



NUS
National University
of Singapore

Department of Industrial Systems
Engineering and Management
College of Design and Engineering

NUS

BUSINESS
SCHOOL

Department of
Analytics and
Operations



IORA

Institute of Operations
Research and Analytics



Operational
Research Society of
Singapore (ORSS)

ANALYTICS FOR 2026

NEXT-GEN SCHOLAR'S SYMPOSIUM

📅 24-25 SEPTEMBER 2026

📍 14.0, 3 RESEARCH LINK · SINGAPORE



Max Shen

The University of Hong Kong



Rene Caldentey

Chicago Booth



Sebastian Pokutta

TU Berlin



Maxim Afanasyev

Google



24-25 September, 2026



09:00 AM - 06:00 PM



NUS Innovation 4.0

More Information: <https://iora.nus.edu.sg/afx2026/>

FOREWORD

Dear Participants,

Welcome to the Analytics for X 2026 conference, a showcase of pioneering research and innovation in analytics. Hosted by the Institute of Operations Research and Analytics (IORA), along with the Department of Industrial Systems Engineering and Management (ISEM), the Department of Analytics and Operations (DAO) at the National University of Singapore and the Operations Research Society of Singapore (ORSS). This conference returns to Singapore on 24-25 September 2026, following a successful 2025 edition.

We continue the Next-Gen Scholars Symposium, which showcases outstanding PhD students in Analytics, Operations Research, and Operations Management. These scholars will share their groundbreaking research and engage with the broader analytics community.

In addition, we continue the ORSS Best Student Paper Competition, held in conjunction with the ORSS. This competition recognizes exemplary student research conducted in Singapore that demonstrates originality, methodological rigor, and significant theoretical or practical impact.

The conference features an impressive lineup of keynote speakers, including Professor Max Shen from the University of Hong Kong, Professor Sebastian Pokutta from TU Berlin, Professor Rene Caldentey from the University of Chicago, and Maxim Afanasyev from Google. Their insights set the stage for diverse presentations.

We extend our gratitude to all contributors and organizers whose efforts make this event possible. We hope the insights shared here will inspire continued exploration and collaboration in the global analytics community.

Thank you for joining us in this exciting journey.

Andrew Lim

Director, Institute of Operations
Research and Analytics

Ng Szu Hui

Associate Professor & Head
Department of Industrial Systems
Engineering and Management at
NUS College of Design and
Engineering

Jussi Keppo

Head and Provost's Chair Professor,
Department of Analytics &
Operations at NUS Business School
Research Director at Institute of
Operations Research and Analytics

TABLE OF CONTENTS

Table of Contents	1
Conference Schedule	2
Abstracts for Keynotes	5
Abstracts for Conference Papers	9
About the Organizers	24
Organizing Team	28
Getting Here	29

CONFERENCE SCHEDULE

DAY 1 - 24 SEPTEMBER 2026

8:30-8:50	Registration
8:50-9:00	Opening Address
9:00-9:30	Real-Time Linear Programming for Sponsored Listings Ranking: Evidence from a Large-Scale Field Experiment Zhu Yuting National University of Singapore
9:30-10:00	Hidden Convexity in Queueing Models Zhao Mingda Georgia Institute of Technology
10:00-10:30	Tea Break
10:30-11:00	The Distributional Behavior of Stochastic Algorithms Zhang Yixuan University of Wisconsin–Madison
11:00-11:30	Beyond Certification: Spatial Spillovers and Compliance Design in Indonesia's Palm Oil Sector Anton Reiling National University of Singapore
11:30-12:00	Conjectural Variations in Competitive Dynamic Pricing: A Learning Foundation via Experimentation Design and Feedback Structure Wang Wenyu National University of Singapore
12:00-13:15	Lunch Break
13:15-13:45	Service-Induced Congestion in Memory-Constrained LLM Serving Ao Ruicheng Massachusetts Institute of Technology
13:45-14:15	The Value of Advance Perishability Information in Dynamic Inventory Control Mao Luanxi National University of Singapore

14:15-15:15	Keynote Address: The Consumer Strikes Back: How Agentic AI Rewrites Market Competition Maxim Afanasyev Google Cloud & National University of Singapore
15:15-16:00	Coffee Break + Poster Session A
16:00-16:30	Maximum Load Assortment Optimization: Approximation Algorithms and Adaptivity Gaps Marouane Ibn Brahim Cornell University
16:30-17:30	Keynote Address: The Agentic Researcher: Agentic AI for Analytics and Operations Research Sebastian Pokutta TU Berlin & Zuse Institute Berlin (ZIB)
18:15-21:00	Dinner Reception (by invitation only) Kent Ridge Guild House 9 Kent Ridge Drive, Singapore 119241

