Decision Analysis Society

Chair: Zhi Chen, National University of Singapore, Singapore, Singapore

1 - Managing Multi-Component Projects: Optimal Stopping and Task Scheduling

Sasa Zorc, University of Virginia, Darden School of Business, Charlottesville, VA, United States, Vladimir Smirnov, Ilia Tsetlin

Many complex projects (e.g., the James Webb telescope, or developing a hydrogen powered vehicle) require multiple component projects all to succeed for the complex project to be successfully executed, where there may be either feasibility or cost uncertainties about the components. We develop a Bayesian model of such a project and consider the questions of a) how should the component projects be prioritized, and b) under which condition should the complex project be terminated in its entirety. We show that if there is feasibility uncertainty, it is optimal to focus exclusively on one component until either that component is successfully executed or the lack of success makes the decision maker pessimistic enough about the prospects to terminate the entire project. However, costs uncertainties can make the decision maker switch back and forth between the projects. Which component one should focus on is a complex question, which we don't resolve in its entirety, but we derive a clean index policy for moonshot projects (ones where the reward for successful completion is extremely high).

2 - Pay Transparency in Heterogeneous Teams: Could Some Fairness Concerns Hurt Productivity?

Lin Chen, INSEAD, Fontainebleau, France, Antoine Desir, Guillaume Roels

In homogeneous teams, pay transparency, which induces agents to compare themselves to others, leads to a higher output (and thus, a higher payoff for the principal managing them), results in no inequalities in equilibrium, and is inclusive. This is true irrespective of the basis of social comparisons: income, utility, or reward. How about heterogeneous teams? Could pay transparency backfire under some types of fairness concerns? We consider a model of a principal and two agents of different abilities who are averse to inequalities between each other under pay transparency. We study three bases of comparison: income, utility, and reward. For each type of fairness concern, we characterize when pay transparency leads to a higher or lower payoff for the principal, results or not in inequalities between agents, and involves only one or both agents in production. We find that income fairness never benefits the principal and might be perceived as unfair and non-inclusive, in contrast to reward fairness, which is always beneficial on all three fronts. Utility fairness lies somewhere in between, offering benefits on all three fronts when agents are not so heterogeneous, but hurting fairness and inclusivity without changing the principal's payoff otherwise. Although pay transparency can help close the gender pay gap, it can hurt agents' productivity, create feelings of unfairness, and result in a non-inclusive team if agents are averse to inequalities in income. In contrast, it can be beneficial on all three fronts if agents are averse to inequalities in utility or reward.

3 - Enhancing Forecast Aggregation with Historical Data from Retired Experts

Leping Tu, Nanjing University of Aeronautics and Astronautics , Nanjing, China, People's Republic of, Zhi Chen, Long Zhao

Expert consultations are typically infrequent due to high costs, leading to sparse historical data and the eventual retirement of experienced experts. This paper proposes a method that leverages the historical forecasts of retired experts to improve the aggregation of predictions from active, younger experts. Empirical validation using a dataset of stock earnings forecasts demonstrates the potential of this approach. Our results offer actionable guidance for forecast aggregation in data-scarce environments.

4 - *Optimal Index Policies for Sequential Search: A Unified Framework Under Bayesian Learning*

Zixin (Stephen) Xu, University of Virginia, Darden School of Business, Charlottesville, VA, United States, Manel Baucells, Sasa Zorc

Sequential search models are fundamental tools in operations management, essential for analyzing scenarios such as job hunting, consumer purchasing, and technological adoption. Traditional approaches assume complete knowledge of the underlying distribution, an unrealistic constraint in dynamic decision-making environments. We propose a comprehensive Bayesian learning framework addressing this gap, applicable across a broad distribution family D with unknown parameters.

Our model categorizes distributions within family D into three distinct subfamilies based on parameter uncertainty: unknown location, unknown scale, and unknown location-scale parameters. Leveraging conjugate priors, we derive closed-form posterior predictive updates, significantly simplifying the dynamic programming recursion through standardized transformations. This reduction yields a universal optimal index policy characterized by structurally intuitive thresholds.

Furthermore, our approach not only disproves Rothschild's (1978) conjecture regarding single-reservation-price strategies in exponential family searches without recall, but also establish the sufficient and necessary conditions for the reservation price property. By incorporating Constant Absolute Risk Aversion (CARA) and Constant Relative Risk Aversion (CRRA) utilities, along with exponential and hyperbolic discounting models, our framework extends naturally to real-world contexts involving risk aversion and behavioral considerations.

Ultimately, this unified methodology provides analytical clarity, computational efficiency, and practical versatility. It accommodates diverse decision-maker preferences, recall scenarios, and time-inconsistent discounting, thereby offering robust insights into optimal sequential search strategies under uncertainty.

5 - Managing Cybersecurity: Data Access & Protection

Marat Salikhov, New Economic School, Moscow, Russian Federation, George Charlson, Ruslan Momot, Oleh Stupak

A standard data access guideline is to limit employee access to only the essential data for their roles, raising the question of how to define these roles and the employees' corresponding data needs. We aim to address this question by considering a game-theoretic model of joint cybersecurity and operational decision-making. The firm chooses the level of data access for each of its employees and the overall level of cyber protection. Providing more data access to employees makes the firm more economically efficient but also more vulnerable to attacks by an adversary who steals the employees' data, inflicting damage that increases with the amount of data stolen. Adversaries may vary in their attack strength and sophistication rate (the ability to pinpoint the most attractive targets). We find that the firm may counter-intuitively decrease its overall protection level when adversarial attacks become stronger and increase its overall access level when the adversaries become more sophisticated.

MA06

Orchard B

Data-Driven Finance: ML, Bitcoin, and Policy Insights in Market Analysis

Invited Session

Finance

Chair: Linlin Niu, A306, Econ Building, Xiamen University, 361005

1 - How Do Trade Conflicts Affect Corporate Carbon Emissions? Evidence from Global Supply Chains

Zeyun Bei, City University of Hong Kong, Hong Kong, China, People's Republic of, Ebenezer Effah, Yaxuan Qi

This paper examines the effect of trade conflicts on corporate carbon emissions through global supply chains. Exploiting the 2018–2019 U.S.-China tariff hikes as a quasi-natural experiment, we find that U.S. firms subjected to higher import tariffs increased Scope 1 and 2 emissions by 15% and 8%, respectively, compared to minimally affected peers. The effect is amplified for firms reliant on green product imports from China, highlighting supply chain disruptions as a key channel. Financial constraints further exacerbate emissions, as tariff-induced cost pressures reduce green innovation and employment. Cross-sectional tests reveal that firms with weaker climate change ideologies and less diversified supply chains exhibit more pronounced emission increases. Overall, our findings demonstrate that geopolitical trade risks undermine firm-level decarbonization efforts, with implications for policymakers and firms navigating climate goals amid rising protectionism.

2 - Predicting China's Treasury Yield Curve with Machine Learning and Shap Interpretation: Context-Specific Nonlinearities in Key Determinants

Zhiwu Hong, China University of Political Science and Law, Beijing, China, People's Republic of, Mucai Lin, Xinda Li

As the pricing benchmark of financial markets and a barometer of macroeconomic trends, accurate prediction of treasury bond yields has long been a critical concern for policymakers and investors globally. This study establishes a novel framework that integrates machine learning models with SHapley Additive exPlanations (SHAP) framework to predict China's treasury yield curve, utilizing multidimensional indicators spanning monetary policy, macroeconomic fundamentals, capital liquidity, market expectations, and international factors, complemented by principal component analysis of yield characteristics. The empirical results demonstrate three key findings: First, the Extra Trees model significantly outperforms both traditional econometric models and alternative machine learning approaches in predictive accuracy. Second, six critical determinants emerge through SHAP analysis: 6-Month SHIBOR, the exchange rate, the RMB Reserve requirement ratio for large deposit, the housing sales price index, the completed fixed assets investment, and the first principal component of yield curve. Third, the study reveals nonlinear relationships where both the magnitude and direction of factor impacts exhibit substantial heterogeneity across different forecasting horizons and bond maturities. These findings provide actionable insights for monetary authorities in policy calibration and offer institutional investors a data-driven framework for yield curve positioning strategies.

3 - Identifying the Forward Guidance Signals in China's Monetary Policy Statements Through Ordered Heteroskedasticity

Mucai Lin, Institution of Quantitive Economics and Statistics, Huaqiao University, Xiamen, China, People's Republic of

China's monetary policy has long utilized dual instruments (quantitative and price-based) targeting complementary objectives. When the People's Bank of China (PBC) adjusts specific policy tools, markets not only respond to explicit headline changes but also implicitly adjust expectations for the complementary target through the central bank's quantity-price coordination mechanism, generating non-headline news shocks. This study develops an ordered heteroskedasticity identification method to disentangle these dual shocks in policy announcements, offering enhanced robustness and reduced data requirements compared to conventional multidimensional approaches.

Empirical analysis demonstrates the dual-shock framework's superior explanatory power for Treasury yield movements during event windows, particularly revealing pronounced non-headline shocks in quantitative policy announcements. These shocks reflect significant market expectation adjustments toward implied price targets and demonstrate predictive capacity for future PBC policy shifts. Moreover, non-headline news shocks exhibit statistically significant forecasting power for subsequent movements in Treasury yields and Shanghai Interbank Offered Rate (Shibor), suggesting their critical role in policy transmission mechanisms.

4 - Treasury Supply Shocks and State-Dependent Market Reactions: High-Frequency Evidence from China's Fiscal Policy Transmission

Haoran Bai, Xiamen University, Xiamen, China, People's Republic of, Mucai Lin, Linlin Niu

This paper examines China’s fiscal policy in the post-COVID era, focusing on Treasury bond issuance as the core instrument and evaluating its short-run effectiveness. Using high-frequency identification, we interpret price variations in Treasury futures around Treasury announcements as shocks to expected bond supply. Empirical analysis shows that during periods of economic stability, these shocks have an insignificant impact on the Treasury yield curve. However, during economic downturns, the shocks cause the yield curve to shift upward and flatten, with rising expected yields outweighing the decline in term premia. Mechanism analysis suggests that markets perceive these shocks as signals of proactive countercyclical fiscal intervention, fostering optimism about future economic fundamentals and prompting expectations of gradual policy rate normalization. Simultaneously, government bonds are reassessed as safe assets, leading to a compression in term premia.

5 - Beyond Conventional Sentiment Indicators: Bitcoin&Rsquo;S Hidden Potential in Vix Forecasting

Juan Lin, Xiamen University, Xiamen, China, People's Republic of, Ming Gu, Siyi Liu

This study introduces a novel decomposition framework aligned with U.S. stock market trading hours, identifying Bitcoin’s overnight returns as a powerful indicator of investor sentiment. By incorporating Bitcoin’s overnight returns into standard forecasting models, we observe a significant improvement in the predictability of

equity market volatility dynamics across both in-sample and out-of-sample analyses. Furthermore, long-short trading strategies based on these enhanced forecasts consistently outperform conventional approaches, generating substantial economic gains. These findings suggest that existing volatility models could benefit from the inclusion of cryptocurrency market signals. Overall, our results highlight the growing integration of cryptocurrency and traditional financial markets.

MA07

Orchard C

Risk Adjustment, Policy, and Data-Driven Innovation in Healthcare

Invited Session

Health Applications Society

Chair: Zhaowei She, Lee Kong Chian School of Business, Singapore Management University, Singapore, Singapore

1 - Design of Medical Reimbursement Policy and Effects of Pooling

Yiwen Shen, HKUST, Hong Kong, Hong Kong, Qingfa Zhang, Ying-Ju Chen

We study the design of medical reimbursement policy using a model with multiple types of agents and services. The central planner decides the reimbursement policy for agents' medical costs to minimize the total utility loss, subject to a budget constraint. The optimal reimbursement policy has a cap structure, in which all excess costs above a threshold are reimbursed. We also solve the optimal structure for the commonly-used ratio reimbursement policy and show that it can be constructed based on a cost index of services. Then, we investigate the effects of horizontal pooling when the central planner manages multiple groups. We show that with the optimal cap policy or under homogeneous cost distributions, it is optimal to integrate the groups together and use a common reimbursement policy for all, suggesting that policy flexibility is not needed. Finally, we consider a dynamic model in which the central planner needs to balance between current and future utilities. In the optimal dynamic policy, the central planner increases the marginal spending rate for current period when the reimbursement scope is enlarged, leading to a piecewise concave structure of optimal spending amount. An extensive numerical study shows the value of horizontal pooling and dynamic policy in medical reimbursement. Our study provides practical insights for the consolidation of medical insurance schemes, which has been an important component of healthcare reform in many regions.

2 - Managing Prediction Errors in Healthcare Capitation: a Two-Stage Risk Adjustment Approach

Zhaowei She, Lee Kong Chian School of Business, Singapore Management University, Singapore, Singapore, Kanix Wang, Yajing Chen

This work addresses a key challenge in capitation-based healthcare payment: prediction errors in risk adjustment that distort incentives and violate budget neutrality across patient risk groups. Motivated by a recent reform proposed by the Centers for Medicare & Medicaid Services (CMS), we develop a novel two-stage algorithm to estimate the HHS-HCC risk adjustment model. Our approach explicitly targets and corrects systematic prediction errors that disproportionately affect specific patient segments, particularly those at the extensive margin. Leveraging administrative claims data, we empirically validate our method and show that it substantially improves predictive accuracy and achieves near-perfect actuarial equivalence across risk classes.

3 - Enhanced Risk Assessment for Autism with Machine Learning and Supersparse Linear Integer Modeling

Yu-Hsin Chen, National Tsing Hua University, Hsinchu, Taiwan, Qiushi Chen

Early diagnosis is crucial for improving treatment outcomes in Autism Spectrum Disorder (ASD), a developmental disorder affecting 1 in 36 children in the US. The Modified Checklist for Autism in Toddlers (M-CHAT) is the most commonly used tool for ASD-specific screening during pediatric well-child visits. However, its effectiveness is limited due to concerns about accuracy. Although previous research has shown that machine learning models using electronic health record (EHR) data can improve ASD risk prediction, these models are challenging to integrate into primary care settings to assist decision-making. This is due to their complex

computations, which make them challenging to interpret and implement, and their lack of direct incorporation of M-CHAT results.

To address these limitations, we first applied machine learning models to EHR data to assess the added predictive value of various variables alongside M-CHAT. We then utilized a Supersparse Linear Integer Model (SLIM) to develop a score-based risk assessment system for straightforward computation and easy implementation. The enhanced risk assessment model demonstrated significant improvements, increasing the area under the receiver operating characteristic (ROC) curve by over 10% with the machine learning model and by 7% with the score-based model compared to M-CHAT alone. This new model complements the existing M-CHAT by simplifying the calculation process and enhancing the overall accuracy of ASD screening.

4 - Gossip Cue: the Impact of Community-Level Rabies Exposure on Influenza Vaccination Among Older Adults

Yang Yue, Xiamen University, Xiamen, China, People's Republic of, Yuyu Chen, Lijing Yan, Meilan Zhang, Xian Zhang

This paper examines how community-level rabies exposure incidents influence influenza vaccination uptake among older adults (aged 60+), analyzing both individual and community-level effects. Using individual-level electronic medical records from 401 communities in X district, China, between 2012 and 2019, we find that whether a rabies exposure incident occurs in a community increases influenza vaccination among older adults in the following month by 28.8\%. At a weekly level, this effect exhibits dynamic patterns over time. Specifically, after a community rabies exposure event, influenza vaccination rates among older adults increase significantly by 13.6\% in the third week, 16.8\% in the fifth week, and 17.7\% in the sixth week. We argue that community-level information diffusion and peer effects are key mechanisms underlying these response patterns. Heterogeneity analysis reveals that the impact of rabies exposure varies with individual socioeconomic characteristics and seasonal factors. Males exhibit stronger responses to rabies exposure compared to females. Higher-educated individuals exhibit earlier responses compared to others. Furthermore, leveraging a dog-restriction policy implemented in certain neighborhoods in 2018 as an exogenous shock, we employ a difference-in-differences (DID) approach to explore variations in rabies exposure and influenza vaccination between treated and control groups, providing additional support for our baseline results. Overall, our findings underscore the crucial role of community-level cues in influencing individual health behaviors, offering actionable insights for targeted public health interventions.

5 - Optimizing Healthcare Utilization Using Conditional Value-at-Risk with Applications in Operating Theatre Planning

Fanwen Meng, National Healthcare Group, Singapore, Singapore, Kiok Liang Teow, Palvannan Kannapiran, John Arputhan Abisheganaden

This study examines healthcare utilization targets across hospital departments with similar resource needs. Given the uncertainty and variability in service times, along with inherent departmental differences, we propose a tailored approach utilizing a quantitative framework based on Conditional Value-at-Risk (CVaR) to evaluate service overruns with a defined confidence level. We then establish optimal service planning for each department and develop specific utilization targets. Our framework includes practical numerical computation methods, particularly for scenarios where service durations follow normal or lognormal distributions. We derive analytical relationships between the mean and standard deviation of service durations and utilization rates under the assumption of normal distributions. The framework is applied to operating theatre planning in a public hospital in Singapore, offering decision-makers insights into the trade-offs involved in setting utilization targets across various departments. Our findings suggest that departments with higher variability in surgery durations may need to operate at lower utilization levels while experiencing a greater likelihood of service overruns.

MA08

Bras Basah A

Recent Advancements in Green Transportation

Invited Session

Transportation Electrification

Chair: Min Xu, The Hong Kong Polytechnic University, Hong Kong, Hong Kong

1 - Revealing Charging Patterns of Electric Vehicles with Intrinsic Usage Heterogeneity

Ruinan Wang, New York University Shanghai, Shanghai, China, People's Republic of, Zhengtian Xu, Zhibin Chen

With the increasing adoption of electric vehicles (EVs), location-based data have emerged as valuable resources for understanding EV usage and charging patterns. Compared to survey data, these datasets provide more comprehensive historical records and offer reliable insights into user preferences. However, due to data collection limitations and privacy protection measures, pinpointing specific user groups remains challenging. As a result, location-based data, often sampled from a general population, exhibit intrinsic and hidden heterogeneity among individuals, making aggregate behavioral analyses less effective. To better support facility planning and policy design, it is essential to cluster individuals based on their EV usage characteristics and differentiate heterogeneous user groups. Addressing this need, our study introduces a two-stage method for identifying trip purposes and classifying EVs based on their usage types, incorporating both vehicle-level and trip-level attributes. Using real-world driving and charging data from a sample of EVs in Shanghai, we categorize vehicles into five usage types: two types of private vehicles and three types of commercial vehicles. We further analyze charging demand variations and user preferences across these categories. This high-granularity classification provides actionable insights for optimizing the deployment of charging infrastructure, while also offering precise data to support power system management and policy development.

2 - An Exact Algorithm for Fleet Co-Deployment and Slot Co-Chartering in A Sustainable Shipping Alliance Under Emission Trading System

Yadong Wang, Nanjing University of Science and Technology, Nanjing, China, People's Republic of, Shenghui Zhu, Cagatay Iris

Shipping alliances have emerged as a cooperation platform between independent shipping companies, aiming to enhance customer satisfaction and exploit the economies of scale through capacity and information sharing. A sustainable shipping alliance should operate in a profitable, fair and environmentally friendly way under emerging Emission Trading Systems (ETS). A non-convex mixed-integer nonlinear programming model is suggested to jointly optimize the fleet co-deployment in the network, sailing speed in each shipping leg, schedule design for each shipping service, and the slot allocation and co-chartering for each alliance member. These decisions ultimately determine each company’s carbon emissions. Under the ETS, companies are charged for emissions that exceed their allowances, while any surplus allowances can be traded for revenue in carbon markets. In addition to maximizing the alliance’s total profit, this study minimizes profit margin variation among members in proportion to their investment, promoting fairness in a novel way. A tailored spatial branch-and-bound (SB&B) algorithm is developed to deliver the global optimal solution for the problem. Novel problem relaxation and branching strategies are suggested based on the structure of the programming model. Results indicate that the SB&B significantly outperforms an existing non-convex nonlinear solver. Compared to case which do not consider slot co-chartering and fairness, our study delivers improved total profit by 3.13%, meets 0.52% more freight demand, and ensures a fairer profit distribution on average. Under ETS, carbon emissions is reduced by 54.3%, with smaller ships being used and average sailing speeds decreasing as emission trading price rises from $0/tonne to $300/tonne.

3 - Electric Vehicle Routing for Vehicle to Vehicle En-Route Charging Services

Min Xu, The Hong Kong Polytechnic University, Hong Kong, Hong Kong

The emerging vehicle-to-vehicle en-route charging (V2V-ERC) service is a promising complementary solution to the conventional charging services at fixed locations. This study explores the V2V-ERC service routing problem that determines the routing and charging strategies for electricity-providing vehicles (EPs) to provide V2V-ERC services to electricity-demanding vehicles (EDs). Each ED could be served collectively by different EPs along their pre-specified traveling routes. Battery degradation costs are incurred for EPs due to frequent charging and discharging cycles. We formulate this problem as a mixed-integer nonlinear programming model and propose an efficient algorithm integrating the adaptive large neighborhood search method and model linearization techniques. Numerical experiments are conducted to evaluate the efficacy of the proposed algorithm and derive managerial insights.

4 - An Exact Algorithm for Passenger Pickup and Delivery Problem with Modular Vehicles

Fangni Zhang, The University of Hong Kong, Hong Kong, Hong Kong, Yizhi Zhen

Modular vehicles (MVs) possess the flexibility to decouple and couple modular units en route to form platoons, offering advantages in energy efficiency and cost efficiency. Routing a homogeneous fleet of MVs for the pickup and delivery problem (PDP) for passengers, while considering the decoupling and coupling operations between MVs, presents a significant challenge. This is because the resource consumption of a MV on an arc is not only linearly related to the distance but also depends on the number of vehicles that choose to pass through the arc (i.e., the platoon size). To address the modular vehicle routing problem (MVRP) for PDP with passenger service time windows, we first formulate an arc-flow model to represent MV movement. We acknowledge that passenger demand appears over time, thus, a rolling horizon framework is introduced to accommodate the dynamic nature of passenger demands. To solve the static MVRP for each horizon, we propose a novel exact method based on a branch-price-and-cut algorithm. Unlike existing branch-price-and-cut algorithms for traditional vehicle routing problems, the pricing subproblem of the MVRP cannot be solved with an exact route cost for each MV individually, as the arc cost is variable and may be affected by the pricing subproblems of other MVs. Therefore, we design a new labeling algorithm with tailored dominance rules for the unique pricing subproblems. We anticipate that the new exact algorithm will be capable of solving the MVRP for small-to-medium-sized datasets within an acceptable computation time for each horizon.

MA09

Bras Basah B

Navigating Autonomy: AI for Ship Operations and Navigation

Invited Session

Maritime AI

Chair: Jingxian Liu, Wuhan University of Technology, Wuhan, N/A

Co-Chair: Yi Liu, Wuhan University of Technology, Wuhan, China, People's Republic of

1 - Applications of Maritime AI Large-Scale Models in Shipping Industry

Jingxian Liu, Wuhan University of Technology, Wuhan, China, People's Republic of

The Maritime AI Large-Scale Models are comprehensive suite of intelligent models designed to support diverse applications in the shipping industry. Rather than a single model, this library integrates multiple specialized modules to address key challenges in maritime operations. These include models for intelligent vessel scheduling, non-boarding pilotage under low-visibility conditions, ship-bridge collision prevention, waterway through-capacity simulation, and under-keel clearance (UKC) prediction. Each model leverages advanced AI techniques—such as Transformer-based forecasting, mixed-integer programming, and real-time data fusion—to enhance safety, efficiency, and operational decision-making. For example, the intelligent scheduling model optimizes Ro-Ro transportation by predicting passenger demand and assigning vessels and berths efficiently. The non-boarding pilotage model uses escort boats and environmental perception systems to ensure safe navigation during fog or other adverse conditions. Meanwhile, the collision prevention system offers early warnings and real-time intervention to reduce bridge-area accidents. The UKC model dynamically calculates safe clearance margins for large vessels by integrating ship behavior, sea conditions, and real-time environmental inputs. Additionally, the waterway throughput model simulates and evaluates traffic capacity under different operational scenarios to inform port planning. By combining domain expertise with intelligent algorithms, the Maritime AI Large Model Library establishes a scalable and modular foundation for smart shipping. It offers port authorities, shipping companies, and regulatory bodies robust tools to support digital transformation and sustainable development in maritime logistics.

2 - Towards Explainable Autonomous Ship Navigation: Evidential Reasoning-Based Ship Trajectory Prediction

Bingxin Liu, Dalian Maritime University, Dalian, China, People's Republic of, Xinyu Zhang

In dynamic maritime environments, accurate ship trajectory prediction is crucial for proactive collision avoidance and safe autonomous vessel navigation. However, trajectory evolution is highly uncertain, challenging traditional model-driven approaches. Existing data-driven methods mainly rely on historical features and often neglect causal relationships, limiting interpretability in autonomous systems. To address this, we propose an evidential reasoning-based trajectory prediction model that combines historical AIS data and typical navigation scenarios to extract ship behavior patterns. These patterns are used to initialize belief rule parameters and construct an initial rule base, modeling the causal relationships between ship behavior and environmental factors. The evidential reasoning method integrates rules to infer the most suitable prediction for the current scenario. To enhance adaptability, a P-CMA-ES-based dynamic optimization mechanism is introduced, allowing the model to update continuously with perception data. To preserve interpretability, human-driver-inspired interpretability constraints are incorporated into the optimization process, ensuring trustworthy and accurate predictions. The model was validated using AIS data in complex scenarios, including cross-channel navigation, emergency collision avoidance, and narrow waterways. Its performance was evaluated against LSTM, GRU, and trajectory clustering methods using RMSE. Results show that our model outperforms others in both accuracy and interpretability. Visualized belief distributions reveal a traceable prediction process, supporting interpretability. Furthermore, comparisons with human driver behavior quantify their consistency, highlighting the model's reasoning capability. These findings demonstrate the model's potential for reliable and interpretable autonomous trajectory prediction.

3 - Enabling Maritime Autonomous Navigation: Sensor Systems and Environmental Challenges

Hyesu Jang, Seadronix, Seoul, Korea, Republic of

The maritime domain represents one of the challenging frontiers of exploration, characterized by its vastness, dynamic conditions, and limited accessibility. Mankind has long pursued technologies to unveil the ocean, and recent advancements in autonomy present new opportunities for maritime exploration. Autonomous navigation systems have shown promising potential; however, their adaptation to maritime environments introduces unique challenges. The sparsity of navigational landmarks, complex water-land boundaries, and the expansive nature of open waters demand sensing solutions with long-range and robust detection capabilities.
This presentation examines the role of sensing as a critical enabler for autonomous maritime navigation. We provide a comprehensive overview of radar technologies across diverse frequency bands, analyze their effectiveness in addressing the complexities of marine perception, and showcase approaches for processing maritime sensor data. The discussion concludes with the future of autonomous navigation in maritime environments, framed through the sight of robotics and intelligent systems research.

4 - A Causal Learning Method for Multi-Ship Interaction Representation Learning and Reasoning

Yanfu Wang, China university of Petroleum (East China), Qingdao, China, People's Republic of, Miaomiao Wang

Accurate understanding of the multi-ship interaction mode is crucial for enhancing situation awareness and intelligent decision support systems. A causal enhanced dynamic representation learning and reasoning method for multi-ship interaction is proposed to achieve interpretable modeling and intent reasoning of ship behavior during the interaction process. The multi-ship encounter scenarios are extracted using the spatio-temporal constraint method, and the dynamic spatio-temporal graphs are constructed to simulate the interaction. The self-supervision task of masking and reconstruction is adopted to identify the causal factors influencing the behavioral decisions of the target ships. Individual behavioral patterns and causal interaction characteristics are integrated to achieve interpretable reasoning about the future behavioral intentions of ships. Behavioral representation is enhanced through ship motion pattern, and the causal attention gating mechanism is used to enhance the attention and modeling of causal features. The research on AIS data from Ningbo Zhoushan Port shows that the proposed method can effectively identify multi-ship interaction patterns with causal characteristics, significantly improve the model's ability to distinguish key interaction intentions, and demonstrate strong robustness and practical interpretability. This method provides theoretical support and technical paths for applications such as ship behavior understanding and intelligent decision-making systems.

MA10

Hullet Room

Trustworthy Optimization and Learning for Electric Grids

Invited Session

Power Grid

Chair: Yue Chen, The Chinese University of Hong Kong, Shatin, NT, Hong Kong, N/A

Co-Chair: Yunhe Hou, The University of Hong Kong, Room 508, CYC Building, , Hong Kong, N/A, Hong Kong

1 - Trustworthy Control for Sustainable Power Grids

Yifei Wu, The Chinese University of Hong Kong, Shenzhen, China, People's Republic of, Jing Yu, Tongxin Li

Predictions are crucial for managing networked systems, especially in dynamic environments like power grids with renewable energy. However, unreliable predictions can destabilize control strategies, necessitating new solutions. This talk reviews existing methods and introduces a novel distributed control strategy where local controllers generate predictions and share information with neighboring nodes. This approach enhances scalability and practicality, overcoming the limitations of centralized systems. A critical discovery is a non-monotonic dependency between communication distance in topology design and the performance gap from centralized control: performance improves with distance up to an optimal point, then declines. This insight offers a new aspect for designing power grid communication networks, balancing local autonomy with global efficiency when prediction error presents. The talk will also explore practical applications and future developments, demonstrating how this method improves reliability and efficiency while managing uncertainties from renewable energy sources.

2 - A Differentially Private Energy Trading Mechanism Approaching Social Optimum

Yuji Cao, The Chinese University of Hong Kong, Shatin, NT, Hong Kong, China, People's Republic of

This paper proposes a differentially private energy trading mechanism for prosumers in peer-to-peer (P2P) markets, offering provable privacy guarantees while approaching the Nash equilibrium with nearly socially optimal efficiency. We first model the P2P energy trading as a (generalized) Nash game and prove the vulnerability of traditional distributed algorithms to privacy attacks through an adversarial inference model. To address this challenge, we develop a privacy-preserving Nash equilibrium seeking algorithm incorporating carefully calibrated Laplacian noise. We prove that the proposed algorithm achieves ε-differential privacy while converging in expectation to the Nash equilibrium with a suitable stepsize. Numerical experiments are conducted to evaluate the algorithm's robustness against privacy attacks, convergence behavior, and optimality compared to the non-private solution. Results demonstrate that our mechanism effectively protects prosumers' sensitive information while maintaining near-optimal market outcomes, offering a practical approach for privacy-preserving coordination in P2P markets.

3 - Resilience-Oriented Planning for Microgrid Clusters Considering P2p Energy Trading and Extreme Events

Zheng Xu, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, Yue Chen, Linwei Sang, Haifeng Qiu, Zhi Wu, Hengqing Ye

Multiple microgrids (MGs) can work more collaboratively and economically thanks to peer-to-peer (P2P) energy transactions. Nevertheless, the more frequent contingencies pose threats to the energy trading of multi-MGs, underscoring the need to strike a balance between economic efficiency and resilience. To safeguard the co-operational security of MG clusters, this paper proposes a resilience-oriented planning method considering extreme events and their destructive impacts on multi-agent energy trading. Firstly, a multi-agent resilient planning model is established, aiming at reducing the planning cost through collaborative energy sharing while ensuring the existence of alternative solutions under contingencies. Both the internal uncertainties of renewable energy sources and load variations and the external stochastic transaction disruptions are considered. To ensure the convexity and continuity of the model, it is proven that the complementary constraints of energy storage in both the normal and contingency conditions can be relaxed without introducing binary variables. Then, an incentive-driven P2P pricing strategy is devised by working on the dual of the social cost minimization problem.

The equality and inequality coupling constraints among multi-stakeholders are decoupled in a distributed manner with limited information exchange. The convergence and optimality of the proposed distributed algorithm are supported with theoretical guarantees, and then it is solved efficiently by incorporating the scenario-decoupled column and constraint generation approach. Numerical experiments on island MG systems validate the efficacy and superiority of the proposed planning method, which is of great significance in improving multi-agent grid profitability and resilience.

4 - Data-Driven Robust Unit Commitment with Statistical Guarantees: An Integrated Forecasting and Optimization Approach

Yue Chen, The Chinese University of Hong Kong, Hong Kong, China, People's Republic of, Rui Xie

The growing uncertainty from renewable power and electricity demand brings significant challenges to unit commitment (UC). While various advanced forecasting and optimization methods have been developed to predict better and address this uncertainty, most previous studies treat forecasting and optimization as separate tasks. This separation can lead to suboptimal results due to misalignment between the objectives of the two tasks. To overcome this challenge, we propose a robust UC framework that integrates the forecasting and optimization processes while ensuring statistical guarantees. In the forecasting stage, we combine multiple predictions derived from diverse data sources and methodologies for an improved prediction, aiming to optimize the UC performance. In the optimization stage, the combined prediction is used to construct an uncertainty set with statistical guarantees, based on which the robust UC model is formulated. The optimal robust UC solution provides feedback to refine the forecasting process, forming a closed loop. To solve the proposed integrated forecasting-optimization framework efficiently and effectively, we develop a neural network-based surrogate model for acceleration and introduce a reshaping method for the uncertainty set based on the optimization result to reduce conservativeness. Case studies on modified IEEE 30-bus and 118-bus systems demonstrate the advantages of the proposed approach.

MA11

Morrison Room

Algorithms to Revenue Management

Invited Session

Revenue Management and Pricing

Chair: Huan Zheng, Shanghai Jiao Tong University, Shanghai, N/A

Co-Chair: Jiannan Ke, Wuhan University, Wuhan, N/A

1 - Mostly Beneficial Clustering: Aggregating Data for Decision Making

Chengzhang Li, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Zhenkang Peng, Ying Rong

With increasingly volatile market conditions and rapid product innovations, operational decision-making for large-scale systems entails solving thousands of problems with limited data. Data aggregation is proposed to combine the data across problems to improve the decisions obtained by solving those problems individually. We propose a novel cluster-based Shrunken-SAA approach that can exploit the cluster structure among problems when implementing the data aggregation approaches. We prove that, as the number of problems grows, leveraging the given cluster structure among problems yields additional benefits over the data aggregation approaches that neglect such structure. When the cluster structure is unknown, we show that unveiling the cluster structure, even at the cost of a few data points, can be beneficial, especially when the distance between clusters of problems is substantial. Our proposed approach can be extended to general cost functions under mild conditions. When the number of problems gets large, the optimality gap of our proposed approach decreases exponentially in the distance between the clusters. We explore the performance of the proposed approach through the application of managing newsvendor systems via numerical experiments. We investigate the impacts of distance metrics between problem instances on the performance of the cluster-based Shrunken-SAA approach with synthetic data. We further validate our proposed approach with real data and highlight the advantages of cluster-based data aggregation, especially in the small-data large-scale regime, compared to the existing approaches.

2 - Dynamic Programs with Shared Resources and Real-TIME Agent Flows

Jingwei Zhang, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Ziyan Zhou

This paper examines a decision-making framework where a decision maker (DM) allocates resources across sequentially arriving agents, subject to shared resource constraints. Such problems are prevalent in various domains including hospital care scheduling and cloud computing. In such problems, the DM must dynamically adjust resource allocations in response to new arrivals, ensuring efficient coordination across diverse agent types. Such problems are naturally formulated as stochastic dynamic programs (DPs) but solving the DP is impractical unless the number of agents is small. In this paper, we develop a dynamic fluid relaxation --- that provide upper bounds on the optimal value function as well as a feasible policy. We develop an efficient way for solving the dynamic fluid relaxation and analyze the performance of the feasible dynamic fluid policy. Our performance analysis implies, under mild conditions, that the dynamic fluid relaxation bound and feasible policy are asymptotically optimal as the amount of resources grows large. We demonstrate the model and results in two applications.

3 - Reformulations of Approximate Linear Programs for Finite-Horizon Markov Decision Processes

Jiannan Ke, Wuhan University, Wuhan, China, People's Republic of, Dan Zhang, Huan Zheng

We establish reformulations of approximate linear programs (ALPs) for a class of finite-horizon Markov decision processes (MDPs). These reformulations are either equivalent formulations or relaxations of the original ALPs. Our main result is established via constraint reformulations and duality arguments. In particular, it allows us to find equivalent reformulations of ALPs by checking a strong duality condition. We apply the framework to network revenue management and the finite-horizon restless bandit problem. For network revenue management with and without customer choice, we provide new, simplified proofs for the main results in Vossen and Zhang (2015) under the usual assumption of a binary consumption matrix. When the consumption matrix is non-binary, we establish new upper bounds under both affine and SPL approximations. For the finite-horizon restless bandit problem, we establish an equivalent reformulation under the assumption of a unit cost for pulling each arm and a relaxed upper bound when the cost is more general. A numerical study shows that the quality of the bounds and heuristic policies of the reformulations are competitive with alternative solution methods but at substantially reduced computational cost. Our work applies to a broad class of finite-horizon MDPs and has potential applications to many practical problems.

MA12

Moor Room

Learning and Optimization for Decision Making

Invited Session

Revenue Management and Pricing

Chair: Jingren Liu, National University of Singapore, Singapore, Singapore

1 - Context-Based Dynamic Pricing with Separable Demand Models

Jinzhi Bu, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, David Simchi-Levi, Chonghuan Wang

Motivated by the empirical evidence observed from the real-world dataset, we consider a context-based dynamic pricing problem with separable demand models. The demand function is endowed with a separable structure in the form of f(p)+g(x), where p and x denote the price and feature vector respectively. Under different configurations of f(p) and g(x), we systematically characterize the statistical complexity of the online learning problem. Specifically, we consider three models: (i) f(p) is linear and g(x) is non-parametric; (ii) f(p) is non-parametric and g(x) is linear; and (iii) f(p) and g(x) are both non-parametric. For each model, we design an efficient learning algorithm with a provable regret upper bound and establish an almost matching regret lower bound.

2 - Satisficing Regret Minimization in Bandits

Qing Feng, Cornell University, New York, NY, United States, Tianyi Ma, Ruihao Zhu

Motivated by the concept of satisficing in decision-making, we consider the problem of satisficing exploration in bandit optimization. In this setting, the learner aims at finding a satisficing arm whose mean reward exceeds a certain threshold. The performance is measured by satisficing regret, which is the cumulative deficit of the chosen arm's mean reward compared to the threshold. We propose SELECT, a general algorithmic template for Satisficing REgret Minimization via SampLing and LowEr Confidence bound Testing, that attains constant satisficing regret for a wide variety of bandit optimization problems in the realizable case (i.e., whenever a satisficing arm exists). Specifically, given a class of bandit optimization problems and a corresponding learning oracle with sub-linear (standard) regret upper bound, SELECT iteratively makes use of the oracle to identify a potential satisficing arm. Then, it collects data samples from this arm, and continuously compares the lower confidence bound of the identified arm's mean reward against the threshold value to determine if it is a satisficing arm. As a complement, SELECT also enjoys the same (standard) regret guarantee as the oracle in the non-realizable case. Finally, we conduct numerical experiments to validate the performance of SELECT for several popular bandit optimization settings.

3 - Approximate Linear Programs for Network Revenue Management: Reduction via Constraint Decomposition and Adaptation for Locally Separable Basis Functions

Chung-seung Lee, SUNY Korea, Incheon, Korea, Republic of, Metin Cakanyildirim, Xiao Zhang

We develop compact approximate linear programs (ALPs) for the network revenue management problem with independent demands. Our approach leverages an intuitive constraint decomposition method, enabling us to derive the compact ALPs in a versatile, primal form. We also propose a *local multivariate* approximation scheme that provides fine-grained control over the trade-off between approximation accuracy and computational tractability. The employed basis functions are multivariate functions that are additively separable in their respective subdomains. To validate the advantage of the proposed method, we conduct an extensive numerical study, demonstrating that the larger the non-separable regions, the tighter the upper bound and the larger the expected revenue, but the longer the solution time.

4 - Shape-Restricted AI-Powered Pricing

Jingren Liu, National University of Singapore, Singapore, Singapore, Hanzhang Qin, Mabel C. Chou

We study a feature-based pricing problem within the predict-then-optimize paradigm. We propose a demand predictor that combines double machine learning (for estimating the price sensitivity) with monotonic neural additive models (MNAM) (for predicting the effect of features) to guide pricing decisions. The incorporation of the monotonicity constraint meets interpretability and fairness requirements, which are important in applications such as personalized lending rate optimization. In terms of statistical properties, we give the first nonasymptotic result on the generalization error of MNAM. For the decision performance, our proposed framework utilizes Neyman orthogonality to lessen the impact of the prediction error on price optimization performance, so that the revenue loss provably converges to zero as the sample size increases, at a rate faster than that of the prediction error.

MA13

Bailey

Operations Management in Healthcare Systems

Invited Session

Simulation and Optimization of Healthcare Systems

Chair: Gang Du, East China Normal University, Shanghai, N/A

1 - A Blessing OR Curse: Teletriage in Healthcare

Xu Guan, Huazhong University of Science and Technology, Wuhan, China, People's Republic of, Krista Li

Teletriage is a telehealth service that uses telecommunication technologies to remotely assess patients' conditions and determine their needs for medical treatment. This paper examines how teletriage affects patients' decisions to seek treatment and when healthcare providers (e.g., medical businesses, hospitals, or clinics) should launch

teletriage services. Our findings reveal that teletriage can encourage more mild patients to seek treatment while discouraging more severe patients. This occurs when teletriage lacks sufficient precision and patients' hassle cost of seeking treatment falls outside a moderate range. Specifically, when hassle costs are low, patients are generally inclined to seek treatment, but teletriage can misclassify severe patients, leading them to avoid necessary care. Conversely, when hassle costs are high, patients tend to stay home, yet teletriage can misidentify mild patients, prompting them to seek unnecessary treatment. As a result, even if teletriage is costless, healthcare providers may abandon this service, even though teletriage always increases patient surplus by improving their assessment of their medical needs. Interestingly, patients can be better off seeking treatment from providers who charge higher prices, as these providers are more likely to adopt teletriage services. These results caution healthcare policymakers that teletriage can exacerbate healthcare inefficiency instead of alleviating it. Nevertheless, to improve patient surplus, the government should sometimes subsidize providers to offer teletriage or allow providers to charge a higher treatment fee.

2 - A Column Generation-Based Heuristic for Various Examination Stages Patient Scheduling Problem

Shaowen Lan, School of Economics and Management, Fuzhou University, Fuzhou, China, People's Republic of

With the increasing number of hospital visits, efficiently scheduling patient examinations within constrained resources has become a critical challenge for hospital management and the enhancement of medical service quality. This study addresses the patient scheduling problem in a hospital setting, where multiple patients undergo examinations simultaneously across various stages. The problem requires a comprehensive consideration of patients' different arrival times, sequential dependencies among examination procedures, resource constraints of examination rooms, and waiting times between examinations. To tackle this challenge, we formulate a mixed-integer programming (MIP) model to minimize the total completion time of all examination stages while reducing patient waiting times, aiming to improve hospital operational efficiency and patient satisfaction. To efficiently solve this problem, a column generation-based heuristic approach is developed which leverages a set partitioning-based reformulation of the original model. A variable neighborhood search algorithm is employed to generate an initial solution, and a heuristic algorithm is designed to exploit the structural properties of the subproblem, facilitating the generation of high-quality columns. Experimental results demonstrate that the proposed approach can generate high-quality solutions, effectively optimizing patient scheduling in a hospital environment.

3 - Multi-TIME Window and Multi-Objective Home Healthcare Scheduling Problem with the Collaboration of Online Diagnosis and Treatment and Onsite Door-to-Door Service

Gang Du, East China Normal University, Shanghai, China, People's Republic of

Driven by technological advances and the active support of national policy for the "Internet + nursing service" model, the service model of home healthcare organizations has become increasingly diversified, not only providing traditional home nursing services but also steadily increasing the demand for online diagnosis and treatment services. Based on this, this paper investigates the scheduling and routing problems of home healthcare workers in the synergy of online diagnosis and treatment services and onsite home care services of home healthcare organizations. The study focuses on three main objectives: minimizing the overall operating cost, improving the satisfaction of healthcare workers, and enhancing patient satisfaction. In order to solve the problem efficiently, this study designs a hybrid non-dominated sorted iterative greedy algorithm based on adaptive destruction operations, an initial population generation method with a hybrid strategy, six knowledge-based multi-objective adaptive destruction operators, a hybrid reconfiguration operation, and a local search strategy in conjunction with the problem characteristics. Extensive numerical experimental results show that the algorithm performs well in terms of diversity and convergence and is able to provide a series of reliable and diversified nondominated solutions for decision-makers to provide strong support for their decisions.

4 - Exfusion: An Explainable Multi-Scale Feature Fusion Framework for Medical Image Processing

Jianjin Yue, Sichuan University, Chengdu, China, People's Republic of, Li Luo, Mengzhuo Guo

In medical image processing, the region of interest (RoI) delineates the location and boundaries of lesions, helping physicians diagnose and prognosis planning. However, current methods struggle to integrate the high-resolution macroscopic features of the RoI with its low-resolution microscopic features. Additionally, they are usually designed separately for image classification or object detection, lacking the versatility to handle diverse tasks simultaneously. Moreover, these methods, often based on deep learning architectures, are frequently

optimized for performance but overlook model explainability, which hinders their broader adoption in real-world healthcare applications. To this end, we propose a new framework, Explainable Multi-scale Feature Fusion (ExFusion), for medical image processing, including an image preprocessing module, an encoder-decoder module, and an explainability module. The image preprocessing module employs various techniques to reduce background noise and emphasize features of RoI. The encoder-decoder module integrates a novel Enhanced Multi-scale Feature Fusion (EFF) block within a U-shaped network architecture. This allows for efficient representation of the RoI's local and global features, enhancing spatial sensitivity and improving the model's ability to capture complex information. The final module focuses on increasing the credibility of the framework's decision-making process by providing post-hoc visualizations that explain how predictions are made. We evaluate the proposed framework through various experiments, demonstrating its superior performance across different tasks compared to existing state-of-the-art models. We invite junior and senior radiologists to evaluate the framework's outputs. Results show ExFusion exceeds the average performance of radiologists, confirming its potential as an effective tool in real-world healthcare.

MA14

Wilberforce

Smart Manufacturing and Modeling

Invited Session

Semiconductor Manufacturing x AI

Chair: Ivan Tan, N/A

Co-Chair: Lin Miao Zhang, Micron Technology, Singapore, N/A, Singapore

1 - Using Virtual Silicon and Machine Learning to Improve Yield During Advanced Semiconductor Manufacturing

Joseph Ervin, Lam Research, Fremont, CA, United States

The rapid advancement of semiconductor technology has ushered in a new era of innovation, driven by the integration of virtual silicon twins and artificial intelligence. Virtual silicon twins, which are digital replicas of physical systems, have become pivotal in addressing high-volume semiconductor manufacturing challenges. This talk will explore how virtual process modeling, virtual metrology, and virtual process optimization with design of experiments (DOEs) are revolutionizing semiconductor production. By leveraging machine learning, these techniques enhance human intelligence and experience, providing real-time feed-forward and feedback optimization. The session will highlight specific examples demonstrating significant improvements in manufacturability and yield, showcasing the transformative impact of virtual silicon on the semiconductor industry.

2 - Evolution of Digital Twin to Accelerate Factory Ramp and Productivity

Madhav Kidambi, Applied Materials, Glen Allen, VA, United States

The Applied SmartFactory® Digital Twin Presentation outlines the current capabilities, evolution of digital twins and roadmap for implementing Digital Twins in manufacturing. A Digital Twin is defined as a model of manufacturing assets or operations that creates predictive and optimized outcomes. The presentation emphasizes the importance of Digital Twins to accelerate ramp of new products & technologies while improving yield and throughput. Various types of Digital Twins examples and capabilities are discussed, including Factory Equipment, Sub fab, Process, and Supply Chain Digital Twins. The paper also shows case studies on how to co optimize factory productivity, output, and yield by integrating digital twins. This presentation shows how AI/ML and cloud computing are helping address scaling challenges. The presentation concludes with a summary of the opportunities and challenges in implementing Digital Twins in the semiconductor industry.

3 - Using Genai to Accelerate Defect Classification and Detection in Semiconductor Fab

Jinhua Ding, Micron Technology, Singapore, Singapore

In semiconductor manufacturing, wafer inspection images play a crucial role in early defect detection and yield analysis. Traditional deep learning models worked well in typical classification, object detection and segmentation use cases, but their need for extensive datasets presents a challenge due to the dynamic nature of

the manufacturing process. The collection of these datasets is a time-consuming and expensive task, often hindering the speed of model deployment and optimization. This presentation shares about a novel approach to address this issue: using GenA to generate synthetic images that accurately reflect real-fab wafer images to boost model training and ultimately achieve faster model deployment and more accurate defect detection. A few use cases will be demonstrated on how GenAI is leveraged to transform defect classification and detection modeling cycle, and benefits for both yield an quality in frontend manufacturing.

4 - How Micron Partner with Key Equipment Suppliers to Harness Data for AI-Driven Innovation and New Manufacturing Capabilities

Yuka Katsuki, Micron , Japan, Japan, Jirapat Puapipat

In the new age of AI, where “data is the new gold,” the advanced semiconductor industry—powering the AI revolution—has developed an insatiable hunger for data. Every wafer processed by an equipment generates billions of bits of data that can be used to improve yield and manufacturing efficiency.

This presentation explores how Micron collaborates with key equipment suppliers to gather, organize, standardize, and process all available data from processing equipment into its cloud data infrastructure (DIMS program). These new datasets, including high-frequency sensor data, OES spectra, and equipment logs, unlock an endless world of new capabilities.

With the help of AI and ML techniques, the presenter will showcase examples of how Micron builds applications to analyze and leverage these datasets, develop ML models that explain and predict inline process variations, and enable new real-time defense-line capabilities. This presentation highlights how enabling new datasets can transform semiconductor manufacturing—enhancing process capability and operational efficiency

5 - Advanced Machine Learning Techniques for Optimizing Defect Review Sampling and Guided Inspection with Metrology, Design, and Process Data

Poh Boon Yong, KLA, Milpitas, CA, United States

Accelerating yield learning for advanced technology nodes in the semiconductor manufacturing process necessitates high sensitivity inspection to pinpoint critical defects. However, higher sensitivity optical inspection introduces increased noise and results in a higher SEM non-visual (SNV) count during review sampling, particularly with limited review SEM budget. Consequently, a high SNV level hinders chip manufacturers from identifying the yield-limiting defects, thereby impeding process optimization and yield improvement. SPOT, KLA’s production-ready machine learning platform, employs analytical predictive techniques to optimize review sampling strategies. SPOT integrates all historical defect information, metrology measurements, design hotspots, and equipment data to identify the critical defects. This paper also explores the use of SPOT for various applications such as Metrology Guided Review, Metrology Guided Inspection, and Hotspot Prediction with design signatures.

MA15

Fairmont Ballroom

Autonomous Navigation and Maritime Logistics

Contributed Session

Transportation Science and Logistics

Chair: Yibai Zhang, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, N/A

1 - Stochastic Ship Repositioning Optimization in Tramp Spot Market with a Cargo Selection Mechanism

Xiaoya LIAO, Department of Aeronautical and Aviation Engineering, The Hong Kong Polytechnic University, Hong Kong, China, People's Republic of, Lingxiao Wu, Shengnan Shu, Hui Zhao

In tramp shipping, the spot market is characterized by intense competition and volatility. Securing profitable spot cargoes for a tramp shipping company is challenging. To handle this issue, this paper addresses a ship

repositioning problem in the dry bulk spot market, aiming to provide short-term operational planning for a tramp shipping company. The goal is to maximize profits by optimizing destination decisions, speed, and cargo selection upon arrival. To handle high competition and market volatility, the cargo selection mechanism considers competitors' movements, a priority principle, and cargo-worthiness, along with uncertainties in future cargo availability and competitors' behavior. We propose a two-stage stochastic programming model to solve this problem. The first stage determines the destination and arrival time for each empty ship, while the second stage involves cargo acquisition based on first-stage decisions, revealed uncertainties, and the cargo selection mechanism. We use the Sample Average Approximation (SAA) method to solve the stochastic model. Extensive experiments validate the SAA method's performance, and the Value of Stochastic Solution (VSS) is calculated by comparing it with a deterministic model. Finally, a sensitivity analysis based on a case study informs the development of operational policies for tramp shipping companies.

2 - Tramp Fleet Renewal and Contract of Affreightment Decisions Under Uncertainty: An Approximate Dynamic Programming Approach

sirui ding, Tsinghua University, Beijing, China, People's Republic of, Xiwen Bai, Liangqi Cheng

Rising geopolitical tensions and economic instability are increasing volatility in tramp shipping markets. It becomes essential for companies to incorporate real-world uncertainty into both long-term fleet renewal decisions at the strategical level, as well as multi-period affreightment contracting at the tactical level. The former represents significant capital budgeting decision while the latter affects profitability by balancing spot and contract cargo under shifting market conditions. While fleet planning has been studied, multi-period contracting under uncertainty remains under-researched, especially integrated decision models. Besides, when modeling market uncertainty, expert insights and the interconnected movements of various market prices are vital. Nevertheless, many studies oversimplify these elements, limiting their practical relevance. Additionally, Sample Average Approximation (SAA), commonly used in stochastic programming, struggles with high-dimensional uncertainty and long planning horizons. To address these gaps, we propose a new stochastic programming model that jointly considers vessel acquisition (newbuild and secondhand), chartering, and cargo contracting decisions. We introduce a hybrid forecasting framework that combines expert-driven qualitative input with a Markov chain-based market-state transition process, generating realistic scenarios through historical data bootstrapping. To manage computational complexity, we reformulate the problem as a dynamic programming and solve it using Approximate Dynamic Programming (ADP) approach with a neural network-based value function approximation. A double value network is proposed to reduce overestimation bias. Numerical experiments based on a crude oil shipping company demonstrate that our approach significantly reduces operational costs and enhances decision quality. The findings provide practical insights and a scalable decision-support tool for managing uncertainty in tramp shipping.

3 - Conflict-Free Path Planning for Multi-Agvs in Automated Container Terminals

Lubin Fan, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China, People's Republic of, Shuai Jia

With the expansion of automated container terminals, collaborative path planning for automated guided vehicles (AGVs) faces challenges such as low efficiency in conflict detection and high computational complexity. This study proposes a hybrid approach combining classical path planning with machine learning to enhance the operational efficiency and stability of multi-AGV systems. Firstly, based on the terminal layout and the directional constraints of AGV lanes, a topological map model is established to describe the features of path network. Secondly, a customized A* (CA*) algorithm framework is designed for multi-AGV path planning. By treating conflict and reservation nodes as dynamic obstacles, and integrating dynamic path replanning and starting point waiting strategies, this framework achieves rapid obstacle avoidance and conflict-free path generation. Thirdly, to improve conflict detection efficiency, a multi-AGV simulation environment is developed to collect trajectory data. Supervised learning is then employed to model AGV movement patterns and train a lightweight conflict prediction module, which is embedded into the CA* framework to predict potential conflicts in real-time during path search. Experimental results demonstrate that the proposed method ensures zero conflicts in large-scale scenarios. This research contributes both theoretical insights and practical guidelines for efficient coordination of large AGV clusters in industrial environments.

4 - Optimizing Alternative-Fuel Tugboat Fleet Towards Decarbonization Goals

Yibai ZHANG, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China, People's Republic of, Jing Wang, Shuai Jia, Hai Yang

Harbor crafts are identified as key contributors to greenhouse gas emissions in marine ports. Renovation of the existing harbor crafts by alternative fuel transitions is highlighted as a significant focus for decarbonization from a range of cleaner energy pathways. Focusing on the tugboat fleet, one of the heavy-duty types of harbor crafts, this paper examines decarbonization strategies that optimize the alternative fuel fleet size under vessel service demand uncertainties. We propose a two-stage robust optimization model for determining the composition of the tugboat fleet powered by selected alternative fuels—biodiesel, liquefied natural gas (LNG), methanol, hydrogen, ammonia, and electricity. An adapted column-and-constraint generation algorithm is developed for solutions with minimized fleet operation cost and emission cost. The model and solution method are applied on instances generated from the tugboat and vessel operations data of the Port of Singapore. We provide a comprehensive analysis of the feasibility of these alternative fuels, in terms of operational performance, technical requirements, environmental impact, and economic viability, for the alternative fuel tugboat fleet. Additionally, with reference to the decarbonization goals set by the International Maritime Organization's (IMO), we discuss the transitions of fuel types for tugboat fleet at different decarbonization milestones, and also reveal the impacts of carbon prices on the tugboat fuel type transitions, based on our computational results.

MA16

VIP A

Empirical Research in Sustainability Topics

Invited Session

Sustainable and Socially Responsible Operations

Chair: Li Ding, Nanyang Technological University, Singapore, N/A, Singapore

1 - Fuel Price Shocks and Truck Driver Behavior: Evidence from a Natural Experiment

Yiming Zhang, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Ye Shi, Xingliao Wan

Truck driver safety and operational efficiency remain critical concerns in transportation management. Among these, fuel prices and fuel consumption are particularly significant to truck drivers' daily driving decisions. This study investigates the short-term effects of fuel price adjustments on truck driver behavior using a natural experimental setting and a difference-in-differences (DID) model. We employ two complementary datasets: one capturing drivers' routine shipment records and the other providing high-resolution trajectory data on drivers' real-time movements during cargo transportation. Our findings reveal that rising fuel prices correlate with an increase in driving speed, suggesting that drivers respond to cost pressures by accelerating. While fuel price hikes induce short-term changes in driving speed, they do not affect drivers' shipment choices. Additionally, we conduct a targeted survey to understand the underlying mechanisms driving these behavioral shifts. The survey insights suggest that economic pressure, operational efficiency, and fuel-saving motives play pivotal roles in explaining drivers' adaptive responses to fuel price changes. This research contributes to our understanding of driver behavior under economic constraints and informs strategies to promote safety and operational stability in the logistics sector amidst fuel price volatility.

2 - Reducing Gender Disparities in Online Labor Platforms: the Role of Performance Feedback and Task Exposure

Li Ding, Nanyang Technological University, Singapore, Singapore, Basak Kalkanci

Despite the speculation of online labor platforms as equalizers of opportunity, gender disparities in job access and income persist among gig workers. These disparities not only hurt worker welfare but also create inefficiencies for platforms. Unlike traditional labor markets, online labor platforms such as Upwork and Prolific offer anonymity, flexibility, and short-term engagements. These differences obscure the factors driving the disparities and complicate efforts to mitigate them. Our paper investigates the behavioral mechanisms underlying these disparities and potential mitigation strategies. We conduct experiments on Prolific, where online workers

are recruited to grade elementary school-level math problems of two difficulty and payment levels. We find that women earn significantly less than men because they are less likely to select the more difficult task compared to men of comparable capabilities, despite a higher payment for the same expected task completion time. This difference in task selection is partially driven by women's negatively biased self-evaluation of performance, particularly for more difficult tasks. Motivated by practice, we then examine two operational levers, performance feedback and task exposure, as potential mitigation strategies. We show that performance feedback alone may not be sufficient to motivate women to pursue the more difficult task later on because women often choose to work on the easier task in the first phase, for which they have a relatively unbiased self-evaluation. Supplementing performance feedback with early exposure to challenging tasks can effectively reduce gender disparities in task selection and significantly boost women's productivity and hourly earnings without adverse effects on men.

3 - Green Image Management in Supply Chains: Strategic Disclosure of Corporate Suppliers

Yilin SHI, Chinese University of Hong Kong, Hong Kong, China, People's Republic of, Jing Wu, Yu Zhang, Yuqing Zhou

This paper examines how firms manage their disclosure of customer-supplier relationships to create a favorable green image in their supply chain. Our study finds robust evidence that firms strategically disclose relationships with environmentally responsible ("good") suppliers while withholding relationships with "bad" suppliers, ceteris paribus. This strategic disclosure is particularly pronounced for firms with a greater concern for their brand image, a higher level of institutional ownership, and a smaller improvement of suppliers' environmental performances in the future. Moreover, the extent of this strategic disclosure decreases as regulations on environmental information transparency become more stringent. Additionally, we observe that companies that strategically disclose ``green'' suppliers have higher stock returns and asset turnover in the short term, although these effects diminish over time. This suggests that it takes time for investors and consumers to fully understand such strategic disclosures.

MA17

VIP B

Technology and Regulation in Social Media and E-Business

Contributed Session

eBusiness and Social Media

Chair: Zishun Qian, Huazhong University of Science and Technology

1 - Blockchain and Fake News

Hang Lee, National Kaohsiung University of Science and Technology, Kaohsiung, Taiwan

In response to the widespread misinformation on social media, many countries have implemented robust strategies. These include media literacy programs, enhanced fact-checking collaborations, and legal repercussions for spreading false information. Despite these efforts, traditional fact-checking struggles with the rapid spread and sheer volume of fake news, leading to increased interest in technological solutions like AI. However, AI's effectiveness is limited by its accuracy and reliability in distinguishing true from false content. This study proposes blockchain technology as a potential solution for curbing fake news on social media by ensuring data authenticity through immutable records. Grounded in the Technology Acceptance Model (TAM), this research aims to assess how blockchain impacts users' trust and willingness to engage with content on social media platforms. The findings are expected to guide future applications of blockchain in enhancing content veracity, offering strategic insights to both academics and industry practitioners.

2 - Activating Innovation Through Enterprise Social Media: the Role of Thriving at Work and Mindfulnes

Shanshan Guo, University of Science and Technology Beijing, Beijing, China, People's Republic of, Jinnan Wu

Abstract: Enterprise social media (ESM) has become an essential tool for organizations seeking to enhance efficiency and foster growth. Despite its growing importance, there is limited research on how to effectively

manage ESM and connect its usage to employees' innovative behaviors. Grounded in the job demands-resources model, this study investigates the effects and underlying mechanisms of different types of ESM use on employees' innovative behaviors, with thriving at work as a mediator and mindfulness as a moderator. Using paired data from 320 participants in two stages, the findings reveal that both work-related (WU_ESM) and social-related (SU_ESM) ESM positively influence employees' innovative behaviors. Additionally, the study demonstrates that thriving at work mediates the relationship between ESM use and innovative behaviors. Moreover, the indirect effects of WU_ESM and SU_ESM on innovative behaviors, via thriving at work, are amplified by mindfulness. These results broaden the understanding of how ESM impacts employees' innovative behaviors from a resource-based perspective and introduce mindfulness into this theoretical framework, offering novel insights for the effective and scientific management of enterprise social media.

3 - Resource Planning for Order Fulfilment in Quick Commerce

Hasmukh Gajjar, Indian Institute of Management Indore, Indore, India

Quick Commerce is rapidly expanding across the globe. The major challenge for quick commerce is to handle the tradeoff between operating cost and delivery speed to fulfil customer orders within a small time-windows, usually within 10 to 15 minutes. TUsing a stochastic optimization approach, this work attempts to optimize the number of pickers and delivery agents to maintain required service level considering variability in customer demand and the order fulfilment processes such as picking, packing, delivery etc.

4 - Communication Intensity and Consumer Purchase: Quantifying the Impact of Permission and Unsolicited Modes

Yihan Liu, Huazhong University of Science and Technology, Wuhan, China, People's Republic of, Jun Yang, Yuanyuan Zhou

The development of an effective communication strategy is a pivotal factor in the management of customer relationships and the assurance of sustainable financial growth. The communication intensity is the primary issue that must be addressed. With the advent of online technologies and social media, permission communication has become a widespread mode for firms to engage with their customers. Our study focuses on the differential impact of permission and unsolicited communication intensity on consumer purchase. To this end, we have analyzed 31 weeks of communication and transactional activity data from a consumer finance firm. The results demonstrate that an inverted U-shaped relationship between the intensity of unsolicited communication and consumer behavior, which aligns with the results of previous research. However, our study reveals a U-shaped effect of permission communication intensity on consumers’ purchasing behavior, and provides inner mechanisms and theoretical explanations for the potential causes based on Social Penetration Theory (SPT) and Cognitive Dissonance Theory (CDT). We then quantitatively explain how the effects of permission communication intensity are mitigated by intensity of broadcast in social network, temporal inconsistency, and consumer loyalty, and reinforced by content richness. We used the control function method to address endogeneity issues and conducted robustness checks to support our conclusions. This study sheds light on the value of firms’ permission communication, offering critical managerial insights on leveraging permissive communication.

5 - Ai in Audiobook Narration: Boon OR Bane?

Sara Vessal, ESSEC Business School, Cergy, France, Debajyoti Biswas

In the last two years, there has been a surge of generative Artificial Intelligence (AI) in different industries across the spectrum, disrupting business models and traditional jobs. AI has started to make an impact in the audiobook industry, with major players like Apple and Google investing in AI narration for audiobooks. Although this makes the listening experience less authentic than the human-narrated version, it eases the barriers to entry for independent publishers to distribute audiobooks by shifting the production costs of audiobooks from the publisher/creator to the audiobooks distribution platform. In this paper, we consider a two-sided platform with publishers sharing their audiobooks on the platform and customers joining the platform to listen to these audiobooks. We compare the platform profits and decisions when it adopts AI technology to when it does not by examining two different settings - (i) monopoly and (ii) duopoly for the subscription-based business model. We find that

We find that in the monopoly setting, the platform would prefer to invest in AI only when the AI technology cost coefficient is sufficiently low, whereas in the duopolistic setting, both platforms would prefer not to invest in AI. In an extension, we examine a special case where both subscription and purchase options are offered by the platform in both monopoly and duopoly settings and find that their preferences for investing in AI depend on the technology cost coefficient.

6 - The Impact of Generative AI Announcements on Suppliers: Evidence from The Stock Market

Zishun Qian, Huazhong University of Science and Technology, Wuhan, China, People's Republic of, Jing Peng, Jianbin Li

The rise of Generative AI (GenAI) technology is revolutionizing firm operations and supply chain management, profoundly influencing firms within the GenAI supply chain network. While a few emerging studies have explored the impact of GenAI initiatives on firms' stock performance, it remains unclear how such disruptive technology initiatives affect the stock price movements of their supply chain partners. This paper empirically investigates the impact of firms' GenAI initiatives on the abnormal stock returns of their suppliers by analyzing data on news, supply chain relationships, and daily stock prices. Our empirical analysis reveals that, on average, firms' GenAI announcements lead to a positive abnormal return of 0.27% for their suppliers on the announcement day, and they also lead to long-term positive effects on suppliers' stock market value. We further investigate how this spillover effect varies across suppliers. The results indicate that the abnormal return is more positive for suppliers with higher levels of R&D intensity, greater sales growth, located closer to their customers, and operating within a less competitive industry. Additionally, we find that the spillover effect on suppliers is stronger when the announcement focuses on product improvement rather than process improvement. Our findings have important implications for managers and investors.

Monday, July 21, 9:30 AM - 10:00 AM

Moor/Morrison/Hullet Foyer

AM Poster Session

Poster Session

1 - Assessing Uncertainty in Stock Returns: A Gaussian Mixture Distribution-Based Method

Yanlong Wang, shenzhen international graduate school, tsinghua university, shenzhen, China, People's Republic of

This study proposes a deep learning framework incorporating Gaussian mixture distributions to model time-varying return uncertainty in China's stock market. Our Transformer-based architecture captures non-normal return characteristics (skewness, heavy tails), demonstrating superior volatility forecasting accuracy compared to GARCH-type models, particularly during high-volatility regimes. We develop stock code embeddings using bag-of-words representations to quantify inter-asset uncertainty correlations, enabling visualization of risk profile clusters for portfolio management. The framework evaluates distributional accuracy via CRPS and volatility predictions through MSE/QLIKE metrics. Empirical tests on 3,200 Chinese equities reveal significant improvements in volatility estimation accuracy, with the model achieving top-tier performance across all evaluation metrics. By simultaneously characterizing return distributions and asset uncertainty linkages, this approach advances financial risk modeling through two synergistic mechanisms: enhanced short-term volatility forecasting precision and systematic identification of cross-asset risk dependencies.

2 - Turning the Tide: Firm-Specific Variables as Moderators Between Economic Policy Uncertainty and Financing Decisions

SHAO-CHI FANG, National Sun Yat-sen University, Kaohsiung City, Taiwan, Shih-Sian Jhang, Chien-Yuan Sher

This study investigates how economic policy uncertainty (EPU)—encompassing tax, regulatory, spending, and monetary policy fluctuations—affects corporate financial leverage. Despite mixed results from previous studies, some studies have associated EPU with a reduction in borrowing due to a higher default risk, while others suggest an increased use of debt when equity financing becomes more expensive. In this study, we utilize the Compustat annual database and cover the period between 2000 and 2023 for U.S. public companies. Empirical

analyses of panel data are conducted using firm-fixed effects and industry-fixed effects models. We contribute to the literature by bridging economics, operations management, and finance, clarifying inconsistencies in findings on EPU's impact on leverage, and offering practical advice to corporate decision-makers as they navigate policy-driven uncertainty.

3 - American Options: A 10-Year Bibliometric Review

Ziyu Xu, Hwa Chong Institution, Singapore, Singapore

American options are popular amongst investors due to their early exercise feature that provides great flexibility. However, to the best of current knowledge, there are no prior studies that provide a comprehensive review on American options specifically. This poster presents a bibliometric review of 160 articles on American options in the past decade, obtained from the Web of Science database. Firstly, this poster identifies three major clusters of research subtopics obtained from a keyword co-occurrence network on Citespace, and provides a detailed content analysis. Secondly, the most cited references and the most cited papers are surfaced in this poster to uncover articles with substantial impact on research. Thirdly, evolution analysis is presented, which acknowledges recent developments and emerging subtopics of research in this field through the timeline and citation burst features on Citespace. Finally, based on observations, several future research directions to bridge existing gaps in literature are recommended.

4 - Price Limit Restriction and Stock Market Efficiency

Xianwei Liu, Harbin Institute of Technology, Harbin, China, People's Republic of

Can easing price limits improve stock market efficiency? The Shenzhen Stock Exchange’s 2020 ChiNext reform provided a unique research setting. Using a difference-in-differences (DID) approach, this study examines the impact on volatility, liquidity, and market efficiency. Results show that relaxed price limits increased volatility (higher daily amplitude and return volatility) but also boosted liquidity (higher turnover, lower Amihud illiquidity). Additionally, it reduced price delay and improved market efficiency (lower SYN measure). The findings assess the reform’s effects and offer policy insights for China’s market regulation.

5 - Collaborative Transfer of Critical Patients Based on Bayes-Seihrd-Mcmc Joint Prediction

Ji Kai, School of Economics and Management, Nanjing University of Science and Technology, Nanjing, China, People's Republic of

This research focuses on resolving the intricate decision-making challenge related to the cross-regional coordinated transfer of critical patients in epidemics. Initially, an infectious disease model (SEIHRD) incorporates a Bayesian framework. Markov Chain Monte Carlo (MCMC) sampling is then applied to dynamically modify the posterior distribution of uncertain parameters. This method assesses the uncertainty in the stochastic demand for critical care caused by the spread of infectious diseases. Subsequently, we build a multi - stage stochastic programming model (MSPM) with time-changing demand scenarios. This model enables a dynamic decision - making process for patient transfers. Furthermore, we detail and contrast four proposed transfer strategies. The results indicate that, compared to three alternative strategies, the cross-regional coordinated transfer strategy achieves the best balance between objective utility and risk. This study gives public health authorities a quantitative basis for formulating hierarchical medical service plans and optimizing regional patient transfer coordination in various emergency situations.

6 - Cost-Effectiveness of a Precision Digital Health Program for Glucose Management in Pediatric Patients with Newly Diagnosed Type 1 Diabetes

Paul Dupenloup, Stanford University, Stanford, CA, United States, Margaret Brandeau, David Scheinker

Background: The Teamwork, Targets, Technology, and Tight Range (4T) Study is a team-based type 1 diabetes (T1D) care model focused on tighter glucose targets, early technology use, and continuous glucose monitoring (CGM)-based remote patient monitoring (RPM). While previously associated with improved glucose management, reduced provider effort, and equitable outcomes, its cost-effectiveness remains unclear.

Methods: We used a Markov model to compare the 4T model with standard care (self-monitoring of blood glucose) and CGM alone in newly diagnosed pediatric T1D patients. We tracked hemoglobin A1C, adverse events, and progression of diabetes complications. Lifetime costs, quality-adjusted life years (QALYs), and

incremental cost-effectiveness ratios were evaluated from a healthcare sector perspective, discounted at 3% annually, using 2024 U.S. dollars. Probabilistic and deterministic sensitivity analyses were performed.

Results: The 4T model is cost-saving and offers greater health benefits.

Conclusion: The technology-enabled 4T model is a cost-saving and beneficial strategy for managing new-onset T1D.

7 - Trauma Care Resource use Optimization: Network Analysis, Sequential Pattern Mining, & Machine Learning

Michael Dorosan, Singapore Health Services Pte Ltd, Singapore, Singapore, Jun Hong Soh, Yichuan Zhang, Wei Qiang See, Chantal Lim, Jen Heng Pek, Shao Wei Sean Lam

Trauma care aims to stabilize, diagnose, treat, and rehabilitate critically injured patients to restore their health and function. Often, such cases require urgent care amidst rapidly deteriorating conditions; resources must be efficiently and appropriately allocated. We focused on an acute hospital's trauma cases seen at the emergency department (ED) within 2019-22. Through network analysis we surfaced varying sequential resource utilization patterns among three severity tiers, signaling differences in utilization patterns. However, the existing tiering system is done almost retrospectively for national-level reporting and monitoring purposes. We then explored two approaches: sequential pattern mining to uncover prefixes that lead to downstream use of often constrained intensive care resource (ICU), and machine learning using patient characteristics at triage to predict ICU use. We uncover rare prefixes that confidently lead to ICU use. Tree-based models struggled with the severed imbalance of the dataset. Future work will explore combining both approaches through Bayesian inference.

8 - Prediction of Intensive Care Unit Admission in Sepsis Patients Based on Activity Sequence from Event Logs

Jongchan Kim, Yonsei University, Wonju, Korea, Republic of

This study proposes a novel approach for predicting ICU admission in sepsis patients. Using machine learning models, we extract patterns from the sequential activities recorded. These sequences provide insights into the progression of the condition in the emergency room and the likelihood of ICU admission. The model is trained on a dataset containing detailed event logs of sepsis patients, where each patient's case includes a series of activities leading to the final outcome—ICU admission or discharge. We compare the performance of various machine learning algorithms, including Random Forest and CatBoost, to evaluate their predictive accuracy. Additionally, transfer learning is applied to improve model performance by leveraging knowledge from similar datasets. The proposed approach not only offers a predictive tool for early ICU admission identification but also demonstrates the potential of using event log data in clinical decision support systems to enhance sepsis management and patient care.

9 - From Ethics to Trust: the Role of Brand Authenticity in Consumer Responses to Suspended Meal Initiatives

Long-Chuan Lu, National Chung Cheng University, Taiwan, Taiwan, Yan-Jun Huang, Yih-Jen Lu, Tzu-Chiao Kuo, Hsi-Chieh Lin, Hsin-Ting Lin

In today's socially responsible climate, ethical behavior has become a vital component shaping consumer perceptions of brands. This study explores how the practice of "Suspended Meals"—prepaid meals donated by customers for people in need—affects consumer evaluations of restaurants. As restaurants increasingly share these initiatives on social media, consumer attention often focuses on the sincerity behind such actions, influencing perceived brand authenticity. Drawing on 849 valid questionnaires and structural equation modeling, this research examines whether perceived ethical behavior enhances brand authenticity and how this authenticity subsequently impacts brand image, recommendation intention, and brand trust. Results indicate that when consumers perceive restaurants as genuinely committed to ethical practices, they are more likely to view the brand as authentic. Furthermore, brand authenticity significantly mediates the relationship between perceived ethical behavior and favorable brand outcomes. These findings offer practical implications for businesses aiming to build trust and loyalty through authentic ethical engagement.

10 - Structure-Aware Reinforcement Learning for Routing Operations Optimization and Its Application to Skilled Nursing Facility Recommendation

hongtao yu, Dalian University of Technology, Dalian, China, People's Republic of, Zachary Hass, Nan Kong

Post-acute care (PAC), including skilled nursing facilities (SNFs), plays a vital role in patient recovery and hospital resource management in the U.S., with Medicare spending exceeding $60 billion annually. However, rising demand, workforce shortages, and concerns over SNF quality highlight the need for improved operations management. This paper proposes a novel decision support framework for PAC facility recommendation that is client-centered and responsive to system dynamics. We develop a utility model using aggregated data to capture individual preferences and SNF characteristics, then embed this into a large-scale Markov Decision Process (MDP) that accounts for service heterogeneity and queueing dynamics. To solve the MDP, we introduce a structure-aware reinforcement learning (SARL) algorithm incorporating monotonicity and superconvexity constraints. Using data from 510 SNFs in Indiana, our model reveals disparities in access based on age and race and demonstrates a 10.7% improvement in client utility and system efficiency over traditional approaches.

11 - Green OR Greenwashed? Examining the Psychological Pathways Linking Sustainability Advertising to Consumer Trust

Long-Chuan Lu, National Chung Cheng University, Minhsiung, Taiwan, Jou-An Su, Yan-Jun Huang

With growing concerns over climate change and environmental degradation, consumers increasingly favor brands that reflect their sustainability values. Green advertising has emerged as a key strategy for companies to promote their environmental efforts and enhance brand image. However, the rise of such advertising has also led to more instances of greenwashing—misleading environmental claims that can erode consumer trust. While prior research has recognized green skepticism and confusion as consumer responses, few studies have explored their mediating roles in the link between advertising message processing, trust formation, and behavioral intention. This study proposes an integrative model to explain how green advertising influences consumer trust and green purchase intention through these psychological mechanisms. The findings contribute to a deeper understanding of consumer responses to sustainability messaging and offer practical guidance for companies seeking to communicate their environmental commitments more credibly and ethically in increasingly discerning markets.

12 - Profit Split OR Risk Share? An Analysis of off-Invoice vs. Scan-Back Promotions Under Demand Uncertainty

Yu Jiang, University of Science and Technology of China, Hefei, China, People's Republic of, Xiang Ji, Jie Wu

This paper studies a manufacturer's optimal decisions on trade promotion and product line design when the manufacturer sells through an independent retailer to an uncertain market. The two major types of trade promotion, off-invoice and scan-back, are considered. We show that discrete demand uncertainty makes the retailer change her ordering strategy, and the retailer may order more products when the market demand volatility greatly. We also show that when consumers have discrete demand uncertainty, a manufacturer may prefer scan-back to off-invoice, whether he is selling a product or a product line. And this kind of preference can be strengthened if the manufacturer adopts a longer product line. In addition, a manufacturer has more incentives to extend his product line under both promotion policies when the fixed cost of per-category is non-negligible.

13 - Disruptive Innovation Reexamined: An Analytical Framework for Market Entry and Segment Targeting

Kilsun Kim, Sogang University, Seoul, Korea, Republic of, Hartanto Wong, Dilip Chhajed

Our paper seeks to address the gap in research focused on ex-ante analysis of disruptive innovations by introducing an analytical model to provide a more proactive framework for identifying the potential advantages and challenges associated with each entry strategy. A critical element of our model is the consideration of the general market preference structure, encompassing both dominating and non-dominating structures. We address the following questions: (1) What are the essential conditions for successful entry when introducing a new

product attribute into a market served by an incumbent? (2) Which market segment(s) should the entrant target? (3) How does the consumer preference structure influence these decisions?

14 - Social Bundling with Status Products: the Role of Social Technology and Reference Group Effect

Mingjun Li, University of Science and Technology of China, Hefei, Anhui, China, People's Republic of

This study examines how social technology and consumer social effects influence bundling strategies in conspicuous product markets. We consider a distribution channel in which a retailer sells a conspicuous product and has the option of bundling it with a manufacturer's product. We study the retailer's optimal bundling strategy of the conspicuous product in the presence of social technology impact. Our findings show that bundling helps retailers of ordinary status products charge higher prices by leveraging followers' imitation behavior. For luxury goods, bundling helps retailers better serve high-end consumers and distinguish their products. Both social technology effect and consumer segments may facilitate (or weaken) retailers' incentives to bundle. Notably, better social technologies do not always lead to higher profits. Even manufacturers of non-status products can benefit from the retailer's bundling strategy, as wholesale prices are influenced by both bundling and social effects, which may reinforce or offset each other.

15 - Store Brand Introduction Strategy Under Cap-and-Trade Regulation

Yelin Hu, University of Science and Technology of China, Hefei, China, People's Republic of

Introducing store brands boosts retailer profitability but raises carbon emissions, conflicting with tightening environmental regulations. With stricter environmental regulations and carbon policies, retailers struggle to introduce store brands while meeting compliance requirements. Against this backdrop, this paper investigates green store brand introduction strategy under cap-and-trade. We design a management system comprising a manufacturer, a third-party manufacture, a retailer, and heterogeneous consumers, and apply a game-theoretic model to study the relevant decision-making processes.Our study demonstrates that retailers choose in-house production when fixed costs are low, but prefer upstream manufacturers collaboration when carbon prices, green levels, or eco-conscious demand rise; otherwise they opt for third-party manufacturers. Additionally, introducing green store brand can reach win-win outcome for retailers and manufacturers.This research advances existing research on the interaction between store brand strategy and cap-and-trade by uncovering strategic insights into store brand introduction, offering policymakers and practitioners actionable insights for regulated markets.

Monday, July 21, 10:00 AM - 11:00 AM

Fairmont Ballroom

AI up in the Sky: The Future of Aviation Security

Invited Session

Plenary

1 - Ai up in the Sky: the Future of Aviation Security

sheldon jacobson, University of Illinois, Champaign, IL, United States

The United States Transportation Security Administration (TSA) secures over 400 commercial airports around the nation to keep the air system secure from threats. They have historically been threat detection centric, working to keep threat items off airplanes. Since the launch of TSA PreCheck program in 2011, a risk-based system that matches security resources with passenger risk, their focus has shifted, working to separate the large pool of benign passengers from the handful of bad actors. This talk provides an overview of airport security in the United States and how biometrics like facial recognition fits into the risk-reduction calculus for airport security. A roadmap for the future of airport security is provided that relies on facial recognition that will transform airport security from a process of stopping prohibited items from entering the air system to reducing air system risk by better knowing the population of air travelers and identifying risky passengers prior to them entering the air system.

Monday, July 21, 11:15 AM - 12:45 PM

MB01

Enterprise

Advances in Decision Analytics

Invited Session

Applied Probability Society

Chair: Long Zhao, NUS DAO, Singapore, Singapore

Co-Chair: Fuqing Li, Nationaly University of Singapore, Singapore, N/A, Singapore

1 - Aggregating OR Interacting?: the Role of Individual Overconfidence

Wenjie Tang, University of Vienna, Vienna, Austria, Yanwei Jia

We compare two approaches to combining individual judgments about an unknown value into a collective point estimate: statistically aggregating individual estimates or asking individuals to interact and reach a joint estimate. The latter interacting approach pools private information from all group members, and thus is superior in this regard to the former aggregating approach where each individual estimate is based solely on the individual's own information. We are interested in whether the aggregating approach is able to produce a more accurate collective estimate than the interacting approach.

Even under perfect group interactions, we show that the aggregating approach can do so when individuals are overconfident, that is, when they misinterpret their private information without being aware of it. In particular, when the level of individual overconfidence is very high (resp., very low), the aggregating approach (resp., interacting approach) outperforms; when the overconfidence is at a moderate level, the performance comparison depends on the group size. Importantly, the finding that the aggregating approach may outperform the interacting approach is not driven by the existence of individuals' misinterpretation errors, but rather by their ignorance of it. Moreover, we show that whereas the accuracy of the aggregating approach always increases in the group size, it is not always optimal to have as many interacting group members as possible when these group members are overconfident. Finally, we provide the theoretical optimal solutions for the joint decision of how to form collective estimates, aggregating versus interacting, and how to determine group size.

2 - An Ambiguity Framework to Inform Optimal Data Allocation for Data-Driven Optimization: Analysis and Application to Agribusiness

Ying He, Southern University of Science and Technology, Shenzhen, China, People's Republic of, Saurabh Bansal, Mahesh Nagarajan

This paper is based on a research collaboration with a large agribusiness firm in North America. The firm invests a considerable sum annually in developing new varieties of seeds for food crops. In many geographical regions, the changing climactic conditions interact with the potential of seeds in in complex ways. As a result, it is no longer accurate to assume specific parametric distributions for yield uncertainty. Instead, the firm now collects yield data during field trials and directly optimizes its large-scale production decision based on these observations as potential scenarios, using sample average approximations in the form of stochastic linear programs. Recent literature provides performance guarantees for these programs as a function of number of observations. These guarantees tighten as data size increases, however collecting observations is also significantly costly. In this paper we develop a framework to optimize this trade-off between the benefit and cost of collecting observations. Furthermore, the firm had a limited budget to collect samples of a portfolio of seeds, and it needed a transparent way to determine the optimal number of observations for the yield uncertainty. We extend our framework to determine the allocation of a limited number of test fields to various seeds. We also provide an analytical expression for ambiguity premium for a canonical data-driven optimization problem in the presence of performance bounds. Subsequently, we discuss the implementation of this protocol at the firm, along with a documentation of the benefits realized.

3 - Individual Differences In Forecasting: a Heterogeneous Bin Approach

Marat Salikhov, New Economic School, Moscow, Russian Federation, Ville Satopaa

We extend the Bias-Information-Noise (BIN) framework for probabilistic forecasting by incorporating both heterogeneity across forecasters and events, and by allowing for general forecast-level covariates that can affect all three components simultaneously. Our model decomposes forecast errors into bias (systematic deviation from

base rates), information processing (over- or under-weighting of available information), and noise (random variation), while accounting for how these components vary with individual characteristics, event features, and forecast-specific factors such as timing and experimental conditions. Using data from the Good Judgment Project, we find that cognitive abilities, personality traits, and collaborative structures influence forecasting performance through different mechanisms. For instance, fluid intelligence reduces bias and improves information processing but increases noise, while cognitive reflection reduces both information sensitivity and noise. Treatment effects reveal that structured collaboration in "superteams" improves all three components simultaneously, with facilitated teams showing the largest benefits. To quantify these effects, we develop a matching-based approach using Shapley values that accounts for the inherent path dependence in decomposing forecast errors. Our findings suggest that effective forecasting systems should be tailored to individual characteristics and that different interventions may be needed to address specific sources of forecast error.

4 - Enhancing Forecast Aggregation with Side Information About Experts

Fuqing Li, National University of Singapore, Singapore, China, People's Republic of, Zhi Chen, Long Zhao

Traditional forecast aggregation methods, such as the linear opinion pool, rely primarily on limited historical forecast errors to assign aggregation weights. However, in many real-world settings, additional background information about the forecasters—such as their professional training or area of expertise—is also available. This paper explores how such background information can be systematically incorporated to enhance the accuracy of aggregated forecasts. Using about 50,000 time series from the M4 Forecasting Competition dataset, we demonstrate that our proposed method often outperforms standard benchmark approaches including equally weighted aggregation, especially when expert background characteristics are indicative of individual forecasting performance.

MB02

Mercury

Air Transport Management

Invited Session

Aviation Applications

Chair: Achim Czerny, Hong Kong Polytechnic University, Hong Kong, VA, Hong Kong

1 - Aviation Ecosystem: Strategic Insights from LLM-Augmented Global Route Analysis

Jinjia Huang, Institute of Operations Research and Analytics, National University of Singapore, Singapore, Singapore, Chaik Ming Koh , Chung Piaw Teo

In partnership with VariFlight and Hangzhou Airport Authority, this study applies large language models (LLMs) to uncover strategic drivers of global inter-city aviation growth. By integrating socio-economic, cultural, and city-specific factors, we identify key patterns shaping route performance—highlighting the lasting influence of historical demand, post-pandemic recovery dynamics, and the interplay between city connectivity, economic ties, and travel appeal. These insights inform targeted route development and support data-driven policy design to strengthen Hangzhou’s position as a globally connected aviation hub.

2 - Hub Competition and Demand Modeling in Long-Haul Aviation Markets

Li Mingmei, Singapore University of Technology and Design (SUTD), Singapore, Singapore, Nuno Ribeiro, Sebastian Birolini

This study advances air transport modeling by proposing an integrated analytical framework that bridges supply-demand interdependencies, airline network optimization, and strategic competition within a systematic methodological structure. We adopt a cutting-edge integrated demand modelling approach, leveraging a two-layer nested logit structure to simultaneously address demand allocation and demand generation. A mixed-integer nonlinear programming (MINLP) model is introduced to optimize long-haul airline network configurations, explicitly incorporating supply-demand interdependencies to refine route network design and

flight frequency planning. Furthermore, strategic interactions among competing airlines are modeled through a game-theoretic framework to capture adaptive decision-making in oligopolistic markets.

Empirical findings demonstrate the framework's efficacy in projecting long-term demand, simulating adaptive strategic responses to competitive disruptions, and identifying optimal airline network configurations under demand and competition uncertainties. By systematically integrating demand dynamics, network optimization, and competitive interactions, the framework offers a robust tool for analyzing the evolution of airline networks and market structures. These insights hold practical significance for airline executives and aviation policymakers in navigating complex market environments.

MB03

Minto

Frontier Topics in Market Design: The View from Economics

Invited Session

Auctions and Market Design

Chair: Kwok Hao Lee, National University of Singapore, Singapore, N/A, Singapore

1 - Platform Design in Dynamic Differentiated Goods Markets - the Case of Airbnb

Wenxuan Xu, Toulouse School of Economics, Toulouse, France

This paper studies the allocation of goods in two-sided, one-to-one matching markets with finite time horizons. In these markets, the timing of seller participation distorts the allocation of goods. This is due to dynamic externalities inherent in one-to-one matching markets -- sellers fail to fully consider how their participation impacts the evolution of the market. Using data on Airbnb, I estimate a dynamic structural model of a two-sided matching market in order to quantify externalities related to the timing of seller entry. I evaluate policies that a platform can implement to improve total surplus on its market. Entry subsidies that encourage later entry reduce total welfare by $3 for every $1 spent, while subsidies that encourage earlier entry increase total welfare by $2 for every $1 spent.

2 - Non-Linear Peer Effects with Conditionally Random Assignments

Suk Joon Son, University of Tokyo, Tokyo, Japan

We reexamine the peer effect on student achievement utilizing two recent econometric advancements. First, we allow the whole distribution function of peers' test scores to affect the outcomes utilizing the functional principal component analysis, allowing for flexible effects. Second, we utilize the conditionally random lottery-based assignments generated by the Deferred Acceptance algorithm as exogenous variation in co-enrolling peer composition. Our findings from New York City middle schools reveal that peer effects are non-linear and non-monotone. While peers with slightly higher baseline scores relative to a student's own baseline scores significantly improve the student's academic performance, peers with very high baseline scores have negative impacts. Aggregating this non-monotone effect, the linear-in-mean model misleadingly estimates the null peer effect.

3 - Content-Hosting Platforms: Discovery, Membership, or Both?

Tat How Teh, Nanyang Technological University, Singapore, Singapore, Ben Casner

Our model classifies platforms in the "creator economy" --- e.g. Youtube, Patreon, and Twitch-into three business models: pure discovery (helping consumers search for creators); pure membership (enabling direct creator-consumer monetization); and hybrid (combining both). Creators respond to platforms' decisions by choosing their content design along a broad-niche spectrum, trading off between viewership size and per-viewer monetization. With a monopoly platform, moving from pure discovery to hybrid increases profit while inducing a more niche content design; moving from pure membership to hybrid can reduce platform profit while inducing broader content design. We show how intense platform competition reverses these trade-offs.

4 - Distributional Effects of Vehicle Ownership Restrictions

Kwok Hao Lee, National University of Singapore, Singapore, Singapore, Tiffany Tsai

We study vehicle population controls, their dynamic effects, alternative policy designs, and implications for efficiency and equity. We leverage data on the universe of vehicles in Singapore, the number of which is limited by the Vehicle Quota Scheme (VQS). Under the VQS, vehicle registrations must be accompanied by certificates allocated in fortnightly auctions; after ten years, certificates expire and may be renewed at the prevailing auction price. First, we observe that the effective tax on vehicle ownership fluctuates by as much as S$10,000 per year per vehicle, possibly mispricing the negative externalities of ownership and use. To explain these large fluctuations in certificate prices, we formulate a dynamic equilibrium model of vehicle ownership and replacement by individuals and firms, and take it to administrative microdata from the Singapore government.

5 - Connect OR Clique: the Role of Agent Networks in the Housing Market

Ruiqi Wang, National University of Singapore, Singapore, Singapore, Hefan Zheng

This paper studies how real estate agents' collaboration networks affect the search behavior and purchase decisions of the consumers they serve. The collaboration network among agents is identified and clustered using unsupervised machine learning algorithms. We develop and estimate a model of dynamic consumer search using complete search records of consumers and the network structures of agents in the Beijing housing market. We find that the size and exclusivity of agents' collaboration networks have significant effects on consumer decisions and welfare. Consumers whose agents belong to larger collaboration networks tend to pay higher prices and search more intensively, while the opposite holds for higher network exclusivity.

MB04

Indiana

Behavioral Operations Management I

Invited Session

Behavioral Operations Management

Chair: Liang (Leon) Xu, Singapore Management University, Singapore, Singapore

1 - To Err is Human: a Field Experiment in Prescription Safety

Xiaodan Shao, Nanyang Technological University, Singapore, Singapore, Vivek Choudhary, Arnab Majumdar, Burhanuddin Pithawala

Can simple, non-mandatory, and cost-effective informational interventions effectively reduce prescription errors, given that mandatory systems often fail due to high overrides by doctors? Do such interventions facilitate proactive learning or merely promote reactive corrections? We address these questions by examining major drug-to-drug interaction (DDI) errors in prescribing, a critical patient safety issue causing significant harm. Using a difference-in-differences model, we analyze data from a randomized field experiment involving 2.81 million prescriptions by ~1,700 doctors in India. Doctors were randomized into a treatment group receiving non-mandatory real-time informational interventions about potential errors or a control group without interventions. Results show that the treatment group prescribed ~8.6% fewer errors compared to controls, potentially reducing hospitalization costs by US$4.8 million and saving an estimated 134 lives annually. Interestingly, doctors exposed to the intervention appear to effectively learn from the interventions, resulting in proactive avoidance of prescribing errors rather than mere corrections. Initially, error reduction is driven by immediate corrections, but sustained improvements result from proactive avoidance facilitated by learning. The effect is more pronounced among doctors with higher workloads. We find that the effectiveness of interventions remains consistent regardless of intervention fatigue and does not negatively impact productivity or care quality. Managers can strategically utilize heterogeneity in treatment effects to target specific doctors and address errors more effectively. The feasibility of proactive learning reduces the reliance on continuous digital support, particularly benefiting healthcare systems in regions undergoing digitalization with limited internet connectivity.

2 - The Impact of Social Comparison on Newsvendor Decisions: An Experimental Investigation

Liang (Leon) Xu, Singapore Management University, Singapore, Singapore, Valery Pavlov

For humans, it is important to perform not just well, but well relative to others. This study examines how individual newsvendor decisions are influenced by social comparison. Specifically, we investigate how the two

comparison mechanisms (profit-based and rank-based) affect the dimensions of social comparison (i.e., ahead-seeking and behind-aversion) that are responsible for different ordering behaviors and whether social comparison could be leveraged for profit improvement. We conducted two controlled experiments. In the first one, simulating typical real-life settings, two newsvendors make inventory decisions simultaneously. We manipulated the comparison by providing decision-makers with different information about their peers: orders only, whose profit proved higher (without exact profit values), and exact profits. The data shows that subjects tend to adjust their orders toward those of their peers, particularly if they underperformed in the previous round. As a result, this peer-chasing behavior prompts learning, virtually non-existent in individual newsvendor experiments, and the performance gradually improves. Because the simultaneous-move game does not allow a clean structural estimation of our behavioral model, we conducted another sequential-move experiment. The estimation results are consistent with the findings of the first experiment. Overall, our results highlight that the psychological discomfort of being outperformed is much stronger than the reward of outperforming others. Additionally, our structural estimation quantifies a novel factor not captured by the theory: subjects derive distinct utility from matching their peers' decisions and thus eliminating any risk of underperforming.

3 - How Do Behavioral Biases Affect Backfiring of Intensified Auditing on Corporate Social Responsibility?

Luyao Li, Tsinghua University, Beijing, China, People's Republic of, Xiaobo Zhao, Wanshan Zhu, Jinxing Xie

Auditing is widely used to motivate suppliers to exert more corporate social responsibility (CSR) efforts. However, trade journals reported that intensified auditing backfired, reducing suppliers' CSR efforts and prompting them to hide violations to pass audits. We conduct an experiment to examine suppliers' behavioral biases and their impact on this "backfiring effect." The experiment provides evidence for two key behavioral biases: loss aversion and probability weighting. Interestingly, the two biases have opposite influences: loss aversion mitigates, but probability weighting aggravates the "backfiring effect." Despite their conflicting influences, our analysis reveals that loss aversion dominates, resulting in an overall alleviation of the "backfiring effect" by behavioral biases. Our findings imply that, in practice, managers can improve CSR by making good use of behavioral biases' positive impact.

4 - Multi-Attribute Sequential Search: Theory and Experiment

Minjiao Cheng, National University of Singapore, Singapore, Singapore, Yaozhong Wu

This paper presents theoretical and empirical findings on the multi-attribute sequential search problem. In this problem, decision makers sequentially evaluate a series of alternatives, each characterized by multiple attributes whose values are unknown but can be learned gradually at a cost. During the search process, decision makers face three options at each step: gather more information about the current alternative by sampling an additional attribute, move to the next available alternative, or accept the current alternative. We characterize the optimal searching policies in a dynamic programming setting, and design an online controlled experiment to study how decision makers make those choices under various experimental conditions. Our findings reveal that decision makers tend to search more extensively than rationally predict when the search cost is high, whereas they tend to search less at low search costs. And it was found that the decision makers tend to oversearch for earlier stopping scenarios but undersearch for later stopping scenarios. These patterns hold for both the breadth of the search in terms of the number of alternatives assessed and the depth of the search in terms of the number of attributes sampled within an alternative. We incorporated bounded rationality into the decision maker's choice model and compared it with alternative behavioral models. Our structural estimation shows that bounded rationality can explain the data well.

MB05

Orchard A

Data-Driven Platform Operations

Invited Session

Decision Analysis Society

Chair: Hao Sun, Shenzhen University, Shenzhen, N/A

1 - Data-Driven Robust Allocation in Cash Transfer Programs

Jiannan Ke, Wuhan University, Shanghai, China, People's Republic of, Zhihang Yuan, Zhiyuan Chen

Cash Transfer Programs (CTPs) are widely used in different regions of the world as an effective approach to fighting extreme poverty. A key issue is how to utilize information from beneficiaries to design an allocation policy if potential beneficiaries cannot be accurately targeted. In this paper, we construct a robust allocation model to minimize the objective of the FGT index with parameter α based on a scenario-wise ambiguity set. To deal with the nonlinearity and infinite number of constraints in the robust model, we employ the linear decision rule and duality theory, which does not lead to a loss of optimality. Specifically, we reduce the robust model into a compact second-order cone program for α=1 (Poverty Gap Index) and a compact mixed-integer second-order cone program for α=2 (Poverty Severity Index). Based on the China Family Panel Study (CFPS) dataset, we divide the beneficiaries into different clusters with selected covariates and estimate the parameters of the ambiguity set from the predicted demand by the regression tree. Our robust allocation policy turns out to have better performance than the robust benchmarks without covariate information and the robust model based on KL divergence in terms of efficiency and equity, indicating the advantage of scenario-based distributionally robust models in moderating conservativeness to balance inclusion and exclusion errors.

2 - From Doorstep to Diet-Related Disease: the Health Impacts of Online Food Delivery Platforms

Yao Li, Southern University of Science and Technology, Shenzhen, China, People's Republic of, Lin Qiu, Liangfei Qiu

Online Food Delivery (OFD), as an emerging sector of the sharing economy, has transformed the food landscape. Unlike traditional food sources, OFD platforms may promote unhealthier dietary choices by facilitating easy access to and digital marketing of unhealthy foods, thereby heightening the risk of diet-related diseases. While prior studies have explored the association between OFD use and diet-related health risks, they have primarily relied on small-scale, self-reported data, leaving the causal impact of OFD on the incidence of major diet-related diseases unestablished. To address this gap, we utilize a comprehensive dataset of hospital outpatient visits across the 12 most populous counties in New York State, examining the health impacts of OFD platforms, specifically focusing on major diet-related conditions. Leveraging staggered OFD entry dates and employing a difference-in-differences methodology, our analysis reveals that the introduction of OFD platforms leads to an increase in hospital visits for overweight and obesity, diseases directly tied to poor diet. However, no significant association was found with hospital visits for conditions such as hyperlipidemia, hypertension, and diabetes, which are less directly influenced by diet alone. We identify the increased availability of unhealthy food options on OFD platforms as a primary driver of elevated obesity-related hospital visits, with further analysis indicating stronger effects within high-frequency OFD user groups and in New York City areas. These findings underscore the need for sustainable operations management in OFD platforms to address health impacts and provide actionable insights for OFD providers.

3 - Gig Worker Social Referrals on An on-Demand Food Delivery Platform

Hao Sun, Shenzhen University, Shenzhen, China, People's Republic of, Hai Wang

Social referral programs are commonly used by online labor platforms to incentivize labor supply through rewarding existing workers for successful referrals. This study investigates the impact of such programs on gig workers' labor supply in online labor platforms using data from an on-demand food delivery platform in Singapore. First, we analyze how gig workers' past labor supply and referral behavior influence the generation and value of social referrals. Second, we assess the benefits that the online labor platform derives from social referrals in terms of gig work labor supply. Finally, we examine whether referral-linked gig workers are more likely to operate in the same or nearby regions compared to non-referral pairs, providing empirical evidence on the spatial effects of social referrals. This research offers insights into the mechanisms driving labor supply dynamics in the gig economy and highlights the effectiveness of social referral programs in shaping worker behavior and enhancing platform performance.

4 - Balancing Supply with Demand on Ride-Hailing Platforms in Markets with Price Regulations: An Operational Approach Operational Approach

Qin Zhou, East China Normal University, Shanghai, China, People's Republic of, Jingqi Wang, Yifan Jiao, Jinzhao Du

Ride-hailing platforms face a supply-demand imbalance. During peak periods, the demand from passengers far exceeds the supply of drivers, while during off-peak periods, there is an abundance of supply but weak demand. The well-studied surge pricing can be challenging to implement in markets where prices are subject to regulation and are inflexible to adjust. We study an alternative operational approach that shifts the supply from off-peak to peak periods to address the supply-demand imbalance. We propose two novel incentive schemes: the qualification scheme and the prioritization scheme. Specifically, the platform sets a work target for the peak period. Under the qualification scheme, the platform assigns off-peak service requests only to drivers who meet the peak period target. Under the prioritization scheme, the platform prioritizes off-peak requests for drivers who meet the peak period target. We analyze the effectiveness of these schemes, considering the openness of the platform's supply system. For platforms with a closed system that only allows full-time drivers to provide service, the qualification scheme improves the total matching volume to a greater extent but hurts full-time drivers more than the prioritization scheme. For platforms with an open system that also allows the abundant part-time drivers to serve off-peakrequests, the prioritization scheme outperforms the qualification scheme in improving total matching volume. Furthermore, the implementation of an incentive scheme in an open system may benefit both the platform and full-time drivers. Our study suggests an alternative to surge pricing for on-demand platforms to address supply-demand imbalance when prices are inflexible.

MB06

Orchard B

Supply Chain Finance

Invited Session

Finance

Chair: peng liu, N/A

1 - Advertising Signaling in New Product Financing: Platform Financing vs. Bank Financing

Yinghui Liu, National University of Singapore, Singapore, Singapore, Gongao (Lucy) Chen

This paper studies the interplay between advertising signaling and financing choices for manufacturers launching new products. We compare the efficiency of platform financing with traditional bank financing under settings of information symmetry and asymmetry. With a signaling game framework, we examine scenarios where the platforms hold an informational advantage over banks. Contrary to conventional understanding, our findings suggest that bank financing, especially under information asymmetry, can enhance overall supply chain performance.

2 - Bayes Asymmetric Similarity Estimation in Binary Classification: a Case Study in Bankruptcy Prediction

Wei Li, Harbin Institute of Technology, Harbin, China, People's Republic of, Wolfgang Härdle, Stefan Lessmann

There has been intensive research regarding machine learning models for predicting bankruptcy in recent years. However, the lack of interpretability limits their growth and practical implementation. This study proposes a data-driven explainable case-based reasoning (CBR) system for bankruptcy prediction. The results show the CBR system achieves superior predictive accuracy in comparison to existing other CBR systems, and the performance is among the best of the prevailing machine learning models. We demonstrate that the asymmetrical structure effectively captures the features of financial data, therefore improving not only the accuracy of predictions but also the explainability. In addition, we delicately examine the explainability of the CBR system in the decision-making process of bankruptcy prediction. Our findings reveal a potential research avenue that a data-driven designed explainable model can reconcile the tension between boosting prediction accuracy and increasing explainability.

3 - Optimizing Digital-Enabled Procurement Financing Under Uncertain Demand and Credit Risk

Jinjia Huang, Institute of Operations Research and Analytics, National University of Singapore, Singapore, Singapore, Hang Xu, Chung Piaw Teo

Procurement payment financing leverages the focal firm's credit and verified transaction data from the digital financial platform, allowing the bank to finance distributors' purchase orders and alleviating their financial burden on inventory management and market expansion. This study investigates how the bank and the focal firm adopt procurement payment financing to create value in the supply chain. We consider a digital financial platform focusing on supply chain transactions between a wholesaler and multiple distributors. First, to address the uncertainties of distributor demand distribution and credit risks, we introduce a robust optimization framework to optimize the bank's decision-making. We find that focal firm management can mitigate distributors' credit risks and improve the effectiveness of procurement payment financing. Moreover, using distribution data from a microprocessor manufacturer, we demonstrate the benefits of procurement payment financing over traditional individual bank financing. These findings provide managerial insights for a profit-seeking bank on financing strategies and offer focal firms valuable guidance on distributor management, thereby enhancing supply chain resilience and competitiveness.

4 - Offline Pricing with Additive Demand

Mabel C. Chou, National University of Singapore, SG, Singapore, Jingren Liu, Hanzhang Qin

We study a personalized pricing problem using observational data, with the objective of finding an optimal personalized pricing policy to maximize revenue. A key challenge in the pricing problems with observational data is the lack of counterfactual outcomes: we only observe customer demands generated under a historical pricing policy, but not under any other policies that could have been used. To address this challenge, we propose a Sample Dependent Adaptive (SDA) estimator that achieves minimax-optimal mean squared error for evaluating the counterfactual expected revenue of the pricing policies of interest. Building on the SDA estimator, we design an effective algorithm for pricing policy optimization.

MB07

Orchard C

Operations Optimization in Healthcare Systems

Invited Session

Health Applications Society

Chair: Zhenzhen Yan, Nanyang Technological University, Singapore, Singapore

1 - Appointment Scheduling with Sequential Servers Using Distributionally Robust Optimization

Zhenzhen Yan, Nanyang Technological University, Singapore, Singapore, Youhui Goh, Zihao Li

This paper addresses an appointment scheduling problem involving multiple sequential servers using a distributionally robust optimization (DRO) approach. Two decisions are optimized: the schedule for patient visits and an adjusting policy to rebalance customers' waiting time across servers. We formulate the distributionally robust appointment scheduling problem in sequential-servers systems using conic optimization, incorporating service time correlations. We find that the traditional cost-minimization approach results in imbalanced waiting times, concentrated at downstream servers. To address this, we propose strategically idling upstream servers, inspired by queueing literature, and develop a DRO model to jointly optimize the schedule and strategic idling (SI) policy. Through extensive numerical studies, we thoroughly examine the role of SI in a system's performance and the effect of correlation information on the optimal schedule and SI policy. Finally, using data from a clinic, we conduct a case study to demonstrate the performance advantage of our SI policy, over existing SI policies in the literature. Managerial implications are summarized below: First, our SI model provides a jointly optimal schedule and SI policy that is effective in balancing waiting times across servers, in some scenarios, and also reduces total waiting time, with the cost borne by the service provider's overtime. Second, the typical dome-shaped schedule provides an advantage in systems with multiple congested servers, as it evenly

distributes congestion across servers. Finally, incorporating correlations lessens the expected cost and improves the patients' waiting experience with balanced waiting time across servers.

2 - Helping the Captive Audience: Advance Notice of Diagnostic Service for Hospital Inpatients

Zheng Zhang, Zhejiang University, Hangzhou, China, People's Republic of, Miao Bai, Nan Liu

Inpatients are often treated as the "captive audience" on-demand for hospital diagnostic services, and they are notified only when service capacity is available. This arrangement causes significant chaos and inefficiencies in hospital operations. We propose an innovative scheduling approach called "advance notice" to manage hospital diagnostic practice. Advance notice is a brand-new scheduling paradigm in between the classic allocation scheduling and advance scheduling. Patients are placed in a common queue waiting to be called for service, and they will be provided both a fixed preparation time and a guaranteed service time window in advance. The advance notice policy enjoys the benefit of allocation scheduling and that of advance scheduling. It calls for two decisions: who to serve now and who to send advance notices to. We formulate a Markov Decision Process model to optimize these decisions dynamically. Via a novel variable transformation, we reveal the hidden antimultimodular structure of the problem and show how the optimal decisions should be adjusted in response to changes in the system load. Our numerical study, populated by real data from a large academic medical center in the United States, demonstrates significant improvement in operational efficiency by switching from current practice to adopting our proposed advance notice policy. The advance notice policy strikes a fine balance between the two classic scheduling paradigms: it gives the provider flexibility in using her capacity as in allocation scheduling, and it reduces patient online waiting as in advance scheduling.