DAY 2 - 25 SEPTEMBER 2026

9:00-10:00	Keynote Address: TBD Max Shen The University of Hong Kong
10:00-10:30	Tea Break
10:30-11:00	Speed of Intervention in Algorithmic Markets: Controlling Collusion and Stability Xie Tong The University of Chicago
11:00-11:30	Multi-Agent Queueing Network: Coordination with Learning Wei Zhaoxuan National University of Singapore
11:30-12:00	Satisficing Regret Minimization in Bandits: Constant Bound and Light-Tailed Distribution Feng Qing Cornell University & Cornell Tech
12:00-13:15	Lunch Break

13:15-14:15	Keynote Address: Information, Forecasting, and the Design of Supply Chain Policies Rene Cardentey The University of Chicago
14:15-14:45	Queueing Causal Models: Comparative Analytics in Queueing Systems Xu Zhenghang University of Toronto
14:45-15:15	AI-Generated Measurements for Identification and Inference with Missing Data: A Weak Shadow Variable Approach Chen Hongyu Massachusetts Institute of Technology
15:15-16:00	Coffee Break + Poster Session B
16:00-16:30	Learning to Rank under Strategic Corruption: A Nonatomic Approach Li Qinzhen National University of Singapore
16:30-17:00	Optimizing Service Operations via LLM-Powered Multi-Agent Simulation Wang Yanyuan The Hong Kong University of Science and Technology
17:00-17:15	Prize Distributions
17:15-17:30	Closing Address
18:15-22:00	Students' Night (by registration only)

KEYNOTE SPEAKERS

KEYNOTE ADDRESS

Maxim Afanasyev

Financial Services Industry Head for Asia Pacific and Japan
Google Cloud
Adjunct Associate Professor
National University of Singapore Business School



The Consumer Strikes Back: How Agentic AI Rewrites Market Competition

ABSTRACT

Traditional economics often models firms as strategic actors and consumers as relatively passive responders. Agentic AI may reverse this asymmetry. Unlike many earlier enterprise technologies, agentic AI has been adopted rapidly by consumers, while firms remain constrained by organizational inertia, legacy systems, compliance requirements, and slow process redesign. This creates a setting in which consumers may become strategically more sophisticated than the enterprises serving them.

We develop a model of this consumer first mover advantage and study how agentic AI increases consumers' ability to self serve, search, negotiate, substitute, and strategically bypass firm controlled processes. The model shows how this shift can reduce demand for some services and alter consumer behavior in ways that differ sharply from the pre agentic AI environment. These effects are particularly relevant in industries such as legal services, wealth advisory, education technology, e commerce, banking, and insurance.

The framework also helps explain several counterintuitive industry observations: AI may shift demand from services toward physical goods, create new financial stability risks, and reduce rather than increase consumers' available leisure time. The results suggest that firms should move beyond viewing AI primarily as a productivity tool and instead treat agentic AI as a source of market structure change, demand disruption, and strategic consumer empowerment.

KEYNOTE ADDRESS

Sebastian Pokutta

Professor, Optimization and Machine Learning
Mathematics and EECS (courtesy), TU Berlin
Vice President, Zuse Institute Berlin (ZIB)



The Agentic Researcher: Agentic AI for Analytics and Operations Research

ABSTRACT

Progress in analytics, operations research, and operations management rests on a steady stream of new models, algorithms, and validated insights. Yet the research pipeline that produces them remains a largely manual process, creating a bottleneck between the questions practitioners and scholars care about and the pace at which we can deliver solutions. In this talk, I will present a practical framework for AI-assisted research that aims to compress this cycle.

The framework turns general-purpose coding agents into autonomous research assistants through a five-level taxonomy of human-AI collaboration, a set of methodological commandments encoding scientific rigor, and a sandboxed environment supporting sessions from quick explorations to multi-day, multi-node experiments. I will illustrate the framework through two optimization case studies: establishing tight lower bounds for Frank-Wolfe methods on uniformly convex sets, where the agent discovered the proof strategy through systematic numerical exploration before formalizing it; and developing multi-variable dual tightening for mixed-integer convex optimization within the Frank-Wolfe-based solver Boscia, where the agent carried out the full cycle from theorem and proof to implementation and computational evaluation. Both examples demonstrate how agentic AI can accelerate the analytics research pipeline, compressing the cycle from idea to validated, deployable methodology.