3 - Leveraging Telemedicine to Tackle Regional Health Disparities via Inter-Hospital Collaborations

Guohua Wan, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Nan Liu, Sergei Savin, Yunting Shi

Rural populations lack access to care compared to their urban counterparts, and health disparities between urban and rural residents are well documented. The rise of telemedicine offers a new avenue for expanding access to quality care in rural regions. This study explores telemedicine-based collaboration between an urban and a rural hospital. Although forging collaboration between the two hospitals can improve the care quality at the rural hospital, it may also be costly for the urban hospital, which needs to divert some of its care capacity to help rural hospital to deliver high-quality care. We develop a three-stage Stackelberg-game model to characterize the interactions among the social planner, the two hospitals, and patients. In our model, the social planner sets the compensation amounts for the two hospitals, the urban hospital proposes a capacity and revenue sharing contract to the rural hospital, and patients choose their care facility. We find that the extent of telemedicine collaboration depends on the reimbursement rates at both hospitals and their capacity and demand dynamics. Notably, a telemedicine collaboration may not always result in improved social welfare. However, a social planner can simply adjust the ratio of per-patient reimbursement rates at both hospitals to induce welfare-maximizing strategies. Our research highlights the significance of both financial incentives and operational characteristics in shaping telemedicine-driven collaboration between the hospitals.

4 - *Resource Allocation in Urban Search and Rescue Operations*

Huiyin Ouyang, The University of Hong Kong, Hong Kong, Hong Kong, Zheqi Zhang, Nilay Argon, Serhan Ziya

This work addresses resource allocation in urban search and rescue during flooding emergencies, inspired by Hurricane Harvey (2017). Using Markov decision processes, I model how to optimally deploy limited rescue assets amid uncertainty. For two-class prioritization, I characterize the optimal policy structure and develop effective heuristics for multi-class, multi-server scenarios. Simulation testing validates these approaches, providing actionable insights for emergency managers while contributing to scheduling theory. The work bridges operational research with practical disaster response, helping decision-makers maximize rescue effectiveness when resources are constrained.

5 - The Impact of Physician Familiarity on Emergency Department Productivity

Yiran Zhang, The University of Hong Kong, Hong Kong, Hong Kong, Huiyin Ouyang, Zhankun Sun

Emergency departments (EDs) frequently face challenges with overcrowding, resource limitations, and inefficiencies. Extensive research aims to enhance ED operational performance by understanding physician practices, particularly focusing on physician productivity and patient handoffs. Our working paper examines the impact of physician familiarity on productivity, handoff practices, and patient outcomes, building on existing literature in worker familiarity and productivity. Pairwise physician familiarity is quantified by the number of shared shifts with active interactions. Physician productivity is measured by new patient pickups and instantaneous service rate. Using two years of data from a Canadian teaching hospital, we aim to provide insights for improving ED efficiency.

MB08

Bras Basah A

Data-Driven Healthcare Analytics and Decision Support

Contributed Session

Healthcare Applications

Chair: Ahmadreza Pourghaderi, National University of Singapore, Singapore, Singapore

1 - A Study of the Economic Linkages of the Health System: Network Analysis of Singapore&Rsquo;S Economic Input-Output Table

Shao Wei Sean Lam, SingHealth Duke-NUS Academic Medical Center; Singapore Management University, Singapore, Singapore, Ashish Kumar

Background

Conventional analysis of forward and backward linkages of Singapore’s input-output table shows that the health services sector has a meagre role in the economy. Modern network analysis represents an input-output table as a weighted directed graph and provides a richer perspective of economic structure. Our aim is to assess the role of health services in Singapore from this perspective.

Methods

We used public data on industry-wise direct requirements from Singapore’s national accounts for the year 2022 to construct an economic network. Using Python’s NetworkX package, we computed betweenness centrality for each industry and used the Louvain community detection algorithm to establish the presence of communities. We used Gephi for network visualization.

Results

Using conventional measures, out of 107 industries, the health services sector ranks 86^{th} on forward linkage and 47^{th} on backward linkage. However, when viewed from a network perspective, it ranks 31^{st} on betweenness centrality. Further, it is part of a clique of 24 industries that contribute to 28% of the nation’s value added.

Conclusions

The multiplier effects of health services are under-stated by conventional measures. Network analysis shows that in addition to being the custodian of population health, the health services sector is an important node in the annual flows of value in the national economy. Investments in health should be viewed in this light. Our work generates new evidence from existing data in support of this view.

2 - Nutriopt: Bridging Pixels and Robust Optimization for Personalized Nutrition Tracking

Haiyan Yu, Peking University Chongqing Big Data Research Institute, Chongqing, China , China, People's Republic of, Kun Zhang, Senlin Li, Renying Xu, Zhiqi Chen, Ehsan Hajizadeh, Tao Huang

Traditional dietary tracking methods—manual logs, recalls, questionnaires—are error-prone and unscalable, undermining personalized healthcare. We introduce *NutriOpt*, an AI-driven framework that transforms meal images into actionable nutritional plans. Leveraging a U-Net model, *NutriOpt* analyzes meal images to identify foods, estimate portions, and compute nutrient profiles with over 90% accuracy. We then bridge computer vision with optimization: a novel goal programming model balances personalized nutrient targets (e.g., protein) against caloric needs, while a robust optimization layer mitigates real-world uncertainties like food variability and measurement noise. Validated on real-world datasets, *NutriOpt* achieves over 10% higher precision than existing tools, reduces tracking time by 50%, and adapts seamlessly to dietary noise. This work pioneers a scalable synergy of deep learning and decision science, offering clinicians and patients a dynamic tool for precision nutrition. Results underscore its potential to reshape chronic disease management, public health strategies, and AI-driven healthcare optimization.

3 - Predicting Device-Related Complications Among Women Undergoing Breast Implant Surgeries Using Machine Learning Techniques

Ahmad Reza Pourghaderi, Monash University, Melbourne, Australia, Arul Earnest, Patrick Garduce, Susannah Ahern

Background: Breast implant surgeries are increasing globally, yet predicting complications remains challenging. While registries like the Australian Breast Device Registry (ABDR) have improved post-market surveillance, traditional methods offer limited predictive power. This study evaluates the utility of machine learning (ML) in predicting device-related complications following breast implant surgeries.

Methods: We analysed primary breast implant surgeries recorded in the ABDR from 2012 to 2023, covering cosmetic and reconstructive indications. An extreme gradient boosting (XGBoost) model was developed to predict revision procedures due to complications, using comprehensive demographic, clinical, procedural, and device-level data. The dataset was split 70:30 into training and testing sets, and class imbalance was addressed using the ROSE technique. Recursive feature elimination guided variable selection, and hyperparameters were optimised via 10-fold cross-validation. Model performance was assessed using AUC, sensitivity, specificity, PPV, and NPV.

Results: The XGBoost models demonstrated strong predictive performance, with AUCs exceeding 0.75 in both cosmetic and reconstructive cohorts. Key predictors included patient age, surgeon type, surgical techniques (e.g. mastopexy, plane of insertion), and device characteristics. Despite moderate predictive value of device-level variables, the model revealed multifactorial drivers of complications.

Conclusions: This study provides the first large-scale application of ML for complication prediction in breast device surgery using national registry data. These models offer a promising framework for real-time risk stratification, enhancing clinical decision-making and regulatory oversight. Future efforts should focus on external validation, integration with clinical workflows, and incorporation into interactive dashboards for precision surveillance.

MB09

Bras Basah B

AI-Driven Innovation in Maritime Traffic Management

Invited Session

Maritime AI

Chair: Ke Wang, N/A

1 - Safety 4.0: AI-Driven Ship Safety Management System

Kum Fai Yuen, Nanyang Technological University, Singapore City, Singapore

2 - AI-driven AIS Data Mining

Zipei Fan, LocationMind Inc, Tokyo, Japan

The Automatic Identification System (AIS) generates vast quantities of real-time maritime traffic data, offering unprecedented opportunities to enhance safety, security, and efficiency in global shipping. Leveraging advanced AI methodologies—including graph neural network, generative deep learning, large language models (LLMs), and anomaly detection algorithms—researchers and practitioners are developing practical tools to extract actionable intelligence from this raw data. This talk will introduce our recent research progress in AI-driven AIS data mining, specifically covering: using diffusion-based generative models to enhance AIS data quality, applying graph neural networks to predict future vessel trajectories and discovering the changes of the global maritime traffic network, and employing LLMs to analyze maritime news and estimate the resulting impact on maritime traffic.

3 - Intelligent Systems for Vessel Traffic Management: Bridging Data, Modeling, and Computing

Zhe Xiao

MB10

Hullet Room

Scalable Control and Operation of Large-Scale Complex Power Systems

Invited Session

Power Grid

Chair: Peng Yang, Xi'an Polytechnic University, Xian, N/A

Co-Chair: Feng Liu, Tsinghua University, Beijing, N/A

1 - Stability Analysis of Large-Scale Power Systems: a Distributed and Equilibrium-Free Paradigm

Peng Yang, Xi'an Polytechnic University, Xi'an, China, People's Republic of

Traditional centralized stability analysis struggles with scalability in large modern power grids with massive heterogeneous devices. This talk proposes a compositional and equilibrium-free approach to analyzing power system stability. We establish the theory that use equilibrium-free local conditions to certificate system-wide stability of power systems with heterogeneous nonlinear devices and structure-preserving lossy networks. This is built on a recently developed notion of delta dissipativity, which enables local stability conditions without knowing the system-wide equilibrium. As a consequence, our proposed theory can certificate stability of equilibria set rather than single equilibrium. We also present methods for applying our theory to large-scale power grids, together with case studies across a wide range of system scales. Our results enable a more scalable and adaptable approach to stability analysis, accommodating variations in system equilibrium. It also sheds light on how to regulate grid-connected devices to guarantee system-wide stability.

2 - Switching Topology Control for Networked Microgrid Systems

Tao Liu, The University of Hong Kong, Hong kong, China, People's Republic of

Networked microgrid (NMG) systems, formed by interconnecting multiple microgrids, offer enhanced reliability and resilience but are inherently prone to stability challenges, particularly under disturbances. While conventional stability control methods rely on inverter and load control, this talk introduces a novel topology control approach that leverages the flexibility of NMG interconnection networks for stabilization. By switching between different network topologies in response to disturbances, the proposed control scheme ensures system stability without causing power supply interruptions. This talk will cover the theoretical foundations, computational methods, and practical implications of topology control, paving the way for a new paradigm in NMG stability management.

3 - Optimal Control in Both Steady State and Transient Process with Unknown Disturbances for Power Systems

Zhaojian Wang, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Ming li

The scheme of online optimization as a feedback controller is widely used to steer the states of a physical system to the optimal solution of a predefined optimization problem. Such methods focus on the optimality in the steady state, without considering the performance during the transient process. In this paper, we simultaneously consider the performance in both the steady state and the transient process of a linear time-invariant system with unknown disturbances. The performance of the transient process is illustrated by the concept of overtaking

optimality. An overtaking optimal controller with known disturbances is first derived to achieve the transient overtaking optimality while guaranteeing steady-state performance. We propose a disturbance independent near-optimal controller, which can achieve optimal steady-state performance and approach the overtaking optimal performance in the transient process. The system performance gap between the proposed controller and the optimal counterpart proves to be inversely proportional to the control gains. A case study on a power system with four buses is used to validate the effectiveness of the proposed controller.

4 - Adaptive Pricing for Optimal Coordination in Networked Energy Systems with Nonsmooth Cost Functions

Yan Jiang, Chinese University of Hong Kong, Shenzhen, Hong Kong

Modern energy systems are undergoing a significant transformation, marked by the increasing prevalence of distributed energy resources (DERs), responsive loads, and the emergence of more autonomous devices. These developments have created opportunities for customers to actively participate in system operations. However, unlike dispatchable resources, customers often cannot be directly controlled by an operator and must be coordinated through some form of incentives. Incentive pricing mechanisms for distributed energy consumption have shown promise in aligning individual benefits with system objectives, especially under rising privacy concerns from customers. Nevertheless, the system cost in serving the demand profile is often determined from a DCOPF problem, which can make the cost depend nonsmoothly on the demand. To overcome this new layer of complexity, we present a generalization of the pricing update rule that leverages the generalized gradients of the system cost function. Borrowing tools from nonsmooth analysis, we prove that the resulting dynamic system asymptotically converges to the unique equilibrium that achieves the social welfare.

5 - Sharing Energy in Wide Area: a Two-Layer Energy Sharing Scheme for Massive Prosumers

Yifan Su, Chongqing University, Chongqing, China, People's Republic of

The popularization of distributed energy resources transforms end-users from consumers into prosumers. Inspired by the sharing economy principle, energy sharing markets for prosumers are proposed to facilitate the utilization of renewable energy. This paper proposes a novel two-layer energy sharing market for massive prosumers, which can promote social efficiency by wider-area sharing. In this market, there is an upper-layer energy sharing market (U-ESM) in the distribution system and numerous lower-layer energy sharing markets (L-ESMs) in local communities. Prosumers in the same community share energy with each other in the L-ESM, which can be uncleared. The energy surplus and shortage of L-ESMs are cleared in the U-ESM. Thanks to the wider-area two-layer structure, the market outcome is near-social-optimal in large-scale systems. However, the proposed market forms a complex mathematical program with equilibrium constraints (MPEC). To solve the problem, we propose an efficient and hierarchically distributed bidding algorithm. The proposed two-layer market and bidding algorithm are verified on the IEEE 123-bus system with 11250 prosumers, which demonstrates the practicality and efficiency for large-scale markets.

6 - A Matrix-Valued Differential Passivity Index for Power System Small-Signal Stability Analysis

Xi Ru, Department of Electrical Engineering, Tsinghua University, Beijing, China, People's Republic of, Feng Liu, Xiaoyu Peng

This paper presents a novel small-signal stability criterion tailored for power systems integrating heterogeneous dynamic devices and distributed generation. The key contribution is the introduction of a matrix-valued differential passivity framework, through which we define a differential passivity index capable of characterizing both passivity excess and deficiency at the component level. The proposed criterion ensures system stability when the collective passivity surplus of the devices compensates for the passivity shortfall introduced by the network interconnections. To facilitate localized analysis and implementation, a distributed formulation of the criterion is developed, employing the minimal eigenvalue of the matrix index as a tractable local stability indicator. This approach generalizes existing differential passivity-based methods and extends their applicability to a broader class of power system models. Furthermore, we show that common components in power systems inherently exhibit the proposed differential passivity properties. An algorithm is provided for computing matrix-valued passivity indices in general device models. The effectiveness of the proposed framework is validated through simulation studies on both a three-bus system and the IEEE 118-bus benchmark.

7 - Towards Impact of Angle-Voltage Coupling on Small-Signal Stability of Power Systems: a Damping Perspective

Xiaoyu Peng, Tsinghua University, Beijing, China, People's Republic of, Cong Fu, Peng Yang, Xi Ru, Feng Liu

The rise of inverters integrated into power systems has raised enormous concerns about potential system instability. This study aims to investigate the correlation between phase-angle and voltage dynamics and their impacts on the stability of power systems, focusing on damping perspectives. To this end, we first formulate a generic model to unify the dynamics of synchronous generators and diverse inverter-interfaced devices. By exploring systems' damping characteristics, we find that angle-voltage coupling tends to undermine stability by redistributing the system's damping in a more dispersed manner, consequently rendering the system more vulnerable to inadequate or even negative damping. Our work also involves analyzing and quantifying the factors affecting coupling strength, including power flow, control parameters of devices, and the number of integrated inverters. A cross-loop control is also introduced to demonstrate the potential benefits of utilizing angle-voltage coupling to enhance stability. We validate these findings through simulations on a modified IEEE 39-bus system with heterogeneous inverters and other standard benchmarks with up to 500 buses. This comprehensive investigation yields insights into the complexities of angle-voltage coupling, which offers useful implications for the analysis and control of power systems.

MB11

Morrison Room

Revenue Management and Innovations in Online and Omnichannel Retailing

Invited Session

Revenue Management and Pricing

Chair: Yimeng Sun, Singapore Management University, Singapore, N/A, Singapore

1 - The Impact of Kol Targeting Strategy in Live Ecommerce: An Empirical Study on Tiktok

Kai Tian, Singapore Management University, Singapore, Singapore, RUIJIE ZHANG, Yun Fong Lim, Kejia Hu

The use of live-streaming advertisements on social media platforms has garnered significant popularity across various industries because of their potential to significantly enhance the marketing performance. In this paper, we adopt propensity score matching (PSM) and R(1) auto-regressive linear models to investigate the impact of targeting the followers of key opinion leaders (KOLs) on advertising outcomes. We identify three categories of marketing outcomes: traffic, purchase intention, and engagement. Our empirical results indicate that a KOL targeting strategy reduces traffic, enhances purchase intention, but has no significant effect on engagement. Specifically, the KOL targeting strategy reduces the number of advertisement clicks and the number live-streaming room entries by 0.552 and 0.490 respectively. On the other hand, it effectively increases slidecart interactions by 0.147 clicks, order creation by 0.027 orders, and order payment by 0.024 orders. To improve the overall marketing performance without incurring additional costs, we further conduct a counterfactual analysis to first identify critical parameters for specific marketing outcomes and then fine-tune non-critical parameters. Our findings offer practical guidance for advertisers seeking to enhance the effectiveness of their live-streaming advertising campaigns.

2 - Multi-Product Pricing and Multi-Channel Budget Allocation for E-Commerce Retailers: a Data-Driven Target-Oriented Robust Optimisation Framework

yue zeng, Singapore Management University, Singapore, Singapore

As online multi-channel marketing becomes increasingly prevalent, e-commerce retailers face challenges in effectively distributing budgets across various channels and setting prices for multiple products to maximize overall returns. To tackle these challenges, we propose a data-driven decision-making framework using robust optimisation to improve an e-commerce retailer’s profitability through optimized multi-product pricing and multi-channel budget allocation under demand uncertainty in an online platform.

Using a world-renowned retailer's skincare data, we empirically construct a demand forecasting model that captures critical spillover and promotion effects. Facing demand uncertainty, we adopt a target-oriented robust optimisation (TRO) approach to absorb as much uncertainty as possible in our solutions while ensuring a target total return. We conduct extensive experiments on both a real dataset and synthetic datasets. The TRO policy demonstrates notable improvements in daily average returns and is especially promising during promotion phases, yielding approximately $16,000 more per day in the real-world data case study. On a series of synthetic datasets designed to test robustness under varying numbers of products and channel structure, the TRO approach consistently outperforms the deterministic policy in both the regular and promotion phases.

We find that considering the randomness of the error terms of the demand model can significantly raise the retailer's return. Retailers can leverage these findings by modestly lowering prices during both regular and promotion phases to boost competitiveness, while in the promotion phases, directing more budget to channels that expand product exposure to potential consumers, instead of marketing channels functioning as search ads.

3 - Joint Pricing and Refund Optimization for Omnichannel Retailing with Product Returns

Yimeng Sun, SUTD, Singapore, Singapore

We consider an omnichannel retailer selling multiple substitutable products through an online channel and a physical store. Online purchases can be returned either by mail or to the physical store. The retailer decides each product's price and refund value for each channel to maximize his expected profit.
We capture a consumer's sequential decisions on making a purchase and potentially returning her product using a generalized Markovian logit choice (MLC) model. We use this model to formulate the retailer's joint pricing and refund optimization problem. If there are constraints on prices and refund values, this problem may be non-convex, and we approximate it using an MILP. Furthermore, we propose a new method to derive the theoretical performance accuracy of this MILP. Otherwise, this problem is convex, and we derive its optimal analytical solution. We estimate the generalized MLC model using transaction data, making our framework applicable to a data-driven setting. Numerical experiments using synthetic data demonstrate that our estimation-and-optimization framework based on the generalized MLC model fits a general data set well and produces competitive results. A case study using data with some unobservable consumer choices from a major fashion retailer demonstrates the flexibility of our framework in incorporating various refund policies.
We find that to benefit from return policies, the retailer needs to guarantee good quality or non-negative online or offline return social welfare for each product. Products with a smaller loss coefficient exhibit a higher return rate but a larger profit.

4 - Can Priced Sample Benefit Sellers and Consumers? How to Signal Quality Through Pricing Samples

Yitong LIU, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Linqiu Li, Yun Fong Lim

Reducing consumers' uncertainty about product quality before purchase is crucial for sellers. This paper examines the priced sample strategy as a quality-informing mechanism, where sellers set a sample price to influence consumer beliefs and purchasing decisions. Using a singnaling game model, we analyze two types of sellers—high-type and low-type—operating under asymmetric information. We identify a dynamic belief updating mechanism: an **ex ante** signaling mechanism, where the high-type seller sets a lower sample price to separate from the low-type seller, and an **ex post** mechanism, where the high-type seller sets a higher sample price to encourage trial consumption and maximize full-size product sales.

Our findings show that when the full-size product margin is high, the seller prefers the **ex post** mechanism to drive sample purchases and repeat sales. However, when the quality gap is large and consumer expectations are low, high-quality seller perfers the separating equilibrium. Notably, a high-quality advantage does not always translate to higher sample price or increased profit. Moreover, compared to the "free-or-no" sample strategy, the priced sample strategy can benefit all parties, achieving a win-win-win outcome for high-type seller, low-type seller, and consumers.

Our study provides managerial insights into optimizing sample pricing to reduce consumer uncertainty. We highlight future research directions, including integrating coupon-based promotions, leveraging advertising and live-streaming to expand consumer reach, and empirically validating our theoretical conclusions.

MB12

Moor Room

New Advancements in Revenue and Supply Chain Management

Invited Session

Revenue Management and Pricing

Chair: Jinzhi Bu, The Hong Kong Polytechnic University, Hong Kong, N/A

Co-Chair: Xingyu Bai, The Hong Kong Polytechnic University, Hong Kong, N/A, Hong Kong

1 - Onshoring or Offshoring? the Role of Government Incentives in Shaping Manufacturing Strategies

Lingyun Chen, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, Xiaomeng GUO, Guang Xiao

Onshoring manufacturing has become a significant concern for many countries, prompting the implementation of various government incentives, including subsidies and tariffs, to accelerate this trend. Motivated by this, this paper investigates the manufacturing decision of global manufacturers between onshoring and offshoring under government incentives. Our findings suggest that both product valuation and competition intensity significantly influence manufacturing decisions under government incentives. Our analysis further reveals that increased government incentives can drive more firms to produce onshore, but this may occur at the expense of firms' profits and may not always benefit consumers. Moreover, compared to a tariff hike, increasing the subsidy is a more consumer-friendly approach, as it never harms consumer surplus.

2 - Spatial Matching with Delays

Yufei Wang, Shanghai University of Finance and Economics, Shanghai, China, People's Republic of

We study a centralized dynamic matching market with spatial operations, where requests arrive stochastically and wait to be matched. Notable examples are carpooling platforms, where passengers share a ride to similar destinations; food delivery platforms, where a driver delivers multiple orders from the same restaurant; and video game matchmaking, which matches players based on skill and playstyle. In such applications, delaying real-time matching increases pending requests, which improves match quality but also increases wait times, impacting user satisfaction. Thus, balancing the trade-off between matching quality and delay is crucial for efficient platform operations. In this paper, we investigate this trade-off and design matching policies that demonstrate strong performance in terms of competitive ratios. With an unknown arrival rate, we propose a greedy policy achieving a constant competitive ratio. With a known arrival rate, we propose the radius, batching, and partition policies, all with strong performance guarantees. Numerical results further validate the effectiveness of the proposed policies and provide valuable managerial insights.

3 - Ride-Hailing Platform Coopetition Under Fairness-Aware Pricing Regulation

Jiazi Yang, School of Business Administration, South China University of Technology, Guangzhou, China, People's Republic of

Ride-hailing aggregators have emerged to integrate multiple ride-hailing companies onto single platforms, allowing customers to choose from various providers. Small companies often rely on these platforms to connect with customers, while large companies may benefit from expanded market share by partnering with aggregators. However, such partnerships could disrupt existing customer bases and intensify price competition. This raises important questions: Should large companies join these platforms? What are the effects on small companies, aggregators, consumers, and society?In this study, we apply a game-theoretical framework to analyze a market involving large and small companies, aggregators, and loyal and disloyal customers, both in non-joint and joint settings, with and without price regulation. Our analysis reveals that without regulation, large companies use price discrimination by offering lower prices on their own app and higher prices on aggregators' platforms, intensifying price competition and benefiting only the large company. To mitigate these negative effects, we propose a fairness-aware pricing regulation, which imposes a price difference constraint across the large company's app and aggregator platform. With sufficient regulation, even with potentially higher prices, all players benefit due to reduced competition and enhanced driver pooling. Our findings challenge the conventional belief that higher prices harm consumers, showing that total consumer surplus can improve under such

regulation. The results suggest that large companies should partner with aggregators, while small companies, aggregators, and customers stand to benefit from advocating for stringent fairness-aware pricing regulations.

4 - Asymptotic Optimality of Semi-Open-Loop Policies in Inventory Models with Stochastic Lead TIMES

Xingyu Bai, The Hong Kong Polytechnic University, Hong Kong, Hong Kong

Inventory models with large and uncertain lead times are notoriously difficult to manage due to the curse of dimensionality. Recent works suggest that in inventory models with large deterministic lead times, semi-open-loop policies are asymptotically optimal. In this paper, we provide a theoretical foundation for the superior performance of semi-open-loop policies in inventory models where the lead times are not only large but also exhibit high variability. In the single-sourcing lost-sales inventory model with divisible products, we show that the optimality gap of constant-order policies decays exponentially fast as the lead time increases. In the single-sourcing lost-sales inventory model with indivisible products, under the assumption that the placed orders cannot cross in time, we propose a bracket policy, which alternates deterministically between two consecutive integer order quantities, and prove that the bracket policy is asymptotically optimal. In the dual-sourcing backlog inventory model with divisible products, we show that a semi-open-loop policy, which places a constant order from the regular supplier in each period, and implements a state-dependent modified base-stock policy from the emergency supplier, is asymptotically optimal, and we also extend our analysis to the joint pricing and inventory model. Finally, we provide a comprehensive numerical study to demonstrate the good performance of the proposed policies, and derive further managerial insights.

MB13

Bailey

Stochastic Optimization and Artificial Intelligence

Invited Session

Simulation and Optimization of Healthcare Systems

Chair: Hao Cao, Fudan University, Fudan University, 220 Handan Road., Shanghai, 200433

1 - A Kernel-Based Stochastic Approximation Framework for Contextual Optimization

Hao Cao, Fudan University, Shanghai, China, People's Republic of, Jianqiang Hu, Jiaqiao Hu

We present a kernel-based stochastic approximation (KBSA) framework for solving contextual stochastic optimization problems with differentiable objective functions. The framework only relies on system output estimates and can be applied to address a large class of contextual measures, including conditional expectations, conditional quantiles, CoVaR, and conditional expected shortfalls. Under appropriate conditions, we show the strong convergence of KBSA and characterize its finite-time performance in terms of bounds on the mean squared errors of the sequences of iterates produced. In addition, we discuss variants of the framework, including a version based on high-order kernels for further enhancing the convergence rate of the method and an extension of KBSA for handling contextual measures involving multiple conditioning events. Simulation experiments are also carried out to illustrate the framework.

2 - Exploring Artificial Intelligence's Role in Improving Supply Chain Efficiency: a Study on Key Mechanisms in the Manufacturing Sector

Wenjia Li, Harbin Institute of Technology, Harbin, China, People's Republic of, Guangxin Jiang

This study investigates the role of Artificial Intelligence (AI) in improving supply chain efficiency in the manufacturing sector. By analyzing corporate annual reports, we examine AI applications across critical supply chain functions, including demand forecasting, procurement, and inventory management. Using a combination of Word2Vec-based keyword expansion and DeepSeek sentiment scoring, we quantify the breadth of AI adoption, its targeted supply chain processes, and the perception of its impact. The results indicate that AI technologies are most frequently applied in demand forecasting and inventory management, with sentiment analysis suggesting a predominantly positive outlook on AI’s contributions. Regression analysis reveals a strong positive correlation between AI adoption metrics—such as technology diversity and supply chain coverage—and firm performance indicators, including revenue growth and profit margins. This research advances the literature by integrating advanced textual and sentiment analysis tools to provide empirical evidence of AI’s transformative potential in supply chain optimization. It also offers actionable insights for businesses striving to

enhance operational efficiency through AI technologies. Future research should address the limitations of relying on corporate disclosures and explore the long-term effects of AI adoption on supply chain sustainability and resilience.

3 - A Structural Forest Estimation of Heterogeneous Newsvendor Models

Mengfei Li, Fudan University, Shanghai, China, People's Republic of, Guihua Wang, Wenchang Zhang, Xiaole Wu

We develop a causal forest method to perform structural estimation for newsvendor models with large heterogeneities. We demonstrate the statistical properties of this method and validate its accuracy through simulation experiments. Applying this method to healthcare data, we analyze surgeons' decision-making in scheduling operating rooms.

4 - Achieving Robust Data-Driven Contextual Decision Making in A Data Augmentation Way

Zhaoen Li, Tsinghua University, Beijing, China, People's Republic of, Zhi-Hai Zhang, Maoqi Li

This paper focuses on the contextual optimization (CO) problem, where the decision is made subject to uncertain parameters, and covariates can be exploited to reduce the uncertainty. We train a robust decision rule (a function of the covariate) for the problem through a distributionally robust optimization (DRO) model, which maximizes the worst-case expected cost over an uncertainty set including all distributions closed enough to a nominal one with the Wasserstein distance. To tackle the computational challenge, a stochastic gradient descent algorithm is developed in a data augmentation way, which iteratively a) bootstraps a sample from the nominal distribution; b) perturbs the sample adversarially according to the DRO model; and c) updates the decision rule. The algorithm is both efficient and effective. From the efficiency aspect, the algorithm saves significant computational time and is not restricted by the specification forms of nominal distributions like other solving approaches. As such, we can apply the DRO model to more general settings like online learning. From the effectiveness perspective, we prove that the algorithm converges to the optimal solution of the DRO model at a rate of $O(1/\sqrt{T})$, where T denotes the number of iterations. Consequently, the performance guarantee of the algorithm to the clairvoyant optimal is $O(1/\sqrt{T})$ plus that of the DRO model, which is $O(1/\sqrt{N})$ when the nominal distribution is empirical with N samples and $O(1/\sqrt{T})$ in the online settings where one sample is obtained in one iteration. Extensive numerical experiments demonstrate the algorithm's superior performance to several benchmarks.

5 - Enhancing Nutritional Tracking: Robust Goal Programming with Patients' Dietary Imagery

Haiyan Yu, Chongqing University of Posts and Telecommunications, Chongqing, China, People's Republic of, Kun Zhang, Senlin Li, Renying Xu, Zhiqi Chen, Ehsan Hajizadeh, Tao Huang, Keyan Chen

Traditional dietary tracking methods—manual logs, recalls, questionnaires—are error-prone and unscalable, undermining personalized healthcare. We introduce *NutriOpt*, an AI-driven framework that transforms meal images into actionable nutritional plans. Leveraging a U-Net model, *NutriOpt* analyzes meal images to identify foods, estimate portions, and compute nutrient profiles with over 90% accuracy. We then bridge computer vision with optimization: a novel goal programming model balances personalized nutrient targets (e.g., protein) against caloric needs, while a robust optimization layer mitigates real-world uncertainties like food variability and measurement noise. Validated on real-world datasets, *NutriOpt* achieves over 10% higher precision than existing tools, reduces tracking time by 50%, and adapts seamlessly to dietary noise. This work pioneers a scalable synergy of deep learning and decision science, offering clinicians and patients a dynamic tool for precision nutrition. Results underscore its potential to reshape chronic disease management, public health strategies, and AI-driven healthcare optimization.

MB14

Wilberforce

Simulation-Driven and AI-Based Optimization Models in Advanced Manufacturing

Invited Session

Semiconductor Manufacturing x AI

Chair: Jr-Fong Dang, National Taiwan University of Science and Technology, Taipei, N/A, Taiwan

1 - Integrating Machine Learning with Stochastic Optimization for Joint Predictive Maintenance and Job Scheduling

Tzu-Li Chen, National Tsing Hua University, Hsinchu, Taiwan

In modern manufacturing, machine health plays a crucial role in determining product yield. Fluctuations in machine health can lead to significant yield losses, which in turn affect production efficiency and order accuracy. However, traditional scheduling methods struggle to cope with these uncertainties, especially in large-scale production environments. This study investigates the issue of yield loss uncertainty caused by machine health conditions. We developed a two-stage contextual stochastic optimization scheduling model, utilizing historical machine condition data to predict yield losses. The aim is to improve accuracy through the use of contextual information, which is then applied to optimize subsequent scheduling and minimize order delay losses. To solve this complex model, we designed a framework that combines residual-based sample average approximation (SAA) and heuristic algorithms. The method integrates data-driven machine physical state information, performing sample average approximation on several random variable samples to estimate the expected value. Then, based on the differences between predicted and observed values, residuals are adjusted to improve the accuracy and robustness of the SAA method. The algorithm then can solve the complex problem.To validate the effectiveness of the model, we plan to conduct experiments using real production data to accurately capture the impact of machine health on yield. We expect the model to significantly improve the accuracy of yield predictions. Finally, the model will be compared with traditional deterministic scheduling methods.

2 - Calibration of Simulator Parameters Using Sampling, Machine Learning, and Optimization

I-Hsuan Hong, National Taiwan University, Taipei, Taiwan, Chiao-Yu Chen, Rudi Nurdiansyah, Jack C. P. Su

Calibration of simulator parameters is essential for improving model accuracy in complex systems, particularly when simulations are computationally intensive. This paper proposes a general framework that integrates sequential D-optimal sampling, surrogate modeling with machine learning, and optimization to efficiently calibrate simulator parameters with reduced computational cost. The framework strategically selects informative samples to construct accurate surrogate models, significantly lowering the number of required simulations. Once the surrogate model is built, optimization algorithms are applied to identify parameter settings that align simulation outputs with target responses. The effectiveness of the proposed approach is demonstrated through a high-cost simulation case, where the framework achieves a comparable prediction error (13.40%) using only 27 samples, compared to 83 samples required in a previous study to reach a similar error of 12.04%. Although the case study involves marine turbine simulations, the methodology is broadly applicable to a wide range of calibration problems in science and engineering. By combining optimal experimental design, machine learning, and optimization, this framework enables scalable and data-efficient parameter tuning, making it especially valuable for simulation-based studies with limited computational resources.

3 - Pareto Set Approximation Using Probabilistic Branch and Bound for Simulation Optimization

Hao Huang, Yuan Ze University, Taoyuan, Taiwan, Shing Chih Tsai, Chuljin Park

This work addresses the challenge of identifying the Pareto optimal set in discrete multi-objective simulation optimization problems, where objective functions are only accessible via stochastic simulation and the solution space may be extremely large. We propose a novel Discrete Multi-Objective Probabilistic Branch-and-Bound algorithm that systematically explores the solution space by partitioning it into subregions. The algorithm allocates simulation effort adaptively, favoring non-dominated subregions through deeper branching and increased evaluations, while pruning dominated areas. We establish theoretical convergence properties of the algorithm and validate its performance through numerical experiments. Results demonstrate the algorithm’s potential in approximating the Pareto set in complex discrete simulation environments.

4 - The Intelligent Defect Detection Framework Integrating Object Detection Algorithm with Large Language Model

Jr-Fong Dang, National Taiwan University of Science and Technology, Taipei, Taiwan

As the demand for mass production continues to rise, traditional manual defect inspection methods have become increasingly labor-intensive and prone to errors due to operator fatigue. To address these challenges, this study proposes an enhanced defect detection approach based on a modified YOLO architecture, aiming to improve inspection accuracy and reduce operational costs. Furthermore, a large language model (LLM) is integrated into the framework to facilitate in-depth root cause analysis, providing valuable insights into the origins of defects and supporting proactive quality control. To validate the effectiveness of the proposed algorithm, real-world manufacturing data from industrial rubber liquefied petroleum gas (LPG) hoses are utilized. The experimental results demonstrate the high capability of the proposed model, achieving an accuracy rate of 99.5%.

5 - Intelligent Defect Detection Algorithm for Gear Inspection

YUAN WANG, National Taiwan University of Science and Technology, Taipei, Taiwan

Gears are vital in mechanical systems, yet side-surface defects, especially in powder metallurgy (PM) small gears are remain challenging due to porosity and misalignment. We propose an intelligent detection algorithm combining a precise image acquisition strategy with deep learning. A custom positioning mechanism aligns gear side views, filtering out misaligned frames and reducing noise. YOLO-based models trained on positioned images converge more rapidly and achieve significantly higher mAP, precision, and recall across multiple variants. This method also halves annotation effort by focusing on well-aligned frames and critical tooth regions. To support statistical analysis, we introduce a Defect Map that visualizes defect distributions, enabling real-time process control and long-term manufacturing optimization. Experimental results demonstrate enhanced accuracy, consistency, and reliability, establishing a data-driven foundation for next-generation gear inspection workflows.

MB15

Fairmont Ballroom

Vehicle Routing and Last-Mile Delivery Innovations

Contributed Session

Transportation Science and Logistics

Chair: ShiZeng Zhao, The Australian National University, Canberra, Australia

1 - Hybrid Matheuristics for Mobile Parcel Locker Delivery with Robots and Drone Resupply

Emrah Demir, Cardiff Business School, Cardiff University, Cardiff, United Kingdom, Cheng Chen

This research introduces a novel vehicle routing problem with time windows in an autonomous last mile delivery system featuring mobile parcel lockers (MPLs), delivery robots, and drone resupply. Customers can choose between home delivery via robots or self-pickup from MPLs stationed at designated parking areas. Delivery robots operate from MPLs, while drones replenish these lockers as needed. We define this problem as the Mobile Parcel Locker Problem with Delivery Robot and Drone Resupply (MPLPDR-DR). To address this problem, we propose a mixed-integer linear programming (MILP) model and develop a hybrid matheuristic approach. Our method integrates a metaheuristic for optimising MPL and robot routing, while a MILP model determines optimal drone resupply operations. The metaheuristic is based on the artificial bee colony framework, incorporating large neighbourhood search, variable neighbourhood descent, and mutation mechanisms to enhance solution quality. Extensive computational experiments demonstrate the efficiency of the proposed approach in solving large-scale MPLPDR-DR instances. The findings provide key managerial insights into optimising autonomous delivery operations, improving efficiency, and reducing last mile logistics costs.

2 - Same-Day Delivery via a Truck-Drone Tandem with Drone Resupply

Shuanglong Liu, Southeast University, Nanjing, China, People's Republic of, Hui Sun

Implementing drones in the context of city logistics has recently received intensive attentions from both practitioners and researchers. With the intention of improving the efficiency of same-day delivery, we propose a new parcel delivery system, in which a truck-drone tandem is employed for home delivery and a drone is used to

resupply newly available parcels from the depot to the truck on its tour. The routes of the delivery vehicles (i.e., the truck and drone of the tandem) and the schedule of drone resupply need to be determined, so that the makespan to complete all deliveries is minimized. A mixed integer programming model (MIP) is presented to describe this flying sidekick Traveling Salesman Problem with drone resupply (FSTSP-DR). In order to find an optimal or near-optimal solution rapidly, we develop a matheuristic approach. This approach iteratively searches for the routes of the delivery vehicles using a metaheuristic algorithm, and determines the optimal resupply schedule for each set of routes found by solving a mathematical programming model. Computational experiments are conducted to test the effectiveness and efficiency of the developed matheuristic. Furthermore, the performance of the new system is compared with that of the delivery system discussed in the literature, where a single truck is used for delivery and a drone is utilized for resupply. Finally, operational insights and guidance are given for the application of the proposed system.

3 - Integer Programming Models for One Commodity Pickup-and-Delivery Flying Sidekick Traveling Salesperson Problem

Soundarya R, Indian Institute of Science, Bengaluru, India, India, Tarun Rambha

Truck-based deliveries face challenges like urban traffic, high fuel costs, and customer expectations for fast deliveries. Drone integration facilitates quick deliveries, traffic avoidance, and access to remote areas. This collaborative truck-drone delivery approach has been studied in the literature as the "Flying Sidekick Traveling Salesman Problem (FSTSP)". In more practical scenarios, logistics operations include pickups and deliveries. To address this feature, we propose a new formulation called the "One Commodity Pickup-and-Delivery Flying Sidekick Traveling Salesman Problem (1-PD-FSTSP)." In our model, the products of the pickup nodes are delivered to the delivery nodes. The goal is to minimize the total completion time while synchronizing the arrival times for both the truck and the drone. The drone is restricted to making deliveries and may wait at customer locations to coordinate with the truck. In compliance with drone regulations, the maximum flight time for a drone is capped. We developed a Mixed Integer Linear Programming model for the 1-PD-FSTSP and utilized a Branch-and-Cut method to find the optimal solution. We conducted tests on randomly generated 15-customer datasets, and the results indicate that employing a drone in cooperation with a truck reduces completion time by 12.02%.

4 - An Exact Method for Vehicle Routing Problem with Lockers and Drones

Huan Tang, Northwestern Polytechnical University, Xi’an, China, People's Republic of, Yuan YUAN, Jiliu Li

This study investigates a truck-locker-drone delivery system. In this system, trucks transport parcels from a warehouse either directly to customers or to transshipment facilities, where lockers and drones are deployed for last-mile delivery. Due to safety concerns, drone operations require human supervision, and each operator can monitor only a limited number of drones. The problem is modeled as a Vehicle Routing Problem with Lockers and Drones (VRP-LD). To address this problem, we propose a mixed integer programming formulation and a logic-based benders decomposition approach. The master problem is a routing-allocation model and determines the assignment of parcels to vehicles and transshipment facilities, then the vehicle routes delivering these parcels. It considers the overall battery availability of drones at each facility without tracking the battery level of individual drones. The subproblems handle drone assignment for transshipment deliveries, explicitly accounting for drone battery capacity constraints. To improve the solution process, several infeasibility and optimality benders cuts are incorporated into the algorithm. Additionally, an initial heuristic is designed to provide a good upper bound, and valid inequalities are introduced to improve the efficiency of solving the master problem. Computational results on a set of newly generated instances demonstrate the effectiveness of proposed methods.

5 - Llm-Driven Deep Reinforcement Learning for Solving Heterogeneous Capacitated Vehicle Routing Problem in Complex Environments

Yongwu LIU, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, Binglei Xie, Gangyan Xu

The heterogeneous capacitated vehicle routing problem (HCVRP) is a critical research topic in logistics, focusing on leveraging vehicles with diverse characteristics to serve dynamic demand while optimizing transportation efficiency and reduce fleet costs, etc. When solving the HCVRP, traditional operations research (OR) methods, such as exact and heuristic algorithms, have limitations when addressing large-scale problems or dynamic scenarios. Deep reinforcement learning (DRL) overcomes these limitations by learning adaptive

policies through environmental interactions, efficiently managing high-dimensional states and constraints. However, designing effective DRL reward functions remains challenging, as manually aligning rewards with evolving task objectives and dynamic environments risks suboptimal exploration or biased policy convergence. To solve these problems, we propose a DRL framework leveraging large language model (LLM) to design adaptive reward functions for HCVRP. Unlike coventional manual reward engineering that struggles with dynamic fleet configurations and evolving parameters, LLM can enable automated reward shaping through their inherent reasoning capabilities, contextual pattern recognition, and application of problem-specific rules. This approach dynamically adjusts rewards by interpreting real-time operational states and complex environmental factors, without requiring exhaustive manual reward specifications. The framework further integrates an attention-based mechanism to coordinate heterogeneous vehicle selection and node allocation, where LLM-designed reward functions guide the learning process toward operationally efficient solutions. Experiments on benchmark datasets demonstrate superior performance over existing approaches, with ablation studies confirming LLM-designed rewards' critical role in accelerating DRL training and improving solution quality. Analysis of different LLM reveals their varying impacts on reward function efficacy, providing practical insights for model selection.

6 - Dynamic Uav Fleet Arc Routing Optimization for Commercial Promotions Integrating Traffic Flow and No-Fly Zones

ShiZeng Zhao, The Australian National University, Canberra, Australia, Yongxi Cheng, Zilu Zhao

Unmanned Aerial Vehicles (UAVs) show potential for commercial promotion over urban areas, yet research in this domain remains limited. This paper explores the arc routing path optimization problem of urban advertising delivery using UAVs, and combines real-world constraints such as heterogeneous UAV fleets, no-fly zones, time windows, and differences in advertising value on different streets. Different from the traditional arc routing path problem, this study constructs an arc routing model for heterogeneous UAV fleets that considers dynamic population density and no-fly zone avoidance. By introducing a dynamic advertising value function, the advertising value of urban streets that changes over time and space is incorporated into the path planning objective, and energy consumption and flight costs are jointly minimized, thereby maximizing advertising revenue. In addition, this paper innovatively uses traffic flow data to characterize spatiotemporal variations in population density. A hybrid heuristic algorithm is proposed, demonstrating both effectiveness and computational efficiency in large-scale case studies. The algorithm's capability to optimize promotional coverage while adhering to operational constraints in complex urban environments is systematically validated, which provides theoretical contributions and practical significance for smart city aerial marketing.

MB16

VIP A

Sustainable Operations

Invited Session

Sustainable and Socially Responsible Operations

Chair: Guang Xiao, The Hong Kong Polytechnic University, Kowloon, N/A

1 - Horizontally Differentiated Product Design Strategies Under Extended Producer Responsibility

Yuxiao Chen, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Chengzhang Li, Qinghua Zhu

Horizontal product differentiation has been widely adopted to cater to increasingly diverse consumer preferences. However, the increased variety of products complicates the recycling process. Under the extended producer responsibility (EPR) regulations, how brand manufacturers should incorporate recycling costs into the configuration of horizontally differentiated products remains unclear. We first develop a stylized model to examine a firm's multi-product design strategies under Extended Producer Responsibility (EPR), where products are differentiated in both price and position, and the per-unit recycling cost depends on the distance between product positions and a recycling cost coefficient. Our analysis reveals the critical roles of the consumer's sensitivity to product distance and the recycling cost coefficient in shaping product design strategies. We next

investigate a decentralized system involving two competing manufacturers and one recycler responsible for product recycling. The firms' product design strategies exhibit a lower likelihood of achieving entire market coverage and are more likely to lead to market segmentation. The competition effect yields lower system profits but results in a lower environmental impact. We adopt a vertical integration strategy and design a revenue-sharing contract to coordinate the system, respectively. While the former improves the system profits at the cost of additional environmental impacts, the latter significantly improves system profits with only a marginal increase in environmental impacts. We analyze extensions with multiple attributes and asymmetric production costs to test the robustness of the main insights. Our study underscores the importance of considering recycling costs in shaping product design strategies under EPR regulations.

2 - From Trade-Offs to Trade: the Effects of Framing and Social Norms on Waste Supply Chain Coordination

XiaoLi Zhang, FuZhou University , FuZhou, China, People's Republic of

The world generates substantial waste that could serve as feedstock, but often inefficiencies in waste exchange arise, caused by a lack of coordination among companies. When trading waste reuse, negotiators must manage the trade-off between supply chain, personal, and environmental outcomes. Unlike studies focusing on monetary incentives(e.g., taxation, subsidies), we aim to identify how nudge techniques impactnegotiation agreements in waste supply chains. We specifically study the effect of various (environmental and economic) frames and (collective and individual) social norms on agreement rates. Additionally, we open the "black box" of the negotiation process by revealing the underlying mechanism from the offer proposer and receiver, respectively. We analyze supply chain coordination in a set of stylized bargaining experiments characterized by conflicting economic and environmental interests. We design different treatments in two studies, varying the environmental and economic frames in Study 1, and the presence of social norms in Study 2. We find that framing environmental externalities using a damage frame results in higher agreement rates compared to a help frame, and framing economic outcomes as an additional cost is equally
effective as a foregone profit frame. However, providing negotiators with individual social norms appears to reduce agreement rates. The receivers' acceptance decision primarily drives these agreement changes, with environmental concerns significantly influencing both proposers' offering and receivers' acceptance decisions. Our findings can guide policy makers to improve inter-firm collaboration in the presence of negative environmental externalities by using frames to nudge supply chain partners.

3 - Impact of Government Subsidy on Green Product Innovation

Lingyun Chen, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, Xiaomeng GUO, Dan Tan, Guang Xiao

This paper captures two key incentives for green technology innovation: subsidy and consumers' willingness to pay a premium. It studies how these incentives affect the firm's innovation decision to identify the optimal subsidy policy for the welfare-maximizing government. Our results offer the sufficient condition for the firm to innovate and reveal the differentiating effect of subsidy policies for various recipients on innovation. In addition, we find the optimal subsidy policy and demonstrate that there exists a portfolio of subsidy policies that can yield greater social welfare than a single subsidy policy.

4 - Recycled Label Design and Green Competition

Xiaoli Zhang, Hong Kong Polytechnic University, Kowloon, Hong Kong, Xiaomeng GUO, Shi Duo, Guang Xiao

Facing eco-conscious markets, a growing number of firms are using recycled materials in their products. The most effective way to credibly convey such green efforts may be getting labeled by a third-party certifier. Across industries, we observe two typical formats for recycled labels: A continuous label precisely displays the percentage of recycled materials in a product. A binary label sets a minimum percentage standard and is issued to a product if the standard is met. This paper investigates how a certifier should design the recycled label to boost an industry's environmental performance. We capture the context with a game-theoretic model comprising a certifier and a duopoly. The firms determine product composition through recycled-technology investments and thereby compete for market share. We focus on three key industry characteristics: competition intensity, average efficiency to invest in recycled technology, and consumer green awareness. We uncover a sandwich principle for

label selection: Considering any characteristic, the certifier should design the recycled label as continuous if its magnitude is medium and as binary if its magnitude is either low or high. We further demonstrate that the certifier's label selection cannot align with both firms' preferences, but can sometimes align with the industry's preference as a whole. Sensitivity analyses reveal that environmental performance and industry profitability change nontrivially along the three metrics. In particular, neither more (less) intense green competition nor more efficient recycled investment guarantees better environmental performance (higher industry profitability).

MB17

VIP B

Digital Marketing and Consumer Engagement in Social Media

Contributed Session

eBusiness and Social Media

Chair: Jingwen Yan, Xi'an Jiaotong University, Xi'an, N/A

1 - Facial and Textual Cues in Social Media Video Cover Images: Implications for User Engagement

ping huang, Xi'an Jiaotong University, xi'an, China, People's Republic of, xinyu sun

The proliferation of social media platforms has heightened the strategic importance of cover images in shaping user engagement. Despite their ubiquity, the interactive effects of facial and textual elements within social media video cover images remain understudied. Drawing on theories of visual salience, visual harmony, and informational value, this research investigates how facial characteristics (e.g., attractiveness, smile intensity) and textual features (e.g., sentiment) influence user engagement, while examining moderation by content type (entertainment vs. informational) and creator follower size.Leveraging a large-scale dataset and employing rigorous methodologies, including DiNardo-Fortin-Lemieux reweighting and Generalized Propensity Score matching, this study presents three key findings. First, facial cues attract intuitive perceptual attention, whereas textual cues provide informational clarity, collectively enhancing user engagement through balanced visual harmony. Second, content type and creator popularity significantly moderate these effects: facial cues more effectively drive engagement in entertainment content, while textual cues bolster credibility and participation in informational content. Smaller creators particularly benefit from visually appealing facial cues, while larger creators leverage textual elements to strengthen authority. Third, critical visual design characteristics, such as facial attractiveness, smile intensity, and textual sentiment, are identified as essential factors in optimizing visual harmony and increasing emotional resonance.This research enriches theoretical understanding by integrating visual salience, harmony, and informational value theories within social media contexts, offering a nuanced comprehension of multimodal cue effectiveness. Practically, the findings provide actionable guidance for content creators and platform designers to strategically design cover images, enhancing user engagement in digital communication ecosystems.

2 - What Affects Consumer Spending - Negative Press OR Social Media Coverage?

Tae Lee, Bayes Business School (UK), London, United Kingdom, Pawel Bilinski, Jay Jung, Caroline Wiertz

Does the negative sentiment of press articles or of social media posts influence consumer purchasing behavior? Utilizing over 2.7 million daily credit and debit card transactions for Starbucks' stores and hand-collected press news and Twitter activity, we find that social media sentiment on the day, not sentiment of press articles on that day, affects consumer spending. This result is explained in two ways. First, even when the average sentiment across all press articles is negative, there remains a substantial volume of positive sentiment articles. This makes press signals ambiguous and costly for consumers to process. In accordance with the dual processing theory, individuals avoid modifying their behavior when confronted with elevated information acquisition and processing costs to resolve signal dissonance. In contrast, social media discourse is readily understandable, prompting more pronounced behavioral shifts in response to negative sentiment conversations. Second, Starbucks' consumers are, on average, more attuned to social media communication than press news, which suggests they tend to acquire information that shapes their spending habits on social media.

3 - Visual Narratives in E-Commerce: How Image Type Shapes Review Informativeness and Consumer Trust

Prabin Panigrahi, Indian Institute of Management Indore, Indore, India, Abhishek Kumar Jha

Online product reviews constitute a critical component of the e-commerce ecosystem, significantly influencing consumer purchase decisions and brand evaluations. While prior research has predominantly focused on textual elements within user-generated content, the abundance of image-enriched reviews necessitates a deeper investigation into the role of visual content. Existing literature has primarily examined the mere presence of images, often overlooking the nuanced effects of image content and type. Drawing from media studies, which suggest that images serve both aesthetic and informational purposes, this study explores the differential impact of *process-oriented* images and *outcome-oriented* images on perceived review informativeness and helpfulness.

Utilizing a machine learning-based classifier with over 85% accuracy, we analyse a large-scale dataset of product reviews from Amazon.com. Our empirical findings reveal that process-oriented images significantly enhance review helpfulness compared to outcome-oriented images ($p = 0.016$), likely due to their immersive and experiential nature, which increases perceived authenticity and credibility. To further validate these findings, we propose a series of controlled experiments designed to isolate the effects of image type while accounting for potential confounding variables. These experiments will assess both the aesthetic and informational value of images, with a particular emphasis on the latter's influence on consumer trust and decision-making.

This study contributes to the literature by advancing a conceptual framework for image-based UGC that distinguishes image types based on their functional attributes. From a managerial perspective, our findings offer actionable insights for e-commerce platforms seeking to optimize visual review strategies to enhance consumer engagement and trust.

4 - Single OR Collage? Visual Complexity Decisions in Social Media Cover Design

Jingwen Yan, Xi'an Jiaotong University, Xian, China, People's Republic of

On social media, the **cover image** plays a crucial role in capturing user attention and driving engagement when users scroll the content feed. Influencers often face a key visual decision when designing a cover: whether to use a single image or a collage composed of multiple images. This choice reflects not only aesthetic preferences but also strategic considerations influenced by content types and cover characteristics. This study investigates the decision-making process behind visual complexity in cover design, focusing on the trade-offs between simplicity and information richness. Grounded in **Visual Information Processing Theory**, the research hypothesizes that collage may lead to higher complexity as well as information richness.To explore how different levels of image complexity affect users' attention, cognitive load, and likelihood to engage, the research collects real content and influencer data from **Xiaohongshu**, one of the largest social media platforms in China, and conduct an empirical study to address the issue. The research aims to provide theoretical insights and practical implications for influencers seeking to optimize their visual strategies and content effectiveness.

Monday, July 21, 1:45 PM - 2:45 PM

Fairmont Ballroom

From Theory to Impact: Building Data-Driven Models for Real-World Optimization

Invited Session

Keynote

1 - From Theory to Impact: Building Data-Driven Models for Real-World Optimization

Teo Chung Piaw, National University of Singapore, Singapore, Singapore

This keynote reflects on a three-stage journey—from foundational work in optimization theory, to building tractable models for complex systems, and finally to embracing data-driven thinking in uncertain, high-stakes environments. At each stage, we explore how theory can be translated into decision tools that deliver real-world impact across sectors like **aviation**, **gaming**, and **supply chains**.

Along the way, we've accumulated a collection of **"war stories"** from the frontlines of applied optimization. In **airline revenue management**, we are developing solutions to tackle real-world challenges—from pricing ancillary products like luggage based on customer features, to designing dynamic fare ladders, steering demand across unbalanced networks, and quantifying the revenue impact of customer experience. Many of these themes will be featured in the sessions under **"Recent Advances in Airline Revenue Management"**, held on **Sunday, July 20 (SC11 and SD11, Morrison Room)**.

Beyond aviation, we've ventured into the complex world of **lottery and gaming**, where analytics must do double duty: boosting operator margins combating the illegals. In these domains, optimization becomes a balancing act between excitement and exposure. In **logistics and public health**, we've seen robust models put to the test—**designing parcel locker networks that keep pace with surging e-commerce**, and **orchestrating nationwide vaccine rollouts under shifting supply constraints**, where every misstep meant delays for millions.

Each story illustrates a simple truth: **rigorous models are only powerful when stress-tested in the messiness of the real world**. These experiences continue to shape how we think about building models that are not just elegant, but consequential.

Morrison Room

Semiconductor Manufacturing Development – past, present and future

Invited Session

Keynote

1 - Semiconductor Manufacturing Development &Ndash; past, Present and Future

Jonathan Chang, VisionPower Semiconductor Manufacturing Company Pte Ltd (VSMC), Tampines, Singapore

Semiconductor manufacturing is a complicated integrated process where even a single step—such as wafer testing involving man, machine, material, and method (4M)—can trigger multiple deviations, impacting key performance metrics like quality, speed, productivity, and cost. A completed wafer fabrication involves hundreds to thousands of steps will produce hundreds of deviations for overall process. Beyond these factors, the growing complexity of advanced process technologies (from 0.35μm to 5nm) has exponentially increased challenges in yield management, driven by interactions between device layers (e.g., OD–Poly–Metal) and production modules (lithography, etch, thin-film, diffusion, and CMP). A high-volume manufacturing facility must simultaneously manage hundreds of thousands of wafers, including production, R&D, and quality assurance lots—while maintaining competitive cycle times (CT) for customers and ensuring optimal machine productivity to control costs. Over past 30 years, the semiconductor foundry industry has demonstrated that effective deviation management is the core of manufacturing capability. This capability ultimately differentiates companies that sustain long-term competitiveness, enabling them to evolve into leading semiconductor technology providers and secure market leadership. This presentation will explore the evolution of semiconductor manufacturing capability, highlighting how industrial engineering (IE) and information technology (IT) establish the foundational capabilities required for world-class semiconductor production.

Hullet Room

Operations Research for Pandemic Preparedness: Lessons from COVID

Invited Session

Keynote

1 - Operations Research for Pandemic Preparedness: Lessons from COVID

Kelvin Bryan Tan, Ministry of Health, Singapore, Singapore

COVID was unprecedented in it's global impact of human lives and impact on healthcare strain, as well as the economic disruption as a result of mitigation measures. Operations Research played an important role in many of the critical aspects in the management of the pandemic by informing resource allocation, logistics and distribution, supply chain management, optimization of testing strategies and decision analysis. How successful

were we and what has been the contribution of Operations Research? More importantly, how can we do better for the next pandemic?

Monday, July 21, 2:45 PM - 3:15 PM

Moor/Morrison/Hullet Foyer

PM Poster Session

Poster Session

1 - Fintsbridge: a New Evaluation Suite for Real-World Financial Prediction with Advanced TIME Series Models

Yanlong Wang, shenzhen international graduate school, tsinghua university, shenzhen, China, People's Republic of

Despite studies have proposed various solutions to address the challenges encountered in time series prediction, aiming to improve forecasting performance. However, effectively applying these time series forecasting models to the financial asset pricing remains a challenging issue. There is still a need for a bridge to connect cutting-edge time series forecasting models with financial asset pricing. To bridge this gap, we have undertaken the following efforts: 1) We constructed three datasets from the financial domain; 2) We selected over ten time series forecasting models from recent studies and validated their performance in financial time series; 3) We developed new metrics, msIC and msIR, in addition to MSE and MAE, to showcase the time series correlation captured by the models; 4) We designed financial-specific tasks and assessed the application potential in important financial problems. We hope these work can provide insights into the effectiveness of advanced AI models in finanical domains.

2 - Enhancing Traveler Demand Forecasting Through Integrated Deep Graph Attention Networks

Shouchun Peng, Hitachi Ltd., Tokyo, Japan

Accurate traveler demand forecasting is essential for optimizing transportation systems and enabling informed urban planning. Graph Neural Networks have become widely used for demand forecasting considering the ability to model spatial dependencies across graph structures. However, most existing GNN-based approaches rely on simplistic graph structures and often overlook deep spatial features that play a critical role in real-world applications.
We propose a Deep Spatial-Temporal Graph Attention Network (DST-GAT) that addresses the challenges by integrating heterogeneous data sources into a unified attention-based architecture. Our model is designed based on a dual-stage attention mechanism that abstracts rich contextual information enables the network to learn local interactions and global latent feature relationships across graphs.
Evaluated on a urban dataset, DST-GAT achieves a MAPE of 9.8, outperforming a baseline GNN model that resulted in MAPE of 11.8. The ablation study demonstrates the effectiveness of capturing fine-grained spatial and semantic information for improving demand prediction

3 - Confidence-Aware Multimodal Learning for Trustworthy Fake News Detection

Shoumeng Ge, Harbin Institute of Technology, Harbin, China, People's Republic of, Ying Chen

With the rapid growth of social media, fake news has become widespread. Recent studies have focused on multimodal information for improved detection, but two issues remain: first, existing methods neglect the effectiveness of multimodal semantic information, leading to noise and irrelevant data; second, existing models lack reliability in output. Although post-hoc confidence calibration methods can be directly applied to existing models to improve the trustworthiness of the detection results, they lack theoretical guarantees for their performance. To address these, we propose a confidence-aware multimodal learning framework with two key contributions: a novel fusion module for effective multimodal integration and a post-hoc confidence calibration method with theoretical guarantees. Experimental results on three public datasets show that our framework outperforms existing models.

4 - Explainability Framework for Intelligent Risk Prediction of Heart Disease Using Multimodal Models for Clinical Decision Support in Low-Resource Health Systems

Ira Promasanti Rachmadewi, Monash University, Tangerang, Indonesia, Taufiq Asyhari

Heart disease remains the deadliest disease globally, with low- and Low-middle-income countries (LMICs) suffering an increasing burden shown by the rising prevalence of heart disease. Lack of cardiologists/specialists, underutilisation of electronic health records (EHRs), lack of trust in artificial-intelligence-based information systems (IS), and resource constraints become the predominant factors of the increasing burden of heart disease management problems. To overcome these challenges, an explainability framework of an intelligent heart disease risk prediction model that integrates multimodal EHR data, numerical, omics, images, and clinical notes, to enhance diagnostic accuracy and improve patient outcomes in resource-limited settings.

To strengthen clinical trust and usability, a Large Language Model (LLM)-based orchestrator enhances the explainability framework with human-AI interactivity, allowing clinicians to query the intelligent risk prediction and align output with clinical reasoning. This research contributes a novel architecture for explainability, AI-driven diagnostics that facilitate human-AI collaboration and support timely decision-making in under-resourced healthcare systems.

5 - How &Ldquo;AI Info&Rdquo; Labels Influence Social Media Posts: Evidence from Instagram

Yuanyuan Wu, Southwestern University of Finance and Economics, Chengdu, China, People's Republic of, Pengkun Wu, Yuanxi Liu

The rapid advancement of artificial intelligence (AI) is transforming content creation on social media. However, the impact of AI-generated content disclosure remains under explored. This study examines how the introduction of the "AI info" label influences content creators' posting behavior and user engagement on Instagram. Drawing on Signaling Theory, the research employs a difference-in-differences (DiD) methodology to analyze Instagram post data before and after the introduction of the "AI info" label. The findings reveal that the "AI Info" label leads to a reduction in content creation frequency, suggesting that AI disclosure may discourage creators from posting as often. Additionally, posts with "AI info" label receive fewer likes but generate more comments, indicating a shift in engagement towards discussions rather than passive approval. These results contribute to the literature on AI transparency and offer valuable insights into the balance between content transparency and user perceptions.

6 - The Impact of AI-Generated Summaries on User Engagement: Evidence from a Natural Experiment on a Video-Sharing Platform

Yingpeng Zhu, University of Macau, Macau, China, People's Republic of, Rongxin Zhou, Junkai Zheng, Jiexin Zheng

While existing studies highlight the efficiency benefits of AI-generated summaries (AIGS), there is a concern regarding their potential to diminish the immersive experience that fosters user engagement. This study investigates the effects of AI-generated video summaries on user engagement by leveraging a natural experiment on Bilibili, a leading Chinese video-sharing platform. Using a difference-in-differences (DID) approach, we compare user comments between videos with and without AI-generated summaries. Our findings reveal that AI summaries significantly reduce user engagement. Our further analyses suggest that this impact is more pronounced for longer videos, those created by more popular creators, and videos that rely heavily on narrative development. Additionally, we find that viewers exposed to AIGS tend to use fewer affective words and fewer words about video content in their comments. Our study contributes to the literature by providing the first empirical evidence of the unintended negative consequences of AIGS on online platforms.

7 - Mitigating Modality Bias in Multi-Label Emotion Recognition

Ying Chen, Harbin Institute of Technology, Harbin, China, People's Republic of, Shoumeng Ge

Modality bias widely exists in multi-modal emotion recognition tasks, resulting in that one modality dominates the class prediction than the others. Even though several latest methods have been developed to improve the model performance of emotion recognition, theoretical analysis of the learning process of a multi-modal model is still lacking. Moreover, these methods do not have theoretical guarantees for their performance. Regarding

these, we pioneer the analysis of the training process of a multi-modal model and find that dominant modalities impede the learning process of other modalities. Based on this analysis, we theoretically propose a condition for multi-modal fusion. Once satisfied, the modality bias will be alleviated. According to these theoretical insights, we develop a new multi-modal fusion framework for the multi-modal multi-label emotion recognition (MMER) tasks. In experiments, we apply the proposed method to three large MMER datasets and show the outstanding performance of our method over existing ones.

8 - Balancing Privacy with Utility in Resource Allocation: Benefits of a Joint Differential Privacy Approach

Du Chen, Nanyang Technological University, Singapore, Singapore, Geoffrey Chua, Victor Jose

Resource allocation is an important problem and recent growth in privacy concerns has made this problem more challenging. Joint differential privacy (JDP), compared to differential privacy (DP), is believed to better balance privacy with utility in this context, though its implications remain underexplored. We show that if agents do not disclose their allocations, JDP provides privacy protection as strong as DP with a substantially lower utility loss. While the theoretical results imply JDP's potential for a better privacy-utility trade-off, existing JDP algorithms can backfire and empirically underperform DP. To bridge this gap, we develop Noisy Dual Mirror Descent (NDMD), a minimax near-optimal algorithm. Simulations on the allocation of Title I US education grants show that shifting from DP to JDP substantially reduces misallocation by as much as 52%. Our study highlights the importance of disclosure discretion, downstream operational tasks, and algorithm design in balancing privacy with utility.

9 - Parallel Machine Scheduling with Sequence-Dependent Setup TIME and Group Constraints for Implantation Processes in Semiconductor Manufacturing

Jun-hee Han, Pusan National University, Geumjeong-gu, Korea, Republic of, Ho-Hyeon Park, Ju-Yong Lee

This study addresses a complex parallel machine scheduling problem in semiconductor ion implantation, incorporating sequence-dependent setup time and energy group constraints. A Mixed Integer Linear Programming (MILP) model is proposed to minimize makespan, accurately reflecting operational rules such as beam tune insertion after consecutive job limits within energy groups. Implemented via IBM CPLEX, the MILP model precisely handles job assignment, sequencing, and setup/beam tune prioritization. Its effectiveness was validated through Gantt chart visualization and comparison against meta-heuristics (Genetic Algorithm (GA), Simulated Annealing (SA), Tabu Search (TS), and Hybrid GA+SA) under various job-machine scenarios Experimental results show that MILP achieved better solution quality and constraint satisfaction than most meta-heuristic approaches. Meta-heuristics offered faster but less accurate solutions. This research demonstrates MILP’s robustness and suitability as a benchmark for industrial scheduling in semiconductor manufacturing.

10 - Optimizing on-Demand Ride-Hailing Services in Two-Sided Coupled Markets with Impatient Riders

Yuanyuan Li, Harbin Institute of Technology, Harbin, China, People's Republic of, Hui Wang, Yang Liu, Xiaowei Hu, Jian Wang

This paper studies the optimal dynamic pricing strategy in a coupled ride-hailing market offering both ridesharing and non-ridesharing services. Riders may switch between alternatives or renege due to long waiting times, causing demand diversion and system imbalance. We develop dynamic meeting models that capture temporal matching patterns while incorporating rider switching behavior and impatience-induced reneging. Conditions on trip fares and commission rates are derived to ensure market existence. Furthermore, we analyze how dynamic pricing influences waiting times, market demand, and matching efficiency. An MDP model is formulated to optimize pricing and matching decisions, solved using Proximal Policy Optimization combined with a Lookahead-based market state estimation method. Extensive numerical experiments are conducted. If riders are more patient with waiting times, the platform operator can improve the trip fares for ridesharing riders, which can enhance the matching rate. Meantime, both the reneging rate and the level of demand diversion can be decreased.

11 - A Real-TIME Vehicle Routing Problem for Cold Chain Logistics Dealing with Sudden Changes in Traffic Conditions and User Needs

Pengkun Wu, Sichuan University, Chengdu, China, People's Republic of, Jiaoyue Ji, Eric Ngai

Cold chain logistics imposes stricter requirements on both delivery time and quantity. Extended transportation durations not only lead to energy waste but also diminish the freshness of cold-chain products, rendering a significant portion of excessive delivered goods unusable after prolonged transit. Consequently, in response to sudden changes in traffic conditions or user demands during delivery, real-time adjustment of vehicle routing becomes critically important. To effectively manage frequent fluctuations in traffic conditions and user needs, this study proposes a two-stage dynamic vehicle routing optimization model by integrating traffic congestion indices to accurately estimate delivery times. In the first stage, the pre-optimized model aims to minimize operational costs, including vehicle dispatch, product damage, and fuel consumption costs. In the second stage, dynamic adjustment strategies are implemented to account for sudden changes in traffic conditions and user needs. An enhanced particle swarm optimization algorithm is developed to address the proposed model.

12 - Valuation Benefit OR Financing Benefit? Modeling Capacity Investment with Technology Choice

Yingying Huang, Harbin Institute of Technology, Harbin, China, People's Republic of, Shu-Jung Sunny Yang, Teck Eng Yong

Motivated by observed practices in capital-intensive industries, this paper examines the asset pricing implications of a firm’s technology-choice and capacity-investment decisions under demand uncertainty in the presence of financial constraints. The firm faces borrowing limits when expanding capacity with either upgraded or conventional technology. We model this process using a mixed-integer nonlinear program (MINLP), integrating its outputs into a random coefficient choice model and a capital asset pricing model to evaluate firm value. A semiconductor sector case study demonstrates the model effectiveness in capturing the trade-offs between financing and valuation benefits. While the conventional technology offers a financing benefit, upgraded technology provides a valuation benefit. Sensitivity analysis reveals that increased intensity of upgraded technology reduces both benefits relative to upgraded technology. These findings suggest that in competitive, capital-intensive industries, the valuation advantage—rather than the financing benefit—of adopting upgraded technology serves as a stronger predictor of firm profitability.

13 - Decision Support in Mental Healthcare: Recognizing Out-of-Distribution Mental Disorders on Social Media

Qing Yin, University of Manchester, Manchester, United Kingdom, Tianrun Cai, Dongling Xu, Xian Yang

As mental health awareness grows, social media has emerged as the platform for individuals to share experiences and self-diagnose, though such diagnoses often lack professional insight and objectivity. AI-based techniques are increasingly replacing user self-diagnosis by automatically recognising mental disorders in user posts and guiding users toward more appropriate support. However, these techniques typically focus on predefined known/in-distribution (IND) mental disorders, neglecting constantly emerging unknown/out-of-distribution (OOD) disorders in the real world. This paper introduces a novel AI-based framework called cascaded mental disorders recognition (CMR), which recognises OOD mental disorders in user inputs to improve decision-making support for mental health conditions. In the CMR framework, a large language model utilises a specially designed prompt to generate names for mental disorders based on the detected OOD samples. The proposed CMR framework was evaluated using a ground-truth dataset from Reddit, demonstrating improved effectiveness in mental health decision support systems.

14 - The Role of Artificial Intelligence (AI) on Supply Chain Efficiency

Meini Han, Harbin Institute of Technology, Harbin, China, People's Republic of, Siqi Xu

This study investigates the role of artificial intelligence (AI) on supply chain efficiency. We find that the efficiency of the supply chain increases when the level of AI in enterprises within the supply chain improves. The magnitude of the effect varies predictably across industries, market positions, regional legal environment, and supply chain concentration, and is stronger in low-tech industries and low supply chain concentrations. By reducing the cost of matching supply and demand, reducing operating expenses, alleviating information asymmetry, and improving enterprise productivity, AI drives the improvement of supply chain efficiency. Based on China's supply chain innovation and application pilot policy, compared with non-pilot enterprises, the pilot enterprises promote the improvement of supply chain efficiency. Additional analyses show AI is associated with

enhanced supply chain resilience. Overall, our study highlights the effect of AI in reducing information asymmetry, transaction costs and improving operational efficiency in supply chain management.

15 - Product Line Design in the Presence of Generative AI Technology

Wei Lu, University of Science and Technology of China, Hefei, China, People's Republic of, Yu Jiang

This paper analyzes the practice of GAI technology in the production process and investigates how GAI technology impacts an upstream supplier's product line decision and a downstream platform's distribution contract strategy. Our work demonstrates that a longer product line consistently incentivizes the supplier to adopt GAI-assisted product design, irrespective of the distribution contract type. We also show that GAI technology motivates the supplier to extend its product line. Our paper is also the first to show that the choice of distribution contract significantly influences the supplier's strategy regarding GAI technology adoption and product line design. Our analysis finally reveals that with GAI technology, the platform favors an agency contract over a wholesale contract when the commission rate is either relatively low or high, the GAI investment cost is small, and the product line extension fee is moderate.

16 - Understanding the Relationship Between Firms&Rsquo; AI Innovation and Consumers Complaints

Yongchao MA, Huazhong University of Science & Technology, Wuhan, China, People's Republic of, Zhongzhun DENG

In the artificial intelligence (AI) age, firms increasingly invest in AI technology development to maintain competitive advantages. However, the relationship between AI technology innovation and consumer complaints remains inadequately explored. This paper draws on protection motivation theory to examine how firms' AI technology innovation affects consumer complaints. Using a multimethod approach, Study 1 examines S&P 500 firms' panel data, Study 2 analyzes Reddit data, and Study 3 includes two controlled experiments. Results show that AI technology innovation significantly increases consumers' privacy concerns. Meanwhile, product innovation (vs. process innovation) increases complaints more. This paper advances our understanding of consumers' psychological response to AI innovation and provides practical implications for managing consumer complaints.

Monday, July 21, 3:15 PM - 4:45 PM

MB02

Mercury

Market Design and Online Marketplaces

Invited Session

Auctions and Market Design

Chair: Bar Light, National University of Singapore, Singapore, Singapore

1 - Alternative Allocation Policies to Improve Deceased Donor Kidney Utilization

Grace Guan, Stanford University, Stanford, CA, United States, Itai Ashlagi, Philipp Strack, Felipe Subiabre

While around 90,000 patients are currently waiting for a deceased donor kidney transplant, almost 30% of recovered deceased donor kidneys go unused due to inefficiencies in the current allocation system. The Organ Procurement and Transplantation Network is currently searching for a structured and equitable way to improve organ utilization. We demonstrate how data-driven insights and simulation can provide a transparent and principled way to identify which organs are at risk of nonuse and how to allocate them.

2 - Optimal Fees in a Two-Sided Market

Yuxin Yuan, National University of Singapore, Singapore, Singapore, Bar Light, Hanzhang Qin

This study examines optimal fee structures for a two-sided online economic market in which buyers and sellers transact via a digital platform. We develop a stylized theoretical model that captures the strategic behavior of both sides, focusing on the platform's decision to

allocate fees between the buying side and the selling side. By analyzing how fees affect equilibrium price, we characterize the fee allocation scheme that maximizes the platform's value while preserving balanced market growth. Our findings offer insights into how platform operators can optimize their pricing strategy to attract high-quality sellers and sustain demand from buyers, ultimately improving market efficiency and long-term revenue.

3 - Platform Competition in Digital Advertising

Kerun Xu, National University of Singapore, Singapore, Singapore, Bar Light, Jussi Keppo, Zhengyuan Zhou, Xiaoyu Fan

Auto-bidding has become prominent in major online advertising platforms. In these settings, advertisers do not place individual bids but instead allocate budgets and set constraints across multiple platforms to maximize returns. This paper studies the allocation problem faced by advertisers and models it as a mean-variance portfolio optimization problem over platforms. We analyze the resulting equilibrium and characterize the platform's optimal mechanism in a competitive environment. We show that this mechanism can be computed via a convex optimization problem. Moreover, it can outperform the First Price Pacing Equilibrium (FPPE), which is optimal in the single-platform setting, by improving advertisers' return on investment and attracting a larger share of their budgets.

4 - Asymptotically Optimal Adaptive A/B Tests for Average Treatment Effect

Vikas Deep, Kellogg School of Management, Northwestern University, Evanston, IL, United States, Achal Bassamboo, Sandeep juneja

In A/B testing, an experiment designer typically aims to identify the better treatment (best treatment identification, BTI). We instead focus on the related, equally important, but more informative problem of estimating the difference between the two means, defined as the average treatment effect (ATE). For computational efficiency, we restrict accuracy to a confidence interval (CI) of width at most ϵ, where the probability of the CI not containing the ATE is at most δ. The goal is to estimate a CI of the ATE with minimum expected sample size, i.e., minimize the expected total number of individuals receiving treatment A or B.

We first establish a lower bound on the expected sample size needed for any adaptive experimental policy that constructs such a CI, framed as a max-min problem for small δ. This also provides the asymptotic fraction of assignments to A and B for any asymptotically optimal policy. To gain insight, we define a notion of information for any distribution and prove that the fraction of treatment assignments to A and B should be inversely proportional to the square roots of the information of their outcome distributions.

Using these insights, we construct an adaptive policy that is asymptotically optimal, i.e., matches the lower bound for small ϵ and δ. We compare our policy with two popular ones: a) randomized controlled policies (equal assignment), and b) Neyman's rule (assignment proportional to standard deviations), and show meaningful gains in sample efficiency.

MC01

Enterprise

Service Operations and Platform Management

Invited Session

Applied Probability Society

Chair: li xiao, N/A

1 - Distance-Based Fee Design of Rapid Delivery

Jianfu Wang, College of Business, City University of Hong Kong, Kowloon, Hong Kong, Manlu Chen, Mingliu Chen, Ming Hu

The rapid delivery service industry has seen significant growth and diversification in pricing and operational strategies. Employing a queueing economics model, we investigate the impact of a distance-based two-part

pricing scheme and the introduction of a self-pickup option on vendor profitability, customer welfare, and social welfare across different stages of market evolution—ranging from low-density introduction to high-density saturation stage. Our findings reveal that two-part pricing enhances vendor profits compared to flat-rate pricing but has nuanced effects on customer welfare: it reduces welfare in low-density markets while improving it in high-density markets. Two-part pricing optimizes delivery efficiency in high-density markets by focusing on nearby customers, benefiting the vendor and customers. However, in low-density markets, the reduction in customer welfare outweighs vendor profit gains, leading to lower social welfare.
Additionally, introducing a self-pickup option creates a mutually beneficial situation for vendors and customers. By allowing nearby customers to collect their orders, self-pickup increases delivery capacity, improves resource allocation, and enhances overall efficiency, especially in high-density markets where the concentration of nearby customers maximizes the positive impact. Our study provides actionable insights for rapid delivery platforms, highlighting the importance of tailoring pricing and operational strategies to market conditions. Early-stage platforms can adopt flat-rate pricing to foster market adoption, while mature platforms can transition to two-part pricing and self-pickup options to optimize efficiency and profitability. These findings underscore the strategic value of two-part pricing and self-pickup as tools for improving operational performance and stakeholder outcomes in the rapidly evolving delivery service industry.

2 - Designing a Forward-Looking Probabilistic Matchingpolicy for En-Route Ridepooling Service

Xiaolei Wang, Tongji University, Shanghai, China, People's Republic of

With the development of information technologies in recent years, dynamic ridepooing services e.g., UberPool, Didi Pinche, has become an increasingly popular choice. In dynamic ridepooling services, serviceproviders respond to on-demand mobility requests immediately, dispatch (vacant or partially occupied)vehicles in real-time, and keep searching for matching orders along the trip. Most existing dispatchingstrategies for dynamic ridepooling give priorities to partially occupied vehicles and form a ridepooling tripimmediately when they find a match. Such vehicle dispatching strategies ignore forthcoming matchingopportunities that may have higher matching qualities, therefore having short-sighted limitations. In thispaper, we begin by modeling the matching process in en-route ridepooling system as a Markov decisionprocess and prove that its long-run average reward, such as the distance saved per unit time, convergesto a constant under any given matching policy. We then, propose a typical type of probabilistic matchingpolicy, under which appeared matching opportunities are accepted with different probabilities, anddevelop a mathematical model that can efficiently approximate the long-term performance of theridepooling system, as measured by total distance savings, under any of these probabilistic matchingpolicies. Based on the model, we formulate the matching policy optimization as a single-stage problemand design an efficient algorithm that optimizes the probabilistic matching policy based on its estimatedperformance. Through simulation experiments on real road network of Haikou (China), we demonstratethe prediction accuracy of our model and the effectiveness of the optimized matching policy in enhancingtotal distance saving and emission reduction.

3 - A Data-Driven TIME-Constraints Order Dispatch Method for on-Demand Food Delivery

Dongling Rong, Xi'an Jiaotong University, Xi 'an, China, People's Republic of, Meilin Zhang, Ruina Yang

The rapidly expanding on-demand food delivery sector confronts a significant challenge in dynamically matching orders to meet strict delivery time constraints. This task is fraught with unpredictability, including varied driving speeds, fluctuating service times, and changing external conditions. The goal is to refine this matching process to consistently achieve high service standards amidst these uncertainties. We first formulated the problem as a target-oriented order-matching stochastic program, designed to balance uncertain delivery times with the delivery service targets, and transformed complex measures into more manageable standard assignment problems utilizing stochastic empirical distributions of uncertain parameters. This encompasses two key service targets: the stringent mandatory delivery target and the stage-wise lateness penalty mechanism. The maximized joint probability model, targeting stringent mandatory delivery targets, is effective in reducing delivery delays and minimizing average lateness. The maximized expected revenue model, focused on the stage-wise penalty mechanism, seeks to optimize profit margins for the platform or drivers. Interestingly, the outcomes of these two models tend to align when service capabilities are robust and external conditions are favorable. Extensive numerical tests with synthetic and industry data quantified enhanced service commitments across various metrics. Compared to benchmarks, our satisficing method proves computationally scalable and generally yields higher returns in the unpredictable on-demand food delivery environment. This method adeptly circumvents

computational challenges typical of stochastic methods in dynamic order matching. Our findings also suggest several managerial implications and extensions, indicating the practical applicability and adaptability of our approach in real-world on-demand delivery scenarios.

4 - The Israeli Queue with a Capacitated Server: Modeling and Approximations

Nir Perel, Afeka Academic College of Engineering, Tel Aviv, Israel, Efrat Perel, Mor Kaspi

The Israeli Queue is a batch service polling system in which a single server attends a collection of queues according to a seniority order. An arriving customer may belong to one of multiple classes. Upon arrival, she either joins a queue with customers from her class or opens a new queue in case she is the first in line from her class. Members of the class with the most senior customer are served together as a batch within a service time that is independent of their amount. This service regime is encountered in applications such as advanced elevator systems and on-demand shared mobility, where passengers with the same destinations may share a ride. However, in many systems, the vehicles' capacities are small and binding. Thus, calling for an implicit modeling investigating the Israeli Queue with a capacitated server (IQCS). In this paper, we provide a formal definition of the IQCS and delve into the challenges associated with formulating a mathematically tractable model to represent it. We approximate the IQCS by formulating a corresponding quasi-birth-death process, and derive approximations for various performance measures. To validate our approach, we implement an ad-hoc simulation model and use it to compare the IQCS, the approximate model, and the Israeli Queue. Our analysis across various settings demonstrates the accuracy of the approximate model. Nonetheless, the presence of a remaining gap underscores the ongoing challenge of precisely and efficiently modeling the IQCS, posing an open question for the research community.

MC03

Minto

Recent Advances in Data-driven Optimization

Invited Session

Computing Society

Chair: Bin Tian, The Hong Kong Polytechnic University, Kowloon, N/A

Co-Chair: Shiyi Jiang, The Hong Kong Polytechnic University, Hong Kong, N/A

1 - *Optimized Dimensionality Reduction for Moment-Based Distributionally Robust Optimization*

Shiyi Jiang, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, Jianqiang Cheng, Kai Pan, Zuo-Jun Shen

Moment-based distributionally robust optimization (DRO) provides an optimization framework to integrate statistical information with traditional optimization approaches. Under this framework, one assumes that the underlying joint distribution of random parameters runs in a distributional ambiguity set constructed by moment information and makes decisions against the worst-case distribution within the set. Although most moment-based DRO problems can be reformulated as semidefinite programming (SDP) problems that can be solved in polynomial time, solving high-dimensional SDPs is still time-consuming. Unlike existing approximation approaches that first reduce the dimensionality of random parameters and then solve the approximated SDPs, we propose an optimized dimensionality reduction (ODR) approach by integrating the dimensionality reduction of random parameters with the subsequent optimization problems. Such integration enables two outer and one inner approximations of the original problem, all of which are low-dimensional SDPs that can be solved efficiently, providing two lower bounds and one upper bound correspondingly. More importantly, these approximations can theoretically achieve the optimal value of the original high-dimensional SDPs. As these approximations are nonconvex SDPs, we develop modified Alternating Direction Method of Multipliers (ADMM) algorithms to solve them efficiently. We demonstrate the effectiveness of our proposed ODR approach and algorithm in solving multiproduct newsvendor and production-transportation problems. Numerical results show significant advantages of our approach regarding computational time and solution quality over the three best possible benchmark approaches. Our approach can obtain an optimal or near-optimal (mostly within 0.1%) solution and reduce the computational time by up to three orders of magnitude.

2 - Optimal Experimental Design with Routing Constraints

Baizhi Song, London Business School, London, United Kingdom, Jean Pauphilet

Data collection in application domains like agriculture and environmental science requires the deployment of sensors in large remote areas. These use cases challenge the traditional optimal experimental design (OED) formulation from statistics by their scale as well as the presence of complex operational constraints, such as that data is collected along the trajectory of a drone. This paper proposes a tailored branch-and-bound algorithm to solve large-scale instances of OED as well as account for routing constraints. Efficient screening rules and first-order methods for solving the relaxations are key to the tractability of our algorithm, which significantly outperforms commercial solvers on both the classical OED problem and OED with routing constraints. For classical OED for example, it increases the size of instances solvable under a minute from n=200 to n=10,000. In the presence of routing constraints, it finds solutions with up to 10% higher objective value and solves twice as many instances to optimality. In use cases where data is collected to construct a predictive model, our approach achieves the same accuracy as alternative collection strategies with 30-50% less samples, hence demonstrating the potential impact of optimization on data collection cost.

3 - The Information Projection in Moment Inequality Models: Existence, Dual Representation, and Approximation

Rami Tabri, Monash University, Melbourne, Australia

Relative entropy minimization is a widely used method in decisions and operations research that incorporates information through constraints on the underlying probability model. The solution is called information projection, and we present new results for its existence, exponential family representation, and approximation in the infinite-dimensional setting for moment inequality constraint sets, nesting both conditional and unconditional moments and allowing for an infinite number of such inequalities. Our approach considers the Fenchel dual of the relative entropy minimization problem, with a key innovation being the exhibition of the dual variable as a weak vector-valued integral, enabling the formulation of a simple approximation scheme. Under suitable assumptions, the values of finite-dimensional convex stochastic programs can approximate the dual problem's optimum value and that, in addition, every accumulation point of a sequence of optimal solutions for the approximating programs is an optimal solution of the dual problem. We illustrate the verification of assumptions and construction of the approximation scheme's parameters for the cases of unconditional and conditional first-order stochastic dominance constraints and conditions that characterize selectionable distributions for a random set, and present numerical experiments demonstrating the simplicity of the approximation scheme.

MC04

Indiana

Behavioral Operations Management II

Invited Session

Behavioral Operations Management

Chair: Fadong Chen, Zhejiang University, Hangzhou, N/A

1 - Could Large Language Model be Used for Worldwide Large-Scale Survey?

Fadong Chen, Zhejiang University, Hangzhou, China, People's Republic of

Large-scale survey is an essential and valuable approach for social scientists. As Large Language Models (LLMs) rise in prominence, some propose that LLMs have the potential to apply in surveys and experiments to simulate humans as subjects aiming to reduce time and monetary costs. This paper leverages a virtual worldwide survey using GPT, based on the framework of the Global Preference Survey (GPS), to predict more than 70,000 respondents' economic preferences across 75 countries in 43 different languages using their actual responses as a benchmark. GPT consistently predict people's economic preferences, with much smaller elasticity. We evaluate GPT survey results by replicating the previous conclusions draw from the GPS dataset. Most findings drawn from individual-level data could be reproduced whereas only one country-level conclusion could. During replication process, we found a side effect that the gender difference gap in LLM predictions is contaminated by

contemporary gender stereotype and gender inequity. Clarifying pros and cons of LLMs will be essential for upcoming transformation of social science research.

2 - Bilateral Pricing and Quality Regulation of An E-Reading Platform

Juan Li, Nanjing University, Nanjing, China, People's Republic of, Zhongyuan Hao, Yue Zhou

We investigate the quality regulation strategies and bilateral pricing decisions of an e-reading platform that connects writers and readers. Writers are classified into two skill levels: high-skill and low-skill. Low-skill writers must exert extra effort to enhance their work from low to high quality. The platform considers two quality regulation strategies: ruling out low-skill writers and excluding low-quality works. We analyze how these strategies interact with writers' time-inconsistent preferences and examine their effects on bilateral pricing, platform performance, and writer abandonment behavior. Our findings indicate that as writers' time-inconsistency weakens, the platform tends to impose higher subscription fees for readers and offer greater compensation for writers. While the platform consistently favors excluding low-quality works, this approach exacerbates writer abandonment. We further explore the role of sunk cost aversion and find that writers exhibiting this bias are less likely to abandon their work, particularly when their time-inconsistency is pronounced. When writers are initially unaware of their sunk cost aversion, the platform benefits from making them cognizant of this behavioral tendency. Moreover, integrating quality regulation with sunk cost aversion strategies emerges as the most effective approach for maximizing platform profit. However, the combined effect on abandonment rate depends on the trade-off between stricter quality enforcement and behavioral adjustments driven by sunk cost aversion.

3 - Manufacturer-Led Supply Chain Collaborative Emission Reduction: Mandatory Targets and Contractual Incentives

Xiaohuan Wang, Northeastern University School of Business Administration, Shenyang, China, People's Republic of, Yujie Li, Zhi-Ping Fan

Reducing direct and indirect carbon emissions from supply chains poses a significant challenge for brand firms (brands) in their pursuit of sustainable development. To mitigate this, brands incentivize their suppliers to work with them on carbon emission reduction (CER). Two incentive measures are commonly used: setting mandatory CER targets for supply chains and providing contractual subsidies to suppliers. In this paper, four incentive strategies are introduced: NN (no target or subsidy), NY (no target but subsidy), YN (target but no subsidy), and YY (target and subsidy). Considering the brand's different internal and external CER conditions, implicit utility and explicit pressure from supply chain CER, this paper explores the brand's optimal incentive strategy and its implementing conditions. The results reveal that, counter-intuitively, if brands can not derive high implicit utility from supply chain CER, they should not provide any incentive; otherwise, NY is suggested. Moreover, setting a CER target under the existing contractual incentive will reduce the brand's utility. Nevertheless, in the cases where the brand must declare a CER target, if the brand is more concerned about market share, YN is optimal; if the brand wants the supplier to reduce emissions to the greatest extent, YY is optimal. Interestingly, implementing cap-and-trade does not affect CER performance in the supply chain where clear CER targets are set. However, it encourages the brand to pass on some of the CER pressure to its supplier, and the market demand will be reduced.

4 - Channel Choice in Live-Streaming Commerce: a Comparative Analysis of Weak vs. Strong Brands

Haozhe Yao, Henan University of Technology, Zhengzhou, China, People's Republic of, Yonghai Li

*Background:*Live streaming commerce has emerged as a critical customer acquisition tool, driving the proliferation of diverse streaming formats. As different modes align with firms' varying brand strengths, selecting an optimal live streaming strategy presents a strategic challenge. *Problem definition:*This paper examines how firms with weak or strong brands choose among three live streaming modes in competitive markets: (1) traditional e-commerce (Channel N); (2) a dual channel structure combining traditional e-commerce with self-run live streaming channel (Channel NS); or (3) a dual channel structure integrating traditional e-commerce with third-party live streaming (Channel NL).*Methodology and Results:* Using a game-theoretic model involving a focal firm (M1, weak/strong brand) and a competitor (M2), we analyze M1s' decisions on streaming modes, pricing, slotting fees (negotiated with third-party streamers), and streaming effort, all

constrained by brand strength. This study leads to three key findings. First, When brand strength is high or consumers exhibit strong live streaming preference, M1s prioritize the NS channel. High extra cost for watching live streaming don't rule out NS/NL as optimal choices. Second, When M1 is weak brand, NS channel is prioritized if amortized cost for setting up the channel are low, and NL channels are favored when the cost are high and third-party streamers have low popularity. Third, When M1 is strong brand and its bargaining position is strong, slotting fees may correlate positively with commission rates.*Managerial Implications:* This study offers actionable insights for firms to align live streaming strategies with brand strength, optimize pricing, and negotiate effectively with third-party streamers.

MC05

Orchard A

Digital Twin in Urban Transportation

Invited Session

Digital Twin and Metaverse

Chair: Lu Hu, Southwest Jiaotong University

Co-Chair: Yan Li, N/A

1 - Dynamic Joint Decision of Matching Parameters and Relocation Strategies in Ride-Sourcing Systems Interacting with Traffic Congestion

Jun Zhang, Southwest Jiaotong University, Chengdu, China, People's Republic of, Lu Hu, Yan Li, Weiyao Xu, Yangsheng Jiang

As the ride-sourcing market expands, ride-sourcing fleets have increased urban traffic congestion, and in turn, road congestion is affecting ride-sourcing operations. It is crucial to incorporate this interaction mechanism into the operational decision of ride-sourcing systems. This mechanism can be described in detail by simulation models. Applying simulation-based optimization methods to the real-time operation of ride-sourcing systems is challenging due to the limitations of simulation efficiency. To tackle this issue, this study first proposes a modular simulation model combining time-driven and event-driven mechanisms. The simulation model comprises two layers: the traffic flow module at the bottom, defined by trip-based and multi-region macroscopic fundamental diagrams and cell transmission models, and the ride-sourcing operation module at the top, which considers the interaction between the traffic congestion and ride-sourcing operation and models the random behavior of passengers and drivers. We verify the efficiency and accuracy of the simulation model. The simulation model is utilized as an evaluator of the objective functions and performance measures. Then, we integrate the Bayesian optimization, parallel sampling, and rolling horizon approaches to develop an efficient and effective simulation-based optimization framework for the dynamic joint decision-making of matching radius, matching intervals, and threshold relocation strategies. Numerical experiments for the traffic scenario in Berlin reveal some interesting findings. The optimal operation strategy considering endogenous congestion not only leads to an increase in the system revenue and efficiency but also leads to congestion relief, even though the latter was not our initial objective.

2 - Real-TIME Operational Decision-Making for Shared Autonomous Vehicles Considering Vehicle-Pile-Road-Grid Interaction

Boyu Hou, Southwest Jiaotong University, Cheng Du, China, People's Republic of

With the rise of the Internet and smart technologies, shared mobility has transformed transportation by optimizing resource allocation. Shared Autonomous Electric Vehicles (SAEVs) are being piloted in multiple cities, but large-scale adoption may worsen traffic congestion and grid load, necessitating efficient scheduling. This study develops a simulation model integrating vehicle-station-road-grid interactions and proposes a strategy for jointly optimizing SAEV charging and operations. Using the O2DES framework, the urban traffic network is modularized, defining static/dynamic attributes and event-driven processes. A three-module SAEV system (vehicles, stations, users) is designed, with operations categorized into service, dispatch, and charging. Demand is classified into intra-/inter-zone trips for better matching, dispatch employs a regional-threshold-based strategy, and charging adjusts dynamic thresholds for efficient energy use. A BO-SPSA hybrid algorithm optimizes threshold parameters. Experiments on Berlin’s morning peak data show a 43.51% profit increase, 20.62% energy

reduction, and peak-hour charging shift, with rolling optimization achieving 95.07% of global optimization performance. SAEVs have minimal impact on traffic delays, and fast chargers become more effective as demand rises.

3 - Optimizing Shared Autonomous Vehicle Systems with Vehicle-to-Grid Interaction: a Multi-Objective Simulation-Optimization Approach

Hao Li, Changsha University of Science and Technology, Changsha, China, People's Republic of, Yan Li, Lu Hu

Shared autonomous electric vehicles (SAEVs) are becoming increasingly prevalent. Integrating vehicle-to-grid (V2G) capabilities offers significant potential to enhance their operational profit, yet it also introduces considerable complexity to service management strategies. Addressing the critical challenge of optimizing SAEV operational decisions while effectively leveraging V2G interactions is paramount. This study proposes a large-scale, multi-objective simulation-optimization framework specifically designed to explore the synergistic advantages of combining V2G functionalities with SAEV services. We developed a logic-based physical simulation model utilizing a modular discrete-event mechanism to accurately replicate SAEV operations. This model encompasses three core modules: a station module simulating user pick-ups/drop-offs and vehicle charging/discharging processes (including V2G); a state-dependent link module capturing the dynamic interplay between traffic congestion and vehicle dynamic; and an operation module managing real-time vehicle relocation, user request assignments, trip pricing, and an inventory-theory-based charging/discharging strategy optimized via threshold updates. Based on this simulation, we established a multi-objective optimization framework utilizing multi-gradient descent and Simultaneous Perturbation Stochastic Approximation (SPSA) algorithms. Targeting Pareto optimal solutions, the framework simultaneously maximizes operator profit and user satisfaction by co-optimizing key decision variables: dynamic pricing schemes, charging/discharging strategy parameters, fleet size, and station capacities. Validation via a large-scale Chengdu case study confirmed the methodology's effectiveness and superiority. The results demonstrate that V2G integration significantly enhances SAEV operational profits compared to traditional charging, while minimally affecting user satisfaction. Sensitivity analyses on critical parameters were also performed.

4 - Simulation-Based Optimization for Electric Vehicle Routing with Multiple Depots and Endogenous Congestion

Juanxiu Zhu, Xihua University, Chengdu, China, People's Republic of

With the development of the logistics industry and the electrification of freight vehicles, large-scale deployment of electric freight vehicles is inevitable in urban distribution systems. However, large fleets of electric freight vehicles are expected to exacerbate congestion in urban road networks and congestion at charging stations, negatively impacting distribution efficiency. In scenarios involving multiple distribution centers and an incomplete road network, this bidirectional interaction becomes particularly pronounced and thus should be incorporated into the electric vehicle routing problem (EVRP). Due to the strict modeling assumptions, traditional mathematical programming approaches face inherent limitations in accurately capturing the aforementioned endogenous congestion. Simulation-based optimization offers a more tractable approach. The challenge of applying simulation-based optimization to solve EVRP lies in achieving a dual emphasis on model accuracy and computational efficiency. In this paper, we first developed a simulation model to capture mesoscopic dynamics in the distribution process of electric trucks in an incomplete road network. This model adopts a modular event-driven discrete-event mechanism, enhancing efficiency while maintaining high fidelity. The simulation model is utilized to evaluate the objective functions and performance measures. Subsequently, we proposed an enhanced Adaptive Large Neighborhood Search (ALNS) algorithm to establish a rapid and effective simulation-based optimization framework for the EVRP. Based on Solomon's benchmark dataset, we designed 29 experimental cases. Results demonstrate that incorporating endogenous congestion in EVRP helps reduce the total logistics distribution cost and this cost-saving effect becomes more pronounced under intensified congestion and an increased number of distribution centers.

5 - Electric OR Hybrid Electric Vehicles in Carsharing: a Large-Scale Multi-Objective Simulation-Optimization Analysis

Yan Li, National University of Singapore, Singapore, Singapore, Lu Hu, Haobin Li, Ek Peng Chew

This work pioneers a general large-scale multi-objective simulation-optimization (MOSO) method to explore the values of deploying hybrid electric vehicles (HEVs) in CSSs. We firstly develop a physically logical simulation model, emulating operations of CSSs and capturing mesoscopic dynamics of shared vehicles in a link-based traffic network. Subsequently, we design a simulation-optimization framework aimed at achieving Pareto optimality by jointly optimizing station capacity, fleet size, and trip pricing. The goal is twofold: to maximize operational profits and to minimize carbon emissions. To tackle the high-dimensional MOSO problem, we introduce the multi-objective optimization into stochastic approximation field by proposing a general algorithm that incorporates the multiple gradient descent algorithm with the simultaneous perturbation stochastic approximation algorithm. Furthermore, we derive its analytical expression for bi-objective optimization problems. The efficiency of this method was validated through large-scale computational experiments conducted in Chengdu, Sichuan Province, involving 66,710 decision variables. These experiments showcased the method's superiority over existing MOSO algorithms. Several groups of sensitivity experiments reveal some interesting findings. (1) Regardless of the increase in travel distances, shared HEVs, which can be viewed as shared EVs (SEVs) without range anxiety (RA), continue to primarily rely on electricity rather than fuel for their operational mileages. (2) Under any traffic condition, the dual-engine feature of SHEVs significantly reduces the number of failed pickups. (3) As travel demand increases, the state of charge for SEVs may rapidly fall below the threshold that triggers RA, whereas SHEVs maintain a more reliable power supply.

MC06

Orchard B

Advanced Analytics and Financial Decision-Making

Invited Session

Finance

Chair: Ben Charoenwong, INSEAD, 1 Ayer Rajah Ave, Singapore, 138676, Singapore

1 - Deep Learning for Solving Dynamic Asset Allocation Models: Application for the Optimal Dynamic Rebalancing Problem

Bong-geun Choi, Yonsei Univeristy, Seoul, Korea, Republic of

Dynamic asset allocation is a cornerstone of modern portfolio management, yet solving high-dimensional, nonlinear dynamic programming problems remains computationally challenging. Recent advances in deep learning offer promising avenues for overcoming the curse of dimensionality inherent in these models. In this paper, we develop a novel deep learning framework for solving dynamic asset allocation models, with a particular focus on the optimal dynamic rebalancing problem. Building on the methodology of Maliar, Maliar, and Winant (2021), we employ neural networks to approximate value and policy functions within a stochastic dynamic programming context. Our approach leverages the flexibility and scalability of deep neural architectures to efficiently handle large state spaces and complex asset dynamics. Through extensive numerical experiments, we demonstrate that our deep learning-based method achieves high accuracy in approximating optimal rebalancing strategies, outperforming traditional grid-based and regression-based techniques in both speed and scalability. The results highlight the potential of deep learning as a powerful tool for dynamic portfolio optimization, paving the way for practical applications in high-dimensional financial environments.

2 - When Robo-Advising Meets Large Language Models

Fangzhou Lu, University of Hong Kong, Hong Kong, Hong Kong

We utilize a robo-advising setup to investigate the impact of financial literacy on investors, specifically focusing on the concepts of beta and alpha. Partnering with a large brokerage firm, we study the effects of educating investors about beta and alpha in mutual funds and stocks. To enhance investors' understanding, we integrate a back-end LLM (ChatGPT/DeepSeek) system with the robo-advisor. Our findings reveal that after learning about beta, investors show a preference for investing in ETFs or index funds rather than directly in active mutual funds or individual stocks. A one-click automatic enrollment mechanism has a much larger impact than a self-assembled portfolio, but the gap is much smaller for LLM-supported investors. This suggests that an LLM-based chatbot exerts a long-term improvement on financial literacy. LLM-based chatbots also generate a spillover effect, prompting investors to pursue international diversification even when these options are not initially recommended by the robo-advisors but are still available in the menu. Overall, investors who receive this

financial education benefit significantly, with their monthly returns increasing by 40 basis points and an average monthly gain of 750 RMB.

3 - Are Asset Pricing Models Sparse?

Shuhua Xiao, City University of Hong Kong, Hong Kong, Hong Kong, Doron Avramov, Guanhao Feng, Jingyu He

Given empirical evidence linking cross-sectional equity returns to high-dimensional characteristics (e.g., size, value) and the growing use of variable selection methods (e.g., lasso), this paper examines whether these characteristic signals are sparse. Specifically, we adopt a Bayesian sparsity prior within the conditional latent factor model (Instrumental PCA, Kelly, Pruitt, and Su, 2019), where conditional alphas (mispricings) and betas (factor loadings) are driven by characteristics. We have several important findings by analyzing monthly U.S. equity data from 1980 to 2023. First, alphas and betas are unlikely to be driven by either sparse or dense models, as supported by statistical and economic evidence. Second, characteristics associated with betas tend to be less sparse than those linked to alphas. Third, sparsity varies with test asset selection. Models estimated with less efficient test assets exhibit lower sparsity, highlighting the critical role of test asset choice. Fourth, we also find notable temporal variation in the degree of sparsity, with a marked increase during economic recessions. Finally, our framework offers a novel method for estimating the conditional CAPM by recovering the unspanned factor through augmented latent factors, providing fresh interpretive insights.

MC07

Orchard C

Healthcare and Socially Responsible Operations: Emerging Research by Rising Scholars

Invited Session

Health Applications Society

Chair: Jiatao Ding, Nanyang Technological University, Singapore, Singapore

1 - Collaborative Intelligence: Shareholder Value of AI and Advanced Analytics Partnerships for Biopharma Firms

Jiatao Ding, Nanyang Technological University, Singapore, Singapore, Michael Freeman, Niyazi Taneri

Artificial intelligence and advanced analytics (AIAA) technologies have emerged as promising solutions for enhancing the sustainability and profitability of the biopharmaceutical (biopharma) industry amid persist research and development (R&D) challenges, operational inefficiency, and heightened market competition. They offer transformative potential across the biopharma value chain, from screening drug candidate and accelerating clinical trials to streamlining supply chains and optimizing commercialization strategies. Recognizing that traditional biopharma firms often lack deep AIAA expertise, many have increasingly pursued strategic partnerships with specialized technology providers. While such partnerships allow biopharma firms to access complementary AIAA capabilities, their management carries inherent complexities and risks. This study empirically investigates the shareholder value implications of AIAA partnership announcements for biopharma firms. Using the event study methodology, we find that, on average, these collaborations generate positive and significant abnormal stock returns, confirming their strategic importance. Moreover, partnerships spanning multiple therapeutic domains or those emphasizing exploratory activities such as early-stage drug discovery yield higher returns, compared to narrowly focused or exploitative ones. This underscores the advantages biopharma firms derive from leveraging the versatility and explorative power of AIAA technologies. Moderation analyses reveal that firm-specific attributes also shape outcomes: R&D-intensive biopharma firms capture greater value from exploratory partnerships, whereas sales-driven firms see comparatively lower benefits from such initiatives. Our findings illuminate the nuanced relationship between AIAA partnerships and shareholder value creation, providing actionable insights for biopharma firms to strategically align these collaborations with their internal capabilities and market opportunities, thereby enabling them to navigate the rapidly evolving technological landscape.

2 - Should I Change My Physician&Rsquo;S Gender? on the Role of Ambiguity in Patient-Physician Gender Concordance

Xiaodan Shao, Nanyang Technological University, Singapre, Singapore, Vivek Choudhary, Anandasivam Gopal

Matching patients with physicians remains an important part of the healthcare system. Literature has established that patients tend to have strong preferences for physician attributes, especially gender. Patients often prefer physicians of the same gender to the extent that they are willing to postpone procedures or incur additional costs. This gender concordance can improve healthcare quality and reduce unnecessary resource utilization through better patient-physician alignment. However, major societal events, such as pandemic lockdowns, may disrupt these established preferences. This study examines how heightened ambiguity during the lockdown influences patients' adherence to or deviation from their typical gender-based physician choices. Using an ambiguity aversion framework and a differences-in-differences model, we investigate whether uncertainty shifted patients' priorities between gender preference and perceived empathy in care. The findings indicate a ~11% decrease in gender concordance for female doctors post-lockdown, driven primarily by male patients increasingly choosing female doctors, while female patients' preferences remain unchanged. This shift may be attributed to the perception that female physicians exhibit greater empathy through spending more time with patients, making them a preferred choice amid the ambiguity introduced by the lockdown. To address these imbalances, policymakers may increase the supply of female physicians through targeted recruitment, expanded admissions, and retention strategies. Healthcare systems can adopt flexible workforce planning and training to accommodate shifting gender-based preferences.