KEYNOTE ADDRESS

Max Shen

Chair Professor in Logistics and Supply Chain Management
Director, Institute of Supply Chain Innovation
Director, Research Institute for Civil Low-Altitude Aviation
Senior Advisor to the President
Dean, Graduate School
The University of Hong Kong



...

KEYNOTE ADDRESS

Rene Caldentey

Eli B. and Harriet B. Williams Professor of
Operations Management
The University of Chicago



Matching under Congestion: Structure and Design in Bipartite Queueing Systems

ABSTRACT

Many service systems –from public housing and healthcare to flexible production and online matching markets– must match heterogeneous customers to heterogeneous servers while managing congestion. Simple first-come, first-served matching rules are attractive because they are transparent and easy to implement, yet the stochastic systems they generate can be surprisingly difficult to analyze.

In this talk, I will survey a line of work on bipartite matching and queueing systems that begins with the analysis of infinite FCFS matching and evolves toward the design of multiclass, multiserver service systems. A central theme is that heavy traffic can reveal a remarkably simple structure beneath an otherwise combinatorially complex system: the matching network decomposes into complete-resource-pooling components connected through a directed acyclic graph, providing sharp characterizations of waiting times and important invariance properties for matching probabilities. I will then discuss how these structural insights can be used to design compatibility networks and service menus that balance congestion and matching quality, including environments in which customers strategically choose among service options. Along the way, we will see that more flexibility—and even more customer choice—need not always improve performance. I will conclude with open questions on matching approximations, optimal flexibility and menu design, and extensions beyond the steady-state heavy-traffic setting.

ABSTRACTS FOR CONFERENCE PAPERS

DAY 1

Real-Time Linear Programming for Sponsored Listings Ranking: Evidence from a Large-Scale Field Experiment

Lu Haihao (MIT)

Zhang Luyang (Carnegie Mellon University)

Zhu Yuting (National University of Singapore)

Contact Email: Zhu Yuting (y.zhu@nus.edu.sg)

Abstract:

Sponsored product advertisements constitute a major revenue source for online marketplaces such as Amazon, Walmart, and Alibaba. A key operational challenge in these systems lies in the Sponsored Listings Ranking (SLR) problem, that is, determining which items to include and how to rank them to balance short-term revenue with long-term relevance and user experience. Industry practice predominantly relies on score-based algorithms, which construct heuristic composite scores to rank items efficiently within strict real-time latency constraints. However, such methods offer limited control over objective trade-offs and cannot readily accommodate additional operational constraints. We propose and evaluate a Linear Programming (LP)-based algorithm as a principled alternative to score-based approaches. We first formulate the SLR problem as a constrained mixed integer programming (MIP) model and develop a dual-based algorithm that approximately solves its LP relaxation within 0.1 second, satisfying production-level latency requirements. In collaboration with a leading online marketplace, we conduct a 19-day field experiment encompassing approximately 329 million impressions. The LP-based algorithm significantly outperforms the industry-standard benchmark in key marketplace metrics, demonstrating both higher revenue and maintained relevance. Mechanism analyses reveal that the performance gains are most pronounced when the revenue-relevance tradeoff is stronger. Our framework also generalizes to settings with inventory, sales, or fairness constraints, offering a flexible and deployable optimization paradigm. The LP-based algorithm was deployed in production at our partner marketplace in January 2023, marking a rare large-scale implementation of a mathematically grounded ranking algorithm in real-world online advertising.

Hidden Convexity in Queueing Models

Xin Chen (Georgia Tech)

Linwei Xin (Cornell University)

Minda Zhao (Georgia Tech)

Contact Email: Minda Zhao(mindazhao@gatech.edu)

Abstract:

We study the joint control of arrival and service rates in queueing systems with the objective of minimizing long-run expected cost minus revenue. Although the objective function is non-convex, first-order methods have been empirically observed to converge to globally optimal solutions. This paper provides a theoretical foundation for this empirical phenomenon by characterizing the optimization landscape and identifying a hidden convexity: the problem admits a convex reformulation after an appropriate change of variables. Leveraging this hidden convexity, we establish the Polyak-Łojasiewicz-Kurdyka (PŁK) condition for the original control problem, which excludes spurious local minima and supports global convergence guarantees for first-order methods. Our analysis applies to a broad class of $GI/GI/1$ queueing models, including those with Gamma-distributed interarrival and service times, as well as $GI/M/1$ queues with log-concave interarrival times. As a key ingredient in the proof, we establish a new convexity property of the expected queue length under a square-root transformation of the traffic intensity.