3 - Optimized Subsidy Design: Application to Food Assistance Program

Zhicong Hu, INSEAD, Singapore, Singapore, Ali Aouad, Kamalini Ramdas, Alp Sungu

In the design of subsidy programs, a significant challenge lies in balancing policy efficiency with the flexibility of individual choice. Traditional paternalistic approaches often prioritize optimal social outcomes, such as promoting the consumption of nutritious but less-popular foods, potentially at the cost of individual preferences. This can lead to low take-up of the subsidy or low compliance, undermining the policy's effectiveness. Conversely, flexible subsidies, such as cash transfers, may result in suboptimal impact on nutrition. This study introduces the *optimal subsidy design problem*, which embeds a probabilistic choice model to represent the beneficiaries' decision to avail of an offered subsidy. Our formulation can optimize over various subsidy mechanisms, including in-kind bundles, assortments, vouchers, or cash-plus-in-kind mixes. After calibrating our model to data from a randomized controlled trial (RCT) on staple-food subsidies in India and synthetic data from AI simulations, we run counterfactuals to identify optimal subsidy designs. The analysis shows that pairing an in-kind staple subsidy with a modest cash top-up can outperform both pure in-kind and pure cash schemes on social outcomes and beneficiaries' choices.

4 - Sharing in Caring: Analyzing Shared Medical Appointments Through An Operations Lens

Sundara Natarajan Panchanatham, HKU Business School, Hong Kong, Hong Kong, Enver Yucesan

Shared medical appointments (SMA) are doctor-patient visits in which groups of patients are seen by healthcare professionals in a concurrent session. The setup can potentially enhance the patient experience, eliminate redundant work, improve compliance, and reduce costs for certain treatments. From an operations standpoint, providers face several key challenges in the smooth conduction of SMAs, including finding the optimal group size and the optimal capacity allocation. In this work, we mathematically analyze the key operational trade-offs involved in SMAs as a service and provide prescriptions that enable providers to choose the optimal decision variables (capacity/payment/group size) for several system configurations.

MC08

Bras Basah A

Innovative Approaches in Healthcare Optimization and Decision-Making

Contributed Session

Healthcare Applications

Chair: Jinxin Zhang, Tuen Mun, N/A

1 - Novel Immunological and Optimization Strategies to Optimize Kidney Exchange

Michal Mankowski, NYU Langone, NYC, NY, United States, Khalid AlMashari

Many patients with end-stage kidney disease (ESKD) are unable to receive a transplant from their willing living donors due to immunological incompatibility. Kidney Exchange Programs (KEPs) offer a solution by matching incompatible patient-donor pairs with others in a similar situation to enable compatible exchanges. This study focuses on enhancing the effectiveness of a nationally coordinated KEP through the integration of advanced immunological matching strategies—specifically high-resolution HLA typing and eplet-based compatibility scoring. These approaches allow for more precise assessments of compatibility, increasing the number and quality of possible matches.

We apply machine learning models to predict graft and patient survival based on immunological profiles, generating estimates of expected health benefits in terms of quality-adjusted life years (QALYs). These predictions are embedded in optimization algorithms that guide match selection, not only maximizing the number of transplants but also improving outcomes and ensuring equity. This is especially relevant for patients who are harder to match, such as those with rare HLA types or heightened sensitization.

While the methods are broadly applicable, the project is grounded in collaboration with the national kidney exchange infrastructure in Saudi Arabia, where expanding access to living donor transplantation is a growing priority. The findings aim to inform data-driven and equitable health policy, improve transplant outcomes, and serve as a model for integrating precision immunological data into transplant logistics within a national program

2 - A Graph-Theoretic Framework for Facility Layout Planning in An Eye Specialist Outpatient Clinic

Fanwen Meng, National Healthcare Group, Singapore, Singapore, Kiok Liang Teow, May Lan Lim, Ling Yang Quek, Xin Jin, Grace Chiang, John Abisheganaden

The facility layout problem is of significant practical relevance in healthcare settings, where spatial efficiency and streamlined patient flow are essential. This study presents a comprehensive framework for facility layout planning in an eye specialist outpatient clinic within a public hospital, utilizing a graph-theoretic approach. Patient visit data were analyzed to identify movement patterns and interaction frequencies among service stations. Based on these insights, a maximal planar graph was constructed using a modified greedy heuristic algorithm. Incorporating input from the hospital team, the resulting graph served as the foundation for generating a block layout design that integrates both functional requirements and design constraints, ultimately contributing to the architectural design of the new eye clinic.

3 - Smart Risk Prediction: Integrating Physician Decision and Patient Adherence for Diabetes Management

Xiner Dai, Zhejiang University, Hangzhou, China, People's Republic of, Zheng Zhang

Risk prediction plays a critical role in diabetes management by guiding clinical decisions. However, conventional risk prediction focuses on identifying the most accurate state possible through patients' diabetes-related features, neglecting the subsequent decisions of physicians and patients. We propose a smart risk prediction framework for diabetes management that prospectively integrates downstream physician decision-making and patient adherence behavior. Our approach models patients' examination actions under different perceived risks, coupling them with Markov chains to capture the dynamics of disease progression. We analyze this model by exploiting a linearized reformulation to incorporate long-term policy evaluation aligned with patient status. Moreover, we introduce a new cross-entropy-based loss function designed to maximize the benefits of prediction-based physician decisions. Finally, we propose a gradient-based optimization algorithm by implementing a differentiable approximation of discrete decision process through a benefit-based probabilistic mapping. We construct and evaluate our model using a large public dataset. Numerical experimental results demonstrate the distinct value of each component: (1) far-sighted risk prediction extended 3.16 quality-adjusted life-years (QALYs) per 1,000 patients, (2) physician decision-based risk warning provided 3.17 incremental

QALYs/1,000 patients, and (3) the interaction value is quite marginal with improvement no more than 0.01 QALYs/1,000 patients.

4 - Leveraging Patient Flexibility: a Short-Notice Policy for Outpatient Scheduling

Jue Wang, Zhejiang university, Hangzhou, China, People's Republic of, Zheng Zhang

In appointment-based outpatient systems, patients often allocate extra time for visits to accommodate uncertainties, creating inherent flexibility in arrival times. This study proposes a threshold-based short-notice policy that dynamically adjusts patient arrival times to enhance operational efficiency. Our approach divides patients into three groups—observation, buffer, and variable—based on their appointment sequence. After observing the first k patients' service durations, the system compares the waiting status against a predefined threshold to determine whether to notify the last m patients to arrive earlier. The buffer phase ensures notified patients have sufficient time to adjust their schedules, while intermediate patients remain unaffected to maintain scheduling robustness. We formulate this problem as a two-stage stochastic programming model, incorporating uncertainties in patient punctuality and service durations, with the objective of minimizing total costs, including waiting time, physician idle time, overtime, and notification efforts. Numerical results demonstrate the effectiveness of our policy.

5 - Is Faster Always Better? Evidence from An Online Medical Consultation Platform in China

JINXIN ZHANG, LINGNAN University, Hong Kong, Hong Kong, Liping Liang

The widespread adoption of online medical consultation platforms (OMCPs) has significantly transformed health service delivery. Despite its growth, the impact of response time, prevalent in online services, on patient satisfaction remains under-explored, particularly across different service stages. This study examines how pre-service and during-service response time separately affects patient satisfaction and uncovers inverted U-shaped relationships. Additionally, it explores how a patient’s regional affiliation, familiarity with a physician, and the physician’s professional status moderate these effects, offering insights for improving patient experiences on OMCPs.

Keywords: Response time; Patient satisfaction; Online medical consultation platforms.

MC09

Bras Basah B

Accelerating Maritime Decarbonization with AI

Invited Session

Maritime AI

Chair: Ran Yan, Nanyang Technological University, Singapore, Singapore

1 - Promotion Strategy for Shore Power Utilization in a Carbon Trading System: Vertical Alliances

Ning Ma, Dalian Maritime University, Dalian, China, People's Republic of, Xiang Zhang, Jingwen Li, Liwen Zhang

The greenhouse gas emissions and air pollution from the shipping industry account for a significant portion of global emissions, with a substantial share concentrated in port areas. Shore power(SP) is an effective measure to reduce CO_2 emissions at ports. However, the high installation and operational costs of SP infrastructure have hindered its widespread adoption.

To address this challenge, mitigate the cost burden of SP usage, and enhance competitiveness in an intensely competitive market, contractual vertical alliances present a viable solution for the shipping industry. This study examines the contractual effects and vertical choices within two parallel shipping supply chains operating under the same carbon trading market. By constructing a contract selection matrix, we derive equilibrium solutions and optimal strategies, identifying revenue-sharing contract conditions that facilitate the adoption of SP while

enhancing the profitability of both ports and shipping companies. Finally, based on the conclusions, we provide insights and recommendations for shipping companies, port authorities, and market participants.

2 - Ai-Driven Dynamic Vessels Speed Optimisation Under Just-in-TIME Arrival Implementation for Fuel-Efficient Maritime Operation

Ran Yan, Nanyang Technological University, Singapore, Singapore, Lei Liu

As a promising operational measure, Just-In-Time (JIT) Arrival concept seeks to reduce fuel consumption and emissions by optimising vessels' arrival time and enabling more efficient speed management during voyages. This study advances the practical implementation of JIT Arrival by developing an integrated prediction–optimisation framework that enables dynamic vessel speed optimisation under real-world operational and meteorological uncertainties. In the first step, a high-resolution deep neural network (DNN) model is developed to predict vessel fuel consumption across diverse operating conditions. The training dataset, which represents unit fuel consumption under various navigational and meteorological conditions, is constructed by integrating noon report and AIS data through a physical resistance-based baseline. In the second step, a multi-objective optimisation approach based on NSGA-II is introduced to dynamically adjust vessel speed across segmented voyage sections. The DNN prediction model is embedded as the fitness function to jointly minimise emissions and arrival time deviations under JIT Arrival requirements. Given that the Port of Singapore launched the JIT Planning and Coordination Platform for container ships in October 2023, this study focuses on a representative container ship calling at the Port of Singapore as a case study. The proposed framework enables vessels to establish dynamic replanning mechanisms that achieve significant emission reductions without compromising schedule reliability. Furthermore, by assessing vessel responsiveness to schedule adjustments based on discrepancies between estimated arrival time and assigned JIT Arrival time, the study provides actionable insights for designing effective port-side strategies to support JIT Arrival implementation.

3 - Evaluation of International Maritime Network Configuration and Impact of Port Cooperation

Shinya Hanaoka, Institute of Science Tokyo, Tokyo, Japan

This study evaluates the impact of port cooperation on port hierarchy through a network analysis approach. Specifically, we apply various complex network measures and link prediction techniques using graph embedding methods. In the context of intensified competition among container ports, achieving a higher position in the port hierarchy is crucial, as such ports play a central role in regional economic development by facilitating container cargo trade. To enhance their competitiveness, ports—particularly those in close geographic proximity—often form cooperative relationships aimed at leveraging synergies such as economies of scale.

We focus on the case of the aggregated Hanshin Port (comprising Kobe and Osaka, Japan) within the Japanese cargo network, and compare its hierarchical position with that of its major regional competitor, Busan Port (South Korea). Our analysis reveals that port cooperation contributes positively to an elevated position within the port hierarchy and leads to significant changes in network configuration, including the formation of more interconnected port communities. However, cooperation alone is insufficient to surpass a dominant competitor in the hierarchy.

These findings offer valuable insights for policymakers and port authorities seeking to enhance port competitiveness through collaborative strategies. In particular, the results are highly relevant for regions characterized by cargo networks with nearby dominant hub ports, such as those found in Japan.

4 - Unlocking Renewable Fuels for Green Shipping Corridors: a Data-Driven Analysis Framework

Jing Wang, Nanyang Technological University, Singapore, Singapore, Shuai Jia

To accelerate maritime decarbonization in alignment with the International Maritime Organization's (IMO) greenhouse gas reduction strategy, green shipping corridors (GSCs) have emerged as a pivotal mechanism for piloting low- and zero-emission solutions. Yet, a comprehensive, data-driven evaluation framework to assess the practical viability of renewable marine fuels in these corridors remains underdeveloped. This study develops a data-driven analysis framework that integrates Automatic Identification System (AIS) data with high-resolution meteorological and oceanographic datasets to refine predictions of vessel-level fuel consumption under real-

world operating conditions. Eight decarbonization pathways are examined, spanning four renewable fuels (compressed/liquefied hydrogen, ammonia, and methanol) deployed via dual-fuel engines and fuel cells systems. The assessment framework is structured around three pillars: vessel operational performance, environmental impacts, and financial implications. Operational analysis focuses on the implications of retrofitting vessels for renewable fuel adoption, particularly with respect to fuel storage requirements and bunkering frequency. Environmental performance is evaluated using a dual-level lifecycle assessment (LCA), incorporating both well-to-tank (WTT) emissions for fuels and cradle-to-grave emissions for supporting infrastructure. The financial dimension includes lifecycle cost (LCC) for fuel producers, total cost of ownership (TCO) for shipowners, and marginal abatement cost (MAC) for policymakers. A case study on the Rotterdam to Singapore GSC reveals that decarbonization potentials range from 72% to full carbon neutrality, depending on the fuel and technology selected. This work provides a structured decision-support framework for evaluating renewable fuels in global shipping industry and contributes to informed planning and investment toward achieving maritime decarbonization goals.

5 - A Bi-Level Robust Optimization Framework for Coordinating Post-Disaster Personnel and Material Aid

Jingwen Li, Dalian Maritime University, Dalian, China, People's Republic of, Xiang Zhang, Ning Ma, Yingming Yao, Jianhao Xu

Effective emergency logistics are essential for safeguarding lives, protecting property, and promoting sustainable recovery following disasters. This study examines the complexities of resource distribution in sudden disaster scenarios, emphasizing the interdependence between personnel and material support. Recognizing the urgency of rapid aid delivery, our research develops a bi-level robust optimization framework that addresses both coordinated allocation and transportation scheduling challenges. The upper tier of the framework minimizes rescue time through optimal routing of support from assisting cities to affected regions, while the lower tier focuses on cost-efficient distribution from assisting cities to distribution hubs. A robust optimization formulation with budget constraints is employed to cope with uncertainty challenges in demand, potential disruptions at distribution hubs, and transportation delays due to secondary disasters. By utilizing the Karush-Kuhn-Tucker conditions alongside the big-M method, the nonlinear bi-level problem is transformed into a mixed-integer linear programming formulation that is computationally tractable. Extensive computational experiments and sensitivity analyses demonstrate that, although the proposed bi-level approach may require longer solution times compared to the leader’s and follower’s perspective models, it yields balanced and effective outcomes for all stakeholders. Notably, the findings reveal that under increasing uncertainty, overall costs and rescue times do not necessarily rise in tandem. Moreover, the approach achieves equilibrated results even under multi-leader multi-follower dynamics, and cities providing similar levels of aid tend to maintain cost stability. These insights offer valuable technical guidance for decision-makers aiming to enhance the resilience and sustainability of emergency logistics systems and disaster response operations.

MC10

Hullet Room

Advances in Optimization for Resilient and Data-Driven Power Systems

Invited Session

Power Grid

Chair: Shunbo Lei, Chinese University of Hong Kong, Shenzhen, Shenzhen

1 - Power Grid Topology: Optimization, Transition and Stabilization

Tong Han, Nanyang Technological University, Singapore, Singapore

This talk provides an overview of recent efforts in developing optimization and control methods to enable better utilization of the topological flexibility of power grids. First, we design learning-based methods for solving common network topology optimization problems, including distribution network reconfiguration and transmission switching. Second, we develop optimization-based approaches to ensure a smooth transition from

the initial topology to the obtained optimal topology. Finally, we explore transient topology control strategies aimed at stabilizing the power grid following large disturbances.

2 - Making Modern Buildings Grid-Responsive Using Data, Optimization, and Control

Maomao Hu, Department of the Built Environment, National University of Singapore, Singapore, Singapore

Building energy flexibility has been increasingly demonstrated as a cost-effective solution to respond to the needs of energy networks, including electric grids and district cooling and heating systems, improving the integration of intermittent renewable energy sources. Our work investigates the demand response potential of large Heating, Ventilation, and Air Conditioning (HVAC) systems based on an extensive set of measurements from several commercial buildings. At the building level, we found that increasing the cooling setpoint by two degrees Fahrenheit led to a 13–28% reduction in daily cooling loads. In addition, we developed a non-intrusive, data-driven load disaggregation method to examine the spatial variability of energy use and flexibility. Our analysis revealed significant heterogeneity across and within buildings: a small number of zones accounted for a disproportionately large share of both energy consumption and flexibility. Interestingly, the most energy-intensive zones were not always the most flexible. These findings enable the electric grid or district energy system operators to regard the controlled energy-flexible entities as a fleet of zones instead of a fleet of buildings and help unlock the possibility for targeted demand flexibility strategies that balance zone-by-zone energy reduction with zone-by-zone costs to occupants. This work offers a blueprint for constructing energy-response functions at both building and zone levels and underscores the importance of repeated testing across diverse weather conditions. Providing statistical performance guarantees will be essential for scaling demand response technologies to reliably meet peak electricity demands.

3 - Fundamental and Complete Resolution of Large Scale Unit Commitment

JIANGHUA WU, Shenzhen Research Institute of Big Data, Shenzhen, China, People's Republic of

Large-scale Unit Commitment (UC) problems in power systems are challenging due to their mixed-binary nature, and the complexity to obtain an optimal solution increases exponentially as the problem size increases, limiting solution quality or problem sizes that can be effectively solved. In this talk, a fundamental and complete resolution of UC problems for near-optimal solutions with quantifiable quality in a computationally efficient way is presented. The novelties include (1) a novel decomposition and coordination approach exploiting the exponential reduction of complexity upon decomposition, with subproblems generally non-NP-hard; (2) an innovative and systematic way to tighten subproblem formulations offline, noting that if the convex hull of a subproblem can be directly delineated by its constraints, then the subproblem can be solved by using linear programming methods without much complexity; and (3) a hybrid learning-optimization mechanism where machine learning predicts subproblem solutions, and optimization or heuristics are applied when predictions are inadequate. Numerical testing demonstrates superior performance of the approach, and points a brand new way to formulate and solve large scale UC problems.

4 - Secured Data-Driven Methods for Operation and Control of Modern Power Systems

Ziming Yan, Nanyang Technological University, Singapore, Singapore

This study presents a secured data-driven framework for the intelligent operation and control of modern power systems, addressing the growing complexity brought by renewable integration, aging infrastructure, and socio-economic uncertainty. The proposed approach establishes a multi-layered architecture that unifies physical modelling with multi-source data across three core dimensions. At the local control layer, frequency regulation is enhanced through the fusion of physical mechanisms and data-driven deep reinforcement learning, bridging the Laplace inverse dynamics and policy reinforcement actor gradients. At the system operation layer, secure AI agents are developed by embedding physical constraints for Lagrangian and KKT conditions into policy optimization, enabling adaptive and safety-aware control. At the strategic planning layer, multi-modal asset risk assessment integrates condition monitoring, degradation models, and external factors to shift asset management from reactive replacement to proactive optimization. Collectively, this work contributes a holistic, explainable, and resilient AI control paradigm that forms a closed-loop from perception to strategic decision-making, ensuring security and adaptability in next-generation power systems.

MC11

Morrison Room

Application of Choice Modeling on Sustainability

Invited Session

Revenue Management and Pricing

Chair: Xiaobo Li, National University of Singapore, Singapore, Singapore

1 - Quantifying Framing Effects: Insights from Joint Assortment and Pricing Experiments

Zhenzhen Yan, Nanyang Technological University, Singapore, Singapore, Changchun Liu, Limeng Liu

In the context of framing effects, the perceived appeal of individual items is not static but rather contingent upon the context in which they are presented. Despite extensive exploration of framing effects in marketing and psychology literature, limited research has addressed their integration into discrete choice models and their implications for revenue management. This paper contributes to the field by introducing a novel utility-based choice model that accounts for framing effects. In our proposed model, the utility derived by a customer from a presented item is influenced by the composition of the offered assortment. We propose a separable representative consumer model that incorporates both position and size to capture framing effects. Leveraging the separable structure, we develop a two-stage approach—a joint assortment and pricing experiment strategy—to acquire market share insights. Specifically, we design a pricing experiment policy to identify the perturbation function and an assortment experiment policy to learn the framing effect. This approach enables effective discernment of the framing effect and perturbation functions, adhering to a predetermined threshold of ϵ with a polynomial number of experiments. Extensive numerical experiments validate the efficacy of our data-driven approach. Our method accurately predicts customer choices and outperforms existing approaches in addressing model misspecification. Additionally, we present a case study in the airline industry, demonstrating that our proposed method effectively learns customer purchasing behavior compared to the spiked MNL model.

2 - A Dynamic Mechanism for Achieving Sustainable Quality Supply

FANG LIU, Durham University Business School, Durham, United Kingdom, Tracy Lewis, Jing-Sheng Song

We model a supply chain where a retailer relies on an economically weaker supplier to supply a high-quality key material. For sustainable supply, the production process must comply with socially and environmentally responsible standards over time. However, compliance costs fluctuate due to an underlying state that is privately observed by the supplier, who may thus fail to comply, risking future supply capabilities. While the retailer can invest in reducing the cost, the outcome remains uncertain. To address these challenges, we introduce a sustainability index to monitor and manage the supplier's compliance within a multi-period, index-based agreement to optimize the supply chain's overall sustainability. The proposed mechanism modifies the sustainability index and corresponding payment terms for each party in each period, ensuring it reflects current circumstances and the worst-case scenarios, and anticipates future changes. This structure promotes voluntary participation, truthful information sharing, nonnegative payment to the supplier, and the first-best investment. Through two examples within our general framework, we illustrate the payment term dynamics and the importance of continuous investment. The examples also highlight how contract length, investment success probability, and the nature of investments are critical for the success of sustainability programs, pinpointing when external support from non-profit organizations is most effective.

3 - Efficient Heuristics for Warehouse Assortment Selection with Performance Guarantees

Hongyuan Lin, University of Science and Technology of China, Hefei, China, People's Republic of, Xiaobo Li, FANG LIU, Sentao Miao

The rapid growth of e-commerce has significantly expanded product variety, intensifying the challenge of efficient order fulfillment. To enable fast delivery, e-retailers strategically position front-end warehouses near customers. However, these warehouses operate under strict capacity constraints, making assortment selection a critical decision in minimizing fulfillment costs. Ideally, each customer order can be fulfilled entirely from a single front-end warehouse. Otherwise, reliance on back-end support increases operational costs, carbon emissions, and delivery lead times, ultimately degrading service quality. In this work, we propose a scalable and

flexible heuristic for warehouse assortment selection in both single- and multi-warehouse settings. The heuristic accommodates diverse warehouse- and order-specific cost structures while providing rigorous performance guarantees. Specifically, we establish that its fulfillment cost remains within a constant factor of the optimal solution, where the approximation ratio depends on the largest order size when fulfillment costs are uniform, and on cost disparities across warehouses when fulfillment costs vary. Extensive numerical experiments further demonstrate that the proposed heuristic consistently outperforms existing methods across both single- and multi-warehouse settings.

4 - Efficiency Versus Accuracy: The Optimal Level of Decentralization in Humanitarian Operations for Sudden-Onset Disasters

Yinuo Lin, National University of Singapore, Singapore, Singapore, Wenjie Huang, Xiaobo Li, Lei Xu

Humanitarian organizations (HOs) face operational challenges during sudden-onset disasters, such as constraints on staff, storage, and transportation, delaying critical donations to people in need (PIN). These inefficiencies create a trade-off: centralized operations enable accurate demand estimation but face bottlenecks, while decentralized donations, where individual donors distribute supplies based on preferences, offer faster responses but risk supply-demand mismatches.

To address this, we propose a model that partially decentralizes humanitarian operations while accounting for random demand. In the deterministic case, the optimal decentralization level reduces to a one-dimensional piecewise-linear convex optimization problem, enabling a novel algorithm to find all solutions efficiently. We prove that the objective function has at most 3n+2 pieces, leading to an improved O(n)-complexity algorithm for n affected zones. These results extend to the stochastic case, accommodating $|\Omega|$ demand scenarios with complexity $O(|\Omega|n)$.

We derive closed-form expressions for the optimal decentralization level, facilitating detailed comparative analyses. Our model effectively incorporates demand uncertainty, donor preferences, and fairness considerations, all of which can be efficiently addressed using the proposed algorithms.

As a case study, we apply this framework to early-stage humanitarian aid distribution in Ukraine during the 2022 Russo-Ukrainian War. Results show that partial decentralization, when HO efficiency discounts are high, is nearly three times more effective than full centralization. A 33% decentralization level can reduce demand shortage costs by up to 80% compared to full decentralization via postal services. This highlights the potential of partial decentralization in improving disaster relief operations.

MC12

Moor Room

Advanced Models in Revenue Management and Pricing

Invited Session

Revenue Management and Pricing

Chair: Guang Li, Queen's University, Kingston, ON, Canada

Co-Chair: Ragavendran Gopalakrishnan, Smith School of Business at Queen's University, 143 Union Street West, Kingston, ON, K7L 3N6, Canada

1 - Add to Order in Contingent Free Shipping on Online Grocery Platforms

Yuxuan Tang, National University of Singapore, Singapore, Singapore, Xiaobo Li, Rowan Wang, Xuan Wang

Problem definition: This paper examines the impact of a newly introduced "add to order" (ATO) feature by online grocery platforms on contingent free shipping (CFS) policies. The ATO feature allows customers to add items to an existing order without incurring additional shipping fees, provided the order preparation has not yet begun. **Methodology/results:** Utilizing a utility-based model and real-world transaction data, we analyze the effects of the ATO feature on customer behavior, customer surplus, platform profitability, and optimal CFS

settings. Our findings reveal that ATO increases customer surplus, boosts conversion rates, and promotes order padding. While ATO can increase platform profits by reducing operational costs through order consolidation, this benefit depends on existing shipping fee structures and consolidation practices. Platforms with low shipping fees and efficient consolidation strategies may actually experience a decline in profitability upon adopting ATO. **Managerial implications:** Our results suggest that the optimal CFS strategy for platforms implementing the ATO option is either to set a very high shipping fee or to impose a modest shipping fee on all customers. Our research provides actionable insights for the strategic design of CFS policies and underscores the complex trade-offs introduced by the innovative ATO feature in online grocery retail.

2 - Assortment Optimization in the Presence of Influencers

Guang Li, Queen's University, Kingston, ON, Canada

We investigate three scenarios of a platform's assortment optimization problem when some products are promoted by influencers. The first scenario employs a mixed logit model, with varying impacts of promotions across consumer segments. In the second scenario , we consider two types of consumer segments: the first type is divided into individual segments, each focusing exclusively on one of the promoted products; whereas, the second type considers the entire product assortment. Lastly, we explore a scenario where the platform makes the assortment decision without prior knowledge of which item will be promoted, with each product having a certain probability of promotion. In each scenario, we establish the structural properties of the optimal assortment and provide efficient or approximation algorithms to address the assortment challenges.

3 - When Strategic Customers Meet Strategic Servers: Individual and Social Optimization in Many-Server Queueing Systems

Ragavendran Gopalakrishnan, Smith School of Business at Queen's University, Kingston, ON, Canada, Yueyang Zhong, Amy Ward

We initiate the study of interactions between strategic customers and strategic servers in many-server queueing systems. In the seminal works of Naor (1969) and Knudsen (1972), arriving customers decide whether to join the system based on a trade-off between their reward from service and cost of waiting, while servers operate at fixed exogenous rates. The joining threshold (maximum queue length) induced by self-interested customers at equilibrium is higher than socially optimal (i.e., they "over-join"), a result known as Naor's inequality. This is because they fail to internalize the congestion externality they impose on other customers. Although Naor's inequality is robust across a wide spectrum of queueing models, its validity when servers are also strategic (in choosing their work-speeds) is unknown. Within a large-system asymptotic framework, we show that strategic servers provide slower than socially optimal service at equilibrium, thereby deterring customers from joining the system, and when this effect outweighs customers' tendency to over-join, Naor's inequality is reversed. Next, we propose a scheme that charges an entry fee from customers and offers a performance-based compensation for servers, which realigns individual incentives with the social optimum. Finally, we analyze how the overall welfare gains at the social optimum (relative to the equilibrium) are distributed across customers and servers, in order to unearth potential equity trade-offs in policy design.

4 - Expanding Platform Capacity Through Asset Rental Programs: Design and Implications

Deyin Ji, Singapore Management University, Singapore, Singapore, Guiyun Feng, Ioannis Bellos

In recent years, the growth of on-demand platforms has been hindered by the lack of available workers. To address this challenge, platforms have been introducing short-term asset rental programs for individuals who are interested in providing services but do not own a qualifying asset or are reluctant to use their own. Different designs of such programs have been adopted with some of them discriminating between workers renting versus using their own assets by either: i) imposing minimum requirements on service time of asset-renting workers, ii) offering different wages to the two pools of workers, or iii) combine the two previous programs. These practices, however, have raised fairness concerns regarding the treatment of workers who do not own an asset. Motivated by these concerns and the divergent platform practices, we attempt to understand the role played by the discriminating rental policies and their implications for workers and the whole system. To that end, we develop a game-theoretic model and analyze different rental program designs. Our study reveals that if a platform discriminates by offering different wages, it should pay higher wages to the workers who rent, which challenges the prevailing practice of offering lower wages. Contrary to fairness concerns, all discriminatory designs can benefit workers and under certain conditions, can also lead to higher platform profit and more customers served.

Our research provides valuable managerial insights for designing effective rental programs and a comprehensive understanding of their social impact.

5 - Platform-Enabled Farmer Cooperation for Sustainable Agriculture: Coalition Formation and Pricing Decisions

Rajeev Ranjan Tripathi, Indian Institute of Management Bangalore, Bengaluru, India

This paper investigates how platform businesses can promote sustainable agricultural practices by managing farmer heterogeneity. It focuses on two distinct types of farmers - those who prioritize sustainability and those who do not - and examines how platforms can design operational strategies that encourage collaboration among them to scale sustainable practices.

The platform considered generates positive cross-side externalities: as more farmers collaborate and form groups, buyers' willingness to pay increases, thereby benefiting the entire ecosystem. The paper makes two key contributions. First, it proposes a cooperative game-theoretic approach to designing a fair pricing mechanism on the platform. This approach involves reframing the pricing problem as a payoff allocation problem and determining fair pricing through the Shapley value that ensures equitable compensation for participating farmers. Second, the paper highlights that farmers prioritizing sustainability form coalitions and take joint actions, unlike the other types of sellers who are purely driven by individual self-interest and remain singleton. The central question addressed is: how can a stable coalition of sustainability-oriented farmers be formed and maintained on the platform? To answer this, the paper uses a cooperative game theory model of coalition formation.

The key findings inform platform design choices that incentivize sustainable practices, align economic interests, and foster inclusive growth in agriculture.

MC13

Bailey

Advanced Process Monitoring, Control, and Optimization

Invited Session

Semiconductor Manufacturing x AI

Chair: Jinhua Ding, Micron Semiconductor Asia, Singapore, N/A, Singapore

Co-Chair: Lin Miao Zhang, Micron Technology, Singapore, N/A, Singapore

1 - Leveraging Evolving Analytical Insights for Enhanced Tool Performance in Semiconductor Manufacturing

Chen Miao, Tokyo Electron, Singapore, Singapore

As the semiconductor industry rapidly evolves with increasing technological complexity, the demand for advanced process control has never been greater. In this presentation, we will explore how TEL is leveraging edge computing systems and extensive sensor data from the equipment to develop cutting-edge solutions aimed at enhancing process control and optimizing chamber-to-chamber matching across our etch systems. By leveraging integrated server systems deployed at customer sites, we are able to continuously integrate real-time analytical insights based on actual performance metrics. This allows us to refine control signals alongside ongoing advancements in processing technology. Attendees will learn about customizable tool health monitoring options and automated recovery strategies designed to achieve productivity goals through versatile data utilization. We will share case studies demonstrating how tailored solutions derived from these insights deliver effective feedback for tools—ultimately enhancing future processes while promoting advanced control measures that boost overall tool performance.

2 - Enhanced Analytics for Lithography: Incorporating Advanced Scanner Models in Semiconductor Manufacturing

Wim Tel, ASML, Veldhoven, Netherlands

Semiconductor device manufacturing seems to be on a collision course between ever more aggressive processing and tightening of specifications. This puts ultimate pressure on correctability and process monitoring. In order to

maximally exploit the correction capability of the lithography tool a new scanner control paradigm is proposed. Novel ways for setup, control, monitor, and deepdive analysis are needed. For this purpose a highly accurate scanner simulation model is introduced: Digital TWINSCAN

3 - Proactive Monitoring and Alarm Aggregation in Fault Detection Systems in Semiconductor Manufacturing

Chunhui Zhou, Micron Semiconductor Asia Pte. Ltd., Singapore, Singapore

Fault Detection (FD) through real-time monitoring of sensory data from tools is essential for quality control in semiconductor manufacturing. With the rapid advancement of data science and artificial intelligence in the semiconductor industry, the number of FD analytics solutions has steadily increased over the years. This growth in FD solutions necessitates more reactions from fab engineers, highlighting the need for a streamlined process to review various FD alarms. Additionally, although FD alarms indicate significant deviations from baseline performance, fab engineers find that some of these alarms are false and have no impact on the product. Therefore, it is desirable to have a solution capable of filtering out noisy signals from FD alarms. This work proposes a methodology to consolidate different FD alarms at various granularities into tool-level alerts. It utilizes a unique logic based on alarm rate monitoring, weighted scoring, and alarm rules to ensure that only the most significant and chronic deviations are flagged. The proposed approach has proven effective in filtering out false alarms and reducing data reviewing efforts from fab engineers at Micron.

4 - Optimizing Post-Maintenance Tool Performance: a Predictive Model for Reducing Qualification Cycles

Yuan Ma, Micron Technology, Sinapore, Singapore

In the realm of advanced manufacturing, maintaining optimal tool conditions post-maintenance is a significant challenge. This study addresses the problem of excessive deviations in tool performance after preventive maintenance, which necessitates multiple qualification cycles to restore target conditions. Utilizing a comprehensive dataset, we developed a predictive model that significantly reduces the number of qualification cycles required. Our approach has achieved a significant reduction in qualification time, enhancing overall efficiency and productivity.

5 - Maximizing Review Sem Effectiveness for Process Variation Monitoring

Poh Boon Yong, KLA, Milpitas, CA, United States, Alex Lim

Over the years, semiconductor manufacturers have invested in a fleet of review Scanning Electron Microscopy (SEM) tools to inspect and classify hundreds to thousands of sampled defects post-wafer inspection. These tools offer higher throughput but lower resolution compared to more expensive metrology tools. Despite their resolution limitations, this paper explores the extension of review SEM tools to monitor a fab’s process performance and variation through a comprehensive set of metrology measurements provided off-tool by KLA's aiSIGHT software. This software leverages design layout and complex models, including deep learning-based image processing, to analyze SEM images. This enables the fab to utilize a broader range of tools for monitoring process variation through overlay, critical dimension (CD), edge placement error (EPE), pitch walk, and pattern placement error (PPE). These measurements facilitate the study of local, across-die, across-reticle, and global effects. Consequently, chip manufacturers can expand wafer coverage by monitoring more sites per wafer and additional wafers, all without impacting wafer throughput. This increased coverage has led to the discovery of more insightful wafer signatures in a timely manner, enabling better corrective decisions and ultimately reducing costs in high-volume manufacturing.

MC14

Wilberforce

Manufacturing Science & Supply Chain

Invited Session

Semiconductor Manufacturing x AI

Chair: Lin Miao Zhang, Micron Technology, Singapore, N/A, Singapore

Co-Chair: Xun Li, 1 North Coast Avenue, 737663

1 - Sequential Models to Optimize Assembly and Test Capacity at Micron Technology

Shantanu Bagchi, Micron Technology, Hyderabad, India, Pauline Romanoff

2 - Learning Vector Representations of Complex Production States for Scheduling in Semiconductor Manufacturing

Christophe Senger, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany

Semiconductor manufacturing scheduling is among the most complex tasks in production planning, due to its stochastic nature, large-scale job shop structure, and overall system complexity. A standard fab typically contains more than 1000 machines and several thousand WIP lots, a situation that is further complicated by frequent reentrant flows, machine failures and maintenance events. Conventional scheduling methodologies, founded upon mathematical programming, encounter challenges in terms of scalability. Conversely, heuristic approaches show low flexibility, resulting in suboptimal performance.

In this context, reinforcement learning has emerged as a promising scheduling alternative, particularly through the training of dispatching agents. While other state modeling techniques may offer more precise representations of the system state, they often require regrouping highly dynamic and heterogeneous elements. This results in variable and high-dimensional input structures, which are complex to handle and difficult to utilize in downstream calculation models. The process of feature engineering becomes excessively complex and use-case specific and fails to provide a consistent input format for training robust models.

This work proposes a novel approach to encode complex production states as graphs and embed them into fixed-sized vectors. These vectors retain rich structural and temporal information, enabling their use for both static evaluation of current production states and predictive analysis of future performance. We show that, even after compressing the complex production state, the resulting representation still contains enough information to estimate both dynamic performance indicators such as throughput, job lateness, and cycle time and static characteristics like WIP distribution and remaining processing time.

3 - Challenge of Supply Chain Adaptability: a Prospective Study of Semiconductor Supply Chains

Shardul Phadnis, Asia School of Business, Kuala Lumpur, Malaysia

Supply chain adaptability (SCA) is critical to maintaining competitiveness in an uncertain future, as the business environment is shaped unpredictably by technological advances, geopolitical tensions, trade wars, demographic shifts, extreme weather events, and more. Yet, scholarly knowledge of SCA relies almost entirely on retrospective studies that benefit from hindsight and surveys assuming prescient respondents. How uncertainty about the future, emerging from a complex combination of driving forces, affects strategic choices about SCA remains poorly understood. This study addresses this gap in the literature by adopting a prospective research design embedded in an immersive industry study of the semiconductor industry and using generative artificial intelligence (GenAI) by making three contributions. One, it shows that a wide variety of driving forces shape the environment in which supply chains operate and need to be considered when making decisions about SCA. Two, it shows that building SCA is likely to require undertaking initiatives in different functional areas both internal and external to the firm, and thus working on multiple fronts. Three, determining the sequence in which these investments should be made is not trivial. This is particularly challenging when budget constraints limit the number of concurrent investments. Furthermore, the interdependence of these investments in assets and capabilities is not obvious. Together, these findings suggest that SCA is far from trivial and requires a systematic process for approaching it. Building on this, we conclude the paper by presenting directions for future research.

4 - Sequential Models to Optimize Assembly and Test Capacity at Micron Technology

Shantanu Bagchi, Micron Technology, Hyderabad, India, Pauline Romanoff

In this article, we describe a sequential stochastic optimization model implemented by Micron Technology to identify optimal assembly and test manufacturing capacity that achieves fastest cycle time and maximum flexibility at minimum cost. Given that Micron Technology operates in the semiconductor industry, it inherits characteristic demand volatility of the market, making flexibility and time-to-market key factors to become successful. This analytical approach comprises of several mathematical models that are executed sequentially to

optimize the global capacity at Micron. First, a “flexible-capacity model” identifies the optimal global capacity which can mitigate the future demand uncertainty. Next, an “insource-vs-outsource capacity model” determines the optimal split of this total global capacity between in-house factories and subcontractors that minimizes the expected cost. Then, more detailed capacity at the equipment level by site are optimized. Subsequently, a “site qualification optimization” model is executed to determine which manufacturing facility(s) should be enabled to manufacture which product(s) so that maximum flexibility is attained at minimum cost. Finally, a “minimum utilization optimization” model is developed to provide operating guidelines to the factories to trade-off between underutilization costs and inventory costs.

We also describe the implementation process where we identify the stakeholders involved in various steps of the capacity optimization problem and how the objectives of the individual stakeholders organically provide the guardrails to the solutions obtained by the methodology. We also highlight the benefits of using of sequential solving of smaller sub-problems over solving a large, consolidated problem in terms of speed, computational complexity and business justification.

MC15

Fairmont Ballroom

Scheduling and Optimization in Transportation

Contributed Session

Transportation Science and Logistics

Chair: Maximiliano Cubillos, Politecnico di Milano, Milan, N/A, Italy

1 - Workflow Scheduling Using Graph Neural Networks

Rakesh Nagi, Singapore University of Technology and Design (SUTD), Singapore, Singapore, Vivek Rajopanth, Asser Tantawi

Workflow Scheduling is a combinatorial optimization problem that aims to allocate jobs with predefined precedence relationships to available resources in the most efficient manner. While heuristic approaches can offer rapid solutions, they often sacrifice solution quality. Mixed Integer Linear Programming (MILP), on the other hand, provides optimal solutions for a given problem instance but is computationally expensive. In many real-world systems, job types tend to repeat within specific time frames, creating an opportunity to exploit historical patterns. In this work, we harness the power of Graph Neural Networks (GNNs) to learn from historical near-optimal schedules (generated via MILP) and predict scheduling patterns for new instances. Specifically, we train the GNN to produce “heatmaps” that guide a novel heuristic approach, effectively combining the precision of MILP with the predictive capabilities of deep learning. Our experiments show that this hybrid method outperforms traditional baselines by achieving superior schedule quality with minimal computational time.

2 - Reverse Logistics Network Design for End-of-Life Power Batteries Considering Different Storage TIME Limitations Under Uncertain Quantity and Quality of Returns

Limin Ai, Chongqing University, Chongqing, China, People's Republic of, Zhang Wu

With the rapid proliferation of new energy vehicles, the number of end-of-life (EoL) power batteries is increasing rapidly. Effectively recycling these batteries has become an urgent challenge that needs to be addressed. Given that the quality condition of EoL power batteries directly affects their storage time at collection stations, we design a reverse logistics network (RLN) that accounts for different storage time limitations. The model also incorporates features such as uncertainties in the quantity and quality of returned EoL power batteries. Furthermore, to enhance the utilization efficiency of collection stations, we further propose a multi-type collaborative collection station (CCS) scheme. Therefore, this study proposes a multi-period, multi-scenario, two-stage hybrid stochastic-possibilistic mixed-integer nonlinear programming (MINLP) model. By employing the Big-M method and the α-acceptability approach, we transform the formulation into a crisp mixed-integer linear programming (MILP) model. Due to the computational complexity of large-scale problems, we develop a hybrid metaheuristic algorithm that integrates Simulated Annealing combined with a Constructive Heuristic algorithm (CH-SA) and an Improved Adaptive Large Neighborhood Search (IALNS). A comparison with the GUROBI solver demonstrates the superiority of our proposed algorithm. Then we apply the model to a

real-world case study in Chongqing, China. Incorporating multi-type CCS scheme reduces the total cost by 8.24%. Finally, we perform sensitivity analyses on several key parameters and provide managerial recommendations for addressing future increases in the number of EoL power batteries and the reduction of limited storage time. These insights advance the design of RLNs for EoL power batteries.

3 - Exploiting the Buffer for Optimizing Integrated Yard and Quay Crane Operations

Qi Wang, Tsinghua University, Shenzhen, China, People's Republic of, Canrong Zhang

Container terminals play an important role in global shipping industry. As container traffic continues to grow rapidly, terminal operators face the challenge of improving throughput efficiency while managing limited space and operational constraints. This paper addresses the coordination of two key terminal operations—quay crane (QC) and yard crane (YC) operations—with an emphasis on managing apron buffers, a strategy that optimizes container flow and minimizes operational delays. We focus on a novel shared buffer strategy, where buffers positioned between QCs and YCs are accessible for tasks from all cranes. We formulate the problem as a mixed-integer linear program (MILP), distinguishing it from traditional approaches that assign dedicated buffers to specific cranes. This shared buffer system maximizes the utilization of limited apron space but introduces new challenges, such as buffer competition and synchronization issues. To address these challenges, we develop a buffer size-driven branch-and-price algorithm that leverages the optimal structural characteristics of the buffer to accelerate the branching process. Additionally, we propose a specialized bidirectional labeling algorithm to further enhance the subproblem pricing procedure.

Drawing on operational practices at Tianjin Port, we conduct extensive numerical experiments to demonstrate the effectiveness of the proposed model and algorithms. For small-scale instances, the buffer size-driven branch-and-price algorithm significantly outperforms commercial MILP solvers in terms of computational efficiency. For larger instances, the method produces high-quality solutions within a reasonable time frame. These results provide valuable insights into the role of shared buffers in container terminal scheduling and offer a promising approach to optimizing port operations.

4 - Dynamic Truck Scheduling Under Partial Arrival Information and Stochastic Service TIMES

Maximiliano Cubillos, Politecnico di Milano, Milan, Italy, Ola Jabali, Justin Goodson

This paper addresses the dynamic scheduling of inbound trucks at a warehouse, leveraging vehicle positioning technology that provides real-time arrival time estimates. Building on recent theoretical work that characterizes uncertainty in these estimates, we focus on practical deployment using real-world data from an industrial partner. The collected data includes trajectories of arrival time estimates, actual arrival times, and service duration. Contrary to standard modeling assumptions that rely on normally distributed noise, our analysis reveals distinct and recurring patterns in the behavior of these estimates. Specifically, we observe prolonged periods without updates, followed by abrupt increases—often caused by scheduled stops at other clients or mandatory driver breaks. These events, which offer lower bounds on en-route delays, exhibit considerable variability and are not captured by conventional models of arrival uncertainty. We also address the uncertainty in service times, which in practice are only revealed during operations. Our model incorporates this feature by dynamically updating beliefs about service times as unloading progresses. To integrate both sources of uncertainty, we formulate the problem as a partially observed Markov decision process aimed at minimizing the expected total waiting time of trucks. We propose a lookahead policy tailored to this setting and evaluate different strategies of leveraging arrival time estimates, including point estimates, anticipation of abrupt increases, Bayesian filtering, and robust formulations. Computational experiments demonstrate that explicitly modeling the evolving nature of arrival and service information leads to significant improvements over practice benchmarks.

MC16

VIP A

Sustainable and Ethical Innovation in Global Supply Chains

Invited Session

Sustainable and Socially Responsible Operations

Chair: Zelong Yi, N/A

1 - Synchromodality Service Planning and Procurement in a Global Logistics Platform: a Bi-Level Programming Approach

Xiangtianrui Kong, Shenzhen University, Shenzhen, China, People's Republic of, Lele Zhang

With the evolution of global supply chains and trade patterns, synchromodality has emerged as a crucial logistics model to meet shippers' demands for reliability, flexibility, and security. However, existing global logistics digital platforms face challenges in balancing supply and demand, decision-making mechanisms, and multi-objective collaborative optimization, which hinder the efficiency of synchromodality. To address these issues, this paper proposes an optimal synchromodality service planning and procurement model based on bi-level programming. The upper level is the synchromodality routing planning (SRP) model, which optimizes transportation mode selection and routing to minimize total transportation time. The lower level is the winner determination problem (WDP) model for transportation service procurement through combinatorial auctions, using the O-VCG mechanism to select carriers to minimize total transportation costs. Numerical experiments demonstrate that the model reduces transportation time by 10.3% and transportation costs by 6.6%, validating its effectiveness. This model provides new insights and tools for theoretical research on synchromodality and the optimization of international logistics platforms.

2 - Blockchain-Enabled ESG Reporting

Yelin Fu, Shenzhen University, Shenzhen, China, People's Republic of, Jian Zhong

The introduction of the Environmental, Social and Governance Reporting Guide (Guide) by HKEx in 2013, and the subsequent upgrade of the Guide’s reporting obligation to “comply or explain” in 2016, have significantly moved the dial for Hong Kong issuers’ ESG reporting. However, current ESG reporting is confronted with many bottlenecks: (1) the ESG data are not collected in a sophisticated manner; (2) the ESG data delivered to ESG reporting agencies may be manipulated; (3) the reporting process is not transparent to investors, NGOs and industry associations. This project proposes to innovate a suite of core technologies to form the foundation for developing a total solution to upgrade the ESG reporting industry. In addition to our previous R&D achievements in IoT technologies, this project will focus on several key innovations: (1) a fact-telling blockchain gateway is developed to evaluate fact level of ESG data and to transfer data into the consortium blockchain; (2) a n-Dimension (nD) consortium blockchain is built for ESG data sharing among different ESG report stakeholders; (3) the versioning smart contract mechanism is proposed to assess the consistency and transparency levels of ESG report in line with annual HKEx ESG reporting KPIs; (4) the token-based incentive mechanism is established to construct the fact index of ESG report; (5) a package of consortium blockchain-based analytics services to facilitate decision-making of ESG report stakeholders.

3 - A Novel Carrier-Shipper Matching Mechanism with Value-Added Services for the Online Logistics Platform

Xuan Yang, Shenzhen University, Shenzhen, China, People's Republic of, Zunxin Zheng, Xiangtianrui Kong

As online logistics platform (OLP) evolves from basic supply-demand matching tool to comprehensive logistics ecosystems, the integration of value-added services (VASs) has become a key factor influencing carrier-shipper matching decisions. Motivated by a real-life OLP case, this paper investigates the carrier-shipper matching problem with value-added service bundling (CSMP-VSB), aiming to jointly optimize matching strategies and service bundling decisions to enhance both matching numbers and user satisfaction. We propose a novel cluster and feature-based genetic optimization (CFGO) algorithm, which combines the strengths of clustering techniques and data-driven learning. Empirical analyses using real-world operational data from our collaborative partner demonstrate that proactive VAS bundling significantly improves matching performance, increasing the number of successful matches by 42.49%, carrier satisfaction by 22.98%, and shipper satisfaction by 17.36%. Sensitivity analyses further reveal diminishing marginal returns with increasing carrier numbers and VAS budgets, highlighting the importance of dynamically scaling resource allocation according to platform growth. Additionally, user preference analysis shows that prioritizing certain attributes, such as completion time, can significantly broaden feasible matches, whereas price-focused preferences tend to constrain matching

opportunities due to additional VAS-related costs. Similar to the evolution of modern e-commerce platforms, the proposed mechanism will effectively catalyze a paradigm shift in the online logistics marketplace — transitioning from conventional information-based models to AI-driven multi-service bundling and next-generation intelligentization.

4 - A Novel Auction-Based Logistics Service Matching Mechanism

Yangdi Wang, Beijing Institute of Technology, Beijing, China, People's Republic of, Suxiu Xu, Xiang Ji, Yu Ning

The emergence of the Cyber-Physical Internet (CPI) has revolutionized the global logistics ecosystem by building a bridge between the physical and virtual worlds. The cyber and physical layers constitute the CPI's connections. The cyber layer integrates, transmits, and handles logistics information in the digital realm, while the physical layer loads, transports, and delivers goods in the real world. This paper proposes efficient auction mechanisms, which introduce weights for the submitted bids, to construct CPI in certain areas across multiple cities. In our auctions, the builder is the auctioneer, and the contractors are the bidders. We also introduce affiliated contractors determined by the builder, who play a role in guaranteeing the execution of the auctions. We construct a graph where the cities are regarded as the vertices, and the connections are considered as the edges. We develop a one-sided weighted VCG (OW-VCG) auction on the graph. The OW-VCG auction achieves efficient allocation, incentive compatibility, and individual rationality. Next, we construct a multi-round ascending auction based on the primal-dual theory on the graph, also called the weighted primal-dual (WPD) auction. We prove that the WPD auction implements the same outcomes as the OW-VCG auction. In addition, we prove that the existence of affiliated contractors can improve the revenue received by the builder from the contractors and provide a strategy on how to deploy them. We finally test the strategy on a graph case. The results demonstrate the superior performance of the affiliated contractors in improving the revenue.

MC17

VIP B

AI in Complex Systems and Decision Making

Contributed Session

Artificial Intelligence

Chair: Stanley Nwoji, Harrisburg University of Science and Technology, 103 Constitution Terrace, Elkton, MD, 21921, United States

1 - Echo Warfare: Strategic Manipulation of Intangible Heritage and Decision-Making Implications

Sara AlNabet, Joint Warfare Training Center, Qatar Armed Forces, Qatar, Qatar

This paper introduces Echo Warfare, a novel non-kinetic doctrine centered on the strategic manipulation of intangible cultural heritage through systematic appropriation, replication, and subtle distortion. Unlike traditional warfare, Echo Warfare avoids direct physical or territorial confrontations, instead targeting societal cohesion, cultural resilience, and collective identities. By recontextualizing culturally significant practices, symbols, and narratives, adversaries subtly influence public perceptions and societal trust, causing profound, long-term generational impacts.

A critical dimension of Echo Warfare lies in its subtle yet profound impact on long-term military and strategic stability. By systematically embedding distorted or appropriated cultural narratives, adversaries gradually shape perceptions of identity and alter societal attitudes over generations. As a result, the cultural distinction between the targeted community and adversary becomes increasingly blurred, weakening overall societal cohesion and indirectly affecting national defense capabilities.

To effectively identify, analyze, and mitigate these subtle cultural manipulations, the research proposes adapting advanced computational modeling, predictive analytics, and scenario-based simulation methods widely employed in tactical and operational military contexts, but underexplored in strategic and political cultural analysis. Leveraging these techniques provides strategists and policymakers with robust analytical tools to foresee and counteract the long-term consequences of Echo Warfare without attributing vulnerabilities directly to any specific group. Ultimately, the paper highlights the necessity of interdisciplinary frameworks to preserve cultural integrity, societal resilience, and national security against sophisticated, non-kinetic threats. This paper

thus opens pathways for future research, encouraging deeper exploration of non-kinetic threats and the development of targeted analytical approaches to counteract their strategic implications.

2 - Utilizing Large Language Models and Decision Models to Reflect Human Thought Processes

Senyao Gu, Tsinghua University, Beijing, China, People's Republic of, Chen Wang

Electronic Health Records (EHRs) incorporate multimodal data, encompassing patient health status and contextual information influencing medical decisions. This study aims to develop a decision-making framework that integrates large language models (LLMs) with traditional decision models to examine the process by which individuals make medical decisions, utilizing objective medical knowledge and scenario-specific information. By training machine learning algorithms on patient health status data from EHRs, we evaluate treatment outcomes and probabilities, simulating lottery choice problems in medical decisions. Following this, we propose the use of LLMs to interpret scenario data within EHRs and the predictive outcomes from machine learning models. LLMs offer an advantage over conventional machine learning techniques by their ability to process unstructured information and infer from multimodal contexts using logical frameworks such as cumulative reasoning. Our research endeavors to refine the Mixture of Theories approach for real-world, multimodal decision-making contexts and investigates the preferences for different decision models in practical applications.

3 - Myerson Additive Explanations on Graphs (Myer): Advancement of Explainable Artificial Intelligence Using Graphical Approach

R K Amit, IIT Madras, Chennai, India, Rushikesh Bhimewar

The field of explainable artificial intelligence has evolved to bridge the gap between complex, obscure black-box algorithms and human users. These methods try to find what matters most to the models and how it matters. Feature importance answers this question, signifying the degree to which each feature contributes to the model’s final output. However, existing approaches for estimating feature importance face limitations by often overlooking intricate relationships between input features. Challenges arise, particularly when input features are correlated, leading to potentially misleading results. To address these challenges, the causal knowledge among input features is leveraged to estimate feature attribution. However, establishing the knowledge of causal graphs may be difficult in real-world applications. Additionally, none of the current methods address the utilization of knowledge about the relationship among input features to handle correlation in cases where causality may not be present. Recognizing this gap, where methods either neglect or explicitly consider causal knowledge among input features, we propose a novel approach: MYerson additive Explanations on gRaphs (MYER). MYER aims to explain machine learning (ML) models by integrating network knowledge among input features without explicitly relying on causal assumptions. Leveraging the Myerson value and the network among input features, MYER calculates feature importance, presenting a framework that builds upon the groundwork laid by SHapley Additive exPlanations (SHAP) in estimating feature importance. The practical application of our proposed approach is demonstrated using feature importance as the indicator of health monitoring
and fault diagnostics for gas turbine engines.

4 - Exploring Continuous-TIME Mean Field Games with Finite State and Action Spaces: Stationary Equilibria and Learning Dynamics

Yunpeng Li, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Antonis Dimakis, Qisheng Huang, Costas Courcoubetis

Mean Field Games (MFGs) provide a foundational framework for modeling interactions in multi-agent reinforcement learning. This paper investigates MFGs with finite state and action spaces, where agents' decision-making processes are governed by continuous-time Markov decision processes (MDPs). Unlike existing approaches that predominantly rely on discrete-time MDPs or continuous-time MDPs with exponentially distributed sojourn times, our study accommodates general sojourn time distributions. We begin by proving the existence of mean field equilibria under mild regularity conditions, utilizing a fixed-point approach tailored to the continuous-time setting. Next, we examine the efficiency loss of decentralized equilibria compared to

centralized optimization by analyzing the price of anarchy (PoA) across diverse sojourn time distributions. Our analysis reveals the maximal and minimal PoA values, along with their corresponding sojourn time distributions, shedding light on how sojourn time distributions influence the efficiency of mean field equilibria. This is the first work to bring together price of anarchy research and distributed robust optimization. Additionally, we introduce and analyze the best-response-based adaptive learning dynamics, naturally suited to the continuous-time MFG framework. This type of learning dynamics, inspired by revision protocols in evolutionary population games, establishes a connection between MFGs and population games. We further identify a general monotonicity condition that enables explicit formulas for the PoA and ensures the convergence of the best-response dynamics.

5 - Investigating Brain Oscillations for Efficient Artificial Intelligence

Yina Wei, Zhejiang Lab, Hangzhou, China, People's Republic of

Recent advances in neuroscience have highlighted the critical role of brain oscillations in coordinating information processing and memory consolidation. These oscillatory mechanisms, particularly during sleep, offer valuable insights for the development of more efficient and robust artificial intelligence (AI). We investigate the computational principles underlying neural oscillations and explore how similar dynamics, such as brain oscillations or spike-timing-dependent plasticity, can be adapted to improve artificial learning systems. By modeling key oscillatory patterns such as slow oscillations, we demonstrate their potential to enhance energy-efficient communication across brain regions, improve robustness, and support continual learning in neural networks. Our findings suggest that incorporating biologically inspired rhythmic dynamics can bridge the gap between brain-like adaptability and machine intelligence, paving the way for next-generation energy-efficient and resilient AI systems.

6 - Ai and Competitive Advantage: Drivers of Heterogeneity OR Homogeneity

Swapnil Garg, Indian Institute of Management , Indore, India, Ananthakrishnan Rajamani

As businesses compete to win the AI adoption race, many lament the missing strategic view. Like electricity and the internet, Artificial Intelligence (AI) is a general purpose technology. AI adoption reduces costs, increases efficiency, and enables development of new products, but doubts remain whether this new technology is going to be a source of sustained competitive advantage for firms. In this research, we theoretically explore the mechanisms that make AI a source of competitive advantage, driving homogeneity or heterogeneity across firms.

We identify six of the key characteristics of AI i.e., Self-learning ability, Training on past data, Fast response, Trust by decision makers, AI-Human ensemble, and Emulating intuition, to yield homogenization pressure and bring heterogenization opportunity on firms. Our key contribution lies in identifying the underlying organizational mechanisms through which AI operates to make firms similar (i.e., cost advantage), and identification of the actions that firms need to take to create unique value (i.e., differentiation advantage). For instance, as the learning ability of AI is restricted to the available explicit data for training, this builds homogenizing pressure. Whereas, training on proprietary data and human prompting can lead to unique value creation.

Our study seeks to build a deeper understanding of how AI is likely to impact organizations strategically, enabling firms to develop a viable AI strategy, moderate expectations from AI, and enable organizations to develop a system based approach of employing AI in the organization, and cut through and see clearly through the haze that surrounds AI developments.

7 - Comparative Analysis of the Influx of Remittances and Foreign Direct Investments (Fdi) to Five English-Speaking West African Countries

Stanley Nwoji, Harrisburg University of Science and Technology, Harrisburg, PA, United States

This study compared the influx of FDI and remittances to five English-speaking West African countries. The goals included investigating if there is a significant difference between the influx of the two foreign funds, formulating econometric relationships between them, and discovering high-performing machine learning and deep learning models in analyzing the influx of remittances and FDI. Another goal was comparing statistical, machine learning, and deep learning models for the influx of remittances and FDI to these countries. The result

showed a statistically significant difference between total remittances and total FDI. Nigeria had a highly significant difference between remittances and FDI, as well as the highest levels of both in the region. Gambia and Sierra Leone also showed a significant difference between remittances and FDIs. However, this was not the case in Ghana and Liberia. We discovered that SARIMA is the best statistical model for FDI prediction, while Random Forest was the best machine learning model. Low-performing machine learning models included Linear Regression, XGBoost, and Support Vector Regression. For the remittances' dataset, Random Forest was the best model. The statistical models (ARIMA and SARIMA) performed lower than all machine learning models. Further tests with deep learning models, like LSTM, GRU, RNN, and Transformer, on the FDI dataset revealed that the RNN was the most accurate for predicting remittances and FDIs. In forecasting FDIs, the Random Forest performed better than all models. However, in the remittances dataset, the RNN performed better than all models.

Monday, July 21, 5:00 PM - 6:30 PM

MD01

Enterprise

Learning and Incentives in Complex Environments

Invited Session

Applied Probability Society

Chair: Yifan Feng, National University of Singapore, Singapore, Singapore

1 - Hybrid Transfer Reinforcement Learning: Provable Sample Efficiency from Shifted-Dynamics Data

Pengcheng You, Peking University, Beijing, China, People's Republic of, Chengrui Qu, Laixi Shi, Kishan Panaganti, Adam Wierman

Online reinforcement learning (RL) typically requires online interaction data to learn a policy for a target task, but collecting such data can be high-stakes. This prompts interest in leveraging historical data to improve sample efficiency. The historical data may come from outdated or related source environments with different dynamics. It remains unclear how to effectively use such data in the target task to provably enhance learning and sample efficiency. To address this, we propose a hybrid transfer RL (HTRL) setting, where an agent learns in a target environment while accessing offline data from a source environment with shifted dynamics. We show that -- without information on the dynamics shift -- general shifted-dynamics data, even with subtle shifts, does not reduce sample complexity in the target environment. However, focusing on HTRL with prior information on the degree of the dynamics shift, we design HySRL, a transfer algorithm that outperforms pure online RL with problem-dependent sample complexity guarantees. Finally, our experimental results demonstrate that HySRL surpasses the state-of-the-art pure online RL baseline.

2 - Experimental Design in the Presence of Supply Chain Spillovers

Yunhong Wang, National University of Singapore, Singapore, Singapore, Hanzhang Qin, Jinglong Zhao

In this paper, we develop a framework to design experiments for a regionalized supply chain, especially in the presence of supply chain spillover effect. Interferences are common in online platforms, and we identify a new interference pattern from the spillover effect of regionalization fulfillment, different from the previous setting in online platforms. We also propose different experimental designs and corresponding standard estimators for estimating the global treatment effect We model customer purchase behavior via time-dependent choice model and use mean-field limit to capture the system dynamics. Finally, we provide sign and bias analysis for the standard inverse probability weighting estimator and propose some alternative estimators that can use more system information than the standard estimator.

3 - Credit Rating Design Under Adverse Selection

Qinzhen Li, National University of Singapore, Singapore, Singapore, Yifan Feng, Jussi Keppo

We study a strategic credit rating model where investors rely on credit rating agencies (CRAs) to assess firms' default probabilities. Firms seeking funding can solicit ratings by paying the CRA, who then optimizes a profit-maximizing rating policy based on firms' quality and solicitation decisions. We find that, in equilibrium, high-

quality firms are more likely to solicit ratings and receive higher ratings compared to low-quality ones. Although solicited ratings are observed to be higher than unsolicited ones, this discrepancy diminishes when accounting for firms' intrinsic qualities. Our model suggests that even if CRA chooses a rating policy that maximizes her own revenue, the reported discrepancy between unsolicited and solicited ratings in the existing literature could be primarily explained by adverse selection on the demand side of the credit rating market.

4 - Revisiting the Principal Agent Problem with Moral Hazard

Hang Jiang, National University of Singapore, Singapore, Singapore

Principal–agent problems with moral hazard form a class of bi-level optimization problems with numerous business applications. They are typically solved by the first-order approach, which replaces the agent's optimization problem with its first-order condition, yielding a relaxed problem for the principal. To justify this approach, the literature usually assumes a fixed outcome support and derives various exogenous sufficient conditions—most notably the Mirrlees–Rogerson Condition—to ensure the concavity of the agent's problem.

In contrast, rather than justifying the first-order approach, we propose a new approach and establish more general conditions under which the infinite-dimensional bi-level optimization problem reduces to a two-dimensional one. We show that these conditions are economically interpretable, relatively easy to verify, and sufficiently flexible to accommodate common instances in which the agent's expected utility is non-unimodal in effort under optimal contracts (in addition to the usual unimodal cases). In particular, under our conditions, the optimal contracts are simple.

Furthermore, our method can be extended to settings with shifting outcome support, where the conventional first-order approach is generally inapplicable. We establish general conditions for simple contract optimality. Notably, our conditions identify scenarios in which the first-order approach fails exactly when the support shifts. To the best of our knowledge, our work is the first to provide exogenous conditions for literature that endogenously assumes the first-order approach is invalid.

MD02

Mercury

Learning and Market Design

Invited Session

Auctions and Market Design

Chair: Jiaqi Lu, the Chinese University of Hong Kong (Shenzhen), Shenzhen, N/A

Co-Chair: Yilun Chen, CUHK Shenzhen, Shenzhen, N/A

1 - Information Deadlock In Matching Markets: a Branching Process Perspective

Jiaqi Lu, the Chinese University of Hong Kong (Shenzhen), Shenzhen, China, People's Republic of, Nicole Immorlica, Yash Kanoria

A key feature of many real-world matching markets is congestion, i.e., market participants struggle to find match partners. We characterize congestion in a model of random matching markets where an agent pair must perform a mutual inspection to verify compatibility prior to matching with each other. We ask when, in large random two-sided markets, will information deadlocks arise in which many agents delay inspections indefinitely awaiting for each other. We characterize the existence and size of information deadlock as a function of simple descriptive market primitives, and find a phase transition from a deadlock-free regime (where a vanishingly small fraction of agents are stuck waiting) to the information deadlock regime as the market primitives change. Our analysis is inspired by the machinery of message passing and density evolution from statistical physics, and the emergence of deadlock corresponds to a certain branching process being supercritical.

2 - Bayesian Dynamic Pricing and the Optimality of Incomplete Learning

Meixi Wu, National University of Singapore, Singapore, Singapore, Yifan Feng, Zhenyu Hu

We consider a setting where a monopolist makes pricing decisions facing a stream of customers over an infinite horizon. The customers are short-lived and homogeneous, all of whom share the same private valuation for the

product. Initiating with a prior on the valuation, the monopolist updates her belief through observation of customers' purchase decisions. By framing this problem as a Bayesian dynamic programming problem, we show that under the optimal pricing policy there exists a region of beliefs, referred to as the “no learning region”, upon which the monopolist ceases learning and adopts a constant price from then on. This cessation hinges on the discount factor being less than one, and, under mild conditions, there is a positive probability of transitioning into this no-learning region. Additionally, we characterize the properties of the pricing dynamics and show that the price may never converge to the customer's valuation. We also compare the Bayes-optimal pricing strategy with the Thompson sampling pricing strategy.

3 - On the Efficiency of Fair and Truthful Trade Mechanisms

Yiding Feng, Hong Kong University of Science and Technology, Hong Kong, Hong Kong, Moshe Babaioff, Noam Manaker Morag

We consider the impact of fairness requirements on the social efficiency of truthful mechanisms for trade, focusing on Bayesian bilateral-trade settings. Unlike the full information case in which all gains-from-trade can be realized and equally split between the two parties, in the private information setting, equitability has devastating welfare implications (even if only required to hold ex-ante). We thus search for an alternative fairness notion and suggest requiring the mechanism to be KS-fair: it must ex-ante equalize the fraction of the ideal utilities of the two traders. We show that there is always a KS-fair (simple) truthful mechanism with expected gains-from-trade that are half the optimum, but always ensuring any better fraction is impossible (even when the seller value is zero). We then restrict our attention to trade settings with a zero-value seller and a buyer with value distribution that is Regular or MHR, proving that much better fractions can be obtained under these conditions.

4 - Recent Developments in Online Resource Allocation: Algorithms and Performance Guarantees

Yilun Chen, CUHK Shenzhen, Shenzhen, China, People's Republic of

The online resource allocation (ORA) problem lies at the heart of many applications in operations research and management science. In this talk, we highlight two recent advances in algorithm design and analysis for ORA, addressing both classical and more general settings.
In the first part, we revisit the Certainty Equivalent (CE) heuristic under the classical i.i.d. model. While previous analyses of CE rely on restrictive non-degeneracy assumptions, we show that CE achieves near-optimal regret—up to polylogarithmic factors in the time horizon—under much milder (continuous) distributional conditions. This challenges the common belief derived in the discrete settings that degeneracy necessitates fundamentally new algorithmic approaches. Our analysis draws on tools from empirical process theory to establish strong concentration results for random linear programs, enabling regret bounds without structural assumptions on the input data.
In the second part, we consider ORA under general correlated uncertainty. We develop a Polynomial Time Approximation Scheme (PTAS) that operates with black-box access to a Monte Carlo simulator. The algorithm scales polynomially with the time horizon and problem parameters, while remaining independent of the size of the underlying state space. Our approach frames ORA as a multistage stochastic integer program and introduces a novel “on-the-fly” stochastic (sub)gradient method for solving large-scale relaxations on scenario trees.
This is joint work with Wenjia Wang, and Sabri Cetin and Dave Goldberg.

MD03

Minto

Data-Driven Optimization in Revenue Management and Transportation

Invited Session

Computing Society

Chair: Long He, George Washington University, Washington, DC, United States

1 - Rational Choice Models: the Temporal Tree Representation

MENGYING XUE, School of Management, Fudan University, Shanghai, China, People's Republic of, Qi Feng, J.George Shanthikumar

Choice models, specifying the consumer's choice probability of an option over a given choice set, are widely studied and applied in many fields. We propose a temporal tree representation of choice that covers all rational choice models. Compared with the existing structural choice models, the tree representation exhibits two major advantages that overcome the major challenges of model identification. First, all rational choice models have a tree representation, and a subclass of tree representation (with set-dependent branching) has a one-to-one correspondence to the rational choice models. This bridges the gap in the existing characterizations of structural models, which are unidentifiable, impose uninterpretable conditions, or do not cover the entire space of rational choice models. Second, the tree representation allows for the flexibility of systematically specifying the choice model structure based on available knowledge and data. In particular, the number of parameters needed to specify a tree representation can be primarily determined by the sufficient knowledge level, which corresponds to a specific layer of the tree branching. The sufficient knowledge level can be empirically determined based on the amount of available data, which in turn determines the number of parameters needed for model estimation. Therefore, the tree representation allows for a natural way of data integration, avoiding misspecification due to restrictive assumptions and overfitting for general models.

2 - Robust Satisficing Vehicle Rental Revenue Management with Substitution and Repositioning

Ning Zhu, University of Science and Technology of China, Hefei, China, People's Republic of, Zhu Meng, Chenyi Fu

This paper investigates the practical inventory management challenges faced by the vehicle rental industry, which arise from the random arrival of orders and lead times before demand occurs, as well as the uncertainty and dynamic nature of cancellations and rental durations after demand is realized. To ensure on-demand rental services while mitigating the impact of randomness on inventory stability, the industry usually employs substitution service and inventory repositioning. We extend the methodology of Bandi et al. (2018) and propose a tractable robust satisficing approach featuring an entropic risk measure. This approach jointly optimizes substitution service and inventory repositioning, addressing the maximization of revenue under varied risk-aversion preferences while safeguarding against multiple uncertainties and misspecification. Numerous numerical experiments constructed by real-world data exhibit that given a certain inventory level, our robust model tends to substitute vehicle types for customers without prior reservations to maintain a stable inventory level to meet the potential demands of reserved customers, achieving 16% higher average weekly revenues compared to the expectation optimization benchmark model. Compared to the sample average approximation model, which disregards the variant nature of probability transitions, the robust model can achieve higher revenue with fewer vehicle inventories. Besides, we also find that further restricting substitution to implement only when the order is placed can reduce in-sample overfitting of the robust model, further improving the out-of-sample revenue and enhancing the interpretability of the substitution service.

3 - Near-Optimal Distributionally Robust Warehouse Assortment Planning

Hao Shen, School of Business, Renmin University of China, Beijing, China, People's Republic of, MENGYING XUE, Zhenyu Gao

E-commerce logistics encounter challenges when distributing items across multiple warehouses in the face of uncertain and fluctuated demand, nonuniform fulfillment costs, and complex operational constraints. This complexity is captured by the distributionally robust (DR) warehouse assortment planning problem, which aims to minimize the total cost of fulfilling customer orders while maintaining sufficient robustness against ambiguous demand distributions. We establish that this problem is NP-hard and cannot be approximated within a factor of 1.5 in polynomial time. These hardness results, combined with the immense complexity of realistic problem instances involving millions of items, make it practically infeasible to derive assortment plans with optimal DR performance.

To address these challenges, we introduce a novel concept termed the "weighted utilization rate under independence (WURI)", which explicitly integrates information on both item popularity and warehouse cost efficiency to evaluate assortment plans. Leveraging the WURI, we show that, under mild assumptions, it is tractable to derive assortment plans with near-optimal DR performance. In particular, for the practically important scenario where assortments are cardinality-constrained, our analysis provides a simple and highly scalable heuristic with an approximation ratio close to 1.5. Our heuristic implies that near-optimal approximation ratio can be achieved by implementing the dual-indexing policy to assortment planning, which assigns more

popular items to warehouses with higher cost efficiency. Numerical experiments and a case study using real-world data from a leading e-commerce company demonstrate that our approach outperforms existing methods in the literature, and significantly reduces fulfillment costs compared to the company's current practices.

4 - Computing Stable Cost Allocations for Heterogeneous Vehicle Routing Game

Zhenzhen Zhang, Tongji University, Shanghai, China, People's Republic of, Lindong Liu

Transportation cooperation is essential for enhancing service efficiency and minimizing operational costs. However, achieving stable cooperation in the heterogeneous vehicle routing problem (HVRP) remains a significant challenge due to the inherent diversity of vehicle types and the varying demands of customers. This study constructs a heterogeneous vehicle routing game (HVRG) grounded in cooperative game theory, aiming to identify a cost allocation mechanism that ensures the stability of the grand coalition. The proposed cost allocation satisfies coalition stability constraints while maximizing the total allocation among all players. To address this problem, we introduce the concept of separable games and propose a row generation algorithm that determines cost allocations for separable games by incorporating multiple coalitions in each iteration, thereby reducing the total number of iterations. Furthermore, we develop several theoretical propositions to improve the computational efficiency of the branch-and-price algorithm used in the row generation subproblem. First, we theoretically establish the equivalence between the capacitated profitable tour problem (CPTP) and the set packing problem (SPP), simplifying the subproblem's solution. Second, by analyzing the relationship between the objective function coefficients and the cost allocation, we predetermine some variables to further reduce the problem size. Numerical experiments on HVRG and VRG instances demonstrate that the proposed algorithm can compute cost allocations for instances with up to 60 players within two hours. The results achieve a gap of less than 1% from the optimal cost allocation, significantly outperforming the 4% gap observed in general row generation algorithms.

MD04

Indiana

Behavioral Research in Supply Chains

Invited Session

Behavioral Operations Management

Chair: Fadong Chen, Zhejiang University, Hangzhou, N/A

1 - Environmental Performance Feedback and Supply Base Stability: a Regulatory Focus Theory Perspective

Yongyi Shou, Zhejiang University, Hangzhou, China, People's Republic of, Ge Wu, Chang Wu, Jinan Shao, Kee-hung Lai

The extant literature has highlighted the pivotal role of environmental performance feedback in affecting firms' strategic behaviors, yet little is known about how it influences firms' behaviors regarding supply base adjustments. To narrow this gap, this study draws upon the behavioral theory of the firm and the regulatory focus theory to investigate how focal firms adjust their supply bases in response to environmental performance feedback and how different regulatory foci moderate the relationship. Using a panel dataset of 2,429 listed Chinese manufacturing firms from 2009 to 2023, our analysis reveals that firms experiencing negative environmental performance feedback tend to increase supply base stability, reflecting a more conservative response . Conversely, firms with environmental performance above the industry average are more likely to actively adjust their supply bases, resulting in decreased supply base stability. Furthermore, under negative environmental performance feedback, a prevention focus triggered by external environmental regulatory pressure mitigates the inertia and forces firms to stabilize their supply bases . In contrast, under positive environmental performance feedback, a promotion focus driven by the CEO's internal expectations, as reflected in firm growth, encourages firms to further actively adjust their supply bases. This study extends the literature by offering novel insights into the differential effects of environmental performance feedback on firms' supply base stability and the crucial moderating roles of firm growth and environmental regulatory pressure.

2 - An Adaptive Supply-Demand Matching Method Considering Service Exceptions in Shared Manufacturing

Chunxia Yu, China University of Petroleum, Beijing, China, People's Republic of

With the integrated development of the industrial internet and the sharing economy, capacity sharing platforms have emerged as key infrastructure for the optimized allocation of manufacturing resources. However, unexpected service exceptions frequently occur during platform operation, leading to chain reactions such as unstable supply-demand matching and delayed task delivery, resulting in economic losses. This study aims to develop an adaptive adjustment decision-making model for supply-demand matching under service exceptions, which can achieve reliable delivery of tasks and robustness of the matching system by flexibly invoking optimal exception handling strategies and collaboratively optimizing multi-dimensional objectives of user satisfaction - economic efficiency - adjustment deviation. First, an exception handling strategies determination method is proposed to accurately identify exception scenarios and match appropriate strategies. Second, the impact of service exceptions on two-sided matching preferences is captured, and a supply–demand adjustment satisfaction evaluation method is developed. Third, an adaptive adjustment model is constructed with the objectives of maximizing the adjustment satisfaction of two-sides, and minimizing the adjustment cost and deviation, the model is solved using an algorithm based on U-NSGA-III to generate optimal adjustment scheme. Finally, the feasibility of the proposed decision model is verified through multi-scenario case studies and comparative experiments. Results show that the proposed method can accurately evaluate exception impacts, avoid unnecessary re-matching, reduce adjustment costs and improve user satisfaction. Moreover, the algorithm has the ability of adaptive multi-strategy invocation, and effectively enhances the quality and efficiency of the solution.

3 - Incorporating Risk Interdependency and Behavioral Heterogeneity: a Genai-Sd Approach for Compound Risk Governance In Critical Infrastructures

Weilan Suo, School of Economics and Management, Beijing University of Chemical Technology, Beijing, China, People's Republic of, Lin Wang

Critical infrastructures (CIs) confront escalating threats from compound risks necessitating adaptive governance strategies. Current governance frameworks exhibit critical limitations in assessing the nonlinear interdependencies of synchronously emerging risks and the behavioral heterogeneity of multi-agent stakeholders, particularly in scenarios dominated by non-traditional risks (e.g., extreme climate events) that exceed the boundaries of historical data and expert knowledge. To bridge these gaps, this study presents a novel framework integrating Generative Artificial Intelligence (GenAI) with System Dynamics (SD). First, we harness multi-source heterogeneous data (including global disaster databases, policy documents, and news reports) to train GenAI models in identifying governance patterns for compound risk events, risk interdependency networks, and behavioral heterogeneity rules among CI stakeholders like operators, regulators, and users. Subsequently, SD models validate these AI-generated insights through causal loop diagrams and stock-flow simulations, enabling the identification of core risk governance pathways and quantitative evaluation of risk mitigation efficacy. Then, empirical validation on power-grid systems under concurrent physical-climate disruptions demonstrates the superior performance of the proposed approach. Noteworthy contributions encompass: (1) dynamic mapping of emerging risk scenarios beyond historical precedents, (2) explicit quantification of the behavioral impact on risk governance outcomes, and (3) policy adaptation through GenAI-SD feedback loops. The symbiosis between GenAI-SD heralds a computational governance model for intricate systems confronting unprecedented uncertainties, furnishing actionable governance strategies for bolstering critical infrastructure resilience amidst new risk landscapes. This paradigm particularly addresses emergent climate-technology interdependencies, where traditional risk assessment frameworks fall short.

4 - Diy vs. Outsource: How Does Manufacturer's Optimal Promotion Strategy Form Sustainable Live Stream Supply Chain?

Weijun Xu, School of Management, Xi'an Jiaotong University, Xi'an, China, People's Republic of

This study analytically examines a manufacturer's promotion strategy choices, focusing on building a sustainable live stream dual-channel supply chain. Developing a game-theoretic model, we define and analyze three sales strategies: traditional, DIY live stream, and outsource live stream. Promotion effort and streamer popularity are incorporated into DIY and outsource strategies, respectively, for sensitivity and comparison analysis. Results show that, within acceptable cost ranges, promotion strategies can create a "win-win" live stream dual-channel supply chain, positively impacting the manufacturer and the reseller on pricing and sales. Streamer popularity's spillover effect decreases with short-term live stream sales but increases with long-term market sales. Strong

popularity spillover drives a "hot-selling" phenomenon, benefiting the reselling channel. We also introduce cost schemes involving channel construction and incremental pit fees to guide optimal promotion strategy selection. Additionally, multiple Pareto optimality regions are identified under varying streamer popularity levels, forming sustainable live stream dual-channel supply chains. We extend our analysis to examine consumer surplus, social welfare and the impact of consumer's acceptance ratio of outsourced live stream to demonstrate the robustness of our model. This study contributes to the literature and provides practical insights for supply chain participants and managers. It assists the manufacturer in choosing optimal promotion strategies from a cost structure perspective to build sustainable live stream supply chains.

5 - Understanding Groupthink in Project Investment Decisions: a Quantitative Approach

Seyed Ashkan Zarghami, Australian National University, Canberra, Australia

Group decision-making plays a vital role in shaping organizations and societies, but the tendency to prioritize consensus over critical evaluation can lead to flawed decisions. This phenomenon, known as groupthink, was first introduced by Irving Janis (1972) and has been widely studied over the past five decades. Despite extensive research on its antecedent conditions and effects, significant gaps remain in quantifying its impact on decision-making outcomes.

This study addresses two key research questions: (1) What is the likelihood of groupthink occurring in project investment decisions under specific antecedent conditions? (2) How does groupthink distort rational decision-making and affect the selection of optimal outcomes?

To answer these, this research develops an integrative measure to estimate the probability of groupthink behavior. It further applies evidence theory to model how groupthink influences decision-making by adjusting utility values and belief degrees for different alternatives. The proposed mathematical model predicts the probability of suboptimal choices in the presence of groupthink, offering a novel approach to quantifying its effects. These insights are crucial for improving investment decision processes, minimizing economic losses, and ensuring rational, data-driven choices.

MD05

Orchard A

Design and Governance of of Web3.0 and Metaverse

Invited Session

Digital Twin and Metaverse

Chair: xiaoni lu, Xi'an Jiaotong University, Xi'an, N/A

1 - Designing a Blockchain-Driven Ai Agent Identity Management Framework

zhiyue yan, Xi'an Jiaotong University, xi'an, China, People's Republic of, Yingying Li, xi zhao

As Artificial Intelligence (AI) agents evolve, they transcend traditional automation paradigms, shifting from simple computational tools to autonomous decision-makers that extend beyond the digital realm into physical society. This evolution underscores the urgent need for AI agent identity management mechanisms to address pressing issues such as identity theft, accountability risks, scalability constraints, and even internal identity confusion among agents. Drawing on the Design Science Research (DSR) methodology and ontological principles, we identify three foundational objectives for AI agent identity management: uniqueness, consistency, and continuity. To achieve these goals, we propose a blockchain-based framework for AI agent identity management, which captures both static identity attributes and dynamic evolution of AI agent identities, enabling immutable record-keeping and transparent verification across stakeholders and temporal dimensions. Based on the perception-cognition-action (PCA) cycle paradigm of intelligent agents, we further map the timeliness evaluation of the dynamic attribute iterations of AI agent identity to the environment coupling of the agent AI. Accordingly, we propose four dimensional indicators: environmental specificity, environmental perception depth, decision adjustability, and response dynamics. To this end, we validate the framework's relevance and effectiveness through surveys and a prototype demonstration. This research will significantly

contribute to advancing the theoretical understanding of AI agent identity, while offering practical insights for building resilient identity management systems in real-world AI applications. Our findings will hold crucial implications for fostering a trusted AI ecosystem and supporting the seamless integration of autonomous agents across industrial and societal domains.

2 - On the Value of Information Sharing in the Presence of Information Errors:&Ndash; a Perspective of Blockchain Adoption

ziming zhang, xi'an jiaotong university, xi'an, China, People's Republic of

Although blockchain technology shows great potential in enabling supply chain visibility and ensuring information security, the information security dilemma still exists and even becomes more hidden and harmful with the enhancement of blockchain technology features. Based on the reality that there is an information gap between the physical ledger and the electronic ledger before supply chain information is collected and uploaded to the blockchain, this paper explores the supply chain visibility decision-making under the influence of “information noise” with the help of the three-party evolutionary game model, and designs targeted mitigation strategies. The study shows that: Under the driving force of “active benefit acquisition” and “passive risk defense”, “information noise” is able to dissolve the value creation logic of supply chain quality visibility. Under the “threshold effect”, the positive pull of required disclosure on corporate disclosure behavior and end-user utility is replaced by the “positive” value dissipation and “negative” value pull

3 - The Impact of Economic Costs on Web3 Governance Participation: An Empirical Study of Dao

Zhecheng Wang, Xi'an Jiaotong University, Xi'an, China, People's Republic of, xi zhao, Gengzhong Feng

Within the blockchain ecosystem of Decentralized Autonomous Organizations (DAOs), effective governance remains a critical yet underexplored area. This study examines the impact of refunding gas fees as an economic incentive to lower participation costs and thereby enhance governance voting rates. Grounded in the Model of Frame Selection (MFS), our research explores how individuals’ decision-making processes are informed by both cognitive and social frames under varying contexts. Previous studies have demonstrated that rational choice theory fails to adequately explain collective behaviors in contexts such as political elections, corporate governance, and online community participation. In the unique environment of DAOs—characterized by anonymity, decentralization, immutability, and consensus-based interactions—the core elements of members’ governance decision frameworks remain ambiguous.

Our empirical findings indicate that reducing the economic burden of participation significantly increases voter turnout within DAOs. Specifically, elements within the cognitive frame, such as voting rights, individual experience, and original membership, have a positive moderating effect on the relationship between reducing the cost of participation and enhancing participation. Conversely, the presence of arbitrage opportunities within the social environment serves to negatively moderate the effectiveness of the gas fee refund incentive. This research not only fills a critical gap in understanding the economic mechanisms underpinning DAO governance but also provides essential insights for the design of more effective incentive policies within blockchain-based systems.

4 - The Impact of Bulk Purchasing on The Nft Market: An Empirical Evidence from Opensea

xiaoni lu, Xi'an Jiaotong University, Xi'an, China, People's Republic of, Xiaoyun Rong, Xi Zhao

Non-fungible token (NFT) marketplaces, the primary platforms for NFT transactions, have increasingly introduced the "Sweeping" —a bulk purchase feature designed to enhance the buying experience and improve market liquidity. While this feature offers several apparent advantages, such as improved operational efficiency for acquiring multiple NFTs and reduced gas fees, its actual impact on the sales performance of NFT collections remains unclear. This study investigates how the adoption of the bulk purchase feature differentially affects sales performance among NFT collections that adopt the feature and those that do not. Leveraging Ethereum on-chain NFT transaction data and NFT collection characteristic information from OpenSea, one of the leading NFT marketplaces, we employ an empirical analysis combining propensity score matching (PSM) and difference-in-differences (DID) to identify causal effects. We find that the bulk purchase feature significantly improves sales performance of NFT collections across four key metrics: (1) the number of NFT sold, (2) total transaction volume, (3) the number of unique buyers, and (4) average NFT price. These results align with value perception

theory. Notably, the feature also generates positive spillover effects, boosts sales performance even in traditional single-item purchase mode. Furthermore, we identify structural market shifts induced by the feature: high-value NFT collections exhibit disproportionate performance gains, and the contribution of high-value, active buyers to total sales increases significantly. Our results provide practical insights for NFT marketplace operators regarding the implementation of bulk trading mechanisms.

5 - Understanding Consumption Values and Purchase Intention in Retailing Metaverse

Jun Yan, Huazhong University of Science and Technology, Wuhan, China, People's Republic of, Tanveer Abbas, Khuram Shahzad, Zhen LI, Fangzhou LI

Retailing metaverse, offering immersive and interactive shopping experience to users, could dramatically change the retail landscape. In retailing metaverse, consumers can shop in an ideal blend of the real and virtual worlds by navigating metaverse platforms with their avatars, interacting with other users (either actual customers or virtual assistants), and improving their overall shopping experience. However, there remains a scant understanding of what drives the users of metaverse to shop inside the metaverse and maybe continue to the offline shopping. This study addresses calls for research on the role of epistemic and conditional values in metaverse retailing. For this purpose, the authors applied the theory of consumption values to examine the factors that influence the consumer’s attitudes and purchase intentions towards metaverse retailing. To validate the proposed framework, the data was obtained from 504 Chinese consumers through online survey.
The findings indicate that (1) five consumption values contribute to user attitude toward retailing metaverse differently. Among them, functional value (“realism, availability, customization”) is the main driver of attitude towards metaverse retailing, followed by epistemic value (“variety seeking”), conditional value (“preference value”), and arousal (part of emotional value) while social value (“social interaction, self-image expression”) and pleasure (proposed as part of emotional value) have insignificant impact. (2) Perceived complexity is found negative moderating effect on the relationship between attitude and purchase intention. The findings offer several critical implications for marketers to increase the behavioral outcomes towards metaverse retailing.

MD06

Orchard B

New Techniques in Financial Mathematics/Engineering

Invited Session

Finance

Chair: Guanxing Fu

Co-Chair: ZHAOLI JIANG, The Hong Kong Polytechnic University, Shenzhen, N/A

1 - Data-Driven Merton&Rsquo;S Strategies via Policy Randomization

Min Dai, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, Yuchao Dong, Yanwei Jia, Xunyu Zhou

We study Merton’s expected utility maximization problem in an incomplete market, characterized by a factor process in addition to the stock price process, where all the model primitives are unknown. The agent under consideration is a price taker who has access only to the stock and factor value processes and the instantaneous volatility. We propose an auxiliary problem in which the agent can invoke policy randomization according to a specific class of Gaussian distributions, and prove that the mean of its optimal Gaussian policy solves the original Merton problem. With randomized policies, we are in the realm of continuous-time reinforcement learning (RL) recently developed in Wang et al. (2020) and Jia and Zhou (2022a,b, 2023), enabling us to solve the auxiliary problem in a data-driven way without having to estimate the model primitives. Specifically, we establish a policy improvement theorem based on which we design both online and offline actor–critic RL algorithms for learning Merton’s strategies. A key insight from this study is that RL in general and policy randomization in particular are useful beyond the purpose for exploration – they can be employed as a technical tool to solve a problem that cannot be otherwise solved by mere deterministic policies. At last, we carry out both simulation and empirical studies in a stochastic volatility environment to demonstrate the decisive outperformance of the devised RL algorithms in comparison to the conventional model-based, plug-in method.

2 - Long TIME Behavior of Optimal Liquidation Problems with Semimartingale Strategies and Regime Switching

Xiaonyu Xia, Wenzhou University, Wenzhou, China, People's Republic of, Xinman Cheng, Guanxing Fu

This paper studies an optimal liquidation problem on infinite time horizon with semimartingale strategies, regime switching and random coefficients. We obtain that the value function and the optimal strategy can be characterized in terms of three multi-dimensional backward stochastic differential equations (BSDEs) which involves a kind of BSDE with stochastic Lipschitz coefficients and infinite time horizon. We establish the existence of the solutions to these systems of BSDEs mainly by means of BMO martingales and comparison theorem for multi-dimensional BSDEs.

3 - Quantitative Propagation of Chaos for Mean Field Randomized Stopping in Continuous TIME

Yuchao Dong, Tongji Univeristy, Shanghai, China, People's Republic of

The classical decision problem of optimal stopping has found an amazing range of applications from finance, statistics, marketing, and phase transitions to engineering. When extend the standard framework to multi-agent setting, one obtains mean field stopping problem when the number of agents goes to infinity. In this paper, we consider a mean field stopping system with a randomized stopping strategy. The wellposedness of such system is given together with some results on quantitative Propagation of chaos.

4 - Dynamic Mean-Variance Efficient Fractional Kelly Portfolios in a Stochastic Volatility Model

ZHAOLI JIANG, Chinese University of Hong Kong, Shatin, N. T., Hong, Hong Kong

In this paper, we improve the mean-variance efficiency of the fractional Kelly strategy, a popular investment strategy in the market, by minimizing the variance of the return of a portfolio with as high expected return as the fractional Kelly strategy. In view of time inconsistency arising from the mean-variance criterion, we consider so-called equilibrium portfolio strategies that can be consistently implemented by the investor. We derive the equilibrium portfolio strategy in closed form and show that it always leads to a smaller variance of return than the fractional Kelly strategy. By calibrating the model parameters to market data, we show that the reduction of variance achieved by the equilibrium portfolio can be economically significant. We also show in an out-of-sample test that the equilibrium portfolio strategy outperforms the fractional Kelly strategy under various performance measures.

MD07

Orchard C

Market Equilibrium and Energy Systems Integration

Contributed Session

Energy, Natural Resources and the Environment

Chair: Paola Labrecciosa, ESSCA School of Management, Paris, N/A, France

1 - Equilibria in Interdependent Natural-Gas and Electric Power Markets: An Analytical Approach

Seyedamirabbas Mousavian, University of New England, Biddeford, ME, United States, Beheshteh Raouf

Natural-gas and electric power systems and their corresponding markets have evolved over time independently. However, both systems are increasingly interdependent since combined cycle gas turbines that use natural gas to produce electricity increasingly couple them together. Therefore, suitable analysis techniques are most needed to comprehend the consequences on market outcomes of an increasing level of integration of both systems. There is a vast literature on integrated natural-gas and electric power markets assuming that the two markets are operated centrally by a single operator. This assumption is often untrue in the real world, which necessitates developing models for these interdependent yet independent markets. In this vein, this paper addresses the gap in the literature and provides analytical Nash-Cournot equilibrium models to represent the joint operation of natural-gas and electric power markets with the assumption that the market participants in each market make their own decisions independently seeking the maximum profits, as often is the case in the real world. We develop an

analytical equilibrium model and apply the Karush-Kuhn-Tucker (KKT) approach to obtain Nash-Cournot equilibria for the interdependent natural-gas and electric power markets. We use a double-duopoly case to study the interaction of both markets and to derive insightful analytical results. Moreover, we derive closed-form analytical expressions for spot-market equilibria in both natural-gas and electric power markets, which are relevant and of practical significance for decision makers. We complement the double-duopoly study with a detailed sensitivity analysis.

2 - Finding Sparse Reactancy Matrix for Optimal Power Systems with Network Flow Model

Aurelio Oliveira, Universidade Estadual de Campinas, Campinas, Brazil, Fernanda Teixeira

In formulating the optimal power flow model with network flows, it is convenient to find the sparsest possible full row rank reactance sub-matrix, reducing the computational effort of numerical methods. Such a sub-matrix comes from the second Kirchhoff loop's law, and the more sparse version corresponds to a minimum cycle basis in graph theory. The existing algorithms for finding this basis are very expensive in terms of time and memory, and in practice heuristics are used to find a sufficiently sparse basis. However, there are applications where the network keeps a fixed topology for long periods of time, and it may be convenient to find the sparsest possible basis, even at a high computational cost. Numerical experiments of the known algorithms and proposed heuristics are performed on real Brazilian power systems and IEEE test systems. The processing time and the sparsity of the sub-matrix found are compared, as well as the total time to solve the optimal power flow. The results are studied concerning application contexts ranging from near real-time problems to system expansion planning problems, indicating the best options in each case.

3 - On the Effects of An Increase in the Number of Abaters in Pollution Abatement Games

Luca Colombo, Rennes School of Business, Rennes, France, Paola Labrecciosa

We study the effects of an increase in the number of abaters in pollution abatement games, first in a static, then in a dynamic (continuous-time) game. In both games, we assume that m countries/agents agree on taking action to reduce the stock of pollution, which is a public bad, whereas n-m countries free ride on the abatement levels of the abaters. Moreover, we assume that abaters can either coordinate on their contributions or not. In the static game, both in the coordination and the non-coordination scenario, an increase in m leads to a decrease in the stock of pollution and to an increase in social welfare. In the dynamic game, instead, both in the coordination and the non-coordination scenario, an increase in m may result in a higher steady-state stock of pollution and a lower social welfare, depending on the "business-as-usual" level of output.

4 - Efficiency-Inducing Lump Sum Taxation in Nonrenewable Resource Industries

Paola Labrecciosa, ESSCA School of Management, Paris, France, Hassan Benchekroun, Luca Colombo

We study efficiency-inducing lump sum taxation within a differential game framework where oligopolistic firms share access to a common-pool nonrenewable resource and use the harvested resource in the production of homogeneous goods to be sold in the marketplace. A lump sum tax refers to a flow payment that firms must make while the resource is exploited. Such a tax is independent of the amount of resource extracted. We show that the socially optimal extraction path can be decentralized either as an Open-Loop or as Closed-Loop (Feedback) Equilibrium. We then provide a characterization of the optimal tax for a class of utility functions that encompasses the sets of utility functions generating the linear and the constant-elasticity demand functions, and show that the 'intuitive' lessons, which can be drawn from the existing literature on per unit taxation and oligopoly exploitation of common-pool resources under laissez-faire, may no longer be valid when the tax depends only on the resource stock.

MD08

Bras Basah A

Consumer Behavior and Decision Support Systems

Contributed Session

Information Systems

Chair: Xingyue Zhao, Xian Jiaotong University, Xian, N/A

1 - Recent OR Helpful? the Recency Gravity Effect of Review Set Sorting on Consumer Purchasing

Yan Zhang, Xi'an Jiaotong University, School of Management, Xi'an, China, People's Republic of, Xinyu Sun

Online platforms commonly sort reviews based on specific criteria to alleviate the information overload, thereby presenting consumers with different sets of reviews to peruse. The set of helpful reviews (sorted by review helpfulness) and recent reviews (sorted by posting time) are two commonly provided options. This study investigates which review set exerts a greater influence on consumers' purchasing decisions when there is an opinion divergence. We propose a theoretical framework to explain how consumers sway between two review sets on the basis of the Information Adoption Model. We conduct an archival analysis and four laboratory experiments for empirical validation. Results show that, 1) consumers' purchasing decisions are surprisingly influenced more by the set of recent reviews rather than helpful reviews, which we term the "recency gravity". 2) The effect stems from consumers' perception that newer reviews are more credible and that the recency-based sorting method is more transparent in comparison to the helpfulness-based one. 3) The "recency gravity" is strengthened when the product price is lower or when the product is hedonic. Our findings challenge the prevalent view that helpful reviews are inherently more persuasive in driving consumer purchases, revealing that in addition to information quality, consumers are sensitive to how the information is framed. This work contributes to a deeper understanding of consumers' purchasing behaviors, the consequence of review sorting methods, and effective approaches for offering online review sets.

2 - Negative Recommendations Using Signed Network

Mukul Gupta, Indian Institute of Management Indore, Indore, India

Online shopping has become a key part of global retail. To help users navigate overwhelming choices, e-commerce platforms use recommender systems that typically focus on suggesting items users may like. However, these systems often ignore negative feedback, such as dislikes or low ratings, while generating recommendations, which can be just as valuable in understanding user preferences. This study addresses the task of generating negative recommendation lists: personalized lists of items a user is likely to dislike. Incorporating such lists can prevent poor recommendations, reduce user discomfort, and build long-term trust. While platforms like YouTube and IMDb allow for negative feedback, most systems do not use this information effectively. Moreover, negative reviews have been shown to influence purchasing decisions more strongly than positive ones. We propose a method that models both positive and negative user preferences using low-dimensional representations within a signed network. This approach identifies items users are likely to dislike, improving the accuracy of positive recommendations as well. We evaluate our method on real-world datasets and compare it with SNBI-2, a recent negative recommendation model. Results show our method outperforms the SNBI-2 approach in identifying disliked items and enhancing recommendation quality.

3 - Platform Pricing and Exclusion Strategies: Balancing Quality, Profitability, and Social Welfare

Yangjie Gao, HKUST(GZ), Guangzhou, China, People's Republic of, Xu Zhang, Xiaotong Sun

This paper investigates how a profit-driven two-sided platform sets its price rate and subsidies for consumers and tailors exclusion strategies for high- and low-quality merchants operating under a fixed commission rate. Consumers are sensitive to price and quality but have incomplete information about individual merchant types. We first analyze the indirect exclusion strategy that requires low-quality merchants to subsidize their consumers while the platform directly provides subsidies on behalf of high-quality merchants, demonstrating that both prices and subsidies increases in the weight placed on quality by consumers. We then examine the partial exclusion strategy where the platform directly excludes a subset of low-quality merchants without subsidizing consumers, and find that this approach lowers prices and potentially the platform's profitability. We further study the complete exclusion strategy, in which low-quality merchants are fully excluded and the platform covers all consumer subsidies. We reveal that, if the quality gap between high- and low- merchants is sufficiently large, complete exclusion can be more profitable than indirect exclusion. Social welfare analyses highlight that indirect exclusion dominates complete exclusion and often outperforms partial exclusion when commission rates are moderate and the proportion of low-quality merchants is high. Taken together, our findings underscore the complexities that platforms face in regulating quality, offering insights into optimal platform strategies for balancing market efficiency, profitability, and consumer welfare.

4 - Algorithmic Superspecies: Generative AI and the Transformation of Consumers and Brands into Simultaneously-Evolving Algorithms

Sunil Erevelles, University of North Carolina, Charlotte, Charlotte, NC, United States, Anjali Sinha, Sahar Maleki, Ganesh Yogeeswaran

Big Data provided us with a richer analysis of consumers beyond anything possible before. Social media algorithms, in turn, leveraged Big Data to target consumers with ever-greater precision. While the Web was designed to connect consumers and businesses worldwide, social media algorithms did the opposite, and divided consumers into micro-segments, each receiving its own unique content. While most social media algorithms are AI-generated, until recently most content was human-generated. Generative AI marks the beginning of a new era. It both exploits Big Data and penetrates consumer minds, capturing their intimate cognitions and emotions. The focus in marketing is thus shifting from attention to *intimacy*. Not only will generative AI transform consumers into algorithms and super-precisely target them, but brand content itself will be simultaneously-generated by AI, precisely tailored and constantly evolving for each individual.

Methodologically, the authors draw on Bayesian, transfer learning and neuro-symbolic theories, using inductive-reasoning, cognitive mapping techniques to address two key questions: (i) Do simultaneously-evolving generative AI consumer-brand algorithms improve the strength of consumer-brand relationships when compared with traditional social media algorithms? and (ii) Do simultaneously-evolving generative AI consumer-brand algorithms possess high emotional intelligence even though they have no emotions? Our answers may suggest that the days of consumer-targeted advertising may be ending and that future brand content may be targeted at agentic AI bots. After all, why would consumers need advertising if their generative AI bots can cognitively and affectively tell them what to buy? Finally, suggestions for the design of self-evolving AI algorithms are presented.

5 - Is Contract Governance Everything? a Study on the Impact of Blockchain Governance on Supply Chain Risk Perception

Xingyue Zhao, Xian Jiaotong University, Xian, China, People's Republic of

The structural imbalance inherent in global supply chains exacerbates vulnerabilities for weaker parties, disproportionately amplifying their perception of risk and consequently destabilizing collaborative ecosystems and sustainability. Although contractual governance has traditionally been employed to mitigate such risks, emerging critiques increasingly question its efficacy, especially under conditions of power asymmetry. Blockchain technology has emerged as an alternative governance mechanism, often praised for potential to recalibrate governance structures through technological efficiency. However, existing research predominantly emphasizes technological capabilities, largely overlooking the socio-structural dimensions inherent in dyadic inter-organizational relationships. To address this research gap, this study draws upon Social Information Processing theory to explore how governance mechanisms in structurally imbalanced supply chains act as critical environmental cues shaping the risk perceptions of weaker parties. We posit that governance mechanisms serve as influential signals. Contractual governance institutionalizes existing power imbalances, paradoxically intensifying the perceived risk of weaker parties. Conversely, blockchain governance fosters equitable power distribution and enhances weaker parties' participation through dual cognitive pathways, thereby significantly alleviating their risk perceptions. Furthermore, blockchain governance also recalibrates the institutionalized power asymmetries previously reinforced by contractual governance, mitigating its adverse impacts. Empirically, this study constructs a PLS-SEM model using SmartPLS 4 to analyze data from 209 survey responses. The results robustly support our hypotheses. Theoretically, this research advances SIP theory by extending its applicability into inter-organizational contexts, critically challenging overly optimistic assumptions prevalent in traditional contractual governance literature. Practically, this study provides actionable insights for supply chain managers by integrating technological adoption with cognitive fairness considerations.

MD09

Bras Basah B

AI for Smart Ports

Invited Session

Maritime AI

Chair: Haobin Li, National University of Singapore, Singapore, N/A, Singapore

1 - A Self-Adaptive Monte Carlo Tree Search Approach for Generalized Quay Crane Scheduling Problem

Lei Hai, Northwestern Polytechnical University, Xi'an, China, People's Republic of, Chenhao Zhou

The quay crane scheduling problem (QCSP) optimizes container handling operations to minimize vessel turnaround time. While existing studies focus on container batch size and safety constraints, they neglect varying container group sizes (20-foot vs. 40-foot occupying 1-2 bays). This work introduces a generalized QCSP model addressing these gaps through a Markov Decision Process (MDP) framework, where tasks are dynamically assigned to cranes to reduce idle time. To solve large-scale instances (10-11 QCs, 52-56 bays, 200-280 tasks), we propose a self-adaptive Monte Carlo tree search (SA-MCTS) method with multi-stage lower bounds. This approach prunes inefficient search paths and prioritizes urgent nodes, enhancing computational efficiency. SA-MCTS outperforms three baseline strategies (best-left, best-right, minimal-distance) and a Bi-criteria Evolutionary Algorithm (BIEA) across small- to large-scale scenarios. In large instances, SA-MCTS achieves average improvements of 20%, 14%, 11%, and 72% over best-left, best-right, minimal-distance, and BIEA respectively. When applied to single-size container groups, SA-MCTS matches the state-of-the-art EXM algorithm (1% gap from lower bounds) but with greater scalability. The algorithm’s adaptive tree search structure and multi-stage evaluation mechanism enable effective handling of real-world complexities, including bidirectional crane movement and safety constraints. Results demonstrate SA-MCTS’s superiority in balancing solution quality and computational resource use for modern ultra-large vessel operations requiring high-volume container handling.

2 - A Digital twin based simulation framework Technology for Operations Optimization in Large-Scale Container Terminals

Yi Ding, Shanghai Maritime University, Shanghai, China, People's Republic of, Youxuan Pan, Tiantian Zhu, Weihao Wang

Container terminals across the globe are confronted with mounting challenges in the form of intricate operations and mounting demands for intelligent decision-making in the domains of scheduling, resource allocation and logistics management. Conventional methodologies are unable to adequately balance real-time responsiveness with optimal decision quality in dynamic port environments. To address these challenges, a DT-ABS-OPT framework is presented, integrating agent-based simulation (ABS), digital twins (DTs), and optimisation techniques specifically for maritime terminal operations. The proposed approach employs DTs to furnish precise data support and system representation, thereby facilitating real-time operational insights devoid of decision-making capabilities. The ABS then analyses the data and drives the decision-making process through optimisation model selection and parameter tuning. The optimisation models, in turn, provide high-quality solutions under agent guidance, thus overcoming the traditional real-time performance limitations through continuous agent operation.As an illustration of the application of this framework, this research developed a simulation and optimisation system for the terminal's vessel operations parts, which is driven by the framework. The system is integrated for three parts: prediction, decision making and action implementation, including: predicting the operational efficiency of vessel and quay cranes (QC), making decisions on QC scheduling plans and container work sequences, and interacting with the container terminal digital twin system to simulate operation implementation. The system is utilised for the purpose of verifying and predicting the rationality of resource allocation in the port for forthcoming daily, in addition to evaluating the effectiveness of prediction and decision algorithms.

3 - A Data-driven Global Port Performance Monitoring System

Dong Yang, Hong Kong Polytechnic University , Hong Kong, Hong Kong

4 - From Data to Foresight: Leveraging AI and Digital Twins to Support Smart Port and Terminal Operations

Carlo Mojica, Moffatt & Nichol, Singapore, Singapore, Alan Zhang

This presentation explores how artificial intelligence (AI) and digital twin technologies are being leveraged to transform port and terminal operations, moving "From Data to Foresight". Drawing on solutions like FlexTerm's simulation software for analyzing the past, monitoring the present, and predicting the future by modeling port

master plans, terminal layouts, operations, and equipment, and NextPort's Operational Digital Twin which provides real-time visualization and utilizes machine learning for proactive decision-making and intelligent tools, this presentation will delve into the ways these technologies enable smarter, more efficient maritime logistics. It will highlight how the integration of data with AI and virtual representations may facilitate enhanced strategic planning, optimized resource allocation, and improved operational performance in the complex port and terminal environment.

5 - Smart Ports for a Green Future: Integrating Spatiotemporal AI

Dohee Kim, Pusan National University, Busan, Korea, Democratic People's Republic of, Minseop Kim, Jaeeun Kwon, Kikun Park, Hyerim Bae

As global supply chains advance, smart ports are becoming vital hubs for fostering sustainable and intelligent logistics ecosystems. Strengthened international environmental regulations are driving demand for eco-friendly developments in maritime and port industries, leading to innovations in both hardware and software. Beyond the adoption of green vessels and alternative fuels, attention is increasing toward technologies and systems that enhance fuel efficiency and reduce greenhouse gas emissions. Accordingly, the necessity for smart port solutions that minimize environmental impact while optimizing operational performance is gaining prominence. Given the multi-stakeholder environment of maritime and port logistics, effective decision-making requires the integration of spatiotemporal data. This chapter examines the application and future potential of spatiotemporal AI technologies in smart ports, focusing on case studies based on real-world data.

Keywords: Smart Port, Spatiotemporal AI, Big Data, Green Port, Environmental Sustainability, Maritime Innovation

MD10

Hullet Room

Building Resilient Power Systems: Adaptive Strategies for Climate-Induced Deep Uncertainty

Invited Session

Power Grid

Chair: Chenxi Hu, The University of Hong Kong, Hong kong, N/A

Co-Chair: Yujia Li, lawrence berkeley national laboratory, Berkeley, CA, United States

1 - Multi-Network Constrained Integrated Community Energy Systems: Techno&Ndash;Economic Modelling and Operational Optimization

Ziqing ZHU, Hong Kong Polytechnic University, Kowloon, Hong Kong

The integrated community energy system (ICES) has emerged as a promising solution for enhancing the efficiency of distribution system by coordinating multiple energy sources. However, the concept and modeling of ICES still remain unclear, and operational optimization of ICES is hindered by the physical constraints of heterogeneous integrated energy networks. This paper, therefore, provides an overview of the state-of-the-art concepts for techno–economic modeling of ICES by establishing a Multi-Network Constrained ICES (MNC-ICES) model. The model underscores diverse energy devices at community and consumer levels and multiple networks for power, gas, and heat in a privacy-protection manner, providing a basis for practical network-constrained community operation tools. The corresponding operational optimization in the proposed model is formulated into a constrained Markov decision process (C-MDP) and solved by a Safe Reinforcement Learning (RL) approach. A novel Safe RL algorithm, Primal-Dual Twin Delayed Deep Deterministic Policy Gradient (PD-TD3), is developed to solve the C-MDP. By optimizing operations and maintaining network safety simultaneously, the proposed PD-TD3 method provides a solid backup for the ICESO and has great potential in real-world implementation. The non-convex modeling of MNC-ICES and the optimization performance of PD-TD3 is demonstrated in various scenarios. Compared with benchmark approaches, the proposed algorithm merits training speed, higher operational profits, and lower violations of multi-network constraints. Potential beneficiaries of this work include ICES operators and residents who could be benefited from improved ICES

operation efficiency, as well as reinforcement learning researchers and practitioners who could be inspired for safe RL applications in real-world industry.

2 - Deep Reinforcement Leaning-based Data Driven Method for Fast Solution of Stability Constrained Optimal Power Flow in Isolated Microgrids

Jun Wang, The Univeristy of Hong Kong, Hong Kong, Hong Kong

3 - A Nested Decomposition Method for the Ac Optimal Power Flow of Hierarchical Electrical Power Grids

Qi Wang, The Hong Kong Polytechnic University, Hong Kong, China, People's Republic of

The interconnected electrical power grids are operated by multilevel control centers with hierarchical architecture. Since large-scale renewables are integrated in different voltage levels, the traditional isolated operation of interconnected power grids is uneconomical and meets operational risk due to lack of coordination. Therefore, the coordination between control centers in different levels and areas are indispensable. However, existing distributed approaches may either encounter numerical problems or converge slowly when the nonlinear AC optimal power flow (ACOPF) models are applied. Moreover, the total iteration number presents an exponential increase with the levels of hierarchical grids. In order to improve computational efficiency, this paper proposes a nested decomposition method for the coordinated operation of multilevel ACOPF problem, which has superlinear convergence and can achieve the optimal solution (KKT point). During each iteration, a projection function, which embodies the optimal objective value of a lower level power grid projected onto its boundary variable space, is computed with second-order exactness. Thus, the proposed method can be applied to nonlinear continuous optimizations with high efficiency. The paper also provides a rigorous proof for its convergence under the condition that the lower level optimization problems are convex with continuously differentiable objectives and constraints. Numerical tests are conducted with three trilevel power grid of different scales, which verify that the computational efficiency and scalability of the proposed algorithm are superior to those of existing methods.

4 - A distributionally robust resilience enhancement strategy for distribution networks considering decision-dependent contingencies

Yujia Li, lawrence berkeley national laboratory, El Cerrito, CA, United States

MD11

Morrison Room

Advances in Online Resource Allocation

Invited Session

Revenue Management and Pricing

Chair: Zihao Li, National University of Singapore, Singapore, N/A, Singapore

1 - Beyond Regularity: Simple Versus Optimal Mechanisms, Revisited

Yiding Feng, Hong Kong University of Science and Technology, Hong Kong, Hong Kong, Yaonan Jin

A large proportion of the Bayesian mechanism design literature is restricted to the family of regular distributions [Mye81] or the family of monotone hazard rate (MHR) distributions [BMP63], which has overshadowed this rich and well-developed theory. We (re-)introduce two generalized families: quasi-regular distributions and quasi-MHR distributions. Likewise, the parameterized families of λ-regular (a.k.a. α-strongly regular) distributions [CR14,SS19], which smoothly interpolate regular distributions and MHR distributions, generalize to λ-quasi-regular distributions.

The significance of our new families is manifold. Firstly, their defining conditions are immediate “economic” relaxations of the original defining conditions (e.g., regularity as monotonicity of the virtual value functions), capturing key economic intuitions. Secondly, they satisfy natural mathematical properties (about order statistics) failed for the original families, thus technically more tractable. Thirdly, numerous results (by [BK96, HR09, CD15, DRY15, HR14, AHNPY19, JLTX20, JLQTX19, FLR19, GHZ19, JLX23, LM24] etc) known merely for the original families now can extend to our new families. Many of these extensions incur no quantitative loss, or

even improve the state of the art for the original families. Finally, beyond the third point, our new families guide us to entirely new perspectives and thus entirely unknown results. For example, regarding revenue maximization for symmetric versus asymmetric regular buyers, we acquire 0.5- versus 0.1908-approximations for the (less-than-)one-sample prophet inequalities, respectively. To the best of our knowledge, such results are blank in the literature, despite their widely-studied welfare maximization counterparts [CDFS19, RWW20, CCES20, CDFSZ21, CCES24].

2 - Fair and Efficient Online Matching with Reusable Resources: a Sequential Choice Perspective

Zihao Li, National University of Singapore, Singapore, Singapore, David Simchi-Levi, Zhenzhen Yan

In online ride-hailing platforms, when a passenger request arrives, the system must assign a suitable driver to fulfill the order. However, drivers may reject requests proposed by the platform, requiring the system to sequentially query multiple drivers until one accepts. Once a driver accepts a request, they temporarily leave the platform to complete the trip and return only after finishing it. Different matching decisions yield different rewards.

In this paper, we model this setting as an online matching problem with reusable resources under a sequential choice model. We incorporate a fairness constraint based on workload balance, while aiming to maximize the total expected reward. We propose a linear programming (LP)-based algorithm and prove that it achieves a competitive ratio of at least 0.5. To further investigate the trade-off between fairness and efficiency, we analyze the \emph{price of fairness} in this context. Finally, we conduct numerical experiments to evaluate the empirical performance of our proposed approach.

3 - Capacity Configuration in Non-Stationary Environments: the Operational Value of Prediction and Dynamic Decision-Making

Hao Su, HKUST, Hong Kong, Hong Kong, Piao Hu, Jiashuo Jiang, Guodong Lyu

We study how volume flexibility can be used to determine the optimal capacity level while satisfying designated performance targets, particularly when the demand is not only random but also exhibits variations in its distribution across the planning horizon. To address this challenge, we develop an online two-stage stochastic optimization framework over a finite planning horizon. During each period of the planning horizon, the planner makes the first-stage decision for capacity configuration prior to demand realization, and then allocates the capacity to fulfill multiple performance targets at the second stage when the demand is realized. We propose an adversarial learning approach to iteratively solve the capacity configuration and allocation problem in an online manner. Our main results show that this approach provides a provably near-optimal performance guarantee while offering great computational efficiency. Furthermore, we demonstrate the value of volume flexibility in capacity configuration to minimize potential losses, particularly when the underlying stochastic inputs follow non-stationary distributions. This extends previous studies that focused on stationary inputs and paves the way for establishing a transparent data-driven oracle that can be applied to a wide range of two-stage online decision-making problems. Through extensive numerical experiments, we underscore the benefits of demand distribution information in non-stationary environments and showcase how volume flexibility, in addition to the classic process flexibility, can help mitigate the misalignment between supply and demand. Furthermore, we apply these techniques and insights to address a crowdsourcing walker deployment problem arising in last-mile delivery operations, using real-world data from a logistics service provider.

4 - Online Stochastic Generalized Assignment Problem with Uncertain and Delayed Demand

Hao Wang, University of Science and Technology of China, Hefei, China, People's Republic of, Zihao Li, Jinming Liu, Limeng Liu, Zhenzhen Yan

We study an online generalized assignment problem under a stochastic arrival model. Items arrive sequentially, following a known independent and identical distribution. Each arriving item has an associated demand, which is stochastic and drawn from an unknown type-specific distribution. However, a key challenge is that the actual demand is only revealed after a delay, occurring some time after the item has been successfully packed into an offline bin. Such delayed demand revelation is common in many real-world applications, such as food delivery, where actual order sizes or customer consumption patterns may only be observed post-allocation.
Upon each item’s arrival, an immediate decision must be made: either pack it into a bin, reducing its capacity, or

reject it. The objective is to maximize the total reward of the packing scheme while adhering to capacity constraints. To tackle the uncertainty introduced by delayed demand revelation, we design an algorithm based on the exploration-exploitation principle, which simultaneously learns about the demand distribution and optimally allocates items. We provide a non-asymptotic parametric guarantee for discrete demand distributions supported on finite rational values and extend our analysis to arbitrary distributions. Finally, we conduct numerical experiments to evaluate the effectiveness of our algorithm across different demand distributions.

5 - Personalization vs. Marketplace Diversity: How Recommendation Algorithms Shape Pricing, Participation, and Consumer Search

Anil Omer Saritac, Singapore Management University, Singapore, Singapore, Nur Kaynar, Ozge Sahin

Recommendation algorithms are pivotal in mediating online marketplaces, generating substantial economic impact (e.g., ~30% page views, +18% consumer welfare). This research examines their simultaneous effects on seller pricing strategies, market participation decisions, and consumer search for suitably matched products. We develop and analyze a tractable model capturing these intricate demand- and supply-side dynamics between heterogeneous strategic consumers who search for products better suited to them and price-setting sellers.

We find that niche sellers, who offer products that are relatively less suitable for the average customer, charge lower posted prices. However, their “effective prices”, defined as the posted price plus the disutility from being less suitable, remain relatively higher. Relatively more personalized recommendation lowers posted prices by heightening competition among sellers that are better suitable to the average customer, even though it weakens competitive pressure from niche sellers. More personalization also narrows the spread of effective prices, facilitating customer search and boosting overall customer welfare. Despite driving down posted prices, increased personalization often raises seller profits by facilitating consumer search and increasing demand.

Building on the results above, although fully personalized recommendations may appear attractive, our findings indicate they can be suboptimal for overall welfare. Instead, preventing user attention from converging on only a few top sellers by recommending niche sellers with small probabilities can benefit both sellers and customers. The result holds true for most common demand functions, with our model identifying specific conditions under which exceptions arise.

MD12

Moor Room

Food Supply Chain Management I

Invited Session

Revenue Management and Pricing

Chair: Xue-Ming Yuan, National University of Singapore, Singapore, N/A, Singapore

Co-Chair: Dikky Indrawan

1 - Image Analysis of the Classification of Banana Quality by the Artificial Neural Network Approaches

Utami Syafitri, IPB University, Bogor, Indonesia, Yessie Sari, Nur Hasanah, Wulan Wahyuni, Faozan Faozan, Dase Hunaefi, Muhammad Syafiq, I Gusti Ngurah Sentana Putra, LIU PENG

The quality of a banana can be seen from its peel. The changing colour of the peel can determine the banana's nutrition and influence people to buy it. Digital data based on images has become popular lately because it can extract valuable information based on visual evidence. The characteristics of the data image are also challenging since it belongs to a large dataset. Machine learning methods can be used to analyze the large dataset and find a model to classify the data. This paper will implement several Artificial Neural Network approaches, such as BP Neural Network, RBP Neural Network, and SOM Neural Network. Before the analysis, data reduction methods will be implemented. The classical one is Principal Component Analysis, which is based on linear dimensionality reduction; however, the T-distributed stochastic neighbor embedding (t-SNE) can be used for more complex data and non-linear dimensionality reduction. Both methodologies will be compared. The research aims to find a suitable combination of reduction methods and a classification model based on the neural

network approach on the banana peel images. The source of data is https://universe.roboflow.com/real-images/shelf-life-of-banana.

2 - Perishable Inventory System and Value of Information

Mabel C. Chou, National University of Singapore, SG, Singapore, Luanxi Mao, Xiaobo Li

We study a periodic-review perishable inventory system that incorporates advanced perishability information (API). Unlike traditional models in the perishable inventory literature, we assume that product lifetimes are stochastic and can be revealed through a spectroscopy-based scanner. In each period, the firm makes replenishment decisions based on both the available inventory and revealed perishability information, with the objective of minimizing the total expected cost, which includes linear ordering costs, holding costs, backlogging or lost-sales penalties, and disposal costs. We develop a general model and derive the structural properties and optimal policies for two specific scenarios: the Baseline Model, where the retailer has only distributional information, and the API Model, where the retailer has perfect perishability information. Our numerical results demonstrate the substantial value of incorporating perishability information, showing significant cost reductions and improved decision-making under the API Model.

3 - The Design of a Banana Harvest Decision Support System Model

Marimin Marimin, Bogor Agricultural University, Bogor, Indonesia, Dikky Indrawan, Handaya Handaya, Utami Syafitri, Dase Hunaefi, Yessie Sari, Asaduddin Abdullah, Thabed Baladraf

Banana is one of the most widely cultivated fruits in tropical regions, contributing significantly to both local consumption and export markets. However, the decision-making process regarding the optimal harvest period is often based on manual observation and experience, leading to inconsistencies in quality and yield. This study proposes the design of a Banana Harvest Decision Support System (DSS) Model aimed at assisting farmers and agricultural stakeholders in making data-driven and accurate harvesting decisions. The proposed model integrates three essential sub-models. The first sub-model focuses on Estimating the Optimal Harvest Time, using agro-climatic data such as planting date, accumulated temperature, and plant age to determine the most suitable harvest window. The second sub-model is dedicated to Ripeness Prediction, which applies image analysis and classification algorithms to evaluate the maturity of bananas through visual features such as color and texture. The third sub-model involves Production Quantity Estimation, utilizing historical yield records, plant density, and current growth indicators to predict the expected output. Machine learning will be integrated in the sub models. These sub-models will be incorporated into a single decision support framework that can be deployed through user-friendly digital platforms. The system aims to enhance harvest precision, reduce losses due to premature or delayed harvesting, and support more efficient post-harvest logistics. The model will be validated using experts judgement and or available real plantation data.

4 - Conditional Generative Learning for Joint Pricing and Inventory Control

Xianghua Jiang, Chinese University of Hong Kong, HongKong, Hong Kong, xiangbin Hu, Xun Zhang, Zhisheng ye

We propose a provable deep conditional generative framework for data-driven joint pricing and inventory control with contextual information, in which the retailer makes pricing and inventory decision based solely on historical data consisting of demand, price and covariates. In contrast to the existing methods that directly learn the optimal policy curve as a function of the covariate, we develop a deep conditional generative learning approach to learn from historical data the whole demand distribution conditional on selling price and covariates. The conditional demand estimator is a deep conditional generator, which is a function of the selling price, covariates and a `noise' variable independent of price and covariates. The special structure of the deep conditional generator allows us to sample the noise easily and compute the gradient of the random cost in closed form, with which stochastic gradient descent can be designed for the subsequent policy optimization under both risk neutral and risk averse settings. We establish asymptotic properties of our data-driven policy including asymptotic optimality and consistency, and we provide a consistent estimator of the optimal expected cost. Through comprehensive simulations, we show that the expected cost evaluated at our data-driven policy is significantly smaller than those from state-of-the-art prescriptive benchmarks. In particular, we apply our method to real meal delivery data and find that our method beats the benchmarks by at least 10\% reduction in out-of-sample costs, and it is at least four times faster than three machine learning-based descriptive approaches.

MD13

Bailey

Recent Simulation Optimization Algorithms and Applications

Invited Session

Simulation Society

Chair: Chuljin Park, Hanyang University, Seoul, Korea, Republic of

1 - Stochastic Kriging Assisted Controlled Random Search for Measuring the Optical Critical Dimensions

Hyungjin Kim, Georgia Institute of Technology, Atlanta, GA, United States, Chuljin Park, Shing Chih Tsai, Aerim Hwang

We consider a simulation optimization problem with continuous decision variables, where the objective function can only be evaluated through noisy simulation. Such problems often appear in natural science and engineering domains and solving them efficiently often requires the use of surrogate models and global search heuristics. However, existing surrogate-assisted optimization algorithms typically update their surrogate models every time new data is observed. While effective in small- to medium-scale settings, this frequent updating of the model becomes computationally intensive in large-scale settings and may also suffer from data imbalance due to excessive exploitation of specific regions.

To address these limitations, we propose a new algorithm called Stochastic Kriging-assisted Controlled Random Search (SKCRS). The SKCRS algorithm combines the strengths of surrogate-based and random search methods. It utilizes a fixed stochastic kriging model, constructed initially, to estimate the objective function values and identify promising solutions among candidates generated by controlled random search. If the model fails to provide useful guidance (e.g., due to noise or model inaccuracy), the algorithm defaults to the original controlled random search to maintain exploration and robustness.

We tested the proposed and existing methods on 12 benchmark functions and a simulation-based case study for optical critical measurement in semiconductor manufacturing. The experimental results show that SKCRS performs comparably to, or better than, existing methods in terms of both solution quality and computational efficiency, particularly in settings where only one or two simulation observations can be obtained at each solution and random noise is relatively small.

2 - Handling a Safety State Constraint over a Finite Horizon via Ranking and Selection

Jaewon Lee, Hanyang University, Seoul, Korea, Republic of, Chuljin Park

We consider multi-period decision-making under uncertainty. In each period, a decision-maker observes the system’s state and then selects an action that directly influences the state transition. We assume that state transitions can be observed through stochastic simulation. Additionally, we aim to avoid situations where the system enters undesirable states, referred to as critical states. In this study, we introduce a safety-state constraint: the probability of being in a critical state needs to be less than or equal to a threshold during each period. Given this setting, we focus on identifying a set of feasible policies (i.e., the collection of actions at each state and each period) that comply with the state safety constraint over a finite horizon. Our primary objective is to solve this problem by reducing computational costs while ensuring the statistical validity of the results.

We propose a new algorithm designed to prune infeasible actions at each period in a forward direction. Specifically, we first employ a fully sequential indifference-zone procedure to assess the feasibility of each action based on the current state. Next, we prune all policies that include the actions identified as infeasible during each period. After this pruning process, the algorithm continues to simulate the system only using the actions identified as feasible. We prove the statistical validity of our algorithm and demonstrate its empirical performance through two numerical examples.

3 - A Fully Sequential Feasibility Check Procedure with Relaxed Tolerance Levels

Chuljin Park, Hanyang University, Seoul, Korea, Republic of, Sigrun Andradottir, Seong-Hee Kim, Yuwei Zhou

We consider the problem of finding feasible systems with respect to stochastic constraints when system performance needs to be evaluated through simulation. Our objective is to solve this problem by enhancing computational efficiency while guaranteeing statistical validity. Existing fully sequential indifference-zone (IZ) procedures introduce a fixed tolerance level, the amount that the decision maker is willing to be off from the threshold in the constraint. These procedures are derived under the assumption that all system means are exactly the tolerance level away from the threshold, causing unnecessary observations for systems whose means significantly differ from the threshold. In contrast, IZ-free procedures eliminate the tolerance level. While these procedures perform well in favorable mean configurations, they may require significantly more observations than IZ procedures when certain systems have means that are close to the threshold. This indicates that completely removing the IZ parameter is not always beneficial.

To tackle these challenges, we propose the Indifference-Zone Relaxation (IZR) procedure to improve the efficiency of IZ feasibility checks. IZR introduces a set of relaxed tolerance levels and simultaneously utilizes two subroutines for each level: one to identify systems that are clearly feasible and another to exclude those that are clearly infeasible. This approach enables feasibility checks at different tolerance levels, allowing for early conclusions for some systems while guaranteeing statistical validity. We mathematically prove that IZR can determine system feasibility with the desired probability and empirically demonstrate its effectiveness through numerical examples.

4 - Genetic Algorithm as An Alternative to A/B Testing for Feature Optimization

Rajesh Sinha, Indian Institute of Management Indore, Indore, India

A/B testing is akin to the experimental approach, where treatments are the elements of ad-copies, product features and likes. A/B testing performs feature comparisons to find superior features from a set of features. A/B testing, although experimental, has some limitations. Algorithmic assignments of treatments under A/B testing may not be random assignments; therefore, there results of A/B testing may not be equivalent to the results that can be obtained when the assignments are actually random. In addition, while conducting A/B testing, only a part of the targetable population may be exposed to the treatments of the A/B testing, making its results unreliable for the full population. Further, under A/B testing an overall optimal version is difficult to identify. In addition, A/B testing can help in customer behaviour optimization; however, the profit optimization goal may not happen in a rather direct way. The research proposes genetic algorithm as a way to overcome the above mentioned limitations with A/B testing. Genetic algorithm can provide a quicker way to find the set of optimal features.

5 - A Unified Robust Optimization Approach for Problems with Costly Simulation-Based Objectives and Constraints

Yanzhan Chen, School of Traffic & Transportation Engineering, Central South University, Changsha, China, Changsha, China, People's Republic of, Liang Zheng

Abstract: This study proposes a unified robust optimization approach to address min-max problems involving expensive simulation-based objectives and constraints impacted by implementation errors and parameter perturbations. This approach optimizes the worst-case scenarios of stochastic simulation responses across multiple evaluation criteria to achieve robust efficient solutions. It integrates multiple objectives and constraints into a cohesive framework, featuring a novel performance metric designed to rigorously assess solution quality. This metric can simplify the inner constrained multi-objective maximization problem into an unconstrained, stochastic, and single-objective minimization problem, based on which a softened condition is provided to identify robust efficient solutions. Then, these neighborhood exploration and robust local move mechanisms leverage infeasible neighbors' information to guide the iterative solution process towards a globally robust efficient point. To mitigate computational costs, surrogate models of simulation-based objectives and constraints are utilized to guide the initial exploration of worst-case neighbors. The proposed approach's effectiveness and

superior performance are demonstrated through test results on four synthetic multi-objective robust optimization problems with constraints. Furthermore, the approach is utilized to design robust traffic signal timing plans under cyber-attacks and uncertain traffic volumes, yielding satisfactory results within limited simulation budgets. This approach presents a promising tool for addressing constrained multi-objective simulation-based optimization problems under uncertainty.

MD14

Wilberforce

Smart Scheduling & Operational Intelligence

Invited Session

Semiconductor Manufacturing x AI

Chair: Zixin Li, Micron, N/A

Co-Chair: Lin Miao Zhang, Micron Technology, Singapore, N/A, Singapore

1 - Balancing Long Term and Short Team Planning by Leveraging Operations Research in Discrete Manufacturing Production

Jeffrey Zhang, Lumileds, Singapore, Singapore, Lei Wang

A global leader in lighting solutions faces significant challenges in managing the complexity of its supply chain due to highly configurable products and a widespread manufacturing network. With thousands of product characteristics and assembly part numbers, the company must efficiently locate materials across numerous international factories and warehouses. Additionally, the company is committed to confirming and scheduling product runs within 48 hours of receiving customer orders. For years, the company has relied on advanced optimization software to manage these tasks. However, the company sought to enhance its planning capabilities by gaining a more strategic, global perspective on its operations. Challenges arose from having different manufacturing methods across sites and dealing with subcontractors, which made it difficult to assess long-term resource needs. To address this, the company implemented a system that enables better scenario analysis, helping planners predict how changes in capacities or resource allocation would impact production and fulfillment timelines. This improved visibility allows the company to plan more effectively for growth and optimize operations across multiple sites.

2 - Support Fab Operations Using Multi-Agent Reinforcement Learning

Jeroen Bedorf, Minds AI, Santa Cruz, CA, United States

Over recent years semiconductor operations have grown in scope, dynamics, and complexity. Consequently, neither advanced scheduling algorithms nor manufacturing engineers can quickly estimate the best-performing schedule. In this presentation we present how machine learning, in real time, can be used to augment and support manufacturing engineers. The presented results, obtained from production deployments in Micron’s fabs, demonstrate that AI-assisted scheduling can improve Key Performance Indicators (KPIs) with little to no downside.

Modern 300MM semiconductor fabrication facilities face challenges such as complex hierarchical structures, high financial stakes, and dynamic processes. These facilities contain thousands of tools with unique capabilities, resulting in intricate dispatching schedules. Frequent schedule updates are required due to maintenance, failures, and shifting priorities. This constant need for updates makes it impractical for Manufacturing Engineers (MEs) to manually optimize scheduling parameters.

This presentation demonstrates how reinforcement learning (RL) can aid manufacturing engineers (MEs) in scheduling decisions. A key goal is to improve MEs' well-being by mitigating time constraints arising from dynamic business priorities. Machine learning assists MEs by handling routine tasks, allowing them to focus on rare events. RL uses simulations and historical data to optimize scheduling based on current conditions, offering MEs suggestions to enhance KPIs like throughput, daily targets, tool idle time, and critical queue times.

3 - Semiconductor Equipment Capacity Optimization Through Operation Research

Chen Yue, Micron Technology, Singapore, Singapore

The semiconductor industry is characterized by its relentless pace of innovation and rapid evolution, with new breakthroughs and advancements occurring regularly. The fast-evolving nature of the industry further complicates capacity planning, as manufacturers must anticipate and adapt to technology shifts while efficiently managing expensive production resources and capital investments. At Micron, capacity planning has become increasingly difficult and challenging due to rapid technological advancements, increasing number of parts and tool types, and demand volatility. These challenges underscore the importance of agile and data-driven capacity planning strategies that can adapt to a rapidly changing business landscape. The conventional way of using Excel spreadsheets in capacity planning, while versatile, has several disadvantages including limited problem size and being prone to human error. To address that, we have developed software which uses Linear Programming (LP) technologies to help Micron deliver consistent and optimize capacity decisions. In this session, we will present the data foundation and LP methodology of the software and its various advantages.

4 - Optimizing Stack Injection for Immersion Lithography: Challenges, Innovations, and Impact

Xun Li, Micron Technology, Inc., Singapore, Singapore, Andrew Wang

The photolithography (Photo) immersion workstation is often identified as the primary bottleneck in fabrication (fab) plants due to the high cost of immersion tools. Enhancing the performance of these tools is crucial for overall fab efficiency. A key factor in optimizing immersion performance is achieving a balanced load across all immersion tools. However, advancements in wafer fabrication technology have led to increased stacking requirements and additional immersion mask layers, complicating the management of immersion loading balance. Traditional methods of managing immersion loading have become less effective, resulting in higher standby rates and wafer output losses due to unbalanced loading.

To address these challenges, our team developed an innovative solution utilizing a machine learning algorithm, specifically a genetic algorithm (GA), combined with artificial intelligence (AI). This system considers multiple constraints in production management and uses the remaining accumulative load index as a fitness measure to resolve the wafer injection problem for the first layer of the immersion stacking group. This approach has significantly improved loading balance across stacking tools and enhanced wafer output capabilities from these constrained workstations.

5 - Innovative Optimized Line Strategies for Enhanced Fab Output and Demand Fulfillment

lina wang, Micron Technology, Inc., Singapore, Singapore

In today’s increasingly complex semiconductor manufacturing environment, delivering consistent fab output while meeting dynamic demand profiles remains a formidable challenge. Traditional production line strategies often fall short in responding to real-time variability in resource availability, demand shifts, and production constraints—resulting in delays, inefficiencies, and unmet customer expectations.

This case study explores Micron’s industrial-scale implementation of innovative line strategies designed to enhance fab performance and ensure demand fulfillment for critical components. The initiative, deployed through Micron’s Line Prioritization system and integrated with Micron’s Scheduling System, introduced two pivotal mechanisms: Optimized Line Strategies, which dynamically adapt routing decisions based on operational targets and bottleneck capacities, and Demand Penalty, which protects high-priority part types from being deprioritized in congested conditions.

Tested across multiple Asia-based fabrication sites from end 2024 to beginning of 2025, the system demonstrated measurable improvements in throughput, responsiveness, and production line flexibility. This study highlights the practical impact of deploying intelligent line strategies within complex manufacturing networks, offering a scalable approach to balancing efficiency with demand agility in high-stakes industrial operations.

MD15

Fairmont Ballroom

Dynamic and Stochastic Logistics Models

Contributed Session

Transportation Science and Logistics

Chair: Chieh-Ju Juan, Republic of China Air Force Academy, Kaohsiung City, Taiwan

1 - Dynamic Matching and Pricing for Crowdsourced Delivery Couriers with Random Utility Theory

Chunwei YANG, City university of hong kong, Hong Kong, Hong Kong, Andy Chow, Yang Deng

Crowdsourced delivery has become an increasingly popular solution in urban logistics, enabled by digital platforms and mobile applications. It allows service providers to reduce operating costs and manage flexible workforces. However, this flexibility introduces uncertainty in courier availability and willingness to accept delivery tasks.

This paper proposes a dynamic matching and pricing framework for on-demand delivery platforms to incentivize courier participation and improve system-wide profitability. Couriers' acceptances of orders are modeled by a random utility function with both observable attributes and hidden factors.

We formulate the problem as a Markov Decision Process (MDP), incorporating the stochastic customer demand and courier behaviors. To overcome the computational challenges, we adopt an Approximate Dynamic Programming (ADP) approach. Specifically, we approximate the long-term value function using a neural network trained via temporal-difference learning. This allows the platform to account for the future impact of current decisions and maintain a forward-looking control policy.

Computational results show that the proposed matching and pricing strategy could enhance courier acceptance rates and hence platform revenue compared to fixed or static strategies. This work contributes to intelligent logistics management by integrating advanced optimization techniques and artificial intelligence into real-time operational decision-making.

2 - Optimizing Freight Assignment in Shared High-Speed Railways Under Wasserstein Ambiguity

Kanglin Liu, Beijing Jiaotong University, Beijing, China, People's Republic of, Zhuolin Wang

High-speed railway express freight services play a crucial role in modern transportation systems. This paper explores a model integrating freight and passenger transportation within high-speed rail systems, addressing uncertainties in both freight and passenger demands. Leveraging a two-stage distributionally robust optimization approach with Wasserstein ambiguity sets, the study aims to optimize freight allocation strategies and quantities. The problem is reformulated into a unified mixed-integer programming framework, enabling the identification of worst-case scenarios and distributions. A decomposition algorithm is employed to solve the problem efficiently, achieving a performance improvement of up to tenfold.

3 - An Adaptive Bus Bridging Strategy Under Rail Transit Disruptions via A Two-Stage Reinforcement Learning Driven Approach

Junyi WANG, City University of Hong Kong, Kowloon Tong, China, People's Republic of, Andy Chow

Bus Bridging is widely adopted during unpredictable disruptions in the Urban Rail Transit (URT) system. However, the adaptability and flexibility of bus bridging service design (BBSD) under a network-wide scope disruption have been bottlenecks in this field. Research in this area often grapples with the trade-off between model complexity and solution efficiency. To address this, we propose a two-stage pre- and post-emergency plan framework that merges the strengths of model-based approaches with reinforcement learning (RL) techniques. Before a disruption occurs, candidate routes are prepared using a column generation (CG) procedure. The adaptive scheduling problem under limited bus vehicles is then formulated as a Markov Decision Process (MDP) and trained by a Dueling Double Deep Q-Network (D3QN) algorithm through iterative system transition simulations. Combining the objective value of the CG’s master problem and the reward of the MDP model, the route choice is made with a systematic consideration that minimizes the sum of the total additive travel time and bus waiting time. When a sudden disruption emerges, first backtrack the pre-emergency plan that matches the

current situation and the artificial neural network (ANN) parameters trained before. Then the agent can synchronically adjust dispatching headway and bus vehicle type on all selected bridging routes regarding real-time stranded passenger arrival rate data. Numerical experiments are conducted with real-world data collected from Hong Kong's light rail system. By benchmarking with a purely model-cased approach, the proposed RL-driven strategy reveals a significant reduction in passenger waiting time and computational time.

4 - Boosting Shared Mobility Efficiency: a Closed-Loop Optimization Approach to Demand Forecasting, Pricing, and Vehicle Distribution

Yufei Li, Tianjin University, Tianjin, China, People's Republic of, Zhaofang Mao

With the rapid development of shared mobility platforms, common issues such as low operational efficiency, underutilization of resources, and poor user experience arise. Particularly during periods of high demand fluctuations and resource scarcity, forecasting user demand, dynamically adjusting pricing, and optimizing vehicle distribution and parking locations have become key challenges in improving platform performance. This paper proposes a closed-loop optimization framework that integrates demand prediction → dynamic pricing → vehicle distribution and parking location optimization, aiming to achieve intelligent scheduling and resource allocation for shared mobility platforms. First, based on the Long Short-Term Memory model, this paper accurately predicts the spatiotemporal demand of the platform, capturing temporal changes and spatial fluctuations in demand patterns. Then, a dynamic pricing mechanism combining Reinforcement Learning is designed, where real-time price adjustments are made to balance supply and demand, alleviating resource congestion during high-demand periods. Finally, based on demand prediction and pricing strategy, a vehicle distribution and parking location optimization model is proposed, combining Genetic Algorithm and Game Theory, aimed at adaptively adjusting vehicle distribution and parking locations to minimize empty mileage, improve resource utilization, and optimize operational performance.Experimental results show that the proposed framework significantly improves platform efficiency in various simulation scenarios. Compared to traditional single-objective optimization strategies, it responds more flexibly to demand fluctuations, reduces empty mileage, enhances user satisfaction, and optimizes overall platform revenue. This integrated approach provides new research insights and practical solutions for intelligent scheduling and resource optimization in shared mobility platforms.

5 - Uses of Generative AI in Supply Chain Decision Making

Mohammad Moshref Javadi, University of Illinois at Urbana-Champaign, Champaign, IL, United States, Amlan Baruah

This paper presents a comprehensive framework for applications of Generative Artificial Intelligence (GenAI) in Supply Chain Management (SCM). The framework is structured across five levels to detail various processes and functions in supply chain. The study investigates the applications of GenAI models within the framework, highlighting the unique contributions of different architectures, such as Generative Adversarial Networks (GANs), Transformers, Variational Autoencoders (VAEs), and flow-based models. After reviewing more than 400 articles, the paper reviews practical use cases and challenges, and emphasizes the need for more research in areas, such as unstructured data analysis, contextual understanding, and scalability of GenAI models for broader adoption in SCM. We present the potential of GenAI to address key challenges in SCM processes and therefore, future research directions in this area.

6 - Dynamic Multi-Period Location-Routing Under Fuzzy Demand with Adaptive Resequencing

Chieh-Ju Juan, Republic of China Air Force Academy, Kaohsiung, Taiwan, Taiwan

Location-routing problems (LRP) play a crucial role in enterprise logistics distribution and route planning. Traditionally, LRP studies consider facility location as a strategic decision, remaining fixed throughout the planning horizon due to high establishment costs, while routing decisions are adjusted tactically based on dynamic demand. This study addresses scenarios involving mobile sales facilities, simultaneously considering dynamic adjustments in location-allocation and routing to respond effectively to fluctuating customer demands. Given the uncertainty inherent in customer demands, this research employs fuzzy credibility theory to model demands as fuzzy variables and formulates the problem using a fuzzy chance-constrained programming

approach. To effectively manage dynamic and uncertain logistics environments, Genetic Programming (GP) is utilized to automatically evolve heuristic methods, improving foundational heuristic strategies, and developing adaptive Routing Policies (RPs). These policies significantly enhance the system's responsiveness and flexibility in dealing with varying demands. Preliminary results indicate that Routing Policies incorporating fuzzy demands effectively adjust facility location and routing decisions in partially dynamic environments, demonstrating promising optimization potential in practical logistics management.

MD16

VIP A

Sustainability and CSR

Invited Session

Sustainable and Socially Responsible Operations

Chair: S Viswanathan, Nanyang Technological University, Singapore, Singapore

1 - Corporate Social Responsibility Committee: International Evidence

Yuxia Zou, Nanyang Technological University, Singapore, Singapore, Xi Li, Jenny Chu

We provide worldwide large-sample evidence of a recent innovation in corporate governance: the voluntary creation of a separate board committee to oversee corporate social responsibility (CSR) activities. Our evidence suggests that a firm’s decision to voluntarily adopt a CSR committee is shaped by external demands from shareholders and other stakeholders, as well as internal needs and costs associated with forming a separate CSR committee. Upon the formation of CSR committees, firms enhance their internal CSR management practices and experience improvements in future environmental and social outcomes. Notably, the environmental improvements are more pronounced among firms operating in industries where environmental concerns are material and in countries with stronger environmental efforts. Overall, our evidence suggests that CSR committees represent a substantive corporate governance mechanism in improving firms’ environmental and social impact.

2 - Algorithmic Nudging for Energy Savings and Environment: Evidence from An Iot Platform

Jacob Zeng, Gonzaga University, Spokane, WA, United States, Nur Sunar, Nil Karacaoglu, Serasu Duran

Excessive energy consumption leads to substantial financial and environmental costs for commercial organizations. While digital nudges have shown promise in influencing individual behavior, the effectiveness of real-time feedback in reducing organizational energy consumption remains underexplored. This study investigates how real-time alerts from an Internet of Things (IoT) energy platform impact energy use in commercial retail settings and identifies the optimal frequency of these alerts to maximize energy savings.

3 - Sustainability-Linked Financing: Bringing Supply Chain Finance and Supplier Responsibility Together

Stuti Arora, Goa Institute of Management, Goa, India, Avijit Raychaudhuri

Several firms, such as Levi's and Puma, have partnered with financial intermediaries like the International Finance Corporation to provide sustainability-linked financing to their suppliers. The primary goals of this kind of supply chain financing are: (i) providing low-cost financing to the suppliers, and (ii) enticing the suppliers to improve their environmental and social responsibility. Under such a financing scheme, the financial intermediary rewards a supplier exerting high responsibility effort with a low financing fee. In this paper, we analyze the effectiveness of sustainability-linked financing by studying the interaction between a profit-maximizing supplier and a social welfare-maximizing financial intermediary. We find that under certain conditions, suppliers with varying responsibility efforts can be charged the same financing fee, and suppliers with the same responsibility effort can be charged different financing fees, which is in contrast to practice. We further find that such a financing scheme effectively achieves both its goals when the financial intermediary's cost of financing is low, and the supplier's cost of exerting responsibility effort is moderately low. On the other hand, neither of the two goals is effectively achieved if the financial intermediary's cost is high and the supplier's cost is moderate. Finally, we observe that improving the supplier's efficiency of exerting responsibility effort does not always improve his profit.

4 - The Impact of Remanufacturing on Leasing and Selling Business Models

S Viswanathan, Nanyang Technological University, Singapore, Singapore, Preethi Kulasekara Manavala Ramanujam, Juan Ramon L Senga

Our study explores leasing business model, often promoted for its environmental and economic benefits. By retaining ownership, firms are incentivized to recover and reuse products. Companies like Xerox, Dell, and HP use leasing to support sustainability efforts. However, there is an ongoing debate about its environmental merits. Some studies suggests that leasing is not always more profitable or sustainable than selling, with outcomes depending on how products are managed throughout their lifecycle.

Building on this, we examine joint adoption of leasing along with remanufacturing. Using a discrete-time, infinite-horizon game between a firm and its customers, we determine optimal production and pricing decisions solved through backward induction. The firm offers new, remanufactured, and used products, considering customer heterogeneity. Our analysis identifies firm's strategies such as zero or full or partial refurbishment and remanufacturing, determined by the cost structures and product utility.

Our findings suggest that high disposal costs encourage full remanufacturing/refurbishment, while lower costs favor partial refurbishment and remanufacturing. Computational experiments show that incorporating remanufacturing generally makes leasing both more profitable and environmentally advantageous compared to selling. However, this advantage diminishes if remanufacturing process has a high environmental impact, or if remanufacturing does not significantly reduce product's emissions during use.

Tuesday, July 22, 8:00 AM - 9:30 AM

TA01

Enterprise

Applied Probability and Learning Theory

Invited Session

Applied Probability Society

Chair: Anand Deo, Indian Institute of Management Bangalore, Bangalore, India

1 - The Scaling Behaviours in Achieving High Reliability via Chance-Constrained Optimization

Anand Deo, Indian Institute of Management Bangalore, Bangalore, India, Karthyek Murthy

We study the problem of resource provisioning under stringent reliability or service level requirements, which arise in applications such as power distribution, emergency response, cloud server provisioning, and regulatory risk management. With chance-constrained optimization serving as a natural starting point for modeling this class of problems, our primary contribution is to characterize how the optimal costs and decisions scale for a generic joint chance-constrained model as the target probability of satisfying the service/reliability constraints approaches its maximal level. Beyond providing insights into the behavior of optimal solutions, our scaling framework has three key algorithmic implications. First, in distributionally robust optimization (DRO) modeling of chance constraints, we show that widely used approaches based on KL-divergences, Wasserstein distances, and moments heavily distort the scaling properties of optimal decisions, leading to exponentially higher costs. In contrast, incorporating marginal distributions or using appropriately chosen f-divergence balls preserves the correct scaling, ensuring decisions remain conservative by at most a constant or logarithmic factor. Second, we leverage the scaling framework to quantify the conservativeness of common inner approximations and propose a simple line search to refine their solutions, yielding near-optimal decisions. Finally, given N data samples, we demonstrate how the scaling framework enables the estimation of approximately Pareto-optimal decisions with constraint violation probabilities significantly smaller than the Omega(1/N)-barrier that arises in the absence of parametric assumptions.

2 - Asymptotic Optimality Theory of Confidence Intervals of the Mean

Vikas Deep, Kellogg School of Management, Northwestern University, Evanston, IL, United States, Achal Bassamboo, Sandeep Juneja

We address the classical problem of constructing confidence intervals (CIs) for the mean of a distribution, given N i.i.d. samples, such that the CI contains the true mean with probability at least $1-\delta$, where $\delta \in (0,1)$. We characterize three distinct learning regimes based on the minimum achievable limiting width of any CI as the sample size $N_\delta \to \infty$ and $\delta \to 0$. In the first regime, where N_δ grows slower than $\log(1/\delta)$, the limiting width of any CI equals the width of the distribution's support, precluding meaningful inference. In the second regime, where N_δ scales as $\log(1/\delta)$, we precisely characterize the minimum limiting width, which depends on the scaling constant. In the third regime, where N_δ grows faster than $\log(1/\delta)$, complete learning is achievable, and the limiting width of the CI collapses to zero, converging to the true mean.

We demonstrate that CIs derived from concentration inequalities based on Kullback-Leibler (KL) divergences achieve asymptotically optimal performance, attaining the minimum limiting width in both sufficient and complete learning regimes for distributions in two families: single-parameter exponential and bounded support. These results extend to one-sided CIs, with the width notion adjusted accordingly. Finally, we generalize our findings to settings with random per-sample costs, motivated by applications such as stochastic simulators and cloud service selection. Under a cost budget C_δ, we identify analogous learning regimes and characterize the optimal CI construction policy.

<div justify-start"="">

3 - Beyond Expected Regret: Characterizing the Reward Distributions in Online Allocation Problems

Buse Unluturk, Singapore University of Technology and Design, Singapore, Singapore, Karthyek Murthy

Considering online resource allocation problems in which irrevocable decisions are made with limited information, we seek to tightly characterize the probability distributions of the cumulative reward and the regret incurred by employing any dual-based algorithm. Existing analyses focus exclusively on expected regret's dependence on the time horizon, leaving other aspects of the regret distribution – such as its variance, skew, tail behaviour, etc. -- and their relationship with other key problem parameters less well understood. To address this gap, we develop a novel Diffusion Approximation Toolbox, which can flexibly characterize the limiting reward & regret distributions resulting from the use of a wide range of statistical estimators available for estimating dual prices. This flexibility enables analysis of both well-studied and unexplored algorithms, broadening performance analysis in online resource allocation. Leveraging the martingale functional central limit theorems, our diffusion approximations reveal a degree of universality in reward accumulation, resource consumption, and regret behaviour, regardless of the underlying problem primitives. The diffusion approximations also lend novel insights into the robust effectiveness of the re-solving heuristic by demonstrating the similarity of its resource consumption dynamics with that of the Brownian bridge tied to zero.

4 - *Sequential Procedures for Early Detection of Machine Faults*

RAJESH SUNDARESAN, Indian Institute of Science, Bengaluru, India

Motivated by the application of detecting rolling element bearing faults, we study a non-Bayesian quickest change detection problem, where the change is from a known pre-change distribution to an unknown post-change distribution, but with parameters belonging to a known set. Recognising the cyclostationary nature of the rolling element bearing vibration signals, which possess correlations in the frequency domain, we construct two cyclic spectrum statistics that detect cyclostationarity in the vibration signal. We then develop sequential decision algorithms based on the above statistics and demonstrate their capability to detect small changes at low signal-to-noise ratio, with reasonable delay to detection, low false alarm rate, and low computational complexity. The talk is based on joint work with Krishna Chaythanya K. V. and Anurag Kumar.

TA02

Mercury

Operational Optimization for Delivery Logistics in Urban Air Mobility and Commercial Aviation

Invited Session

Aviation Applications

Chair: Ang Li, Hong Kong Polytechnic University, Kowloon, CA, Hong Kong

1 - On Coordinated Drone-Courier Logistics for Intra-City Express Services

Shuiwang Chen, The Hong Kong Polytechnic University, Hong Kong, Hong Kong

Drones, despite being acknowledged as a transformative force in the city logistics sector, are unable to execute the last-meter delivery due to airspace restrictions and safety concerns. To leverage advancements and overcome the limitations of drones in providing intra-city express services, we introduce a coordinated drone-courier logistics system where drones operate within a closed network among vertiports, while couriers connect customers to the drone delivery system. This paper aims to shed light on this coordinated system in terms of system feasibility, network interactivity, and long-term sustainability. We develop an integrated optimization model to optimize the network planning of the coordinated logistics system. The interplay between network planning and tactical operations is mirrored by a queueing network model, resulting in the nonlinear and nonconvex feasible region of the optimization model. An iterative exact algorithm that tightens lower and upper bounds by adaptively refining the linear approximations of nonlinear constraints is developed to provide optimality-guaranteed solutions with finite convergence. The computational experiments demonstrate the scalability and robustness of our algorithm across various network configurations and scenarios. The case study, based on a real-world dataset from SF Express validates that the coordinated logistics system efficiently attains cost and time savings by leveraging the effective turnover of drones and the coordination between drones and couriers. The optimal network design features a concentrated structure, streamlining demand consolidation and reducing deadhead repositioning. Moving forward, maintaining enduring cost and time savings necessitates a dual focus on advancing drone technologies and reducing courier delivery expenses.

2 - Efficient Hangar Slot Scheduling with Branch-and-Price

Haiyang KONG, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, Lingxiao Wu

The rapid growth of global flight demand has driven the development of the aircraft maintenance repair and overhaul (MRO) industry, prompting airlines to increasingly outsource heavy maintenance operations. Aircraft heavy maintenance requires the coordination of long-term tasks, specialized workforce and shift-based operations to minimize costs while pursuing on-time delivery. However, due to limited hangar slots and maintenance technicians, MRO companies have to optimize their maintenance plans to allocate maintenance resources effectively and maximize the completion of maintenance orders. Maintenance scheduling models in previous studies often fall short in addressing multi-skilled team allocation, shift continuity, and dynamic penalties for delivery delays. This study proposes an integrated optimization framework for shift-based maintenance scheduling to allocate hangar slots and maintenance shifts, and to form maintenance teams with precise skill combinations without hierarchical replacement. The MILP model formulated incorporates shift-level decision variables for task-to-hangar allocation, employee assignment, and delay penalties, subject to constraints on hangar capacity, skill matching, shift continuity, workforce availability, and delivery time windows. The objective minimizes total costs incurred for serving aircraft, encompassing hangar usage, labor expenses, and delay penalties. To solve large-scale instances, we reformulate the model as a set-partitioning model and design a tailored branch-and-price algorithm. Experimental results on real-world instances demonstrate the framework’s efficacy in reducing costs compared to MILP solved by a commercial solver, achieving balanced resource utilization and on-time delivery. This work provides a practical tool for aviation operators to harmonize complex maintenance logistics, workforce availability and skill requirements, and cost objectives.

3 - Dynamic Speed Optimization and Berth Reallocation for Autonomous Vessels Under Sailing TIME Disturbances

Tingting Chen, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, Lingxiao Wu, Shuaian Wang

Liner vessels, which operate on predetermined routes, need to maintain scheduled arrival times and berth allocations at a sequence of ports. Sailing time disturbances can lead to deviations from these schedules, resulting in arrival time delays and increased handling costs. Adaptive adjustments to sailing speeds are often necessary to mitigate these disturbances. The advent of autonomous vessels (AVs) has enabled dynamic, remote,

and precise speed adjustments. Operators can now make real-time speed adjustments for AVs in response to on-route disturbances, thus ensuring greater adherence to the planned schedules. This paper presents a dynamic optimization approach for AV speed optimization and berth reallocation under the collaboration between liner shipping companies and port operators to address operational disturbances effectively. The problem is addressed by using a rolling horizon framework, where the decision problem in each decision period is formulated as a mixed integer nonlinear programming (MINLP) model. A tailored branch-and-cut algorithm is developed to solve the model efficiently. Extensive numerical experiments demonstrate the approach's efficacy in solving the problem. The study further investigates the impacts of varying operational factors, such as the number of visiting ports per decision period, sailing time disturbance, penalty cost for the arrival time delay, fuel price, and sailing speed, providing strategic managerial insights for the optimization of AV operations.

4 - A Reliable Predict-Then-Optimize Approach for Minimizing Aircraft Fuel Consumption

Ziming WANG, The Hong Kong Polytechnic University, Hong kong, China, People's Republic of, Dabin XUE, Lingxiao Wu, Ran Yan

Achieving sustainability in aviation necessitates optimizing loaded fuel to reduce both financial costs and environmental impact, as loaded fuel directly affects aircraft weight, which in turn influences fuel consumption throughout the flight. This study develops a reliable predict-then-optimize approach for minimizing aircraft fuel consumption. First, artificial intelligence-based models are developed to predict fuel consumption rates using Quick Access Recorder data. Then, based on accurate fuel consumption predictions, a data-driven optimization model is further established to determine the minimum loaded fuel, assisting dispatchers in airlines with flight planning. We rigorously prove that under mild assumptions, the approach can return the minimum loaded fuel with given reliability within polynomial times. Experiments were conducted using the four most widely used aircraft models, i.e., A320, A321, B737, and B738. The results show that optimized loaded fuel can achieve an average fuel consumption reduction of 3.67% compared to actual consumption.

TA03

Minto

Computing for Discrete Optimization Problems

Invited Session

Computing Society

Chair: Rui Chen, Chinese University of Hong Kong, Shenzhen, Shenzhen, N/A

1 - Accelerating Benders Decomposition for the P-Median Problem Through Variable Aggregation

Rick Willemsen, Erasmus University Rotterdam, Rotterdam, Netherlands, Daniel Aloise, Raf Jans

The p-median problem is a classical location problem where the goal is to select p facilities while minimizing the sum of distances from each location to its nearest facility. Recent advancements in solving the p-median and related problems have successfully leveraged Benders decomposition methods. The current bottleneck is the large number of variables and Benders cuts that are needed. We consider variable aggregation to reduce the size of these models. We propose to partially aggregate the variables in the model based on a start solution; aggregation occurs only when the corresponding locations are assigned to the same facility in the initial solution. In addition, we propose a set of valid inequalities tailored to these aggregated variables. Our computational experiments indicate that our model, post-initialization, provides a stronger lower bound, thereby accelerating the resolution of the root node. Furthermore, this approach seems to positively impact the branching procedure, leading to an overall faster Benders decomposition method.

2 - Inverse-Feasible Regions In Discrete Optimization: Approximations and Applications

Jerry Chua, NUS Business School, Singapore, Singapore, Ian Zhu

Given a feasible solution to an optimization problem, the inverse-feasible region (IFR) describes the set of objective function parameters for which this solution remains optimal. The inverse-feasible region lies at the core of a wide range of problems spanning operations research and mathematical programming. In this talk, we derive new insights for approximating IFRs for discrete optimization problems by leveraging the linear relaxation of the discrete feasible region. Specifically, we introduce a novel concept of a restricted linear

relaxation, provide analysis on their ability to approximate the IFR, and use them to develop new algorithms for solving inverse optimization problems.

3 - Inexact Column Generation for Causal Structure Learning

Rui Chen, Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Yiran Yang

Causal structure learning, the task of inferring causal relationships from data, is computationally challenging due to its combinatorial nature. State-of-the-art integer programming formulations suffer from exponential growth in the number of variables and constraints, while traditional column generation approaches struggle with the complexity of solving mixed-integer nonlinear programming (MINLP) pricing problems. In this talk, we propose an inexact column generation method that reformulates the pricing problem as a difference of submodular minimization problem. This approach significantly reduces the computational complexity of pricing subproblems while maintaining the quality of generated columns. Empirical results demonstrate that our method outperforms pure integer programming approaches, such as GOBNILP, in both solution quality and computational efficiency.

4 - Information Theoretic Subset Selection

Junkai Zhao, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Taotao He, Jun Luo

Information-theoretic subset selection methods, which employ criteria from information theory to identify informative features or data subsets, are widely adopted in large-scale data preprocessing. However, the inherent non-convexity of such problems—often arising from fractional or logarithmic terms in their formulations—poses significant challenges in obtaining globally optimal solutions. This paper focuses on two representative problems: maximum relevance minimum redundancy (mRMR) feature selection and Kullback-Leibler (KL) divergence-based review selection. To address these challenges, we propose exact mixed-integer linear programming (MILP) reformulations via convex relaxation. Leveraging modern optimization techniques, our method transforms the original non-convex problems into tractable MILP models while preserving optimality guarantees. Numerical experiments demonstrate that the proposed framework solves information-theoretic subset selection problems hundreds of times faster than global optimization methods while achieving significantly superior objective values compared to heuristic approaches, revealing the dual advantages of our methods in computational efficiency and solution quality.

TA04

Indiana

Operations Research based on Machine Learning

Invited Session

Computing Society

Chair: Zhao-Rong Lai, Jinan University, Guangzhou, N/A

1 - Total Variation Based Invariant Risk Minimization and An Out-of-Distribution Generalization Scheme

Zhao-Rong Lai, Jinan University, Guangzhou, China, People's Republic of

Invariant risk minimization (IRM) is an arising approach to generalize invariant features to different environments in machine learning. While most related works focus on new IRM settings or new application scenarios, the mathematical essence of IRM remains to be properly explained. We verify that IRM is essentially a total variation (TV) of the learning risk with respect to the classifier variable. Moreover, we extend IRM-TV to a Lagrangian multiplier model named OOD-TV-IRM, which is essentially a primal-dual optimization model. The objective is to reach a semi-Nash equilibrium where the balance between the training loss and OOD generalization is maintained. We also develop a convergent primal-dual algorithm that facilitates an adversarial learning scheme.

2 - Almost Minimax Optimal Best Arm Identification in Piecewise-Stationary Linear Bandits

Zixin Zhong, Hong Kong University of Science and Technology (Guangzhou), Milan, Italy, Yunlong Hou, Vincent Y. F. Tan

We propose a novel piecewise stationary linear bandit (PSLB) model, where the environment randomly samples a context from an unknown probability distribution at each changepoint, and the quality of an arm is measured by its return averaged over all contexts. The contexts and their distribution, as well as the changepoints are unknown to the agent. We design Piecewise-Stationary ε-Best Arm Identification+ (PSεBAI+), an algorithm that is guaranteed to identify an ε-optimal arm with probability $\geq 1-\delta$ and with a minimal number of samples. PSεBAI+ consists of two subroutines, PSεBAI and Naïve ε-BAI (NεBAI), which are executed in parallel. PSεBAI actively detects changepoints and aligns contexts to facilitate the arm identification process. When PSεBAI and NεBAI are utilized judiciously in parallel, PSεBAI+ is shown to have a finite expected sample complexity. By proving a lower bound, we show the expected sample complexity of PSεBAI+ is optimal up to a logarithmic factor. We compare PSεBAI+ to baseline algorithms using numerical experiments which demonstrate its efficiency. Both our analytical and numerical results corroborate that the efficacy of PSεBAI+ is due to the delicate change detection and context alignment procedures embedded in PSεBAI.

3 - ***4M: Machine Learning Based Marketing Mix Model***

Kitae Kim, Korea Institute of Science and Technology Information, Daejeon, Korea, Republic of, Minhyung Lee, Sung-Hyuk Park

Marketing Mix Models (MMMs) are essential tools for optimizing budget allocation, but traditional approaches struggle with the medium-day-level granularity common in digital marketing. They also have difficulty incorporating modern machine learning (ML) and deep learning (DL) methods. Typically, existing MMMs use weekly aggregated data and rely on gradient-based optimization, limiting their suitability for daily-level insights and compatibility with advanced models. This study introduces the **ML-based Marketing Mix Model (4M),** a novel framework that integrates ML/DL techniques for ad response prediction and employs gradient-free optimization—such as evolutionary strategies—for effective budget allocation. 4M is designed to handle granular datasets and navigate the complexity of multi-channel digital campaigns. Through evaluations using real-world marketing data, 4M consistently outperformed the existing models, including Google and Meta. It demonstrates higher predictive accuracy across various metrics and achieves superior budget allocation outcomes, with up to 6.2% uplift compared to the lowest-performing baseline in certain scenarios. Additionally, 4M operates at a significantly lower computational cost than Bayesian MMMs and maintains robustness as the number of media channels increases. Our findings emphasize the link between predictive accuracy and allocation performance and contribute to the advancement of MMM methodologies. 4M offers practitioners a reliable, scalable, and computationally efficient tool for data-driven marketing optimization and strategic decision-making in increasingly complex digital environments.

TA05

Orchard A

Digital Twin in Logistics

Invited Session

Digital Twin and Metaverse

Chair: Yi Ding, Shanghai Maritime University, Shanghai, N/A

1 - Agentblock.Net: An Online Modular Platform for Agent-Based Modeling and Simulation in OR/MS Complex Systems Research

Zhou He, University of Chinese Academy of Sciences, Beijing, China, People's Republic of

AgentBlock.net is an online platform for operations research (OR) and management science (MS) scholars to address challenges in agent-based modeling and simulation (ABMS) for complex systems. It tackles ambiguities in model description, ad-hoc rule design, and technical barriers through a modular architecture that combines customizable agent modeling, reusable decision modules, and low-code simulation tools.

The platform hosts a knowledge graph-driven library of validated decision modules, integrating theories from game theory, behavioral economics, and optimization. A visual interface enables drag-and-drop model construction, reducing coding effort by 80% while ensuring reproducibility through standardized documentation.

Cloud-native scalability supports large-scale simulations, ideal for OR/MS scenarios like supply chain resilience, policy evaluation, and urban resource allocation.

The collaborative ecosystem of AgentBlock.net encourages global participation: researchers can develop custom modules, share validated models, and leverage cross-disciplinary insights. Prioritizing theoretical rigor and accessibility, the platform empowers scholars to focus on hypothesis testing rather than technical implementation, accelerating progress in modeling complex adaptive systems.

With multilingual support and a focus on OR/MS integration, the platform aims to become a foundational tool for translating complex system theories into actionable simulations, bridging academic research in OR/MS, and practical problem-solving.

2 - A Digital Twin Based Simulation Framework Technology for Container Terminal&Rsquo;S Vessel Operations Simulation and Optimization

Yi Ding, Shanghai Maritime University, Shanghai, China, People's Republic of, Youxuan Pan, Tiantian Zhu, Weihao Wang

Container terminals across the globe are confronted with mounting challenges in the form of intricate operations and mounting demands for intelligent decision-making in the domains of scheduling, resource allocation and logistics management. Conventional methodologies are unable to adequately balance real-time responsiveness with optimal decision quality in dynamic port environments. To address these challenges, a DT-ABS-OPT framework is presented, integrating agent-based simulation (ABS), digital twins (DTs), and optimisation techniques specifically for maritime terminal operations. The proposed approach employs DTs to furnish precise data support and system representation, thereby facilitating real-time operational insights devoid of decision-making capabilities. The ABS then analyses the data and drives the decision-making process through optimisation model selection and parameter tuning. The optimisation models, in turn, provide high-quality solutions under agent guidance, thus overcoming the traditional real-time performance limitations through continuous agent operation.As an illustration of the application of this framework, this research developed a simulation and optimisation system for the terminal's vessel operations parts, which is driven by the framework. The system is integrated for three parts: prediction, decision making and action implementation, including: predicting the operational efficiency of vessel and quay cranes (QC), making decisions on QC scheduling plans and container work sequences, and interacting with the container terminal digital twin system to simulate operation implementation. The system is utilised for the purpose of verifying and predicting the rationality of resource allocation in the port for forthcoming daily, in addition to evaluating the effectiveness of prediction and decision algorithms.

3 - Toward Intelligent Seaport Technological Evolution: From Real Ports to Physical Simulation and Digital Twins

Ang Yang, Dalian Maritime University, Dalian, China, People's Republic of

4 - Design and Development of An O2des-Powered Airport Surface Simulation Platform

Hongjin Zhang, Northwestern Polytechnical University, Xi'an, Shaanxi, China, People's Republic of, Chenhao Zhou

This paper presents the design and development of an O2DES-powered airport surface simulation platform aimed at addressing the growing complexity of airport surface operations. As global air traffic increases and airport infrastructures expand, traditional manual scheduling struggles to cope with dense flight volumes, varied aircraft types, and frequent disruptions such as adverse weather. These challenges significantly impact the safety and efficiency of airport surface operations.
To tackle these issues, we propose a discrete-event simulation platform powered by O2DES, developed within the .NET framework. The platform models three critical airport components—runways, taxiways, and aprons—and simulates core processes including aircraft takeoff, landing, taxiing, and parking. In particular, a structured taxiway network is constructed within the simulation environment, enabling the implementation of path planning algorithms that take into account a wide range of operational factors, including aircraft type, wake turbulence separation requirements, routing constraints. Additionally, conflict-avoidance strategies and optimized apron allocation methods are incorporated to ensure operational safety and efficiency.

The platform supports flexible model customization, visualizes aircraft movement in real-time, and provides integrated tools for data collection and performance analysis. Using operational data from a major Chinese airport, we validate the simulation results against actual performance metrics, confirming the model's accuracy and robustness. Furthermore, three distinct apron assignment schemes are proposed and evaluated. Through comparative analysis of key indicators such as delay time and turnaround cycles, the platform reveals the trade-offs of each strategy, offering data-driven decision support for airport management.

TA06

Orchard B

Frontiers of Digital Finance and Innovation

Invited Session

Finance

Chair: Ying Chen, National University of Singapore, N/A

Co-Chair: Lei Zhou, National University of Singapore, I4 Building, 3 Research Link, 04-03, 117602

1 - Investor Sentiment and Cryptocurrency Returns

yinggang zhou, Xiamen University, Xiamen, China, People's Republic of, Jian Chen, Zeyun Bei

Little is known about how investor sentiment predicts individual cryptocurrency returns. Utilizing coin-level sentiment measures, we show that sentiment predicts subsequent cryptocurrency returns positively in the short run but negatively in the long run. In cryptocurrencies preferred by short-term traders, high sentiment initially boosts prices but eventually crowds out investor participation, suppressing long-term prices. Conversely, in cryptocurrencies favored by long-term holders, high sentiment persistently attracts investor adoption and maintains its positive return predictability. This sentiment-based mechanism primarily explains momentum and reversal anomalies in cryptocurrency markets. Moreover, sentiment mainly arises from extrapolating past returns and user participation.

2 - Hfts and Dealer Banks: Liquidity and Price Discovery in Fx Trading

Shihao Yu, Singapore Management University, Singapore, Singapore, Wenqian Huang, Peter O'Neill, Angelo Ranaldo

This paper examines the complementary roles of dealer banks and high-frequency traders (HFTs) in foreign exchange markets. Using unique data from dealer-to-client (D2C) and dealer-to-dealer (D2D) segments, we find that dealers retain an informational advantage, while HFTs mitigate adverse selection by leveraging speed and technology. Dealers drive price discovery through informed trades, whereas HFTs enhance it via rapid quote adjustments. HFTs provide stable liquidity during market-wide volatility spikes, while dealers sustain liquidity before macroeconomic announcements. Our findings highlight how dealers and HFTs contribute to the quality and resilience of OTC markets.

3 - Innovation Value Discrepancy and Its Role In Shaping Firms' Short-Term Gains and Sustainable Growth

Lei Zhou, National University of Singapore, Singapore, Singapore, Ying Chen, Wenhan Gao, SyBor Wang

This research explores the relationships between technological innovations and firm growth trajectories, examining both immediate and long-term development. By analyzing the divergence between scientific and economic innovation value, we demonstrate how the timing and nature of market recognition play a pivotal role in shaping a company's development and long-lasting success. Our analysis spans 7,506 firms with 2,522,601 granted patents from 1976 to 2022. It reveals that firms with substantial scientific contributions alone often realize considerable long-term growth due to gradual market appreciation of their innovations. However, firms that display strong economic appeal but lack substantial scientific depth typically experience fleeting short-term gains. Those that excel in both scientific and economic dimensions display steady but moderate growth, potentially limited by increased competition. In contrast, firms deficient in both metrics consistently perform

poorly. We further construct trading strategies via both predicted and ex-post data to test the portfolio return within the scientific and economic value discrepancy. These findings offer strategic implications, emphasizing the importance of balancing innovation's scientific and economic aspects to foster sustained growth and long-term competitive advantage, offering insights for executives in intangible asset management.

4 - Neural Tangent Kernel in Implied Volatility Forecasting: a Nonlinear Functional Autoregression Approach

Hannah L. H. Lai, National University of Singapore, Singapore, Singapore

Forecasting implied volatility across different levels of moneyness and maturity is crucial yet challenging due to the high dimensionality of the Implied Volatility Surface (IVS) and the nonlinearity that characterizes its temporal dependence. We adopt a Nonlinear Functional Autoregressive (NFAR) framework to a sequence of IVS and employ neural networks that admit a Neural Tangent Kernel (NTK) parametrization to capture nonlinear interactions between surfaces. We illustrate the theoretical and numerical advantages of the proposed functional NTK (fNTK) estimator and establish a link to functional kernel regression. Our empirical analysis includes over 6 million European call and put options from the S&P 500 Index, covering January 2009 to December 2021. The results confirm the superior forecasting accuracy of the fNTK across different time horizons. When applied to short delta-neutral straddle trading, the fNTK achieves a Sharpe ratio ranging from 1.30 to 1.83 on a weekly to monthly basis, translating to 90% to 675% relative improvement in mean returns compared to forecasts based on functional Random Walk model.

TA07

Orchard C

Systems Analytics for Future Health Systems

Invited Session

Health Systems Modelling with AI

Chair: Sean Lam, Singapore Health Services/ Duke-NUS Medical School

Co-Chair: Ashish Kumar, Singapore Health Services/ Duke-NUS Medical School, Singapore, N/A, Singapore

Co-Chair: Thomas Campion, Weill Cornell Medicine, 575 Lexington Ave, New York, 10022, United States

1 - Coordinated Transfer of Critical Patients and ICU Beds Expansion: a New Epidemic-Logistics Optimization Model

JI KAI, School of Economics and Management, Nanjing University of Science and Technology, NanJing, China, People's Republic of

Enhancing critical care medical services is a key measure to improve emergency response capabilities during major public health crises, to safeguard lives and public health. This study addresses the complex decision-making problem of cross-regional coordinated transfer of critical patients and the dynamic expansion of Intensive Care Unit (ICU) beds during such emergencies. Firstly, a Bayesian framework is embedded within an infectious disease model, employing Markov Chain Monte Carlo (MCMC) sampling to dynamically update the posterior distribution of uncertain parameters. This approach quantifies the uncertainty associated with stochastic demand for critical care due to infectious disease transmission. Based on this, we construct a multi-stage stochastic programming model (MSPM) with time-varying demand scenarios, enabling a dynamic decision-making mechanism for patient transfers and ICU beds expansion across hospitals. We further describe and compare four candidate transfer strategies. Finally, based on problem characteristics, a customized Lagrangian relaxation algorithm solution is designed. The results demonstrate that, compared to three alternative strategies, the cross-regional coordinated transfer strategy achieves the best balance between objective utility and risk. Additionally, the quadratic utility function ensures greater fairness than the exponential utility function. The MSPM outperforms the deterministic model (DM) in both effectiveness and robustness. This study provides a quantitative reference for public health authorities in developing hierarchical medical service plans, optimizing regional patient transfer coordination, and expanding ICU beds under different emergency scenarios.

2 - A Hybrid Rule-Based and LLM Approach for Automated Hospital Operational Dashboard Analysis

Cuong Hoang, National Healthcare Group, Singapore, Singapore, Steven Miller, Kiok Liang Teow, Palvannan Kannapiran, Lynette Ong Wan Kee, Kirk Chuan Wong, Joe Hau Chee Seng, John Abisheganaden

Background:
This study investigates the automation of hospital operational dashboard analysis, addressing limitations of Large Language Models (LLMs) in mathematical computations and their opaque nature, which challenge accurate data interpretation.

Methods:
We propose a two-stage hybrid approach, dividing dashboard analysis into a pre-processing stage using reliable statistical computing libraries and knowledge-based rules for computation and a second stage an LLM for reasoning and summarizing. Utilizing the National Healthcare Group's (NHG) monthly hospital dashboard—featuring key metrics on demand, capacity, and utilization—the numerical computing and rule-based processing in the initial stage identifies essential patterns, including trends, significant changes, and threshold breaches by employing Python to compute relevant statistical measures and generate preliminary insights. The LLM component then processes these observations to deliver deeper insights, using specially engineered prompts with hospital knowledge and Chain-of-Thought reasoning instructions.

Results:
Our results show that preprocessing data through the statistical computation and rule-based component yields consistent and traceable outcomes, capturing many observations noted by senior analysts in official reports. However, to enhance the quality and completeness of the text summary of items of interest in the dashboard, further refinement of rules and LLM prompts is required to identify edge cases and produce more concise comments.

Conclusion:
This exploration advances the field of automated decision support in healthcare operations by integrating traditional basic statistical computations and rule-based methodologies with innovative Generative AI capabilities. Future work will focus on enhancing the rules and prompts and developing specialized collaborative agents for diverse analytical tasks within hospital workflows.

3 - An Automated Health Record Redaction Tool Trained on A Non-Health Domain

Christian Alis, Analytics, Computing, and Complex System laboratory (ACCeSs@AIM), Makati, Philippines

Secondary use of health data is typically done on aggregates due to data privacy concerns associated with the use of patient-level data as these contain sensitive personal information. To address this challenge, automated redaction tools have emerged but are still quite limited due to the complex and diverse structures of health records. Accordingly, this study demonstrates the cross-domain capability of KALASAG (Kaakibat sa LAyuning Siguraduhing ligtas Ang inyong Gintong Impormasyon, Filipino for “Companion for Safekeeping your Valuable Information”), an automated personally identifiable information (PII) identification and redaction system that uniquely integrates state-of-the-art transformer-based language models and rule-based algorithms, in redacting semi-structured health records despite it being developed using data from a non-health domain.
KALASAG, tested on a manually annotated unstructured non-health dataset from a Philippine-based organization, performs with a 99.52% recall score. This was then tested on real-world health records from a Philippine-based business process outsourcing company which caters to home health care providers in the United States. Preliminary results show that KALASAG can redact PII in health records without the need to specify all the fields that contain PII.
Our results show that KALASAG demonstrates cross-domain capability in PII redaction, representing a significant advancement in health data redaction. This offers a solution for safeguarding sensitive information without the need for a custom-built model, and subsequently allowing for the use of patient-level data in building data science and artificial intelligence models that can potentially improve the healthcare landscape.

TA08

Bras Basah A

The Power of Data and Generative AI: Impacts on Work, Consumers, and Society

Invited Session

Informations Systems

Chair: Eric Kwon, Nanyang Technological University, Singapore, Singapore

1 - Harnessing Big Data for Urban Planning: a Social Science Perspective

Zhichao Hu, Nanyang Technological University, Singapore, Singapore, Singapore

In the urbanization process, city clusters create remarkable economic dynamism through positive economic externalities but also face systemic costs like epidemic risks, social instability, pollution, and traffic congestion. As urban systems evolve under development pressures, structural cracks appear. There's an urgent need for a forward-looking monitoring mechanism to detect urban operational issues and trigger early warnings for policy-making and spatial planning. This study aims to develop an urban system that integrates multi-source data and social science theories for detecting and predicting urban problems. It should monitor known unresolved urban issues and identify unknown emerging anomalies, providing comprehensive decision-support for urban governance. This research expands information systems paradigms through three key contributions: 1) establishing an urban anomaly detection theoretical framework; 2) creating a replicable cross-sectoral data integration technical architecture; 3) developing a theory-based and multi-source data-driven urban system application for solving complex social problems. The research offers an early-warning decision-support system for urban governance and broadens information science theory in urban governance contexts.

2 - Fit for the Task: Aligning Dynamic Capability Development with Team Function

Anisha Sharma, Nanyang Business School, NTU, Singapore, Singapore, Dr. John Dong, Dr. Haoyuan Liu

In an era of relentless market disruptions, many organizations struggle to maintain competitiveness and resilience. The pandemic in the past further exposed vulnerabilities and underscored the necessity for agile strategic responses. Dynamic capabilities (DC) sensing, seizing, and transforming have thus become indispensable tools for sustained performance. This short paper explores how aligning DC development with specific team functions (market-facing vs. technology-facing) can boost both resilience and overall performance. While prior research has largely centered on conceptual insights, we offer fresh empirical evidence that validates the merit of targeted DC interventions at the team level, shedding new light on the microfoundations of dynamic capabilities literature.

3 - When Generative Artificial Intelligence Meets Human Reviews: How Does It Change Consumer Purchases?

Shiyi Wang, Nanyang Technological University, Singapore, Singapore, Siliang Tong, John Dong, Xuan Huang

Digital platforms face the growing challenge of helping consumers navigate the overwhelming volume of user-generated reviews. This study investigates the impact of generative artificial intelligence (GAI) summaries on consumer purchase decision-making. Using a field experiment conducted on Trip.com, we assess how a GAI summary feature influences hotel bookings. Our findings reveal that, while the GAI summary feature does not significantly increase consumer demand, as measured by total bookings or spending, it enhances decision efficiency, reflected in higher purchase rates. Further analysis reveals that linguistic characteristics of GAI summaries—such as subjectivity, sentiment, and content quality—play a critical role in shaping consumer decision efficiency. Consumers show a clear preference for summaries that are objective, concise, and easy to read. Moreover, the study highlights the complementary role between GAI summaries and human reviews in facilitating purchase decisions. GAI summaries that provide unique, non-redundant insights are particularly well-received by consumers. These findings offer timely and actionable guidance for digital platforms seeking to leverage GAI to improve user experience in filtering online reviews and to support informed consumer decision-making.

4 - The Generative AI Effect: How Productivity Gains from Generative Ai Reshape Daily Life and Socioeconomic Behavior

KAYOUNG SHIN, Nanyang Technological University, Singapore, Singapore, Hyeokkoo Eric Kwon

This study examines how generative artificial intelligence (GAI) influences individual productivity and generates broader socioeconomic spillover effects. Unlike traditional AI systems deployed within organizational contexts, GAI—exemplified by tools such as ChatGPT—offers general-purpose capabilities that empower creative and skilled individuals. Adoption is increasingly user-driven, bypassing institutional deployment and reshaping labor, income, consumption, and leisure at the individual level.

While previous research has explored GAI's impact on personal productivity, most studies rely on qualitative methods or focus narrowly on task efficiency. This study distinguishes itself by using a novel real-world transactional dataset to conduct a quantitative analysis. It explores not only direct productivity outcomes but also spillover effects across working time, consumption behavior, and leisure participation.

Using transactional and demographic data from a mobile wallet application, we proxy GAI adoption via ChatGPT subscription status. We apply a Difference-in-Differences (DiD) framework combined with Coarsened Exact Matching (CEM) to compare adopters and non-adopters. The analysis covers changes in income, secondary income generation, consumption, and time use, while also examining demographic and occupational heterogeneity.

We further investigate whether productivity gains arise from time savings, improved work quality, or both, and assess whether GAI adoption may exacerbate income inequality by favoring those with higher technological readiness. Clustering techniques are used to identify distinct benefit profiles, offering insights for targeted policy and strategy.

By focusing on bottom-up adoption dynamics, this study contributes theoretical and practical insights into how GAI transforms work, lifestyle, and consumption beyond traditional AI paradigms.

5 - Who Leads the Way? Buyer vs. Supplier Initiatives in Supply Chain Carbon Footprint Reduction

Tarkan Tan, University of Zurich, Zurich, Switzerland, Elif Kuscu, Ozgen Karaer

As firms face mounting regulatory and consumer pressure, supply chain collaboration has become vital for reducing emissions beyond internal capabilities. We study two collaboration models, Buyer-Led and Supplier-Led, in a setting where the buyer bears a carbon price, and the supplier controls most emissions. Using game-theoretic modeling, we show that supplier-led collaboration, where the supplier proposes a cost split and commits to emission reduction, consistently improves both environmental and financial outcomes when clean technologies are accessible. In contrast, buyer-led collaboration, where the buyer offers incentives for emission cuts, becomes essential in markets with limited technology access. Our results guide managers in tailoring collaboration strategies to market conditions.

6 - Eliciting Supplier Cooperation for Value Chain Decarbonization: a Field Experiment with Smallholder Farmers in India

Sukti Ghosh, INSEAD, Fontainebleau, France, Jasjit Singh

Firms are increasingly under pressure to reduce greenhouse gas (GHG) emissions not just within their operations but across their value chains. However, value chain decarbonization is typically not a priority for suppliers, and aligning their goals through formal contracts is often also impractical. This study examines the effectiveness of a firm employing relational strategies as a way to address this dual challenge. In a research collaboration with a Fortune 500 firm seeking to reduce GHG emissions in its agricultural supply chain in India, we designed a field experiment to examine the effectiveness of combining decarbonization training for the firm's supplier farmers with personalized agricultural support to boost the economic value that they derived from their relationship with the firm. Specifically, we examined two interventions that differed in the nature and extent of this personalized support: a lower investment intervention that provided personalized support only for the crop the supplier farmers grew specifically for the firm's value chain, and a higher investment intervention that additionally provided personalized support for broader agricultural practices relevant for the other crops the farmers grew. Relative to a control group that was only exposed to decarbonization training and not given any personalized support, both interventions improved the farmers' adoption of the firm- recommended climate-friendly practices.

The higher investment intervention produced greater environmental impact in terms of emissions reduction per farmer as well as emissions reduction per dollar invested, while also leading to better business outcomes in terms of expected retention of farmers in the sourcing program.

7 - Are Sustainable Supply Chain Policies Beneficial for Firms?

Ali Shantia, HEC Paris, Jouy en Josas, France, Sara Vessal, Felix Papier

Recent regulations, such as the EU Corporate Sustainability Due Diligence Directive (CSDDD), require firms to ensure sustainable practices throughout their supply chains. These mandates compel companies to invest in Supply Chain Due Diligence (SCDD) through audits, supplier development, and compliance monitoring. Yet, many firms hesitate due to high costs and the possibility of competitors benefiting from their investments without contributing—known as free-riding.

We develop a game-theoretic model to study this dynamic in a setting with two buyers and a shared supplier. Each buyer's sustainability effort generates private benefits (e.g., reduced risk, improved supplier performance) and spillover benefits to the other buyer. We introduce a regulatory mechanism combining a minimum standard with penalties for non-compliance and analyze how these instruments influence firm behavior.

Our analysis reveals three key insights. First, without strong penalties, both buyers may underinvest in SCDD due to spillover-induced free-riding. Second, appropriately calibrated penalties can reduce free-riding and lead to higher compliance. Third, under specific conditions, stronger penalties can paradoxically increase firm payoffs—when spillover benefits and improved supply chain performance outweigh the costs of compliance.

These results highlight the need for careful policy design. Regulations must strike a balance: too weak, and they fail to incentivize investment; too strong, and they risk overburdening firms. When calibrated effectively, SCDD regulations can enhance both sustainability outcomes and firm performance. Our findings also suggest that industry-specific factors and the structure of supply networks should inform policy development.

TA09

Bras Basah B

Algorithmic and AI-Driven Marketing Strategies

Contributed Session

Marketing

Chair: Jingman Cao, National University of Singapore, Singapore, Singapore

1 - Popularity Bias in Digital Platforms: the Impact of "Un" Fair Recommendation Algorithms on Content Creation

Hengming Peng, School of Management, Xiamen University, Xiamen, China, People's Republic of, Zhaowei Miao, Yunchuan Liu, Yongquan Lan

Many digital platforms face issues of popularity bias in their recommendation algorithms, where popular content is frequently recommended while less popular is rarely, if at all. In this paper, we develop a game-theoretic model to examine the impact of (un)fair recommendation algorithms on creators' content creation, profits, as well as consumers and social welfare. We consider two algorithmic fairness policies: no fairness restriction and equal opportunity constraint, and we explicitly characterize the equilibrium strategies of creators' content creation effort under both policies. Counterintuitively, our results show that under no fairness restriction, the small creator always benefits from its content creation efforts, while the large creator may actually suffer from its efforts. Moreover, we find that a platform can benefit from implementing algorithmic fairness, as a fair recommendation algorithm encourages content creators to exert creation effort when the cost of content creation effort is low, algorithmic accuracy is high, and the follower base of small creators is relatively high. We refer to this phenomenon as a quality spillover effect. Our study also suggests that algorithmic fairness has a double-edged impact on both consumers and society, potentially yielding either positive or negative outcomes.

Nevertheless, under certain conditions, the quality spillover effect can create a win-win-win scenario for the platform, consumers and society.

2 - The Power of Linear Programming in Sponsored Listings Ranking: Evidence from a Large-Scale Field Experiment

YUTING ZHU, National University of Singapore, Singapore, Singapore, Haihao Lu, Luyang Zhang

Sponsored listing is one of the major revenue sources for many prominent online marketplaces. When consumers visit a marketplace's webpage for a specific item, in addition to that item, the marketplace might also display a ranked listing of sponsored items from various third-party sellers. These sellers are charged an advertisement fee if a user purchases any of the sponsored items from this listing. Determining how to rank these sponsored items for each incoming visit is a crucial challenge for online marketplaces, a problem known as sponsored listings ranking (SLR). The major difficulty of SLR lies in balancing the trade-off between maximizing the overall revenue and recommending high-quality and relevant ranked listings. Due to the latency requirement and historical reasons, many online marketplaces use score-based ranking algorithms for SLR optimization. Alternatively, recent research also discusses obtaining the ranking by solving linear programming (LP). In this paper, we collaborated with a leading online global marketplace and conducted a series of field experiments to compare the performance of the score-based ranking algorithms and the LP-based algorithms. The field experiment lasted for 19 days, which included 329.3 million visits in total. We observed that the LP-based approach improved all major metrics by 1.80% of revenue, 1.55% of purchase, and 1.39% of the gross merchandise value (GMV), compared to an extremely-tuned score-based algorithm previously used in production by the marketplace.

3 - How Many Sub-Segments? Counting Product Variants for Customer Welfare Optimization

Rajesh Sinha, Indian Institute of Management Indore, Indore, India

Market segmentation is the primary marketing strategy exercised to address the problem of customer heterogeneity. Because customers are heterogeneous, to ensure their satisfaction, relatively homogenous segments are crafted out, and thereafter, the selected ones among them are targeted. To further enhance their satisfaction levels, either personalization can be deployed, or sub-segments can be created, which can be targeted with different variants of the offering. Full or partial personalization cannot be deployed in every situation, particularly when the offering involves physical product or physical product parts. Therefore, how many variants should be offered to a segment being targeted, is an interesting question to be addressed (alternatively, how many sub-segments should be crafted out, one variant for each of the sub-segments). This research, instead of analyzing for seller’s profit optimization, focuses on customers’ overall welfare optimization as the objective for determining the optimal number of sub-segments. Using simulation method, the research investigates with different numbers of sub-segments and their influence on the overall customer welfare.

4 - Advertisement Personalization with Multimodal AI

Rajhans Mishra, Indian Institute of Management Indore, Indore, India, Pankaj Mishra

With the rising influence of data driven activities, personalization of advertisement has become an important area in the marketing domain and being categorized under data-based marketing activities. Over a period of time massive augmentation has happened in terms of how and what kind of data is being used for the personalization considering advancements in predictive analytics, generative AI, and the rise of first-party data from sources like social media, e-commerce, and IoT. However, challenges such as the decline of third-party cookies, complex omnichannel environments, strict privacy laws (e.g., GDPR, CCPA), and inconsistent metadata standards limit effective personalization. Traditional ad personalization methods often rely on online behavioral data, neglecting can important content-based personalization and dynamic metadata critical for scalability and contextual relevance.

This research proposes a novel framework leveraging dynamic metadata—AI-generated, adaptive content descriptors rooted in social cognition, multimodal semantic interpretation, and real-time user context. Unlike static metadata, dynamic metadata evolves continuously to reflect shifting consumer perceptions and trends. It incorporates insights from long-term engagement strategies, large language models, and multimodal discourse analysis to support adaptive ad targeting.

The resulting AI-powered framework enables marketers to deliver hyper-personalized, context-aware ad recommendations. It enhances content strategies by integrating descriptive, structural, and interactional metadata, improving real-time decision-making, A/B testing, and adaptability during crises. Offering value to academia, organizations, and advertisers, this interdisciplinary approach advances personalized advertising across cultural, social, and temporal contexts.

5 - When AI Writes Alike: Implications for Job Matching Efficiency

Jingman Cao, National University of Singapore, Singapore, Singapore, YUTING ZHU

Online platforms are introducing various kinds of generative AI tools after the emergence of ChatGPT. Although these tools could improve work productivity, they could also introduce issues of writing content convergence, which could result in increased screening difficulty and lower matching efficiency. In this paper, we study Job Post Generator, a generative AI tool launched by a freelance platform, Upwork. After its introduction, the average pairwise cosine similarity of job postings increased by 18% from 2023 March to 2024 September. We find that clients who use the Job Post Generator match with freelancers more quickly on average. However, those who rely heavily on AI-generated content with minimal modifications experience longer matching times. In contrast, clients who actively edit and customize AI-generated job postings experience significant improvements in matching efficiency, benefit from lower payment costs, and report higher satisfaction with freelancers' performance. Our paper provides evidence that homogenization of job postings due to AI-driven writing can negatively impact online job matching efficiency. Our findings suggest that online platforms offering generative AI tools should encourage users to modify and personalize AI-generated content to achieve better outcomes.

TA10

Hullet Room

Advanced Modeling and Algorithms for Enhancing Resilience in Interdependent Energy Networks

Invited Session

Power Grid

Chair: Yujia Li, lawrence berkeley national laboratory, Berkeley, CA, United States

Co-Chair: Jun Wang, The Univeristy of Hong Kong, Hong Kong, Hong Kong

1 - Risk-informed Resilience Planning of Transmission Systems Against Ice Storms

Chenxi Hu, The University of Hong Kong, Hong kong

2 - Deep Reinforcement Leaning-based Method for Fast Solution of stability enhancement approach in seaport energy logistic systems

Jun Wang, The Univeristy of Hong Kong, Hong Kong, Hong Kong

3 - Multi-Objective Stochastic Planning and Analysis for Multi-Energy Microgrid Considering Unscheduled Islanded Operation

Yan Cao, School of Electrical and Electronic Engineering (EEE) Nanyang Technological University (NTU), Singapore, Singapore

The multi-energy microgrid (MEMG) improves the energy supply economy through a multi-energy coupling operation. However, due to faults or maintenance, outages may occur in the main grid, forcing the MEMG switch to the islanded mode, which heavily impacts the energy supply reliability. This talk introduces multi-objective stochastic planning and analysis for a MEMG considering unscheduled islanded operation. A planning-operation co-optimization model is established with two objective functions: the minimum total planning cost and the minimum expected economic loss caused by islanded operation. The model incorporates equipment capacity allocation constraints, operation constraints in grid-connected and islanded modes, and load curtailment constraints in the islanded mode. In addition to considering the uncertainties of load and renewable energy, the model of the uncertain islanded mode is established using four correlated random variables, and the scenario tree is used for islanded scenario generation. The adaptive weighted sum (AWS) method is used to solve the multi-

objective model and obtain an optimal and evenly distributed Pareto frontier. On this basis, the technique for order preference by similarity to an ideal solution (TOPSIS) is employed to determine the final planning scheme. Numerical results illustrate the effectiveness of the proposed method, which makes a trade-off between planning economy and energy supply reliability.

TA11

Morrison Room

Advancements in Revenue Management and Discrete Choice Models

Invited Session

Revenue Management and Pricing

Chair: Yunlong Wang, Shanghai Jiao Tong University, 1954 Huashan Road, Shanghai, N/A

Co-Chair: Rui Zhang, University of Colorado Boulder, 995 Regent Dr, Boulder, CO, 80309, United States

1 - A Bilevel View for Fluid Stockout-Based Substitution

Jingwei Zhang, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Yuan Guo, Chenhao Wang

This paper revisits the classic joint assortment and inventory planning problem under stockout based substitution. We focus on the case where total demand is a bounded, integer-valued random variable– a setting that better reflects the challenges faced by real-world retailers, yet presents substantial analytical difficulties. Fluid relaxation is commonly employed for such problems, and we provide a performance guarantee that improves previous results in Honhon and Seshadri (2013). However, solving it directly remains NP-hard. To address this, we propose a bilevel reformulation that integrates inventory decisions with a lower-level dynamic assortment problem, where nature seeks to maximize sales. This formulation requires the consistent nested assortment property, for which we provide a sufficient condition. On the other hand, we show that the problem simplifies to a (stochastic) bilevel linear program under several commonly used choice models, including the multinomial logit model, the general attraction model and the Markov chain choice model. We further show that the upper-level objective function is Lipschitz continuous. Additionally, we propose a scalable upper bound inspired by the bilevel structure. We also provide a simple policy for the fluid relaxation with performance guaran tee. Numerical experiments suggest that bilevel-related approaches are effective for tackling dynamic substitution problems. We find that solving the bilevel reformulation directly works well at smaller problem scales, while for larger scales, the derived upper bound and the simple heuristic provide strong approximations.

2 - Assortment Optimization Under the Nested Markov Chain Choice Model

Yicheng Liu, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Zizhuo Wang, Lianmin Zhang, Wanting Yang

This paper proposes a Nested Markov Chain (NMC) choice model to capture the customer’s switching and purchasing behavior among different brands/nests of products, especially in the online retailing setting. Under the NMC model, upon arriving at some brand, the customer (she) chooses among the set of products offered by the brand according to the multinomial logit (MNL) model. The customer will leave if she makes a purchase. Otherwise, if no purchase decision is made, she will transit to other brands (or leave) according to certain transition probabilities. We show that the NMC model is equivalent to the Markov Chain (MC) choice model in terms of model representability. However, the NMC model has a unique advantage in terms of the tractability of the decision problem. For the unconstrained assortment optimization problem, we show that there exists an optimal revenue-ordered assortment in each nest. Based on this property, we can transform the assortment problem into a pricing problem under the MC model with a finite choice of prices. Leveraging an equivalent MDP reformulation of the latter problem, we propose an efficient policy-iteration algorithm to solve the unconstrained problem. We also study the cardinality-constrained assortment optimization problem under the NMC model. We find that the algorithm utilized for the unconstrained case can be extended to the constrained case by collecting all the non-dominated sets for each nest, which can be efficiently identified under cardinality constraints. Several extensions are studied and extensive computational experiments are conducted to test the fitness of the proposed model.

3 - Assortment Optimization for Online Video Games

Yunlong Wang, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Fan You, Rui Zhang, Thomas Vossen

We consider an assortment optimization problem for a class of online video games where the in-game virtual store has a unique structure with two sections: *Featured* and *Just For You (JFY)*. All customers (players) are offered the same Featured section assortment, whereas the JFY section is used for personalized recommendations. Our paper is the first to study assortment optimization for the gaming industry under discrete choice models; it is also the first to devise both exact and approximate solution approaches for the constrained mixture-of-nested-logit model. We model customer choice under a constrained mixture-of-nested-logit model and propose a number of solution methods for the resulting assortment optimization problems, including an exact mixed-integer nonlinear programming (MINLP) formulation and three approximate solution methods: a fully polynomial time approximation scheme (FPTAS), a mixed-integer linear programming (MILP) formulation, and a fast heuristic algorithm. We provide theoretical performance guarantees for the FPTAS, MILP formulation, and heuristic algorithm. These experiments demonstrate that the MINLP can obtain optimal solutions efficiently. Our work provides guidance for how online video game stores can manage assortments for effective revenue maximization. Furthermore, we conduct simulations based on ground-truth instances to demonstrate that considering the nests in the mixture-of-nested-logit model can better capture consumer behavior. Most importantly, the assortments obtained by the proposed solution methods can lead to significant gains in expected revenue. Including the nests leads to over 25% gains on average. The extra gains average over 50% compared to the current practice of our industrial partner.

4 - Opt95: Squeezing 95Th Percentile Billing Using Online Primal-Dual Algorithm

Li zhuo Xie, Shanghai Jiao Tong University, shanghai, China, People's Republic of, Baoxiang WANG

Rising bandwidth costs present a significant challenge for cloud service providers. While previous studies have explored traffic allocation strategies for cloud edge networks, most of them lack robustness in real-time scheduling due to the volatility of traffic data and the complexity of the 95th percentile pricing scheme.
In this work, we present an online algorithm with the primal-dual framework to optimize cloud edge traffic. Our approach introduces an implicit price for free slots. Inspired by the shadow price concept in economics, the algorithm adapts to the fluctuating demand by: i) adjusting the shadow price for free slots and ii) reweighting the importance of historical data. Offline scheduling solutions based on historical data are incorporated into the online process to further enhance performance and robustness. We show that our algorithm incurs a regret at most $O(\sqrt{T})$ under a stochastic input model.

TA12

Moor Room

Food Supply Chain Management II

Invited Session

Revenue Management and Pricing

Chair: Xue-Ming Yuan, National University of Singapore, Singapore, N/A, Singapore

Co-Chair: Dikky Indrawan

1 - Joint Assortment and Inventory Optimization for the Front-End Warehouse

Kexin Wu, National University of Singapore, Singapore, Singapore, Xiaobo Li, Xue-Ming Yuan

This study examines assortment and inventory optimization in a two-warehouse e-commerce system that seeks to balance delivery speed with cost efficiency. The system consists of a front-end warehouse, positioned near customers to support rapid order fulfillment but limited by capacity constraints, and a back-end warehouse, which holds a larger inventory but entails higher fulfillment costs. We model this system as a joint assortment and inventory optimization problem, where initial assortment and inventory decisions are made for the front-end warehouse prior to customer demand realization. Once orders are placed, the e-commerce firm must decide on the quantity of each order type to fulfill from the front-end warehouse, with the remaining quantities fulfilled from the back-end warehouse. Due to the high costs associated with order-splitting, each order must be fulfilled entirely by one location. Our approach uses a two-stage stochastic programming model to minimize total fixed

assortment costs and non-fulfillment costs. We derive the closed-form optimal policies for specific order structures, including distinct single-item sales with non-overlapping bundles, as well as configurations like series and star cover bundles. For more general cases, we develop an approximation algorithm and establish a tight performance guarantee bound. Furthermore, we demonstrate that the joint assortment and inventory optimization problem is NP-hard, even for simplified scenarios with deterministic demand and single-item order types. To address this complexity, we introduce a heuristic algorithm that enables efficient assortment decision-making. Numerical experiment results show that our solutions achieve a good balance between optimality and computational efficiency.

2 - Mapping Value Distribution and Governance Mechanisms in Different Indonesia's Banana Variety

Asaduddin Abdullah, SB-IPB, Bogor, Indonesia

Indonesia, a top global banana producer, exhibits a significant gap between its production volume and its export performance to key markets like Singapore, and faces price competitiveness issues. This research investigates the underlying factors contributing to this situation by focusing specifically on value chain governance and cost structures, which are critical determinants of efficiency and fairness in supply chain operations. Employing a comparative case study approach, the study analyzes the value chains for Cavendish and Mas Kirana bananas destined for modern markets in two distinct Indonesian regions: West Java and Lampung Province. Through detailed mapping of governance structures, identification of key actors and their roles, and a thorough analysis of cost accumulation along the chain, this research contrasts how different governance models influence overall cost structures and competitiveness. The findings aim to elucidate the critical links between specific governance arrangements and operational efficiency as well as equitable value distribution within the chain. These insights are crucial for developing targeted interventions to improve the economic performance and export viability of Indonesia's Cavendish and Mas Kirana banana sectors, contributing to knowledge on enhancing efficiency and fairness in agricultural supply chain management.

3 - Supply Network Modelling and Optimization of Banana Supply Chains in Indonesia

Heti Mulyati, IPB University , Bogor, Indonesia, Dikky Indrawan, Nurul HIdayati, Achmad Fadillah, Asaduddin Abdullah, Alim S.Slamet

The Indonesian banana supply chain faces significant challenges including fragmented structures, logistical inefficiencies, high perishability, and market access limitations, hindering its potential despite Indonesia being a major producer. This study aims to develop a supply network model for the Indonesian banana supply chain, encompassing key stakeholders from smallholder farmers to end consumers, considering the unique characteristics of banana production and distribution in the Indonesian context. Furthermore, it will explore optimization strategies to enhance the efficiency, resilience, and sustainability of these supply networks. The modeling approach will integrate factors such as geographical dispersion, transportation infrastructure, post-harvest handling practices, market demand dynamics, and the impact of perishability. Optimization techniques, including network flow optimization, multi-objective programming, and simulation-based analysis, will be investigated to identify optimal configurations for minimizing costs, reducing waste, improving delivery times, and ensuring product quality. The findings of this research are expected to provide valuable insights for stakeholders, policymakers, and agribusinesses to improve the performance and competitiveness of the Indonesian banana supply chain, contributing to enhanced farmer incomes, reduced food loss, and greater market access. The model and optimization strategies will serve as a framework for future research and practical implementation towards a more efficient and sustainable banana supply chains n Indonesia.

4 - Modeling Price Determinants and Dynamic Pricing Strategies for Indonesia&Rsquo;S Banana Commodities

Linda Karlina Sari, IPB University, Bogor, Indonesia, Zenal Asikin, Noer Azam Achsani, Sendy Watazawwadu Ilmi, Ade Holis

In Indonesia, selected banana products play a significant role in household consumption and local agribusiness, with relatively stable price fluctuations compared to other horticultural commodities. However, despite this stability, a substantial knowledge gap persists in understanding how prices are determined throughout the supply

chain—from farm gate to retail level. Limited transparency in price-setting mechanisms and a lack of systematic analysis contribute to inefficiencies, uneven value distribution, and weak bargaining power for small-scale producers.

This research aims to examine the price-setting practices of selected banana commodities in Indonesia and to propose a dynamic pricing framework that promotes transparency, accountability, and fairness within the supply chain. Drawing on the theoretical foundations of supply chain governance, vertical price transmission, and behavioral pricing theory, the framework incorporates key variables such as production costs, market structure, negotiation styles, demand elasticity, and information asymmetry.

The study adopts a mixed-methods approach. Econometric time-series techniques, will quantify the relationships between farm-gate, wholesale, and retail prices. Complementary qualitative interviews with farmers, traders, and retailers in major banana-producing regions will provide insights into informal practices and negotiation dynamics.

The expected outcome is a clearer understanding of the factors influencing banana pricing and the development of a data-driven dynamic pricing model tailored to Indonesia's banana supply chain. This model aims to support fairer value distribution, enhance producers' market position, and increase overall supply chain efficiency. Ultimately, the findings will help improve the competitiveness of Indonesian bananas in international markets and support broader national export diversification strategies.

5 - Data-Driven Demand Forecasting for the Banana Supply Chain in Southeast Asia

Anggi Mayang Sari, IPB University, Bogor, Indonesia, Toni Bakhtiar

The need for accurate demand forecasting in agricultural supply chains has become increasingly urgent due to the volatility of market demands, climate variability, and the risk of food loss. This research proposes the development of a data-driven, intelligent forecasting model specifically designed for the banana supply chain in Southeast Asia. The study seeks to address inefficiencies in harvest planning and distribution by integrating machine learning techniques into forecasting practices. Using secondary data from FAOStat (2010–2022), including indicators such as production, trade flows, consumption, and losses from countries such as Indonesia, the Philippines, and Vietnam, the research will explore both statistical time-series models (ARIMA, SARIMA) and advanced methods (LSTM neural networks). The proposed model will incorporate key demand-influencing variables such as seasonality, weather patterns, market trends, and consumer behavior through feature engineering. Forecast accuracy will be evaluated using metrics such as RMSE, MAE, and MAPE. This proposal aims to invite constructive feedback on model design, data integration strategies, and the relevance of forecasting tools for decision-making among farmers, distributors, and policymakers. Insights from the conference will be essential for refining the methodology and ensuring the practical impact of the research.

TA13

Bailey

Data-driven Supply Chain Management

Invited Session

Supply Chain and Inventory Management

Chair: Yuanguang Zhong

Co-Chair: Weili Xue, Southeast University, Nanjing

1 - Distributionally Robust Spare Parts Inventory Optimization Considering Substitutability

Debiao Li, Fuzhou University, Fuzhou, China, People's Republic of

The electronics industry struggles with spare parts inventory management due to rapid product updates and intermittent demand. To address this, we propose a distributed robust optimization approach considering demand uncertainty and spare parts interchangeability. A single-stage model (1-DRO) is developed, but its conservativeness leads to a two-stage model (2-DRO). We solve the NP-hard problem using a linear decision rule (LDR) and a column-and-constraint generation algorithm (C&CG). For large-scale problems, the DBSCAN clustering algorithm is applied to classify and solve spare parts efficiently.Results show the 2-DRO model

outperforms the 1-DRO and LDR models, reducing costs by 10.1% and 8.7%, respectively. The C&CG algorithm solves 115 spare parts types in one hour, outperforming Gurobi (55 types). DBSCAN expands solvable parts to 255 types, saving 80% computation time with a 4.5% optimization loss.Experiments indicate that solution time is linear with sample size. For limited spare parts, more historical data improves reliability. In high price fluctuation scenarios, classification strategies remain stable. When holding costs exceed 24% of unit price, inventory or substitution strategies control costs. If substitution penalties are higher than unit prices, increasing purchase quantities is recommended. Higher stockout costs require more inventory to reduce customer loss risk. Flexible inventory strategies based on stockout costs balance cost control and customer satisfaction.

2 - Carbon Emissions Dynamic Simulation and Carbon Peak Path of China's Logistics Industry

Hongyan Liang, Fuzhou university, Fuzhou, China, People's Republic of

The IPCC "bottom-up" method is used to measure the carbon emissions of China's logistics industry from 2006 to 2022, and the GRA-Lasso model is used to identify its key influencing factors, and accordingly, the baseline scenario and four optimization scenarios are constructed. Monte Carlo simulation and PSO-LSSVM model are used to dynamically simulate carbon peaking trajectories under different scenarios from 2023 to 2035. The results show that transportation structure, technological progress, industrial agglomeration and energy structure are the key influencing factors. At the national level, the carbon emissions of China's logistics industry will peak in 2032 with a value of 1.801 billion tons under the baseline scenario, and can reach the peak in advance under the four optimization scenarios, among which the effect is the most ideal under the synergistic optimization scenario, which will peak in 2027 with a value of 1.146 billion tons. At the provincial level, under the baseline scenario, only 10 provinces, including Tianjin, Hebei and Jiangsu, can reach the peak ahead of schedule. Under the four optimization scenarios, three provinces, including Qinghai, Ningxia and Xinjiang, can reach the peak in advance only under the synergistic optimization scenario, Tianjin does not reach the peak as scheduled under the structural optimization scenario, Inner Mongolia and Guizhou do not reach the peak as scheduled under the agglomeration development scenario, and the other provinces can reach the peak before 2030 under all four optimization scenarios. Based on the scenario peaking patterns of each province, region-specific carbon peaking pathways are proposed.

3 - Spare Parts Provisioning Strategy of Warranty Repair Demands for Capital-Intensive Products

Wei Xie, Fuzhou University, Fuzhou, China, People's Republic of

The inventory cost of stocking spare parts is a nonnegligible expenditure of providing after-sales service for the manufacturers making capital-intensive products, such as electric vehicles. Especially, for warranty repair service, it is important to manage the spare stock appropriately to satisfy the warranty claims of customers as well as reduce the associated inventory costs. In this paper, we investigate the spare parts inventory issue related to a critical component for under-warranty units of a product. In particular, under the free-replacement warranty policy, failed component will be replaced by a new one by consuming the spare stock. According to the field claim data, we find that the general trend of warranty claims is nonstationary, which will be affected by the product sales and under-warranty failures. Thus, we first propose a model to forecast the time-varying warranty repair demand by explicitly considering the randomness from two major sources, i.e., product sales and under-warranty failures. Under the assumptions of Poisson sales process and exponential failure distribution, the closed-form expressions of mean and variance of cumulative warranty repair demand over time are obtained. Because the number of warranty claims in each period is a one-time data, the associated distribution information is unavailable. Then, based on the properties of the demand statistics, we derived a worst-case upper bound for the associated inventory cost and formulate a three-phase finite-horizon spare parts inventory model, which can be used to appropriately address the time-varying warranty claims. Finally, numerical experiments are conducted to investigate the impacts of key parameters.

4 - Assortment and Price Optimization Under Brand Spillover Effect

Chenbin LUO, Southeast University, Nanjing, China, People's Republic of

In retail management, the umbrella-brand strategy—wherein a single brand name is applied across multiple product categories, whether complementary, substitutive, or independent—is widely adopted to enhance market share. This approach gives rise to demand spillover effects, whereby consumer preferences for a product are

influenced by their usage or perception of other products under the same brand. In revenue management, such effects are typically modeled through cross-category choice models. However, analyses focusing solely on complementarity or substitutability often fail to capture the demand-side implications of spillover effects in a decision-oriented manner. To address this, we extend the classical MNL model to explicitly incorporate demand spillovers via a two-stage choice framework. Within this structure, we investigate the retailer's optimal assortment and pricing strategies under demand interdependence. Our results demonstrate that accounting for spillover effects yields significantly higher revenues. Furthermore, we show that as the spillover intensity increases, retailers are incentivized to lower prices for products in the initial stage to capture greater market share, thereby generating mutually beneficial outcomes. For the assortment management, our findings reveal that indiscriminately expanding assortment size can be detrimental, as it can cannibalize the market for products with strong spillover effect. Considering that solving the assortment problem is NP-hard, we propose an approximation algorithm with provable performance guarantees.

5 - Blessing OR Curse? Strategic Quality Communication in the Presence of Social Learning

Feiyang Shen, Southeast University, Nanjing, China, People's Republic of, Weili Xue, Yimin Yu, Han Zhu

This paper examines the impact of social learning on a firm’s optimal quality communication strategies---voluntary disclosure versus price signaling---in the presence of social learning. Consumers can be either sophisticated or naive on the basis of their ability to deduce product quality from price signals. We find that social learning can render price signaling effective by reducing the incentive for low-quality firms to mimic high-quality firms. The proportion of sophisticated consumers and quality dissimilarity play key roles in a firm's quality communication strategy through the interaction with social learning. When the proportion of sophisticated consumers is moderate, price signaling emerges; otherwise, firms either disclose quality (if dissimilarity is high) or pool (if low). Contrary to intuition, social learning can reduce market-level firm profit and has ambiguous effects on information transparency. When the quality dissimilarity is high, the transparency of the quality information exhibits a U-shape; otherwise, it exhibits an inverted U-shape in the proportion of sophisticated consumers. Interestingly, when social learning improves information transparency, it adversely affects market-level firm profit, consumer surplus, and social welfare. Thus, firms must be careful in deciding whether to facilitate social learning.

TA14

Wilberforce

Supply Chain Financing and Risk Management

Contributed Session

Supply Chain Management

Chair: Qing Zhu, Nanjing University of Science and Technology, Nanjing, N/A

1 - Balancing Logistics and Liquidity: a Game-Theoretic Analysis of 3PL-Led Financing with Buyback Commitments

Huili Han, School of Management, Xiamen University , Xiamen, China, People's Republic of, Tong Shao, Zhaowei Miao, Yongquan Lan

To address the challenges of high financing costs and difficulties faced by downstream small and medium enterprises (SMEs) in the supply chain, this paper investigates a control-role third-party logistics (3PL) provider in offering financing services and examines the impact of buyback contracts on overall supply chain performance. We consider two supply chain configurations: a short supply chain without 3PL financing involvement and a long supply chain with 3PL participation. Then we develop four decision models, namely, the Short Supply Chain (SS), Short Supply Chain with Buyback Contract (SSB), Long Supply Chain (LS), and Long Supply Chain with Buyback Contract (LSB). Our results indicate that the inclusion of a cost-efficient control-role 3PL can mitigate the double marginal effect inherent in decentralized decision-making, thereby potentially enhancing overall supply chain profitability. In particular, although buyback contracts stimulate higher order quantities and partially offset financing costs, their benefits depend on specific capital cost parameters and logistics expenses. Furthermore, variations in supply chain length are primarily influenced by logistics costs,

while capital costs indirectly affect the choice of supply chain length through their impact on buyback contracts. Notably, when the 3PL provider holds advantages in both logistics and capital costs, both manufacturers and retailers are more inclined to transition from a short to a long supply chain.

2 - Does Smart Contract Add Value to a Retailer with Buyback Risk?

Zhengping Wu, Syracuse University, Syracuse, NY, United States, Ao Zhuo, Wanshan Zhu

Retailers face buyback risk in supply chains when manufacturers fail to honor their buyback commitments. A smart contract can mitigate this risk by reserving a portion of the manufacturer's wholesale revenue in a dedicated account exclusively for buyback purposes. Intuitively, the smart contract benefits the retailer by reducing buyback risk. However, our analysis shows that the smart contract adds value to the retailer only if she is highly risk-averse and the smart contract must be carefully designed. In certain cases, the retailer suffers from the implementation of a smart contract because it allows the manufacturer to charge a higher wholesale price and to benefit consistently. Therefore, a "win-win" situation can be achieved for the supply chain only with a well-designed smart contract and a risk-averse retailer. In addition, a supply chain with a low-margin product is more likely to benefit from the deployment of a smart contract. Furthermore, an attractive investment opportunity can lead to a "lose-lose" situation when the buyback risk arises from investment constraints, as the smart contract prevents the manufacturer from pursuing a more lucrative opportunity. Although smart contracts are reliable technology for streamlining transactions, our research advises supply chain managers to carefully assess their implementation to avoid potential losses. When well designed, smart contracts significantly enhance overall supply chain efficiency.

3 - Regulating Collapse Risk and Dynamic Pricing in Buyer-Backed Supply Chain Finance

Jia Yu, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, Chi To NG

Supply chain finance (SCF) has become an essential tool for easing capital constraints faced by upstream suppliers. However, the inappropriately designed SCF mechanisms may amplify financial risk and undermine supply chain stability. The introduction of risk regulations Basel III underscores the need to design risk-sensitive financing mechanisms that align with regulatory requirements and supply chain resilience. To address these issues, we analyze two financing mechanisms, including the uniform financing without Basel III and the dynamic financing with Basel III, for the buyer-backed supply chain finance (BSCF) scheme. We also investigate how two distinct pricing models, including a high interest rate paired with a low guarantee, and a low interest rate with a high guarantee, for financial service providers (FSPs) under the two financing mechanisms influence the supply chain efficiency. Our analysis establishes conditions where the dynamic scheme outperforms the uniform scheme in terms of supply chain stability and surplus. To prevent collapse of the supply chain system, SCF should be withheld from systems exhibiting high operational risks and low returns. Building on these results, we propose an integrated framework for risk management and financing decisions that incorporates Basel III regulatory requirements and accounts for the risk preferences of both financial service providers and supply chain participants. Our research highlights the importance of risk-differentiated pricing and guarantee mechanisms by providing policymakers a framework to design resilient, cost-effective financing strategies aligned with supplier-specific risks, enhancing supply chain and financial system resilience while increasing supply chain surplus.

4 - The Impact of Supply Chain Structure on Risk Hedging: Procurement Contract Selection Under Monopoly, Retailer Competition, and Manufacturer Competition

Qing Zhu, Nanjing University of Science and Technology, Nanjing, China, People's Republic of, Yudong Wang, Geng Cui, Stephen Disney

We explore how manufacturers choose procurement contracts as a means to mitigate the risks associated with price fluctuations and to maximize profits under different supply chain structures, taking into account factors such as interest rates and product substitution rate. We consider a two-stage supply chain in which contract selection is modeled as a Stackelberg game, with the manufacturer acting as the leader and the retailer as the follower. Our analysis covers three distinct supply chain structures: monopoly, retailer competition, and manufacturer competition. For each competitive setting, we consider two types of competition—Bertrand and Cournot. We examine the contract preferences of both manufacturer(s) and retailer(s), and whether the underlying reasons for these preferences differ. The results show that under a monopoly structure, medium spot

market price volatility affects manufacturer and retailer's contract choice in a similar way. Under retailer competition, spot market price volatility affects manufacturer and retailers' preference similarly in Bertrand and Cournot competition. Under Cournot competition, spot market price volatility influences supply chain profits through the same mechanism as in a monopoly setting. In manufacturer competition, whether under Bertrand or Cournot competition, manufacturers adopt the same contract and achieve the same contract equilibrium. However, whether a retailer shifts its preferred contract depends on the level of market demand uncertainty. The study reveals that increased product substitution rate reduces profits across different market structures, while higher interest rates consistently diminish manufacturers' profits.

TA15

Fairmont Ballroom

Sustainable and Green Logistics

Contributed Session

Transportation Science and Logistics

Chair: Mingda Liu, Tsinghua University, Beijing, N/A

1 - The Location-Routing Research of Hazardous Materials Transportation Considering Risk Equity

Liping Liu, East China University of Science and Technology, Shanghai, China, People's Republic of

Hazardous materials (hazmat) pose significant risks during transportation and storage due to their dangerous characteristics. Accidents can lead to catastrophic consequences, making risk control essential. Effective risk assessment is the foundation of risk control; however, existing studies often overlook the impact of emergency response capabilities on accident outcomes. This study employs the Gaussian plume model to simulate hazmat diffusion following an incident, quantifying initial risk by estimating the affected population. Emergency response capabilities are then integrated into the risk assessment process.

Beyond risk control, risk equity—ensuring fair risk distribution—is crucial. Current research typically addresses risk control and cost separately from governmental and enterprise perspectives. However, hazmat transportation routes can result in uneven spatial risk distribution, leading to perceptions of unfairness among communities in higher-risk areas. This paper introduces the average deviation index to measure road segment risk equity.

Furthermore, the location of hazmat storage facilities influences transportation routes. Jointly considering warehouse locations and route planning can better manage population exposure. This study proposes a bi-level programming model: the upper-level model (government) minimizes a weighted combination of total risk and risk equity; the lower-level model (enterprise) minimizes transportation cost. A weighting parameter facilitates trade-offs between risk control and equity. To address nonlinearity from the equity index, the model is linearized using auxiliary variables. Risk thresholds are also imposed on storage facilities to ensure equitable risk distribution. Case studies demonstrate the model’s effectiveness in balancing governmental and enterprise objectives.

2 - The Impact of Transportation Infrastructure on Inclusive Growth: a Cross-National Analysis of Developed and Developing Countries

Jinchan Liu, Sichuan University, Chengdu, China, People's Republic of, Chuan Chen, Wenting Zhan

Amidst widening development disparities, this study advances empirical understanding of how transport infrastructure subsectors differentially affect inclusive growth across development contexts. Analyzing panel data from 20 countries (2010-2021), we employ fixed-effects regression with Inequality-adjusted HDI as the outcome variable, controlling for GDP per capita. Three key subsectors are examined: railway passenger volume demonstrates robust positive associations with inclusive growth suggesting its unique capacity to enhance equitable opportunity distribution. Air transport shows limited statistical significance, potentially reflecting its premium service orientation. Port infrastructure exhibits context-dependent impacts—significant in developed economies but insignificant in developing countries, highlighting institutional and supply chain maturity thresholds. By distinguishing the heterogeneous effects of transport subsectors across development contexts, this study advances the theoretical understanding of infrastructure’s multidimensional role in inclusive development.

The findings offer empirical insights to guide transport investment strategies and support policy efforts aimed at fostering equitable and sustainable economic outcomes.

3 - Flexible Transshipment for Collaborative Freight: a Two-Stage Optimization Approach for Small-Scale Cross-Docks

Timo Lehmann, Karlsruhe Institute of Technology, Karlsruhe, Germany, Sebastian Koehler, Alexander Ernst

Small and medium-sized freight forwarders in the European Union face persistent inefficiencies due to limited network density and a high share of less-than-truckload (LTL) operations, resulting in substantial empty mileage. While collaborative transport strategies offer significant potential to increase vehicle utilization and reduce costs, traditional transshipment facilities are often inflexible, capital-intensive, and unsuited for dynamic inter-company collaboration.

This paper introduces a novel framework for enabling cost-effective freight collaboration through minimal-scale, bilateral transshipment points, ranging from simple forklift-operated spots to autonomous crane-equipped hubs. We develop a two-stage optimization model that addresses the strategic placement of transshipment facilities and the operational routing of vehicles within a split delivery vehicle routing problem with transshipment (SDVRP-T) context.

The first stage solves a location planning problem with a variable number of facilities using an exact approach. The second stage applies a Simulated Annealing metaheuristic to optimize daily routing under real-world constraints, such as vehicle load levels and reloading efforts. Our model is evaluated using data from an industry partner and compared across multiple scenarios, including direct shipments, manual and autonomous transshipment points, and hybrid setups.

Results show that our approach enables efficient collaboration without the overhead of large-scale cross-docking centers, offering both economic and environmental benefits. This work closes key gaps in the literature by integrating flexible transshipment location planning with practical routing considerations in collaborative transport.

4 - Motorail Loading Problem with Two-Dimensional Irregular Constraints: Formulation, Optimization and Implications

Mingda Liu, Tsinghua University, Beijing, China, People's Republic of

Railway transportation, renowned for its safety, efficiency, and environmental friendliness, is favored by transportation companies for long-haul automobile distribution, necessitating excellence in loading management and operations. However, manual scheduling, particularly for JSQ6 wagons, suffers from inefficiencies and suboptimal wagon utilization, leading to idle space and revenue loss. To address these challenges, we propose a mixed-integer programming model to account for irregular loading constraints, drawn from the operational practices of our industry partner. Our research develops a novel branch-and-price algorithm to produce exact solutions, leveraging variable neighborhood search metaheuristics for initial warm-up and integrating tailored pricing problems for different deck positions. Compared to manual scheduling, our approach substantially improves resource utilization and operational performance, particularly in large-scale instances. The findings also offer valuable managerial insights in terms of objective selection and vehicle composition for better decision making.

Keywords: OR in service industries, motorail loading problem, irregular automobile vehicles and JSQ6 wagons, variable neighborhood search, branch-and-price.

TA16

VIP A

Agri-Production Management and Food Supply Chain

Invited Session

Sustainable and Socially Responsible Operations

Chair: Indira Roy, IIT Hyderabad, A-425 A Block IITH Girls hostel, IIT Hyderabad, Sangareddy, 502285, India

1 - Farmer Panic and Disaster Response Strategies in Contract Farming

Jun Yang, Huazhong University of Science and Technology, Wuhan, China, People's Republic of, Xue Zhou, Quan Zhou, Hongchen Duan

Contract farming plays a critical role in modern agricultural value chains, but its sustainability is increasingly threatened by agricultural disasters. Such events cause uncertain yield losses and may lead to farmer withdrawal, a phenomenon that can be intensified by farmers' panic. To address disaster risks, companies typically adopt either an information strategy to reduce farmers' uncertainty or a compensation strategy to mitigate their losses. This study develops a stylized model to examine the company's optimal decisions and the relative effectiveness of these strategies. We model farmers' panic as arising from both individual loss concern and social learning effect, and analyze how panic influences strategic decisions and value creation for both the company and the farmers. Our analysis yields several key findings. First, the optimal levels of information and compensation depend jointly on disaster uncertainty and mean yield loss, with farmers' panic playing a significant role. Second, the company's information strategy may inadvertently harm farmers due to the panic amplification effect of information. Third, strategic alignment between the company and farmers arises only under a compensation strategy when disaster uncertainty is either high or low. Finally, while a hybrid strategy can enhance the company's profit, it also entails higher investment costs. We suggest that governments offer loan programs to support companies facing budget constraints when disaster uncertainty is moderate or very low, thereby improving outcomes for both parties. These findings contribute to a theoretical understanding of disaster response in contract farming and offer insights for developing resilient agricultural supply chains.

2 - The Economics of Agricultural Carbon Credits: Implications for Subsidies, Farmers&Rsquo; Income, and Consumer Welfare

Lalbabu Nishad, Indian Institute of Management Udaipur, Udaipur, India, Govind Kumawat

Governments have traditionally relied on agrochemical subsidies to meet crop production targets. However, subsidy-driven conventional farming alone fails to simultaneously achieve the dual objectives of production targets and greenhouse gas (GHG) emission reductions. Recent government policies, such as PM-PRANAM and Farm to Fork Strategy, aim to reduce agrochemical subsidies and promote carbon farming practices to curtail GHG emissions while ensuring sufficient production. Using a Stackelberg game framework, we examine the strategic interactions between agrochemical subsidies and carbon credit markets to optimize the government's budget allocation and farmers' profit while balancing the dual objectives. Our findings reveal that without carbon credit markets, subsidies increase planting and profits for low-yield farmers but reduce both for high-yield farmers. However, when a carbon credit market is available, subsidies primarily benefit low-yield farmers by increasing their planting and profits, while high-yield farmers' planting and profits depend on price sensitivity. Importantly, high-yield farmers transition to carbon farming. We also identify a productivity threshold above which farmers adopt carbon farming, which brings three key benefits: sufficient production to meet targets, reduced GHG emissions and improved farmer incomes. Additionally, the widespread adoption of carbon farming reduces per-unit subsidies and improves budget stability. In markets with very low average crop productivity, carbon credits can serve as a viable mechanism to enhance farmers' overall profitability. The government can also achieve the same levels of consumer welfare, farmer profitability, and social welfare at a reduced budget by introducing carbon credit markets, though at the cost of increasing income inequality.

3 - Responsible Omnichannel Fresh Food Distribution: Retail Operations, Food Redistribution, Sustainability

Indira Roy, Indian Institute of Technology Hyderabad, Sangareddy, India, Lohithaksha Maniraj Maiyar

As the fresh food retail landscape transitions from multi-channel to omnichannel retailing, retail operations have become increasingly complex. The complexities arise from huge-fragmented distribution networks, inefficient utilization of resources to curb fresh food wastages and lack of awareness towards the hidden cost of fresh food losses. With sustainability as the rising priority of today, investigation of responsible operations is critical to

meet consumer satisfaction levels, addressing environmental and societal concerns while maintaining sufficient profit margin. Hence, this study develops a multi-objective optimization model to maximize the retail profit margin, customer satisfaction, and alignment towards government sustainable development goals (SDGs). The first objective maximizes the retailer profit margin influenced by costs associated with procurement, logistics, freshness-keeping efforts, fresh food donation and fresh food loss. Customer satisfaction is aimed at maximizing the economic and psychological utility of loss-averse consumers who make their purchase decisions influenced by sensitivity to price, freshness levels, discounts and sustainability efforts of retailers. The alignment towards SDGs is measured in terms of quantities of fresh food donations, maintaining fair trade prices and carbon emission reduction efforts. These aims are realized by studying the interplay between key decision variables of the optimal selling price, procurement price, freshness-keeping effort level, carbon emission reduction effort level, discounted quantities for the secondary market and donation quantities. The model is solved and validated on real data sets from the Indian fresh food scenario using evolutionary solution methodologies.

Tuesday, July 22, 10:00 AM - 11:00 AM

Fairmont Ballroom

Transforming Mobility: SMRT's Journey in Data and Digital Innovation

Invited Session

Plenary

1 - Transforming Mobility: SMRT's Journey in Data and Digital Innovation

Hoon Ping Ngien, SMRT Corporation Ltd, Singapore, Singapore

In today's fast-evolving landscape, digital transformation is not just an option – it is a strategic imperative. SMRT is driving digital transformation by leveraging analytics, artificial intelligence, management science and digital technologies to address complex operational and business challenges.

Drawing on practical examples from SMRT and beyond, this session will showcase how data-driven decision-making, predictive analytics and digital solutions are improving operations effectiveness, safety, and customer-centricity in the fields of public transportation, mobility, and supply chains. Building a future-ready organization starts with putting data and digital at the heart of the business.

Tuesday, July 22, 11:15 AM - 12:45 PM

TB01

Enterprise

Applied Probability and Modern Machine Learning: Applications to LLM and Beyond

Invited Session

Applied Probability Society

Chair: Wang Chi Cheung, National University of Singapore, Singapore, Singapore

1 - Roste: An Efficient Quantization-Aware Supervised Fine-Tuning Approach for Large Language Models

Chung-Yiu Yau, The Chinese University of Hong Kong, Hong Kong, Hong Kong, Hoi To Wai

Supervised fine-tuning is a standard method for adapting
pre-trained large language models (LLMs) to downstream tasks.
Quantization has been recently studied as a post-training technique for
efficient LLM deployment. To obtain quantized fine-tuned LLMs,
conventional pipelines would first fine-tune the pre-trained models,
followed by post-training quantization. This often yields suboptimal

performance as it fails to leverage the synergy between fine-tuning and quantization. To effectively realize low-bit quantization of weights, activations, and KV caches in LLMs, we propose an algorithm named Rotated Straight-Through-Estimator (RoSTE), which combines quantization-aware supervised fine-tuning (QA-SFT) with an adaptive rotation strategy that identifies an effective rotation configuration to reduce activation outliers. We provide theoretical insights on RoSTE by analyzing its prediction error when applied to an overparameterized least square quantized training problem. Our findings reveal that the prediction error is directly proportional to the quantization error of the converged weights, which can be effectively managed through an optimized rotation configuration. Experiments on Pythia, Qwen and Llama models of different sizes demonstrate the effectiveness of RoSTE. Compared to existing post-SFT quantization baselines, our method consistently achieves superior performances across various tasks and different LLM architectures.

2 - One-Pass Bandit Learning Under Nonlinear Reward and Heavy-Tailed Noise

Peng Zhao, Nanjing University, Nanjing, China, People's Republic of

Stochastic linear bandits (SLB) is a powerful and fundamental model for interactive online learning and decision making. To better capture real-world complexities, important extensions of SLB have been developed; two notable examples are generalized linear bandits (GLB), which incorporate nonlinear link functions, and heavy-tailed linear bandits (HvLB), which allow for potentially heavy-tailed noise. While optimal regret bounds have been established for GLB and HvLB, the computational efficiency remains unsatisfactory—these optimal algorithms require storing all historical data and performing a full pass over them at each round, which is evidently not feasible for online learning scenarios. To address the issue, we have developed “one-pass” algorithms based on the Online Mirror Descent framework, a textbook-standard approach for regret minimization in online optimization. We will demonstrate its effectiveness in statistical parameter estimation, which further enables the design of UCB-based algorithms with O(1) computational cost per round, while still maintaining optimal regret guarantees. The results can be further extended to decision-making scenarios such as online RLHF, MDPs with function approximation, and adaptive linear control.

3 - Stochastic Optimization Schemes for the Performative Prediction Problem with Nonconvex Loss

Qiang Li, The Chinese University of Hong Kong, Hong Kong, Hong Kong, Hoi To Wai

The Performative Prediction problem describes a scenario in which a trained model influences the outcome it predicts, creating a feedback loop that can impact the stability of optimization algorithms like stochastic gradient descent (SGD). While this phenomenon is well-studied in policy-making, it has been overlooked in supervised learning. In this work, we present the first study of performative prediction with smooth but potentially non-convex loss functions. Note that in the literature, SGD with different deployment schemes is often studied with strongly convex losses. We first propose the definition of a stationary performative stable (SPS) solution by relaxing the popular performative stability condition. We then demonstrate that SGD with a greedy deployment scheme converges to a biased SPS solution in expectation, where the bias is proportional to the degree of distribution shift. Additionally, we extend our analysis to a lazy deployment scheme under non-convex settings, showing that it can converge with reduced bias. This is joint work with Prof. Hoi-To Wai, and has been published in NeurIPS 2024 (Arxiv Link: https://arxiv.org/abs/2405.17922).

4 - Banditspec: Adaptive Speculative Decoding via Bandit Algorithms

Yunlong Hou, National University of Singapore, Singapore, Singapore

Speculative decoding has emerged as a popular method to accelerate the inference of Large Language Models (LLMs) while retaining their superior text generation performance. Previous methods either adopt a fixed speculative decoding configuration regardless of the prefix tokens, or train draft models in an offline or online manner to align them with the context. This paper proposes a *training-free* online learning framework to

adaptively choose the configuration of the hyperparameters for speculative decoding as text is being generated. We first formulate this hyperparameter selection problem as a Multi-Armed Bandit problem and provide a general speculative decoding framework BanditSpec. Furthermore, two bandit-based hyperparameter selection algorithms, UCBSpec and EXP3Spec, are designed and analyzed in terms of a novel quantity, *the stopping time regret*. We upper bound this regret under both stochastic and adversarial reward settings. By deriving an information-theoretic impossibility result, it is shown that the regret performance of UCBSpec is optimal up to universal constants. Finally, extensive empirical experiments with LLaMA3 and Qwen2 demonstrate that our algorithms are effective compared to existing methods, and the throughput is close to the oracle best hyperparameter in simulated real-life LLM serving scenarios with diverse input prompts.

TB02

Mercury

Aviation Topics I

Invited Session

Aviation Applications

Chair: Lishuai Li, City University of Hong Kong, Kowloon, Hong Kong

1 - A Scalable Utm Simulation: Assessing Operational Concepts for Large-Scale Uav Management

Shuangxia BAI, City University of Hong Kong, Kowloon, Hong Kong, China, People's Republic of, Xinyu He, Lishuai Li

The integration of Unmanned Aerial Vehicles (UAVs) into urban airspaces for applications such as delivery, surveillance, and emergency response offers transformative potential across logistics, mobility, and public service domains.

Despite this promise, large-scale UAV deployment demands a robust Unmanned Aircraft System Traffic Management (UTM) system. However, current Concepts of Operations (ConOps) and airspace designs remain largely exploratory, with a lack of scalable simulation tools hindering systematic evaluation and comparison.

To address this gap, we introduce an extensible UTM simulation designed for high-density UAV traffic in complex urban environments. This simulation captures the unique characteristics of low-altitude UAV operations, including interactions with urban infrastructures (e.g., buildings, dynamic airborne obstacles) and localized atmospheric disturbances (e.g., turbulence, wind shear). It incorporates essential modules such as environmental parsing, Communication–Navigation–Surveillance (CNS) emulation, and UAV flight controllers. The platform supports both 3D and 4D trajectory execution and accommodates heterogeneous UAV dynamics, enabling realistic simulations of up to 1000 simultaneous flights.

To ensure consistent and meaningful evaluation, we conduct a performance assessment encompassing four critical aspects: safety (e.g., conflict rates), business performance (e.g., delivery success), operational efficiency (e.g., path optimality), and UTM system performance (e.g., computational overhead). This structure facilitates fair benchmarking across diverse ConOps and provides actionable insights for system designers and policymakers.

Overall, our work delivers a foundational tool for accelerating the design, validation, and deployment of scalable UTM systems—advancing the safe integration of UAVs into future urban airspaces.

2 - Airport Accessibility Assessment and Shuttle Bus Scheduling Across Socioeconomic Groups

Tinghe Zhang , Hong Kong Polytechnic University, Kowloon, Hong Kong, Ang Li

This study presents an equity-integrated optimization framework for airport shuttle bus scheduling that aims to improve both transit accessibility and socioeconomic fairness. Building on heterogeneous travel behaviors and constraints across income groups, we propose the Equitable Airport Accessibility Metric (EAAM), a utility-based measure that captures spatial, temporal, and behavioral components of access. This metric underpins a

Two-Stage Airport Bus Scheduling and Allocation Problem, formulated from both government-led and transit operator-led perspectives. In the government-led model, Stage 1 maximizes accessibility across income groups, followed by cost minimization with equity constraints in Stage 2. Conversely, the transit operator-led model reverses this order, focusing first on cost-efficiency and then improving accessibility under budgetary limits.

Using a generalizable numerical experiment combining real-world data with a synthetic setup, our experiment analysis reveals a complex trade-off between cost, accessibility, and equity. Sensitivity analysis of the accessibility weighting parameter (μ) identifies an optimal balance at (government-led) and (operator-led), where accessibility gains are maximized with minimal cost deviation. Further experiments adjusting income group weights (λ) show that prioritizing low-income groups does not necessarily improve total accessibility due to behavioral constraints, while high-income prioritization improves efficiency but deepens equity gaps. Targeted fare policies—such as late-night zero-fare buses and income-based subsidies—demonstrate meaningful accessibility gains without requiring substantial operational changes. These findings provide policy-relevant insights for designing socially inclusive, cost-aware airport transit systems.

3 - An Interactive Decision Making Framework Design for the Outsourcing Cooperation Between the Service Provider and the Airline: An Exact Bilevel Method

Xuting Sun, Shanghai University, Shanghai, China, People's Republic of

In reality, more and more airlines outsource maintenance tasks to the MRO service providers to achieve cost-effective operational planning. However, in actual operations, the MSP may not be able to provide “mutually beneficial” maintenance services to its customers due to its resource scarcity and the lack of information sharing. Therefore, in this paper, based on industry practices we model an aviation service supply chain consisting of one MSP and one airline, in which the MSP may provide maintenance services to multiple airlines with limited spatiotemporal resources with different prices. A bilevel mixed integer programming (MIP) model is formulated to characterize the interactive decision-making structure between the MSP and the airline that aims to help increase the operational efficiency of both the MSP and the airline. An exact bilevel solution framework is developed which is verified to achieve convergence and obtain the bilevel optimality. Through the proposed algorithm, the computational study is further conducted based on the actual operational data from one of the biggest Chinese airline companies. The results verify that, compared to the benchmark, the proposed bilevel model can achieve Pareto improvement. In addition, the sensitivity analyses regarding pricing, flexible maintenance resources and opportunity cost coefficients verify the effectiveness and robustness of the proposed bilevel model, and further derive some useful managerial insights for both the MSP and the airline at strategic and operational levels.

4 - From Pre-Planned Air Corridors to Dynamic 4D Trajectories: An Comparative Study of Utm Conops

Lishuai Li, City University of Hong Kong, Kowloon, Hong Kong, Shuangxia BAI, Xinyu He

Drone delivery and urban air mobility systems offer advantages over ground logistics, including improved accessibility, agile supply chain support, and reduced traffic congestion. Scaling these operations requires efficient Unmanned Aircraft System Traffic Management (UTM). Existing UTM Concepts of Operations (ConOps) propose diverse airspace structures, ranging from predefined corridors to unrestricted free-flight zones, yet systematic comparisons of their safety and scalability remain limited.

This study evaluates three ConOps with varying airspace organization: **ConOps A** (non-overlapping predefined corridors), **ConOps B** (intersecting predefined corridors), and **ConOps C** (free-flight with dynamic routing). A simulation framework assesses their performance across safety, operational cost, scalability, and flexibility metrics under diverse urban environments (e.g., high-rise vs. low-rise zones) and delivery network topologies (hub-to-spoke vs. hub-to-hub).

Results demonstrate inherent trade-offs: ConOps A achieves optimal safety (zero corridor conflicts) but suffers from low efficiency and adaptability due to rigid routing. ConOps C maximizes operational flexibility and throughput but incurs high computational demands and escalating collision risks (15–22% higher than ConOps A) as traffic density increases. ConOps B balances moderate safety (manageable intersection conflicts) and efficiency but struggles with scalability in complex airspaces.

This analysis provides empirical insights for UTM policymakers and operators to select context-appropriate airspace architectures, emphasizing the need for hybrid approaches that dynamically adapt to traffic demand and urban geographies.

TB03

Minto

Numerical Algorithms and Applications of Optimal Transport

Invited Session

Computing Society

Chair: Zeyi Chen, INSEAD, Singapore, Singapore

Co-Chair: Zeyi Chen, INSEAD, Singapore, Singapore

1 - A Nonparametric Approach with Marginals for Modeling Consumer Choice

Yanqiu Ruan, Singapore University of Technology and Design, Singapore, Singapore, Xiaobo Li, Karthyek Murthy, Karthik Natarajan

Given data on the choices made by consumers for different offer sets, a key challenge is to develop parsimonious models that describe and predict consumer choice behavior while being amenable to prescriptive tasks such as pricing and assortment optimization. The marginal distribution model (MDM) is one such model, which requires only the specification of marginal distributions of the random utilities. This paper aims to establish necessary and sufficient conditions for given choice data to be consistent with the MDM hypothesis, inspired by the usefulness of similar characterizations for the random utility model (RUM). This endeavor leads to an exact characterization of the set of choice probabilities that the MDM can represent. Verifying the consistency of choice data with this characterization is equivalent to solving a polynomial-sized linear program. Since the analogous verification task for RUM is computationally intractable and neither of these models subsumes the other, MDM is helpful in striking a balance between tractability and representational power. The characterization is then used with robust optimization for making data-driven sales and revenue predictions for new unseen assortments. When the choice data lacks consistency with the MDM hypothesis, finding the best-fitting MDM choice probabilities reduces to solving a mixed integer convex program. Numerical results using real world data and synthetic data demonstrate that MDM exhibits competitive representational power and prediction performance compared to RUM and parametric models while being significantly faster in computation than RUM.

2 - Feasible Approximation of Matching Equilibria for Large-Scale Matching for Teams Problems

Qikun Xiang, Nanyang Technological University, Singapore, Singapore, Ariel Neufeld

We propose a numerical algorithm for computing approximately optimal solutions of the matching for teams problem. Our algorithm is efficient for problems involving large number of agent categories and allows for non-discrete agent type measures. Specifically, we parametrize the so-called transfer functions and develop a parametric formulation, which we tackle to produce feasible and approximately optimal primal and dual solutions. These solutions yield upper and lower bounds for the optimal value, and the difference between these bounds provides a sub-optimality estimate of the computed solutions. Moreover, we are able to control the sub-optimality to be arbitrarily close to 0. We subsequently prove that the approximate primal and dual solutions converge when the sub-optimality goes to 0 and their limits constitute a true matching equilibrium. Thus, the outputs of our algorithm are regarded as an approximate matching equilibrium. We also analyze the theoretical computational complexity of our approach. In the numerical experiments, we study three matching for teams problems: a business location distribution problem, the Wasserstein barycenter problem, and a large-scale problem involving 100 agent categories. We showcase that the proposed algorithm can produce high-quality approximate matching equilibria, provide quantitative insights about the optimal city structure in the business location distribution problem, and that the sub-optimality estimates computed by our algorithm are much less conservative than theoretical estimates.

3 - Provably Convergent Stochastic Fixed-Point Algorithm for Free-Support Wasserstein Barycenter of Continuous Non-Parametric Measures

Zeyi Chen, INSEAD, Singapore, Singapore, Ariel Neufeld, Qikun Xiang

We propose a provably convergent algorithm for approximating the 2-Wasserstein barycenter of continuous non-parametric probability measures. Our algorithm is inspired by the fixed-point iterative scheme of {\'A}lvarez-Esteban, Del Barrio, Cuesta-Albertos, and Matr{\'a}n (2016) whose convergence to the 2-Wasserstein barycenter relies on obtaining exact optimal transport (OT) maps. However, OT maps are typically only computed approximately in practice, and exact computation of OT maps between continuous probability measures is only tractable for certain restrictive parametric families. To circumvent the intractable computation of exact OT maps between general non-parametric measures, we develop a tailored iterative scheme that utilizes consistent estimators of the OT maps in place of the exact OT maps. This gives rise to a computationally tractable stochastic fixed-point algorithm which is provably convergent to the 2-Wasserstein barycenter. Our algorithm remarkably does not restrict the support of the 2-Wasserstein barycenter to be any fixed finite set and can be implemented in a distributed computing environment, which makes it suitable for large-scale data aggregation problems. In our numerical experiments, we propose a method of generating non-trivial instances of 2-Wasserstein barycenter problems with known ground-truth barycenter measure. The results showcase the capability of our algorithm in developing high-quality approximations of the 2-Wasserstein barycenter in both low and high dimensional settings, and its superiority over the benchmark generative methods in terms of accuracy, stability and efficiency.

4 - Semidefinite Relaxations of the Gromov-Wasserstein Distance

Binh Nguyen, National University of Singapore, Singapore, Singapore

The Gromov-Wasserstein (GW) distance is an extension of the optimal transport problem that allows one to match objects between incomparable spaces. At its core, the GW distance is specified as the solution of a non-convex quadratic program and is not known to be tractable to solve. In particular, existing solvers for the GW distance are only able to find locally optimal solutions. In this work, we propose a semi-definite programming (SDP) relaxation of the GW distance. The relaxation can be viewed as the Lagrangian dual of the GW distance augmented with constraints that relate to the linear and quadratic terms of transportation plans. In particular, our relaxation provides a tractable (polynomial-time) algorithm to compute globally optimal transportation plans (in some instances) together with an accompanying proof of global optimality. Our numerical experiments suggest that the proposed relaxation is strong in that it frequently computes the globally optimal solution.

TB04

Indiana

AI-driven Modeling of Complex Diseases and Drug Discovery

Invited Session

Computing Society

Chair: Qingpeng Zhang, The University of Hong Kong, Hong Kong, Hong Kong

1 - Optimized Patient-Specific Immune Checkpoint Inhibitor Therapies for Cancer Treatment Based on Tumor Immune Microenvironment Modeling

Qingpeng Zhang, The University of Hong Kong, Hong Kong, Hong Kong

Enhancing patient response to immune checkpoint inhibitors (ICIs) is crucial in cancer immunotherapy. We aim to create a data-driven mathematical model of the tumor immune microenvironment (TIME) and utilize deep reinforcement learning (DRL) to optimize patient-specific ICI therapy combined with chemotherapy (ICC). Using patients' genomic and transcriptomic data, we develop an ordinary differential equations (ODEs)–based TIME dynamic evolutionary model to characterize interactions among chemotherapy, ICIs, immune cells, and tumor cells. A DRL agent is trained to determine the personalized optimal ICC therapy. Numerical experiments with real-world data demonstrate that the proposed TIME model can predict ICI therapy response. The DRL-derived personalized ICC therapy outperforms predefined fixed schedules. For tumors with extremely low CD8

+ T cell infiltration (‘extremely cold tumors’), the DRL agent recommends high-dosage chemotherapy alone. For tumors with higher CD8 + T cell infiltration (‘cold’ and ‘hot tumors’), an appropriate chemotherapy dosage induces CD8 + T cell proliferation, enhancing ICI therapy outcomes. Specifically, for ‘hot tumors’, chemotherapy and ICI are administered simultaneously, while for ‘cold tumors’, a mid-dosage of chemotherapy makes the TIME ‘hotter’ before ICI administration. However, in several ‘cold tumors’ with rapid resistant tumor cell growth, ICC eventually fails. This study highlights the potential of utilizing real-world clinical data and DRL algorithm to develop personalized optimal ICC by understanding the complex biological dynamics of a patient’s TIME. Our ODE–based TIME dynamic evolutionary model offers a theoretical framework for determining the best use of ICI, and the proposed DRL agent may guide personalized ICC schedules.

2 - Optimization of Clinical Decision-Making for Vasopressors in Shock Treatment

Na Geng, Shanghai Jiao Tong University, Shanghai, China, People's Republic of

Vasopressor treatment strategies are essential in treating shock patients, yet determining optimal type, dosage, and timing of vasopressors remains challenging given variable clinician expertise and patient conditions. Applying existing reinforcement learning (RL) algorithms in treatment decision making risks Q-value overestimation, and ignores the gap between AI-driven recommendations and established clinical practices. This work introduces an offline RL algorithm called Safe Conservative Q-learning (SafeCQL), integrating conservative regularization and clinician-informed safety constraints. Patient trajectories from the MIMIC-IV database are modeled as a Markov Decision Process (MDP) to optimize vasopressor administration. Off-policy evaluations with model-based and model-free estimation methods demonstrate the superior performance of SafeCQL over clinician policies and existing RL models in improving survival rates and policy robustness. Findings show that SafeCQL improves the survival rate by 5% and raises expected returns from 45.78 to 69.15 (model-based) and 44.33 to 61.24 (model-free). SafeCQL can derive a treatment policy that aligns closely with clinician preferences while surpassing clinician expertise in decision quality. This work offering a deployable solution for personalized vasopressor management in dynamic critical care environments.

3 - Ai-Driven Influenza Surveillance, Early Warning, and Intervention in Macau's Primary and Secondary Schools

Fengshi Jing, City University of Macau, Macao SAR, China, People's Republic of, Hao Ren, Wenjian Liu, Weibin Cheng, Qingpeng Zhang

This study will present the current status of the influenza surveillance system in Macau's primary and secondary schools, share a machine learning-based influenza early warning approach, and introduce the research design of an AI-powered intervention system featuring computer vision-based smart hand hygiene monitoring.

4 - AI-driven Drug Discovery and Clinical Research for Rare Diseases

Xiangyun Mao, Tsinghua University, Beijing, China, People's Republic of, Guanqiao Li

tbd

TB05

Orchard A

From LLM to Digital Humans: Reshaping Corporate Communication, Disclosure, and Branding

Invited Session

Digital Twin and Metaverse

Chair: Mijia Gong, Northwestern polytechnical university, Xi'an, N/A

1 - Llms and Corporate AI Initiatives

FUJING XUE, Business School, Sun Yat-sen University, Shenzhen, China, People's Republic of, Nan Hu

The rise of large language models (LLMs) has accelerated firm engagement with AI technologies, prompting strategic shifts in AI-related initiatives. This study develops an LLM-based textual analysis framework to identify and evaluate firm-level AI initiatives, including strategic intent, application practices, and governance, based on corporate disclosures and executive communications. We examine how these initiatives evolve over time and explore the internal and external factors that shape firms’ responses to LLM-driven technological

change. Furthermore, we assess the implications of shifting AI initiatives for firm performance and capital market outcomes. Our findings offer novel insights into how organizations adapt strategically to generative AI and contribute to the literature on technology adoption, strategic disclosure, and organizational adaptation.

2 - Ceo Characteristics and Corporate Disclosure Decision on Data Asset Capitalization: Evidence from LLM-Based Experiments

Ting ZHANG, Nanjing University of Information Science & Technology, Nanjing, China, People's Republic of

In recent years, data resources have played an increasingly vital role in corporate operations and value creation, demonstrating inherent asset-like attributes. However, challenges such as property rights confirmation and valuation complexities have hindered the recognition of data resources as "genuine" accounting assets. Despite these barriers, China has pioneered regulatory frameworks for data asset capitalization, allowing enterprises to disclose data resources as assets in financial reports starting in 2024. This study investigates how CEO characteristics influence corporate disclosure decisions on data asset capitalization. We implemented a questionnaire-based experimental design leveraging large language models (LLMs). Results reveal that firms are more likely to capitalize data assets when CEOs are male, younger, possess professional accounting backgrounds, or have overseas experience. This research contributes to upper echelons theory by identifying executive-level determinants of corporate disclosure strategies, while also provides methodological insights into the application of LLMs in management research.

3 - Knowledge Graph-Enhanced Reciprocal Matching for Influencer-Brand Collaborations

Ruibin Geng, Northwestern Polytechnical University, Xi'an, China, People's Republic of, Jing Wang

Influencer marketing has emerged as a transformative strategy in digital advertising, yet brands face significant challenges in identifying suitable influencers who align with their target audiences and marketing goals. Existing approaches predominantly adopt a brand-centric perspective, overlooking the reciprocal nature of influencer-brand collaborations and failing to capture the multi-dimensional relationships among stakeholders. This study introduces a novel framework that integrates knowledge graph embeddings, Gale-Shapley matching theory, and reinforcement learning to enhance influencer-brand collaboration. Our approach addresses three critical limitations in current methods: (1) it enhances recommendation interpretability by providing clear paths and justifications for suggested collaborations; (2) it mitigates biases toward popular influencers by promoting diversity in selection, ensuring consideration of long-tail influencers; and (3) it addresses cold-start challenges by leveraging relational structures encoded in the knowledge graph. Unlike conventional methods that focus solely on first-order connectivity or oversimplify relational dependencies, our framework incorporates relational paths into representations, preserves dependencies to improve performance and interpretability, and extends to solve the reciprocal recommendation problem by accommodating preferences and constraints of both parties. Empirical evaluations demonstrate that our approach not only improves matching accuracy but also facilitates strategic, sustainable partnerships that adapt to evolving market conditions, offering a comprehensive solution for optimizing influencer-brand collaborations in the dynamic digital marketing ecosystem.

4 - Executives' Digital Human and Corporate Communication

Mijia Gong, Northwestern Polytechnical University, Xi'an, China, People's Republic of

With the development of artificial intelligence, executives’ digital humans have become a new avatar-mediated communication between firms and stakeholders because of their strong working ability, high efficiency, and low cost. Based on executives in the real world, executives’ digital humans represent or assist executives in conveying firm financial and nonfinancial information to stakeholders. Current research mainly focuses on digital humans in the field of marketing and information systems, while paying little attention to the role of executives’ digital humans. Considering that executives’ digital humans are different from virtual influencers in marketing because of the difference in motivation, appearance, and behaviors., this research seeks to explore how executives’ digital humans improve firms’ communication with stakeholders such as investors, consumers, employees, and the public. This research has practical implications for firms to accelerate the application of digital humans, and enrich stakeholder communication channels efficiently. In this way, we can help deepen the implication of AI in corporate communication and help build Metaverse.

5 - Digital Human: the Pandora's Box of Modern Business- Unleashing Efficiency While Trapping Ethics

SiJia Wang, Northwestern Polytechnical University, Xi'an, China, People's Republic of

As a revolutionary frontier, digital humans are reshaping business models, societal structures, and human-digital interactions, and also opening up a plethora of possibilities, demonstrating great potential across education, marketing, finance, and beyond. In practice, digital humans—including Singaporean educator Andrew Chin's multilingual one and those associated with Deloitte, NVIDIA's Jensen Huang, Trip.com Group's James Liang, and JD.com's Richard Liu—actively engage stakeholders across educational and business sectors.

It is essential to understand both the opportunities digital humans present and the challenges they pose. For example, digital humans benefit firms in terms of their cost-effectiveness, scalability, and customizability. Specifically, digital humans deliver competitive advantages via cost reduction, revenue growth, personalized stakeholder engagement, and operational efficiency. For instance, their human-like appearance and real-time interaction capabilities offer users a more engaging and personalized experience in customer service. However, the rise of digital humans also brings several challenges that need to be addressed. The primary risks of digital humans stem from distrust fueled by privacy and data security vulnerabilities inherent in their reliance on personal information collection, coupled with ethical challenges surrounding transparency and consent in human-AI interactions. As digital humans continue to evolve, striking a balance between innovation and ethical responsibility will be crucial.

TB06

Orchard B

Game Theory in Supply Chain and Ownership Structures

Contributed Session

Game Theory and Applications

Chair: Guanhua Huang, N/A

1 - Dual Sourcing with Inherent Product Uncertainty: Will It Compensate for Heterogeneous Consumer Utility

Fanzhuo Zeng, Southampton Business School, Southampton, United Kingdom, Qin Zhou

Numerous brands confronting consumers' heterogeneous utility are adopting a novel strategy referred to as "Dual Sourcing with Inherent Product Uncertainty" (DSPU Strategy hereafter), under which the components are sourced from different suppliers but are interchangeably used in production. The brand sells final products to consumers by disclosing the mixed-use of components rather than the specific component source. Clearly, dual sourcing may induce supply competition and then reduce the procurement cost, consumers' utility may be further hurt because of the disappointment when they receive products containing low-quality components. We hence examine whether it is really beneficial for a brand to adopt DSPU Strategy and if the answer is "yes", under what conditions. Interestingly, we find that the value of DSPU Strategy is beyond procurement cost reduction, but the flexibility to determine the procurement proportions allocated to competing suppliers. We further reveal that an increasing procurement proportion of high-quality components may lower consumers' overall expected value, particularly when they are highly sensitive about products with low-quality component (referred to as Consumers' Utility Paradox in this paper). Therefore, when product differentiation is strengthened and product competition is softened—either due to the use of low-quality components or consumers' significant aversion towards disappointment—brands will be more inclined to favor DSPU Strategy. Our results are robust when the high-quality component suffers from either supply capacity constraint and/or supply disruption risk.

2 - The Implications of Power Structure on Pay-Per-use Business Model

Jiahong Hu, School of Management, Lanzhou University, Lanzhou, China, People's Republic of, Kangzhou Wang

Some dealers within distributed channels adopt the Pay-Per-Use (PPU) model, an innovative business paradigm characterized by pricing based on customer usage. The pricing mechanism not only influences manufacturers' wholesale prices but also interacts with the channel power structure, which determines who has the first-mover advantage to announce a high marginal profit, thereby affecting overall channel pricing decisions. However, the

impact of the power structure on pricing and profit under the PPU model remains unclear. This study investigates the joint impacts of channel structure and the PPU pricing mechanism on channel pricing and profits across three power structures: manufacturer first-mover advantage, dealer first-mover advantage, or neutral first-mover advantage. We find that the PPU price is highest when the dealer has the first-mover advantage, implying that the pricing mechanism enhances the degree of first mover advantage. Specially, the manufacturer even gives up the first move advantage if potential usage is low, which improves the follower's profit. While the PPU model may increase the dealer's profit, it can undermine the first-mover advantage's benefits for him, as he captures a smaller share of channel profits compared to the Selling model. Furthermore, the power structures impact the profit superiority of the PPU model over the Selling model, with the greatest superiority observed when the manufacturer has the first-mover advantage and the least when the dealer does. This study enriches the research on channel pricing under the PPU model, particularly regarding the role of channel power structures.

3 - Cooperative Mnl Model in the Assortment Game

Guanhua HUANG, City University of Hongkong (Join Training with SUSTech), Hongkong, China, People's Republic of

In our work, we study the collaboration among multiple retailers in an assortment planning model. In order to analyze the interaction between the decision maker and his environment, we formulate a choice model to calculate the possibility of a certain choice. Multinomial Logit (MNL) Model is widely used in consumer choice modeling. When the number of retailers extend to n, there will be some cooperation, including product transshipment and demand switching. We assume that all the retailers are controlled by a central decision maker, who will allocate the product type of retailers. Next, we define the notion of balanced game and provide this model's connection with core non-emptiness. The core of TU game (N, v) is nonempty if and only if the game is balanced. When the game (N, v) has a nonempty core, then a unique allocation called nucleolus is always in the core. These theories will be proved whether still exist in our cooperative assortment MNL model. Finally, we apply the Shapley value to this multiple assortment MNL model's allocation problem.

4 - Coalitional Manipulations and Immunity of the Shapley Value

Frank Huettner, SKK GSB, Jongno-Gu, Korea, Republic of, Christian Basteck

We consider manipulations in the context of coalitional games, where a coalition aims to increase the total payoff of its members. An allocation rule is immune to coalitional manipulation if no coalition can benefit from internal reallocation of worth among its subcoalitions (reallocation-proofness), and if no coalition can benefit from a lower worth (weak coalitional monotonicity). Replacing additivity in Shapley's original characterization by these requirements and disposal-proofness (an individual player does not benefit from disposing private worth) yields a new foundation of the Shapley value, i.e., it is the unique efficient and symmetric allocation rule that awards nothing to a null player and is immune to such manipulations.

TB07

Orchard C

HAS - Healthcare Operations and Analytics

Invited Session

Health Systems Modelling with AI

Chair: Guang Cheng, National University of Singapore, Singapore, Singapore

Co-Chair: Joel Goh, NUS Business School, Singapore, Singapore

1 - Optimizing Colorectal Cancer Screening Policies: when Another Fit Before Colonoscopy is Beneficial

Xiaodong Wang, Singapore Management University, Singapore, Singapore, Sarah Yini Gao, Zhichao Zheng, Mabel C. Chou

Current colorectal cancer screening guidelines for average-risk individuals follow a two-stage process: an annual fecal immunochemical test (FIT) in the first stage, and if positive, a second-stage physician consultation, which typically recommends a colonoscopy. However, in practice, some individuals report being advised to undergo another FIT before proceeding to colonoscopy or even wait for the next annual FIT, given the possibility of false positives in FIT screening. Motivated by this discrepancy between guidelines and actual practice, we investigate when and for whom incorporating an additional FIT is beneficial. To address this, we develop a partially observable Markov decision process (POMDP) model to optimize recommendations following FIT results, providing a structured guideline for incorporating another FIT within the screening protocol.

2 - The Role of Contextual Information in Customer Reviews: a Field Experiment on a Mental Health Platform

Guang Cheng, National University of Singapore, Singapore, Singapore, Sidika Candogan, Joel Goh, Bilal Gokpinar

Online mental health platforms can increase the accessibility of mental health care by allowing customers to select a therapist from a large pool of vetted professionals. However, customers may find it difficult to choose a therapist best suited to their needs because therapists typically list multiple specializations that tend to overlap significantly. This makes it challenging for online mental health platforms to attract new customers and retain existing ones. Motivated by this challenge, we conducted a randomized field experiment in partnership with Singapore's largest online mental health platform. In this experiment, we empirically investigate whether a simple intervention on the customer reviews—revealing contextual information about the review writers—can improve the platform's ability to attract and retain customers. We find that, relative to the control, treated customers were 42% more likely to book at least one therapy session, and they booked 87% more sessions in total during the experiment. We also find that customers in the treatment group spent more time searching on the platform, which explains the improvement in the platform's ability to attract and retain customers. Furthermore, treated customers who have the first therapy session booked 36% more sessions with their first chosen therapist and were twice as likely to write a positive review. Overall, our results suggest that revealing contextual information about the review writers improves the informativeness of customer reviews on mental health platforms, and it is a simple and effective way to increase successful matches and foster long-term relationships between customers and the platform.

3 - Admission Control Under Imperfect Information

Yu Luo, The University of Hong Kong, Hong Kong, Hong Kong, Huiyin Ouyang

We consider a queueing system where customers seeking service arrive at a gatekeeper according to a Poisson process. Customers belong to one of two different types, but their types are not directly observable upon arrival. Instead, each customer provides a signal that represents the probability of being of a more important type. Such signals are measured based on the limited information available. Different customer types have different holding costs and rewards. The service provider utilizes these signals to make admission decisions with the objective of maximizing the long-run average reward. We first show the existence of the optimal policy and completely characterize its structure, which is state-dependent on both the signal and the number of customers in the system. We then propose several heuristic policies based on analytical analysis that can be implemented in practice. Through numerical experiments with various settings, we demonstrate that one of the proposed policies, Dynamic Threshold Policy, performs consistently and substantially better than others and is quite close to the optimal one under certain conditions. Furthermore, we use a real-world dataset of ICU admission for Covid-19 patients to conduct a simulation study. The simulation results corroborate the superiority and practicality of the proposed Dynamic Threshold Policy.

4 - Predicting Vaccine Acceptance and Timing of Uptake Using the Understanding America Survey

Sze-chuan suen, University of Southern California, Los Angeles, CA, United States, Citina Liang

Demographic factors may be correlated with attitudes towards vaccines, which may in turn predict vaccine acceptance and uptake timing. Prediction of whether and when acceptance occurs is therefore useful for infectious disease control efforts to identify demographic groups with high vaccine hesitancy. However, in pandemic scenarios, factors for prediction may only be collected and revealed over time, complicating prediction efforts.

We use the longitudinal Understanding America Survey (UAS) and Los Angeles vaccine registry data to build, test, and validate predictive models of COVID-19 vaccine uptake and timing. The UAS collected repeated measures of vaccine attitudes, demographic characteristics, and COVID-19-related health behaviors (e,g., hand washing, etc.) before and after vaccines became available in Los Angeles in 2020, allowing for analysis of whether and how attitudes changed over time.

We use regression and machine learning methods to predict vaccine uptake using only data observed prior to vaccine availability and test our models on the vaccine registry data. Preliminary findings show that vaccine recipients held more positive COVID-19 vaccine attitudes than non-recipients, with the gap widening over time. Among non-recipients, average attitude scores declined from 2.1 to 1.4, while recipients maintained a stable, higher score (2.2-2.5). Moreover, each one-point increase in early vaccine attitude score was associated with a 50-day reduction in vaccine uptake delay after adjusting for demographic characteristics.

Acknowledgements: This work is supported by the Merck Investigator Studies Program (MISP), Merck & Co., Inc.

TB08

Bras Basah A

AI and Machine Learning in Safety and Risk Management

Contributed Session

Informations Systems

Chair: Saurabh Kumar, Indian Institute of Management Indore, Indore, N/A, India

1 - From Noise to Strength: Strategic Noise Composition in LLM-Rag Frameworks

Dayoung Kim, Pohang University of Science and Technology, Nam-Gu, Pohang-Si, Korea, Republic of, Young Myoung Ko

Retrieval-Augmented Generation (RAG) has emerged as a technique for improving the performance of large language models (LLMs), where the inclusion of noisy documents is unavoidable in real-world scenarios. Although recent studies have explored noise within RAG frameworks, most have treated it as detrimental and focused on filtering it out, lacking systematic analysis of its inherent effects. In this study, we empirically demonstrate that certain types and combinations of noisy documents can improve LLM-based QA performance and explore “beneficial noise” by introducing noise types into the LLM-RAG framework. We categorize noise across diverse scenarios using six document categories and conduct experiments on multiple datasets, systematically controlling variables such as document count, position, and noise ratio to evaluate their impact. To investigate underlying causes of performance changes, we employ quantitative metrics to offer insights into how noise contributes to both improvement and degradation. Furthermore, we validate that in the biomedical QA domain, specific combinations of noise types and document configurations can positively affect LLM performance, supporting the generalizability of our findings. While prior studies have attributed performance improvements to attention diversity from added noise, our study extends this view by considering noise diversity and positional randomness as key contributing factors. Based on these findings, we propose a first approach for strategically mixing different types of noisy documents to enhance LLM performance. We believe our study provides new insights for RAG designers, suggesting that noise can be leveraged as a strategic component of document composition to improve robustness and effectiveness.

2 - Designing Explainable Artificial Intelligence for Service Workers: a Representation Theory Perspective

Yueyue Zhang, University of Nottingham Ningbo China, Ningbo, China, People's Republic of, Qi Zhang, Cheng Zhang

Artificial intelligence (AI) is notable for its ability to transform real-world complexities into computational models and differs from traditional information technology (IT) by its autonomy in learning the real world. To explore the design features of AI systems that improve employees’ effective use of AI in a service context, we first conducted a multifaceted study comprising a survey of 252 salespeople, interviews with 10 stakeholders, and a three-day field observation at a well-known beauty company in China. We found that AI’s effective use

depends on its representational fidelity— the degree to which it accurately embodies the structural, supportive, and strategic knowledge relevant to the service context. Expanding upon these initial discoveries, we designed four explainable AI schemes and evaluated their effects on employees' work efficiency via a field experiment across 10 service stores with 134 employees over a month. Our findings indicate that AI systems with higher representational fidelity improve effective AI use by bolstering user trust. Moreover, we identified that the effect of representational fidelity on trust is influenced by users' perception of AI autonomy and their resistance to change. Further analysis of users' post-adoption behaviors suggests that supportive and strategic knowledge presented by AI systems improves effectiveness in customer service behaviors. This research broadens the literature on the effective use of IT by revealing the distinct nature of AI systems. Furthermore, our work bridges representation theory and the emerging research on the design of explainable AI, providing insights for managers aiming to optimize AI deployment among service staff.

3 - Deep LSTM-Driven Tornado Simulation and Loss Modeling for Long-Term Catastrophe Risk Assessment

Chirag Subramanian, Georgia Institute of Technology , Atlanta, GA, United States, Dr. Siamak Daneshvaran

This study presents my independent research initiative with Dr. Siamak Daneshvaran showcasing a deep learning–driven tornado simulation and loss estimation framework using Long Short-Term Memory (LSTM) networks. The model simulates 4,000+ years of synthetic tornado events across the continental United States for use in catastrophe modeling and disaster risk analysis. Using data from the Storm Prediction Center (1950–2021), we corrected early underreporting via statistical detrending and simulated additional events. Seven deep LSTM networks retrained 250 times with varied learning rates generated 5.83 million tornado events capturing start/end coordinates, path length, width, bearing, and magnitude.

A novel geospatial technique mapped each tornado's footprint onto a national 5 km × 5 km grid to compute hit probabilities. Due to the scale (5.83 million polygons × 700,000 cells), spatial intersections were executed using GPU-accelerated vectorized joins and batch processing, reducing runtime from 3 months (CPU) to 2 hours (NVIDIA A100 GPU). The model preserved the historical distribution of tornado intensities, achieving strong realism in magnitude simulation—an essential input for estimating damage and insured loss.

The model estimates an Average Annual Loss (AAL) of USD 6 Billion, aligning with industry expectations for this return period. This work demonstrates the potential of deep recurrent neural networks in forecasting extreme weather events and generating high-fidelity stochastic event sets. Its significance lies in enhancing scenario-based catastrophe modeling, enabling better preparedness, infrastructure planning, and insurance solvency assessments in disaster-prone regions. This approach sets a precedent for the use of deep learning in long-term, high-resolution natural catastrophe simulation.

4 - LLM and Automl for Safety Management and Risk Mitigation: from Safety Algorithms to Edge AI

Ke Wang, Agency for Science, Technology and Research (A*STAR) , Singapore, Singapore, Xiupeng Shi

This study introduces innovative applications of Large Language Models (LLMs) and Automated Machine Learning (AutoML) for maritime safety enhancement. We propose a dual-system architecture integrating LLM-based cognitive reasoning with hardware-aware AutoML, enabling concurrent complex decision modeling and edge-optimized inference capabilities. The methodology is validated through benchmark applications in intelligent transportation systems, achieving automated risk prediction for autonomous vehicle control and unsupervised clustering for smart road diagnostics. Subsequently, the framework is extended to maritime safety system through the implementation of vessel-infrastructure collision risk detection. Furthermore, we investigate FPGA-accelerated edge AI implementations with hardware-software co-design strategies, specifically addressing distinct maritime challenges in sensor fusion and computational constraints. Emerging paradigms like LLM2AutoML are implemented to achieve codeless adaptive modeling. These advancements collectively form an autonomous intelligent system for enhancing safety, responsiveness, and operational efficiency in marine traffic management.

5 - The Sustainability Challenges of Large Language Models and Proposed Solutions

Saurabh Kumar, Indian Institute of Management Indore, Indore, India

The rapid adoption of Large Language Models (LLMs) has transformed industries, yet their extensive energy consumption poses serious sustainability concerns. Training and deploying these models require vast computational power, leading to high carbon emissions and significant water usage for data center cooling. Studies indicate that training a single LLM can consume thousands of megawatt-hours of electricity, exacerbating environmental impact.

A recent example is the generation of **Studio Ghibli-style images using AI**, which showcases the creative potential of LLMs but also highlights sustainability issues. AI models running on power-intensive GPUs contribute to rising operational costs and environmental degradation. If unchecked, this trend may undermine efforts toward a greener digital economy.

To mitigate these concerns, multiple strategies can be used:

- **Energy-Efficient Hardware** – Using specialized AI processors can reduce power consumption.
- **Renewable Energy Integration** – Transitioning data centers to solar, wind, or hydroelectric power can cut fossil fuel reliance.
- **Optimized Architectures** – Techniques like pruning and quantization help decrease model size while maintaining performance.
- **Efficient Training Practices** – Transfer learning and fine-tuning smaller models lower energy usage.
- **Policy and Regulation** – Transparent reporting of AI energy consumption and environmental impact is essential for sustainable AI development.

Addressing these challenges requires a collaborative approach between researchers, policymakers, and industry leaders. By implementing sustainable practices, AI can continue to innovate without compromising environmental responsibility.

TB09

Bras Basah B

Consumer Behavior and Sustainable Market Practices

Contributed Session

Marketing

Chair: Rajesh Sinha, Indian Institute of Management Indore, Indore, India

1 - Research on the Effects of Subscription Models on Consumer Behavior and Profitability of Delivery Platforms

Chi-Cheng Wu, National Sun Yat-sen University, Kaohsiung City, Taiwan, Chia-Yun Yang

With the rapid advancement of the digital age, online platforms have become essential, particularly in delivery services. The market for platforms like UberEats and Meituan has grown significantly, with many adopting subscription models to enhance user engagement and stabilize revenue. While these models can increase engagement, they may also alter consumer behavior, as the sunk cost effect can lead to more frequent usage. However, they may negatively impact platform profits, especially with free delivery, which reduces income from delivery fees and shifts consumer preferences from higher-priced to lower-priced meals.

This study investigates the impact of subscription services on consumer behavior and platform profits. Although previous research has examined reasons for using delivery platforms (e.g., Chen et al. 2020, Guo and Liu 2023), there is limited focus on subscription pricing strategies and their effects on profits. This research aims to fill this gap by examining how subscriptions influence consumer choices and establishing a game-theoretic model for optimal market strategies.

The study consists of two parts: first, in-depth interviews to assess behavioral changes under subscription models; second, developing a game theory model to analyze the effects on consumer behavior and platform profits. Findings suggest that subscription models increase service usage frequency and influence meal selection based on delivery costs. We incorporate these findings into the second part of the game-theoretic model to derive the optimal pricing strategy for platform service providers. These insights will inform optimal pricing strategies for platform providers and contribute to the literature on subscription models in delivery services.

2 - When to Offer Price Match Guarantee: the Blessing of Supply Scarcity

Qian WANG, Lingnan University, Hong Kong, Hong Kong

Price match guarantee, whereby a seller promises to refund the price difference to a customer who purchases a product when the customer finds a qualified lower price elsewhere, is commonly adopted by major retailers. In this paper, we explore the role of supply scarcity within the context of price matching. We consider a competitive market with two firms selling identical products. Each firm is equipped with given supplies, used to serve segmented customers in the market, including the loyal customer who only purchase from the specific firm and switchers who purchase from the firm with a lower price. The firms first determine whether to adopt the price matching policy and then make the price decisions to serve one or two segments of customers. Interesting results emerge from our analysis of various marketing scenarios. First, effective price matching equilibrium emerges in multiple market scenarios, in which one firm offers price matching and charges a high price while the other firm does not offer price matching and charges a low price. Second, fake price matching euiqlibrium emerges while at least one firm can satisfy the total demand from customers. In this equilibrium, one firm offers price matching at a low price while the other does not offer price matching and charges a high price. Finally, our results suggest that the supply scarcity, whether it affects a single firm or is prevalent throughout the market, often prompts the affected firm or the competitor to offer price matching.

3 - Exploring the Future of Intelligent Unmanned Convenience Stores in India: Opportunities and Challenges

TANU PRIYA KOHLI, XLRI Jamshedpur, Jamshedpur, India, Prof. Smitu Malhotra, Prof. Venugopal Pingali, Prof. Rajeev Sharma

Intelligent Unmanned Convenience Stores (IUCVSs) represent a retail model that leverages advanced automation technologies to enable customers to complete transactions independently, without waiting in queues and eliminating the need for any on-site service personnel. Despite the innovative nature of IUCVSs, academic investigations into this retail model remains limited, largely due to its relatively recent emergence in the marketplace. While this retail model has gained traction in regions such as Taiwan and the United States, its overall growth trajectory remains relatively modest in these regions. Its adoption in India brings forth a unique set of challenges, stemming from the country's vast demographic diversity, socioeconomic inequalities, and inconsistencies in technological infrastructure. To investigate the consumer adoption of these stores in India, the study adopts a mixed-methods approach—integrating qualitative content analysis of YouTube video transcripts utilizing the NVivo software with quantitative evaluation utilizing the Cognitive-Affective-Normative (CAN) framework. This investigation empirically assessed the influence of cognitive determinants—specifically customer convenience, trust, and perceived price value—alongside affective determinants, particularly hedonic motivation, and normative determinants, including social influence, on consumer attitudinal formations towards these stores. These attitudinal constructs subsequently shape consumer intentions to utilise these stores. Furthermore, the analysis revealed that demographic variables—encompassing age, gender, income stratification, and educational attainment—serve as significant moderating factors in the relationship between the aforementioned antecedents and consumer attitudinal dispositions towards these stores. The findings contribute substantive theoretical insights while simultaneously providing actionable recommendations that retail management practitioners may leverage in the strategic development and implementation of these stores.

4 - Optimizing Return Policy Duration: a Bayesian Learning Approach

Bretislav Hajek, National University of Singapore, Singapore, Singapore, Jussi Keppo

The rise of e-commerce has driven the adoption of policies allowing customers to return products shortly after the purchase. Free return policies are intended to both increase the demand and protect customers who rely on the truthfulness of the online seller. In this paper, we study how to optimally design these return policies by balancing the benefits of an extended return window against the associated costs. We adopt a Bayesian learning framework in which customers dynamically update their beliefs about product quality based on usage

experience. We derive an explicit formula for a single customer's optimal return policy duration, maximizing the expected seller's profit. Further, we extend our analysis to a heterogeneous population of customers, demonstrating the trade-off between customer acquisition and return-related losses. Finally, we illustrate the practical applicability of our model by fitting it to an e-commerce dataset, showing that our approach accurately captures consumer behavior and yields actionable insights for profit-maximizing return policy design.

5 - Estimation of Customer Satisfaction Shifts: Addressing the Limitation of Upper Ceiling Point

Rajesh Sinha, Indian Institute of Management Indore, Indore, India

Customer satisfaction is seen as the goal of marketing. Often it is measured with point scale (i.e., 5-point scale, 7-point scale) based response feedbacks. The average of satisfaction scores from customers, during a particular period, is considered as the measure of customer satisfaction. The convention of the average of point scores has some limitations. For example, a small number of outliers can influence the average in a significant way because satisfaction scores may not be normally distributed, resulting in bias. Further, because there exists a natural upper point in the point-based scale, a customer giving extreme scores can be considered as a saturated customer. However, adaptation level theory suggests that humans adapted to make extreme as normal, when the extreme happens regularly. In addition, sellers' keep updating themselves, and markets are dynamic places, making the limited point scales incapable to capture the updates and dynamism in the environment. To address the above mentioned problems, drawing form business cycle prediction literature, the research proposes an extensive shift index to gauge the shifts in the overall satisfaction level. The proposed index will work with a customer panel. Instead of the point response scale, the shift scores will drive the index score. Index score will provide the overall shift in the satisfaction level. A high index score will be preferred over a low index score. The gauged shifts will also provide a confidence interval to indicate the uncertainty associated with the estimated shift in the overall satisfaction.

TB10

Hullet Room

Quantum Algorithms for Operations Research with Provable Guarantee

Invited Session

Quantum Computing and OR

Chair: Tongyang Li

1 - Quantum Algorithms for Bandits with Knapsacks with Improved Regret and TIME Complexity

Yuexin Su, Peking University, Beijing, China, People's Republic of, Ziyi Yang, Tongyang Li, Yinyu Ye

Bandits with knapsacks (BwK) constitute a fundamental model that combines aspects of stochastic integer programming with online learning. Classical algorithms for BwK with a time horizon T achieve a problem-independent regret bound of $\mathcal{O}(\sqrt{T})$ and a problem-dependent bound of $\mathcal{O}(\log T)$. In this paper, we investigate the BwK model in the quantum setting, where both reward and resource consumption can be accessed via quantum oracles. We prove that quantum algorithms can improve the problem-independent bound by a factor of $\tilde{\mathcal{O}}\left(\sqrt{m\mathrm{OPT}_{\mathrm{LP}}} \right)$, where m is the number of arms in BwK and $\mathrm{OPT}_{\mathrm{LP}}$ denotes the optimal value. Additionally, we obtain a quadratic improvement in the coefficients of the problem-dependent $\mathcal{O}(\log T)$ bound. Compared to previous works on quantum multi-armed bandits, our study is the first to consider bandit models with linear resource constraints. Furthermore, we improve the time complexity of solving BwK instances using quantum linear programming (LP) solvers from $\tilde{\mathcal{O}} \left(md + \frac{(m+d)^2}{\epsilon_{\mathrm{LP}}^2} + \left(\frac{m+d}{\epsilon_{\mathrm{LP}}^{2}} + d\right) T \right)$ to $\tilde{\mathcal{O}} \left(md + \frac{(m+d)^{1.5}}{\epsilon_{\mathrm{LP}}^{2.5}}+ \left(\frac{\sqrt{m+d}}{\epsilon_{\mathrm{LP}}^{2.5}} +d\right) T \right)$, where d is the number of constraints in BwK and ϵ_{LP} denotes the accuracy of LP solver. we provide a detailed robustness analysis of approximately solving LP problems that may be of independent interest.

2 - Logarithmic-Regret Quantum Learning Algorithms for Zero-Sum Games

Minbo Gao, Institute of Software, Chinese Academy of Sciences, Beijing, China, People's Republic of, Zhengfeng Ji, Tongyang Li, Qisheng Wang

We propose the first online quantum algorithm for zero-sum games with logarithmic regret under the game setting. Moreover, our quantum algorithm computes an ε-approximate Nash equilibrium of an m×n matrix zero-sum game in quantum time $O(\sqrt{(m+n)}/\varepsilon^{2.5})$, yielding a quadratic improvement over classical algorithms in terms of m, n. Our algorithm uses standard quantum inputs and generates classical outputs with succinct descriptions, facilitating end-to-end applications. Technically, our online quantum algorithm 'quantizes' classical algorithms based on the optimistic multiplicative weight update method. At the heart of our algorithm is a fast quantum multi-sampling procedure for the Gibbs sampling problem. This procedure leverages the recently developed quantum approximate k-minimum finding algorithm, a technique that may be of independent interest.

3 - Quantum Walk-Based Algorithm Design

Guanzhong Li, Sun Yat-sen University, Guangzhou, China, People's Republic of, Lvzhou Li, Jingquan Luo

Quantum walk, the quantum counterpart of classical random walk, has gradually developed into an important paradigm in quantum algorithm design. This talk gives an overview of the quantum walk model and its value in solving various problems, focusing on three problems that offer quantum speedups from polynomial to super-exponential: spatial search on graphs, the welded tree problem, and finding the center of a spherical surface. Derandomization of quantum walk-based algorithms is also discussed, which is the process to increase the success rate to 100% while maintaining the original quantum speedup. The problems studied include the famous element distinction problem and triangle finding in quantum computing.

TB11

Morrison Room

Frontiers in Data-driven Decision-making Methods

Invited Session

Revenue Management and Pricing

Chair: Xiaojie Mao, Tsinghua University, Beijing, China, People's Republic of

1 - Reducing Conservativeness of Data-Driven Chance Constraints via a Conformal Statistical Framework

ZHIYUAN HUANG, Tongji University, Shanghai, China, People's Republic of, Jeff Hong, Henry Lam

Chance-constrained optimization is a popular approach to obtain decisions that balance objective performance and risk, by requiring the solution to satisfy safety constraints with a high probability. In the data-driven context, the predominant statistical framework is to attain this high-probability feasibility at a confidence level with respect to past data. Such a confidence-based guarantee, while natural, is primarily suited for decisions applied repeatedly over a long lifespan without the access to new data. However, in many operational settings where decisions are updated periodically based on recent data, this framework becomes overly conservative. In this talk, we propose an alternative statistical framework for data-driven chance-constrained optimization based on conformal prediction guarantees on the obtained decision. Our framework aims to provide less conservative and more adaptable solutions for decision-making over rolling windows. Methodologically, we demonstrate how to "conformalize" a range of approaches for data-driven chance constrained problems, including scenario generation, sampling-and-discarding, FAST, and learning-based robust optimization, which leads to smaller sample size requirements and improved objective performance for short lifespan problems than the classical framework.

2 - Statistical Properties of Robust Optimization Under Distribution Shift

Zhiyi Li, The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong, Xiaojie Mao, Yunbei Xu, Ruohan Zhan

We propose a theoretical framework for evaluating two robust methods called Distributionally Robust Optimization (DRO) and Robusyt Satisficing (RS) under natural distribution shifts. This analysis is performed

via a target excess loss metric, yields key advances over previous results: For DRO and RS, we establish the first generalization bounds that explicitly quantifies the trade-off controlled by its hyperparameter parameter respectively and breaks the curse of dimensionality. We theoretically demonstrate that DRO introduces regularization against distributional shifts, while RS adaptively adjusts robustness via a satisficing term. Further, we provide principled guidelines for selecting hyperparameters for two robust models under partial shift information, supported by simulations comparing their merits. This work bridges theoretical gaps in distributionally robust learning and offers practical insights for algorithm design under distribution shifts.

3 - Contextual Airline Fleet Assignment

Shuming Wang, University of Chinese Academy of Sciences, Beijing, China, People's Republic of, Yi Zhao, Zhe Liang, Zhenzhen Zhang

We investigate an airline fleet assignment problem under demand uncertainty, where the airline aims to maximize its expected total profit by determining the aircraft type assigned to each flight leg and the seat capacity reserved for each itinerary. We study this problem in a data-driven setting, leveraging historical demand data and associated covariates, where the fleet assignment decision itself serves as a key covariate affecting the demand. To account for the effects of contextual information, possible correlation, and decision dependency across itineraries, we employ a demand prediction model of multivariate regression with fixed designs (e.g., price and fleet assignment) and random covariates (e.g., seasonality). Furthermore, we construct a \textit{decision-dependent} predicted demand distribution utilizing the estimated demand model and estimation residuals, and develop a robust-control model with Wasserstein distance to hedge against the distributional ambiguity induced by the predicted demand distribution. We establish finite-sample performance guarantee and asymptotic optimality of the model while considering price-deviation effects, under several regularity conditions. We also analytically measure the value of decision dependency via analyzing the performance gap induced by the decision-omitted bias in the demand model, which captures the effects of price, seat capacity, and the assignment coefficients in the demand model. Computationally, leveraging a reformulation of the robust-control model, we show the proposed contextual robust airline fleet assignment model can be transformed into a mixed-integer linear program, without inducing additional integers to the assignment decisions. Finally, extensive numerical experiments validate the effectiveness of our proposed model from various perspectives.

4 - A Fenchel-Young Loss Approach to Data-Driven Inverse Optimization

Xiaojie Mao, Tsinghua University, Beijing, China, People's Republic of, Zhehao Li, Yanchen Wu

Data-driven inverse optimization seeks to estimate unknown parameters in an optimization model from observations of optimization solutions. Many existing methods are ineffective in handling noisy and suboptimal solution observations and also suffer from computational challenges. In this paper, we build a connection between inverse optimization and the Fenchel-Young (FY) loss originally designed for structured prediction, proposing a FY loss approach to data-driven inverse optimization. This new approach is amenable to efficient gradient-based optimization, hence much more efficient than existing methods. We provide theoretical guarantees for the proposed method and use extensive simulation and real-data experiments to demonstrate its significant advantage in parameter estimation accuracy, decision error and computational speed.

TB12

Moor Room

Food Supply Chain Management III

Invited Session

Revenue Management and Pricing

Chair: Xue-Ming Yuan, National University of Singapore, Singapore, N/A, Singapore

Co-Chair: Dikky Indrawan

1 - Fragility in Food Supply Chain

Eugene Ang, National University of Singapore, Singapore, Singapore, Prasanta Bhattacharya, Chaik Ming Koh, Andrew Lim

The global food supply chain is a complex network that is increasingly vulnerable to disruptions such as natural disasters or geopolitical conflicts. These disruptions can result in extensive shortages, supply delays, and higher consumer costs, thus underscoring the need for effective risk management techniques. In this paper, we propose an empirical framework for modeling global trade flows and simulating the impact of various trade-related shocks on constituent countries. We then draw on this framework to propose a new fragility metric to assess the resilience of the food supply chain network based on the countries' risk of food shortages. We demonstrate the utility of our framework through an illustrative study where we derive the comparative statics of nodal and edge disruptions, and conduct simulation-based experiments using real-world food trade and population data. The results highlight that the fragility index can track changes in the resilience of countries following major shocks to the trade network, with some country-level variance in accuracy. Our proposed framework and the country-level fragility indices provide a sandbox that can help stakeholders better plan, manage, and mitigate the effects of supply chain disruptions, especially benefiting countries at risk of chronic shortages.

2 - The Impact of Blockchain on Profit and Waste in Food Supply Chains

FAN SI, National University of Singapore, Singapore, Singapore, Gongao (Lucy) Chen

Food supply chains can incur significant waste due to different quality standards adopted by supply chain members. In this paper, we investigate the impact of blockchain technology on food waste and profitability in a supply chain with one supplier and one retailer. We find that adopting blockchain technology always improves the supplier’s profit but can hurt the retailer’s and the supply chain’s profitability. We also find that when the supplier’s quality standard is very high, the implementation of blockchain is more likely to reduce the waste in the supply chain.

3 - Integrating Food Values into Consumer Preference Analysis for Banana Purchasing: Implications for Food Supply Chain Management

Alim Setiawan Slamet, IPB University, Bogor, Indonesia

Bananas, a crucial tropical fruit commodity in Indonesia, are experiencing increasing consumer demand and heightened awareness of food-related values. This study aims to explore consumer preferences in banana purchasing by incorporating various food values, including health, sustainability, convenience, taste, safety, and origin. The research will involve a consumer survey in selected urban markets, utilizing the collected data to analyze using a Multinomial Logit (MNL) model. This model will estimate the impact of different product attributes and food values on purchasing decisions. The anticipated results will provide strategic insights for actors across the food supply chain, from producers to retailers, enabling them to develop value-driven supply chain strategies that align with consumer expectations. By integrating consumer value orientation into supply chain decisions, this research seeks to enhance efficiency, responsiveness, and sustainability in the banana supply chain.

TB13

Bailey

Inventory Management

Invited Session

Supply Chain and Inventory Management

Chair: Chaolin Yang, SUFE, Shanghai

1 - Full Supply Chain Flexibility Through Commit-Then-Adjust Strategies

Yi Yu, Shanghai University of Finance and Economics, Shanghai, China, People's Republic of, Tianhu Deng, Jeannette Song

Inspired by companies such as Shein that have achieved strategic advantages by integrating multiple forms of order flexibility, this paper introduces a novel "commit-then-adjust" strategy within a retail-supplier supply chain and a multi-shipment setting, featuring two-directional flexible adjustments in time (order deferring and expediting) and in quantity (total volume increasing and order canceling). Using a dynamic programming model that tracks pipeline shipments through a high-dimensional state variable, we show that the optimal two-directional adjustment follows a target-interval policy with two limits: adjusting upwardly to the lower limit or

adjusting downwardly to the upper limit as close as possible. This finding further provides fresh insights into how various flexibilities reshape decision-making and operational efficiency. In a centralized setting, contrary to the conventional wisdom that various flexibilities act as substitutable strategies, time and quantity flexibilities can be strategic complements. For instance, increased flexibility in cancellation enhances the potential benefits of expediting. Moreover, our findings counter the existing viewpoints that greater flexibility always prompts retailers to order less. Instead, we demonstrate that increased flexibility on expediting or canceling incentivizes retailers to commit to greater total order volumes. In decentralized systems, the retailer tends to order more while the supplier prefers to produces less. Properly designed subsidy-based mechanisms can align the retailer and supplier's decisions.

2 - Multi-Agent Deep Reinforcement Learning for Decentralized Transshipment

jiayu xu, Shanghai University of Finance and Economics, Shanghai, China, People's Republic of, Qiushi Han, Wanyu Zhang, haoyue tang, Chaolin Yang

This study explores decentralized inventory transshipment in multi-retailer networks under complex conditions, including multi-period, non-stationary demand, random lead times, and stochastic yields. Traditional approaches struggle due to the problem's NP-hard nature and the more complex inter-agent interactions caused by decentralization. We use a multi-agent deep reinforcement learning (MADRL) framework with the Heterogeneous-Agent Proximal Policy Optimization (HAPPO) algorithm. Each retailer acts as an agent, making autonomous decisions via a game-theoretic mechanism that incentivizes collaboration. We also design tailored techniques for MADRL in inventory management, such as exogenous information-dependent baselines, reward shaping based on exogenous demand, and an agent-specific central state representation, supported by theoretical proofs.
In such a decentralized problem, we explored using reinforcement learning to design mechanisms, and examined what kind of mechanisms can better promote cooperation among agents, finding that decentralized mechanisms achieve comparable performance to centralized ones due to market self-regulation. Moreover, as the system scale increases, the gap between decentralized and optimal strategies narrows because the market's regulatory ability is enhanced. Additionally, we observed that adjusting the transshipment price influences the activity level of the transshipment market, thereby impacting the overall profitability of the system. Furthermore, our results demonstrate that HAPPO outperforms heuristic algorithms in both synthetic and real-world datasets. Overall, our work not only demonstrates applicability but also offers valuable managerial insights for decentralized inventory transshipment problem.

3 - An Online Mirror Descent Learning Algorithm for Multiproduct Inventory Systems

Sichen Guo, Shanghai University of Finance and Economics, Shanghai, China, People's Republic of, Cong Shi, Chaolin Yang, Christos Zacharias

We address a canonical inventory control problem: multiproduct periodic-review lost-sales inventory system with a warehouse capacity constraint. We study this well-researched problem under the lens of demand learning from censored data. Unlike traditional literature, we do not assume that demand distributions are known a priori. Instead, the decision-maker has access only to observed sales data, while the quantity of lost sales remains unobserved. Existing online learning algorithms bear limitations in providing good-quality solutions to inventory systems offering a large variety of products. We employ and innovate mirror descent with cyclic update techniques to address the challenge of high-dimensionality in product menus. We prove theoretically that our algorithm’s regret bound exhibits a logarithmic dependence on the number of products. This constitutes a significant improvement compared to the square-root regret bound established in existing literature. Based on empirical data, we implemented our methods to assess their practical merit and expose additional managerial insights. Our numerical study confirms that indeed our methodology produces inventory policies superior to existing state-of-the art solutions, especially when managing a large menu of products. Drawing from our numerical observations and theory-informed insights, we provide clear guidelines for practical implementation, along with fine-tuning recommendations.

4 - When Actions and Rewards Mismatch in Intermittency: Imitation and Event-Driven Temporal Difference

Wenfan Ou, Shanghai University of Finance and Economics, Shanghai, China, People's Republic of, Sheng Bi

Mismatches between actions and rewards in reinforcement learning (RL) often arise in complex environments and tasks, e.g., intermittency. Recognizing that general deep reinforcement learning (DRL) algorithms struggle with intermittency, we identify two key sources that causes the mismatches: *overwhelming effect* and *lasting effect*. To mitigate the impact of such mismatches, we develop a novel algorithmic framework, **EV**ent-driven temporal differenc**E** with imitatio**N** **T**raining (**EVENT**), which integrates deep learning (DL)-based imitation learning with DRL-based event-driven temporal difference (TD). This novel TD scheme features adaptive step-sizes, triggered by the event signals of interest. In this paper, we investigate a notorious problem of inventory control for long-tail items, where demand is inherently intermittent. As the optimal solution to this problem remains elusive, we propose a state-dependent heuristic as the target for imitation learning and apply EVENT to address it. We validate our framework using a heavy machinery parts industry dataset. Our algorithm design is deeply inspired the structural analysis from classic operations management methods, and superior numerical results highlight the value of structural knowledge in DRL algorithm design.

TB14

Wilberforce

Game Theory and Contract Design in Supply Chains

Contributed Session

Supply Chain Management

Chair: Weixin Shang, Lingnan University (Hong Kong), SEK212/3, Department of ORM, Faculty of Business, Lingnan University, Hong Kong, N/A, Hong Kong

1 - Analysis of Trade-in Decisions Under Price Commitment

Chenmeng Tian, School of Management, Xiamen University, Xiamen, China, People's Republic of, Zhaowei Miao

Trade-in is an important strategy for promoting the circular economy and stimulating consumption, while price commitment provides consumers with information on trade-in subsidies during the initial purchase phase, thereby enhancing their willingness to buy. The combination of these two strategies not only reduces consumer hesitation and lowers the cost of upgrading but also promotes resource recovery, achieving a win-win situation for both economic and environmental benefits. This paper constructs a game-theoretic model between a manufacturer and a retailer, considering two scenarios under the trade-in policy where either the manufacturer or the retailer provides a trade-in subsidy. For each scenario, we analyze the optimal decisions of supply chain members under both price commitment and no price commitment strategies. The results show that:(1) Regardless of whether a price commitment strategy is implemented, providing a trade-in subsidy helps improve the profits of the subsidizing party.(2) When both the cost of new products and the residual value of old products are low, price commitment benefits the subsidizing party but reduces the profits of the non-subsidizing party; in other cases, price commitment can lead to mutual profit gains.(3) Price commitment increases the first-period sales of new products, while no price commitment results in a higher quantity of traded-in products.

2 - Regulating Counterfeits in Live Streaming Commerce

Jie Zhang, Guangdong University of Finance and Economics, Guangzhou, China, People's Republic of

In live shopping, some internet celebrities mix counterfeit and genuine products to deceive consumers, which can be considered a form of probabilistic selling, with the products being referred to as probabilistic products. Once the live streamer's sale of counterfeit goods is discovered, they will be punished by the live e-commerce platform. Accordingly, consumers can be divided into two categories: expert consumers and general consumers. Expert consumers can 100% identify counterfeit products and return them in time to avoid losses. We construct a game theoretic model to study the platform's regulatory strategies in the presence of such probabilistic sales behavior, and further analyze the impact of the platform's punishment, the live streamer's influence, consumer categorization, and the commission ratios of the suppliers and the platform on the equilibrium.

3 - The Entrepreneur Offshore Contract Manufacturing: when Does It Make Sense to Hire An Agent?

Qiong Chen, Southwestern Univerity of Finance and Economics, Chengdu, China, People's Republic of, Gulru Ozkan-Seely, Johnny Rungtusanatham, Aleda Roth

Consider the case of an entrepreneur with a newly-developed product to be offered to the marketplace with known demand. The entrepreneur has already decided to offshore the manufacture of its product to an overseas manufacturer and now faces the decision of whether or not to engage the services of an intermediary or agent to identify, qualify, and manage the selected overseas contract manufacturer. In this paper, we examine four factors that arguably influence this decision - one pertaining to agent capabilities (i.e., breadth of service offerings), one pertaining to the product (i.e., pace of future product changes), and two pertaining to attributes of the entrepreneur making the decision (i.e., familiarity with country to which production is being offshore outsourced and prior experience in overseas contract manufacturing). We develop a game-theoretical model to derive analytical results regarding dynamics across the four factors. We then analyze experimental data collected using a scenario-based role-playing experiment with entrepreneurs to empirically assess our analytical insights.

4 - Vertical Integration in Multi-Component Supply Chains: the Role of Quality Substitution

Jeongeun Sim, Seoul National University, Department of Agricultural Economics and Rural Development, Seoul, Korea, Republic of, Bosung Kim

While the value of vertical integration is well established in single-component settings, in practice, most products consist of multiple components. This study examines a multi-component supply chain in which a product consists of two components, each provided by a distinct supplier, and the final product quality is determined by the qualities of these components. To accommodate a broad class of quality relationships, we adopt a general quality function based on the constant elasticity of substitution (CES) specification.

We show that, in a multi-component supply chain, the value of vertical integration depends on the degree of quality substitutability between components: vertical integration is profitable only when substitutability is sufficiently high. When suppliers differ in quality capability, the value of vertical integration depends on both the degree of quality substitution and the extent of supplier heterogeneity. We discuss the underlying mechanisms and strategic implications.

5 - Information Sharing and Manufacturer Rebate Competition

Weixin Shang, Lingnan University (Hong Kong), Hong Kong, Hong Kong, Albert Ha, Yunjie Wang

We investigate the incentive for a retailer to share private demand information with two rebate-offering manufacturers who sell substitutable products through the retailer. We show that the retailer’s incentive to share information depends on the proportion of rebate-sensitive consumers, the competition intensity, and whether the retailer can charge a side payment for sharing the information. When the retailer cannot charge a side payment, we show that he will not voluntarily share information with a monopolistic manufacturer, but he may do so with none, one or both of the manufacturers when there is competition. Interestingly, we find that more intense competition or a smaller proportion of rebate-sensitive consumers may benefit a manufacturer if it induces the retailer to share information with her. When the retailer can charge a side payment, we consider the two cases when he either contracts concurrently or sequentially with the manufacturers for sharing the information. We show that the retailer always prefers concurrent contracting, which induces the system-optimal information sharing decision, over sequential contracting.

TB15

Fairmont Ballroom

Traffic Management and Intelligent Systems

Contributed Session

Transportation Science and Logistics

Chair: Haokun Yu, 21 Lower Kent Ridge Road,, 119077

1 - Dynamic Prioritized Coordination in Mixed Traffic: a Game-Theoretic Reinforcement Learning Framework for Multi-Scenario Adaptation

Dan Wang, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, Gangyan Xu, Changsheng Qu, Zhizhou Wu

As Autonomous Vehicles (AVs) become increasingly integrated into traffic systems, effective interaction strategies between AVs and Human-driven Vehicles (HVs) are essential for ensuring safety and efficiency across diverse traffic scenarios. This paper extends our proposed Game-Theoretic Reinforcement Learning (GTRL) framework to optimize vehicle interactions in mixed traffic environments, including merging and lane-changing on highways, and intersections and roundabouts in urban traffic. We formulate the problem as a partially observable Markov decision process (POMDP) and develop a bi-level reinforcement learning approach, where interactions between AVs and HVs follow Stackelberg leader-follower dynamics, and homogeneous vehicle groups (e.g., AV-AV or HV-HV) employ Nash cooperative strategies. Through extensive simulations in a gym-based environment, we evaluate the performance of the GTRL method across multiple scenarios, demonstrating its superiority over traditional rule-based algorithms and other multi-agent reinforcement learning approaches like MADQN in terms of adaptability, traffic flow improvement, and collision risk reduction. This research further derives scenario-specific priority configuration and optimal strategy selection, providing practical insights for traffic management in mixed autonomy. By integrating game-theoretic principles with reinforcement learning, this work advances the optimization of mixed traffic, offering a scalable and robust solution for future autonomous transportation networks.

2 - A Hybrid Physics-Based and Data-Driven Approach for Car-Following Behavior Modeling and Analysis

Qixiu Cheng, University of Bristol, Bristol, United Kingdom, Yang Zhang

This paper proposes a hybrid car-following (CF) model that combines the strengths of physics-based and data-driven approaches to improve prediction accuracy and ensure collision avoidance. Physics-based models offer interpretability and safety guarantees, while data-driven models capture complex human behavior but depend heavily on data quality. We adopt Multi-Output Gaussian Process Regression (MOGPR) as the data-driven component and integrate it with two physics-based models: the Intelligent Driver Model (IDM) and the extended s-shaped three-parameter (ES3) model. Two combination schemes are designed to merge predictions across time steps. Evaluations using the NGSIM and Zen Traffic Data show that the hybrid models outperform both individual approaches, particularly with limited training data. Additional safety experiments confirm the hybrid model's ability to maintain safe gaps in critical scenarios.

3 - Robust Nonlinear Reduced-Order Framework for Freeway Network Control

Haokun Yu, Institute of Operations Research and Analytics, National University of Singapore, Singapore, Singapore, Qiqing Wang, Kaidi Yang

Freeway network control (FNC) involves managing traffic flows across interconnected freeway segments. This problem is typically formulated by discretizing time and space, e.g., using cell transition models (CTMs), which can result in high-dimensional, non-convex system representations. To address these complexities, some studies have theoretically improved the solvability of the FNC problem by converting the CTM into an exact convex program through hypograph relaxation and cumulative flow reformulation.But such relaxation explodes the number of decision variables and accompanying constraints, creating severe computational complexity that can completely offset the advantages of convex optimization, particularly for large scale traffic networks.

To address this scalability issue, we introduce a Robust Nonlinear Reduced Order Freeway Network Control (RO FNC) framework. Leveraging Spectral Submanifolds (SSMs), we distill the nonlinear, high dimensional CTM into a compact low dimensional manifold. By retaining the original set of junction flow control and optimizing them on the reduced order dynamics, we directly minimize the Total Time Spent (TTS) of the reduced-order model. Unlike conventional convex relaxation approaches, our framework preserves the control-input dimension yet projects the high-dimensional dynamics onto a low-dimensional manifold, thereby drastically lowering computational complexity. At the same time, it provides rigorous error bounds between the reduced and full order models, guaranteeing robust controller performance. We validate the framework on two case studies—a real world section of the Grenoble freeway and a synthetic network—and benchmark it against state of the art relaxation techniques, highlighting its efficiency and accuracy.

TB16

VIP A

Sustainable Operations Management, Strategy, and Ecosystem

Invited Session

Sustainable and Socially Responsible Operations

Chair: Yu-chen Yang, National Sun Yat-sen University, Kaohsiung, Taiwan

1 - Platform Strategy and Environmental Policy: a Model of Quality-Based Leasing and Sales Decisions

Shan Xiao, Huazhong University of Science and Technology, Wuhan, China, People's Republic of, Ziyi Wang, Yu-chen Yang, Hao Ying, Xiaowei Mei, Luxi Zhang

This study investigates how car manufacturers strategically allocate high-quality and low-quality vehicles between sales and rental services under carbon emission policies. By developing a monopolistic platform model, we analyze two scenarios: a benchmark with carbon tax costs and an extension incorporating government regulations (penalties and subsidies). Key findings reveal that manufacturers prefer leasing low-quality vehicles and selling high-quality ones when residual costs are low and carbon tax rates are moderate. Conversely, mixed strategies (leasing high-quality and selling low-quality) become optimal under specific quality gaps and residual cost conditions. Carbon policies amplify these dynamics, where emission intensity and regulatory thresholds critically influence strategy selection. Notably, stringent carbon quotas drive manufacturers toward new-energy vehicles to avoid penalties. The results highlight the interplay between regulatory environments, product quality, and profitability, offering insights for automakers to align operational decisions with sustainability goals. This research bridges gaps in literature by integrating quality differentiation, carbon policies, and market segmentation, providing actionable guidance for green transitions in the automotive industry.

2 - The Government Tax and Subside Police in Response of Foreign Government Tariff Policy

Guan-Cheng Zhou, National Sun Yat-sen University, Kaohsiung, Taiwan, Chieh Lee, Yu-chen Yang

This study explores the scenario where, against the backdrop of foreign governments implementing tariffs or non-tariff trade barriers, the local government intervenes in the market through two means: taxation or subsidies. Suppliers, in turn, act as followers and decide whether to make investment improvements and set the prices of their goods. The research employs the framework of game theory and utilizes the social welfare to solve for the strategies aiming to identify the key decision-making factors in supply chain investment.

It analyzes how the government balances the profitability of suppliers with consumer surplus while pursuing the maximization of social welfare. Additionally, it compares different combinations of intervention strategies and evaluates their impacts on suppliers' profits, production efficiency, overall social welfare, and the resilience of the supply chain. The study also delves deeply into the decision-making basis for suppliers to accept or reject government intervention and assesses the influence of investment in upgrading on long-term innovation, profitability, and market competitiveness.

The analysis demonstrates that the provision of government subsidies does not unequivocally enhance supplier profitability in scenarios where investment decisions are made. Furthermore, under taxation-oriented policies, aggregate social welfare may deviate from the socially optimal benchmark. Particularly under external tax pressures, a policy framework characterized solely by domestic taxation without corresponding subsidies and met with supplier non-investment is shown to be unsustainable. These insights underscore the critical importance of incorporating supplier behavioral dynamics and system-wide welfare considerations into the design of effective policy interventions.

3 - Unraveling the Link: How Supply Chain Relationships Shape Corporate Financing Decisions

Shih-Sian Jhang, National Sun Yat-sen University, Kaohsiung, Taiwan

Over the past decade, firms have moved toward fewer key trading partners, resulting in a growing interest in supply chain finance, a field traditionally devoted to working capital decisions, such as trade credit and inventory management. However, capital structure decisions (e.g., debt-to-equity mix) remain largely

unexplored in this context. Choosing long-term financing is an essential decision, affecting a firm's value, growth, risk, and default.

The purpose of this study is to examine how supply chain management affects the capital structure of suppliers. The first stage examines this link through the lens of resource dependence theory. A second study examines the effects of operational slack on the relationship between supply chain structure and capital structure. This study was conducted using data from WRDS and hand-collected 10-K filings of U.S. public companies (1980-2023). A variety of regression models are included. By integrating perspectives from operations, supply chain, and financial management, this study addresses gaps in supply chain finance literature, while also helping stakeholders gain a deeper understanding of the financial implications of supply chain structures.

4 - Balancing Privacy and Profitability: Data Disclosure Strategies in Digital Platforms

Chenghan Wu, Department of Industrial and Information Management, National Cheng Kung University, Tainan City, 701, Taiwan

Ensuring the sustainability of digital platforms requires careful management of user data, which has become a key driver of service optimization, personalization, and monetization. However, the strategic use of such data introduces fundamental trade-offs between profitability and user privacy. This study investigates how a platform can design data disclosure policies that balance economic incentives with privacy concerns. By modeling the interactions between user data provision and firm strategy within a static framework, we analyze how different disclosure schemes influence user participation and platform outcomes. The results reveal that user agency in data sharing plays a critical role in shaping disclosure strategies. Our findings offer theoretical insights into the challenges platforms face and provide practical guidance for responsible and long-term data governance.

5 - Temporal Dynamics in Product Launch: Balancing Innovation and Market Timing for Revenue Management

Wei-Shiun Chang, National Cheng Kung University, Tainan City, Taiwan, Sarah Ereneo

In today’s dynamic technology market, consumers are increasingly informed about product specifications, forcing them to choose between purchasing an existing model or waiting for a superior upgrade. Instead of perceiving these informed consumers as obstacles, manufacturers can harness insights into consumer decision-making to refine product launch strategies. To analyze consumer behavior, we apply a behavioral economics framework centered on temporal preferences, employing a quasi-hyperbolic discounting model to compare the utility of immediate consumption with that of a higher-quality future upgrade. Building on this approach, we develop a profit function model that integrates consumer patience and anticipated product quality. Our numerical analyses illustrate how varying levels of impatience and present bias influence the optimal quality thresholds manufacturers should target to maximize profits. Our findings reveal that patient consumers allow for steady, incremental quality improvements over time, while highly impatient customers necessitate rapid high upgrades that may render delayed launches unprofitable. Furthermore, by extending our model to include marketing investments, we demonstrate that promotional efforts can effectively lower the required upgrade threshold and enhance profitability. However, marketing efforts are only effective at shorter launch intervals. Overall, this study provides actionable insights for firms striving to balance quality innovation with market timing, advancing our understanding of how temporal consumer preferences shape product launch strategies in an evolving technological landscape.

Tuesday, July 22, 1:45 PM - 2:45 PM

Morrison Room

Quantum Computing and Operations Research: Potential, Challenges, and the Path Forward

Invited Session

Keynote

1 - Quantum Computing and Operations Research: Potential, Challenges, and the Path Forward

Stefan Wörner, IBM, Zurich, Switzerland

Quantum computing is a new computational paradigm with the potential to transform many disciplines, with optimization often being mentioned as a very promising candidate. In this talk, I will discuss the state of quantum computing today and explore both its potential and challenges for optimization. I will highlight opportunities for quantum advantage, as well as how classical optimization can support progress in quantum computing. A critical step towards achieving quantum advantage is identifying the right problems to target. To this extent, I will discuss efforts in systematic benchmarking and outline new research directions for advancing quantum optimization.

Fairmont Ballroom

Supply-Demand Matching in Space

Invited Session

Keynote

1 - Supply-Demand Matching in Space

Ming Hu, University of Toronto, Toronto, ON, Canada

This talk provides an overview of recent efforts in building models and theories for supply and demand matching over a space. First, we will model recent applications with spatial movements in the on-demand economy and global supply chains to extend the scope of urban operations research/supply chain management and draw managerial insights on decisions such as pricing and matching in managing those systems. Second, we will develop theories on spatial resource allocation and general choice modeling by drawing connections with the classic optimal transport (OT) problem.

Hullet Room

Bridging the Divide: Advancing Knowledge Through Industry-Academia Collaboration

Invited Session

Keynote

1 - Bridging the Divide: Advancing Knowledge Through Industry-Academia Collaboration

Sharon Ng, Nanyang Technological University, Singapore, Singapore

In an era marked by rapid technological transformation and increasingly complex societal challenges, the need for robust collaboration between academia and industry has never been more critical. I will address the enduring tension and transformative potential of bridging the academic-industry divide. While scientific rigor forms the bedrock of credible knowledge production, relevance to real-world application ensures that such knowledge translates into meaningful societal and economic impact. Though achieving both simultaneously is inherently challenging, neither domain can thrive in isolation. Progress demands deliberate alignment and sustained partnership. I would like to advocate for deeper, more structured collaborations that not only enrich scholarship but also enhance innovation and inform practice. Only by co-creating knowledge that is both rigorous and relevant can we address the pressing needs of our time and shape a more sustainable and inclusive future.

Tuesday, July 22, 3:15 PM - 4:45 PM

TC01

Enterprise

Hypothesis Testing Problems and Multi-Armed Bandits

Invited Session

Applied Probability Society

Chair: Shubhada Agrawal, Indian Institute of Science, Bengaluru, NA, India

1 - Locally Minimax Optimal and Dimension-Agnostic Discrete Argmin Inference

Ilmun Kim, Yonsei University, Seoul, Korea, Republic of, Aaditya Ramdas

In this talk, we study the discrete argmin inference problem in high-dimensional settings. Given n observations from a d dimensional vector, the goal is to test whether the rth component of the mean vector is the smallest among all components. We propose dimension-agnostic tests that maintain validity regardless of how d scales with n, and regardless of arbitrary ties in the mean vector. Notably, our validity holds under mild moment conditions, requiring little more than finiteness of a second moment, and permitting possibly strong dependence between coordinates. In addition, we establish the *local* minimax separation rate for this problem, which adapts to the cardinality of a confusion set, and show that the proposed tests attain this rate. Empirical results illustrate the strong performance of our approach in terms of type I error control and power compared to existing methods.

2 - Tight Lower and Upper Bounds for Non-Parametric Hypothesis Testing Problems and Multi-Armed Bandits

Shubhada Agrawal, Indian Institute of Science, Bengaluru, India

KL-inf, a certain infimum of KL-divergences between probability measures, appears naturally in the lower bounds for multi-armed bandit problems. Algorithms that match these lower bounds exactly (even in the multiplicative constants) have been developed for very general classes of distributions. Such optimal sequential algorithms typically rely on the concentration of appropriate empirical versions of KL-inf.

In this talk, we will consider a simple 1-armed bandit problem or a specific sequential hypothesis testing problem (possibly composite null and composite alternative). We will then see a lower bound on the average number of samples required to ensure delta-correctness in this setup, where delta is a small known constant that is a bound on the error probability. This lower bound will involve the KL-inf. I will then present a sequential test that achieves this lower bound in the limit of small delta, and conclude with some discussion about the KL-inf statistic and optimality.

3 - Quickest change point detection with measurements over a lossy link

RAJESH SUNDARESAN, Indian Institute of Science, Bengaluru, India

Motivated by Industry 4.0 applications, we consider quickest change detection of an abrupt change in a process when its measurements are transmitted by a sensor over a lossy wireless link to a decision maker. The loss probability of link also changes after the abrupt change. We will highlight how to combine information from both measurements and losses, and how to handle the losses, in the asymptotics of low false alarm rate. We will compare service disciplines like last-come-first-served and first-come-first-served in terms of their detection delay. We will share some insights to optimise system design parameters in the nonasymptotic regime. The talk will be based on joint work with Krishna Chaythanya KV, Saqib Abbas Baba, Anurag Kumar, and Arpan Chattopadhyay.

4 - Semi-Bandit Learning for Monotone Stochastic Optimization

Arpit Agarwal, Indian Institute of Technology Bombay, Mumbai, India, Rohan Ghuge, Viswanath Nagarajan

Stochastic optimization is a widely used approach for optimization under uncertainty, where uncertain input parameters are modeled by random variables. Exact or approximation algorithms have been obtained for several fundamental problems in this area. However, a significant limitation of this approach is that it requires full knowledge of the underlying probability distributions. Can we still get good (approximation) algorithms if these distributions are unknown, and the algorithm needs to learn them through repeated interactions? Particularly, can we design an online learning algorithm whose performance smoothly approaches the performance of an (offline) optimal algorithm which has knowledge of these distributions?

In this work, we resolve this question for a large class of “monotone” stochastic problems, by providing a generic online learning algorithm with \sqrt{T log T} regret relative to the best approximation algorithm (under known distributions). Importantly, our online algorithm works in a semi-bandit setting, where in each period, the algorithm only observes samples from the r.v.s that were actually probed. Our framework applies to several fundamental problems in stochastic optimization such as prophet inequality, Pandora’s box, stochastic knapsack,

stochastic matchings and stochastic submodular optimization.

In this talk, I will first introduce the multi-armed bandits framework and describe the popular algorithmic principle of "optimism in the face of uncertainty". I will then give a brief introduction to stochastic optimization. Finally, I will describe our result on conversion from offline to online semi-bandit algorithms for stochastic optimization.

This talk is based on a FOCS'24 paper with Rohan Ghuge and Viswanath Nagarajan- https://arxiv.org/pdf/2312.15427.

5 - Self-Structuring Importance Sampling for Chance-Constrained Optimization

Anand Deo, Indian Institute of Management Bangalore, Bangalore, India

Chance constraints emerge naturally in the context of enhancing the resilience of engineering systems, particularly those exposed to high-impact but rare events such as natural disasters, infrastructure failures, supply disruptions, and financial crashes. However, the rarity of such events renders solving chance-constrained models using direct Monte Carlo methods highly impractical. Importance sampling can be a viable solution method in such cases, as it employs an alternative probability distribution to generate samples more frequently from the rare-event regions. The key challenge, however, lies in informing how to select alternative distributions that remain effective across the decision space and in diverse problem settings. Our work employs self-structuring importance samplers to develop a new sampling-based solution procedure for a broad class of chance-constrained optimization problems. The algorithm transforms samples generated through direct Monte Carlo into importance samples which accurately represent the rare region behavior and can be applied uniformly alike across a wide range of constraint functions and probability distributions. Our results demonstrate that the proposed algorithm offers asymptotically optimal variance reduction and can substantially outperform Monte Carlo sampling in terms of both feasibility and optimality, even with small sample sizes.

TC02

Mercury

Aviation Topics II

Invited Session

Aviation Applications

Chair: Cheng-Lung Wu, UNSW Sydney, Kensington, Australia

1 - Flight Departure Delay Prediction Model in Data-Sparse Scenarios Based on Transformer

Jiuxia Guo, Civil Aviation Flight University of China, Chengdu, China, People's Republic of, Jingyuan Li, Siying Xu, Bo Wu, Haigang Chen

With the expansion of air transportation, traditional departure delay prediction models face limitations in handling challenges such as sparse aviation data, extreme delay scenarios, and multi-source data integration. To address these issues, this study proposes an improved time-fusion-based flight chain delay prediction model (TFT-DCP). First, an embedded historical data retrieval mechanism is introduced to dynamically match similar flight operation scenarios, mitigating the problem of data sparsity. Second, a temporal convolutional network (TCN) is employed to capture multi-scale temporal dependencies, while a gated residual network (GRN) encodes static features of sample data to facilitate the integration of static and dynamic temporal features. Finally, a flight chain delay propagation mechanism is designed, utilizing exponentially decaying weights to quantify the propagation effect of preceding flight delays, thereby characterizing dynamic inter-flight relationships. Experimental results demonstrate that the improved model achieves MAE, RMSE, and R^2 values of 3.23 minutes, 6.25 minutes, and 0.92, respectively, outperforming the original TFT model by 8.8%, 26.0%, and 5.75%. Notably, in extreme delay scenarios, prediction errors are reduced by up to 55.77%. This model effectively addresses the prediction biases of traditional methods that arise from neglecting dynamic dependencies and extreme events in complex flight delay scenarios. It enhances the accuracy and reliability of flight chain delay predictions, providing a theoretical foundation for intelligent flight scheduling.

2 - Aircraft Recovery Problems: A Two-Stage Reinforcement Learning Approach

Hong Lin, UNSW Sydney, Kensington, Australia, Cheng-Lung Wu

Aircraft Recovery Problems (ARP) in aviation is a critical operation challenge faced by airlines worldwide, with irregular operations costing billions annually industry wide. Delays, cancellations and other unexpected disruptions can cascade through the network and affect aircraft routing, crew scheduling and passenger itineraries. The traditional optimisation approaches to ARP struggle with computational complexity and require explicit mathematical formulations that may not capture the dynamics of the real airline's operational environment.

A two-stage Reinforcement Learning (RL) approach was developed. First, the Q-Learning algorithm is constructed for fleet assignment, which aims to learn effective recovery strategies when an airport suddenly becomes unavailable. Second, based on the recovered timetable generated in the first stage, an Q-Learning algorithm is developed for aircraft routing aims to reroute the aircraft that are affected by disruptions. Next, we will integrate neural networks and reward shaping strategies into the model to provide more efficient solution approaches.

The two-stage RL approach is expected to show comparable results in terms of costs to the traditional optimisation approaches and to demonstrate significant improvement in computational efficiency and adaptability to dynamic conditions.

This research addresses the limitations of traditional optimisation methods for ARP by developing an RL framework with a reward shaping mechanism. By developing reward functions that balance operational costs with schedule integrity, our approach guides the agent toward solutions that align with the airlines' priorities. To validate our approach, we utilised schedule data from ROADEF Challenger 2009, which provides a test bed for evaluating our models.

3 - Should We Fly Domestic OR International? Impact of Corsia on International Aviation Network

Ankit Sharma, Indian Institute of Management Amritsar, Amritsar, India

The aviation industry, like many others, is under increasing pressure due to global climate concerns and the enforcement of stringent carbon emission regulations. Simultaneously, unpredictable factors such as fluctuating fuel prices and global crises like the COVID-19 pandemic have further heightened the sector's financial vulnerabilities, with unstable revenues and rising operational costs. This study explores the intersection of these two challenges by examining how the formation of a green coopetitive network between two airlines—each based in different continents—can influence performance under environmental regulations. Specifically, it assesses the impact of the International Civil Aviation Organization's Carbon Offsetting and Reduction Scheme for International Aviation (CORSIA) on this intercontinental network.
To do this, a duopoly-based coopetitive multi-objective hub location problem is developed, focusing on profit maximization while considering carbon emissions. The problem is solved using the augmented ε-constraint method. Multiple scenarios are constructed to observe how the network structure and emission levels evolve under varying conditions. The study reveals that to reduce emissions, long international routes would be strategically eliminated through cooperation. Furthermore, fuel efficiency emerges as a critical determinant in the competition for market share. The results illustrate distinct behavioral patterns between fuel-efficient and less eco-friendly airlines operating under the CORSIA regime, shedding light on the strategic implications of environmental policy in airline network design and competition.

4 - A Comprehensive Evaluation of Air Traffic Network Intelligence

Mengyuan Sun, Nanjing University of Aeronautics and Astronautics, Nanjing, China, People's Republic of, Yong Tian, Cheng-Lung Wu, Jiangchen Li, Junze Sun

Reducing fuel consumption represents a paramount objective for the aviation industry, yet current research lacks comprehensive methodologies for adaptive fuel flow rate assessment across distinct flight phases. This study presents an integrated neural network framework designed to address this gap through three key contributions. First, we develop a comprehensive evaluation framework that simultaneously considers economic and environmental factors in air traffic networks. Second, we introduce an enhanced neural network architecture capable of accurate fuel flow rate prediction throughout different flight phases. The framework further

incorporates graph convolutional networks to model spatial relationships within route networks and employs attention mechanisms to dynamically adjust feature importance according to flight phase characteristics. Collectively, these components form an intelligent decision-support that significantly advances fuel efficiency optimization in aircraft operations.

5 - Dynamic Prediction of Key Milestone TIMES in Ground Handling for Turnaround Flights

Wenying ZHU, Nanjing University of Aeronautics and Astronautics, Nanjing, China, People's Republic of

To address the challenge of coarse-grained monitoring in ground handling process during the implementation of airport collaborative decision making (A-CDM) system at mega-hub airports with annual passenger volumes in the tens of millions, this study transcends the prevailing limitations of static prediction models centered on single target nodes. A dynamic prediction approach for key ground handling milestones is proposed, leveraging a parallel-cascaded random forest architecture. Initially, a comprehensive feature engineering framework is developed based on operational data from a representative high-traffic airport, incorporating a multi-source heterogeneous feature system and a hierarchical temporal information transmission mechanism. Subsequently, an innovative metric - effective remaining buffer time (ERBT) is introduced to quantify the dynamic temporal margins within ground handling procedures. Finally, by coupling parallel feature selection with a cascaded prediction architecture, this study enables progressive dynamic prediction across eight key milestones, ranging from in-block time to estimated off-block time. The research results show that in the initial stage, the prediction accuracy of over 50% of the key nodes exceeds 80% within a ±3-minute error margin. As the handling process advances, the final prediction accuracy of over 50% of the nodes increases to 85%. This study provides a more forward-looking and fine-grained predictive tool and practical paradigm for multi-agent collaborative operations in flight ground handling.

TC03

Minto

Applications of Optimization in Complex Systems

Invited Session

Computing Society

Chair: Bin Tian, The Hong Kong Polytechnic University, Kowloon, N/A

Co-Chair: Kai Pan, The Hong Kong Polytechnic University, Hung Hom, Hong Kong

1 - A Polyhedral Study on Unit Commitment with A Single Type of Binary Variables

Bin Tian, The Hong Kong Polytechnic University, Hung Hom, Hong Kong, Kai Pan, Chung-Lun Li

Efficient power production scheduling is a crucial concern for power system operators aiming to minimize operational costs. Previous mixed-integer linear programming formulations for unit commitment (UC) problems have primarily used two or three types of binary variables. The investigation of strong formulations with a single type of binary variables has been limited, as it is believed to be challenging to derive strong valid inequalities using fewer binary variables, and the reduction of number of binary variables is often accompanied by a compromise in tightness. To address these issues, this paper considers a formulation for unit commitment using a single type of binary variables and develops strong valid inequality families to enhance the tightness of the formulation. Conditions under which these strong valid inequalities serve as facet-defining inequalities for the single-generator UC polytope are provided. For those large-size valid inequality families, the existence of efficient separation algorithms for determining the most violated inequalities is also discussed. The effectiveness of the proposed single-binary formulation and strong valid inequalities is demonstrated through computational experiments on network-constrained UC problems. The results indicate that the strong valid inequalities presented in this paper are effective in solving UC problems and can also be applied to UC formulations that contain more than one type of binary variables.

2 - Robust Nesting Controls for Single-Resource Revenue Management

Tao Zhang, The Hong Kong Polytechnic University, Hong Kong, China, People's Republic of, Miao Song

Nesting control is a popular capacity allocation strategy in revenue management. In nesting policy, products are ranked based on their value and higher-ranked products have access to all the capacity reserved for lower-ranked products. Nested booking limits are then used to control the availability of these products. In this paper, we propose a static single-resource allocation model for computing booking limits in a nesting control scheme. We show that the nested model, which is the minimal of exponentially many convex functions if we simply expand it, can be reformulated as the minimal of linearly many convex functions thanks to our findings. Furthermore, to immunize the model from the inaccurate estimates of the underlying probability distribution, we propose a Wasserstein distributionally robust approach and show that the robust version of our nested model can be reformulated as a mixed-integer convex program with infinitely many constraints. We also develop a constraint generation algorithm to solve this problem.

3 - Prompt Engineering as the Key: Unleashing the Potential of Large Language Models for Heuristic Optimization in Project Planning

Yaning Zhang, The Hong Kong Polytechnic University, Hong Kong, China, People's Republic of, Xiao Li, Yue Teng, Qiping Shen, Sijun Bai

Large language models (LLMs) have demonstrated exceptional capabilities in text processing, creative tasks, and even mathematical reasoning. However, their performance in addressing complex optimization problems—commonly encountered in the planning phase of project management—has remained unsatisfactory. To bridge this gap, we propose Project Planning Prompt Engineering (PPPE), a novel framework that leverages prompt engineering to unlock the potential of LLMs in solving intricate optimization challenges in project planning. PPPE harnesses the textual processing prowess of LLMs by integrating heuristic algorithm programs and structured prompt chains into its design. This approach enables LLMs to autonomously explore and generate new heuristic algorithms for optimization problems. We validate the effectiveness of PPPE using two NP-hard problems in project planning: the classic Resource-constrained project scheduling problem (RCPSP) and the recently emerging Work package sizing problem (WPSP). In RCPSP, the heuristic algorithms discovered by LLMs guided by PPPE outperform all preset heuristic algorithms embedded in the prompts. In the best-performing instance, the solutions obtained are approximately 10% closer to the theoretical lower bound. In WPSP, the heuristic algorithms identified by LLMs with PPPE match state-of-the-art methods and, in some cases, surpass the currently best-known solutions, with improvements approaching the lower bound by approximately up to 3%. PPPE demonstrates the ability to unleash the potential of LLMs in solving optimization problems with minimal computational overhead. By bridging the gap between LLMs and operations research, PPPE paves the way for innovative applications of LLMs in tackling real-world optimization challenges.

4 - Layout Optimization for a Large-Scale Grid-Connected Solar Power Plant

Chong Wang, Huazhong University of Science and Technology, Wuhan, China, People's Republic of, Qinghua Wu, Kai Pan, Zuo-Jun Shen

A solar power plant provides green electricity to the public via a power grid. As governments worldwide have pledged to reduce carbon emissions and achieve carbon neutrality, large-scale grid-connected solar power plants are booming. Developing such a plant requires significant investment, a large proportion of which covers construction costs and critically depends on the plant's layout.

The layout planning involves a series of complex optimization problems, such as district partitioning, photovoltaic component location, and cable routing problems in a solar power plant. These problems have received limited academic attention and are highly challenging because they involve large-scale instances, complex design principles, and complicated physical constraints. Motivated by our collaborative projects with an electrical engineering company in China, this paper specifically focuses on the integrated location and routing problem, which involves locating service ways, inverters, combiner boxes, and routing cables.

We develop exact algorithms to effectively solve the integrated location and routing problem via a decomposition framework, which is proven to produce an optimal solution. We also develop an exact branch-and-cut scheme to solve each subproblem in the decomposition framework by incorporating cutting planes and separation algorithms. Our solution approach is evaluated on 50 real-world data instances via extensive

numerical experiments. Compared with the manual method used in practice, our approach reduces the total cost by approximately 20%. Our decomposition method also achieves an average gap of 0.02% between the obtained lower and upper bounds, significantly smaller than the 16.08% gap achieved with the traditional Benders decomposition.

5 - Gpu-Accelerated Branch-and-Bound with Best First Search: a Framework for Discrete Optimization Problems

Rakesh Nagi, Singapore University of Technology and Design (SUTD), Singapore, Singapore, Samiran Kawtikwar, Izzat El Hajj

Branch-and-Bound has been established as a fundamental method for solving NP-hard optimization problems due to its ability to systematically explore the solution space and find provably optimal solution(s). Among its many variants, Best First Search (BeFS) stands out for its efficiency in guiding the search towards promising regions. The exponential growth of computational effort in Branch-and-Bound for large problems has led to the exploration of GPU-based acceleration. Existing GPU-accelerated solutions employ easy to parallelize search order strategies such as Depth First Search or Breadth First Search. This work presents the first GPU-accelerated Branch-and-Bound framework that employs BeFS, leveraging GPU's high computational power while addressing the challenges of prioritization and load balancing in BeFS. The key contribution of our framework is the design of a concurrent priority queue within a single GPU kernel using a master-worker paradigm. This design enables efficient communication and work distribution among thread blocks, which function as independent workers, thereby ensuring both load balancing and BeFS order.

Using the Resource Constrained Assignment Problem (RCAP) as a test case, we achieve substantial speedups over a single-threaded CPU implementation. We show that our parallel implementation does not explore many additional nodes compared to sequential BeFS, thereby retaining its effectiveness in exploring the "fewest" number of nodes. The framework offers a flexible design with high degree of control over parallelism and includes offloading mechanisms for handling problems that exceed GPU memory capacity. This way, it unlocks efficient GPU-accelerated solutions to general discrete optimization problems.

TC04

Indiana

AI-Powered Decision Optimization and Personalization: Advances in Intelligent Modeling across Domains

Invited Session

Computing Society

Chair: Zhang Yige, NUS, Singapore, N/A, Singapore

1 - Optimizing AI Intervention in Product Customization: Evidence from a Field Experiment

Yifan Ren, CUHK Business School, The Chinese University of Hong Kong, Hong Kong, Hong Kong, Tonglin Zhang

Through a randomized controlled experiment conducted on a proprietary Print-on-Demand e-commerce platform, we examine the causal effect of AI intervention intensity on customer decision-making and purchase behavior. We find that higher levels of AI assistance in design matching and product configuration significantly reduce customer decision time, increase conversion rates, and enhance willingness-to-pay, especially among users who value convenience and experience decision friction in the customization process. Furthermore, we formalize a decision cost–value tradeoff framework to theoretically characterize how AI-driven reductions in perceived decision costs enable more precise customer segmentation, optimal AI configuration, and strategic service-level adjustments, thereby providing concrete operational guidance for customization platforms

2 - Adoption of AI Technology and Multi-Product Duopoly

Weiming Luo, Shanghaitech University , Shanghai, China, People's Republic of, Wenhao Wu, Xingcheng Xu, Yanqing yang

Our work investigates the adoption of AI technology by firms in a duopoly competitive environment with multiple products, particularly focusing on the optimal pricing strategy for AI companies when their technology could significantly impact the downstream market. AI technology enhances product quality, influencing consumer willingness to pay, but also increases unit costs for users and affects competitive dynamics across market segments. We explore a scenario where a technological leader and follower engage in multiproduct duopolistic competition, analyzing the effects of AI adoption on market shares and competitive behavior in segmented markets. The study reveals that AI adoption can either increase or decrease baseline consumption, depending on the AI usage fee, and that a single company's adoption of AI can lead to a competitive advantage. The optimal pricing strategy for AI services is examined in the context of downstream firms' adoption and production strategies, highlighting a trade-off between high AI service prices and widespread adoption among downstream companies. This research contributes to the AI adoption literature by offering a nuanced understanding of how AI technology intersects with multiproduct competition and pricing strategies.

3 - Learning to Schedule for the Dynamic Flexible Job-Shop with Multiple Objectives

Jinsheng Chen, Advanced Remanufacturing and Technology Centre, Singapore, Singapore, Jiao Liu, Meng Xu

We study the dynamic flexible job-shop scheduling problem where the objective is to optimize profit and sustainability-related criteria. We propose a new unsupervised learning method based on genetic programming that can learn good interpretable priority rules that are parametrized by the system state and other problem parameters. We show that training is efficient and that the output policy outperforms benchmarks used in literature and in practice.

4 - Fedmoe a Federated Mixture of Experts Framework for Secure and Efficient Credit Scoring

Linhui Wang, Xiamen university, 厦门市, China, People's Republic of, Jianping Zhu

The increasing accessibility of financial data offers great potential to enhance machine learning models. However, due to the sensitive and high-value nature of credit data, integrating datasets from multiple institutions risks privacy breaches. To address this, we adopt a horizontal federated learning paradigm to protect individual data privacy while collaboratively building a robust global model. Despite its promise, federated learning faces challenges from heterogeneous data distributions, leading to suboptimal model performance. To mitigate this, we propose a novel Federated Mixture of Experts (FedMOE) approach with a multi-stage training framework. First, a global expert model is trained using standard distributed learning to improve generalization across the full data distribution. Second, an inverse distribution dataset is created to enhance recognition of minority classes, boosting recall and addressing class imbalance. Third, local expert models are trained on device-specific data using a shared feature extractor. A dynamic router then combines outputs from the global, class imbalance, and local experts to optimize predictions. This design enhances both global generalization and local accuracy. We evaluate our method on four credit datasets using four performance metrics. Results show that FedMOE enables secure, multi-source credit data integration while significantly improving credit scoring performance. These outcomes support the broader application of FedMOE in credit risk assessment.

TC05

Orchard A

Data-Driven Techniques for Intelligent Systems

Invited Session

Digital Twin and Metaverse

Chair: Haidong Li, University of Chinese Academy of Sciences, Beijing, N/A

1 - Discount Signaling when Customers Search

Yi Chen, Hong Kong University of Science and Technology, Hong Kong, Hong Kong, Jing Dong, Fanyin Zheng

This paper studies the impact of promotion display at online retailing platform when customers search. Promotions are ubiquitous in online retailing and influence the markets profoundly. The platform can choose to either display the promotions at homepage directly or hide such information so that customers can only reveal them after clicking the links of specific products. We use a large-scale dataset that contains detailed individual level data to estimate a simultaneous search model where customers search for promotions and we find that the search cost is considerable. One challenge in our dataset is that we do not observe the category of each product. Actually, it is also hard to define or record appropriate product category for online retailing due to the massive types of products. Utilizing the behavior data, we recover the product category through community detection algorithms. In counterfactual analysis, we show that when no-purchase option and search behavior exist, hiding promotions harms the total market shares of purchase. However, it is possible to increase the revenue by hiding the promotions of some low-profitable products. We also propose an easy-to-implement policy to determine which promotions should be hidden. Based on estimation results, via simulation, we demonstrate that our policy improves the revenue of platform considerably.

2 - Optimizing Location and Supply Strategies for Electric Vehicle Charging Stations with Heterogeneous Operators

Zihan He, University of Chinese Academy of Sciences, Beijing, China, People's Republic of, Haidong Li, Yu An

The rapid adoption of electric vehicles (EVs) necessitates strategic planning of charging infrastructure to balance accessibility, profitability, and competition among heterogeneous operators. This study proposes a two-stage optimization framework aimed at resolving the dual challenges of charging station location and capacity allocation in the first stage, followed by supply strategy optimization in competitive settings during the second stage. In the first stage, a distributionally robust optimization (DRO) approach is developed to handle uncertainties in both the form and parameters of demand distributions, coupled with a bisection method and iterative elimination algorithm to rationalize location and capacity decisions under budget constraints. In the second stage, operators' interactions are modeled as a stochastic non-cooperative game, wherein simultaneous perturbation stochastic approximation (SPSA) is employed for gradient estimation and mirror descent methods facilitate convergence to a Nash equilibrium. The analysis reveals that operators' risk sensitivities—particularly regarding overage and underage costs—significantly influence their strategic preferences for location preferences and supply levels. Numerical simulations demonstrate that the proposed framework not only achieves a system-wide equilibrium but also ensures scalability and robustness. These findings provide valuable insights for policymakers and operators in aligning competitive dynamics with infrastructural efficiency, thereby advancing the sustainable expansion of EV networks.

3 - Explore Then Confirm: Investment Portfolios for New Drug Therapies

Zaile Li, INSEAD, Fontainebleau, France, Stephen Chick, Shane Henderson

Medical innovations require several risky hurdles, such as multiple phases of clinical trials, to be passed before market access and reimbursement is validated. A mechanism proposed to reduce the financial risk of such development, while also improving the chances of identifying a successful treatment, is a portfolio of clinical trials. We propose a stylized model of a portfolio of projects, each of which must successfully pass two phases of clinical trials before market access is possible. Using ideas from Bayesian sequential optimization, we study the value of running response-adaptive clinical trials to flexibly allocate resources across clinical trials in a portfolio. We find there is value in response-adaptive trials, suggest some potentially useful heuristics for the design of response adaptation across trials in a portfolio, and point to areas of further work.

4 - Rate-Optimal Online Learning for Dynamic Assortment Selection with Positioning

Yiyun Luo, School of Statistics and Data Science, Shanghai University of Finance and Economics, Shanghai, China, People's Republic of, Will Wei Sun, Yufeng Liu

In online retailing, the seller aims to offer assortment of items with maximized expected revenue. In this work, we introduce a new online learning problem called Dynamic Assortment Selection with Positioning (DAP) that additionally investigates the positioning of items within the assortment. Specifically, the customers make purchases based on the item attractiveness as the product of the position effect and unknown preference parameter through a multinomial logit choice model. Our objective is to maximize the revenue over a finite

horizon. We first demonstrate that any assortment-only algorithm that neglects position effects results in linear regrets. To address this gap, we propose the Truncated Linear Regression Upper Confidence Bound (TLR-UCB) policy. TLR-UCB utilizes a novel geometric linear-bandit-type feedback structure to construct upper confidence bounds (UCB) for unknown preference parameters, accounting for random and adaptive position effects. To ensure the validity of UCB construction, TLR-UCB adopts a truncation technique for conditional geometric responses before applying linear regression. In theory, we establish a regret upper bound of $\tilde{O}(T^{1/2})$ for TLR-UCB, matching our derived $\Omega(T^{1/2})$ regret lower bound. Moreover, to tackle an extension to unknown position effects, we develop a new Joint Learning and Exploitation (JLE) policy. It involves a joint estimation procedure for unknown preferences and position effects with theoretical guarantees. Extensive experiments demonstrate the superior performance of TLR-UCB and JLE by incorporating position effects into dynamic assortment selection.

5 - Robust Vessel Traffic Scheduling with Uncertain Berth Service TIMES in a Seaport

Shuai Jia, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China, People's Republic of, Lubin Fan, Shuqin Li

We consider a planning horizon during which a set of vessels visit a seaport for cargo transshipment. To access the designated berths, vessels should travel from the anchorage ground to the port basin by passing through a navigation channel. As soon as the vessels have completed the cargo transshipment, they need to travel from the port basin back to the anchorage ground through the navigation channel again. Navigation channel traffic is affected by the tidal effect and is bottlenecked by the limited capacity. During these operations, the port operators need to assign tidal windows for vessels to travel into or out of the port, as well as the berthing and unberthing times of vessels, in order to minimize the overall operating cost. We formulate the problem as a two-stage robust optimization model, considering the uncertain vessel service times at berths. By exploiting the problem structure, we develop an adapted column and constraint generation algorithm framework, where the second-stage problem is solved by an enumeration-based method for generating candidate vessel service sequences and a dynamic programming algorithm for allocating the uncertainty budgets to vessels. The computation experiments show that our proposed algorithm generates optimal solutions within acceptable computation times (less than 30 s), and performs better than well established benchmark methods in terms of both worst-case performance and mean performance metrics. Several managerial insights can be derived from our experimental results to guide port operations in terms of the application of the robust models and benefits to the industry.

TC06

Orchard B

Platform Ecosystems and Digital Markets

Contributed Session

Game Theory and Applications

Chair: Yingchun Nie, School of Economics and Management, University of Science and Technology Beijing, N/A

1 - Pricing Strategy for 3D-as-a-Service: the Impact of Customer Expertise

Donglong Liao, School of Management, Lanzhou University, Lanzhou, China, People's Republic of, Kangzhou Wang

Device providers currently rent 3D printers to customers via the 3D-as-a-Service (3DaaS) model. Providers can adopt one of three pricing strategies: fixed-fee, pay-per-build, or a hybrid of both. Customers who rent 3D printers customize product designs and produce products independently. Their expertise in using 3D printers influences their cost of producing desired products, which in turn affects their printer usage and providers' profits. This paper develops an analytical model to investigate the impact of customer expertise on providers' optimal pricing strategy choices in both monopoly and duopoly markets. We model customer expertise as the ability to successfully print a product. We find that in the monopoly market, when customers' expertise level is high, the provider adopts the hybrid pricing strategy; otherwise, it adopts the fixed-fee pricing strategy. The monopolistic provider always benefits from a higher level of customer expertise. In the duopoly market, we find that in equilibria, symmetric providers adopt asymmetric pricing strategies to focus on different usage frequency

of customers and soften inter-provider competition. Specifically, one provider adopts the fixed-fee pricing strategy, and the other adopts the pay-per-build pricing strategy. Under the varying levels of customer expertise, the fixed-fee provider may obtain higher or lower profits than the pay-per-build provider. Interestingly, an increase in the level of customer expertise may lead to decreases in the profits of both providers. We further examine the issue that providers offer technical training services to customers to improve their expertise. We find that such training services do not always benefit providers.

2 - Repositioning, Ride-Matching, and Abandonment in on-Demand Ride-Hailing Platforms: a Mean Field Game Approach

Yunpeng Li, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Costas Courcoubetis, Antonis Dimakis

The on-demand ride-hailing industry has experienced rapid growth, transforming transportation norms worldwide. Despite improvements in efficiency over traditional taxi services, significant challenges remain, including drivers' strategic repositioning behavior, customer abandonment, and inefficiencies in dispatch algorithms. To address these issues, we introduce a comprehensive mean field game model that systematically analyzes the dynamics of ride-hailing platforms by incorporating driver repositioning across multiple regions, customer abandonment behavior, and platform dispatch algorithms. Using this framework, we identify all possible mean field equilibria as the Karush-Kuhn-Tucker (KKT) points of an associated optimization problem. Our analysis reveals the emergence of multiple equilibria, some of them characterized by drivers pursuing distant requests, i.e., exhibiting "Wild Goose Chase" behavior, leading to suboptimal system performance. To mitigate these inefficiencies, we propose a novel two-matching-radius nearest-neighbor dispatch algorithm that eliminates undesirable equilibria and ensures a unique mean field equilibrium for multi-region systems. The algorithm dynamically adjusts matching radii based on driver supply rates, optimizing pick-up times and waiting times for drivers while maximizing request completion rates. Numerical experiments and simulation results show that our proposed algorithm reduces customer abandonment, minimizes waiting times for both customers and drivers, and improves overall platform efficiency.

3 - Managing Review Inflation over E-Commerce Platforms

Ayesha Arora, Indian Institute of Management Calcutta, Kolkata, India, Tarun Jain

Online sellers often inflate their product ratings over e-commerce platforms. We analyze platform’s efforts towards combating such review inflation. We also analyze various mechanisms, such as incentive and liability mechanisms implemented by social planner to manage review inflation. Our research reveals some interesting managerial insights for the platforms and the regulators. We find that under certain conditions, seller may set a higher review inflation level under liability mechanism compared to incentive mechanism. Furthermore, we find that while implementing regulatory policies to combat review inflation, social planner may need to be more vigilant towards online portal that offers third-party seller’s product compared to a portal where platform offers its own product (private label).

4 - Platform Manipulation in Competitive Crowdfunding

Xinru Hu, Huazhong University of Science and Technology, Wuhan, China, People's Republic of, Jianbin Li, Lijian Lu

Crowdfunding has recently emerged through internet platforms as an alternative to traditional financial intermediaries and retail channels. To meet funding goals, innovative products compete fiercely on the platforms, exacerbating quality issues caused by consumers' weakness in product information. This paper develops a model including both crowdfunding and regular selling stages to investigate the platform's incentive in manipulating product attractiveness (e.g., by presenting an inferior product as superior, or vice versa, through dishonest endorsements) when two creators with different product qualities compete on it. The results indicate that such manipulation alters creators' prices in the crowdfunding stage and affects competition in the regular selling market, thereby incentivizing the platform under specific conditions that depend on creators' funding goals, product quality homogeneity, and the demand of regular selling market. Moreover, creators' static pricing broadens these conditions by amplifying both the price and competition effects. Even the superior creator may benefit from being manipulated to appear less attractive due to the competition effect, leading to a mutually beneficial outcome for creators and the platform. We further demonstrate the robustness of our findings by

extending the model to various variants. Our findings suggest that the platform's incentives should be carefully considered when regulators address information manipulation.

5 - Platform Ecosystem Charging Mode and Pricing Strategy Considering Complementor Innovation

Yingchun Nie, School of Economics and Management, University of Science and Technology Beijing, Beijing, China, People's Republic of

Complementor innovation, defined as the process by which complementors enhance platform functionality through product or service innovation, has become a key driver of value creation in platform ecosystems. This study investigates the platform enterprise’s choice of charging models and pricing strategies aimed at incentivizing complementor innovation. We develop decision models under two charging schemes—fixed transaction fees and registration fees—to analyze the strategic interactions between the platform and its complementors. The study compares the platform’s profitability under each model and explores the optimal charging model selection based on profit maximization. The results indicate that: (1) When the strength of cross network externalities of complementor to users is higher than a certain threshold, the platform enterprise tends to adopt free or subsidized pricing strategies for complementors under both charging models; otherwise, it prefers charging strategies. (2) Under the fixed transaction fee model, if the strength of cross network externalities of complementor to users is higher than a certain threshold, complementors charge users; otherwise, their pricing decisions depend on the innovation cost coefficient. Under the registration fee model, the complementor’s pricing strategy toward users is similarly influenced by the innovation cost coefficient. (3) When the optimal registration fee falls below a specific threshold, the fixed transaction fee model yields higher profits for the platform firm, in which case the fixed transaction fee model should be chosen; otherwise, the registration fee model is more profitable, which should be chosen.

TC07

Orchard C

Enabling Learning Health Systems with AI

Invited Session

Health Systems Modelling with AI

Chair: Thomas Campion, Weill Cornell Medicine, 575 Lexington Ave, New York, 10022, United States

1 - LLM for Clinical Evidence Extraction, Retrieval, and Summarization

Yifan Peng, Weill Cornell Medicine, New York, NY, United States

In medicine, our understanding of diseases, treatments, and preventive measures is perpetually advancing. Despite this, the existing framework of evidence-based medicine (EBM) faces formidable challenges from information overload, as no clinician can catch up with the explosive publication paces, and lacks a robust mechanism for the continuous integration of new evidence. The advent of Generative AI, exemplified by large language models (LLMs)signifies a potentially transformative development in the methodologies employed to collect, appraise, disseminate, and interact with clinical evidence. In this talk, I will discuss our research on addressing the "lost-in-the-middle" issue in evidence retrieval and the fine-tuning of LLMs for effective medical evidence summarization.

2 - Optimizing Child Mental Health Through Multimodal Data and System-Level Interventions

Yunyu Xiao, Weill Cornell Medicine, New York, NY, United States

Youth suicide is a growing global public health crisis. While traditional prevention strategies often focus on individual-level clinical risk factors, emerging research underscores the importance of broader social and behavioral influences, including social determinants of health (SDOH) and patterns of social media use. This presentation will explore how artificial intelligence (AI), machine learning, and multimodal data can be leveraged to inform suicide prevention within learning health systems.

Drawing from large-scale datasets—including electronic health records, community-level SDOH indicators, and digital behavioral data—I will demonstrate how advanced analytics can identify high-risk youth populations, model suicide risk trajectories, and uncover causal pathways to suicidality. I will present several studies using multimodal data collected from large longitudinal surveys, EHRs, autopsies and public organizations for SDOH needs. Emphasis will be placed on integrating these insights into system-level interventions that support early identification and targeted youth suicide prevention strategies.

Real-world applications will be discussed, such as AI-enabled early warning systems and population-level interventions designed to promote mental health equity. The presentation will also address the ethical and practical considerations involved in implementing these innovations across diverse clinical and policy settings. By aligning predictive analytics with actionable policies and care pathways, this work contributes to the broader goal of enabling learning health systems that are adaptive, data-informed, and equity-focused.

3 - Evaluation of Artificial Intelligence-based Clinical Decision Support

Shyam Visweswaran, University of Pittsburgh, Pittsburgh , PA, United States

Artificial Intelligence (AI) has emerged as a transformative force in healthcare, offering unprecedented opportunities to enhance clinical decision-making processes. Integrating AI-based clinical decision support (AI-CDS) systems into healthcare settings aims to improve patient outcomes by leveraging the capabilities of AI algorithms. Despite these promising benefits, the rigorous evaluation of AI-CDS remains in its infancy. Although the US Food and Drug Administration has approved more than 1,000 AI algorithms for medical use, there has been comparatively little formal assessment of their real-world safety, effectiveness, and overall utility in diverse clinical settings. To address this gap, a systematic and multifaceted approach - similar to the protocols used in evaluating a new medication - can provide a robust path forward. Before a drug is widely adopted, it undergoes a series of clinical trials to determine its safety and effectiveness across different populations. A parallel structure for AI-CDS would involve rigorous testing, from initial feasibility assessments and pilot studies to larger-scale trials and post-marketing surveillance. Each step would ensure that the AI-CDS performs reliably under realistic conditions and maintains an acceptable level of clinical risk. Translation of AI algorithms created offline to AI-CDS systems that apply AI algorithms online in real-time is critical for both evaluation and implementation in clinical practice. This presentation will provide a framework for evaluating AI-CDS, describe ongoing AI-CDS projects at the University of Pittsburgh, and compare the data flows needed for various types of CDS systems.

4 - PCORnet and INSIGHT: Operations Research Opportunities at the Intersection of Artificial Intelligence and Learning Health Systems

Thomas Campion, Weill Cornell Medicine, New York, NY, United States

Funded by the Patient-Centered Outcomes Research Institute (PCORI), PCORnet® is a network-of-networks integrating electronic health record (EHR) data and expertise from physicians, scientists, and patients from across the United States. Of the eight networks comprising PCORnet®, the INSIGHT Clinical Research Network (CRN), which has a coordinating center at Weill Cornell Medicine of Cornell University, brings together real-world data (RWD) for more than 20 million patients from academic medical centers across New York City and Houston, Texas. Patient demographics, visits, procedures, diagnoses, medications, laboratory results, and other clinical details are available in the Observational Medical Outcomes Partnership (OMOP) and PCORnet Common Data Models (CDMs) for data scientists and other researchers to analyze. In exploring use of EHR data from PCORnet® and INSIGHT for investigations ranging from target trial emulation to health services research to pragmatic clinical trials among other areas, this talk will highlight opportunities for operations research at the intersection of artificial intelligence and learning health systems.

TC08

Bras Basah A

Digital Transformation and Governance

Contributed Session

Informations Systems

Chair: DILUPA NAKANDALA, Western Sydney University, 169 Macquarie Street, PARRAMATTA, 2150, Australia

1 - Ai in Power Systems - A Strategic Perspective

Swapnil Garg, Indian Institute of Management Indore, Indore, India, Sachin Garg

Artificial Intelligence is pervading all aspects of modern life. This article looks at how this will play out in the power sector from a strategic perspective.

2 - Cognitive Warfare: Challenges and Strategic Countermeasures

Sara Alhajri, Joint Warfare Training Center, Qatar Armed Forces, Doha, Qatar

This paper explores cognitive warfare, a sophisticated and non-kinetic form of modern conflict that strategically targets human cognition to influence and reshape perceptions, attitudes, and behaviors within populations. Cognitive warfare leverages psychological manipulation, misinformation, and advanced digital technologies—including artificial intelligence, social media platforms, and digital communication channels—to destabilize societies by creating confusion, mistrust, and societal polarization.
The study identifies critical short-term and long-term challenges posed by cognitive warfare, such as undermining public confidence in institutions, weakening societal cohesion, and disrupting effective decision-making processes. Furthermore, the research discusses various theoretical perspectives and strategic methodologies relevant to comprehending the underlying psychological and social mechanisms of cognitive warfare.
Additionally, this paper highlights the necessity of employing advanced analytical frameworks, including computational modeling, data-driven analysis, and strategic simulations, to proactively detect, analyze, and respond effectively to cognitive warfare tactics. These analytical tools and methodologies will enable policymakers, defense specialists, and researchers to better understand, anticipate, and counter cognitive threats. The paper proposes strategic countermeasures and practical interventions to effectively combat cognitive warfare, including enhanced information literacy programs, public awareness initiatives, and coordinated cross-sectoral collaboration. It further suggests integrating specialized cognitive threat analysis capabilities within existing defense and policy institutions to reinforce national resilience and societal stability.

3 - Financial Constraints and the Msme Operational Performance

Hariprasad Bellamkonda, Indian Institute of Management-Indore, Indore, India

Financial flexibility plays important role in enhancing the operational flexibility for the MSMEs thereby improving their financial performance. The aim is to examine the impact of digital finance and the risk perception on the performance of the financially constrained MSMEs. The results support the view that digital finance significantly impacts the MSME performance while the financial constraints do not have significant effect on the performance. It is further observed that risk perception among investors significantly affects the operational performance. Overall, the digital financial services improves MSME operational performance through improved business efficiency and market competitiveness.

4 - Leveraging Unstructured Data for Chatbots Morphing

Khai Chiong, University of Texas at Dallas, Richardson, TX, United States, Ryan Dew

Conversational AI systems now mediate millions of customer interactions each day. Yet most chatbots are stylistically static; limited in their abilities to adapt and morph to the style of individual customers with the goal of optimizing persuasion and satisfaction. How can we personalize not only what an AI chatbot says, but how it says it? This study introduces a general Bayesian framework that “morphs” an AI chatbot’s conversational style to individual users in real time by exploiting unstructured consumer-side data – such as textual queries and prior chat logs – captured as high-dimensional text or image embeddings. Our contributions are threefold. First, we introduce a general method for incorporating consumer-side unstructured data into choice models. Existing choice models incorporating unstructured data have focused on the product-side, such as product images and descriptions. Second, we provide a novel empirical application of our model to a conjoint-like dataset of chatbot usage in the wild. Third, our method enables firms to build their own in-house customer-service

and sales chatbots that morph and adapt to the styles of individual consumers. Beyond chatbots, the technique extends to any setting where consumer-side unstructured data can inform individual-level preference heterogeneity.

5 - Big Data Analytics Resource and Integration Capabilities for Competitive Performance: the Role of Dynamic Capabilities

DILUPA NAKANDALA, Western Sydney University, Parramatta, Australia

Increasing investments in big data analytics capability (BDAC) development demonstrate their strategic and operational potential for organisations; however, our understanding of how BDAC leads to improved competitive performance is far from clear. This study takes a capability approach to investigate how organisations achieve competitive performance from their investments in BDAC. Responding to the known ambiguity and vagueness of the BDAC concept, It conceptualises two distinct concepts of *big data analytics resource capability* (BDARC) as a combination of technological and human resources and *big data analytics integration and coordination capability* (BDICC), recognising the actionable role of organisational processes and routines and tests their effects on the competitive performance of organisations and the mediating role of dynamic capabilities. Based on data collected from an online survey and interviews with managers from 375 Australian firms and analysed using the PLS-SEM method, it finds empirical evidence for the effects of potential and realistic BDAC on competitive performance, with the latter having stronger effects. It also finds that dynamic capabilities fully mediate the relationship between potential capabilities and competitive performance but partially mediate the relationship between realised capabilities and competitive performance. It informs the practitioners of the importance of dynamic capabilities in translating their investments in BDA resources and developing realised BDAC for competitive performance.

TC09

Bras Basah B

Sustainable Decision-Making in Environmental and Social Contexts

Contributed Session

Multiple Criteria Decision Making

Chair: Calvin Wah On CHENG, College of Professional and Continuing Education, The Hong Kong Polytechnic University, Hong Kong, N/A

1 - An Intuitionistic Fuzzy Analytical Hierarchy Process-Based Model for Environmentally Sustainable Port Performance

S Sarifah Radiah Shariff, Universiti Teknologi MARA, Shah Alam SELANGOR, Malaysia, Muhamad Safuan Shamshol Bahri , Nazry Yahya, Siti Meriam Zahari

Ports are pivotal in global supply chains but face significant environmental challenges, including emissions, waste, and resource overuse. This study introduces an advanced decision-making framework utilizing the Intuitionistic Fuzzy Analytical Hierarchy Process (IF-AHP) to evaluate environmentally sustainable port performance. By incorporating membership, non-membership, and hesitation, IF-AHP overcomes the limitations of traditional and fuzzy AHP models, enabling a nuanced approach to decision-making under uncertainty. The proposed framework holistically integrates the landlord, regulatory, operator, and community functions of port authorities, ensuring comprehensive sustainability evaluations. The research focuses on Malaysian ports, leveraging expert insights and Intuitionistic Fuzzy Numbers to prioritize key performance indicators across diverse sustainability criteria. Sensitivity and comparative analyses demonstrate the framework's robustness, highlighting critical sub-factors like minimizing operational impacts and enhancing energy efficiency as primary drivers of sustainability. Results reveal the framework's capacity to balance environmental objectives with operational efficiency, offering a scalable tool for global adoption in maritime sustainability. This study contributes to advancing sustainable port management by addressing gaps in existing evaluation methodologies and emphasizing the importance of integrating multiple functional perspectives in environmental decision-making.

2 - Technology-Driven Transformation of Vaccination Systems: Enhancing Equity, Resilience, and Public Trust

Daniela La Marca, CQ University, Melbourne, Australia

The global response to recent pandemics has exposed the vulnerabilities of vaccination systems, while simultaneously accelerating their digital transformation. As health threats increase in complexity, technologies such as predictive analytics, blockchain-enabled logistics, AI-driven surveillance, and integrated digital health platforms are no longer peripheral tools—they are becoming core enablers of system-wide innovation.

This presentation explores how these technologies are reshaping the performance, responsiveness, and transparency of national vaccination programs across high-, middle-, and low-income settings. Drawing on comparative case studies, it examines how countries are embedding digital innovation into their operational strategies to improve vaccine coverage, enhance system adaptability, and build public trust.

A core contribution of this work is the development of a strategic decision-support model that aligns digital interventions with institutional capacities and public health goals. The model identifies critical enablers and constraints, including feedback loops, cross-sector coordination, and governance mechanisms, that influence implementation success.

By integrating operations research and public health analytics, this study offers a systems-level perspective on how digital transformation can drive equitable and sustainable vaccination outcomes. It contributes to the discourse on optimizing complex health systems and supports evidence-informed policymaking that balances innovation with resilience.

3 - Esg-Constrained Portfolio Choice with Estimation Risk

Zhenzhen Huang, The Ohio State University, Columbus, OH, United States, Pengyu Wei, Chengguo Weng

Environmental, Social, and Governance (ESG) investing has emerged as a global trend, offering sustainable benefits to investors and financial institutions. This study integrates an ESG constraint into the classical mean-variance optimization framework, while accounting for the estimation risk associated with the first two moments of asset returns. We begin by examining the problem in the absence of estimation risk and deriving the optimal portfolio characterized by three-fund separation. To address estimation risk, we propose a combined three-fund portfolio, with components based on a plug-in ESG portfolio. The optimal combination coefficients are derived by maximizing the expected out-of-sample mean-variance utility. Furthermore, we provide a comparative performance analysis on the ESG-constrained portfolios involved in the study. Extensive simulations and empirical analysis demonstrate that the combined portfolio outperforms the plug-in ESG portfolio in terms of certainty equivalent return.

4 - Integrating Fuzzy Delphi&Ndash;Topsis for Sustainable Pharmaceutical Waste Treatment Decisions

NURAKMAL BINTI AHMAD MUSTAFFA, Universiti Utara Malaysia, Sintok, Malaysia

harmaceutical waste management has emerged as a critical sustainability challenge due to its environmental and health impacts. This study proposes an integrated fuzzy multi-criteria decision-making (MCDM) framework combining Fuzzy Delphi and Fuzzy TOPSIS to support sustainable treatment selection. The Fuzzy Delphi method is used to capture expert consensus and refine evaluation criteria covering environmental, technical, economic, and social dimensions. Fuzzy TOPSIS is then applied to rank treatment alternatives based on their closeness to an ideal solution under uncertainty. A case study involving Malaysian pharmaceutical experts evaluates five common treatment methods: incineration, autoclaving, microwave treatment, chemical disinfection, and encapsulation. Results indicate that autoclaving is the most sustainable option, balancing ecological impact and operational feasibility, followed by chemical disinfection. Sensitivity analysis confirms the model's robustness under varying weight assignments. The hybrid approach offers a transparent and adaptable decision-support tool for policymakers, healthcare institutions, and environmental agencies aiming to implement green practices in waste management. By integrating expert knowledge with fuzzy logic, the model enhances decision quality in complex, uncertain environments. This contribution is particularly relevant to developing nations where sustainable infrastructure is evolving, offering a replicable framework for prioritizing eco-efficient pharmaceutical waste treatment strategies.

5 - Comparing Generation X and Generation Y on Their Green Food Product Purchase Behaviour

Calvin Wah On CHENG, College of Professional and Continuing Education, The Hong Kong Polytechnic University, Hong Kong, China, People's Republic of, Simon Sheung Man YUEN, Helen Shun Mun WONG, Joseph Lok Man LEE

Green food products refer to food products which are produced in an eco-friendly way and harmless when consumed. With the raise in health and environmental concerns, purchasing green food products has become increasing popular among the consumers in recent years. Numerous studies have examined the buying behaviour of consumers towards green food products, but no study has investigated generation differences in green food product purchase. This study compared the differences in green food product purchase behaviour between Generation X and Generation Y using the theory of planned behaviour (TPB). A sample of 137 consumers was recruited in Hong Kong through an online survey. Descriptive analysis, t-test and PLS-SEM analysis were applied to analyse the data. The results showed that both Generation X and Generation Y had positive attitude and purchase intention towards green food products. In comparison Generation Y had significantly lower attitude and purchase intention than Generation X. The study also found that for Generation X, only attitude and perceived behavioural control positively and significantly affected their purchase intention towards green food products. For Generation Y, attitude, subjective norm and perceived behavioural control all significantly affected purchase intention. Besides, while perceived behavioural control was the most influential factor affecting purchase intention of Generation X, attitude was the most critical predictor of purchase intention for Generation Y. The study findings contribute to the literature regarding green food products and advance our understanding of the consumer behaviour of Generation X and Generation Y.

TC10

Hullet Room

The Quest for Quantum Advantage in Optimization

Invited Session

Quantum Computing and OR

Chair: Narendra Hegade, Kipu Quantum GmbH, N/A

1 - Transforming the Path to Quantum Advantage: Unlocking New Possibilities in Optimisation for Critical Industries

Samuel Marsh, Q-CTRL, Sydney, Australia, André Carvalho

Quantum computing holds immense promise but the limited scale and error susceptibility of quantum hardware has long delayed its practical benefits. In this talk we show how quantum control delivered by software is helping accelerate the adoption of quantum technologies by dramatically improving hardware performance and algorithmic outcomes. While this is not yet superior to classical algorithms, we continue to drive the scale of effective calculation towards this advantage threshold to empirically assess the impact of quantum computing on high-value commercial problems. We will highlight how our quantum-hardware-optimized algorithmic solvers are allowing us to tackle complex optimization challenges in transport, defense logistics, and network security at world-record scales.

2 - Runtime Advantage in Digital Quantum Optimization

Narendra Hegade, Kipu Quantum, Berlin, Germany

We experimentally demonstrate that bias-field digitized counterdiabatic quantum optimization can surpass both simulated annealing and CPLEX in terms of time-to-approximate solutions for higher-order binary optimization problems, utilizing IBM 156-qubits Heron devices. We suitably select the optimization instances that are still solvable by classical methods but running over timescales of several minutes. Runtime advantage is then accomplished via enhanced counterdiabatic protocols that take few seconds on quantum hardware. A scaling analysis conducted on these instances reveals a performance enhancement that becomes more pronounced with increasing system size. These results indicate that available digital quantum processors, when combined with appropriately tailored quantum algorithms, can exhibit signatures of quantum advantage under specific conditions even in the absence of quantum error correction.

3 - Quantum Optimization Techniques for Solving Combinatorial Optimization Benchmark Problems &Ndash; a Comparative Study

Hoong Chuin Lau, Singapore Management University, Singapore, Singapore

Quantum optimization holds promise for addressing classically intractable combinatorial problems, yet a standardized framework for benchmarking its performance—particularly in terms of solution quality, computational speed, and scalability—is still lacking. In this talk, I will introduce a comprehensive benchmarking framework designed to systematically evaluate a range of quantum optimization techniques against well-established NP-hard combinatorial optimization problems. Our framework focuses on key problem classes, including the Multi-Dimensional Knapsack Problem (MDKP), Maximum Independent Set (MIS), Quadratic Assignment Problem (QAP), and Market Share Problem (MSP). Our study evaluates gate-based quantum approaches, including the Variational Quantum Eigensolver (VQE) and its CVaR-enhanced variant, alongside advanced quantum algorithms such as the Quantum Approximate Optimization Algorithm (QAOA) and its extensions. To address resource constraints, we incorporate qubit compression techniques like Pauli Correlation Encoding (PCE) and Quantum Random Access Optimization (QRAO). Experimental results, obtained from simulated quantum environments and classical solvers, provide key insights into feasibility, optimality gaps, and scalability. Our findings highlight both the promise and current limitations of quantum optimization, offering a structured pathway for future research and practical applications in quantum-enhanced decision-making.

TC11

Morrison Room

Applications of Robust Decision Models

Invited Session

Revenue Management and Pricing

Chair: Zhenyu Hu, National University of Singapore, Singapore, Singapore

Co-Chair: Ruiqin Wang, National University of Singapore, Singapore, Singapore

1 - Robust Pricing with Consumer Search

Hao Qiu, Shanghai University of Finance and Economics, Shanghai, China, People's Republic of, Ying-Ju Chen, Pin GAO, Simai He

This study investigates a seller's revenue-maximizing pricing decision when faced with a representative customer who can purchase directly or after a costly search. With precise knowledge of customer valuation, the seller's optimal decisions exhibit the following characteristics in the low (high) search cost regime: (1) the consumer consistently searches before purchasing (purchases without searching); (2) both the seller's optimal revenue and price are non-increasing (non-decreasing) with the search cost. However, when only knowing the distribution's mean-support information, we demonstrate that under worst-case revenue maximizing policy, both the seller's optimal objective and price are non-decreasing in the two cost regimes.

2 - Cost-Sensitive, Envy-Resistant Consensus in Group Decision Making

Jing Tang, Southeast University, Nanjing, China, People's Republic of, Yilin Xue, Rujeerapaiboon Napat, Xinwang Liu

Consensus reaching processes (CRPs) require every decision-maker to adjust their original opinion, incurring both cost and disutility. Traditional CRPs focus solely on minimizing cost while overlooking decision-makers' discontent. However, repeated and excessive discontent can threaten the sustainability of CRPs, especially when certain decision-makers are consistently disadvantaged (i.e., required to make significant adjustments) compared to others, leading to envy and, ultimately, disengagement. To address this, we introduce a new CRP that balances minimizing both cost and envy. Incorporating an envy objective introduces significant computational challenges, as envy is captured by a non-convex function, even when disutility is convex. Fortunately, many non-convex optimization problems, such as bilinear and mixed-integer linear programs, can now be solved efficiently using off-the-shelf solvers. Leveraging this capability, we enhance the generalizability of our cost-sensitive, envy-resistant CRP by assuming that cost and disutility are polyhedral functions, allowing us to reformulate our CRP as a bilinear program. We demonstrate our model using real data from RateMyProfessors.com, which includes student ratings on the quality and difficulty of courses taught by professors.

3 - Estimation and Prediction Procedures for Unified Robust Decision Models

Zhiyuan Lou, Technical University of Munich, Heilbronn, Germany, Zhi Chen, Melvyn Sim, Jingui Xie, Peng Xiong

Over the past two decades, robust optimization techniques have efficiently addressed decision problems under uncertainty, offering high assurance of feasibility without being overly conservative. However, research on estimating parameters for these robust decision models from data has been lacking. In this paper, we focus on a unified framework for robust decision models that integrate robust optimization and robust satisficing paradigms. In particular, we identify two layers of uncertainty: outcome uncertainty, involving deviations from specified inputs based on historical data, and estimation uncertainty, addressing deviations from latent input vectors, such as the unobservable true mean. We introduce estimation and prediction procedures tailored to reduce conservativeness while enhancing feasibility within this unified robust decision framework. The concept of *minimum volume confidence set* is introduced, which minimizes the size of the outcome confidence set while considering the likelihood of infeasibility, thereby improving the precision of uncertainty characterization. This method also accommodates asymmetric uncertainty by adjusting the confidence set accordingly. Additionally, our framework incorporates an affine predictive model that leverages side information to improve input vector predictions, seamlessly integrating into robust decision models. Our method has been implemented in the algebraic modeling toolbox, RSOME, facilitating practical applications.

4 - Equitable Auction Design: with and Without Distributions

Ruiqin Wang, National University of Singapore, Singapore, Singapore, Cagil Kocyigit, Napat Rujeerapaiboon

We study a mechanism design problem where a seller aims to allocate a good to multiple bidders, each with a private value. The seller supports or favors a specific group, referred to as the minority group. Specifically, the seller requires that allocations to the minority group are at least a predetermined fraction (equity level) of those made to the rest of the bidders. Such constraints arise in various settings, including government procurement and corporate supply chain policies that prioritize small businesses, environmentally responsible suppliers, or enterprises owned by historically disadvantaged individuals. We analyze two variants of this problem: stochastic mechanism design, which assumes bidders' values follow a known distribution and seeks to maximize expected revenue, and regret-based mechanism design, which makes no distributional assumptions and aims to minimize the worst-case regret. We characterize a closed-form optimal stochastic mechanism and propose a closed-form regret-based mechanism, and establish that the ex-post regret under the latter is at most a constant multiple (dependent on the equity level) of the optimal worst-case regret. We further quantify that this approximation constant is at most 1.31 across different equity levels. Both mechanisms can be interpreted as set-asides, a common policy tool that reserves a fraction of goods for minority groups. Furthermore, numerical results demonstrate that the stochastic mechanism performs well when the bidders' value distribution is accurately estimated, while the regret-based mechanism exhibits greater robustness under estimation errors.

TC12

Moor Room

Pricing, Incentives, and Dynamics

Invited Session

Revenue Management and Pricing

Chair: Xu Sun, Miami Herbert Business School, Miami, FL, United States

Co-Chair: Xiaohan Zhu, University of Florida, Gainesville, FL, 210004, United States

1 - To Reject OR to Match: a near-Optimal Policy for Online Multi-Way Matching with Reusable Resources

Wang Jiuhong, Dalian University of Technology, Dalian, China, People's Republic of, Mu Du, Xu Sun

In instant delivery systems, a courier or unmanned vehicle can serve multiple orders within a single tour, significantly improving time and cost efficiency compared to single-order dispatching. We model each feasible group of orders as a matching combination, characterized by distinct service rates and revenues. The platform faces an online sequential decision problem: at each decision epoch, it may reject incoming orders and/or assign available servers to combinations, aiming to maximize cumulative revenue while minimizing queuing penalties.

This gives rise to an online multi-way matching problem with reusable resources and two main challenges. First, the heterogeneity in order arrival rates and combination profitability requires balancing immediate matching gains with potential future rewards. Second, servers are reusable but temporarily unavailable during service, making it essential to ensure high utilization and prioritize high-revenue combinations. This calls for maintaining a sufficient queue buffer to support profitable and flexible matching decisions. We propose a two-step policy guided by insights from a fluid approximation model. The first step dynamically controls rejections by managing a queue with optimally designed lengths; the second employs a dual-based, state-dependent greedy heuristic that solves a simple real-time optimization problem. Despite its simplicity, the algorithm achieves near-optimal performance in a large-market asymptotic regime, with a gap of relative to the fluid benchmark, where denotes the number of servers in the system. We further show that maintaining a non-trivial queue is crucial for enabling high-quality matching opportunities in such multi-way matching system.

2 - Optimal Contracting in a Make-to-Order Queueing System with Order Cancellations

Yuting Yan, Nanjing University, Nanjing, China, People's Republic of

This paper investigates optimal contract design between a principal and an agent in a make-to-order queueing system that allows for order cancellations. The agent controls the production rate by exerting privately observed effort, while the principal provides compensation based on the observed queue length—that is, the number of outstanding orders over time. Using dynamic programming, we derive the associated Bellman equations to characterize the optimal contract, explicitly capturing the impact of order cancellations. Because the Bellman equation involves both the control and its first derivative, we apply the calculus of variations to obtain an explicit solution. Our analysis reveals that the optimal incentive payment is non-monotonic with respect to the queue length. We further extend the model to accommodate generally distributed customer patience times, beyond the standard exponential assumption. Numerical experiments support our theoretical results and offer valuable managerial insights. Specifically, we find that when there is no order cancellation, the incentive payment first decreases and then increases with the queue length level. However, as the cancellation rate increases, the incentive payment becomes strictly decreasing with the queue length level.

3 - Customer Retention Campaigns of Punch-in Mechanism

chenxi zhu, Zhejiang University, Hangzhou, China, People's Republic of, Wei Zhang

With the rise of mobile applications, customer retention has become increasingly pivotal for business success. This paper investigates the optimal design of a “punch-in” mechanism—a continuous login reward system aimed at fostering sustained app usage and customer engagement. We propose a dynamic principal-agent model in which customers’ habit formation evolves over time, while the firm orchestrates an incentive structure to maximize retention and profits. Our model captures how customers’ private habit levels influence their punch-in decisions and, in turn, how these decisions shape future habit formation. Under full information, we show that the optimal policy follows a threshold structure, whereby customers continue engaging only when their habit level surpasses a certain threshold. We then characterize the optimal compensation scheme and demonstrate how it adapts to customers’ evolving habit levels. These results offer practical guidance for firms developing customer retention programs and extend the growing literature on dynamic incentive design in digital platforms.

4 - Optimal Caseload Control in Emergency Department Management

Chengkai Wang, Shanghai Jiaotong University , Shanghai, China, People's Republic of, Ran LIU, Ershun Pan

Hospitals' emergency departments are experiencing serious congestion, urgently requiring efficient physician scheduling to improve the service capacity for patients with multiple priority levels and patient return. This study examines the relationship between the number of patients a doctor is responsible for at a hospital and the hospital's service capacity, and proposes corresponding optimization solutions.
Flexible caseload limit setting strategy is an important approach to improve patient service quality. In this paper, we consider managing the patient flow through caseload control, which involves determining whether the patient should be assigned to the physician immediately after triage. First, we analyzed the impact of the caseload limit on the doctor's service efficiency and quality, explaining why a decreasing caseload setting process occurs when hospitals set caseload limits. Next, for a single physician's service system, we established an evaluation method for the system state when the caseload limit is time-varying. We proposed a set of methods for setting the

caseload limit. Finally, we integrated the caseload limit setting with flexible scheduling for multiple doctors, discovering that under different scheduling scenarios, the proposed caseload setting is also flexible and can adapt to various schedules.
Our findings call for immediate attention from healthcare practitioners to carefully balance the trade-off between emergency care efficiency and quality, emphasizing the necessity of choose the optimal caseload for physician.

5 - Why the Best Machine May Not be the Best: Incentivizing Human-Machine Collaboration

Anyan Qi, The University of Texas at Dallas, Richardson, TX, United States, Xiaotong Guan, Shouqiang Wang

Humans and machines need to work collaboratively to make operational decisions such as demand forecasts. In this paper, we study how the machine should be designed to incentivize human-machine collaboration and obtain a rich set of results.

TC13

Bailey

Data Driven Decision Making Under Uncertainty

Invited Session

Supply Chain and Inventory Management

Chair: Caihua Chen, Nanjing University, Nanjing

1 - Value of Moment Information in Wasserstein Distributionally Robust Optimization

Tianxiao Wang, Nanjing University, Nanjing , China, People's Republic of, Caihua Chen, Yinyu Ye

Data-driven optimization often relies on limited sample data, making the trade-off between numerical tractability and information richness a critical consideration. In this paper, we propose a composite ambiguity set that integrates Wasserstein distance with first-order moment information. We begin by quantifying the volume of this ambiguity set under a family of Gaussian distributions. Then, through a non-asymptotic analysis, we evaluate out-of-sample performance, showing that incorporating moment information enhances the robustness of Wasserstein-based distributionally robust optimization. Our findings suggest that with carefully selected parameters, a tighter upper bound can be achieved even with limited data. Finally, we develop a customized algorithm for this framework. Numerical experiments in portfolio selection with real-world data validate the effectiveness of our approach.

2 - Joint Pricing and Inventory Decisions Under Unknown Demand: a Satisficing Approach with Shape Information

Jiawen Yang, Nanjing University, Nanjing, China, People's Republic of, Caihua Chen, Zhen Wang

In this paper, we study a joint pricing and inventory decision problem where the exact price-demand relationship is unknown, but historical data and certain shape information (monotonicity, convexity, elasticity, etc) are available. Unlike traditional estimate-then-optimize approaches, we propose an end-to-end satisficing model that directly integrates historical data and shape constraints to avoid overly conservative choices of demand function and achieve robust profitability against demand estimation error. To address the original nonlinear model with infinitely many constraints, we reformulate the problem into a bilevel optimization model. We derive tractable solutions via second-order cone programs for lower-level subproblems and design efficient algorithms under continuous and discrete pricing sets for the upper-level problem, respectively. Crucially, our model yields some key management insights. First, the value of shape constraints depends on the decision-maker’s profit targets, highlighting a counter-intuitive fact that more information does not necessarily guarantee superior outcomes. Moreover, we validated by experiments that the end-to-end model consistently outperforms estimate-then-optimize methods in key profit metrics. Finally, empirical results demonstrate the strength of the satisficing structure: superior profits alongside robust risk control.

3 - Distributionally Robust Wind Farm Planning Under Decision-Dependent Uncertainty

Zhiqiang Chen, Nanjing University, Nanjing, China, People's Republic of, Caihua Chen, Jingshi Cui, Qian Hu, Wei Xu

As wind energy development accelerates, the demand for multi-farm planning has grown significantly. Integrated with operational considerations, it aims to determine the optimal number of wind turbines at each location to minimize total costs, including investment and operational expenses. Due to the uncertain nature of wind power resources, random fluctuations in the output of dispersed wind farms can partially counterbalance one another. Fully leveraging this smoothing effect requires strategic capacity allocation, which introduces decision-dependent uncertainty into the planning process. To address this, we propose a decision-dependent Wasserstein ambiguity set, where the distribution and radius are modeled as functions of planning decisions, with formulations specified based on the problem's structure. The joint planning and operations problem is then formulated as a two-stage distributionally robust optimization model and reformulated as an equivalent mixed-integer second-order cone programming, with the optimal objective value providing a probabilistic guarantee on the out-of-sample performance. To enhance computational efficiency, we develop a constraint generation-based solution framework that can accelerate the solution procedure by hundreds of times. Numerical experiments using synthetic and real data validate the effectiveness of the solution framework and demonstrate the superior performance of the proposed model.

4 - Distributionally Robust Chance Constrained Optimization for Resource Constrained Project Scheduling

Tianqi Cai, Nanjing University, Nanjing, China, People's Republic of, Xiaodi Fang, Qian Hu

This work presents the first investigation of the Distributionally Robust Chance-Constrained Programming (DRCCP) formulation for the Resource-Constrained Project Scheduling Problem (RCPSP) with uncertain activity durations. Unlike existing approaches, our model explicitly accounts for distributional ambiguity, providing probabilistic guarantees under a family of possible distributions. To solve the resulting challenging two-stage optimization problem, we propose a novel Benders decomposition algorithm, enhanced with a Branch-and-Bound algorithm to accelerate the solution process.

We design several new classes of Benders cuts, which are tailored to the DRCCP-RCPSP structure, and rigorously prove their theoretical validity and strength. Furthermore, to assess the performance of the proposed method, we conduct comprehensive computational experiments, comparing the DRCCP model against two well-established alternatives: the Sample Average Approximation (SAA) model and the Robust Optimization (RO) model. Results show that our approach not only achieves better trade-offs between feasibility and conservativeness but also offers clearer managerial implications in project scheduling under uncertainty.

5 - On Distributionally Robust Resource Constrained Project Scheduling Problem

Xinhao Huang, Nanjing University, Nanjing, China, People's Republic of, Zhixing Luo

We study the resource constrained project scheduling problem with uncertain processing time from a distributionally robust perspective, and we name it DRRCPSP for short. Specifically, we consider two typical ambiguity sets including a Wasserstein one and an event-wise form. And One of those contributions we made is for both problems we propose an equivalent reformulation to avoid the bilinear terms, significantly reducing computational complexity of solving DRO. Taking advantage of this technique, the DRRCPSP could be turned into a deterministic MILP and be solved by general-purpose solvers. Furthermore, we design a brand-new Branch-and-Bound algorithm to solve the problems mentioned above even faster. To the best of our knowledge, this is the first work that uses BB methods on RCPSP with uncertain processing time. Dominance rules and bidirectional search are proposed to make the algorithm more efficient. The approaches are tested via computational experiments leveraging randomly generated datasets, and the computational results validate the good performance in terms of the solution’s quality and exhibits high computational efficiency of our methodology.

TC14

Wilberforce

Innovation and Optimization in Supply Chains and Logistics

Contributed Session

Supply Chain Management

Chair: Xiaoqi Lyu, Tianjin University, Tianjin

1 - Supply Chain Forecasting with Digital Twins in Retail or Food Industry

Dhanashree Pokale, Domino's Pizza Inc., Ann Arbor, MI, United States

Digital Twins concept is gaining popularity as the advancements in artificial intelligence, cloud infrastructure and computing power is on the rise. Supply chains have been taking gradual steps in silos towards enabling digital twin system as such technologies are evolving. It is crucial for data scientists and data leaders in supply chain forecasting space to have holistic understanding of how digital replica will enhance supply chain forecasts.

This talk will focus on below:

1. What is digital twin for supply chain and what are its benefits?

2. What value does digital twin bring to demand forecasting and supply forecasting space?

3. What does the forecasting process look like in a digital twin setup?

Besides, the talk will discuss challenges in current processes and why businesses should employ technologies that assist in digital twin implementation.

The presentation aims to serve as a thought starter and provide directions on how existing forecasting spaces can evolve in coming future.

2 - A Triadic Perspective on Supply Chain Quality Management: Theory Development from Multiple Case Studies

EBRAHIM SOLTANI, HBMSU, Dubai, United Arab Emirates, Yingying Liao

Previous research on supply chain quality management (SCQM) has predominantly centered on (upstream) dyadic interactions within national boundaries, with an emphasis on quantitative micro-theory generalization, while giving sparse attention to the intricate relationships across the entire supply chains and associated theoretical explanations. Drawing on social network and strategic triangle theories, this study moves beyond the traditional dyadic relationship and considers more complex, uncertain and risk-laden supplier-manufacturer-distributor triadic relationship in managing quality within the context of global supply chain. In doing so, it aims to uncover the essential dynamics that characterize the triadic relationships (including any embedded dyads) and assess their impacts on end-product quality. Using theory building through a sample of nine cases in a high growth global industry (shoe), we delineate the dynamics inherent in triadic relationships in terms of supplier strategic positioning, market demand awareness, and supply chain norm compliance, and their consequential effects on product quality. We further explain these relationships across amity (positive), enmity (negative), and ideal status and elucidate their broader implications for end-product quality. The paper concludes with a discussion of implications and directions for future research in SCQM.

3 - Deep Reinforcement Learning Reproduces Order Smoothing in An Interactive Two-Stage Supply Chain

Geng Cui, The University of Tokyo, Tokyo, Japan, Naoto Imura, Katsuhiro Nishinari

We study a two-stage supply chain where decisions are made by two separate deep reinforcement learning (DRL) agents, reflecting how firms often operate independently in practice. Each stage only uses local data and has no prior knowledge of the demand pattern, learning purely from past demand. We look at two settings: (1) a decentralized supply chain, where each stage aims to maximize its own profit, and (2) a collaborative supply chain, where both stages work together to improve total profit. Instead of direct coordination, collaboration is encouraged by shaping the reward function within the DRL framework, allowing the agents to align their goals without sharing full information or relying on central control. Our results show that DRL can find solutions close

to the best possible, and that collaboration helps lower overall costs. We also find that the learned behavior resembles order smoothing, supporting the idea that DRL can discover effective strategies on its own.

4 - How Contractor Network Structural Characteristics Affect Contractor Network Resilience: the Mediating Role of Resource Reconfiguration

Yiwei Wang, Tianjin University, Tianjin, China, People's Republic of, Wenxue Lu, Wenqian Guo

In today's dynamic business environment, resilience has increasingly become a critical capability for inter-organizational project networks to cope with unexpected events. However, prior research has paid limited attention to the relationship between network characteristics and resilience, particularly lacking quantitative empirical examination of network structural characteristics. Based on social network theory, this study investigates how the contractor's centrality and structural holes affect contractor network resilience, while also examining the moderating role of network density. Additionally, drawing on the dynamic capability view, this study further explores the mediating role of resource reconfiguration. Questionnaire data were collected from Chinese general contractors. The results reveal that both centrality and structural holes enhance network proactive resilience and network reactive resilience. Resource reconfiguration fully mediates the positive effect of structural holes on network reactive resilience and partially mediates the remaining relationships. Furthermore, network density negatively moderates the positive effect of centrality on network reactive resilience, whereas positively moderates the effect of structural holes on network reactive resilience. This study contributes empirical evidence on how contractor network structural characteristics affect network resilience. It also extends the application of dynamic capability theory in project resilience research, thus offering new insights into the antecedents of project resilience.

5 - Improved Integrality Gap for Half-Integral Degree Cuts of the Tsp

Toshiaki Yamanaka, Johns Hopkins University, Baltimore, MD, United States

Determining the integrality gap of the linear programming (LP) relaxation of the metric traveling salesman problem (TSP) remains a long-standing open problem. Recently, a 4/3-approximation algorithm was developed for half-integral cycle cut instances—instances in which there are only cycle cuts and every feasible LP solution takes values in {0, 1/2, 1}—emphasizing the necessity of a deeper understanding of degree cuts to make progress on the overall half-integral case (Jin, Klein, and Williamson, *Mathematical Programming* 210: 511-538, 2025). In this work, we present a novel approach that improves the upper bound on the integrality gap for foundational half-integral degree cuts.

6 - The Impact of Supply Instability on Product Quality Failures

Xiaoqi Lyu, Tianjin University, Tianjin, China, People's Republic of, Min Tian, qiyuan zhang, baofeng Huo

As global supply chains become increasingly unstable due to rising geopolitical risks, companies face significant challenges in maintaining production continuity and ensuring product quality. The supply instability amplifies the likelihood of product recalls, making it essential for firms to understand the underlying drivers of such disruptions. This study examines how supply instability influences product recalls, using data from global supply chains spanning 2008 to 2022, and explores the internal and external factors within firms. Our findings reveal that supply instability significantly increases the probability of product recalls. Specifically, firms' flexibility and digital capacity play crucial roles in enabling firms to reduce the impact of supply disruptions. The experience of a firm's board of directors also acts as a protective factor. However, board heterogeneity, while offering diverse perspectives, may exacerbate the negative effects of supply instability on recalls. These results suggest that firms must carefully balance their governance structures to ensure that diversity does not hinder quick decision-making in the face of global supply chain instability.

TC15

Fairmont Ballroom

Urban and Public Transportation Systems

Contributed Session

Transportation Science and Logistics

Chair: Zhen Tan, University of Nottingham Business School China, Ningbo, N/A, China, People's Republic of

1 - Managing Modular Vehicle Operations in a Hub-and-Spoke Network

Yuxi LIU, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China, People's Republic of, Shuai Jia

Modular vehicles (MVs) represent an innovative transportation paradigm that reduces operational costs and enhances passenger convenience by dynamically adjusting vehicle capacity through module coupling/decoupling and enabling en-route transfers. This work formulates the MV operations as an extended pick-up and delivery problem integrating capacity adjustments and transfer coordination within a hub-and-spoke network. We develop a mixed-integer linear programming (MILP) model accounting for the operational details, with an objective to minimize the overall system operational cost. However, the problem’s complexity challenges commercial solvers in obtaining viable solutions within acceptable computation times, even for toy-sized instances. To address this difficulty, we propose a two-stage matheuristic that synergizes Adaptive Large Neighborhood Search (ALNS) with MILP, accelerating computation and enabling scalability to real-world scenarios. Numerical experiments on small-scale instances validate the accuracy of the matheuristic solution and demonstrate its superior convergence compared to standalone Gurobi solution. For large-scale cases, the proposed method efficiently generates high-quality solutions within practical timeframes, while Gurobi fails to produce viable results. Sensitivity analysis and comparative experiments with fixed-capacity on-demand buses (FCOB) in real-world settings further reveal optimal operational configurations and underscore the advantages of MVs in a hub-and-spoke network, emphasizing their potential to balance cost-efficiency and service flexibility.

2 - Resource Allocation and Route Generation for Frequent Urban Parking Data Collection

Elham Heydari, York University, Guelph, ON, Canada, Mehdi Nourinejad

On-street parking monitoring helps alleviate traffic congestion and improve urban mobility. Mobile sensing platforms, such as dashcams and vehicle-mounted ultrasonic sensors, offer a scalable and cost-effective alternative to traditional systems including in-ground sensors and fixed cameras. However, optimizing the deployment of mobile sensors is critical to balancing data quality, spatial coverage, and operational cost. This paper introduces an optimization framework for parking data collection that jointly determines the optimal number of sensing agents (SAs), i.e., vehicles equipped with sensing devices, and their routing paths. The framework integrates sensor allocation and route planning while enforcing headway constraints to maintain data recency and avoid redundancy. The city is modeled as a graph, where nodes represent parking locations and edges denote travel paths, with the objective of minimizing the total travel time of SAs. We develop an exact mixed-integer programming model and a scalable heuristic algorithm, the latter incorporating Lagrangian-inspired refinements to enhance solution quality via node reassignments. Numerical experiments validate the heuristic performance against the exact model, showing that it yields near-optimal solutions with substantially lower computational cost. Results indicate that travel time decreases with more routes and looser headway constraints, while the required number of SAs declines with headway but varies non-monotonically with the number of routes. Finally, we apply our approach to parking facilities in the City of Toronto, demonstrating its scalability and real-world applicability.

3 - Generating and Leveraging Crowdsourced Travel TIME Information

Elham Heydari, York University, Guelph, ON, Canada, Mohammadamir Ahmadian, Sina Bahrami, Mehdi Nourinejad

Crowd-generated travel data (CTD) offers a scalable alternative to traditional methods for estimating travel times and informing routing decisions. Leveraging voluntarily shared travel experiences, CTD provides extensive coverage and real-time responsiveness. However, its effectiveness depends on user participation and timing, specifically the contribution ratio (the proportion of travelers who share data) and the observation window (the period over which data is collected). This paper examines how these factors influence route decisions, estimation accuracy, and overall network performance. We study a two-route network composed of a safe road with fixed travel time and a random road with stochastic travel time, modeled as an M/M/1 queue. The expected travel time on the random road is estimated using the Central Limit Theorem, based on CTD collected during the observation window. A decision model captures traveler choices under uncertainty and risk sensitivity. The

results highlight five key insights. First, overall travel time is minimized at an interior contribution ratio, where information accuracy is balanced against congestion. Second, the share of travelers choosing the random road increases with the square root of the travel time on the safe road. Third, introducing risk aversion reduces the probability of choosing the random road and increases reliance on the safe alternative. Fourth, when real-time travel information is available for a fee, travelers' willingness to pay depends on both the access cost and the reliability of free CTD. Fifth, offering rewards for data sharing lowers the required contribution ratio and reduces total travel time relative to the baseline model.

4 - Coordination and Pricing Strategy in Mobility as a Service (Maas) System

Jie Zhang, University of Bristol, Bristol, United Kingdom, Min Ye, Maria Battarra, Gilbert Laporte, Meng Meng

Mobility as a Service (MaaS) is an innovative concept, that integrates various transportation services into a unified platform, enhancing accessibility, reducing congestion and improving efficiency. In the evolving market, MaaS operators and traditional service operators coexist, serving all customers while fostering collaboration between transportation operators. This study explores coordination and pricing strategies between operators in a competitive market with stochastic demand. We model the problem as a Stackelberg game where the traditional operator acts as a leader, setting capacity and pricing, while the MaaS operator, as a risk-averse follower, optimizes resource purchases. We analyse decentralised, centralised, and coordinated decision models, introducing a joint contract mechanism that aligns incentives between operators. The coordination framework incorporates wholesale pricing, revenue sharing, and risk-return contracts to mitigate demand uncertainty. Numerical experiments highlight the impact of market share, demand sensitivity, and risk aversion on optimal pricing and capacity decisions. Results show that a joint contract can enhance system efficiency and promote sustainable collaboration. These findings provide managerial insights for promoting operational efficiency and cooperation in MaaS ecosystems.

5 - Matheuristic for Feeder Network Design Problem with Optional Paired Demands and Split Deliveries

Mohammad Moshref Javadi, University of Illinois at Urbana-Champaign, Champaign, IL, United States, Anandhu Dileep, Ahmad Hemmati, Rahul Nath

We study a complex variant of the feeder network design problem from the field of maritime transportation, where weekly scheduled ship routes must be designed to transport containerized cargo between different ports based on a weekly demand. In some scenarios, this is done using disjoint networks of ships services, called feeder networks. The feeder networks are connected to a main port, with which most of the demands are associated, by long ship services called mother routes. Some of the paired demands are delivered by the mother routes themselves, while the remaining demands are delivered using daughter routes, which are shorter ship services that branch off from some of the ports in the mother routes.

So, the problem involves designing a set of mother routes and associated daughter routes, called feeder networks, to satisfy the paired demands. We include the additional considerations of split deliveries of the demands and the rejection of demands with a penalty, and many realistic considerations such as a heterogenous fleet with weekly ship renting costs, speed optimization of ships, loading-unloading costs and times, holding costs and time constraints on the routes and the delivery of the demands.

We develop a matheuristic framework to tackle this problem. For exploring the space of possible networks, we use ALNS and all other decisions are made by the mathematical model and use a matheuristic framework to combine the two, resulting in a promising solution approach.

6 - Data-Driven Arrival TIME Suggestion for Drivers in Trucking Platform of Export Containers

Zhen Tan, Nottingham University Ningbo China, Ningbo, China, People's Republic of

Trucking platform has emerged in recent years for improving export freight logistics. As a bilateral platform, it needs to balance the satisfaction of exporters and truck drivers. A typical scenario is that given the required start loading time of the container at the exporter's warehouse, the platform can make an arrival time suggestion to the assigned driver so as to reduce the expected time cost for both the exporter and the driver. This decision

problem is nontrivial due to uncertainties in drivers' punctuality and stay times at the warehouses as well as partial schedule and truck information possessed by the platform. We propose a model which computes the suggested arrival time in real-time by solving a stochastic program once a new pick-up task is assigned. We propose two data-driven strategies and prove properties of both single- and bi-objective formulations with efficient algorithms for typical punctuality distributions. Experiments on a real dataset show that the proposed model can potentially reduce the average total time cost by 3~8% against historical decisions over various parameter settings. Optimal decision patterns are affected by different task-, driver- and warehouse-specific features. Impact of key parameters of the algorithm and the model are also discussed with managerial insights.

TC16

VIP A

Environmental and Social Impact: Strategic Insights in a Changing Market

Invited Session

Sustainable and Socially Responsible Operations

Chair: Yue Dai, Fudan University, Shanghai, N/A

1 - Cost Increases and Consumer Fairness in Channel Efficiency: Conflict OR Synergy?

Yue Dai, Fudan University, Shanghai, China, People's Republic of, Bin Dai, Yuwen Du

When confronted with cost increases, a supplier typically raises prices and passes the burden onto its downstream buyer. However, in the presence of consumer fairness concerns, this strategy exacerbates the double marginalization effect, ultimately reducing the overall profit of the supply chain. This paper examines the intertwined impact of cost increases and consumer fairness concerns for three types of supply chains (agency selling, reselling, and hybrid). We find that, in the presence of cost increases, fairness concerns could benefit both the supplier and the platform, due to overcharged wholesale price and the dual-channel co-operative relationship. This is a blessing in disguise for firms since either of the two factors can be detrimental when the other is absent.

2 - Addressing Supply Chain Carbon Emissions: Measurement, Reporting, and Reduction Strategies

Yilin SHI, Chinese University of Hong Kong, Hong Kong, China, People's Republic of, Jing Wu, Yi Liao

This paper aims to honor Professor Christopher S. Tang's contributions to sustainable operations management by examining the pivotal yet under-appreciated role of supply chains in global carbon emissions, highlighting the significant challenges in measuring, reporting, and reducing these emissions. Despite the critical importance of supply chains in contributing to overall carbon footprints, current efforts to address these emissions are often overshadowed by more visible direct emissions (Scope 1) and energy-related emissions (Scope 2). We review existing methodologies for carbon measurement, such as the Environmentally Extended Input-Output (EEIO) model and product-specific life cycle assessments (LCAs). In particular, we introduce two novel empirical papers on supply chain emissions and selective disclosure of corporate suppliers. We also review existing literature on reduction strategies in global supply chains, including supply chain mapping, collaborative efforts, and the impact of government policies like carbon border adjustments. Additionally, we present case studies of leading companies, Lenovo and Siemens, to illustrate practical approaches to global supply chain decarbonization and the regulatory landscape, emphasizing the urgent need for accurate measurement and transparent reporting to drive meaningful supply chain climate action.

3 - Social Labeling: a Retail Perspective

Xiajun Pan, University of Florida, Gainesville, FL, United States, honggnag Hu, Quan Zheng, Asoo Vakharia

Labeling, widely accepted by consumers, communicates corporate social performance. Manufacturers increasingly seek social labels, often certified by impact-driven NPOs. However, many products are sold via profit-driven retailers, some of whom now act as certifiers themselves. This paper explores the implications of social labeling when dominant retailers—rather than NPOs—certify product impact, revealing new strategic dynamics.

4 - Inventory Management in a Multi-Echelon Supply Chain Under Supply Disruption Risk

Zhiyuan Zhu, Fudan University, Yangpu

Tuesday, July 22, 5:00 PM - 6:30 PM

TD01

Enterprise

Asymptotic Analysis and Operation Management for Queues

Invited Session

Applied Probability Society

Chair: Yongjiang Guo, Beijing University of Posts and Telecommunications, Beijing

Co-Chair: Jian Cao, Nanjing University of Finance & Economics, Nanjing, N/A

1 - Optimal Pricing for Queueing Systems with Loss-Averse Customers

jian cao, Nanjing University of Finances & Economics, School of Applied Mathematics, Nanjing, China, People's Republic of, Hongfan Chen, Yongjiang Guo, Junfei Huang

This paper investigates how loss aversion, characterized by emotions of disappointment or regret, influences the optimal pricing strategies in a service system with uncertain service rates both from the perspectives of both revenue maximization and social welfare maximization. We use a single-server queueing model with uncertain service rate. We respectively develop the state-independent and state-dependent models to capture customers' disappointment-based loss-averse (DLA) behavior and regret-based loss-averse (RLA) behavior. We explore the implications of DLA and RLA behaviors in terms of customers' decisions, the system manager's pricing strategies, and other performance metrics. We further explore the differences in these implications. Given that customers' DLA behavior discourages customers from joining the system, common intuition suggests that the manager should reduce the price to attract more customers into the system. However, we observe that the revenue-optimal price may actually rise as customers' DLA level increases. In contrast, we demonstrate that customers' RLA behavior has no impact on their decisions, the system manager's pricing strategies, or other performance metrics in the context of revenue maximization, despite that RLA customers experience disutilities. From a welfare-maximization perspective, in the presence of DLA and RLA behaviors, we demonstrate that interestingly, there may be no efficiency loss when the system manager implements the revenue-maximization pricing decision. Compared to revenue maximization, the system manager should consistently raise prices to mitigate the effects of customers' loss aversion. In queueing service systems, it is crucial for the system manager to accurately identify the customers' loss aversion behaviors when determining prices.

2 - A Strong Approximation Approach for the Asymptotic Variability Analysis for Queue

Yongjiang Guo, Beijing University of Posts and Telecommunications, Beijing, China, People's Republic of, Juan Yang

In this talk, we present a strong approximation (SA) approach for the asymptotic variability analysis for queues, which is characterized by the law of the iterated logarithm (LIL) and the functional LIL. For a given stochastic process, both the LIL and the functional LIL tell us the maximal value of the fluctuations around its fluid approximation process, which is obtained by the law of the large numbers in general. We apply the SA method to the GI/G/1 queue and its some generalized queueing networks, and can find the asymptotic variability for the following performance measures under some conditions: the queueing length, workload, busy time, idle time and departure processes, etc. However, there are many situations the SA method does not work. Through a tandem queue, we point out the difficulty that the SA method is used to deal with the LIL of the queue length process.

TD02

Mercury

Aviation Topics III

Invited Session

Aviation Applications

Chair: Darryl Teo, Singapore University of Technology and Design, Singapore, Singapore

1 - World Airline Safety: Darker Days Ahead?

Arnold Barnett, Massachusetts Institute of Technology, Cambridge, MA, United States

Events that cause passenger deaths on scheduled flights worldwide are at the brink of extinction: a traveler who took one flight each day could on average go for 68,000 years before succumbing to a fatal crash.. But many a near-extinct species has staged a major comeback. Through scrutinizing relevant data for the years 1946-2025, we explore how air-passenger mortality risk has evolved over time and why it has done so. Among other things, we take note of the fact that flying in some countries has always been orders of magnitude safer than in others. We also offer data-based reasons that flying in the future might not be as safe as flying today, meaning that the golden age of aviation safety could be coming to an end.

2 - Product Line Design Under a Context-Dependent Mnl Model: An Application to Airline Fleet Assignment and Timetabling

Keji Wei, Tongji University, Shanghai, China, People's Republic of, Yufeng Cao, Xiaoyun Niu

Flight timetabling impacts an airline's operating profit. We consider a joint flight timetabling and revenue management problem with passengers choosing among multiple itineraries. The classic multinomial logit (MNL) choice model does not well depict the interaction between parallel itineraries. So, we introduce a context-dependent MNL model where the attraction value of one product is contingent on the offer set. The context-dependent MNL model extends common choice models and provides alternative explanations to several empirical observations. We formulate the joint flight timetabling and revenue management problem as a mixed integer programming model. We present a two-stage decomposition framework to compute the network timetable within a practical time budget. Numerical results suggest that the combination of the formulation and the solution approach leads to profit improvement for airlines.

3 - Beyond the Average: Uncovering Grading Behaviors with Data Analytics

Yuanzhe Ying, National University of Singapore, Singapore, Singapore, Jussi Keppo, Hong Ming Tan

As part of a broader initiative to enhance pilot training through data-driven interventions, this study focuses on understanding grading patterns and improving the grading informativeness of pilot instructors. Under the widely adopted Competency-Based Training and Assessment (CBTA) framework, instructor grading frequently clusters around average scores, resulting in limited variance. This reduces the ability to differentiate pilot performance levels and increases the risk of overlooking early signs of underperformance. Applying explainable machine learning frameworks such as SHAP, we analyze pilot training data to identify key factors that may influence instructor grading behavior. A randomized controlled trial is currently underway in collaboration with a major airline to evaluate the effectiveness of a redesigned instructor workshop informed by these insights in an operational setting.

4 - Estimating Benefits of Information Sharing and Collaboration Between Countries, Airports and Airlines

Darryl Teo, Singapore University of Technology and Design, Singapore, Singapore, Peter Jackson, Nuno Ribeiro, RAKESH NANDI, Gengling Dai

Air Traffic Flow Management (ATFM) tactical measures such as the imposition of acceptance rates and calculated takeoff times require information sharing and collaboration between Air Traffic Control sectors. In a decentralized system such as Southeast Asia this is a challenge. To evaluate the potential benefits of different measures under different information sharing regimes, it is useful to construct a rolling horizon simulation of the decentralized system. An essential component of such a simulation is a model of each Flight Information Region (FIR) as a separate decision-making entity. We formulate the FIR decision problem as a discrete event simulation using First-Come-First-Served logic, which is to be solved for each FIR at each time step of the simulation using

the local information available to each FIR at that time. We contrast this with a collaborative decision making framework modelled as First-Scheduled-First-Served with shared calculated take-off times (CTOT). Using this model, we estimate the benefits of using information sharing and collaborative ground delays under different levels and scope of participation, including collaboration between regions, airports, and airlines.

TD03

Minto

Robust Prescriptive Analytics

Invited Session

Computing Society

Chair: Li Chen, The University of Sydney, CHIPPENDALE, Australia

1 - Robust Drone Delivery with Weather Information

Chun Cheng, Dalian University of Technology, Dalian, China, People's Republic of, Yossiri Adulyasak, Louis-Martin Rousseau

Drone delivery has garnered significant attention recently due to its potential for faster delivery at a lower cost relative to other delivery options. When scheduling drones from a depot for delivery to various destinations, the dispatcher must consider the uncertain wind conditions, which affect the delivery times of drones to their destinations. To mitigate the risk of delivery delays caused by wind uncertainty, we propose a two-period drone scheduling model to robustly optimize the delivery schedule. In this framework, the scheduling decisions are made in the morning, with provision for different delivery schedules in the afternoon that adapt to updated weather information available by midday. Our approach minimizes the essential riskiness index, which can limit the probability of tardy delivery and the magnitude of lateness. Using wind observation data, we characterize the uncertain flight times via a cluster-wise ambiguity set, which has the benefit of tractability while avoiding overfitting the empirical distribution. The cluster-wise ambiguity set enables us to adapt the intraday delivery schedule depending on which cluster on the wind vector chart the observed morning wind vector belongs to. A branch-and-cut algorithm is developed for this adaptive distributionally framework to improve its scalability.

2 - Adaptive Emergency Operations: A Three-Stage Robust Optimization Model

Qinghe Sun, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, Li Chen, Tao Feng, Lingxiao Wu

Emergency operations inherently involve the sequential realization of uncertainties, as disruptions and demand fluctuations materialize at different stages. Traditional two-stage optimization models struggle to capture adaptive decision-making across these evolving uncertainties, limiting their effectiveness in time-sensitive responses and proactive preparedness. This study introduces a *three-stage* robust optimization model to improve emergency preparedness, adaptive planning, and response under sequential disruptions and demand ambiguity. The model addresses the *unpredictability* of emergency events, such as sudden facility closures, using a discrete budget uncertainty set, and demand ambiguity using a mean-covariance ambiguity set. The three-stage structure enables decision-makers to adjust actions as uncertainties unfold. However, its *min-max-min-max-min* formulation presents significant computational challenges. To address this, we develop a solution approach based on backward analysis. First, we approximate the inner max-min structure using quadratic recourse adaptation, reformulating the problem as a mixed-integer two-stage robust optimization model with a conic recourse function and binary uncertainty sets. To further enhance tractability, we design an exact *branch-check-and-Benders-cut* algorithm, integrating Benders decomposition within a branch-and-check framework, enhanced by customized acceleration techniques. Experiments on real-world instances confirm the computational efficiency of our approach, outperforming the widely used column-and-constraint generation method. Using a humanitarian food relief case in Malawi, we benchmark our model against a two-stage counterpart, demonstrating superior cost efficiency and reduced budget overrun risks. These results underscore the value of explicitly modeling sequential uncertainties in emergency operations.

3 - Explainable Data-Driven Share-of-Choice Product Line Design Optimization

Maoqi Liu, Shandong University, Jinan, China, People's Republic of

The share-of-choice (SOC) problem is a widely studied problem for product line design (PLD), where representative customers are sampled from a target population, and the percentage of customers who choose the offered products over outside options is maximized. The individual choices are captured by the utility maximization framework, where utilities are associated with the products and a customer will choose the one with the highest utility. Since these utilities are not directly observable, they must be estimated from other primitive data such as those from conjoint analysis. With the collected data, various methods, such as the hierarchical Bayesian and polyhedral estimation methods, are employed to estimate the utilities. However, this “estimate-then-optimize” procedure disconnects the decision-making process from the primitive data and may suffer from model misspecifications. To fill the gap, we integrate the polyhedral estimation method, which constructs a polyhedron set containing all utilities consistent with the data, into the SOC problem by maximizing the share-of-choice calculated by the worst-case utility over the polyhedron. The integration brings twofold advantages compared to the classic ``estimate-then-optimize'' approach. First, the model provides a way to attribute solution to the primitive data. Second, it allows quantifying the influence of the number of sampled customers as well as that of estimation accuracy on the out-of-sample performance guarantee. Through extensive numerical results, we show the effectiveness of the proposed model and adaptive survey design approach.

4 - Joint Optimization of Maintenance and Repositioning for Dockless Bike-Sharing Systems

Qiuzhuang Sun, University of Sydney, Sydney, Australia

We study data-driven management of dockless bike-sharing systems (BSSs), which have two main challenges: (i) unusable bikes due to use and vandalism; (ii) mismatch between bike supply and travel demand. A fleet of trucks can transport damaged bikes to repair stations and reposition functional bikes to better serve demand. We optimize the truck routes to jointly solve both problems. We develop a novel data-driven approach that captures the dynamics of BSSs. Our approach is highly scalable, as demonstrated on a real network with approximately 4,000 nodes. We partition the urban area into disjoint regions in a data-driven manner from the spatial distribution of bikes. A stochastic network flow model is formulated, based on which the bike failure probabilities can be estimated as new data become available. The optimal bike allocation scheme is then determined according to the demand data. We synthesize the previous components to formulate a routing problem for joint maintenance and repositioning. We show that bike failure probabilities can be approximately computed recursively. We also show that the estimated bike failure probability converges to one almost surely after a bike fails, with a linear convergence rate that depends on the demand distribution in the region where the bike is parked. Using data from one of the largest BSSs in Asia, we show the proposed methodology outperforms several benchmarks. Extensive numerical experiments are conducted to explore managerial implications such as the choice of fleet size and truck capacity for repositioning.

TD05

Orchard A

Advancing Socially Responsible Operations: Fairness, Privacy, and Social Interactions

Invited Session

Efficiency and Fairness in Supply Chain and Service Operations Management

Chair: Yuanzheng Ma, Shanghai, N/A

1 - An Algorithmic Approach to Managing Supply Chain Data Security: the Differentially Private Newsvendor

Du Chen, Nanyang Technological University, Singapore, Singapore, Geoffrey Chua

Data is now unanimously considered a key firm asset for enabling better operational decisions. However, data-driven decisions can inadvertently expose private data, leaving firms vulnerable to unforeseen danger. How to manage data security risks by protecting data from being inferred from observable decisions thus becomes an important question. In this paper, we focus on data security in supply chains due to their data-intensive nature. Specifically, we examine a data-driven contextual newsvendor problem. To quantify and ensure data security, we adopt the notion of differential privacy, a mathematically rigorous measure of data security that limits an attacker's inference accuracy. Employing convolution smoothing and noise injection, we propose several

differentially private algorithms that provably guarantee both data security and asymptotic optimality with (near) optimal rates. In the non-asymptotic regime, we further identify three drivers of the cost of data security; namely, dataset size, context, and number of products. This finding suggests that gathering more data, collecting detailed context, and pooling data from multiple products can lower data security costs. Lastly, we examine the impact of a newsvendor's private algorithms on supply chain partners. We discover additional distortion to the demand signaling process and lower profit share for an upstream supplier.

2 - Achieving Ex-Ante and Ex-Post Fairness: Design of Allocation Rules for Mobile Food Pantries

Yuanzheng Ma, Huazhong University of Science and Technology, Wuhan, China, People's Republic of, Jinpeng Liang, Guodong Lyu

Mobile food pantries (MFPs) have increasingly expanded in recent years to help alleviate world hunger by allocating surplus food to underserved communities that face challenges in accessing conventional food banks. As non-profit programs, MFPs prioritize operational efficiency and minimizing food waste. However, this approach can result in unequal distribution of services among beneficiaries from different communities. To address this issue, we design a sequential food allocation framework that focuses on serving the sequentially revealed demand of each targeted community during the food delivery trip. This framework aims to maximize the expected minimum fill rate across communities (ex-post fairness criteria), while simultaneously meeting the predetermined expected fill rate target (ex-ante fairness requirement) for each community.
We show that this particular type of food allocation problem, constrained by the ex-ante fairness considerations, can be simplified by solving a series of stochastic dynamic program problems with randomized coefficients in the objective function. The online convex optimization tool and Fenchel duality are applied for analysis. Furthermore, we show that the optimal solution to the stochastic dynamic problem, under the ex-post fairness maximization criteria, exhibits a two-threshold structure. This structure explicitly reveals the trade-offs between achieving optimal performance in ex-post fairness while ensuring fair food delivery service to all communities in an ex-ante manner. Extensive numerical experiments with both synthetic data and real-world data showcase the promising benefits of our allocation rules over existing benchmarks.

3 - When Workload Overburdens: An Empirical Study of Multi-Task Effort Allocation in Livestream E-Commerce

Xinmeng Xu, Tianjin University, Tianjin, China, People's Republic of, Xiaojing Feng, Ning Jia

Livestream e-commerce has experienced prosperous growth, with streamers playing a pivotal role. This research explores the relationship between streamer workload and sales performance in livestream e-commerce. Utilizing nine months of granular minute-level data of over one million observations from five thousand livestream sessions, we employ the fixed effect models with the instrumental variables to address endogeneity concerns. We uncover an inverted U-shaped workload-revenue effect: moderate workload boosts revenue through exerting more efforts on both product promotion (up-selling and wide-selling) and interaction facilitation (lower churn rate and greater fans growth), while excessive workload increases efforts on wide-selling and interaction but reduces up-selling, diminishing overall revenue. Heterogeneity analyses indicate that the lucky draws weaken low-level workload benefits while amplifying high-level drawbacks, whereas the involvement of streamer assistant mitigates the negative effect of a high-level workload. Robustness checks and managerial implications are included.

TD06

Orchard B

Operational Strategies and Risk Management

Contributed Session

Game Theory and Applications

Chair: Xiaoning Shi, German Aerospace Center, Berlin , N/A, Germany

1 - Research on the Risk Strategy of Electric Vehicle Manufacturers Considering Recycling to Deal with the Risk of Supply Disruption

Yike Li, University of Science and Technology Beijing, Beijing, China, People's Republic of

This study addresses electric vehicle (EV) manufacturers' supply chain resilience against power battery shortages through an integrated circular economy strategy. We propose a novel operational framework combining dual-source procurement, horizontal patent competition, product innovation, and retired battery recycling. Developing a stochastic mathematical model that distinguishes between regular operations and supply disruption scenarios, we analyze how critical factors—including supply discontinuity, product differentiation, price elasticity, and battery recovery rates—influence manufacturers' operational decisions.

Three results emerge. First, during stable supply conditions, primary suppliers maintain the lowest optimal pricing while backup channels exhibit price parity at premium levels, regardless of total market demand. Second, under disruption scenarios, the comprehensive strategy demonstrates significant risk mitigation effectiveness, with backup suppliers' order quantities and pricing proving sensitive to both recovery rates and price elasticity coefficients. Third, probabilistic analysis shows manufacturers strategically absorb higher supply risks through increased primary orders when facing substantial product quality differentials between suppliers.

The model advances sustainable operations research by formalizing patent-competitive recycling dynamics within stochastic dual-sourcing systems. Our results suggest that tiered pricing strategies coupled with recovery rate optimization can simultaneously enhance supply chain robustness and circular economy performance.

2 - Research on Financing Decision-Making of Dual-Channel Supply Chain Considering Capital Constraints and Risk Aversion Under Cvar Standards

Rui Shao, University of Science and Technology Beijing, Beijing, China, People's Republic of

This paper considers the dual-channel supply chain where supply chain members are risk-averse and retailers have capital constraints, and retailers can alleviate the problem of capital constraints through three financing models: deferred payment financing, bank loan financing, and equity financing, and use the CVaR criterion to construct a dual-channel supply chain game model with capital constraints, analyze the impact of key factors such as risk aversion and interest rate under different financing modes on the optimal pricing decision, and compare the profits of retailers in the three financing models to obtain the optimal financing strategy of retailers. The results show that when the retailer has sufficient funds or the retailer adopts deferred payment and bank loan financing, the retailer's retail price, the manufacturer's direct selling price and the wholesale price are positively correlated with the manufacturer's risk aversion range, the retailer's retail price is positively correlated with the retailer's risk aversion range, and the manufacturer's wholesale price is negatively correlated with the retailer's risk aversion range. Comparing the profits of retailers under the three models, the equity financing model is the highest, followed by the deferred payment financing model, and the bank financing model is the lowest.

3 - Macroeconomic Implications of ESG Investing in Financial Markets: Information Distortions and Mutual Learning

Ge Zhou, Zhejiang University , Hang Zhou, China, People's Republic of, Liyan Yang

This paper explores the macroeconomic implications of information distortions induced by ESG (Environmental, Social, and Governance) investing in financial markets when mutual learning occurs between financial markets and the real economy. We develop a tractable dynamic macroeconomic framework that captures the interdependence of information between these two sectors. While both traditional traders and ESG traders receive identical signals about monetary and non-monetary shocks to financial assets, their heterogeneous preferences lead them to trade in opposite directions and make the price noisier to each other. As a result, multiple financial equilibria emerge, accompanied by self-fulfilling aggregate economic uncertainty. Based on this dynamic macroeconomic framework, we examine the macroeconomic implications of these information distortions from both a static comparative perspective and a business cycle perspective.

4 - Analytical Framework on Investing Future (Port) Infrastructures

Xiaoning Shi, German Aerospace Center, Berlin , Germany

This research is investigating an analytical framework by deploying game theory, and the framework assumes that a bunch of investors aim to maximize the expected profits and potential benefits of the (port) infrastructures, given limited budgets over the time horizon and provided estimated risks.

TD07

Orchard C

AI for Disease Prediction and Precision Medicine

Invited Session

Health Systems Modelling with AI

Chair: Ahmadreza Pourghaderi, National University of Singapore, Singapore, Singapore

Co-Chair: Ahmadreza Pourghaderi, National University of Singapore, Singapore, Singapore

Co-Chair: Arul Earnest, Monash University, 553 St Kilda Road, Melbourne 3004, Victoria, Australia, Victoria, 3004, Australia

1 - Ai in Clinical Registries

Arul Earnest, Monash University, Victoria, Australia

Background:

Clinical quality registries are rich sources of real-world data that can be harnessed to improve patient outcomes through advanced analytics. Artificial intelligence (AI), particularly machine learning (ML) techniques, offers powerful tools to uncover patterns and predict risks in complex clinical scenarios. This presentation explores the application of supervised ML models to predict key outcomes in two Australian registries: the Victorian Lung Cancer Registry and the Australasian Diabetes Data Network.

Methods:

In the lung cancer study, data from 11,601 patients across 40 hospitals in Victoria (2011–2016) were used to predict adherence to four timeliness-of-care indicators. ML models including random forests, neural networks, and support vector machines were compared against traditional logistic regression using out-of-sample AUC metrics.

In the diabetes study, data from 13,761 individuals with type 1 diabetes across Australia and New Zealand were used to predict diabetic ketoacidosis (DKA) and poor glycaemic control (HbA1c >7%). Nine ML models, including Fast Large Margin and Deep Learning, were evaluated.

Results:

The random forest model achieved an AUC of 0.98 for all lung cancer timeliness indicators, outperforming logistic regression (AUC=0.72). For diabetes outcomes, the Fast Large Margin model predicted DKA with an AUC of 0.859 and perfect sensitivity, while Deep Learning predicted HbA1c >7% with an AUC of 0.863.

Conclusion:

AI techniques applied to large-scale clinical registries demonstrate strong predictive capabilities and have the potential to identify high-risk patients early, enabling targeted interventions and improving care delivery in real-world registry settings.

2 - Predicting Postoperative Outcomes in Rhegmatogenous Retinal Detachment (Rrd) Surgeries

Mohammad amin Honardoost, Dep Public health and Preventive medicine, Monash University, Melbourne, Australia, Arul Earnest, Ahmadreza Pourghaderi

Background:
Rhegmatogenous retinal detachment (RRD) is a vision-threatening condition that often require surgical intervention. Despite modern surgical techniques yielding high anatomic and visual success, yet outcomes may vary considerably depending on patient-specific and surgical factors. Developing an accurate preoperative risk

model is vital to optimize surgical planning, improve patient counselling, and facilitate personalized treatment strategies, while also enabling robust performance benchmarking across surgeons and institutions.

Methods:
In this non-interventional, multinational, multicenter registry study, data was extracted from the Australian and New Zealand Society of Retinal Specialists Retinal Surgery Registry for patients undergoing primary RRD repair. Strict data-quality assessments and imputation methods addressed inconsistencies, outliers, and missing values. Various machine learning (ML) algorithms—including generalized linear models, random forests, support vector machines, gradient boosting, and neural networks—were implemented via the mlr3 package in R. An integrated approach combining pre-processing and feature selection was applied with 5-fold cross-validation to optimize predictive performance for both primary (retinal re-detachment) and secondary (change in best-corrected visual acuity letter score at 3 months) outcomes.

Results:
Total of 9,394 primary RRD surgeries recorded into ANZSRS registry from 2000 to 2023 were included into study. Key preoperative predictors were identified, and the ML models consistently outperformed traditional regression benchmarks as measured by metrics such as AUC-ROC, F1-score, mean squared error, and R^2.

Conclusion:
ML-driven predictive modelling offers a promising clinical decision-support tool to enhance personalized surgical planning and improve postoperative outcomes in RRD surgeries.

3 - Predicting Cardiorenal Complications in Diabetes Using Machine Learning

Ahmadreza Pourghaderi, Monash University, Melbourne, Australia, Hanieh Jamshidian, Hossein Nejati, Arul Earnest, Ella Zomer, Sophia Zoungas

Machine learning (ML) techniques were applied and compared to predict cardiorenal complications in adults with diabetes, utilising data from the Australian National Diabetes Audit (ANDA). Clinical and demographic data from 16,494 patients, collected between 2015 and 2022, were analysed. Several robust ML models, including Random Forest, Bagging, and Gradient Boosting, were trained and evaluated to identify the most effective approach for clinical prediction. Feature selection techniques were employed to enhance model interpretability and highlight the most predictive clinical indicators. Performance was assessed using metrics such as accuracy and F1-score to ensure both reliability and clinical relevance. Exceptional predictive performance was demonstrated by the Random Forest and Bagging models, with over 98.0% of patients with or without cardiorenal complications being correctly identified. After further optimisation, F1-scores of up to 98.9% were achieved. Urinary protein levels, creatinine, and antihypertensive therapy were identified as key predictive features, aligning with known clinical risk factors. The value of interpretable ML approaches in supporting early identification of high-risk individuals and enhancing decision-making in diabetes care was illustrated by these findings. The integration of such models into clinical workflows is expected to improve patient outcomes and optimise healthcare delivery.morbidities, thereby aiding in accurate diagnosis and personalized treatment planning.

Despite these advancements, several challenges hinder the widespread clinical implementation of ML in differentiating psychosomatic from organic diseases. These include the necessity for large, high-quality datasets, potential algorithmic biases, interpretability of ML models, and ethical considerations surrounding patient data privacy.

4 - Leveraging Machine Learning to Differentiate Psychosomatic Disorders from Organic Diseases: Concepts, Applications, and Challenges

Ahmadreza Pourghaderi, National University of Singapore, Singapore, Singapore, Mohammad Reza Sharbafchi, Farzad Goli

Accurately distinguishing psychosomatic disorders from organic diseases remains a significant challenge in clinical practice due to overlapping symptoms and complex etiologies. Recent advancements in machine learning (ML) offer promising avenues to enhance diagnostic precision in this domain. This presentation

provides an overview of ML concepts pertinent to this differentiation, examines current applications, and discusses associated challenges.

Machine learning algorithms, including support vector machines, random forests, and deep neural networks, have been employed to analyze diverse datasets encompassing patient demographics, clinical symptoms, imaging, and biomarkers. These algorithms can identify subtle patterns and correlations that may elude traditional diagnostic methods. For instance, in gastroenterology, ML techniques have been applied to differentiate Irritable Bowel Syndrome (IBS), a common psychosomatic disorder, from organic gastrointestinal (GI) diseases based on symptom profiles. Studies have demonstrated that ML models can effectively classify IBS patients by analyzing clinical symptoms and psychological co-morbidities, thereby aiding in accurate diagnosis and personalized treatment planning.

Despite these advancements, several challenges hinder the widespread clinical implementation of ML in differentiating psychosomatic from organic diseases. These include the necessity for large, high-quality datasets, potential algorithmic biases, interpretability of ML models, and ethical considerations surrounding patient data privacy. Addressing these challenges requires collaborative efforts among data scientists, clinicians, and ethicists to develop transparent, reliable, and ethically sound ML applications.

This presentation aims to elucidate the current landscape of ML applications in differentiating psychosomatic disorders from organic diseases, highlighting both advancements and areas necessitating further research.

TD09

Bras Basah B

Advanced MCDM Methodologies and Theoretical Innovations

Contributed Session

Multiple Criteria Decision Making

Chair: xiaoqing Hou, University of Science and Technology Beijing, Beijing, N/A

1 - Advancing Smart Healthcare: An Integrated Decision Programming for Aiomt Business Model Portfolio Optimization Using MCDM and Activity-Based Life Cycle Costing

CHIH-HAO YANG, Ming Chuan University, Taipei, Taiwan

Digital innovation is the critical driver for the transition of smart healthcare to reach digital innovation and realize health system sustainability development. In particular, AIoMT innovative business model measurement planning is necessary to satisfy healthcare quality assurance. This study focuses on the critical indicators for the AIoMT innovative business model decisions and priority alternatives from the Digital Governance, Industrial Symbiosis, Financial Sustainability, and Artificial Intelligence Technology views. Further, it incorporates Activity-Based Costing techniques and resource programming into the multi-criteria decision-making (MCDM) approach, which integrates the Decision Making Trial and Evaluation Laboratory (DEMATEL), the Analytic Network Process (ANP), the VlseKriterijumska Optimizacija I Kompromisno Resenje (VIKOR), and Zero-One Goal Programming (ZOGP) to achieve optimal AIoMT innovative business model portfolio decisions. The results show that Industrial Symbiosis and Artificial Intelligence Technology play an essential role in the decision process of the AIoMT innovative business model. Moreover, Government Regulation is crucial for evaluating an AIoMT innovative business model. Specifically, optimal AIoMT innovative business model portfolios, including Electronic medical and healthcare records, Biomedical Health Data Platform, Emerging Energy Recycling Mechanism, and Medical Waste Recycling Management, can prove resource requirement feasibility and reasonability. This study provides insight into various MCDM approaches, considering AIoMT resource requirement planning over traditional decision processes. This study provides valuable scientific guidance for optimizing AIoMT innovative business model planning while benefitting Taiwan's smart healthcare industry transformation for sustainable development.

2 - Ceo Real Option Mindset and Managerial Styles

Oupin Tang, Xi'an Jiaotong-Liverpool University, Suzhou, China, People's Republic of, Tarik Driouchi, Lenos Trigeorgis

We develop a firm–year–level composite indicator of CEOs' real-options mindset (ROM) using earnings-call transcripts and investigate its empirical implications. By capturing CEOs' real-options reasoning across three behavioral dimensions—analytical thinking, forward-looking orientation, and semantic adaptability—ROM is positively associated with higher managerial ability and risk aversion; greater investment in option-embedded intangible assets; lower dividend payouts; more efficient investments; increased organisational flexibility; higher return skewness; more diverse and independent boards; and lower relative CEO pay.

3 - A Novel Causal Model: Verification-Based Dematel Type I and Type Ii

Shan-Lin Huang, National Taiwan Normal University, Taipei, Taiwan, Chin-Yi Fang, Lilian Liang, Shiyu Zhao, Hanhan Chang, Gwo-Hshiung Tzeng

The DEMATEL method is a key technique in MCDM that identifies causal relationships among criteria, enabling systematic and visualized analysis. It is widely applied in the social sciences due to its ability to generate reliable insights with relatively low data collection costs by leveraging expert judgment. However, the traditional DEMATEL method presents several limitations in the data collection and analysis process. First, an excessive amount of information regarding interrelationships among criteria may hinder users' decision-making. Second, expert opinions may not fully reflect the perspectives of the general population. Third, the traditional approach fails to distinguish between positive and negative causal relationships. To address these issues, this study integrates two extended models previously proposed in the literature, termed "Verification-based DEMATEL Type I" and "Verification-based DEMATEL Type II". Considering that DEMATEL is essentially an exploratory analytical tool, these two new models retain its strength in domains lacking strong theoretical foundations, while enhancing its structural verification capabilities—thus the designation as "verification-based" models. The main distinction between the two lies in their data sources: Type I is based on secondary databases and emphasizes objective structural verification, whereas Type II relies on questionnaire surveys and incorporates public perspectives into the validation process. Ultimately, both models enhance external validity and effectively overcome the limitations of traditional DEMATEL in terms of information overload and the inability to identify the directionality (positive or negative) of causal relationships, offering a more robust tool for causal analysis.

4 - A New Cross-Efficiency Aggregation Method Considering Decision Maker's Preference

xiaoqing Hou, University of Science and Technology Beijing , BEIJING, China, People's Republic of, feng He

Abstract: The cross-efficiency evaluation model is an extension of the classical Data Envelopment Analysis (DEA) model. The main difference with other models is that the cross-efficiency evaluation takes the self-evaluation value and the other-evaluation value as the final efficiency score, and then sort the decision making units (DMUs) completely. However, traditional cross-efficiency models face challenges such as non-unique weight determination and failure to achieve Pareto optimal. In addition, the simple average aggregation of cross-efficiency values ignores the relative weights of DMUs,leading to loss of decision-making information and questionable credibility of results.This paper focuses on the evaluation process of cross-efficiency, reveals the internal evaluation process of cross-efficiency based on the perspective of attribute utility characteristics, and the psychological characteristics of decision-makers (DMs) when they face risk decision-making, to extend the cross-efficiency aggregation method from a new perspective. This paper first uses the entropy weight method to aggregate the index weights according to the consistency of other-evaluated attributes to generate a set of common weight models. Then building on the TODIM method's principles, this study treats DMUs as DMs, calculates their relative dominance scores through risk attitude integration, and ultimately achieves effective evaluation and precise ranking of DMUs. Finally, the improved cross-efficiency is used to evaluate the efficiency of scientific and technological innovation in 30 provinces of our country, This paper provides a new evaluation method for the empirical analysis of provincial science and technology innovation efficiency in our country.

Keywords: DEA cross-efficiency, TODIM method, Relative advantage

TD10

Hullet Room

Renewable Energy and Sustainable Technologies

Contributed Session

Energy, Natural Resources and the Environment

Chair: Kim Haneul, Korea University, Seoul, Korea, Republic of

1 - Fostering Green Technology Innovation Through Green Knowledge Management : the Roles of Dynamic Capabilities, Green Organizational Identity, and Incentive Environmental Regulation

Wen yi Luo, Huazhong University of Science and Technology, Wuhan, China, People's Republic of, Zongjun Wang

Green knowledge management, as an essential strategic asset for firms, provides relevant knowledge resources for green technology innovation. This study draws upon the dynamic capabilities theory and organizational identity theory to develop a moderated mediation model that explores the influence of green knowledge management on green technology innovation through dynamic capabilities, green organizational identity, and incentive environmental regulation. The study's hypotheses are tested using survey data from 358 Chinese firms. Hierarchical regression and PROCESS are employed to validate the research model. The findings confirm a significant relationship between green knowledge management and dynamic capabilities, and also between dynamic capabilities and green technology innovation. Specifically, absorptive capability and transformative capability of dynamic capabilities function as key mediators in this relationship. Moreover, green organizational identity positively moderates the relationship between green knowledge management and dynamic capabilities, while incentive environmental regulation positively moderates the relationship between dynamic capabilities and green technology innovation. This study provides new insights into the relationship between green knowledge management and green technology innovation, thereby enriching the existing literature and offering practical implications for policymakers and firm practitioners.

Keywords Green technology innovation, Green knowledge management, Dynamic capabilities, Green organizational identity, Incentive environmental regulation,Chinese firms, sustainability

2 - Safer and More Resilient Hurricane Insurance: Insurability, Government Mitigation, and Multi-Region Pricing Strategies

Dahui Liu, Donghua Universiy, Shanghai, China, People's Republic of, Linda Nozick, Meghan Millea, Jamie Kruse, Rachel Davidson, Junkan Li

Hurricanes impose substantial harm on homeowners, not only through physical destruction but also through long-term financial burdens stemming from increased insurance premiums, declining property values, and repair costs. This study addresses two key challenges: first, finding the interrelated decisions of the government mitigation funding of residential acquisitions and retrofit subsidies and of price restrictions on the insurance market in eastern North Carolina to determine the financial effects on stakeholders; second, examining how to effectively price hurricane insurance since those prices must be constructed to match large sporadic outlays in spatially compact areas with many relatively small premium payments. To address these issues, we develop an integrated framework, the Stakeholder-Based Tool for the Analysis of Regional Risk, which allows us to identify cost-effective levels of government spending that can significantly reduce residential losses from hurricane damage. Additionally, we explore multi-region pricing strategies within an oligopolistic hurricane insurance market using a heuristic-based search method to determine equilibrium prices under complex demand and cost conditions. Our findings show the tradeoff between maintaining insurer solvency and expanding homeowner insurance accessibility. Results also indicate insurers often offer lower rates per dollar of expected loss in the most vulnerable coastal areas. Moreover, increased market competition—through the entry of new insurers—leads to reduced policy prices and higher coverage uptake. While this expanded coverage benefits homeowners, it may also place financial pressure on insurance providers.

3 - Power Restoration After Natural Disasters Based on Discrete-Event Optimization

Kim Haneul, Korea University, Seoul, Korea, Republic of, Taesu Cheong

The increasing frequency of natural disasters is causing a significant impact on power infrastructure, leading to prolonged outages and substantial economic losses. Existing power restoration methods primarily focus on network reconfiguration, repair crew scheduling, and mobile power source dispatching optimization. However, most studies use fixed time interval-based decision models, which require a predefined discrete time space, often resulting in unnecessary decision variables and increased computational burden. To address these limitations, this study proposes a Mixed-Integer Linear Programming based discrete event co-optimization model for network reconfiguration, repair crew scheduling, and mobile power source dispatching. The proposed method makes decisions based on the occurrence of events, such as the completion of facility repairs. Moreover, it considers realistic human resource operations by reflecting the off-hour scheduling for repair teams, which enhances the feasibility of long-term recovery planning. To validate the proposed method, we conduct experiments comparing it with an existing fixed time interval-based decision model. The analysis also examines how incorporating off-hours scheduling affects the optimal recovery plan. Results show that this approach yields higher-quality solutions than the existing model. Moreover, the optimal plan is influenced by whether off-hours constraints are taken into account.

TD11

Morrison Room

Pricing and Revenue Optimization in the Renewable Energy Industry

Invited Session

Revenue Management and Pricing

Chair: Lun Ran, Beijing Institute of Technology, N/A

Co-Chair: Zhiyuan Wang, Tsinghua University, Beijing, N/A

1 - Citywide Orchestration of Electrical Vehicle Battery Swapping

Xian Guo, Beijing Institute of Technology, Beijing, China, People's Republic of, Zhiyuan Wang, Lun Ran, Wei Qi

To efficiently integrate the operation of electric vehicles (EVs) and battery swapping stations, we propose an orderly swapping strategy in which EVs carry out battery swaps at designated swapping stations in response to swapping demands, with real-time scheduling managed by the service platform. Based on this strategy, we characterize the battery swapping station location and capacity planning problem as a two-stage distributionally robust optimization model. This model presents challenges with bilinear terms involving decision variables and uncertain variables, and sums of piecewise linear functions in the second-stage problem. Notably, we demonstrate that the second-stage problem is convex concerning the uncertain variable and reveals a distinctive structure in the optimal solution of its dual problem. These model properties guide us in designing an efficient column-and-constraint generation algorithm for solving it. We also propose two alternative strategies designed for specific scenarios, reformulating them as mixed-integer second-order cone programmings. In a numerical study, we validate our algorithm's efficacy using classic transportation data and compare the performance of the three strategies with real transportation data from Beijing. Our findings indicate that the Orderly strategy surpasses the other two strategies regarding flexibility and demonstrates economies of scale as the battery swapping scale increases. This flexibility becomes more pronounced as the service rate decreases and the grid capacity increases. Furthermore, the enhanced grid capacity results in an imbalanced battery distribution and the significant technological advancement eliminates the possibility of battery shortages.

2 - Toward Resilient Green Cloud Computing: Joint Operations of Energy Storage and Spatial Task Allocation

Xiaoxin Xie, Beijing Technology and Business University, Beijing, China, People's Republic of

In recent years, the surge in advanced internet computing workloads in data centers, driven by applications like large language models and the Internet of Things, has intensified the challenge of ensuring energy efficiency while maintaining stable and resilient computational services. In practice, commonly used integrated centralized work scheduling and energy management systems, which aim to optimize power distribution and computational

efficiency, face significant sociotechnical challenges that hinder their effective implementation. To address these issues, we integrate reserve regulation and task allocation to further optimize the operations of multiple data centers in a computational resource-sharing platform. We model the problem as a two-stage stochastic integer program. In the first stage, we optimize backup battery capacity for each data center, and in the second, we introduce a non-preemptive M/M/1 queue with task priorities. We then analyze stationary processing times and apply second-order cone programming for improved computational tractability. We develop an outer-approximation algorithm to improve computational efficiency for large-scale problems, while our integrated strategy balances cost efficiency, sustainability, and resilience at minimal costs. A case study demonstrates that the integrated strategy for multiple data centers reduces total costs by 10.39% and 10.88% compared to task allocation and backup-battery reserve regulation strategies, respectively. Task priorities save 3-4% in operational costs, while the strategy ensures stable operations during power interruptions. The outer-approximation algorithm outperforms commercial solvers by 30%, and task replication improves energy efficiency in smaller data centers. These findings highlight the potential of our strategy to enhance data center efficiency and sustainability.

3 - Assessing Economies' Resilience of International Liquefied Natural Gas Trade Network in the Presence of the Ripple Effect

Min Zhu, Beijing Institute of Technology, Beijing, China, People's Republic of

In the International Natural Gas Trade Network (INGTN), global trade disruptions are caused by fluctuations in major exporters' supply quantities, with economies' capacities for resistance and recovery being crucial. The resilience of economies within the INGTN is assessed using an indicator system based on robustness and recoverability. Economies are classified into four groups according to their robustness and recoverability thresholds. The results of simulations reveal that import volumes decrease significantly after the fifth step, with economies demonstrating higher robustness and recoverability experiencing smaller fluctuations. Economies with both high robustness and recoverability show the highest resilience, while those with both low thresholds exhibit the lowest. Among economies where one value is high and the other is low, those with higher recoverability show better resilience. These findings emphasize the importance of enhancing import competitiveness, particularly for economies with lower recoverability, to improve their ability to withstand external shocks. By applying these insights, economies can better prepare for and mitigate the impacts of supply shocks, thus maintaining stability in global trade.

TD12

Moor Room

Strategic Decision-Making under Uncertainty

Invited Session

Revenue Management and Pricing

Chair: Guangyan Gan, NTU, Singapore, Singapore

Co-Chair: Jingren Liu, National University of Singapore, Singapore, Singapore

1 - Maintaining Supply Network Resilience via Information Design

Qiaochu He, Southern University of Science and Technology, Shenzhen, China, People's Republic of, Peifeng Li, Jie Wang

Advances in virtual and cloud-based manufacturing have enabled platform-based supply networks that algorithmically match interconnected suppliers with upstream and downstream partners. We study the profitability and resilience of such networks, identifying a critical threshold in matching probability: above this point, the network remains stable and resilient; below it, it is prone to collapse.

To address fragility, we propose an information design approach where the platform strategically adjusts the matching algorithm to correct suppliers' underestimation of positive network externalities. This, in turn, enhances suppliers' willingness to participate, amplifying the self-reinforcing nature of network effects. We also explore the joint design of matching algorithms and pricing strategies to maximize platform profits. In the context of rising trade barriers and global supply disruptions, our findings underscore the vulnerability of

interconnected supply chains and highlight the strategic role of information in maintaining platform network stability.

2 - Relative vs. Absolute Volume-Based Export Restriction: Choice, Design, and Effectiveness

Shuguang Zhang, NUS Chongqing Research Institute, Chongqing, China, People's Republic of, Lucy Gongtao Chen, Aditya Vedantam, Srinagesh Gavirneni

Governments commonly implement export restrictions to manage price volatility and safeguard domestic consumers. While existing research mostly focuses on exploring the social welfare impacts of export restrictions, there is a lack of analytical work examining their effects on firms' supply chain decisions and profitability. We address this gap by studying a firm's optimal production and export decisions under yield uncertainty and two types of volume-based export restrictions: absolute export restriction where the government imposes a fixed limit on export quantities, and relative export restriction where a proportion of the realized yield is allowed for export. We employ a game theoretical framework to study the interactions between a firm and a government, who make (production and restriction, respectively) decisions either sequentially or simultaneously. We show that under sequential decision, relative restriction generally leads to higher domestic supply and social welfare. In contrast, absolute restriction tends to be more favorable for the firm's profitability. Under simultaneous decision, however, relative restriction becomes less attractive to the government, while absolute restriction becomes less appealing to the firm. This makes it more likely for the government and the firm to be aligned on their preferred type of export restriction. We also find that preferences for export restrictions are notably different with and without yield uncertainty, and that a higher yield uncertainty increases the likelihood of misalignment between the firm's and government's preferences.

3 - Intertemporal Price Discrimination In Two-Sided Markets

Vinaysingh Chawan, IIM Indore, Indore, India

Amazon and Flipkart conduct many festive sales all through the year. During these festive sales, the platform reduces its platform fees to induce sellers to offer discounted prices to buyers, thereby increasing sales volume. In this scenario, the benefits to buyers are clear. Since the sale dates are known well in advance, this may result in a learned response from the buyers. For the sellers, this is an intertemporal price discrimination opportunity and they need to decide the availability and prices to maximize their profits. The platform needs to consider the behavior of both the buyers and the sellers to decide the frequency and extent of festive sales. We first study the intertemporal price discrimination problem from the seller's perspective and get insights for the seller's pricing problem. Then we look at the platform's problem and offer suggestions for managing festive sales.

TD14

Wilberforce

Sustainability and Circular Economy in Supply Chains

Contributed Session

Supply Chain Management

Chair: Heyin Hou, Southeast University, Nanjing, N/A

1 - Shall Ye (Not) Do It Thyself? Retailer Strategies to Validate Sustainability Attributes of Product

Raunak Joshi, Jindal Global Business School, Sonipat, India, Ranit Saha

With the enhanced awareness of environmental and ethical concerns of product manufacturing, consumers are becoming more conscious of the sustainability features of the products they purchase. Loss of natural habitat, naturally-produced ingredients/raw materials, product recyclability, greenhouse emissions, fair working conditions for laborers, and packaging are some of the factors that influence consumer decision-making. Purchasing decisions of more than 30% of global consumers reflect these factors, with over 60% of consumers seeking sustainable products in the cosmetics segment. Although awareness about sustainability has increased, there are certain challenges: (i) consumers find it difficult to distinguish sustainable products, and (ii) although consumers are willing to purchase the product at a premium, the prices charged are higher, dissuading many consumers. Motivated by the above dilemma, we evaluate the strategies a retailer should adopt to authenticate (or not) the product's sustainability attributes. Our analysis shows that when consumers are less suspicious about the claims, the retailer derives a greater value from partial monitoring, even when the monitoring costs are high.

Surprisingly, when suspicion is high, the retailer either monitors all the claims or does not monitor them at all. In addition, we find that verification of such claims may lead to environmental damage and reduction in the social welfare if consumers are not very suspicious.

2 - Product Reviews and Hybrid Selling Formats for Competing Manufacturers and a Retail Platform

WENJUAN YANG, Nanjing University of Finance and Economics, Nanjing, China, People's Republic of

This paper examines how product reviews affect the pricing decisions and profitability of competing manufacturers and the retail platform under hybrid selling formats: PR (high-quality product via platform, low-quality product via reseller) and RP (roles reversed). Using pure reseller (RR) and platform (PP) as benchmarks, we construct a game-theoretic model to analyze equilibrium outcomes. First, the results show that the precision of product reviews in the quality dimension affects both manufacturers similarly across pure and hybrid selling formats, but its impact on the retail platform differs between the pure and hybrid formats. Under hybrid formats, the retail platform's gains or losses depend on whether it can leverage the perceived quality difference through wholesale margins or commission earnings. Moreover, review informativeness in the fit dimension consistently influences the high-quality manufacturer's profit under pure and hybrid formats. However, its impact on the low-quality manufacturer and the retail platform varies across hybrid and pure formats. Second, selling formats and product reviews jointly affect the profitability for both manufacturers. The high-quality manufacturer reaches the highest profits in PR or RR formats, but the low-quality manufacturer may attain the highest profits in RP or RR formats. Additionally, hybrid selling formats PR and RP can outperform the pure formats RR or PP, potentially achieving a win-win result. Third, for the retail platform, the highest profits can be achieved in RR, PR, or PP formats. Notably, hybrid format PR can generate a win-win-win outcome, outperforming pure format PP for both manufacturers and the retail platform.

3 - Multicommodity Spatial Price Equilibrium Networks with Arc Multipliers, Capacity Constraints and Demand Uncertainty

Heyin Hou, Southeast University, Nanjing, China, People's Republic of

In this paper, we construct and analyze a multicommodity spatial price network model involving arc multipliers, capacity constraints and demand uncertainty. This model aims to explore how the vendor side (i.e., manufacturers) and the buyer side (i.e., demand markets) achieve market equilibrium through strategic interaction in a complex market environment, and pays special attention, through numerical simulations, to the influence of excess supply/demand penalty weight on market equilibrium results.

TD15

Fairmont Ballroom

Data-Driven and AI Applications in Logistics

Contributed Session

Transportation Science and Logistics

Chair: Xinyi Ye, King's college london, 14 Edenham Way, London, N/A, United Kingdom

1 - Predictive Analytics Approach in Determining the Input and Output Variables for Measuring Efficiency of the Indian Urban Road Transport Organizations

M. Mathirajan, Indian Institute of Science, Bengaluru, India, Rupkatha Ghosh

The analysis of the existing literature indicated that no analytical approach exists to determine the ideal input and output variables for measuring efficiency of the Urban Road Transport Organizations (URTO) using Data Envelopment Analysis (DEA). This paper addresses this research gap using predictive analytics approach, particularly a Correlated Random Effects integrated with Quantile Regression (CRE-QR) model. We initially collected data spanning 2011-2020 for 10 Indian URTO (I-URTO) and identified eight input and five output variables from literature and government reports. To remove the inflationary effect, monetary data was converted to 2011-12 constant prices using appropriate price indices. The missing data was handled using the Random Forest-based Multivariate Imputation by Chained Equations because (a) the missing data was Missing at Random and (b) the scatterplots indicated non-linear relationships between the dependent and independent

variables. The final dataset was a compact sample consisting of 90 observations. The descriptive statistics revealed a right-skewed distribution among the variables. A multi-collinearity test was performed to remove the problematic variables with a high variance inflation factor. Using Levin-Lin-Chu and Augmented Dickey-Fuller unit root tests, the data was found to be first difference-stationary. Consequently, the Johansen cointegration test was applied to identify cointegrating relationships among the variables. The Granger causality test resulted in the exclusion of one input variable. The CRE-QR model was then applied to the dependent and independent variables, finally identifying two input variables – labor cost and fuel cost – and one output variable, operating revenue, for assessing the efficiency of I-URTO using DEA.

2 - Structural Transformation and Resilience in India&Rsquo;S Supply Chain: An Empirical Analysis of E-Way Bill Data and Economic Indicators Before and After COVID-19

Snigdha Muddam, Mahindra University - School of Management , Hyderabad, India, Dr.Debopam Chakrabarti, Kirit Ghosh, Medhansh Bairaria, Smriti Madineni

The structural transformation and resilience of India’s supply chain ecosystem in the context of the COVID-19 pandemic were lauded internationally. Through a comparative analysis of pre-and post-pandemic trends (2018–2024), the research highlights how Goods and Services Tax (GST) reforms and the e-way bill system streamlined logistics and enhanced transparency across state boundaries. However, the pandemic exposed systemic vulnerabilities, causing significant supply movement and sectoral output disruptions. Employing correlation analyses between e-way bills, asset values, State Gross Domestic Product (SGDP), PMI, and IIP, the paper identifies patterns of sectoral recovery, regional dominance, and the adaptive responses of state economies. The results underscore the importance of digital infrastructure, targeted policy interventions, and intermodal logistics networks in fostering supply chain resilience. The study further recommends the integration of blockchain and AI for predictive logistics, developing sustainable cold chain systems, and harmonizing logistics policies to mitigate future disruptions. These insights contribute to the broader discourse on economic resilience in emerging markets and offer actionable strategies for policymakers to enhance efficiency, reduce risk, and promote inclusive, sustainable supply chain growth in India.

3 - Applying Two-Stage Heuristics to Optimize and Route Re-Design for the Animal Husbandry Supply Chain

Sanjay Choudhari, Indian Institute of Management Indore, Indore, India

This work extends our earlier project related to applying the VRP (Vehicle routing problem) problem in designing routes for managing the inventory of liquid Nitrogen at various livestock locations in the Animal Husbandry Department. The livestock supply chain manages the semen doses at many locations for the artificial insemination of animals. These doses must be kept and preserved at a cryogenic temperature. Traditionally, these doses are stored in liquid nitrogen worldwide. Hence, a significant part of the livestock supply chain distribution is to ensure liquid nitrogen availability at each livestock location. Our earlier work was designing route plans for optimizing the cost for replenishment of liquid nitrogen across all livestock locations, which resulted in a large saving using VRP heuristics. We redesign these routes by applying two-stage heuristics as a new extension of the same problem. First, we form k clusters of all the livestock locations and then apply the VRP problem to the clusters in order to replenish the liquid nitrogen at the cluster level. In consultation with the animal husbandry department, the new route redesign is expected to make significant savings, saving logistics costs and reducing liquid nitrogen inventory across various livestock locations. The contribution of the work is the practical application of these concepts in real-life situations as a part of redesigning a new supply chain network for the Animal Husbandry Department, which is an extension of our earlier published work in the INFORMS journal.

4 - Analyzing Stability in Makespan Scheduling Algorithms

Xinyi Ye, King's college london, London, United Kingdom, Dimitrios Letsios

In real-world scheduling applications such as logistics and delivery, small delays in job processing times are common, yet fully recomputing optimal schedules in response is often impractical. While disruption management has been widely studied, most existing work focuses on preserving optimality, leaving a gap in understanding the stability of approximate methods commonly used in practice. This work studies the stability of some solution methods for the classical parallel-machine scheduling problem ($P \mid\mid C_{max}$) when job processing times are changed. We present an approach to measure how stable scheduling algorithms are under such

changes, using ideas from counterfactual analysis. We define the concept of stability radius and give exact algorithms to compute it for the Longest Processing Time (LPT) and the Local Search (LS) algorithm. Our analysis describes which perturbations to job times do not affect the overall schedule and shows when LPT keeps the same job-to-machine assignments. For LS, we define stability in terms of keeping the result locally optimal. We also study how both algorithms behave when all jobs are delayed equally (uniform perturbations), and give clear conditions under which their output does not change at all.

TD16

VIP A

Technology and Data-Driven Optimization

Contributed Session

Business Analytics

Chair: Imcha Longchar, No 1. Dasyue Rd, East District, 701

1 - A Comparative Study on the Competitive Advantage of the 5G Industry from the Perspective of Patent Value: A Case Study of China, the United States, Japan and South Korea

lina xiao, University of Science and Technology Beijing, Beijing, China, People's Republic of, Wu Yang

Enhancing the competitive advantage of the 5G industry is crucial for accelerating the development of new quality productive forces and fostering China's independent scientific and technological innovation capabilities. Based on the patent data of China, the United States, Japan and South Korea in the field of 5G technology from 2005 to 2023, this study constructs a three-dimensional model of "technology-economy-law" from the perspective of patent value. Building on this framework, an evaluation system for the 5G industry competitive advantage index is established. The index weights are determined using the analytic hierarchy process and entropy weight method. Subsequently, the composite index method is employed to obtain the 5G industry competitive advantage index for China, the United States, Japan, and South Korea. The results show that from the perspective of the overall index, China surpassed the United States in 2018 to become the country with the largest advantage in the 5G industry. From the perspective of different dimensions, China's competitive advantages in the technical and economic dimensions are relatively prominent, but it is still necessary to further promote the commercialization of 5G technology, accelerate the construction of the 5G application ecosystem, and continuously improve the level of patent protection to obtain sustainable competitive advantages.

2 - Global-Deep Synthetic Control: a Heterogeneous Treatment Effect Estimator for Shared Market

Klaus Ackermann, Monash University, Melbourne, Australia, Priscila Grecov, Christoph Bergmeir

Estimating heterogeneous treatment effects in time series data is crucial across numerous real-world scenarios. Nonetheless, ensuring the no-interference assumption for control units can often be challenging when gauging treatment effects via a counterfactual and synthetic control approach in a shared market environment. To address spillover effects that can bias the results, we introduce “Global-Deep Synthetic Control” (Global-DeepSC), which uses globally trained deep learning forecasting models to estimate the treatment effects on multiple treated instances. Global-DeepSC addresses the violation of the non-interference assumption by forecasting counterfactual outcomes using only pre-intervention data for learning. Such an approach curtails the reliance on post-intervention disturbances, thereby mitigating the interference of spillovers and enhancing the accuracy of estimations. Moreover, our framework incorporates data-driven clustering and partitioning stages to purge the modeling of control units affected by spillovers and pinpoint variables responsible for heterogeneous impacts. Experiments on synthetic and real-world datasets show that our method out- performs well-known structural causal models regarding treatment effect estimation and statistical significance tests. An experimental study on the effect of sales promotions on supermarket sales found that younger, lower-income individuals should be the primary target audience for the marketing strategy, as estimated by conditional treatment effect analysis.

3 - Carbon Trading Price, Green Innovation and Financialization: the Role of Performance Feedback

Ailin WU, Korea University, Seongbuk-Gu, Korea, Republic of, Daeki KIM

As a key driver of green transformation, carbon pricing mechanisms significantly influence corporate strategic decisions between green innovation and financialization. Based on the price theory and corporate behavior theory, this study takes China's A-share listed companies from 2014 to 2023 as a sample and systematically examines the impact mechanism of carbon trading prices on corporate green innovation and financialization through fixed-effect models and moderating effect models, with a special focus on the moderating role of negative financial performance feedback. The results show that, first, the increase in carbon trading prices significantly stimulates green innovation investment, while inhibiting the degree of corporate financialization. Second, negative financial performance feedback weakens the effect of carbon trading policy. When firms' actual performance is lower than expectations, they are more likely to increase their financial asset allocation to maintain liquidity, while reducing green-related investment. Third, heterogeneity analysis reveals that the role of carbon trading policy in promoting green innovation is more significant in state-owned enterprises (compared with non-state-owned enterprises) and non-heavy polluting industries (compared with heavy polluting industries). This paper's research findings enrich the effectiveness of carbon trading policies in fostering corporate development, highlight the key moderate role of corporate performance feedback in the transmission of environmental policies. It also extends the application of behavioral theory in environmental economics, providing a new perspective for understanding corporate sustainable development decisions.

4 - Data-Driven Optimal and Myopic Policies for Inventory Systems with Demand Covariates

Kevin Shang, Duke University, Durham, NC, United States, Jingkai Huang, Yi Yang, Weihua ZHOU

We study a multi-period backorder inventory system in which the demand in each period depends on exogenous random covariates. The objective is to develop a data-driven policy that minimizes the total expected inventory cost in a finite horizon. The policy is developed through a two-step process. First, we assume the distributions of the covariates are known and derive the optimal policy by solving the dynamic program of the problem. Because it is computationally infeasible with a large number of covariates, we propose two heuristic policies, namely, the myopic and semi-myopic policies, and derive their performance guarantees. Second, we develop a data-driven method for each inventory policy using machine learning tools and sample-average approximation. Specifically, we estimate the demand function with machine learning algorithms and calculate the residual function in each period. We then apply the sample-average method to generate the data-driven solution from the corresponding empirical dynamic program. We derive finite-sample performance bounds for these data-driven inventory policies. In addition, we show that when the sample size is sufficiently small, the data-driven myopic and semi-myopic policies are identical to the optimal policy. These findings along with the numerical study assist managers in deciding when to implement these data-driven policies based on the sample size and system parameters. We extend the main results to the system with correlated demand covariates.

5 - Financial Literacy, Access to Finance, and the Moderating Role of Government Support: Implications for Sme Performance in Cambodia

Imcha Longchar, National Cheng Kung University, Tainan, Taiwan, Min-Hsien Chiang, Sotheara Oam

SMEs constitute the backbone of Cambodia's economy, contributing significantly to employment and GDP. Despite their economic and social contributions, nearly half of these firms fail within five years of operation, with limited access to finance identified as one of the key challenges. Drawing on the Upper Echelons Theory and the Pecking Order Theory, this study explores how financial literacy influences SMEs' access to finance, and how access to finance subsequently affects SME performance. Furthermore, the study examines the moderating role of government support programs in these relationships. The proposed relationships were empirically examined using partial least squares structural equation modeling (PLS-SEM). The results reveal that financial literacy is positively associated with access to finance, and that access to finance, in turn, has a positive effect on SME performance. Moreover, government support programs positively moderate the relationship between access to finance and SME performance, but not between Financial Literacy and access to finance.The implications of these findings support the Upper Echelon Theory, suggesting that enhancing SME managers' financial literacy can improve access to finance. This also aligns with the Pecking Order Theory, which posits that financially literate managers are better positioned to make informed financing decisions based on the firm's internal and external funding priorities. Furthermore, while government support programs enhance the effect of access to finance on SME performance, they do not significantly improve the financial literacy and access to finance,

suggesting a need for more integrated support systems that combine financial assistance with skills development and advisory support.

Wednesday, July 23, 8:00 AM - 9:30 AM

WA01

Enterprise

New Insights in Stochastic Systems

Invited Session

Applied Probability Society

Chair: Yunfang Yang, National University of Singapore, Singapore, Singapore

Co-Chair: Yunfang Yang, National University of Singapore, Singapore, Singapore

1 - Economies of Operational Hybridity

Yunfang Yang, National University of Singapore, Singapore, Singapore, Shuangchi He

Large-scale service systems often need to serve customers with different waiting time requirements. For example, the inpatient department of a general hospital is usually required to make beds available within hours for patients transferred from the emergency department, whereas it may provide beds in days or weeks for elective patients. In this study, we consider a many-server queue with two classes of customers that are managed under two distinctive regimes: the waiting times of the first-class customers should be on a lower order of magnitude than those of the second-class customers (e.g., hours versus days). There are cross-trained servers that can serve both classes. We first propose a sequencing policy for cross-trained servers in order to meet service-level requirements for the two classes. Then, we identify the minimum number of cross-trained servers that is required for the sequencing policy to minimize the overall staffing level. We prove that under the proposed sequencing and cross-training scheme, the service system can achieve "economies of operational hybridity"—that is, one may dimension the system as if it operates under the more efficient regime that the second-class customers require, while a small fraction of cross-trained servers are sufficient to fulfill the first-class customers' waiting time requirements.

2 - Data-Driven Matching for Impatient and Heterogeneous Demand and Supply

Xun Zhang, Southern University of Science and Technology, Shenzhen, China, People's Republic of, Weiliang Liu, Amy Ward

Motivated by two-sided service platforms, we study the problem of matching heterogeneous demand (customers) and supply (workers) that arrive randomly over time. Customers and workers arrive with a randomly sampled patience time, and are lost (renege) if forced to wait longer than that time for a match, creating a fundamental tradeoff between immediate matching and waiting for better pairings. We consider a data-driven setting where the system operator designs matching polices based only on historical data on inter-arrival and patience times, without any knowledge of their ground-truth distributions. Finding an optimal policy is intractable even with full information due to potentially non-Markovian arrivals and renegings. To address this, we formulate a deterministic, data-driven fluid matching problem (DDMP) that serves as an approximation to the original matching problem when demand and supply rates grow large. We establish finite-sample statistical guarantees on the DDMP based on a novel moderated uniform error bound for estimating the patience time quantile function---results that may be of independent interest. Furthermore, we show that the DDMP can be approximately solved via a mixed-integer linear program; we quantify the resulting approximation gap and analyze its impact on policy performance. Overall, this paper provides a framework for integrating statistical learning with queueing asymptotics to develop effective data-driven policies under unknown model primitives.

3 - Blessing of Order Crossover in Inventory Systems

Zhaoxuan Wei, Institute of Operations Research and Analytics, NUS, Singapore, Singapore, Yunfang Yang

We revisit classic single-product inventory systems with backlogs and stochastic leadtimes. Though providing analytical tractability, we find sequential supplies may prevent the inventory systems achieving economies of scale---the optimal inventory cost grows linearly with the demand rate λ unless the leadtime is constant. In contrast to sequential supplies, order crossover renders systems with parallel supplies generally intractable. However, order crossover allows new orders to shorten the remaining time until supply units arrive. We show this is the key benefit of order crossover, reducing the optimal cost scale from $\Theta(\lambda)$ to the order of $\Theta(\lambda^{1/3})$ and we refer to such a phenomenon as the blessing of order crossover. Furthermore, we demonstrate that policies with only two base-stock levels suffice to achieve this optimal rate. Our results highlight the distinct nature between systems with sequential and parallel supplies, provide quantitative insights into the impact of stochastic leadtimes and order crossover, and offer general guidance on policy designs for inventory systems with order crossover. We further extend our analytical framework to account for varying unmet demand assumptions and partial information, proving the (near) optimality of our proposed policy in partial backlogged, lost sales, and partially observed partial backlogged systems.

4 - Dynamic Factor Models with Forward-Looking Views

Anas Abdelhakmi, National University of Singapore, Singapore, Singapore, Andrew Lim

We develop a continuous-time framework for integrating noisy expert views of the future values of covariates in a dynamic factor model. By exploiting the graphical structure linking factors, asset prices, and expert views, we derive closed-form conditional dynamics for both the factor and price processes. While it is expected that views will alter the factor dynamics, we show that they also change the price process dynamics when factor and asset innovations are correlated. For linear factor models, the price process becomes a time-inhomogeneous affine process with a new covariate formed from the views, and time-varying coefficients determined by the structure of the views and factor-asset correlations. We establish a novel theoretical connection between forward-looking views and a process we call a Mean-Reverting Bridge (MrB), an extension of the classical Brownian bridge. We derive the investor's optimal portfolio strategy and show that views influence the strategy exclusively through the intertemporal hedging demand. Investors with more precise views are more sensitive to factor movements and hedge more aggressively. More generally, our framework offers a generalizable approach for embedding forward-looking information about covariates in a dynamic factor model.

WA02

Mercury

Aviation Topics IV

Invited Session

Aviation Applications

Chair: Yuan Yuan, Northwestern Polytechnical University, Xi’an, N/A

1 - Adaptive Decision Intelligence: AI-Powered Robust Optimization for Aviation Systems and Crisis Response

Rajesh Kumar Korupalli V, Woxsen University, Hyderabad, India, Raul V Rodriguez

Modern aviation systems and emergency logistics operations demand rapid, data-driven decision-making under conditions of uncertainty, incomplete information, and dynamic constraints. Traditional optimization models, while theoretically sound, often lack the agility and resilience required in these volatile environments. This study presents an adaptive decision intelligence framework that integrates *AI-driven scenario modeling* with *distributionally robust optimization (DRO)*, enabling intelligent decision-making using limited statistical descriptors such as means and variances. Grounded in the principles of *information systems design* and *prescriptive analytics*, the proposed architecture supports real-time optimization even when full probability distributions are unavailable or unreliable. In aviation, the model enables dynamic scheduling and pricing strategies by learning from fluctuating passenger behavior and demand uncertainty. In crisis response scenarios, it aids in robust planning for resource allocation and last-mile delivery where infrastructure and supply chains are unstable. The system architecture incorporates machine learning for pattern recognition and anomaly

detection, feeding into an optimization core that ensures solutions remain feasible and effective across a range of uncertain conditions. This work showcases how the fusion of AI, robust optimization, and intelligent information systems can power resilient, scalable decision tools for mission-critical operations across diverse sectors.

2 - Capturing Tacit Knowledge with Artificial Intelligence: a Conceptual Foundation for Airline Operations

Matthias Bickel, ETH Zurich, Zurich, Switzerland, Torbjørn Netland

Can artificial intelligence (AI) capture tacit knowledge? In a research collaboration with an airline company, we are developing a decision-support system for airline operations that learns from expert decisions made during flight disruptions. This paper provides the groundwork for such a system by exploring how AI can support the capturing and sharing of tacit knowledge, a key concept in Knowledge Management. Unlike explicit knowledge, tacit knowledge cannot be easily articulated by its holders, making it challenging to manage and utilize effectively, especially in dynamic and complex environments like airline operations. Disruption management in airlines, which involves the operational handling of delayed flights and crews, is currently managed manually. Decision-makers rely heavily on tacit knowledge to find feasible solutions that optimize various aspects, such as reducing penalty fees and minimizing cascading delays. However, due to its intangible nature, tacit knowledge cannot be effectively utilized outside of the individuals who possess it, limiting the efficiency of these processes. To overcome this issue, we observe expert decision-makers to identify the variables they consider and the decisions they make, which allows us to model the system of the decision-making process accurately. Subsequently, we suggest applying specific AI technologies to partially externalize the tacit knowledge utilized by these decision-makers. This externalized knowledge is fed back into the decision-support system, enhancing its efficiency and utilization. The managerial impact of implementing AI in this context is that it enhances decision-support systems and improves understanding of decision-makers' reasoning, leading to more robust and consistent airline operations.

3 - Three Stages of Departure Aircraft: Influencing Factors and Prediction of Taxi-Out TIME

tongtong li, Nanjing University of Aeronautics and Astronautics, Nanjing, China, People's Republic of

Accurate prediction of taxi-out time is critical for optimizing pushback scheduling and improving airport surface efficiency. Existing studies predominantly treat surface taxiing as an undifferentiated process, lacking a refined prediction of the taxiing process. Therefore,a multi-stage taxi-out time prediction method based on control handover is proposed. Firstly, the taxiing process is decomposed into three operationally distinct phases through jurisdiction analysis: 1) pushback phase, 2) apron taxiing phase under ramp control,3) movement area taxiing phase under tower control.We innovatively introduce the gate cluster concept and develop phase-specific feature . Then LightGBM-based predictive model is built to verify the prediction effect of multi-stage division on taxi-out time. Taking the actual operation data of Shanghai Pudong International Airport as an example, Empirical results demonstrate that the phased prediction approach achieves significant accuracy improvements over monolithic models, increasing the proportion of flights with prediction errors ≤1 minute by 4.88 percentage points (pp) and ≤3 minutes by 6.46 pp. Compared with the overall prediction, and the characteristics of the taxi stage make a greater contribution to the prediction accuracy of each stage.

WA03

Minto

Recent Advances in Distributionally Robust Optimization

Invited Session

Computing Society

Chair: Chin Pang Ho, City University of Hong Kong, Hong Kong

1 - Contextual Chance-Constrained Programs: a Lifted L^\Infty Optimal Transport Approach

Jun Li, National University of Singapore, Singapore, Singapore, Adam Ng, Shuming Wang

Chance-constrained programming (CCP), as an important class of optimization-under-uncertainty models, has been employed in many high-stakes decision-making problems. In this paper, we study a contextual robust model of joint CCP with piece-wise convex constraints of uncertain coefficients and binary decision variables. To harness contextual information for boosting decision performance with appealing statistical guarantees, and to improve solution scalability, we develop a predictive ambiguity set that lifts an L^\infty optimal transport (OT) counterpart and houses a nonparametric kernel regresssion model. Under several regularity conditions, we show that the solution of the proposed lifted L^\infty-OT based contextual model remains feasible with high probability under the true distribution of uncertainty. Optimization-wise, we develop exact solution methods for the proposed framework, tailored to commonly used OT cost functions. These include a reformulation scheme of mixed-integer conic program that exploits linearization with derived valid inequalities, and a Benders decomposition scheme that exploits the decomposability of the lifted L^\infty-OT model to enhance the scalability. Finally, we validate the model's robustness, the value of capturing contextual information, its finite-sample performance, and its computational efficiency through a redundancy allocation testing problem using both simulated and real-world datasets.

2 - A Robust Satisficing Multi-Objective Optimization Approach for Bike-Sharing Systems with Heterogeneous User Types

Qingxin Chen, Tianjin University, Tianjin, China, People's Republic of, Shoufeng Ma, Chenyi Fu, Ning Zhu, Qiao-Chu He

Diverse subscription policies have been employed to cater to the commuting needs of users in bike-sharing systems. Frequent riders choose to purchase subscriptions to reduce their commuting costs. Thus, operators should improve their service level to maintain the market share. In addition, a small proportion of customers (e.g., tourists) use the system occasionally but contribute to the operational profits. Thus, providers are willing to fulfill more customer demand to ensure high profitability. However, jointly optimizing these objectives for the heterogeneous users in bike-sharing systems is challenging, especially under demand uncertainty. To tackle these concerns, this study presents a robust satisficing framework focusing on both service level and profit under demand uncertainty, which frees the decision-makers from formulating the demand ambiguity. This study further defines a novel risk measure, which is formulated according to the φ-divergence probability distance. The proposed risk measure aims to minimize the violation risk of the service level and profit targets. To avoid overconservation of solutions, side information like weather and weekends is integrated into the model. A tailored local search algorithm is proposed to address large-scale problems, which is capable of integrating rebalancing routes and addressing decentralized demand dynamics. Extensive numerical experiments show that the proposed model achieves higher robustness, lower violation probability and degree, and higher average-case out-of-sample performance than other benchmarks. Managerial insights are also concluded for the operators.

3 - First-Order Methods for Contextual Stochastic Bilevel Optimization

Zhichun Yang, The University of Hong Kong, Hong Kong, China, People's Republic of, Yves Rychener, Yifan Hu, Man-Chung Yue

We consider contextual stochastic bilevel optimization (CSBO), a new bilevel optimization model, where the upper-level problem is to minimize an expectation with respect to the contextual information and the lower-level problem is to minimize another expectation conditioned on the contextual information and upper-level decision, possibly constrained. CSBO is very attractive from a modeling perspective and finds many applications, including meta-learning, end-to-end learning, and Wasserstein-distributed robust optimization with side information. Nevertheless, CSBO is computationally challenging. Most existing algorithms either rely on the Hessian inverse of the lower-level loss function and do not scale well or cannot handle lower-level constraints. In this talk, we propose and analyze a novel fully first-order method for solving CSBO where the lower-level problem is subject to linear constraints. Using the Multilevel Monte-Carlo technique, we also present an accelerated version of our algorithm that enjoys a better complexity.

4 - Distributionally Robust Federated Learning: An Admm Algorithm

Yi WONG, City University of Hong Kong, Hong Kong, Hong Kong, Wen Bai, Xiao Qiao, Chin Pang Ho

Federated learning (FL) aims to train machine learning (ML) models collaboratively using decentralized data, omitting the need for centralized data aggregation. Standard FL models often assume that all data come from the same unknown distribution. However, in practical situations, decentralized data frequently exhibit heterogeneity. We propose a novel FL model, Distributionally Robust Federated Learning (DRFL), that applies distributionally

robust optimization to overcome the challenges posed by data heterogeneity and distributional ambiguity. We derive a tractable reformulation for DRFL and develop first-order methods to solve this problem in a federated setting. Our experimental results demonstrate that DRFL outperforms standard FL models under data heterogeneity and ambiguity.

WA04

Indiana

Mathematics of Machine Intelligence

Invited Session

Data Mining

Chair: Subhroshekhar Ghosh, National University of Singapore

1 - Determinantal Point Processes in Machine Learning: Smaller Coresets via Negative Dependence

Hoang Son Tran, National University of Singapore, Singapore, Singapore

Determinantal point processes (DPPs) are random configurations of points with repulsive behavior. Originally introduced in quantum mechanics, DPPs have a long history in both mathematics and physics. In the big data era, DPPs have attracted the attention of the machine learning community as they are natural candidates for sampling diverse subsets from large datasets. Empirical evidence from several recent works suggests that DPPs can outperform traditional subsampling methods that rely on statistical independence; however, theoretical support remains limited. In this talk, we will discuss using DPPs to construct coresets, a fundamentally important task in machine learning. In particular, we will provide a theoretical guarantee for the statement that *DPPs can provably outperform independently drawn coresets*, which has been empirically verified in the literature.

2 - Learning Causal World Models from Acting and Seeing Using Score Functions

Karthikeyan Shanmugam, Google Deepmind India, Bengaluru, India, Burak Varici, Emre Acarturk, Abhishek Kumar, Ali Tajer

In causal inference, true causal order and the graph of causal interaction can be uniquely determined if you have sufficient interventional data. Interventions are local (randomized control trials) RCTs done where different variables, taken few or one at a time, in a causal graph are randomized. We consider a harder problem when the causal variables are not directly observed and are "latent". Instead we observe a high dimensional transformation (as images etc.) of the true causal variables. Central problem in causal representation learning is to invert the unknown transformation between true causal variables and the observations upto coordinate wise scaling and permutation. We show that this is possible with enough interventional diversity by exploiting two key ideas: a) Represent interventional distributions in terms of their scores (gradient of likelihoods). b) The encoder-decoder pair that minimizes reconstruction loss and sparsifies the score difference in the latent space is the optimal pair. We show various versions of these results for linear transforms and general transforms with mild regularity assumptions on the diversity of interventions. We also will discuss empirical results on some simple image datasets. Time permitting this will be followed by some preliminary results from ongoing work on extending this to robotic simulators.

3 - Statistical-Computational Trade-Offs for Recursive Adaptive Partitioning Estimators

Yan Shuo Tan, National University of Singapore, Singapore, Singapore, Jason Klusowski, Krishnakumar Balasubramanian

Models based on recursive adaptive partitioning such as decision trees and their ensembles are popular for high-dimensional regression as they can potentially avoid the curse of dimensionality. Because empirical risk minimization (ERM) is computationally infeasible, these models are typically trained using greedy algorithms. Although effective in many cases, these algorithms have been empirically observed to get stuck at local optima. We explore this phenomenon in the context of learning sparse regression functions over d binary features, showing that when the true regression function $f*$ does not satisfy Abbe et al. (2022)'s Merged Staircase Property (MSP), greedy training requires $\exp(\Omega(d))$ to achieve low estimation error. Conversely, when $f*$ does satisfy MSP, greedy training can attain small estimation error with only $O(\log d)$ samples. This dichotomy mirrors that

of two-layer neural networks trained with stochastic gradient descent (SGD) in the mean-field regime, thereby establishing a head-to-head comparison between SGD-trained neural networks and greedy recursive partitioning estimators. Furthermore, ERM-trained recursive partitioning estimators achieve low estimation error with O(logd) samples irrespective of whether f∗ satisfies MSP, thereby demonstrating a statistical-computational trade-off for greedy training.

4 - Subspace Embedding and Linear Regression Problems

Yi Li, Nanyang Technological University, Singapore, Singapore

In this talk, I shall give a brief overview of the subspace embedding problem and related linear regression problems from a theoretical computer science perspective. Specifically, the ℓ_p subspace embedding problem asks to find a matrix S for a given $n \times d$ matrix A (where $n >> d$) such that $\|SAx\|_p \approx \|Ax\|_p$ for all d-dimensional vectors x, assuming $p >= 1$ is constant. This matrix S is often constructed using Lewis weights sampling. Using subspace embeddings, one can reduce the scale of the standard linear regression problem $\min_x \|Ax - b\|_p$.

5 - Learning and Steering Game Dynamics Towards Desirable Outcomes

Antonios Varvitsiotis, Singapore University of Technology and Design, Singapore, Singapore

Game dynamics, which describe how agents' strategies evolve over time based on past interactions, can exhibit a variety of undesirable behaviours including convergence to suboptimal equilibria, cycling, and chaos. While central planners can employ incentives to mitigate such behaviors and steer game dynamics towards desirable outcomes, the effectiveness of such interventions critically relies on accurately predicting agents' responses to these incentives -- a task made particularly challenging when the underlying dynamics are unknown and observations are limited. To address this challenge, this work introduces the Side Information Assisted Regression with Model Predictive Control (SIAR-MPC) framework. We extend the recently introduced SIAR method to incorporate the effect of control, enabling it to utilize side-information constraints inherent to game-theoretic applications to model agents' responses to incentives from scarce data. MPC then leverages this model to implement dynamic incentive adjustments. Our experiments demonstrate the effectiveness of SIAR-MPC in guiding systems towards socially optimal equilibria, stabilizing chaotic and cycling behaviors. Notably, it achieves these results in data-scarce settings of few learning samples, where well-known system identification methods paired with MPC show less effective results.

WA05

Orchard A

Frontiers in Resource Allocation Problems

Invited Session

Efficiency and Fairness in Supply Chain and Service Operations Management

Chair: Lyu Hongyao, Shanghai Jiaotong University, N/A

1 - Differential Privacy too Costly for Resource Allocation? Benefits of a Joint Differential Privacy Approach

Du Chen, Nanyang Technological University, Singapore, Singapore, Geoffrey Chua, Victor Jose

Resource allocation is an important and pervasive problem in management science. Recent growth in privacy concerns has made this problem more challenging because agents’ data might be inadvertently revealed through the optimal allocation decisions. While differential privacy (DP) limits information leakage, it often lowers decision utility substantially, making it impractical for resource allocation. Joint differential privacy (JDP), another data privacy framework, is believed to better balance privacy with utility in this context, though its implications remain underexplored. We show that JDP provides privacy protection as strong as DP if agents do not disclose their allocations. Moreover, JDP may achieve this with substantially lower utility loss, with per-agent loss vanishing as the system scales. These theoretical results imply JDP's potential for a better privacy-utility trade-off, but existing JDP algorithms can backfire and empirically underperform DP. To bridge this gap, we exploit the optimization model structure to develop Noisy Dual Mirror Descent (NDMD), a minimax near-

optimal algorithm for a large class of convex resource allocation problems. NDMD offers a practical solution for privacy-preserving allocation, particularly for a pressing and long-debated public policy problem, namely, the allocation of Title I US education grants with privacy. Simulations with real-world data show that shifting from DP to JDP substantially reduces nationwide misallocation by as much as 52% and significantly mitigates the disproportionate impact of privacy protection on historically disadvantaged and underrepresented minority groups in this context. Our study highlights the importance of disclosure discretion, optimization models, and algorithm design in balancing privacy with utility.

2 - The Service-Centric Vehicle Routing Problem with Crowdshipping

Bingjie Zhou, Analytics and Operations Research, SINGAPORE, China, People's Republic of

Last-mile delivery services worldwide have embraced crowdshipping, which involves both regular and occasional drivers to reduce transportation costs and potentially ensure timely deliveries. However, real-world uncertainty in travel times leads to delays in deliveries. Motivated by empirical studies on customer impatience with late deliveries, this paper focuses on a service-centric Vehicle Routing Problem with Crowdshipping (VRPC) under uncertain travel times. Apart from traditional lateness measures, such as on-time arrival probability and expected lateness duration, we also consider scenarios where customers show exponential impatience towards lateness. We introduce a novel approach to calibrate the disutility, leading to salient managerial implications and probabilistic insights.
We develop an exact branch-price-and-cut algorithm for the deterministic VRPC and a route enumeration-based exact algorithm for the non-convex and discontinuous service-centric VRPC. Numerical studies based on existing instances validate the computational efficiency of the developed algorithms and the efficacy of the newly proposed lateness measures in mitigating the risk of late deliveries.

3 - Robust Parallel Machine Selection and Scheduling with Uncertain Release TIMES

Linyuan HU, Beijing Institute of Technology, Beijing, China, People's Republic of, Zhang Yuli

This paper studies a parallel machine selection and scheduling (PMSS) problem with uncertain release times. To handle uncertain release times, we propose a two-stage robust PMSS model where the release time deviation (RTD) is characterized by a budget uncertainty set. In the first stage, machine selection and job assignment decisions are made to minimize startup costs before the uncertainties are revealed. In the second stage, once release times are known, job sequences are optimized to minimize the makespan on each machine. Robust constraints are introduced to ensure that the worst-case minimum makespan on each machine does not exceed a pre-specified due date. The proposed model is a tri-level min-max-min optimization problem with mixed-integer recourse decisions, which cannot be solved efficiently by existing algorithms. To this end, we propose a novel logic-based Benders decomposition (LBBD) algorithm with strengthened Benders cuts and speed-up techniques. Specifically, we first provide an equivalent mixed-integer linear programming reformulation for the max-min subproblem by analyzing an optimality condition of the worst-case RTD. Second, we design novel combinatorial and analytical Benders cuts, which dominate cuts found in the literature, and we further strengthen them by lifting procedures. Third, we design a relaxation-and-correction procedure and a warm-start procedure to speed up the LBBD algorithm. Numerical experiments show the proposed robust model greatly reduces job tardiness compared with the deterministic model. The proposed cuts efficiently reduce the runtime, and the LBBD algorithm is at least three orders of magnitude faster than the state-of-the-art column-and-constraint-generation algorithm.

4 - Online Reusable Resource Allocations with Multi-Class Arrivals

Tianming Huo, National University of Singapore, Singapore, Singapore, Wang Chi Cheung

We study three problem formulations under the general theme of online reusable resource allocations with adversarial multi-class arrivals. In the case of single resource, we consider an admission control model and an advance booking model. In the case of multiple resources, we consider an assortment planning model. All models involve customers who are heterogeneous in their usage durations and rewards. For the admission control problem, we propose an online algorithm which achieves the best possible competitive ratio (CR) among all online policies, under mild conditions including a monotonicity conidition on the usage durations. We show that the problem becomes more difficult without the monotonicity condition, but our algorithm can still achieve a non-trivial CR. For the advance booking problem, we show that our policy achieves at least half of the best possible CR when customers have identical usage duration. For the assortment planning problem, our policy achieves at least half of the best possible CR under a divisible duration condition. The CRs are established with a

novel Auxiliary Algorithm (AA). The AA uses the whole realized online dynamics to construct a feasible solution to a linear program, which upper bounds the optimum.

5 - On Fairness of Process Flexibility

Lyu Hongyao, Shanghai Jiaotong University, Shanghai, China, People's Republic of

In the process flexibility literature, it is well known that "a little flexibility goes a long way," (Jordan and Graves 1995). This means that a sparse process flexibility structure, such as long chain, can capture nearly all the benefits of full flexibility in terms of demand satisfaction and/or capacity utilization. While efficiency metrics have been well assessed, the performance of sparse structures concerning fairness measures remains under-explored. Fairness is a crucial objective in resource allocation, driven by companies' social responsibility mandates and the needs of humanitarian operations. In this study, we investigate the worst-case performance of sparse flexibility structures in terms of the max-min demand fill rate. Through primal-dual analysis, we show that the max-min fill rate of a structure is determined by the demand nodes subset with the minimum ratio of neighboring capacity to demand. We further characterize the worst-case max-min fill rate of k-chain and ψ-expander structures. The results suggest that a more connected structure is needed to achieve outcomes close to full flexibility concerning fairness. Additionally, we demonstrate that there exists an allocation that can achieve both efficiency (i.e., max-flow) and fairness (i.e., max-min fill rate). Based on our theoretical findings, we propose a computationally efficient heuristic, the Cover Ratio heuristic, to generate effective flexibility structures.

WA06

Orchard B

Group Decision Making for Regional Sustainable Development

Invited Session

Group Decision and Negotiation

Chair: Gang Liu, Tianjin University, Tianjin, N/A

Co-Chair: Yi Xiao, Tianjin University, Tianjin, N/A

1 - Research on Water Resource Allocation Based on Gradient Lean Regulation

Xiaoying Lai, Tianjin University, Tianjin, China, People's Republic of, Yejun Xu

Due to the multi-subjectivity of regional water use (such as the "production-living-ecology" water use) and the multi-dimensional nature of the water use process (natural and social attributes and their coupling), the water use efficiency is low and the water resources are seriously wasted. In this study, based on the identification of the binary dynamic supply and demand relationship of water resources composed of water supply uncertainty and spatiotemporal heterogeneity of water demand, we redefined the basic elements and supply-demand balance relationship of water use among multi-agent under two-way constraints. From a comprehensive perspective driven by principles, data, and management, we proposed a regional water resource allocation mechanism based on the concept of lean management, and we established a gradient lean regulation model. Then we constructed a prediction model and analytical method of dynamic water availability and water demand, which forms the decision-making basis for water resource allocation. The expected results demonstrate certain theoretical significance and application value for solving the problem of water resources involving the dual relationship between supply and demand.

2 - A Coupled Model for Water-Saving Decision-Making by Intelligent Farmers Under Policy-Driven Conditions

Yi Xiao, Tianjin University, Tianjin, China, People's Republic of

Farmers’ interests and decision-making processes are crucial for the successful and effective implementation of water-saving policies considering their essential roles in agricultural production, irrigation, and making water-saving decisions. In this work, we developed a coupled socio-hydrological model based on the integration of a multi-agent modeling and hydrological modeling, which simulates the complex human-water interactions by coupling socio-economic behavior and physical hydrological processes. In the multi-agent model, each farmer is regarded as an independent decision-making unit, which makes irrigation decisions based on their own economic

conditions, water acquisition costs, policy environment, and other factors, following their respective water-saving behavior-driven paths. Furthermore, by simulating the interactions between individuals and their adjustments to environment, the overall system effects emerge, forming dynamic collective water-saving behavior. Meanwhile, the hydrological model provides a detailed physical basis to simulate natural processes such as water flow, evaporation, and precipitation. By coupling the multi-agent model with the hydrological model, it is possible to specifically analyze how agricultural activities such as farmers' water-saving practices affect groundwater levels, river flow, and the overall water cycle, and how these hydrological changes in turn affect farmers' irrigation strategies. By simulating these interactions and analyzing the diverse responses of individual farmer to various policy measures (such as subsidies for water-saving technologies or changes in water prices) from a micro-perspective, we can better understand the potential effects of various incentive policies, thereby helping to identify effective water resource management strategies that encourage farmers to make water-saving decisions.

3 - An &Ldquo;Initial-Engineering-Enduse&Rdquo; Quasi-Market Pricing Mechanism Under Cascading Risk for Cross-Regional Water Rights Trading

Gang Liu, Tianjin University, Tianjin, China, People's Republic of

Water rights trading, operating under a market economy, is recognized as an effective approach to enhance water use efficiency and address water scarcity. It remains a crucial challenge to design a water rights pricing mechanism that incorporates the spatial heterogeneity of water resources. This study aims to address this issue by proposing an "Initial-Engineering-Enduse" quasi-market pricing mechanism based on an empirical study of cross-regional water rights trading in Pingxiang, Jiangxi, China. The proposed mechanism consists of three stages, accounting for associated risks for each stage. Three pricing models, namely stochastic differential game, incomplete contract theory, and the Nash-bargaining model, are calibrated for the three stages based on our study. The interplay among the three equilibrium water prices in each stage is examined to understand cascade effects. It is worth noting that the local government of Pingxiang has adopted and implemented this pricing mechanism. Empirical results demonstrate the following key findings: (1) The quasi-market pricing mechanism effectively reveals the determination of enduse water prices. (2) The government's initial investment can mitigate trading risks and enhance the stability of water rights trading. (3) Neglecting the value of ecosystem services results in an undervaluation of both initial water rights prices and enduse water prices. Furthermore, the government should consider making initial investments in water conservancy projects to minimize risks. Overall, this study provides a practical and promising approach to pricing cross-regional water rights trading in both water-rich and water scarce areas.

4 - Can Ecosystem Services be Effectively Perceived by Local Residents：Linking Macro-Es and Micro-Individual Perception in the Yellow River Basin

Yejun Xu, Tianjin University, Tianjin, China, People's Republic of, Huimin Wang

Residents' perception of ecosystem services (ES) is a key factor in judging the effectiveness of ecological governance, especially for the ecological protection and high-quality development of the Yellow River Basin (YRB). However, existing research have not been effective in revealing the match between watershed ES supply and perceptions of local people. To address this gap, this paper develops an ES supply and individual satisfaction perception database of YRB at the city scale, constructs a method to unveil the "supply-perception" matching mechanism of ES based on subjective and objective perspectives, and puts forward a four-quadrant classification governance strategy of the basin city scale. The results show that: (1) Different levels of supply exist for various ES, and there is a spatial clustering on provisioning and regulating services. (2) There are significant differences in residents' satisfaction perception of ES: provisioning > cultural > regulating services. (3) There are pronounced synergistic effects of individual perceptions, however, at the city scale, the correlation between macro-ES supply and micro-individual perceptions is not significant. (4) Marked mismatches are identified between ES supply and individual resident satisfaction perceptions, particularly concerning water and cultural heritage. The results indicate that in most cities, ES are not effectively perceived by local residents. Accordingly, we design adaptive strategies for different urban ecological governance, around two dimensions: ES supply and residents' perceptions. Findings provide a feasible reference for formulating and implementing tailored strategies for coordinated human and geographical development, which is conducive to regional sustainable development.

WA07

Orchard C

Health System Simulation in the Age of AI

Invited Session

Health Systems Modelling with AI

Chair: Ashish Kumar, Singapore Health Services/ Duke-NUS Medical School, Singapore, N/A, Singapore

Co-Chair: Michael Dorosan, Singapore Health Services Pte Ltd, Singapore, Singapore

1 - Trauma Care Pathway Resource Optimization: Through the Lenses of Network Science, Predictive Modelling, and Simulation-Based Evaluation (An Ongoing Study)

Michael Dorosan, Singapore Health Services Pte Ltd, Singapore, Singapore

Trauma care aims to stabilize, diagnose, treat, and rehabilitate critically injured patients to restore their health and function. Often, such cases require urgent care amidst rapidly deteriorating conditions; resources must be efficiently and appropriately allocated. Focusing on an acute hospital's trauma cases at the emergency department (ED) within 2019-22, we surfaced varying resource utilization network structures among three severity tiers, signaling differences in utilization patterns. However, the existing tiering system is done almost retrospectively for national-level reporting and monitoring purposes. The network analysis results, therefore, support the need to develop a tiering system at the point of first presentation in ED. Currently, we are developing decision-support models that determine the downstream use of critical resources to improve their timely allocation. The impact of such predictive models will be evaluated through discrete event simulation of the trauma care pathway, supported by our trauma resource utilization network analysis.

2 - Can Naive Models Beat Health Professional Intuition? Improving Visit Intensity Projections for Home Health Care Services Using Machine Learning

Christian Alis, Asian Institute of Management, Makati City, Philippines, Kristine Ann Carandang

Doctors typically determine the home health care (HHC) plans of their patients, and these plans include the visit intensities (VI) patients may require per discipline [home health aide (HHA), medical social work (MSW), nursing (SN), occupational therapy (OT), physical therapy (PT) and speech therapy (ST)]. However, these plans may be revised on the recommendation of allied health professionals (AHPs) after conducting a more standardized assessment. Here, we investigate the predictive performance of these initial projections and then improve them using machine learning (ML)-based VI prediction models (VIPMs).

Using anonymized real-world health records (N=1,054) from HHC agencies in the United States, we found that doctor-ordered VI performed worse in predicting actual VI rendered for Nursing, OT and PT compared to naïve models in which the mean actual VI per discipline was used as the predicted value. These three services were involved in majority of the episodes (77%, 53% and 85%, respectively), unlike HHA, MSW and SN which were only ordered for less than 16% of the episodes. We then developed VIPMs by training various ML models using different sets of training features: data gathered by (1) doctors or (2) AHPs, and (3) a combination of both.

Our results show that naïve models had better prediction performance than health professionals, all VIPMs outperformed the performance of health professionals, and better estimates of visit intensities can be made using data collected and documented by the doctors and by the AHPs. We also show how changing visit intensities may affect the health outcome.

3 - Causal Effects and Optimal Policy Learning for Intensive Care Unit Discharge Decisions to Solve Hospital Process Bottlenecks

Justus Vogel, University of St. Gallen, St. Gallen, Switzerland, Johannes Cordier, Miodrag Filipovic

Intensive care units (ICUs) operate with fixed capacities and face demand uncertainty, triggering early discharges to free up beds. Such discharges can increase ICU readmissions, negatively impacting patient

outcomes and aggravating ICU bottleneck congestion. We investigate how ICU discharge timing affects readmissions, with the goal of developing policies that minimize readmissions.

Our sample comprises 12,950 ICU stays from the Cantonal Hospital of St. Gallen. To define a binary treatment, we randomly assign hypothetical discharge days to patients, comparing these with actual discharge days to form intervention and control groups. We apply two causal machine learning techniques and assume unconfoundedness, leveraging thousands of patient characteristics as sufficient covariates. For scenarios where unconfoundedness might fail, we discuss an instrumental variable approach. We further develop optimal decision policies based on individualized average treatment effects (IATEs).

We find that for 72% of our sample, a discharge at point in time t as compared to t+1 increases patients' readmission risk while 28% of cases profit from an earlier discharge. 91.4% of IATEs range between -0.05 and 0.05, i.e., earlier discharge causes a -55% to 55% relative change as compared to the average readmission rate of 9.04%. To develop decision policies, we will rank patients according to their IATEs and compare IATEs of optimal and actual discharges. Finally, we will assess the potential reduction in readmissions and saved bed capacities under optimal policies in a simulation.

We aim to provide a novel approach and blueprint for similar operations research and management science applications in data-rich environments.

4 - Revolutionize Rural Healthcare: Multimodal Network Design for Uav-Assisted Medical Delivery Service

Wenzhao Dong, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, Gangyan Xu

To enable timely delivery of medical supplies in rural areas where healthcare access is often limited, Medical Delivery Services (MDS) have undergone significant revolutionization through the development and use of Unmanned Aerial Vehicle (UAV). By overcoming geographic barriers, UAVs can transport critical supplies such as blood, vaccines, and medical equipment in minutes rather than hours. Although UAVs are with such advantages, the role of UAVs in supporting MDS requires further exploration. This study designs and evaluates four multimodal MDS networks connecting urban and rural areas: truck-UAV (T-U), truck-human (T-H), UAV-human (U-H), and UAV-only (U-U), with a focus on efficiency and cost trade-offs. Through simulation-based comparative analysis, we identify the optimal MDS network for different scenarios and specify the conditions under which UAV-assisted network outperform traditional methods. Our findings demonstrate the substantial potential of UAVs in enhancing MDS for rural healthcare and provide managerial insights for policymakers and healthcare providers to enhance medical supply chain efficiency.

5 - An Integrated Approach for Case Mix Inspired Patient Admission Scheduling

Zhaosheng Zhu, Tsinghua University, Beijing, China, People's Republic of, Xiaolei Xie

As hospitals face growing pressure to deliver high-quality, efficient care with limited resources, effective patient admission scheduling has become increasingly vital — it plays a crucial role in optimizing bed utilization, managing case mix, and balancing resource supply and demand in downstream inpatient processes. To address the challenges caused by siloed and myopic decision-making in current practice, which often leads to unbalanced bed usage and misaligned resource allocation, this study proposes a case mix inspired patient admission scheduling model that not only integrates admission scheduling with case mix optimization but also explicitly incorporates preoperative examination resource alignment. An integrated approach combining prediction, optimization, and simulation is then established, in which machine learning forecasts the preoperative examination resource needs of patients, an integer programming model optimizes admission scheduling by jointly considering case mix and resource constraints, and a discrete-event simulation evaluates the model performance using real hospital data. Simulation experiments comparing five scheduling strategies demonstrate that the proposed model significantly improves preoperative resource matching, reduces waiting time, and enhances overall scheduling efficiency. Furthermore, the model demonstrates robustness to prediction inaccuracies and practical applicability, underscoring its potential for real-world deployment in optimizing admission scheduling through case mix integration.

WA08

Bras Basah A

AI, Incentives, and Decision-Making in the Digital Economy

Invited Session

Informations Systems

Chair: Varun Karamshetty, National University of Singapore, Singapore, Singapore

1 - Feedback Design and Equilibrium Knowledge-Sharing

Xinbo Wang, National University of Singapore, Singapore, Singapore, Nan Chen

Feedback systems are widely used by digital platforms. We study how feedback design shapes equilibrium content provision in the context of Stack Overflow. Using a regression discontinuity design based on answer timing, we show that competition shapes feedback and influences contribution. We develop and estimate an empirical model that links feedback design to equilibrium outcomes among heterogeneous contributors. The model allows the platform to continuously adjust leniency (the tendency to give positive feedback) and harshness (the tendency to give negative feedback). We decompose feedback's impact into a direct effect on individual contributors and an indirect competitive effect. We find that although a contributor favors more leniency and less harshness when considered alone, she may prefer less leniency and higher harshness in equilibrium. This reversal arises from competition and depends on user quality, with higher-quality contributors being more likely to alter their preferences under equilibrium incentives. We further show how the platform's targeted quality levels, its ability to distinguish contributor quality, and the distribution of contributor skill shape competition and inform optimal feedback design. Finally, we consider the case of endogenous feedback preferences, where contributors' valuations of feedback depend on the equilibrium quality induced by the mechanism.

2 - The Value of AI-Generated Metadata for Ugc Platforms: Evidence from a Large-Scale Field Experiment

Xinyi Zhang, National University of Singapore, Singapore, Singapore, Chenshuo Sun, Renyu Zhang, Khim-Yong Goh

AI-generated content (AIGC), such as advertisement copy, product descriptions, and social media posts, is becoming ubiquitous in business practices. However, the value of AI-generated metadata, such as title, remains unclear on user-generated content (UGC) platforms. To address this gap, we conducted a large-scale field experiment on a leading short-video platform in Asia to provide about 1 million users access to AI-generated titles for their uploaded videos. Our findings show that the provision of AI-generated titles significantly boosted content consumption, increasing valid watches by 1.6% and watch duration by 0.9%. When producers adopted these titles, these increases jumped to 7.1% and 4.1% respectively. This viewership boost effect was largely attributed to the use of this generative AI (GAI) tool increasing the likelihood of videos having a title by 41.4%. The effect was more pronounced for the groups more affected by metadata sparsity. Mechanism analysis revealed that AI-generated metadata improved user-video matching accuracy in the platform's recommender system. Interestingly, for a video for which the producer would anyway have posted a title, adopting the AI-generated title will decrease its viewership on average, implying that AI-generated titles may be of lower quality than human-generated ones. However, when producers chose to cocreate with GAI and significantly revised the AI-generated titles, the videos outperformed their counterparts with either fully AI-generated or human-generated titles, showcasing the benefits of human-AI co-creation. This study highlights the value of AI-generated metadata and human-AI metadata co-creation in enhancing user-content matching and content consumption for UGC platforms.

3 - How Does Dynamic Algorithmic Pricing Affect Tipping Behavior on Two-Sided Platforms? Empirical Evidence from Uber

XIAOHUI LIU, National University of Singapore, Singapore, Singapore, LIANG ZHANG, Yujing Ma, Wei Ma, Hock-Hai Teo

Tipping on two-sided platforms not only plays a crucial role in workers' overall earnings but also serves as a social incentive to enhance their job satisfaction and motivation. While prior research on tipping behavior has predominantly examined the scenarios of fixed pricing, little is known about how dynamic pricing determined by the algorithm affects customers' tipping behavior. To bridge this research gap, this study empirically investigates the externalities of algorithmic pricing on customers' tipping behavior in the context of ride-hailing services on the Uber platform in New York City. The empirical evidence shows that a 1% increase in algorithm-adjusted fares is associated with a 0.209% decrease in the tip ratio from passengers. Two possible mechanisms explain this outcome: (1) The lack of algorithm transparency in the dynamic pricing may lead passengers to express dissatisfaction with the platform through reduced tip ratios toward drivers; and (2) Passengers, as the payers of both algorithm-determined fares and voluntary tips, may adjust their mental accounting in response to the algorithmic pricing. Our findings suggest that Uber's profit-driven pricing strategy may inadvertently harm drivers' interests, as it fails to account for these indirect effects on tipping. This study provides important insights for understanding the roles of on two-sided platforms, particularly for preventing inequitable exploitation of drivers' income as algorithmic pricing becomes increasingly prevalent.

4 - Enhancing Stock Return Prediction with LLM-Extracted News Events: Accurate and Interpretable Insights

Mingxuan ZHENG, National University of Singapore, Singapore, Singapore, Gang Li, Dandan Qiao

News-driven stock market prediction has demonstrated promising potential in recent studies, suggesting that such documents reflecting the changing states of the world help infer the asset market dynamics. However, existing models based on word embedding and sentiment analysis suffer from noise and bias in the materials. At the same time, these methods show weakness in interpretability, limiting their practical application for decision-makers. To address these challenges, we proposed a method utilizing large language models (LLMs) to extract cleaner and more interpretable event factors from news articles, along with a novel interpretable asset return prediction framework driven by event factors. Our method not only provides precise asset return predictions but also offers insights into the underlying relationships between key events and asset returns. Specifically, we first utilize a pre-trained LLM to extract key events from news texts, for which targeted prompts and algorithms are designed. Secondly, we propose the event-driven asset return prediction framework(EDARPF), which integrates news event factors and essential context features as inputs, and incorporates the multi-head self-attention mechanism to enhance model performance and interpretability. Finally, we conduct experiments on a real-world dataset from the Nasdaq stock market, which verify the prediction performance and the portfolio profit performance of the proposed method, while also providing interpretations for stock movements based on news narratives. Our study highlights the advantages of using LLMs in financial applications, showing that effective extraction and interpretation of news can lead to more informed and accurate decision-making support.

WA09

Bras Basah B

Sustainable and Environmental Optimization

Contributed Session

Optimization

Chair: Juan Carlos Espinoza Garcia, Tecnologico de Monterrey, Epigmenio Gonzalez 500, Queretaro, 76130, Mexico

1 - Developing a Centralized Data Envelopment Analysis Model with Nonseparable Variables to Assess the Allocation Planning of Air Pollution Emissions Abatement Among Fossil Fuel Power Plants

Li-Ting Yeh, Feng Chia University, Taichung, Taiwan

With heat, power generation, and air pollution emissions being jointly produced from fossil fuel consumption, previous research has typically employed standard Data Envelopment Analysis (DEA) models with nonseparable variables. However, since these power plants are under centralized management, it is crucial for decision-makers to optimize the overall efficiency of all fossil fuel power plants, rather than focusing solely on the individual efficiency of a single fossil fuel power plant. Therefore, this study proposes a centralized DEA model with nonseparable variables. By applying the model proposed in this study, Taiwanese decision-makers can optimally

allocate strategies for air pollution emissions abatement and resource input reduction across its 57 fossil fuel power plants. The results of this study are expected to provide Taiwanese decision-makers with strategies to enhance the overall efficiency of fossil fuel power plants.

2 - Air-Quality and Crowd-Aware Pedestrian Routing: A Multi-Criteria Optimization Approach

Josep Maria Salanova Grau, Centre for Research and Technology Hellas, Thessaloniki, Greece, Thomas Dimos, Fokion Karavasilis

As urban populations grow, the need for healthier pedestrian mobility becomes increasingly critical. Long-term exposure to air pollution poses serious health risks, particularly for vulnerable groups including children, the elderly, and individuals with respiratory conditions. This study presents an air-quality and crowd aware pedestrian routing framework that helps city residents navigate urban areas while minimizing exposure to pollutants.

The framework utilizes high-resolution pollution maps to identify areas of increased exposure and support environmentally informed route choices. Separately, pedestrian density data is estimated from telecom operator data, enabling the system to estimate crowd levels across the network. These inputs inform distinct routing objectives—minimizing exposure to pollutants or avoiding congested walkways—depending on user needs and system conditions.

Routing decisions can be optimized to either minimize individual pollutant dose, avoid crowded streets or reduce the marginal cost each user imposes on the network. This formulation enables the system to assign users to paths that balance user-optimal and system-optimal objectives, ensuring equitable use of walking infrastructure. The optimization challenge is addressed using a modified Dijkstra's algorithm with time-dependent costs, enhanced by a weighted objective function that integrates multiple criteria of full and marginal costs—pollution exposure, travel time, and crowd density—for efficient and adaptive pedestrian routing.

The implementation in Thessaloniki, Greece, demonstrates how environmental and behavioral data can be leveraged to inform healthier urban mobility strategies. Results highlight the value of combining environmental modeling with optimization theory to support data-driven pedestrian planning that promotes public health and network-wide efficiency.

3 - Locating Shared Wastewater Treatment Plants Under Flow Uncertainty

Amit Vatsa, Indian Institute of Management Indore, Indore, India, Saurabh Chandra

Pollution of freshwater sources is an enormous issue in many developing nations. Government regulations mandate that firms using water for industrial purposes treat effluent before releasing it and recycle used water as much as possible. Large enterprises can manage setting up a water treatment plant. However, this is financially and technically challenging for medium and small-scale enterprises. Government policies promote the creation of a set of shared water treatment plants (SWTPs) that can service all the enterprises in an industrial cluster and make it more economical to operate. We formulate the problem of locating SWTPs of appropriate sizes and allocating polluting firms to them as a Mixed Integer Non-linear Program (MINLP). The model minimizes the setup and operating costs for the network of SWTPs.

The operating cost, dependent on the wastewater volume and pollutant concentration, is a non-convex function, as evidenced in various sectors and literature. We provide an exact convexification strategy for this MINLP. We extend our model to consider the uncertainty in the volume and concentration of the wastewater generated by the polluting firms. We study a stochastic programming version of this problem. We propose a Generalized Benders decomposition based algorithm with many refinements. We exploit the structure of the problem to make solving the sub-problem very efficient. Our solution method could solve the practical size instances of this problem, which could not be accomplished using commercial solvers.

4 - A Hybrid Ga-Lns Metaheuristic for Sustainable Vehicle Routing with Hybrid Fleets

Juan Carlos Espinoza Garcia, Tecnologico de Monterrey, Queretaro, Mexico, Rafael Batres, Beatriz Murrieta Cortes

Urban freight logistics is a key contributor to pollution, congestion, and road degradation. The Vehicle Routing Problem (VRP) addresses the efficient distribution of goods, but exact solutions become impractical for complex, large-scale scenarios due to computational limitations. As a result, heuristic and metaheuristic algorithms have gained prominence for their ability to generate high-quality, near-optimal solutions within realistic timeframes. At the same time, as environmental concerns—such as rising emissions, poor air quality, and climate change—grow more urgent, integrating sustainability into routing decisions has become essential, particularly since transportation is the most polluting link in the supply chain.

This study proposes a novel hybrid algorithm to address the VRP within the context of a hybrid fleet composed of electric and combustion vehicles. Recognizing the operational complexities associated with electric vehicles —such as limited range, lower capacity, and infrastructure dependency—we developed an integrated metaheuristic combining Genetic Algorithms (GA) and Large Neighborhood Search (LNS). The LNS is embedded as a mutation operator within the GA, leveraging the Clarke-Wright savings method to restructure routes dynamically.

The proposed approach was tested across diverse customer scenarios, demonstrating superior performance in reducing both environmental impact and operational costs compared to conventional GA solutions. This contribution underscores the potential of hybrid algorithmic frameworks to simultaneously advance computational efficiency and support sustainable urban logistics.

WA10

Hullet Room

Dynamic Pricing and Competition

Invited Session

Revenue Management and Pricing

Chair: Weifen Zhuang, Xiamen University, Xiamen

1 - Dynamic Pricing of Limited Inventories with Word-of-Mouth Effect

Yan Liu, University of Science and Technology of China, Hefei, China, People's Republic of

We explore a finite-horizon, finite-capacity dynamic pricing model where the arrival customer is affected by the word-of-mouth (WoM) effect. In particular, the existing buyers serves as an advertising that attracts more customers into the system through the WoM effect. Hence the seller faces an exploration–exploitation trade-off between enhancing the demand through the WoM effect and optimizing price for revenue maximization. We first study the fluid model by suppressing the randomness of demand, and find that the seller employs a penetration-pricing policy under WoM—that is, increasing the selling price as time elapses. Moreover, he balances the underlying exploration–exploitation process by also keeping the sales rate gradually increasing. We further characterize the dynamics of time and inventory due to stochastic demand, both analytically and numerically. To our surprise, we find that the classic time-monotonicity rationale does not hold. In particular, given inventory level, we find that the seller would first lower the price but then raise it as the selling time becomes shorter, which are driven by traditional time-scarcity effect and penetration-pricing of WoM, respectively. More interestingly, we find that the marginal value of inventory is negative at the end of the selling season. This is because inventories sold to customers can attract more customers under WoM but have no value at the end.

2 - Dynamic Price Competition of Substitutable Perishable Products Among Multiple Sellers

Boqian Song, Southwestern University of Finance and Economics, Chengdu, China, People's Republic of, Michael Z.F. Li, Surui Wang, Weifen Zhuang

In this paper, we study a multi-period price-setting stochastic game involving multiple sellers, each offering a perishable product with limited capacity over a fixed time horizon. Employing a stylized linear demand rate model to capture consumer choice behavior, we show that any sub-game perfect Nash equilibrium depends solely on state variables through the reservation costs of sellers with stock on hand. When the Nash equilibrium

is interior for sellers with stock on hand, the equilibrium price set by a seller increases with all reservation costs. However, the corresponding equilibrium demand rate decreases with the seller's reservation cost while increases with competitors' reservation costs. To demonstrate the complexity in the equilibrium behavior with respect to state variables in the stochastic game, we further examine its deterministic counterpart. Counterintuitively, in the equilibrium of the deterministic game, the prices offered by sellers holding relatively high initial inventories can increase with the initial inventories of their competitors and decrease with the horizon length. In situations with incomplete real-time inventory information, we propose four pricing heuristics---EIP, QRIP, OLP and r-OLP---which estimate competitors' inventory levels using the up-to-date sales and price path data. These heuristics then set each seller's price by leveraging equilibrium results from the stochastic or deterministic game under complete information. Extensive numerical studies demonstrate the effectiveness of EIP and QRIP, with QRIP being robust across various market conditions. Notably, EIP and QRIP generally outperform OLP and r-OLP, demonstrating the value of exploiting reservation costs under incomplete inventory information.

3 - Dynamic Pricing Competition Among Multi-Flight Airlines

Yifan HE, Nanyang Technological University, Singapore, Singapore

The airline industry operates in a highly competitive environment, where passengers can compare fares and schedules in real-time across numerous flights. In response, airlines rely on dynamic pricing to adjust fares based on evolving market conditions, demand fluctuations, and competitor actions. While existing revenue management (RM) systems leverage machine learning and artificial intelligence to optimize real-time pricing decisions, they primarily adopt a reactive approach, adjusting fares in response to immediate conditions without considering long-term strategic interactions among competitors. This short-term focus may lead to suboptimal pricing decisions over the entire selling horizon, ultimately reducing profitability.

To address this gap, we develop a competitive dynamic pricing model in which multiple airlines, each operating several flights on the same route, adjust fares dynamically over a finite selling horizon. Our model captures both internal (within-airline) and external (cross-airline) competition, establishing the existence of equilibrium and characterizing optimal pricing strategies shaped by the marginal value of inventory. Through numerical experiments, we uncover complex price and profit dynamics, showing that optimal fares evolve non-monotonically with time and inventory, and that competition influences pricing trajectories in unexpected ways. Our results show that dynamic pricing consistently outperforms fixed pricing strategies, particularly under conditions of high inventory scarcity and long selling horizons. By providing insights into optimal price adjustments under competition, our study advances the theory of dynamic pricing in airline revenue management and offers managerial implications for pricing optimization in competitive markets.

4 - Strategic Implications of Free Cancellation Policies in Competitive Airline Markets

Surui Wang, CUHK Shenzhen, Shenzhen, China, People's Republic of, Zizhuo Wang

We developed a Nash equilibrium model to analyze the impact of free cancellation policies on airline competition. We study a duopoly market with two airlines offering fixed fares across two booking periods: an early period with lower fares and cancellation uncertainty, and a late period with higher fares and certain travel needs. Passengers choose between booking early, considering potential cancellations, or waiting until their travel plans are confirmed. Airlines strategically determine their booking limits and whether to provide refundable tickets. Using a utility-based economic framework, we explicitly model passenger decisions based on expected utility maximization. We derive closed-form equilibrium booking limits and refund policies. Our analysis shows that in moderately competitive markets without extensive customer spillover, airlines adopt free cancellation policies as a dominant strategy, leading to higher early bookings and improved revenues. However, under intense competition with substantial passenger spillover, refundable policies may trigger aggressive early bookings that reduce overall profitability. This study provides theoretical insights into how free cancellation policies strategically impact airline revenue management and competitive interactions.

WA11

Morrison Room

Information Design and Searching Problems

Invited Session

Revenue Management and Pricing

Chair: Dongwook Shin, HKUST Business School, Clear Water Bay, Hong Kong

Co-Chair: Wenjuan Li, The Hong Kong University of Science and Technology, Hong Kong, N/A

1 - Sequential Vs Parallel Search Under Competition

Yufei Zhu, National University of Singapore, Singapore, Singapore, Zhi Chen, Zhenyu Hu

Concept generation and new product development (NPD) are essential for business success, as they enable firms to stay innovative in a highly competitive market. However, due to the inherent uncertainty and complexity of the innovation landscape, it is essential that firms thoroughly search across the space of possible concepts, thereby increasing their chances of discovering a winning idea. In practice, sequential and parallel search strategies represent two fundamental approaches to concept generation and NPD. However, firms oftentimes do not conduct concept generation internally but instead delegate it to multiple external agents, such as suppliers or service providers. In this paper, we ask: in such a competitive, multi-agent setting, should the buyer prefer each supplier to conduct a sequential or parallel search? We employ a contest approach to answer this question. Our analysis reveals that when the unit search cost is high or the maximum number of evaluation opportunities is large, the buyer prefers sequential search. Otherwise, parallel search is preferred.

2 - Wait TIME Announcements to Optimize Customer Satisfaction

Jeunghyun Kim, Korea University, Seoul, Korea, Republic of, Laurens Debo, Robert Shumsky

We develop a model to optimize the provision of queue-position information to waiting customers. A customer’s emotional response, and satisfaction with the wait, may be driven by loss aversion in expected waiting time and/or by disutility due to wait-time uncertainty (we call this latter effect ‘uncertainty awareness’). We consider a class of announcement strategies characterized by frequency of updates between two announcements one upon arrival and the other upon entering service. We control the frequency by the number of queue-position advancements needed to hear the next announcement and optimize it. We find that when customers are purely loss-averse, a strategy that never provides intermediate updates (arrival-only) is optimal while with purely uncertainty-aware customers, a strategy that updates customers with every single step (full information) is optimal. For customers affected by both loss aversion and uncertainty, there is an interior solution and the degree of customer satisfaction improvement (relative to the arrival-only and full information) becomes more substantial as the system becomes more loaded.

3 - Dynamic Persuasion Strategies for Mitigating the Spread of Fake Content

Wenjuan Li, HKUST Business School, Hong Kong, China, People's Republic of, Dongwook Shin

This paper examines how digital content platforms can design signaling mechanisms to shape consumers' engagement with potentially fake content. Initially, the platform and consumers face uncertainty about content veracity, but the platform can conduct random inspections to verify it. Based on inspection results, the platform strategically sends signals to consumers. We investigate both public and private signaling policies to maximize revenue. Our findings show that both policies disclose inspection results with strategic *delays*, which enhances platform revenue but also increases consumers' exposure to fake content. Furthermore, unlike the optimal public policy, we find that the optimal private policy selectively discloses the inspection results to some consumers; consequently, the optimal private policy generates higher revenue than the optimal public policy but significantly undermines consumer welfare. We explore various strategies to reduce the spread of fake content while ensuring profitability. We find that enhancing the platform's inspection capability boosts the platform's profits but fails to curb the spread of fake content. Alternatively, our analysis suggests that implementing subscriptions and government penalties can effectively reduce fake content while enhancing the platform's revenue.

4 - Information Design and Pricing In Lead-Selling Platforms with Supply Competition

Niloofar Zamani Foroushani, Imperial College London, London, United Kingdom, Yanwei Sun, Zhe Liu, Jiahua Wu

We analyze a type of platform common in the service sector, where customers can freely post tasks, and service providers pay a fee to purchase job leads and unlock the ability to communicate with these customers. We adopt a Bayesian persuasion framework to explore how the platform can leverage its informational advantage about the provider base and pricing to strategically influence the decisions of service providers, aiming to maximize its revenue. We begin the analysis by assuming that prices are exogenously given and characterize the structure of the optimal signaling mechanism. Specifically, the platform recommends that all service providers participate when the pool size is below a certain threshold. When the pool size exceeds this threshold, the platform randomizes its recommendations between two consecutive numbers, independent of the pool size. This behavior is driven by a unique feature of our model, where the beliefs of the platform and service providers are non-identical: a provider's knowledge of their own participation inflates their belief about the total number of service providers. Next, we endogenize the platform’s pricing decision and consider the policy space where the signaling mechanism and (possibly randomized) state-dependent pricing are jointly optimized. Notably, within this space, we find that the optimal signaling mechanism with fixed pricing is optimal. This result implies the superiority of the optimal signaling mechanism over the commonly adopted contingent pricing policy.

We show that the performance discrepancy between these two policies is particularly pronounced when supply competition is intense.

WA12

Moor Room

Transportation Electrification

Invited Session

Transportation Electrification

Chair: Prateek Bansal, National University of Singapore, Singapore, 129790, Singapore

1 - A Multi-Phase Stochastic Planning of Charging Infrastructure for Electric Vehicles

Qiming Ye, Singapore ETH Centre, Singapore, Singapore, Prateek Bansal, Bryan Adey

The rapid electrification of ride-hailing mobility requires that fast-charging infrastructure be ready to support the dynamic charging needs of shared electric vehicle (SEV) fleets. Deploying city-wide fast-charging stations requires determining the timing, locations, and number of chargers at each station over multiple planning phases, while managing stochastic charging demand, induced demand from enhanced services, and constraints like grid capacity. In contrast to conventional deterministic and myopic methods that optimise charging stations' configurations based on static demand and immediate benefits, this study formulates the problem as a stochastic sequential decision problem and applies reinforcement learning (RL) to approximate the optimal deployment plan considering long-term benefits and uncertainties. An agent-based model representing an AI gym environment is developed to train the RL model, simulating SEV interactions with charging stations to capture stochastic charging demand across space and time. RL maximises the expected charging service efficiency with the minimal number of fast chargers over the planning horizon. The model reduces the total number of chargers by 30.5% compared to the myopic benchmark, while achieving a 10.9% higher charging service efficiency. This multi-phase fast charger planning approach offers engineers and planners a novel tool for efficient long-term infrastructure allocation and policy development under stochastic charging demand.

2 - Operations Management and Optimization for Demand-Responsive Autonomous Electric Taxis System

Simon Hu, Zhejiang University, Hangzhou, China, People's Republic of

As urban mobility systems evolve, taxi services are increasingly integrating autonomous electric taxis (AETs) alongside human-driven electric and gasoline vehicles, creating complex mixed fleets. Traditional operations management approaches—designed for homogeneous systems—struggle to address dynamic, real-time demands and interactions within such heterogeneous environments. This presentation introduces an integrated operations management and optimization framework tailored for demand-responsive mixed taxi fleets, focusing on AETs.

We develop an integer linear programming model to optimize vehicle assignment and scheduling, aiming to maximize system-wide profitability while considering customer service responsiveness, energy costs, and travel

efficiency. The model supports both centralized and decentralized control strategies, reflecting the operational characteristics of autonomous and human-driven vehicles. To evaluate real-time system performance, an agent-based simulation platform is employed, capturing the dynamic interactions among vehicles, passengers, and charging infrastructure. Case studies utilizing real-world data highlight the operational advantages of AETs, particularly in terms of lower running costs and improved environmental outcomes. Impact assessments demonstrate that AETs significantly reduce carbon emissions per kilometre and per ride while maintaining competitive service quality. This work offers actionable insights for platform operators and urban mobility planners seeking to optimize operations and accelerate the transition toward sustainable, demand-responsive transport systems.

3 - Can Decoy Nudge Preferences for Ridehailing Drivers Toward Electric Vehicle?

Prateek Bansal, National University of Singapore, Singapore, Singapore

Robust policy portfolios are established to encourage the adoption of electric vehicles (EVs) worldwide to realize ambitious net-zero emission targets. However, the EV adoption rate remains low, highlighting the urgent need for more innovative marketing strategies. We investigate whether a decoy effect can nudge consumer preferences toward EVs. The decoy effect occurs when introducing a new alternative to an existing choice set increases the attractiveness of the target alternative. Based on the street-intercept survey of Singaporean ride-hailing drivers, we demonstrate that the decoy effect encourages them to rent EVs over internal combustion engine vehicles, with a more prominent effect among young drivers. We also identify the appropriate model to explain the preference shifts induced by the decoy alternative. Our results show that the multi-alternative Decision by Sampling (MDbS) model generally outperforms other sequential sampling models and utility-based models. Therefore, we use MDbS parameters to estimate the optimal configuration of the decoy to maximize its effect on preference shifts. Our study presents the first real-world application of the decoy effect in the car rental market, demonstrating its potential to nudge preferences for EVs by integrating MDbS into recommender systems.

4 - Simulation-Based Evaluation for Ev Charger Infrastructure Planning

Muhamad Azfar Ramli, Institute of High Performance Computing, Singapore, Singapore

Electrification is a key strategy to help cities decarbonise the transport sector. However, transitioning from combustion engine vehicles that refuel at petrol/diesel stations to battery-electric vehicles (EVs) that receive energy from the electrical distribution network, transport and energy urban planners face significant challenges in ensuring the country’s charging infrastructure can cost-effectively meet its transportation needs.

In this presentation, I will showcase our data-driven approach to help urban planners tackle key issues such as determination of optimal placement of chargers to meet mobility demand, and help energy planners plan for sufficient grid infrastructure upgrades to support the resultant growth of electricity demand from these charging events.

SITEM's advanced integrated modelling framework combines four key aspects of transport electrification, namely a) behavioural modelling and driver decisions, b) geospatial optimisation of mobility demand, c) high performance discrete event simulation of multimodal transport across the entire city and d) high-fidelity model of the distribution grid network.

Through a novel integrated combination of these four sub-models, we demonstrate our capability that is able to validate the adequacy of chargers across the city of Singapore and determine the amount of upgrades required to support EV charging demand for various different modelled electrification scenarios.

Our approach has also been extended and adapted to help plan the charging infrastructure electrification for local harbourcrafts, as well as plan for chargers to support electrification of the fleet of airside vehicles required to service flights in Singapore’s Changi Airport.

WA13

Bailey

Complex Systems in Transition: Modeling, Migration, and Organizational Dynamics

Contributed Session

Complex and Dynamic Systems

Chair: Xian Chen, Shanghai University, Shanghai, N/A, China, People's Republic of

1 - Coexistence of Platforms: a Niche Perspective

Jin Xu, Liaoning University, Shenyang, China, People's Republic of

Although platform competition is the focus of platform literature, the dynamic interactions of platforms remain underexplored. Drawing from niche theory, this paper discusses the simultaneous coexistence of platforms within a highly competitive environment. Given that niche theory's expressions and calculation methods can be simplified only by the quantity of the research entities in most related research, this study develops a meticulous approach to express and describe niches through data mining and content analysis techniques. Furthermore, we analyze the most recent longitudinal variations in the niche indices of these platforms. The framework and methodology proposed in this paper demonstrate that niche theory holds significant potential for exploration within the context of the platform economy. Through the content analysis of users' comments, a platform can realize satisfaction differences across platforms, therefore recognizing the gap between itself and its competitors. Platforms can strategically position themselves in a niche with the appropriate width and overlap and optimize their service or product design through niche analysis. By conducting an in-depth investigation into the coexistence of platforms, this research offers insights into the survival and growth of platform enterprises as well as the sustainable symbiotic development of the platform industry.

2 - Resilience Analysis of Shared Manufacturing Systems in Industrial Iot Environments: a Hypernetwork Modeling Approach

Jiangnan Cheng, University of Science and Technology Beijing, Beijing, China, People's Republic of

In the deeply integrated manufacturing environment of Industrial IoT and Cyber-Physical Systems (CPS), Shared Manufacturing Model (SMM) significantly enhances supply chain efficiency, while its inherent uncertainties in resource coordination and dynamic scheduling pose critical challenges to system resilience. To address this issue, this research constructs a dual-layer interdependent hypernetwork model to systematically investigate resilience influencing factors and optimization strategies. First, a service-task bilayer hypernetwork architecture is established by integrating heterogeneous manufacturing services and real-time scheduling requirements. Second, an innovative composite resilience metric combining functional and structural resilience is proposed, accompanied by a dynamic cascading failure analysis model. Finally, simulation experiments evaluate the system's resilience under three scenarios: random attacks, targeted attacks, and hybrid disruptions. The experimental results reveal that the system exhibits distinct failure mechanisms under different attack scenarios, accompanied by pronounced phase transition characteristics during performance degradation. This research provides theoretical foundations for optimizing resource allocation and strengthening anti-interference capabilities in shared manufacturing systems, advancing the design of self-adaptive manufacturing systems.

3 - Managing Organizational Change Using a Systems Approach

Harshal Lowalekar, Indian Institute of Management-Indore, Indore, India, Shelja Kuruvilla, Durba Banerjee

This study investigates the use of systems thinking (ST) to support organizational change. While ST has been praised for fostering stakeholder buy-in, there is limited empirical evidence of its effectiveness. Through qualitative case research in diverse organizational settings, the study assesses ST's ability to achieve consensus and enhance performance. A detailed longitudinal case study is paired with cross-case comparisons. The findings show that systematic use of ST tools promotes change acceptance and delivers measurable performance gains, offering strong empirical support for ST as an effective change management methodology.

4 - Detecting Multiple Changepoints by Exploiting Their Spatiotemporal Correlations: a Bayesian Hierarchical Approach

Xian Chen, Shanghai University, Shanghai, China, People's Republic of, Hai Jiang

Capturing the non-stationarity of spatiotemporal data over time via changepoints has received increasing attention in various research fields. Although extensive studies have been conducted to investigate changepoint detection with spatiotemporal data, research on detecting multiple clusters of spatiotemporally correlated changepoints has remained unexplored. In this paper, we propose a multi-layer Bayesian hierarchical model: The first layer uncovers the spatiotemporal correlations of changepoints based on multiple propagation binary variables, which describe the occurrences of change propagations. The second and third layers compose non-homogeneous hidden Markov models to capture time series data and their state sequences, in which changes of states signify changepoints. We perform Bayesian inference for changepoints and change propagations via a forward-backward algorithm that combines recursion and Gibbs sampling. Based on the experiments with simulated data, we show that our method significantly improves the detection accuracy toward spatiotemporally correlated changepoints. A real-world application to bike-sharing data also demonstrates the effectiveness of our method. This research has significant relevance to companies operating systems across geographical regions, as it enables a more robust understanding of emerging trends and shifts in spatiotemporal data.

WA14

Wilberforce

Global Supply Chain Dynamics and Policy Impacts

Contributed Session

Supply Chain Management

Chair: Shanshan Zhang, South China University of Technology, Guangzhou, N/A, China, People's Republic of

1 - Product Market Concentration and Supply Base Complexity: a Network Community Perspective

Kedong Chen, Rensselaer Polytechnic Institute, Troy, NY, United States, Xiaojin Liu, Pankaj Kumar, Feng Cheng

Firms face a strategic dilemma in designing their supply bases: whether to consolidate suppliers to enhance efficiency and coordination or to diversify them to improve resilience and foster innovation. While prior research has examined this trade-off through the lenses of risk management, competition, and innovation, the role of product market concentration in shaping supply base design remains underexplored. This study addresses this gap by empirically examining how product market concentration influences three dimensions of supply base complexity -- supply base size, industry diversity, and network community diversity. We construct a unique dataset of over 4,900 U.S. public firms from 2003 to 2019 by integrating FactSet Supply Chain Relationship data, TNIC-3 product market data, and Compustat data. Using machine learning to detect supply network communities and an instrumental variable approach with a control function to identify causal effects, we find that firms operating in more concentrated markets tend to reduce supply base complexity: they work with fewer suppliers, source from fewer industries, and concentrate relationships within fewer network communities. These patterns suggest that firms in concentrated markets prioritize consolidation over diversification to enhance efficiency and strengthen inter-organizational relationships. This study offers novel insights into supply base design and its antecedents by linking macro-level market structures with micro-level sourcing decisions and introducing network community into the concept of supply base complexity.

2 - Partisan Supply Chains: the Impact of Political Ideology on Global Sourcing

Ben Charoenwong, INSEAD, Fontainebleau, France, Jie Peng, Jing Wu

In today's politically polarized and interconnected global economy, firms face escalating risks from ideological shifts in foreign governments. This study addresses the critical question of how political ideology influences global sourcing strategies, particularly focusing on the impact of ideological differences between firms and foreign governments on sourcing decisions. Using a quasi-experimental design centered on foreign elections, we find that an increased ideological distance between a U.S. firm and a foreign government leads it to reduce imports from those countries. Firms with high R\&D intensity are particularly sensitive to these ideological changes, while those with long-established supply chains demonstrate greater resilience. The impact of ideological distance is amplified in countries with strong institutions, where ideologically driven regulations are more likely to be enforced. Ex-post analysis reveals that firms continuing to source from ideologically divergent

countries face heightened ESG-related risks, including anti-corruption challenges and environmental violations. Our findings offer actionable guidelines for managers navigating politically diverse supply chains. Firms should proactively assess the alignment of their values with sourcing countries to mitigate operational and reputational risks, particularly in high R\&D sectors where intellectual property protection is critical. Regularly evaluating sourcing strategies in light of ideological landscapes can help firms avoid costly disruptions and protect their public image in an era of increasing geopolitical uncertainty.

3 - The Role of Public Policy in Enhancing Supplier Entry and Supply Chain Resilience

Ben Charoenwong, INSEAD, Singapore, Singapore, Christopher Chen, Alan Kwan, Jing Li

Using data on customer searches and product listings from North America's largest online platform for supplier sourcing, we study supplier responses to demand in the form of entry into the market for COVID-19 products during the recent pandemic. While many suppliers produced things closely related to COVID-19 products, only a fraction entered the market. We study the effects of government policy on demand responsiveness, measured by searches on the platform for specific products. Firms in states with more fiscal support or search behavior by government agencies respond more to demand, particularly for firms with prior government contract experience. We find that suppliers responded more to demand in states with more flexible procurement practices generally and whose state procurement authorities reduced supplier requirements during the pandemic. In addition, we find firms are more demand-responsive in states with more business-friendly infrastructure (lower taxes, better government IT). These results suggest a strong role for government buying during a pandemic.

4 - Cascading Effects of Regulatory Changes on Supply Chain Disruptions Amid International Environmental Turbulence

Xiaohui Jia, Dalian Maritime University, Dalian, China, People's Republic of

Amid international environmental turbulence, those low-likelihood and high-impact “black swan” events can no longer be overlooked as critical sources of unexpected supply chain disruptions. While existing research has extensively addressed strategies for enhancing supply chain resilience against the disruptions, empirical evidence remains limited regarding the origins of such disruptions and the mechanisms through which they propagate across supply chain networks. We propose that changes in national policies, regulations, and laws are emerging as critical drivers of supply chain volatility. Adopting a risk interaction perspective, we examine how regulatory changes initiate cascading disruptions throughout supply chains. Using Bayesian Nonlinear Structural Equation Modeling to analyze empirical data from industry experts, we uncover the propagation pathways and intensity of low-likelihood but high-impact risks triggered by regulatory changes. We provide robust empirical support for cascading effects of regulatory changes across supply chain networks. Moreover, we identify segments of the supply chain that are particularly vulnerable to the cascading effects, offering targeted insights for risk mitigation. By shedding light on the complex, nonlinear propagation of the risks, the findings contribute to the growing body of knowledge on supply chain resilience and assist in managing the effects of regulatory changes in the context of global turbulence.

5 - The Configuration Analysis of Internal R&D and Supply Chain Structure in Influencing Innovation

Shanshan Zhang, South China University of Technology, Guangzhou, China, People's Republic of, Jinyu Yang

Purpose: The purpose of this paper is to investigate how the impact that the internal R&D has on the innovation performance is dependent on the level of supply chain concentration. This study proposes a moderation model to test the strategic fit relationship in the Chinese Listed companies. Specifically, this study tests the moderation effect of supplier concentration and customer concentration on the relationship between R&D intensity and innovation performance.

Design/methodology/approach: Drawing upon the theory of strategic fit (Venkatraman, 1989) , the proposed relationships are tested with the Chinese listed companies from 2007 to 2021. Fixed effects (within-group) panel regression models are used to test the hypotheses.

Findings: Supplier concentration weakens the positive relationship between internal R&D and innovation performance. However, customer concentration strengthens the positive relationship between internal R&D and

exploitative innovation.

Relevance/contribution: Based on the argument of strategic fit theory, this study provides novel insights into the issue of how internal R&D may impact innovation performance. It investigates the effects of fit between R&D strategy and supply chain structure. First, this study indicates that the effect of internal R&D on exploratory and exploitative innovation is negatively contingent upon firms' supplier concentration. Second, this study suggests that the effect of internal R&D on exploratory and exploitative innovation is positively contingent upon firms' supplier concentration.

WA16

VIP A

Reliability Engineering and Maintenance Optimization

Contributed Session

Quality, Statistics and Reliability

Chair: Ipek Kivanc, N/A

1 - Warranty Cost Prediction for *K*-Out-of-*N* Systems with Non-Homogeneous Sales and Failure Processes

Huang Yichen, Sino-US Global Logistics Institute, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Xin Wang

Warranty cost forecasting is a critical pillar of effective warranty management, essential for ensuring both the financial resilience and operational reliability of manufacturing enterprises. However, accurate warranty cost estimation is fundamentally challenged by three intertwined sources of uncertainty: the stochastic nature of product sales, the randomness of product failures, and the increasing complexity of product systems. To address these challenges, this study develops a *k*-out-of-*n* system-based model that incorporates non-homogeneous sales and failure processes, effectively capturing the time-dependent uncertainties inherent in warranty cost prediction. The k-out-of-n system serves as a flexible and generalized reliability model, encompassing series systems and parallel systems as special cases, thereby enabling unified analysis across a broad range of system configurations. Closed-form expressions for the expected value and variance of cumulative warranty costs over time are derived, enabling efficient and reliable estimation of warranty obligations and their variability. To enhance practical applicability, we further extend the framework to accommodate five classic system configurations: series, parallel, series-parallel, parallel-series, and bridge systems, providing closed-form solutions for warranty costs across these architectures. The proposed model offers a comprehensive approach for cumulative warranty cost prediction, adaptable to complex system structures and varying time horizons. Numerical analyses demonstrate the model's flexibility in estimating warranty costs across diverse product types and its utility in formulating both short-term and long-term reserve planning strategies aligned with managerial objectives.

2 - Optimal Replacement of Mission-Based Systems with Partially Observable Missions

Bora Cekyay, Yildiz Technical University, Istanbul, Turkey, Suleyman Ozekici

We consider an optimal replacement problem for a mission-based system with a partially observable mission, where the phases of the mission modulate the system's degradation process. The system is partially observed at predetermined time points, where perfect information related to the system's degradation level (fully observable) is obtained, and only partial information related to the phase of the mission (partially observable) is obtained. We formulate the optimal replacement problem as a partially observable Markov decision process to minimize the long-run expected total discounted cost. We provide sufficient conditions to guarantee a control-limit type optimal replacement policy.

3 - Integrated Optimization of Cloud Service Pricing and Acquisition

Yuewei Zhang, Tsinghua University, Beijing, China, People's Republic of, Steel Chen, Taiqiang Lv, Zhi-Hai Zhang

In this research, we conduct an in-depth analysis of the integrated optimization of cloud service pricing and acquisition. A mathematical model tailored for multi-stage scenarios is formulated, and it is demonstrated that

the problem falls within the category of NP-hard problems. For 2 short-period special cases, closed-form solutions are derived. Following this, an efficient heuristic algorithm is developed. Additionally, the optimal solution structures for cases with fixed probability structures and base demand are characterized, thereby revealing that partial replenishment strategies previously proposed in the literature are suboptimal. Subsequent numerical experiments, which utilize synthetic data reflecting the common demand patterns of Tencent Cloud and real-world data from Aliyun, corroborate the efficacy of our modeling approach and the proposed algorithm. These experiments also uncover business insights that deviate from conventional inventory management practices.

4 - Sustainable Maintenance Strategies for Multi-Component Systems

Ipek Kivanc, Eindhoven University of Technology , Eindhoven, Netherlands, Claudia Fecarotti, Neomie Raassens, Geert-Jan van Houtum

Original equipment manufacturers are increasingly shifting from product-centric models to service-oriented strategies, placing greater emphasis on sustainable maintenance practices that minimize unplanned interventions. In this context, preventive maintenance plays a crucial role in improving operational efficiency, reducing costs, and meeting customer expectations. When properly implemented, such strategies can significantly reduce the need for corrective maintenance, enhancing system availability, performance, and workforce well-being. Despite these advancements, maintaining high-tech systems remains a complex task. These systems comprise multiple heterogeneous components, each with distinct degradation and failure behaviors, as well as varying levels of criticality to system functionality. Consequently, effective maintenance often requires a combination of age-based, condition-based, inspection-based, and failure-based maintenance approaches, tailored to each component's criticality.

To manage the maintenance of these complex systems, we propose and evaluate four integrated maintenance policies that combine scheduled preventive, semi-urgent, and opportunistic maintenance actions. We distinguish between critical and semi-critical components—where the failure of the former halts the system, while the latter only reduces performance. Semi-urgent maintenance visits are introduced to promptly address failed semi-critical components or those nearing critical degradation. These visits are arranged within a few business days to reduce disruption and workforce stress. Additionally, opportunistic maintenance allows for the simultaneous replacement of multiple components during these visits. Routine scheduled visits remain but may be postponed if a recent semi-urgent visit has occurred. Through numerical experiments, we assess the effectiveness of each policy in minimizing costs and corrective actions, highlighting their potential to improve maintenance outcomes in high-tech systems.

Wednesday, July 23, 10:00 AM - 11:30 AM

WB01

Enterprise

Strategic Decision-Making and Optimization

Invited Session

Applied Probability Society

Chair: Jinting Wang, Central University of Finance and Economics, 39, Xueyuan South Road, Haidian District, Beijing, 100081

1 - Strategic Joining in a Single-Server Batch Arrival Queue with Different Information of Batch Size

Jinting Wang, Central University of Finance and Economics, Beijing, China, People's Republic of

This talk focuese on the strategic behavior of customers in a single-server batch arrival queue, where the batch size is a random variable. Each customer makes a decision to either join or balk under a linear reward-cost structure. The utility of each customer is contingent on her individual decision, and the decisions of her companions. The research considers two scenarios distinguished by the level of information: the batch size

observable case, where customers are aware of their batch size information, and the batch size unobservable case, where customers lack information upon arrival. In both cases, a unique Nash equilibrium joining strategy and a socially optimal joining strategy are derived. To bridge the gap between individual equilibrium joining strategy and the socially optimal strategy, a proposed fee imposed on customers is introduced. Numerical experiments are conducted to present the impact of batch size information and distribution on customers' joining strategies, admission fees, throughput and social welfare. The findings of this paper offer managerial insights into customer strategic behavior and optimal information-disclosure policies for service providers.

2 - Scheduling Policies in Stochastic Models with Abandonment

Gang Chen, Guangzhou University, Guangzhou, China, People's Republic of

We consider a stochastic scheduling problem where a job is characterized by a general service time distribution, an exponentially distributed lifetime, an abandonment cost and a waiting cost. A job abandons the system if its waiting time in the queue is larger than its lifetime. The objective is to minimize the total expected cost. We show that waiting costs and abandonment costs can be treated in a unified manner. We also show that the minimization problem is equivalent to a maximization problem where a reward is earned each time a job is served. For these equivalent problems, we provide a set of sufficient conditions under which an index policy is optimal. For the special case with two types of jobs with identically distributed service times, we establish that two simple conditions are necessary and sufficient for the optimality of an index policy. Finally, in a numerical study, we show that the necessary and sufficient conditions are satisfied by 83% of instances for three service time distributions (exponential, uniform and deterministic). For other instances under consideration, index policies perform very well with a maximum error of 3.5% and an average error of 0.21%

3 - Queueing-Based Optimization for Electric Vehicle Charging Stations: Dynamic Resource Allocation and Behavioral Strategies

Zhiyuan Zhang, Yanshan University, Qinhuangdao, China, People's Republic of, Wei Sun, Shiyong Li

The widespread adoption of electric vehicles (EVs) has introduced new challenges to charging station operations, including capacity constraints, unreliable chargers, customer impatience, and customer interjections. To address these issues, this study integrates queueing theory, dynamic resource allocation, and behavioral economics to develop an optimized operational framework for EV charging stations. First, from the perspective of station management, a multi-server finite-capacity queueing model is formulated to evaluate system performance while accounting for server reliability and customer impatience. Service level adjustments are explored to minimize overall operational costs. Second, from the infrastructure perspective, a two-stage tandem queueing model is proposed to characterize the transition from constant-current to constant-voltage charging. A dynamic server allocation strategy optimizes resource distribution based on real-time state-of-charge variations. Additionally, customer decision-making under different information scenarios is analyzed to refine operational strategies. Third, from the user perspective, a queueing system incorporating customer interjections and loss aversion theory is developed. A dynamic pricing strategy is designed to mitigate disruptive behaviors and enhance overall revenue, balancing service efficiency with user satisfaction. By leveraging numerical optimization techniques such as genetic algorithms and reinforcement learning, the proposed framework improves operational efficiency, minimizes costs, and enhances the customer experience. The findings provide both theoretical insights and practical strategies for optimizing EV charging station management through an integrated queueing-based approach.

4 - Strategic Green Innovation Licensing and Carbon Reduction Policy in Shipping

Shan Su, University of Macau, Macau, China, People's Republic of, Xiaozhi Feng, Ying Shi, Jin Li, Zhaotong Lian

Sustainable low-carbon shipping has become a key driver of the shipping industry's transformation. As global concerns over carbon emissions intensify, different regulators have implemented a variety of carbon tax policies, each shaped by distinct environmental goals and economic considerations. In response, shipping carriers have increasingly invested in research and development or purchased advanced low-carbon technologies in order to meet new compliance standards and maintain market competitiveness. Against this backdrop, this study first analyses, from the perspective of shipping carriers holding carbon-reduction technology patents, how to formulate optimal technology licensing strategies under a variety of regulatory environments to enhance

operational resilience and maintain a competitive edge. Furthermore, the study explores, from the perspective of regulators, how to promote the widespread adoption of carbon-reduction technologies through reasonable carbon tax policies, without hindering market prosperity.

WB02

Mercury

Humanitarian Operations: Decision-Making, Disaster Response, and Social Impact Optimization

Contributed Session

Humanitarian Applications

Chair: Yujia Wang, The Hong Kong Polytechnic University, Hong Kong, N/A

1 - Dynamic Evacuation Simulation in Urban District Ginza with TIME-Expanded Network and Bpr-Based Rerouting

Yafei Wang, Kansai University, Osaka, Japan, Junpei Tokunaga, Hiroyuki Ebara, Naonori Ueda

Urban disasters such as tsunamis demand evacuation strategies that can adapt not only to spatial constraints but also to dynamically changing crowd conditions. In densely populated areas like central Tokyo, the need for responsive and scalable evacuation planning is particularly acute. This study presents a simulation framework that enhances pedestrian evacuation efficiency by integrating time-expanded networks with congestion-sensitive routing.

Focusing on Ginza district in Tokyo, the model uses OpenStreetMap data to simulate the movement of 50,000 individuals. The urban network is unfolded over discrete time steps to create a time-expanded network, enabling detailed tracking of individual trajectories across both space and simulation time. Unlike static routing models, this structure allows for adaptive re-routing in response to evolving congestion. Pedestrians select paths using a shortest-path algorithm guided by a dynamic cost function. Travel times are updated iteratively using the Bureau of Public Roads (BPR) function, which reflects real-time congestion effects based on pedestrian density and road capacity at each time step.

Simulation results show that the proposed method significantly shortens total evacuation time and promotes a more balanced use of urban pathways, even under severe congestion. By incorporating temporal dynamics and feedback mechanisms, the framework offers a practical and transferable tool for evacuation planning. It is well-suited for application in other metropolitan areas and can be adapted to a wide range of disaster scenarios, contributing to the resilience of urban infrastructure and public safety systems.

2 - Or/MS for Social Good

Srinivasa Gopal, Ramanujan Society for Academic Research and Promotion of Science, Chennai, India

OR/MS for Social Good

Operations Research and Management Science (OR/MS) can be a powerful tool for driving positive social change. This topic focuses on applications of OR/MS in non-profit, humanitarian, and social impact contexts.

Subtopics:

1. *Humanitarian Logistics*: Optimizing disaster response and relief efforts using OR/MS techniques.

2. *Public Health Interventions*: Applying OR/MS to disease prevention, detection, and treatment programs.

3. *Sustainable Development*: Using OR/MS to address environmental and socio-economic challenges in developing communities.

4. *Food Security and Nutrition*: Optimizing food distribution, production, and consumption to address hunger and malnutrition.

5. *Refugee Resettlement and Immigration:* Applying OR/MS to optimize refugee placement, resettlement, and integration processes.

Real-World Examples:

1. *The Red Cross's Blood Distribution System*: Optimizing blood distribution to hospitals using OR/MS techniques.

2. *The WHO's Vaccine Distribution Program*: Using OR/MS to optimize vaccine distribution in developing countries.

3. *The UN's Humanitarian Aid Distribution*: Applying OR/MS to optimize aid distribution in response to natural disasters.

Benefits of OR/MS in Social Good:

1. *Improved Efficiency:* Optimizing resources and processes to maximize impact.

2. *Data-Driven Decision Making:* Using data analytics to inform policy and program decisions.

3. *Scalability and Sustainability:* Developing solutions that can be scaled up and sustained over time.

By applying OR/MS principles and techniques to social good challenges, we can create more effective, efficient, and sustainable solutions that drive positive change.

3 - Dynamic Vehicle Location and Relocation Under Supply and Demand Uncertainties in Disaster Response

Yuying LONG, The Hong Kong Polytechnic University, Hong Kong, China, People's Republic of, Zhengxiong ZHU, Gangyan Xu

The unpredictable nature of disasters creates uncertainty in rescue demands, while their large-scale impact necessitates the use of volunteer vehicles, leading to an uncertain supply of emergency vehicle resources. This imbalance between available emergency vehicles and actual rescue needs can result in delayed rescue efforts, unmet demands, and wasted vehicle resources in regions with surplus vehicles. However, current research predominantly focuses on efficient deployment and dispatch of pre-positioned vehicles, overlooking the spatiotemporal imbalance between vehicle supply and uncertain demands, as well as the effective utilization of volunteer vehicles. To tackle these issues, this study proposes an adaptive approach for dispatching and relocating heterogeneous vehicles—including both pre-positioned and volunteer units—while strategically locating temporary facilities to enhance response efficiency. The problem is formulated as an extension of the Location Routing Problem (LRP), incorporating multiple visits and open routes. A rolling horizon-based dynamic solution framework is developed to update plans as new information becomes available and manage uncertainties. Experimental results reveal that: (1) In small-scale scenarios, the proposed strategy underperforms compared to traditional methods due to increased scheduling complexity and resource dispersion. Conversely, in larger-scale scenarios, it enhances response efficiency by reducing vehicle idle distances through temporary facility locations. (2) Vehicle relocation between permanent facilities proves ineffective in disaster response. When vehicles are scarce, they are fully utilized to meet demands, and when abundant, relocations become redundant. In contrast, the introduction of temporary facilities helps bridge the spatiotemporal imbalance between vehicle supply and demand, ultimately enhancing overall response efficiency.

4 - Cost-Sharing Mechanism for Collaborative Prepositioning of Emergency Supplies

Yujia Wang, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, Jian Wang, Gangyan Xu, Xiaowei Hu, Shi An

We study the collaborative partnership involving the joint prepositioning of emergency supplies in regions frequently struck by natural disasters, such as typhoons and earthquakes. These regions, consisting of multiple

cities, face varying impacts across disasters, creating opportunities for significant benefits through risk pooling enabled by centralized horizontal collaboration.

Assuming the presence of an intergovernmental agency tasked with reasonably allocating the costs necessary to operationalize the collaborative partnership, we present a spontaneous cost-sharing mechanism based on insurance theory. A cost-sharing mechanism is considered spontaneous if it satisfies the individual rationality of all participants. This mechanism, which is formulated as a linear programming model, allocates costs to participants by considering both the expected value and the standard deviation of their respective risks. To determine the costs associated with collaborative prepositioning of emergency supplies, we develop a two-stage stochastic programming model to identify the optimal locations and quantities of emergency supplies to preposition.

Using real-world data from the Great Bay Area, we develop a realistic dataset to test our model. Our results demonstrate that the proposed cost-sharing mechanism achieves superior outcomes compared to alternative approaches, particularly in terms of individual rationality for participating cities. These findings highlight the practicality and effectiveness of the spontaneous cost-sharing mechanism in enhancing collaborative disaster preparedness.

WB03

Minto

Robust Optimization

Invited Session

Computing Society

Chair: Qinshen Tang, Nanyang Technological University, Singapore, Singapore

1 - Strategic Optimal Locations of Self-Collection Lockers for Medication: a Feature-Driven Robust Optimization Approach

LI NINGXIN, Southwestern University of Finance and Economics, chengdu, China, People's Republic of

Healthcare systems worldwide are challenged by aging populations and complex patient needs. In this context, home delivery of medication (HDM) is gaining momentum. Self-collection locker, a proven e-commerce solution, offers receivers flexibility and reduces senders' delivery trips. However, the application of a self-collection locker for medication is scant. In this paper, motivated by our partnered hospital's practical need, we study a location problem of setting up a certain number of capacitated self-collection lockers for HDM, which minimizes the costs for patients obtaining medications. Given medications' storage temperature requirements, the heterogeneity of lockers should be considered in location decisions. In reality, the further demands for medication are uncertain. Patient features such as age, clinic description, and residential location are crucial to improving the model fidelity of uncertain medication demands. Using feature data, we re-weight the empirical distribution using machine learning. To hedge against the distribution's ambiguity with limited data, we propose a feature-driven distributionally robust optimization model. That said, we err on caution and minimize the cost under the worst-case distribution. To efficiently solve the problem, we develop a column-and-constraints generation algorithm. The numerical experiment, conducted using real data, identified the predictive value of feature information and demonstrated the effectiveness of the feature-based ambiguity set. Furthermore, it suggested possible directions for improvement.

2 - Robust Network Design and Long-Term Planning for Carbon Capture and Storage

Bei Lin, Renmin University of China, Beijing, China, People's Republic of, Xiaoyu Ji, Yingtong Wang, Yingfu He

In the context of carbon neutrality, Carbon Capture, Utilization, and Storage (CCUS) has emerged as an effective carbon reduction method. This study addresses the challenge of large-scale clustered CCUS network layout and long-term planning problems considering uncertainties in carbon emissions. To achieve carbon reduction targets under such uncertainty, an adaptive robust optimization framework is proposed to strategically locate hubs,

establish pipelines, select sinks, and adjust transportation and storage schemes. The first stage involves determining the suitable pipeline installation and sink location, and the second stage focuses on developing an effective long-term transportation allocation and storage plan. This study presents an enhanced column-and-constraint generation algorithm (ECCG), which combines the column-and-constraint generation algorithm and a modified outer approximation. The ECCG algorithm effectively deals with the uncertain variables with budgeted uncertainty sets. Furthermore, the validity of our model is evaluated through numerical experiments with real data from energy companies. This framework facilitates CCUS projects in making strategic decisions under diverse policies and market contexts, and provides policymakers with targeted recommendations for adjusting subsidies to various sinks. These insights incentivize more companies to deploy CCUS projects and provide decision support for both companies and the government to promote carbon neutrality.

3 - Max-Min-Max Framework and Its First-Order Algorithm for Large-Scale Robust Optimization

Kai Tu, Shenzhen University, Shenzhen, China, People's Republic of, Zhi Chen , Man-Chung Yue

Robust optimization (RO) is a powerful paradigm for decision making under uncertainty. Existing algorithms for solving RO, including the reformulation approach and the cutting-plane method, do not scale well, hindering the application of RO to large-scale decision problems. In this work, we devise a first-order algorithm for solving RO based on a novel max-min-max perspective. Our algorithm operates directly on the model functions and sets through the subgradient and projection oracles, which enables the exploitation of problem structures and is especially suitable for large-scale RO. Theoretically, we prove its global convergence theorem under some suitable conditions. Furthermore, we prove that the oracle complexity of our algorithm for attaining an ε-approximate optimal solution is $O(\varepsilon^{-3})$ or $O(\varepsilon^{-2})$, depending on the smoothness of the model functions. The algorithm and its theoretical results are then extended to RO with projection-unfriendly uncertainty sets. We also show via extensive numerical experiments that the proposed algorithm outperforms the reformulation approach, the cutting-plane method and two other recent first-order algorithms.

4 - Robust Contextual Portfolio Optimization with Gaussian Mixture Models

Ling Dai, City University of Hong Kong, Hong Kong, China, People's Republic of, Yijie Wang, Grani Adiwena Hanasusanto, Chin Pang Ho

We address the portfolio optimization problem with contextual information that is available to better quantify and predict the uncertain returns of assets. Motivated by the distinct regimes for the finance market, we consider the setting where the uncertain returns and the contextual information jointly follow a Gaussian Mixture (GM) distribution. We establish that this problem is equivalent to a nominal portfolio optimization problem where the mean and the covariance matrix are adjusted by the contextual information. To reduce the sensitivity of the model performance with respect to the inherent model parameters within the Gaussian Mixture Model (GMM), we propose the robust contextual portfolio optimization problem. By considering a projection of the ambiguity set, a tractable formulation is derived to approximate the exact model. We conduct numerical experiments in both US markets and the global Exchange-Traded Funds market, and the results demonstrate the advantage of our proposed model against other benchmark methods.

WB04

Indiana

Charging Infrastructure and Pricing Strategy in EV Transition

Invited Session

Transportation Electrification

Chair: Lingling Shi, McMaster University, Hamilton, ON, Canada

1 - Robust Charging Infrastructure Planning for Passenger Vehicles Electrification: Considering Scheduling and Abandonment

Thomas Magnanti, MIT, Cambridge, MA, United States, Yifu Ding, Ruicheng Ao

Transportation electrification is a key strategy for reducing the carbon footprint of heavy industrial sectors like the mining sector. Deploying battery-electric vehicles (BEVs) and their associated charging infrastructure requires a thorough understanding of fuel economy and charger design. Additionally, charging infrastructure planning must consider the operational behaviors of passenger vehicles, such as driving scheduling and

abandonment of EV charging. This study presents a robust, data-driven framework for optimizing charging infrastructure for electric passenger vehicles, incorporating charging allocation and abandonment due to wait times. Case studies were conducted using empirical vehicle data from a mining site, analyzing various scenarios, including the use of backup trucks and different waiting times.

2 - Pricing and Producing Green Products Under Subsidy Expiration, Learning-by-Doing and Escalating Competition

Lingling Shi, McMaster University, Hamilton, ON, Canada, Metin Cakanyildirim, Suresh Sethi

Government subsidies are critical for attaining green product adoption targets. By increasing sales, successful subsidies can reduce production costs through manufacturers' learning-by-doing and escalate their competition. Such a subsidy becomes a victim of its own success and eventually expires. We develop a two-period Stackelberg-Nash game between a government and two manufacturers, considering subsidy expiration, learning-by-doing, and escalating competition. We find that each manufacturer charges a lower (higher) price in the first period than in the second if the subsidy is less (more) than a threshold, and we relate this to the penetration (skimming) pricing strategy. Each adopts a zero-inventory strategy. The subsidy acts as a substitute for competition. Moreover, the effective price (with subsidy) rises over time, eliminating the waiting incentive of the first-period customers. We provide several insights and suggestions to a policymaker. The government incurs less subsidy expenditure under the qualified-sales rule than the all-sales rule. We identify cases in which the government should allow a merger of the manufacturers to reduce subsidy expenditure or foster group learning to increase manufacturer profits. We also find that green product subsidy pulls (effective) market prices down, benefiting consumers of not only green products but also non-green products.

3 - CCP Estimation of Dynamic Discrete Choice Demand Models with Segment Level Data and Continuous Unobserved Heterogeneity: Rethinking EV Subsidies vs. Infrastructure

Tim Derdenger, Carnegie Mellon University, Pittsburgh, PA, United States

4 - Charging Station Location Optimization Based on Queuing Theory Under V2g Mode

Yanzi Zhang, Beijing Institute of Technology, Beijing, China, People's Republic of

As the scale of electric vehicles continues to expand, the demand for charging infrastructure is also increasing. Existing charging facilities face challenges such as supply-demand mismatches and long user waiting times. This paper studies the charging station location optimization problem under the Vehicle-to-Grid (V2G) mode. Based on queuing theory and optimization theory, this problem is formulated as a mixed-integer nonlinear programming problem, leveraging its nonlinear characteristics and solving it using the commercial solver Gurobi. To enhance solution efficiency, a Transformer-based solution approach is proposed and validated through numerical experiments, demonstrating its superiority in handling large-scale charging station location problems. Additionally, this study deeply analyzes the impact of key factors such as distance weight, the number of charging station levels, and charging service efficiency on location decisions. The aim is to improve the profitability of facilities and the user charging experience, achieving maximized social benefits

WB05

Orchard A

Emerging Customer Behavior and Revenue Management

Invited Session

Efficiency and Fairness in Supply Chain and Service Operations Management

Chair: Yufeng Cao, Shanghai Jiao Tong University, Shanghai, N/A

1 - Product Line Design Under a Contextual Choice Model: An Application to Airline Fleet Assignment and Timetabling

Xiaoyun Niu, Tongji University, Shanghai, China, People's Republic of, Yufeng Cao, Keji Wei

On a busy afternoon at San Francisco International Airport, four nearly identical flights to Los Angeles depart within a 30-minute window. Same price, same duration—yet passengers still face a choice. What drives their decision? Traditional models, like the Multinomial Logit (MNL), assume each option has a fixed appeal, independent of the others. But in reality, when options are highly similar, context shapes decisions.

Motivated by this observation, we introduce the Context-Dependent Multinomial Logit (CMNL) model, extending MNL to capture how competing options influence perceived attractiveness. CMNL achieves this by incorporating a decay function, which adjusts a product's appeal based on the presence of close substitutes.

We apply CMNL to the Product Line Design (PLD) problem, where businesses must strategically select offerings to maximize revenue while accounting for competition. Our approach formulates PLD as a mixed-integer convex program via linearization techniques. We prove that solving PLD with a specific type of context effects is NP-hard but develop a scalable solution framework.

Empirical validation on a publicly available hotel dataset demonstrates that CMNL significantly improves predictive accuracy, offering a more realistic model of consumer decision-making. These insights have broad implications for revenue management, product assortment, and industries where context-driven choices play a critical role.

2 - Optimization of Courier-Order Matching and Delivery Performance for in-TIME Delivery Systems

Yi Fan, Tongji University, Shanghai, China, People's Republic of

Meal delivery services are experiencing surging demand, making it vital for delivery platforms to enhance both their efficiency and promptness. We address this challenge by introducing the Adaptive Delivery Matching with Agent-based Courier Reinforcement Optimization (AD-MACRO) framework, which formulates the match-delivery decision problem for such platforms. We utilize a two-layer decision framework to effectively manage multi-agent scenarios and develop distributed reinforcement learning and optimization-based methods to tackle the decision-making challenges. To manage the complexity presented by a large cohort of agents and capture the nuances of platform-courier interactions, we propose the Grouped-Actor Policy Proximal Optimization (GAPPO) algorithm. This algorithm reduces the decision-making dimension by grouping similar agents, thereby accommodating the inherent heterogeneity between platform-employed and crowdsourced couriers. We validate our approach using current dispatch data from Meituan, demonstrating its comprehensive evaluation of delivery performance. Furthermore, to increase the practical relevance of our work, we incorporate couriers' rejection behaviors using fixed probabilities and economic incentives. By explicitly integrating these behavioral aspects, along with considerations of transportation equity, our framework can achieve optimized order assignments and delivery performance from both efficiency and equity, e.g. improved delay rate (-77.61%) and increased leisure time (+50.52%). This results in a rare win-win scenario for couriers, platforms, and customers alike.

3 - Service Allocation for Heterogeneous Delay-Sensitive Customers

Yuting Liu, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Yufeng Cao

We consider an on-demand service platform that matches randomly arrived service providers and heterogeneous delay-sensitive customers. The platform offers two types of services, a regular service and a Wait-and-Save (W&S) service, to serve customers with differentiated delay costs. Under the regular service, a customer pays a higher price and receives an immediate service if a service provider is available. Under the W\&S option, a customer pays a lower price and receives service within a guaranteed waiting period. If a W&S customer remains unserviced by the end of the guaranteed period, the platform incurs a penalty cost, e.g., due to the loss of goodwill. Such a promise in waiting time distinguishes the W&S service from traditional priority-based services. Customers make informed choices based on their delay sensitivity, and service providers remain on the platform until being assigned to a service request. We model the platform's service allocation problem as a long-run average reward Markov decision process (MDP) over an infinite planning horizon. We characterize properties of the optimal dynamic allocation policy. However, due to the curse of dimensionality, it is intractable to solve the problem for practically-sized real-world instances. Therefore, we propose a threshold-type heuristic policy based on the problem structure. We assess the performance of the heuristic policy via numerical experiments. Analytical results such as asymptotic optimality and characterization of the order of loss ratio are expected.

4 - Revenue Management Under a Markovian Search and Choice Model

Junjia Zheng, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Yufeng Cao, Wenhui Zhao

We introduce the Markovian search and choice model (MSCM), a framework for capturing customer searching, waiting and purchasing decisions in e-commerce settings. Leveraging historical search path and transaction data,

the MSCM accurately infers customer preferences. We study the two-product case analytically and show the mechanisms that make the difference between multi-stage and single-stage choice model. We further study a pricing problem under the Markovian search and choice model and prove the monotonicity of the optimal pricing policy. To address computational complexity, we propose a deterministic approximation heuristic that incorporates the demand dependency arising from waiting customers in previous periods. We further show the revenue gap between multi-stage and single-stage choice model in model misspecification setting. Numerical studies based on synthetic and real data suggest that our approach outperforms the MNL model in terms of out-of-sample predictive accuracy.

WB06

Orchard B

Emerging Challenges and Innovations in Transportation Electrification and Automation

Invited Session

Transportation Electrification

Chair: Sen Li, The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong

1 - Integrating Valet Charging with Mobile Discharging Services: a Strengthened Integer L-Shaped Method

Sen Li, The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong

Valet charging is an emerging business model where a platform recruits couriers to assist electric vehicle (EV) owners in recharging their vehicles. While the service provides a convenient charging solution, particularly for those without residential chargers, its adoption has been limited due to inadequate affordability. To address this, we propose integrating valet charging with mobile discharging services, allowing customers to receive discounted prices by permitting their EVs to participate in discharging tasks. The success of this integration hinges on incentivizing customers to extend their vehicle return deadlines, despite uncertainties regarding their willingness to delay. We model this incentive design problem as a two-stage stochastic program. In the first stage, the platform designs a deadline-differentiated price menu, offering multiple return time options, each with an associated price. In the second stage, customers choose the deadline that maximizes their utility, after which the platform assigns EVs to discharging jobs and schedules charging operations. To tackle the problem, we first devise an exact linear reformulation that preserves decomposability without any loss of optimality. We then adapt the integer L-shaped method to this reformulation, constructing strengthened optimality cuts that encode the intrinsic interdependence between first- and second-stage decisions. These cuts are proved to dominate those used in the standard integer L-shaped method. Experimental results indicate that our proposed method significantly outperforms commercial solvers and the integer L-shaped method in its standard form, achieving enhanced efficiency and scalability. Furthermore, the findings underscore the mutual benefits of the integrated service for the customers and the platform.

2 - Real-time Capacity Control in Railway Systems with Ticket-to-Seat Assignment and Cancellations

Linhuan Zhong, National University of Singapore, Singapore, Singapore, Xiaolei Zhu, Yang Liu, Guangming Xu

Real-time capacity control is vital to railway revenue management but is complicated by uncertain passenger demand and frequent cancellations. This study proposes a learning-enhanced hierarchical optimization framework for multi-train, multi-stop systems, explicitly incorporating ticket-to-seat assignments and cancellations to maximize ticket revenue. The upper-level model proactively sets order acceptance ratios for different origin-destination pairs, while the lower-level model assigns specific requests to seats using a mixed-integer linear program (MILP).

To model system dynamics and reduce complexity, the timetable is divided into non-overlapping grids, forming a grid-based multi-agent Markov decision process (MDP). The upper-level decisions are optimized using a multi-agent deep deterministic policy gradient (MADDPG) algorithm enhanced by mean-field approximation. The lower-level MILP incorporates upper-level decisions as constraints to ensure real-time responsiveness, and a tailored decomposition algorithm solves it efficiently. Feedback from the lower-level model is used to improve

the upper-level policy.
Experiments on the Nanjing–Hangzhou high-speed railway demonstrate that this approach outperforms benchmark strategies, including greedy and static capacity controls. Notably, performance gains increase with greater uncertainty, highlighting the adaptability and effectiveness of the proposed method.

3 - A Learning-Informed Optimization Framework for Dynamic Balancing-Charging Management of Shared Autonomous Electric Vehicle Systems

Xiaolei Zhu, National University of Singapore, Singapore, Singapore, Xindi Tang, Jiaohong Xie, Yang Liu

Balancing supply and demand in large-scale Shared Autonomous Electric Vehicle Services (SAEVSs) remains a critical yet challenging task for real-time operations. In this study, we examine the dynamic balancing-charging (BC) management problem with a focus on the joint optimization of both fleet management and charging decisions under stochastic demands and operational constraints. This necessitates advanced modeling and foresight analytics to capture system dynamics and the long-term impact of decisions, enabling real-time, informed decision-making. To this end, we propose a novel learning-informed optimization framework that augments the precision strength of an optimization approach with the foresight capability of Deep Reinforcement Learning (DRL). Specifically, we develop a multi-agent DRL model, functioning as a "Manager", which learns a dynamic BC strategy at the grid level amidst system uncertainties. The information about the underlying demand distribution and intricate interactions between current decisions and future system dynamics is preserved in the DRL Manager model to inform the optimization of operational decisions via BC strategy. Supporting the vehicle-level decisions, we propose a customized space-time-battery network flow model, referred to as a "Worker". This model follows the far-sighted BC strategy developed by the Manager model to optimize real-time vehicle assignments. Nevertheless, learning the optimal BC strategy is highly non-trivial for combinatorial decision-making in large-scale SAEVSs. To tackle the learning challenges, we establish a synergistic DRL-based algorithm to solve the proposed framework. This study also lays a methodological basis for further exploring the integration of DRL and optimization techniques, aiming to enhance decision-making capabilities in urban mobility systems.

4 - Safe and Verifiable Control of Mixed-Autonomy Platoons with Control Barrier Functions

Kaidi Yang, National University of Singapore, Singapore, Singapore, Jingyuan Zhou

Mixed-autonomy platoon control aims to enhance the efficiency of vehicle platoons comprising connected and automated vehicles (CAVs) and human-driven vehicles (HDVs). Among existing control strategies, Multi-Agent Reinforcement Learning (MARL) appears to be promising because it can manage complex scenarios in real time. However, existing MARL approaches address safety by penalizing safety violations in the reward function, thus lacking theoretical safety guarantees due to the black-box nature of RL. Moreover, few studies have explored the cooperative safety of multi-CAV platoons, where CAVs can be coordinated to further enhance the system-level safety involving the safety of both CAVs and HDVs. To bridge the research gaps, we propose a safe MARL framework for mixed-autonomy platoons. Specifically, this framework (i) characterizes cooperative safety by designing a cooperative Control Barrier Function (CBF), enabling CAVs to collaboratively improve the safety of the entire platoon, (ii) provides a safety guarantee to the MARL-based controller by integrating the CBF-based safety constraints into MARL through a differentiable quadratic programming (QP) layer, and (iii) extends the CBF to a neural CBF to make the safety guarantees less conservative and applicable to scenarios with stochastic and complex vehicle dynamics. Simulation results show that our proposed control strategy can effectively enhance the system-level safety through CAV cooperation of a mixed-autonomy platoon with a minimal impact on control performance.

5 -

An LLM Approach to Traffic Assignment and Its Application to Electric Vehicle Behavior Modeling

Jiayang Li, The University of Hong Kong, Hong Kong, China, People's Republic of

In transportation, a fundamental task is predicting travelers’ route choice in a general congestible network, which is known as traffic assignment. A standard approach to traffic assignment is based on certain equilibrium concepts, such as user equilibrium or stochastic user equilibrium. Recently, generative artificial intelligence (AI), particularly large language models (LLMs), has emerged as a powerful tool for simulating human decision-making processes. In this talk, I will present an integrative approach to traffic assignment that combines the power of LLMs with classic reinforcement learning models from behavioral economics. Additionally, I will

discuss the application of this approach to scenarios involving electric vehicles (EVs) to highlight its potential to address emerging challenges in sustainable transportation.

WB07

Orchard C

AI for Epidemic Control

Invited Session

Health Systems Modelling with AI

Chair: Ling Yin, Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, Shenzhen, N/A

1 - Optimizing Spatiotemporal Control for Urban Epidemics with AI Models

Ling Yin, Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, Shenzhen, China, People's Republic of, Ting Hu, Yuxiao Luo, Xuan Li, Shang Wang

In urban areas, high population density and extensive mobility facilitate the rapid spread of emerging infectious diseases, particularly acute respiratory infections, posing significant public health challenges and widespread societal impacts. Determining when, where, by whom, and which disease control measures should be implemented to achieve optimal control at minimal cost constitutes a multi-objective spatiotemporal optimal control problem requiring real-time interactive decision-making. Since real-world experimentation is infeasible, we rely on epidemic spread simulators to model this optimization problem. Real-world-inspired simulation models for guiding disease control are typically highly nonlinear complex systems, making conventional mathematical approaches inadequate for solving such optimal control problems. Artificial intelligence (AI)-based methods, particularly reinforcement learning (RL), have emerged as a promising solution for complex optimal control tasks in recent years. This study tackles urban epidemic optimal control by developing a series of RL-based spatiotemporal intervention strategies. For example, we introduce the concept of "spatiotemporal orderliness of control measures" to minimize societal disruption while maintaining effectiveness. We also develop approaches to address the challenges of generating personalized intervention planning for large-scale urban populations and robust decision-making under partial observation due to incomplete disease surveillance. By integrating these innovations, we provide a suite of AI-driven tools to enhance epidemic control strategies.

2 - A Data Driven Clinical Algorithm for Differential Diagnosis of Respiratory Viruses Using Machine Learning and Natural Language Processing

Mingqing Xie, the School of Public Health, Fudan University, Shanghai, China, People's Republic of, Ye Yao

Background
The accurate and timely diagnosis of respiratory viruses remain a major clinical challenge due to high diagnostic costs, and evolving post-pandemic epidemiological trends. This study aimed to develop an efficient and cost-effective diagnostic approach using natural language processing and machine learning.

Methods
We collected and analyzed data from Huzhou City for four common respiratory viruses: SARS-CoV-2, influenza, respiratory syncytial virus (RSV), and adenovirus (ADV). Natural language processing techniques were employed to extract structured symptom features from unstandardized diagnostic statements which were utilized to train machine learning classifiers.

Results
XGBoost and LightGBM were selected as the final models given that their superior performance during hyperparameter tuning. The models achieved strong discriminative performance, especially for SARS-CoV-2 and RSV(AUC:SARS-CoV-2 0.855, 95% CI 0.829-0.881; influenza 0.736, 95% CI 0.712–0.759; RSV 0.801, 95% CI 0.744–0.857; ADV 0.781, 95% CI 0.747–0.815). Subgroup analyses showed consistent model performance across sexes, while particularly high diagnostic accuracy was observed in pediatric or afebrile patients. The AUC of the SARS-CoV-2 predictive model reached 0.881 (95% CI: 0.843–0.919) among patients aged 0–18, and 0.866 (95% CI: 0.803–0.929) in the afebrile group, both exceeding the overall

performance(0.855, 95% CI 0.829-0.881). The SHAP analysis indicated that for all four pathogens, age was the most significant predictor.

Conclusion

This study demonstrates the feasibility of integrating NLP and machine learning for the accurate identification of respiratory viruses based on clinical symptoms alone, and offers an affordable alternative to PCR testing, which can reduce reliance on resource-intensive testing and enhance early detection.

3 - Lstcnet: Real-TIME Influenza Forecasting in China Using a Hybrid LSTM-Tcn Neural Network with Baidu Search Index Data

Wei He, Macau University of Science and Technology, Macau, China, People's Republic of

Influenza represents a major global public health challenge, leading annually to significant morbidity and mortality. In China, the current influenza surveillance system experiences reporting delays of one to two weeks, impeding timely public health interventions. To mitigate this issue, we propose LSTCNet, a novel hybrid deep-learning framework combining Long Short-Term Memory (LSTM), Temporal Convolutional Networks (TCN), and spatial-temporal attention mechanisms to provide real-time predictions of influenza positivity rates. This model integrates Baidu Index data, filtered through Shapley Additive Explanations (SHAP) for feature selection, to capture regional-specific online search behaviors, thereby enhancing predictive accuracy for both Northern and Southern China.

The proposed LSTCNet architecture employs bidirectional LSTMs to model long-term sequential dependencies and leverages multi-scale TCN layers to detect short-term local temporal fluctuations. Dual attention modules further refine the extracted features by adaptively emphasizing critical temporal channels and spatial dimensions. To rigorously evaluate its performance, we benchmarked LSTCNet against five state-of-the-art forecasting models (LSTM, GRU, CNN-LSTM, TCN, and Transformer) using influenza surveillance and Baidu Index datasets from January 2011 to April 2024. Results indicate superior predictive performance by LSTCNet, achieving peak R^2 scores of 0.8487 for the Southern region, significantly outperforming baseline models.

Our findings demonstrate the efficacy of integrating advanced neural network architectures with high-dimensional search-engine data to enable timely and accurate influenza forecasting, facilitating proactive public health decision-making, targeted vaccination efforts, and improved seasonal epidemic preparedness.

4 - Assessing Large Language Models in Disaster Behavior Science: Predictive Validity for Risk Perception and Preventive Intent During Warnings

Yinghui Li, Tianjin University, Tianjin, China, People's Republic of, Huibin Du, Hongyang Zou, Ling Yin

In recent years, climate change has made typhoon disasters more severe and unpredictable. Public risk perception and preparedness behavior regarding typhoon warnings are crucial for reducing disaster losses. However, traditional emergency management lacks effective methods to predict people's disaster risk perception and preparedness intentions, hindering early prevention and real-time decision-making. The rapidly developing large language model (LLM) offers a new approach for disaster behavior prediction. It can complete traditional survey experiments in a short time. Currently, the social science application of LLMs mainly focuses on verifying classic studies in psychology and economics, while exploration in disaster behavior science is still relatively limited. This study aims to explore the potential of LLMs in simulating human decision-making in natural disaster risk management, using typhoon disasters as the experimental context. We designed two warning levels (high and low) and two subjects (LLM and human) experimental scenarios, collecting data on risk perception and preparedness behavior intentions reported by participants in different scenarios. We compared results to examine whether factors affecting human preparedness behavior, such as disaster warning level, influence LLM decision-making. The results indicate that in the typhoon disaster warning scenario, at the group level, the LLM can more accurately predict individuals' preventive behavior intentions. A higher warning level also leads to increased risk perception and preparedness intentions in the LLM; however, it tends to overestimate people's risk perception and underestimate their willingness to help others. At the individual level, the LLM demonstrates low accuracy in predicting each person's risk perception and preparedness behavior.

5 - The AI-Driven Organizational Adaptation in Pandemic-Emergency Transition: a Parallel Simulation System for Collaborative Health Crisis Governance

Ke Qiao, Tianjin University, Tianjin, China, People's Republic of, Hongyang Zou, Huibin Du

The structural tension between institutionalized health governance systems and evolving biological crises poses fundamental challenges to global epidemic preparedness. As global health systems confront evolving biological threats, artificial intelligence (AI) offers transformative potential for enhancing collaborative decision-making in crisis scenarios. This study develops an AI-powered multi-agent simulation framework, integrating epidemiological datasets from 2015-2022 outbreaks, to systematically analyze coordination mechanisms in public health emergencies.

Through computational experiments combining health organization network analysis with AI-driven agent modeling, three principal findings emerge: First, modularization of health agencies reduces cross-departmental coordination paths by 2.3 steps on average, yet reveals systemic bottlenecks where specialized response teams require 3.7 steps to access central decision-makers. Second, the low betweenness centrality (0.32) of non-governmental health organizations quantifies institutional barriers to collaborative efficiency. Third, AI simulations uncover a U-shaped evolution of governmental centrality ($\rho = -0.67$ with NGO collaboration intensity), reflecting resource allocation tensions during epidemic waves.

The framework demonstrates 84.3% efficiency in crisis resource utilization while maintaining 35-minute response thresholds, providing operational guidelines for building scalable health emergency systems. These findings advance AI applications in pandemic modeling, offering methodological innovations for optimizing multi-stakeholder coordination in disease containment strategies.

WB08

Bras Basah A

Supercharging GenAI with Hyper-Human Interactions: Is AI a Blessing or a Menace?

Invited Session

Informations Systems

Chair: Ben Choi, Nanyang Technological University, Singapore, Singapore

1 - Cognitions and Emotions in Human-AI Collaborations with Augmented Reality: a Dual-Coding Perspective

Zeyuan (Stephen) Hong, Nanyang Technological University, Singapore, Singapore, Ben Choi, Wai Boh

Artificial intelligence (AI) is expected to revolutionize the existing methods of augmented reality processing and natural voice interactions, through which humans collaborate with AI agents to complete specific tasks. In contrast to human-human digital collaborations, which facilitate mutual understanding and synchronized attention, human-AI collaborations focus on optimizing the comprehensibility of AI instructions. Drawing on the instructional design literature, this study considers two augmented reality features facilitating progressive, multiple-round human-AI collaborative diagnosis (i.e., verbalized instructions and visual illustrations) and examines how verbalized instructions and visual illustrations affect diagnosis effectiveness and efficiency. We enrich the instructional design literature by integrating dual-coding theory to elucidate the interaction effect of the two features and reveal the underlying mechanisms through the mediating effects of cognitive effort and confused emotion. We conducted two laboratory experiments involving augmented reality devices and neuroscience techniques. The results show that inductive instructions with schematic illustrations elevated diagnosis efficiency. Meanwhile, inductive instructions with structural illustrations led to significantly higher diagnostic effectiveness. Furthermore, cognitive effort and confused emotions simultaneously mediated the effects of augmented reality on diagnostic performance with humanoid verbalization. Confused emotions were in play with robotic verbalization. Overall, this study contributes to the human-AI collaboration literature and dual-coding theory by unraveling the effects of instructional designs in augmented reality on human-AI collaborative diagnosis.

2 - When Intelligent Agents Sound like My Friend: Understanding Transference Effects in Voice-Based Interactions

Fenfen Zhu, NTU, singapore, Singapore

With the advancement of artificial intelligence and speech synthesis technologies, individuals increasingly interact with voice-based intelligent agents (e.g., Apple Siri, OpenAI ChatGPT, watsonx Assistant, and XiaoAi) for tasks like information retrieval, online shopping, and navigation. Traditionally, these agents use standardized human-like voices, but a growing trend allows users to personalize them by cloning the voices of significant others. While voice customization enhances user engagement and experience, it also raises ethical concerns, particularly regarding risks associated with deepfake. Given these complexities, a deeper understanding of its psychological and behavioral effects is essential. This study aims to fill this gap by examining the effects of audio customization. Specifically, drawing on transference literature, we explore how personalized voice interactions influence users' perceptions of intelligent agents and their subsequent behaviors. To test our hypotheses, we will conduct a lab experiment in a controlled human-agent interaction setting.

3 - I Remember What You Told Me: Investigating the Effects of Chatbot Memory on Self-Disclosure

YAN XING, Zhejiang University, Hangzhou, China, People's Republic of, Xixian Peng, Ben Choi, Qiuzhen Wang

With advancements in generative AI, social chatbots can now simulate memory functions, enhance personalized interactions, and raise privacy concerns. Drawing on the APCO model, this study examined how chatbot disclosure influences user self-disclosure by investigating the impact of information retention on users' willingness to share personal information. Building on social penetration theory, we further explored the moderating roles of chatbot interaction reciprocity and information intimacy in shaping these effects, evaluating both the benefits (perceived rapport) and costs (perceived privacy invasion) associated with this process. To investigate these dynamics, we conducted two experiments: a longitudinal behavioral study to track disclosure patterns over time, and a neuroscientific experiment incorporating psychological and physiological measures to uncover the underlying mechanisms. By integrating insights from privacy research and interpersonal relationship theories, we provide a nuanced understanding of the trade-off that users navigate in chatbot interactions. Our findings offer important implications for designing socially intelligent chatbots, highlighting how information retention can be optimized to balance personalization with user privacy concerns, thus fostering more effective human-chatbot interactions.

4 - Reverse Discrimination in Genderized Artificial Intelligence

Yulia Lunina, Nanyang Technological University, Singapore, Singapore

Traditionally, defensive driving education and assessments have been administered within driving schools, with certain professionals mandating routine safety training. However, the efficacy of traditional driving education and assessment is limited by the availability of human resources. Artificial intelligence (AI) offers promising avenues for overcoming this personnel bottleneck. This short paper outlines a study that contributes to the discourse on stereotypes and biases beyond the foundational aspects of artificial intelligence. Prior studies have extensively investigated and illuminated issues concerning gender and racial biases inherent in AI recommendations and analyses. However, a critical underlying cause of these challenges is that AI is engineered to emulate the human decision-making processes. This study offers a novel perspective by considering potential constructive applications of such inherent biases. Furthermore, by systematically examining the generalizability of non-human agents, this study aims to raise awareness, if not underscore, of the significance of these fundamental human limitations.

5 - When Virtual Characters Surprise: Exploring Conversational Experiences with AI Companions

Yiyao Chen, The Hong Kong Polytechnic University, Hong Kong, Hong Kong, Fei Liu, Yijing Li, Yongqiang Sun, Chee-wee Tan

Virtual characters powered by generative Artificial Intelligence (AI) can provide companionship through immersive conversation and role-playing experiences. This novel form of entertainment captivates users’

imaginations and continues to gain popularity. However, sustaining such experiences remains challenging, as users tend to lose interest when conversations become stale and predictable. To address this issue, the present study seeks to uncover moments during interactions with AI companions that surprise users. As a well-established psychological response triggered by deviations from expectations, surprise can stimulate emotional arousal and renew user interest. Drawing on interpersonal relationship theories, we posit that the approach in which AI companions respond, either with intimacy-oriented (i.e., being friendly) or power-oriented (i.e., being dominant) approaches, is crucial for evoking surprise. Furthermore, we theorize that anthropomorphism, which refers to the degree to which users perceive AI companions as human-like, moderates the impact of these response approaches on user surprise. We empirically examined our hypotheses using a preliminary dataset from a subreddit dedicated to AI characters. Our dataset comprises 542 posts by 481 users sharing screenshots of their conversations with AI companions. Results indicate that both intimacy-oriented and power-oriented responses, as well as perceived anthropomorphism, can evoke user surprise. Interestingly, while higher levels of anthropomorphism weaken the effect of intimacy-oriented responses, they do not diminish the impact of power-oriented responses. These findings contribute to the human-AI interaction literature by shedding light on how surprising moments emerge in casual conversations with AI companions, an emerging research context dinstict from the sevice-orietned interactions with AI agents.

WB09

Bras Basah B

AI-Driven and Computational Optimization

Contributed Session

Optimization

Chair: Jie Deng, N/A

1 - An Optimization Framework for Efficient Variable-Length Anomaly Detection in TIME Series

Hendri Sutrisno, National Taipei University of Technology, Taipei, Taiwan

Detecting anomalous patterns in time series data is a fundamental task in data mining, with applications spanning finance, healthcare, manufacturing, and beyond. However, traditional approaches often assume fixed-length subsequences and struggle to effectively identify anomalies of varying durations and structural complexity. This limitation becomes more prominent in real-world scenarios where anomalies may be subtle, diverse in length, and context-dependent. In this study, we propose a general-purpose, two-stage optimization framework, MP-OPT for detecting variable-length anomalous subsequences in time series. The framework begins with Matrix Profile (MP), which efficiently scans the time series to detect fixed-length subsequences that are least similar to others, serving as initial anomaly candidates. While MP is highly effective for rapid discovery, it is limited by its fixed-length constraint. To overcome the issue, the second stage of our method refines these candidates by optimizing both their starting positions and lengths. This refinement is formulated as a bounded optimization problem and solved using a metaheuristic approach. To evaluate candidate anomalies during optimization, we introduce a localized comparison strategy based on Latin Hypercube Sampling, which selects a representative subset of nearby subsequences from a broader context region. This strategy ensures reliable anomaly scoring while reducing computational costs. Our framework addresses key challenges in anomaly detection, including variable-length analysis, computational scalability, and context-aware evaluation. Experimental results across diverse datasets confirm that MP-OPT effectively identifies complex anomalous patterns with minimal assumptions, offering a robust and flexible tool for time series data mining.

2 - Optimizing Multi-Drone and Vehicle Coordination: a Flexible Launch and Recovery Vehicle Routing Model

Zhang Wu, Chongqing University, Chongqing, China, People's Republic of, Linmin Ai

This paper investigates the vehicle routing problem with drones based on flexible launch and recovery operations (VRPD-FLRO). Unlike traditional approaches, VRPD-FLRO allows vehicles to launch and recover drones at any point along their routes, beyond just the depot and customer locations. We develop a Mixed Integer Programming (MIP) model to minimize the total distribution time. To solve this model, we propose a Hybrid Ant

Colony Algorithm (HACA). Numerical experiments are conducted on a variety of instances, demonstrating that VRPD-FLRO significantly enhances the flexibility of drone launch and recovery operations. The drones are able to launch from their paired vehicles at optimal points along the routes, serve customers, and return to optimal rendezvous points. This flexibility effectively reduces drone flight distance, improves coverage, and minimizes synchronization time between vehicles and drones. The HACA efficiently allocates delivery tasks among vehicle-drone combinations, enabling the planning of VRPD-FLRO routes under tight time constraints and significantly enhancing delivery efficiency.

3 - Optimization of Urban Air Mobility Vertiport Location with Two-Stage Stochastic Programming

Wenjing LIU, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China, People's Republic of, Ziyu WANG, Shuai Jia

This study presents a stochastic optimization framework for vertiport location planning in urban air mobility (UAM) systems, addressing weather-induced flight route disruption uncertainties. We formulate the problem through a two-stage stochastic programming model that integrates randomly generated meteorological scenarios to characterize spatiotemporal connectivity patterns in vertiport network operations. To enhance computational tractability, the proposed methodology introduces a novel preprocessing framework for flight route feasibility screening in stochastic environments, contrasting with conventional exhaustive variable enumeration approaches. The study makes two contributions: (1) integrating weather uncertainties into vertiport location optimization, thereby enhancing the robustness of planning under dynamic conditions, and (2) developing a preprocessing methodology that significantly improves computational efficiency without compromising solution quality. This study provides detailed comparative analysis and demonstrates that the proposed preprocessing strategy achieves superior computational performance through reduced variable dimensionality. Experimental results validate the effectiveness and practical applicability of the models, offering insights into stochastic programming applications in UAM infrastructure planning. The analytical results highlight operational insights: (1) weather uncertainty consideration reduces network vulnerability across tested scenarios, and (2) preprocessing-driven dimension reduction maintains solution quality while decreasing model complexity. Finally, the study concludes with policy recommendations for resilient UAM infrastructure development, advocating for the integration of meteorological uncertainty into strategic planning processes and the adoption of computational preprocessing frameworks. This study advances stochastic programming applications in transportation network design under uncertainty, providing both theoretical and practical foundations for UAM systems.

4 - Mathematical Models for the Leather Nesting Problem

Jie Deng, Xiament University, Xiamen, China, People's Republic of, Sifan Cai

This paper studies the leather nesting problem derived from the manufacturing industry, which requires to optimally arrange leather patterns of irregular shapes and sizes to maximize the material utilization. Specially, limited by the raw materials, the boards may be irregular or even defective. Despite a large number of studies on two-dimensional irregular packing problems, in most of these problems the boards are rectangular and defect-free, which are quite different from our studied problem. To pack a set of irregular pieces into an irregular and defective board, we propose the neighbor points of the vertices for overlapping constraints and build the mathematical model that can be solved directly by commercial solvers. In order to obtain different solutions, we further introduce special points and line segments and add them as constraints to the model. Thus, four different models are constructed and investigated in total. The set of benchmark data are generated based on cutting bag-making material from crocodile leather, and the computational results demonstrate that the model can achieve good solutions on these small-scale instances.

WB10

Hullet Room

Modern Online Methods for Efficient Operations Management

Invited Session

Revenue Management and Pricing

Chair: Hanzhang Qin, National University of Singapore, Singapore, Singapore

1 - Learning Newsvendors: the Price of Algorithmic Instability

Jiaqi Lu, The Chinese University of Hong Kong (Shenzhen), Shenzhen, China, People's Republic of, Yilun Chen

Learning algorithms are increasingly central to how digital platforms allocate demand across service-side agents. In these environments, platforms often rely on bandit-style algorithms to explore agent quality and optimize long-run performance. However, these algorithms introduce stochasticity in how demand is allocated during learning — especially when agent performance is similar — resulting in high variance in agent exposure. In many real-world platforms, such variance imposes consequential agent behavior and costs. Unlike passive arms in classical bandit settings, these agents must frequently make one-shot, pre-committed resource decisions — such as setting availability schedules, ordering inventory, submitting proposal budgets, or selecting shifts — in anticipation of future platform demand. When realized traffic deviates from expectations due to platform-side exploration, agents experience idle time, wasted effort, or costly missed opportunities.

We propose a general framework for agent-aware platform learning, in which agents are modeled as strategic participants who make capacity planning decisions in the newsvendor fashion based on the random demand they expect. The platform allocates sequentially arriving demand via some learning algorithm, and wishes to minimize the cost of learning that comes from two aspects: 1) the expected allocation of demand to suboptimal arms due to incomplete information, i.e., the traditional regret; 2) the total deviation of the agents' endogenous capacity planning under random demand from the full information benchmark.

We show that there is a fundamental tension between these two terms, and propose a simple policy to optimally balance the tradeoff.

2 - Lean-LLM-OPT: a Lightweight Multi-Agent Llm Framework for Large-Scale Optimization Model Auto-Formulation

Hanzhang Qin, National University of Singapore, Singapore, Singapore, Kuo Liang, Yuhang Lu, Jianming Mao, Shuyi Sun, Chunwei Yang, Congcong Zeng, Ruihao Zhu, Chung Piaw Teo

As modern business problems grow increasingly complex, large-scale optimization modeling becomes a critical backbone of decision-making. However, traditional optimization modeling is often labor-intensive, which can be both costly and time-consuming. We address this by proposing LEAN-LLM-OPT, a novel Lightweight Multi-Agent LLM framework for Large-Scale Optimization Model Auto-Formulation, that takes a query (a problem description and associated datasets) as input and orchestrates a team of LLM agents to output the optimization formulation. LEAN-LLM-OPT innovatively uses few-shot learning to demonstrate how to use customized retrieval-augmented generation (RAG) tools to enhance LLM agents' capabilities. Specifically, upon receiving a query, a problem classification agent first determines the type of the problem. Then, a few-shot example generation agent consolidates a set of reference examples to demonstrate how optimization models are built for problems of same type. Finally, a model generation agent learns from these examples to extract relevant information from the input datasets using RAG tools and generate the final optimization model. As an intermediate step, we construct the first large-scale optimization model formulation benchmark Large-Scale-OR based on real-world datasets across different application domains. Extensive simulations demonstrate that \agent achieves state-of-the-art modeling accuracy compared to existing methods, especially on large-scale optimization modeling. Additionally, we validate its effectiveness through a Singapore Airlines flight scheduling use case. Our results thus offer novel insights on LLMs' capability in optimization modeling as few-shot learners. A demo of \agent is available at https://lean-opt-llm.streamlit.app/.

3 - Thompson Sampling for Repeated Newsvendor

Weizhou Zhang, National University of Singapore, Singapore, Singapore, Chen Li , Hanzhang Qin, Yunbei Xu, Ruihao Zhu

In this paper, we investigate the performance of Thompson Sampling (TS) for online learning with censored feedback, focusing primarily on the classic repeated newsvendor model--a foundational framework in inventory

management--and demonstrating how our techniques can be naturally extended to a broader class of problems. We model demand using a Weibull distribution and initialize TS with a Gamma prior to dynamically adjust order quantities. Our analysis establishes optimal (up to logarithmic factors) frequentist regret bounds for TS without imposing restrictive prior assumptions. More importantly, it yields novel and highly interpretable insights on how TS addresses the exploration-exploitation trade-off in the repeated newsvendor setting. Specifically, our results show that when past order quantities are sufficiently large to overcome censoring, TS accurately estimates the unknown demand parameters, leading to near-optimal ordering decisions. Conversely, when past orders are relatively small, TS automatically increases future order quantities to gather additional demand information. Extensive numerical simulations further demonstrate that TS outperforms more conservative and widely-used approaches such as online convex optimization, upper confidence bounds, and myopic Bayesian dynamic programming.
This study also lays the foundation for exploring general online learning problems with censored feedback.

4 - Online Resource Allocation with Limited Flexibility Under Adversarial Arrivals

Lyu Hongyao, Shanghai Jiaotong University, Shanghai, China, People's Republic of

An effective approach to mitigating demand uncertainty and minimizing supply-demand mismatches is process flexibility. In this paper, we investigate a class of online resource allocation problems with limited flexibility under adversarial arrivals. Our focus is on a specific type of structure—the long-chain design. This design has been both theoretically validated and practically implemented, making it a cornerstone in the development of sparse and flexible process configurations. We adopt an adversarial model in which the adversary, with full knowledge of the algorithm, generates sequences aimed at minimizing total fulfillment. we focus on a class of online greedy allocation policies. Our analysis demonstrates the effectiveness of the long-chain design in mitigating supply-demand mismatches. Specifically, we derive the tight lower bound of its worst-case performance, highlighting its resilience to supply-demand mismatches. Moreover, our findings offer actionable insights into the design and implementation of flexible processes, emphasizing the interplay between structural flexibility and policy effectiveness in real-world operational settings.

5 - Market Design of Contest Platforms

ZONGSEN YANG, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Xiao Lei, Pin GAO, Zizhuo Wang

We study the market design problem in an online contest platform (e.g., GGPoker) that repeatedly organizes large-scale contests to optimize long-term player retention. We build a novel dynamic model where the players' active probability depends on their skill levels and engagement states. Active players decide whether to participate based on their gains in each round. The platform must incentivize user participation and match them to maximize player engagement. When players' incentives are disregarded, we propose an outcome-based matching scheme that reallocates the number of winning points among players by matching high-level players with low-level ones. This scheme departs from traditional room-based approaches and can be framed as a downstream network flow problem. We prove that the common skill-based matching is optimal when games are highly competitive/non-competitive, or engagement states are limited. However, it could perform poorly compared to the optimal cross-level matching. We also find that when players suddenly disengage after loss streaks, the optimal policy would transfer winning points from the most and least engaged players to the second-least engaged players. However, the optimal policy may completely sacrifice the interests of low-level players and discourage them from joining in the contests. To resolve this, we design a token-based mechanism to incentivize players and recover the first-best outcome. Our results remain valid in the contexts of limited transfer or online arrivals. Our analysis provides actionable insights and practical solutions for contest platforms seeking to better entertain players.

WB11

Morrison Room

Data Analytics in Revenue Management

Invited Session

Revenue Management and Pricing

Chair: Chengzhang Li, Shanghai Jiao Tong University, Shanghai, N/A

Co-Chair: Lei Li, The Hong Kong Polytechnic University, M507n, Hong Kong, N/A

1 - Infrequent Resolving Algorithm for Online Linear Programming

Jingwei Zhang, The Chinese University of Hong Kong, Shenzhen, Shenzhen, China, People's Republic of, Guokai Li, Zizhuo Wang

Online linear programming (OLP) has gained significant attention from both researchers and practitioners due to its extensive applications, such as online auction, network revenue management, order fulfillment and advertising. Existing OLP algorithms fall into two categories: LP-based algorithms and LP-free algorithms. The former one typically guarantees better performance, even offering a constant regret, but requires solving a large number of LPs, which could be computationally expensive. In contrast, LP-free algorithm only requires first-order computations but induces a worse performance, lacking a constant regret bound. In this work, we bridge the gap between these two extremes by proposing a well-performing algorithm, that solves LPs at a few selected time points and conducts first-order computations at other time points. Specifically, for the case where the inputs are drawn from an unknown finite-support distribution, the proposed algorithm achieves a constant regret (even for the hard "degenerate" case) while solving LPs only O(loglogT) times over the time horizon T. Moreover, when we are allowed to solve LPs only M times, we design the corresponding schedule such that the proposed algorithm can guarantee a nearly O(T^{(1/2)^{M-1}}) regret. Our work highlights the value of resolving both at the beginning and the end of the selling horizon, and provides a novel framework to prove the performance guarantee of the proposed policy under different infrequent resolving schedules. Numerical experiments are conducted to demonstrate the efficiency of the proposed algorithms.

2 - Product Line Design and Pricing Under Mixed Multinomial Logit Model

Zhuodong Tang, Antai College of Economics and Management, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Guillermo Gallego, Anran Li

Product line design plays a crucial role in retailing, where optimizing both product configurations and pricing can maximize expected revenue. This paper examines a special case of the Mixed Multinomial Logit (MMNL) model in which the intrinsic utility of the outside option follows a specific distribution. We begin by analyzing the single-segment problem, which is NP-hard in general. To address this, we consider two families of utility functions: the finite elasticity family and the linearly compensable valuation family. For the finite elasticity family, where the elasticity of attraction is bounded, we develop a fully polynomial-time approximation scheme (FPTAS). For the linearly compensable valuation family, we show that the problem can be solved in polynomial time using the K-shortest path algorithm. Extending to multiple customer segments, we demonstrate that when the intrinsic utility of the outside option follows certain distributions (e.g., Gumbel distribution), our approach achieves an approximation ratio of 0.462. Finally, we conduct a case study using real-world laptop data, showing that our product line design algorithm improves expected revenue by 15%–44%.

3 - Algorithmic Collusion Meets Strategic Consumers: A Regulatory Perspective on Price Discrimination

Jiajing Li, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Yufeng Cao

As algorithms are increasingly employed to set prices dynamically, there is a growing concern about potential collusion among firms, leading to price discrimination. We analyze how strategic consumers, aware of such pricing algorithms, may alter their behavior in response to perceived or real algorithmic collusion. The study further examines the challenges this poses for regulators tasked with safeguarding fair competition and protecting consumer interests.

4 - Data-Integrated Capacity Planning and Allocation with Operational Data Analytics (Oda)

Lei Li, The Hong Kong Polytechnic University, Hong Kong, China, People's Republic of, Chengzhang Li

Decision-making with limited data and a lack of statistical characterization is challenging. Using a one-warehouse multi-retailer inventory system as an example, we study a data-driven procurement and allocation problem aimed at maximizing total expected profit across the warehouse and retailers, where the procurement lead time is one period and the allocation lead time to retailers is zero. Unlike existing literature that focuses on known demand distributions or asymptotic regimes, we concentrate on constructing data-integrated solutions under small samples by adopting the Operational Data Analytics (ODA) framework. This framework aims to identify the appropriate direct mapping from data to decision by balancing data integration and decision validation.

We consider this data-integrated procurement and allocation problem with location-scale demand and linear/newsvendor profit functions. Note that, if demand distributions are known, the theoretically optimal procurement and allocation decisions are equivariant in scale and shift to some extent. We explore the properties of procurement and allocation decisions to construct efficient data-integration models (solution classes) and validating models (performance measure), to determine proper data-integrated solutions in the solution classes. We demonstrate the performance of the proposed solutions both theoretically and numerically. Our approach underscores the role of domain knowledge and the structural relationships between the data and the decision in designing efficient data-integrated solutions under small samples.

WB12

Moor Room

Platform Service

Invited Session

Service Science

Chair: Hui Xiong, Donghua University, Shanghai, 200051

1 - Matching Strategy and Product Line Design on Home Service Platforms

Lu Hsiao, National Chung Hsing University, Taichung City, Taiwan, Hui Xiong, Ying-Ju Chen

In recent years, we have witnessed the growth of home service platforms. These platforms use commission-based contracts to recruit suppliers, provide multiple products to serve consumers, and match suppliers with the products. The platforms must optimize their matching strategies, commission-based contracts, and product quality decisions to improve the profit. Therefore, we build a model to investigate the matching strategies, the contract design, and the quality decisions in which suppliers with different service efficiency provide products for heterogeneous consumers. The consumers with relatively high (low) willingness to pay for quality belong to the high (low) segment. When low-segment consumers become very important (i.e., they take a large proportion, and their willingness to pay for quality is high), the platform may induce high-efficiency suppliers to offer low-quality products. Moreover, the platform is more likely to adopt this strategy when the matching probability for low-quality products is higher. However, when the heterogeneity between two segments is sufficiently small, the platform intends to induce high-efficiency suppliers to offer high-quality products and distort the quality level upward for the high segment. In this case, the quality difference between the two products increases in low-segment consumers’ willingness to pay for quality.

2 - Concerns for Others: Acts of Goodwill or Strategic Ploy

Hai Li, Central South University, Changsha, China, People's Republic of

The pursuit of social good by organizations has become increasingly common across various industries. This paper investigates whether the emphasis on consumer welfare by dual-purpose firms is simply an act of goodwill or a strategic maneuver to secure a competitive edge in competition with traditional for-profit firms. Our analysis shows that in efficient markets, where firms can increase their efforts cost-effectively, a dual-purpose firm may altruistically enhance consumer surplus at the expense of its own profit. However, this occurs only if its competitor is efficient; against an inefficient rival, this focus harms both the firm and consumers. In inefficient markets, where effort escalates less effectively, the dual-purpose firm acts altruistically or self-servingly with a less efficient rival, yet benefits itself while hurting consumers when facing an efficient one. Four extensions—covering general investment efficiency, spillover effects, imperfect substitutes, and parsimonious dual-purpose firms—confirm the robustness of these findings.

3 - Live Streaming Selling Under Consumers&Rsquo; Limited Attention

Jie Liu, Soochow University, Suzhou, China, People's Republic of, Hui Xiong, Ying-Ju Chen

In recent years, many sellers sell products through third-party live streamers and share revenue with these streamers. In the live streaming channel, a streamer can provide informative content about product attributes in real-time. Besides, the streamer can exert effort in non-informative content about social interactions to entertain consumers, and obtain revenue from consumers' gift-giving behavior. In this paper, we consider consumers'

limited attention to informative and noninformative content, and characterize the bargaining process between the seller and the streamer in game-theoretic models. In the scenario with an exogenous commission rate, we find that if the social interaction or gift-giving is not efficient, the increase in consumers' attention to social attention content leads the streamer to increase the social interaction effort, while the demand decreases. In this case, the reduction in attention to the quality causes a loss in the consumers' utility, which exceeds the gain from social interaction. Moreover, our finding shows that consumers' gift-giving behavior may hurt the streamer, which occurs when the effort cost exceeds the commission from product sales. Furthermore, if thecommission rate is negotiated by the seller and streamer, the commission rate decreases in the streamer's bargaining power when the streamer's effort efficiency is high. The concession on the commission rate encourages the seller to set a lower retail price, which leads to more sales.

4 - Platform&Rsquo;S Store Brand, Value-Added Service, and Demand Information Sharing

Hui Xiong, Donghua University, Shanghai, China, People's Republic of, Hai Li, Yifan He, Ying-Ju Chen

After decades of development and prosperity, many platforms, such as Amazon.com and JD.com, began to launch their store brands to improve their profits. To promote the sales of store brands, these platforms bundle value-added services to their store brands and use their superior demand information from serving third-party sellers to optimize the production quantity decision. However, it is common to see that the platforms share the information and provide value-added services for their competitors (third-party sellers). The incentive of information sharing and the influence of the value-added service on the information sharing decision are unclear. In this paper, we investigate the platform’s information sharing decision, the third-party seller’s choice of value-added service and its impact on the information sharing decision. We find that the platform may hide and take the advantage of the superior information when the commission rate and the substitutability of the store brand are low and the third-party seller does not purchase the value-added service. Then, the adoption of the value-added service increases the platform’s likelihood to share the demand information. Moreover, the third-party seller may purchase the value-added service when the service efficiency is high enough to increase the market size. However, although the service efficiency is sufficiently high, the third-party seller may abandon the value-added service when the platform shares the demand information and the demand information’s accuracy is relatively high.

5 - Monetization Strategies of Video-Sharing Platforms: the Role of the Platform-Creator Relationship

Haoyu Liu, University of Technology Sydney, Ultimo, Australia, Cheng-Hsun Hsieh, Lu Hsiao

Video-sharing platforms have seen remarkable growth in recent years. A key to sustaining this growth lies in how these platforms generate revenue. Common revenue-generating methods include sponsored advertisements, subscriptions, and commission fees. To study the monetization strategies of these platforms, we develop a game-theoretic model considering viewers have varying valuations of the videos and differing degrees of advertising aversion. Either the platform or the content creator decides the level of advertising and the pricing for a two-tier subscription scheme, depending on the platform-creator relationship. We find that when the creator decides the level of advertising and designs the subscription plan, with the platform taking a commission, the level of advertising rises, the video price decreases, but the ad-free video price increases. Further, the profits of the platform under different scenarios are compared. Findings from this paper provide valuable implications for platforms to select their monetization strategy based on their relationship with creators.

WB17

Bailey

Advanced Scheduling: Algorithms, AI, and Adaptive Optimization in Complex Systems

Contributed Session

Complex and Dynamic Systems

Chair: Rakesh Nagi, Singapore University of Technology and Design (SUTD), 8 Somapah Road, #07-101, Singapore, 487372, Singapore

1 - Enhancements in Job Shop Scheduling Problem Through Vector-Based Recalculations and An Early Termination Mechanism in a Perturbed Schedule

Golshan Madraki, University of New England, Biddeford, ME, United States, Seyedamirabbas Mousavian

Most heuristic and meta-heuristic algorithms solve scheduling problems by iteratively perturbing a trial schedule and recalculating a performance measure in each iteration to find an optimal/acceptable value for the performance measure. In the literature, Structural Perturbation Algorithms (SPAs) have been used to accelerate the calculation of performance measure in each iteration by reusing calculations from the previous iteration. However, SPAs have two limitations: they can only focus on a single performance measure (usually makespan), and they cannot terminate calculations when it is determined that a better value for a measure is unattainable. To overcome these limitations, we propose innovative Evaluation and Propagation Algorithms that can enhance any SPA. Our algorithms introduce two key features: a vector-based data structure that allows concurrent calculation of multiple performance measures, and an early termination mechanism that revolves around identifying triggering conditions. While the upgraded SPAs that result from our algorithm can reduce the computational complexity of the scheduling heuristics, they can also improve the quality of solution by considering multiple performance measures simultaneously. This study is theoretical, presented in generalized pseudocode and applied to numerical examples (instead of specific heuristics), emphasizing our algorithm’s applicability to a wide range of heuristics that employ graph-based data structures

2 - Mitigating Information Asymmetry in Project Scheduling: a Signaling Game Framework for Adaptive Buffer Sizing

Xi Wang, University of Science and Technology Beijing, Beijing, China, People's Republic of

Addressing information asymmetry between project managers and executors regarding resource allocation and risk perception, this study proposes a signaling game-based method for critical chain project buffer determination. A dynamic signaling game model is developed to simulate strategic interactions and information transmission mechanisms between managers and executors under asymmetric information. By incorporating managers' trust preferences toward executors and analyzing Bayesian equilibrium outcomes, the method enables adaptive buffer calibration and real-time adjustments. Monte Carlo simulations demonstrate that the proposed approach enhances comprehensive project performance, improving schedule adherence and reducing buffer overconsumption compared to conventional methods. The framework provides quantitative decision support for managers to evaluate executors' task reports and optimize buffer strategies in uncertainty-driven environments. Theoretical contributions include a novel integration of game theory with critical chain scheduling, while practical implications offer actionable guidelines for dynamic buffer management in construction, R&D, and IT project domains. Validation across simulated scenarios confirms the method's robustness against strategic misinformation and its scalability for multi-phase projects.

3 - Optimizing Multi-Stage Inventory Management Through Deep Reinforcement Learning

Tolunay Alankaya, Eindhoven University of Technology, Eindhoven, Netherlands

Effective inventory management is crucial when decisions must be made repeatedly over time. We examine a multi-stage inventory setting where a manufacturer relies on dealers to serve end customers. Instead of selling directly, the manufacturer provides order suggestions to dealers, who can accept or reject them. This interplay between manufacturer and dealers creates a complex environment shaped by demand variability, lead times, and cost trade-offs.

Our main goal is to minimize lost sales while managing inventory expenses, striking a balance between product availability and overstocking risks. To address the high dimensionality and uncertainty in these systems, we propose a deep reinforcement learning approach that adapts decision policies over time. By modeling state transitions, reward functions, and action spaces at each stage, our method captures the stochastic nature of demand and incorporates decentralized decision-making across multiple dealers.

We detail the key elements of our framework, including how centralized guidance aligns with local choices. Although we have not yet obtained final results, initial insights suggest that advanced machine learning holds promise for optimizing multi-stage inventory. Our future work will involve rigorous simulation studies to validate the framework and compare its performance to conventional inventory policies. By embracing a data-driven approach, this research aims to enhance supply chain resilience, cost efficiency, and service quality.

4 - Learning Heuristic Selection for Cloud Workflow Scheduling with a Deep Neural Network

Rakesh Nagi, Singapore University of Technology and Design (SUTD), Singapore, Singapore, Jungeun Shin

Efficient workflow scheduling is essential for optimizing resource utilization and minimizing job completion times in cloud-native environments, where diverse and complex workflows must be managed dynamically. Traditional rule-based scheduling heuristics, while interpretable and easy to implement, often lack the adaptability required to respond to variations in workflow structures and cloud resource availability. More recently, learning-based scheduling algorithms have been proposed to address these limitations; however, they often suffer from poor generalization to unseen workflow types, limited interpretability, and high retraining costs when adapting to changing system conditions.

To overcome these challenges, we propose a deep reinforcement learning (DRL)-based hyper-heuristic framework that autonomously selects the most effective job priority heuristic based on real-time cloud status and workflow characteristics. Our approach integrates graph neural networks (GNNs) within the policy network to process Directed Acyclic Graph (DAG) representations of workflows, enabling efficient extraction of structural information for scheduling decisions. By leveraging embedded inputs such as workload durations, task dependencies, and critical paths using GNNs, the DRL model learns robust scheduling policies that generalize across diverse and unpredictable workload scenarios. Experimental results demonstrate that the proposed method significantly improves resource utilization and reduces job completion times, offering a scalable and adaptable solution for modern cloud-native computing.

WB18

Wilberforce

Platforms, Innovation, and Engagement

Invited Session

Informations Systems

Chair: Qian Tang, Singapore Management Unviersity, Singapore, Singapore

1 - From Bytes to Bites: Examining the Online and Offline Spillover Impacts of Online Food Delivery Platform Subscriptions

Keongtae Kim, Chinese University of Hong Kong, Shatin, Hong Kong, Jaeho Myeong, Yongjin Park, Francis Joseph Costello, Jae-Hyeon Ahn

With the rapid growth of online food delivery (OFD) platforms and intensifying competition, platform subscriptions have emerged as a key strategy to retain consumers and influence their behavior. This study explores how OFD subscriptions affect consumer ordering patterns and offline dining behavior in a developed Asian market. Using a Difference-in-Differences approach with a large-scale individual-level transaction dataset, we find online spillovers between the focal and competing platform after the release of the subscription together with lock-in effects on the focal platform depending on the subscription upfront costs incurred. However, alongside this shift, we did not find cannibalization, rather the overall OFD market share increased, leading to offline spillovers that enhanced local dining and consumption of proximally close offline non-food products and services. These findings show how platform subscriptions in location-based services (LBS) like OFD reshape consumer mobility patterns and stimulate not only the focal sector but also non-competing sectors within the local economy. The results provide actionable implications for platform operators to refine subscription strategies, for policymakers and urban planners to manage and embrace LBS platform creation and proliferation thanks to local compensatory offline spillovers leading to greater vitality, and for local restaurant owners to capitalize on increased patronage opportunities via platform exposure.

2 - Power Dependence in Bilateral Exchange in Live Streaming Platforms

Chong Wang, City University of Hong Kong, Hong Kong, Hong Kong, Ningzhe Zhou, Xinyan Yao, Hongju Liu

Live streaming features distinctive exchange relationships between streamers (content creators) and viewers (consumers) through pay-what-you-want (PWYW) pricing strategies in real-time, one-to-many online interactions. Drawing on power-dependence theory, we examine how power structure of the streamer-viewer relationship – specifically, relative and joint power– influence voluntary payments (tips) and other user engagement. By analyzing session-level interactions from over three million streamer-viewer pairs on a Chinese entertainment live-streaming platform over two weeks, we find that viewers with greater relative power tend to tip less, while higher joint power increases tipping. Furthermore, both streamers and viewers exhibit adaptive withdrawal behaviors when experiencing power disadvantages: streamers reduce their streaming frequency, and viewers decrease their watching time. Our research extends the understanding of the live streaming market by highlighting the critical role of streamer-viewer exchange relationships and providing empirical evidence on how power structure influences tipping and user engagement.

3 - New Ideas vs. Novel Recombinations: Unpacking Their Impacts on Innovation Performance in the Movie Industry

Qian Tang, Singapore Management University, Singapore, Singapore, Shaoyi Liaw, Heli Wang

Innovation can stem from groundbreaking new ideas, novel recombination of existing ideas, or a combination of both. This paper investigates how the introduction of new ideas and the recombination of existing ideas jointly drive innovation success within the dynamic landscape of the movie industry. Leveraging human annotation on the platform of TvTropes.org, we decompose each movie into a set of tropes used, including new tropes debut by the movie and existing ones by previous movies. With such decomposition and movie data from IMDb and Box Office Mojo, we analyze how introducing new tropes combined with existing ones affects the quality and economic success of a movie. We find that new ideas positively affect a movie's economic success, while recombination positively affects product quality in terms of movie ratings and award nomination. We also find that recombination negatively moderates the effect of new idea on product quality. Insights from this analysis can provide recommendations to filmmakers and industry stakeholders to strategically leverage new and existing idea combinations for product innovation to foster creativity, enhance audience engagement, and improve the success of their products.

4 - Unpacking the AI Transformation: the Impact of Ai Strategies on Firm Performance from the Dynamic Capabilities Perspective

Jaecheol Park, University of British Columbia, Vancouver, BC, Canada, myunghwan lee, J. Frank Li, Gene Moo Lee

Artificial intelligence (AI) technologies hold great potential for large-scale economic impacts. Aligned with this trend, recent studies explore the impact of AI technologies on firm performance. However, they predominantly measure firms' AI capabilities with input (e.g., labor/job posting) or output (e.g., patents), neglecting to consider the strategic direction toward AI in business operations and value creation. In this paper, we empirically examine how firms' AI strategic orientation affects firm performance from a dynamic capabilities perspective. We create a novel firm-year AI strategic orientation measure by employing a large language model to analyze business descriptions in 10-K filings and identify an increasing trend and changing status of AI strategies among U.S. public firms. Our long-difference analysis shows that AI strategic orientation is associated with greater operating cost, capital expenditure, and market value but not sales, showing the importance of strategic direction toward AI to create business value. By further dissecting firms' AI strategic orientation into AI awareness, AI product orientation, and AI process orientation, we find that AI awareness is generally not related to performance, that AI product orientation is associated with short-term increased operating expenses and long-term market value, and that AI process orientation is associated with long-term increased costs and sales. Moreover, we find the moderating effect of environmental dynamism. This study contributes to the recent AI strategy and management literature by providing the strategic role of AI orientation on firm performance.

5 - Perspective Matters: How Visual Perspective, Follower Count, and Sponsorship Disclosure Shape Video Engagement

Yan Lin, University of Macau, Macao, Macau, Yutong Ma, Yongdan Liu, Matthew LIU

The proliferation of user-generated content platforms has enabled millions of ordinary people to create and share videos online. Since these videos are often captured by handheld devices like smartphones and action cameras

like GoPro and Insta360, it's not uncommon to see videos captured in first-person perspective. This study explores the differential impact of visual perspectives (first-person vs. third-person) on video engagement for macro- and micro-influencers. Drawing from the processing fluency theory, we hypothesize that congruence between the video's visual perspective and the influencer's follower size results in better viewer engagement. Using TikTok data, we find that first-person perspective videos are associated with increased viewer engagement for micro-influencers but decreased engagement for macro-influencers, after controlling for confounding factors and potential selection issues. We conducted two controlled experiments to substantiate the findings and explore the underlying mechanism. Our results suggest that processing fluency explains why first-person perspective videos have opposing effects on viewer engagement for micro-influencers versus macroinfluencers. Moreover, this effect is less pronounced for videos with explicit sponsorship disclosure, as the presence of sponsorship information leads to lower processing fluency due to incongruence created by high psychological distance and profit motive. These findings deepen our understanding of how visual perspective shapes viewer responses and provide actionable insights for crafting more effective social media content.

Wednesday, July 23, 11:45 AM - 1:15 PM

WC01

Enterprise

Probability and Optimization for Statistical Learning

Invited Session

Applied Probability Society

Chair: Sarat Moka, The University of New South Wales, Sydney, Australia

Co-Chair: Anant Mathur, University of New South Wales (Sydney), Sydney, Australia

1 - Sampling with Symmetric Minibatch Splitting Langevin Dynamics

Daniel Paulin, Nanyang Technological University, Singapore, Singapore, Singapore

We propose a scalable kinetic Langevin dynamics algorithm for sampling parameter spaces of big data and AI applications. Our scheme combines a symmetric forward/backward sweep over minibatches with a symmetric discretization of Langevin dynamics. For a particular Langevin splitting method (UBU), we show that the resulting Symmetric Minibatch Splitting-UBU (SMS-UBU) integrator has bias $O(h^2 d^{0.5})$ in dimension $d>0$ with stepsize $h>0$, despite only using one minibatch per iteration, thus providing excellent control of the sampling bias as a function of the stepsize.

We apply the algorithm to explore local modes of the posterior distribution of Bayesian neural networks (BNNs) and evaluate the calibration performance of the posterior predictive probabilities for neural networks with convolutional neural network architectures for classification problems on three different datasets (Fashion-MNIST, Celeb-A and chest X-ray). Our results indicate that BNNs sampled with SMS-UBU can offer significantly better calibration performance compared to standard methods of training and stochastic weight averaging.

2 - A Fast Mm Algorithm for Group Lasso

Anant Mathur, University of New South Wales (Sydney), Sydney, Australia

The group lasso regression is a popular method for performing model selection on groups of predictors with a natural structure. When each group is orthogonalized, the optimization yields the widely-used block coordinate descent (BCD) algorithm. Despite its simplicity, the BCD algorithm converges slowly when the matrix of features is poorly conditioned or when the solution is not highly sparse. This talk presents a novel iterative algorithm for the group lasso based on the majorize-minimize (MM) principle. Unlike previous algorithms, the proposed MM algorithm optimizes a reparameterized objective for the group lasso that maintains the same Karush-Kuhn-Tucker (KKT) conditions as the original problem. This change of variables enables an adaptive MM algorithm that, at each iteration, prioritizes the groups that are furthest from meeting the KKT conditions.

As a result the convergence rate of the MM algorithm is accelerated with minimal additional cost. We establish the global convergence of the MM algorithm to the group-lasso KKT point. Using real data and simulations, we demonstrate numerically that the MM algorithm converges up to an order of magnitude faster than the BCD algorithm when computed over a regularization path.

3 - Improving the Convergence of Markov Chains via Permutations and Projections

Michael Choi, National University of Singapore, Singapore, Singapore

This talk aims at improving the convergence to equilibrium of finite Markov chains via permutations and projections. First, we prove that a specific mixture of permuted Markov chains arises naturally as a projection under the KL divergence or the squared-Frobenius norm. We then compare various mixing properties of the mixture with other competing Markov chain samplers and demonstrate that it enjoys improved convergence. This geometric perspective motivates us to propose samplers based on alternating projections to combine different permutations and to analyze its rate of convergence. We give necessary, and under some additional assumptions also sufficient, condition for the projection to achieve stationarity in the limit in terms of trace of the transition matrix. We proceed to discuss tuning strategies of the projection samplers when these permutations are viewed as parameters. Along the way, we reveal connections between the mixture and a Markov chain Sylvester's equation as well as assignment problems, and highlight how these can be used to understand and improve Markov chain mixing. We provide two toy examples as illustrations. This is based on joint work with Max Hird and Youjia Wang.

WC03

Minto

Computational Optimization under Uncertainty

Invited Session

Computing Society

Chair: Sarah Yini Gao, Singapore Management University, Singapore, Singapore

1 - Data-Driven Markov Decision Processes with Heterogeneous Groups: Regret Guarantees and Fairness Improvement via Robust Optimization

Yuzhou Liu, Singapore Management University, Singapore, Singapore, Sarah Yini Gao, Zhichao Zheng

This paper investigates the efficiency and fairness implications of robust optimization through a finite-horizon Markov Decision Process perspective. We consider heterogeneous groups that differ in rewards and transition kernels, where the underlying parameters for each group can only be estimated from historical data. Additionally, the data available across groups is imbalanced, with significantly fewer samples from the minority group compared to the majority group. We first derive a probabilistic regret bound for the policy obtained from empirical estimations, indicating that the transition uncertainty dominates the deviation. Subsequently, we illustrate the fairness benefit of leveraging abundant pooled data, which improves decision-making for the minority group.

To address parameter uncertainty, we further employ a distributionally robust model using a type-1 Wasserstein ambiguity set. We theoretically demonstrate that adopting such robust approaches does not lead to a worse regret bound, provided the level of robustness is moderate. Intriguingly, we also establish that increased robustness within a certain threshold can monotonically enhance overall performance. This result analytically reinforces the fundamental logic of robust optimization —that an appropriate degree of robustness effectively mitigates the parameter uncertainty risk. Moreover, we reveal that the fairness benefit is amplified by incorporating robustness, becoming especially prominent when data imbalance is severe, thus underscoring the inherent fairness-enhancing effect of robust optimization.

Numerical experiments conducted on a medical extubation problem substantiate our theoretical findings, showcasing the robust model's superiority in both collective performance enhancement and fairness improvement.

2 - Managing Liability Limits in Volatile Markets: the Multi-Product Newsvendor Framework with Random Payouts

Libo Sun, Hefei University of Technology, Hefei, China, People's Republic of, Chung-Piaw Teo, Chengyu Wang, Leemon n.a., Rahul Kataria

This paper explores optimizing product liability limits against rare but substantial payouts, such as credit defaults, insurance claims, and lottery prizes. We propose a risk-averse multi-product newsvendor framework that integrates random demands and payouts, introducing an *Excess Payout Risk* measure to align actual payout ratios with expectations. To tackle the problem's inherent non-linearity, we develop a Variability Risk Optimization (VRO) model, which typically requires integer solutions for practical decision-making. We demonstrate that the VRO model supports an efficient *Mixed-Integer Linear Programming* (MILP) reformulation for uniform margin systems, where revenues and payout ratios are consistent within markets. For general systems, we extend the model to a *Mixed-Integer Second-Order Cone Programming* (MISOCP) framework, augmented with *Sample Average Approximation* (SAA) to manage uncertain demands. Synthetic experiments validate the MISOCP method in multi-market, multi-product scenarios.

Backtesting this methodology for an Asian lottery operator showed the potential to increase legal revenues by tens of millions USD. The change was implemented in 2022, following the COVID-19 pandemic. Post-implementation analysis revealed that more bets could be accepted by the legal operator, thereby diverting more than 2.6% of sales that might have been lost due to various factors, including illegal competitors.

3 - Improved Convergence Rate of Noisy Stochastic Gradient Descent for Interacting Langevin Dynamics

Zhi Wang, University of Texas at Austin, Austin, TX, United States, Rui Gao, Shuang Li

Mean-Field Langevin Dynamics provides a powerful framework for analyzing interacting particle systems in overparameterized machine learning models that are usually highly nonconvex. While prior work has established convergence rates for discrete-time noisy stochastic gradient descent in finite-particle systems, existing results often require large mini-batches to achieve nearly-optimal rates. Our analysis establishes a nearly optimal convergence rate of $O(1/\epsilon \log(1/\epsilon))$, improving existing results that achieve this rate only with large mini-batch sizes and otherwise degrade to $O(1/\epsilon^2 \log(1/\epsilon))$. This improvement stems from our novel discrete-time approach, which directly examines per-step behavior rather than relying on continuous-time interpolation between discrete-time updates. This general framework provides a global convergence solution to many applications in OM field, e.g. estimating mixed logit choice model.

4 - A CoP-Cp Approach to Max-Min Bilinear Conic Problems

Sarah Yini Gao, Singapore Management University, Singapore, Singapore, Yancheng Yuan, Xindong Tang

We study a max–min bilinear optimization problem over completely positive cones—a formulation that naturally arises in robust optimization and two-person zero-sum games involving mixed-binary nonconvex quadratics. We show this problem can be reformulated as a single-stage linear program over the Cartesian product of the copositive (COP) cone and the completely positive (CP) cone, known as a COP-CP program. We then present a moment- and sum-of-squares–based relaxation framework for solving the COP-CP program, along with conditions under which these relaxations are guaranteed to be exact.

WC04

Indiana

AI in Maritime Logistics, Supply Chains, and Global Trade

Invited Session

Maritime AI

Chair: Mikael Lind

1 - Tramp Fleet Renewal and Contract of Affreightment Decisions under Uncertainty: An Approximate Dynamic Programming Approach

Sirui Ding, Tsinghua University, Beijing, China, People's Republic of

With rising geopolitical tensions and economic instability, tramp shipping markets are facing heightened volatility. This makes it essential for shipping companies to incorporate real-world uncertainty into long-term fleet renewal strategies—a cornerstone of maritime transportation planning. Furthermore, equally critical is multi-period affreightment contracting, which affects profitability by balancing spot and contract cargo under shifting market conditions. While fleet renewal has received some attention, research on multi-period contract decisions under uncertainty remains limited. Besides, when modeling market uncertainty, expert insights and the interconnected movements of various market prices are vital. Yet, many studies often neglect this knowledge, limiting their practical relevance. Moreover, most rely on Sample Average Approximation (SAA) to solve stochastic programs, which becomes computationally intractable when dealing with high-dimensional uncertainty and long planning horizons. To address these gaps, we propose a stochastic programming model that jointly considers vessel acquisition (newbuild and secondhand), chartering, and cargo contracting decisions. We introduce a hybrid forecasting framework that combines expert-driven qualitative input with a Markov chain-based market-state transition process, generating realistic scenarios through historical data bootstrapping. To manage computational complexity, we reformulate the problem as a dynamic programming and solve it using Approximate Dynamic Programming (ADP) with neural network-based value function approximation. A double value network is employed to reduce overestimation bias. Numerical experiments based on a Chinese crude oil shipping company demonstrate that our approach significantly reduces operational costs and enhances decision quality. The findings provide practical insights and a scalable decision-support tool for managing uncertainty in tramp shipping.

2 - Reducing port city congestion through leveraging data and artificial intelligence

Wolfgang Lehmacher, On behalf of Research Institutes of Sweden (RISE), Gothenburg, Sweden

3 - The Future of Port Operations: How Digital Twins and AI Analytics Enable Autonomous Decision-Making

Mengya Liu, National University of Singapore, Singapore, Singapore

4 - Early Indicators of Port Congestion for Facilitating Maritime Shipping: Evidence from Big Maritime Data

Jiaxiang Cai, Institute of High Performance Computing, ASTAR, Singapore, Singapore, Liangbin Zhao, xiuju Fu, Zheng Qin

Port congestion presents persistent challenges to global maritime shipping, resulting in delays, increased operational costs, and disruptions across supply chains. This study uses big maritime data to identify early indicators of congestion, offering predictive insights that support proactive operational decisions. We analyze approximately 18 months of Automatic Identification System (AIS) data, spanning from 5 October 2023 to 20 April 2025, with a particular focus on the 2024 congestion episode at the Port of Singapore. A spatial-temporal clustering algorithm is developed to identify and summarize vessel movements in a region centered on Singapore. Several key indicators are examined, including vessel arrival patterns (port, anchorage, and berth), anchorage waiting times, berth/port dwell times, and seasonal traffic variations. Our findings highlight that the 90th percentile of waiting times before the first berth and average berth/port turnaround times are the most reliable early warning indicators, typically offering a lead time of 3–7 days before congestion peaks. These advance signals enable stakeholders to implement mitigation strategies such as rerouting, schedule adjustments, or negotiating priority berthing. The 2024 surge in container vessel arrivals is identified as the primary driver of congestion, while port infrastructure capacity and responsive policies critically shape the speed of recovery. This research provides a practical, data-driven framework for early congestion detection, contributing to more resilient maritime logistics. It offers actionable intelligence for port authorities, shipping companies, and logistics providers aiming to reduce congestion-related disruptions and improve overall port efficiency.

WC05

Orchard A

System Risk Quantification and Simulation

Invited Session

Efficiency and Fairness in Supply Chain and Service Operations Management

Chair: Kun Zhang, Renmin University of China, Beijing, N/A, China, People's Republic of

1 - End-to-End Portfolio Selection: Leveraging Deep Learning in Smart Predict Then Optimize

Wenchu Zhang, Tongji University, Shanghai, China, People's Republic of, Zhaolin Hu, Haiqing Yin

The end-to-end prediction and optimization paradigm based on Smart ``Predict, then Optimize'' (SPO) method provides a new approach for portfolio selection. In this paper, we study the end-to-end portfolio selection by leveraging deep learning in the SPO framework. We consider several deep learning models including two linear models, two recurrent neural network models, and two transformer-based models for financial time series forecasting, and incorporate them in the SPO loss functions in the mean-variance models. We train the end-to-end models using the Adam algorithm based on the specific structure of the SPO+ function to directly obtain the optimal portfolio weights. We compare via the real test the end-to-end paradigm with the two-stage prediction and optimization paradigm that uses the conventional Mean Squared Error (MSE) loss function, and illustrate the competitiveness of the former. Our tests also suggest that the iTransformer model appears to hold some superiority among the various deep learning models in this end-to-end portfolio selection practice.

2 - A Correlation-Induced Finite Difference Estimator

Kun Zhang, Renmin University of China, Beijing, China, People's Republic of, Guo Liang, Guangwu Liu

Finite difference (FD) approximation is a classic approach to stochastic gradient estimation when only noisy function realizations are available. The optimal FD estimator is constructed assuming known optimal perturbation, which is rarely the case in practice. This paper establishes a novel sample-driven method that leverages bootstrap and regression techniques to estimate the optimal perturbation based on all simulated samples. We then normalize and transform these samples according to the estimated optimal perturbation, leading to correlated samples. Using these correlated samples, we propose an efficient FD estimator. Theoretical analyses of both the perturbation estimator and the FD estimator reveal that, *surprisingly*, the correlation enables the proposed FD estimator to achieve a reduction in variance and, in some cases, a decrease in bias compared to the optimal FD estimator. Numerical results confirm the efficiency of our estimators and align well with the theory presented, especially in scenarios with small sample sizes. Finally, we apply the estimator to solve derivative-free optimization (DFO) problems, and numerical studies show that DFO problems with 100 dimensions can be effectively addressed.

3 - Derivative-Free Optimization via Finite Difference Approximation: An Experimental Study

Du-Yi Wang, City University of Hong Kong, Kowloon, Hong Kong, Guo Liang, Kun Zhang, Guangwu Liu

Derivative-free optimization (DFO) is vital in solving complex optimization problems where only noisy function evaluations are available through an oracle. Within this domain, DFO via finite difference (FD) approximation has emerged as a powerful method. Two classical approaches are the Kiefer-Wolfowitz (KW) and simultaneous perturbation stochastic approximation (SPSA) algorithms, which estimate gradients using just two samples in each iteration to conserve samples. However, this approach yields imprecise gradient estimators, necessitating diminishing step sizes to ensure convergence, often resulting in slow optimization progress. In contrast, FD estimators constructed from batch samples approximate gradients more accurately. While gradient descent algorithms using batch-based FD estimators achieve more precise results in each iteration, they require more samples and permit fewer iterations. This raises a fundamental question: which approach is more effective—KW-style methods or DFO with batch-based FD estimators? This paper conducts a comprehensive experimental comparison among these approaches, examining the fundamental trade-off between gradient estimation accuracy and iteration steps. Through extensive experiments in both low- and high-dimensional settings, as well as hyperparameter optimization adjusting the penalty parameter in ridge regression, we demonstrate a surprising finding: when an efficient batch-based FD estimator is applied, its corresponding gradient descent algorithm generally shows better performance compared to classical KW and SPSA algorithms in our tested scenarios.

4 - Optimal Hedging for Automated Market Maker Liquidity Provision with Asymmetric Ranges and Stop-Loss Strategies

Jiancheng Qin, City University of Hong Kong, Kowloon, Hong Kong, Guangwu Liu

This study introduces a framework for liquidity provision in concentrate liquidity market maker (CLMM), leveraging asymmetric price ranges and stop-loss thresholds to enhance risk management. We employ a stochastic price model to derive analytical expressions for expected fee income, terminal position value, and rebalancing costs across a liquidity provision cycle. To mitigate risk, we incorporate a convex risk measure and propose a hedging strategy comprising static components (e.g., options or forwards) and dynamic components (continuous trading adjustments). The framework optimizes expected yield—defined as profit per unit time—by determining optimal range and stop-loss parameters. This optimization balances fee income, position value fluctuations, and transaction costs, enabling liquidity providers to maximize returns while minimizing impermanent loss. Our approach offers a practical tool for decision-making in decentralized finance, contributing to the literature on automated market maker protocols.

WC10

Hullet Room

Analytics for On-demand Retailing and Service Platforms

Invited Session

Revenue Management and Pricing

Chair: Mingzheng Wang, Zhejiang University, Hangzhou, N/A, China, People's Republic of

Co-Chair: Chenrui Li, Zhejiang University, Hangzhou, N/A

1 - Pricing for Hybrid Cloud Provisioning with Small-Scale and Dual-Source Data

Guodong Yu, Shandong University, Jinan, China, People's Republic of, Huiping Sun

This paper considers the supplier's pricing problem for customized upfront reservation discount (URD) contracts, where a newsvendor buyer reserves resources in advance and then adjusts his purchases based on realized end-market demand, e.g., either procuring additional products at market prices to address shortages or returning surplus inventory through buyback provisions. Specifically, we focus on the supplier's pricing challenge of data scarcity due to information asymmetry, as the supplier possesses an inferior knowledge of end-market demand compared to her buyer. To address this issue, we propose a refined data-driven Wasserstein distributionally robust optimization model for URD pricing—a bilevel model that aims to maximize the supplier's worst-case profit. Beyond the conventional reliance on historical demand data, our approach leverages past transaction records involving supplier-buyer interactions, allowing the supplier to better infer true demand patterns. This is achieved through incorporating inverse optimization underlying the first-order conditions of the buyer's newsvendor behavior into a Wasserstein minimization problem. We explicitly demonstrate improvements in minimizing the Wasserstein distance and stabilizing the supplier's profit, theoretically proving the benefits of our method as long as the buyer retains an information advantage over the supplier, regardless of whether the buyer is clairvoyant or non-clairvoyant. To tackle the computational challenges in estimating end-market demand via Wasserstein minimization with dual-source data, we develop a tractable linear model via set-partitioning and Lagrange duality techniques. Extensive theoretical and empirical results in cloud reservation services demonstrate the effectiveness of our approach。

2 - Governmental Incentive Policies for Developing Renewable Energy with Storage Facilitates

Yun Zhou, Zhejiang University of Technology, Hangzhou, China, People's Republic of, Menghai Li, Jian Cao, Mingzheng Wang

Developing renewable energy with storage facilitates is important to increase the penetration of renewable energy and achieve the renewable energy consumption target. Because of its high investment cost, most of renewable energy producers can not or are unwilling to install energy storage facilities. The government is faced with a challenge about how to implement incentive policies in the most cost-effective manner. This paper explores how would the incentive policies affect the installation of renewable energy generators and energy storage facilities, the realization of social welfare through Stackelberg games. The results reveals that the optimal incentive policy depends on the target of renewable energy consumption and unit construction cost of energy storage facilities: (1) if the target is relatively small, no external incentive is needed; (2) as the target increases, it is optimal for the government to implement investment policy to drive the construction of energy storage facilities invested by renewable energy producers; (3) when the target and construction cost coefficient reach a sufficiently high level, it is optimal for government to simultaneously implement investment policy and feed-in-tariffs(FITs). Considering provision of capacity-based tariffs(CBTs) is not optimal for government because there is a substitution effect of FITs on CBTs. Moreover, there may be a substitution or complementary effect among the three incentive policies. We employ the real-world data from China to empirically verify the results. Our study resonates with the policy evolution in China and can serve as guidelines for governments when developing incentive policies to support the achievement of targets.

3 - Managing Hybrid Workforce and Hybrid Demand: Optimal Order Assignment Policy for on-Demand Delivery Platform

Mingzheng Wang, Zhejiang University, Hangzhou, China, People's Republic of, Haohan Li, Yun Fong Lim, Hai Wang

We study an order assignment problem with hybrid demand in a hybrid workforce on-demand delivery platform. We develop an infinite horizon discrete-time event-based Markov Decision Process (MDP) to maximize the expected total discounted profit of the platform. We prove this MDP has favorable properties such as submodularity, subconcavity, and concavity. Leveraging these properties, we propose an optimal order assignment policy using rationing and priority thresholds which are dependent on the number of busy drivers. Numerical studies show significant improvements of our policy using synthetic and real-world data. The average improvement is 12.34\%, while the maximum improvement can reach 30.87\%.

4 - Robust Facility Location and Logistics Design with Drones Under Disruption

Zhi Pei, Zhejiang University of Technology, 杭州, China, People's Republic of

As a responsive and flexible way of transportation, drones are employed to conduct logistics distribution activities in daily life. Since the drones considered here are station- based, the location of station is a vital decision to make, taking into account the risk of facility disruptions. In the present study, we apply a two-stage distributionally robust optimization framework to depict the facility location and transportation problem with drones, where the risk of facility disruptions and the uncertainty of demand are simultaneously considered. The aim is to achieve the best decision on the quantity of facilities established ex ante, as well as the allocation of drones to match the customer demand. To capture the practical drone delivery features, the energy level and payload capacity of drones are incorporated, and the split delivery is assumed. Two exact algorithms with valid inequalities are introduced, namely, the nested and the enumeration-based column-and-constraint generation. Numerical experiments show that the latter strategy performs much better in terms of efficiency and solvability. To showcase the importance of including risk of facility disruptions and uncertainty of demand, the overall costs obtained under various problem settings are listed in the sensitivity analysis. And in the robustness testing, it is observed that the two- stage distributionally robust optimization approach performs better than the two-stage robust optimization method and the deterministic optimization method.

WC11

Morrison Room

Data-driven Pricing and Assortment

Invited Session

Revenue Management and Pricing

Chair: Hailong Sun, Shanghai Jiao Tong University, Shanghai, N/A

Co-Chair: Hongyuan Lin, N/A

1 - Fairness-Aware Dynamic Assortment Selection

Guangyan Gan, NTU, Singapore, Singapore, Hanzhang Qin, Zhenzhen Yan

Traditional assortment planning typically aims to optimize the selection of items offered to customers, often with the primary goal of maximizing revenue. However, such approaches may inadvertently compromise the equality of opportunity among items, leading to diminished content diversity on the platform. To address this challenge, we introduce and analyze a fair dynamic assortment planning problem that explicitly ensures items of comparable quality receive comparable exposure. We develop a novel algorithm that balances fairness with performance, and empirically evaluate its effectiveness through a case study using the real dataset. Our results highlight the algorithm's ability to enhance fairness while maintaining competitive performance, and shed light on the inherent trade-offs between fairness and utility in dynamic assortment settings.

2 - Robust Pricing with Points: the Role of Randomization

Hongyuan Lin, University of Science and Technology of China, Hefei, China, People's Republic of, Mingyang Fu, Chenyu Xue

This paper addresses a data-driven pricing problem where a seller aims to optimize the sale of a single item to a buyer using historical data consisting of observed price and conversion rate pairs. The objective is to design a pricing strategy that is robust across multiple criteria in the worst-case scenario, given demand uncertainty. We establish a simple and easily verifiable necessary and sufficient condition for the existence of a randomized pricing mechanism that satisfies performance targets across three key metrics: worst-case expected revenue, worst-case ratio, and worst-case regret. We propose a novel class of randomized pricing mechanisms that can balance all three metrics simultaneously, and any attainable combination of them can be achieved through this mechanism. In addition, we characterize the Pareto frontier for these three metrics, providing insights into the trade-offs between them and offering valuable guidance for sellers facing pricing decisions under uncertainty. Our analysis also demonstrates that even slight randomization, with a modest sacrifice in worst-case revenue, can significantly improve robustness regarding ratio and regret. Furthermore, we identify and compare the optimal robust mechanisms for each criterion, providing practical guidance for single-criterion selection.

3 - Pricing Analytics for Primary and Ancillary Products Using Conversion Rate Data

Hailong Sun, Shanghai Jiao Tong University, Shanghai, China, People's Republic of, Changchun Liu, Chung Piaw Teo

Ancillary products, such as in-flight WiFi, hotel upgrades, and extended warranties, contribute significantly to revenue across various industries. However, pricing them effectively is challenging due to the interdependence between primary and ancillary purchases and the complex substitution patterns in customer choice behavior. Traditional pricing methods often fail to capture these dynamics, leading to suboptimal revenue outcomes. This paper introduces a data-driven approach to pricing that leverages conversion rate data to optimize the pricing of both primary and ancillary products. The approach is based on an additive perturbed utility model (APUM), which systematically accounts for how customers evaluate and substitute between different product options. Leveraging the separable structure of APUM and the hierarchical nature of product choices, we design a polynomial number of price experiments using a conditional gradient algorithm, and demonstrate that APUM can be calibrated with high accuracy using conversion data from these experiments. By employing piecewise linear approximation techniques, we further show that the resulting data-driven pricing problem can be efficiently solved as a mixed-integer program. The results highlight that firms can significantly enhance profitability by integrating machine learning and optimization techniques into their pricing decisions. By continuously learning from real-world data, this data-driven framework ensures that pricing strategies remain adaptive and aligned with

evolving market conditions. Our research provides a scalable solution for businesses seeking to optimize revenue in environments where customer choice patterns are complex and dynamic.

4 - Attribute-Based Pricing: A Nested-Logit Approach

Mengzhenyu Zhang, University College London, London, United Kingdom, Christopher Ryan, Wei Sun, Shivaram Subramanian, Markus Ettl

Attribute-based pricing---giving a price to potential product attributes individually and allowing customers to choose the attributes that form the final product---has been shown to improve customer satisfaction in the hospitality industry. In this paper, we consider the problem of finding optimal attribute prices to maximize the expected revenue from selling to a customer who chooses one product from a set of products that differ by only a few attributes. We model customer preferences using a nested logit model that can handle dependencies between products sharing common attributes but introduces tractability challenges. Because of complicated substitution effects among the final products that share common attributes, expected revenue is not concave in attribute prices. We provide two algorithms to solve the attribute pricing problem and show that they converge to a stationary point that provides a high-quality solution to the problem. {In contrast, popular approaches including gradient ascent and its variants, are not guaranteed to converge in this setting. Through numerical experiments, we show our algorithm is, on average, ten times faster than gradient-based methods, both in terms of runtime and number of iterations. We also extend our algorithm to a setting where attribute prices are constrained through linear inequalities and prove convergence to a stationary point. We implement our algorithm on a real hotel data set and demonstrate its revenue benefits.