The Distributional Behavior of Stochastic Algorithms

Zhang Yixuan (University of Wisconsin–Madison)

Qiaomin Xie (University of Wisconsin–Madison)

Contact Email: Zhang Yixuan(yzhang2554@wisc.edu)

Abstract:

Stochastic algorithms are fundamental tools in modern machine learning, optimization, and control. Traditionally, they have often been viewed as noisy counterparts of deterministic algorithms, with stochasticity treated primarily as a nuisance to be controlled. Accordingly, classical theory largely focuses on controlling stochastic fluctuations through algorithm design (e.g., the use of diminishing stepsizes) and analysis (e.g., bounds on mean-squared error).

In a recent line of our work, we depart from this paradigm and instead develop a probabilistic framework for obtaining a finer-grained understanding of the distributional behavior of stochastic algorithms by viewing their iterates as a stochastic process. Compared with classical error-bound analyses, this perspective characterizes not only the magnitude of the error, but also the structure of the fluctuations, the emergence of bias, and the distinction between transient and long-run behavior. Such information can in turn be exploited to improve algorithm design and enable principled statistical inference.

In this talk, I will present several of our recent results in this direction. We establish finite-time distributional convergence of the iterates, characterize their steady-state fluctuations and bias, and leverage these characterizations for stepsize selection, variance reduction, bias correction, and statistical inference. Our primary focus will be on the constant-stepsize regime, which is widely used in practice. I will conclude with applications to asynchronous Q-learning.

Beyond Certification: Spatial Spillovers and Compliance Design in Indonesia's Palm Oil Sector

Dan A. Iancu(Stanford University)

Jussi Keppo(National University of Singapore)

Bar Light(National University of Singapore)

Anton Reiling (National University of Singapore)

Contact Email: Anton Reiling (anton.reiling@u.nus.edu)

Abstract:

Eco-certification schemes like the Roundtable on Sustainable Palm Oil (RSPO) aim to reduce deforestation in commodity supply chains. Their system-wide effectiveness, however, may be limited when conservation gains in certified areas are offset by deforestation in nearby or economically connected locations. This work analyzes the effectiveness of certifications granted by RSPO in Indonesian Sumatra and Kalimantan between 2001 and 2021. We find that, while RSPO certification successfully reduces deforestation of primary forest on certified land, there is a significant spillover effect leading to increased deforestation in adjacent areas. We further find that certification decisions are driven by stark heterogeneity: while the intention of RSPO certification is to reduce deforestation, most certification happens on already developed land with little remaining forest, attenuating the marginal effect of certification. Further, spillover effects are stronger in forest-frontier areas. Since reduced-form estimates suffer from SUTVA violations and cannot simulate counterfactual regulatory regimes, we develop a structural model, where palm oil mills make sequential decisions of land expansion to balance clearing costs and expected revenue. We estimate the latent parameters of this model and find them coherent with our previous findings. In a model-based counterfactual analysis, we subsequently estimate the net effect of RSPO certifications and find that past RSPO practices have led to significantly more primary forest cover loss compared to a no-certification counterfactual scenario.

Conjectural Variations in Competitive Dynamic Pricing: A Learning Foundation via Experimentation Design and Feedback Structure

Bar Light (National University of Singapore)

Wenyu Wang (National University of Singapore)

Contact email: Wang Wenyu (wang.wenyu@u.nus.edu)

Abstract:

We study competitive dynamic pricing among multiple sellers, motivated by the rise of large-scale experimentation and algorithmic pricing in retail and online marketplaces. Sellers repeatedly set prices using simple learning rules and observe their own realized demand, while possibly observing only a subset of rivals' prices, even though demand depends on all sellers' prices and is subject to random shocks. Each seller runs local price experiments, such as switchback-style designs, and updates a focal price using a linear demand estimate fitted to its own demand data and the competitor prices it observes. Under certain conditions on demand, the resulting dynamics converge to a Conjectural Variations (CV) equilibrium, a classic static equilibrium notion in which each seller best responds under a conjecture that rivals' prices co-move systematically to changes in its own price. Unlike standard CV models that treat conjectures as behavioral primitives, we show that these conjectures arise endogenously from the interaction between the feedback structure and the correlation structure of experimentation. When a seller does not observe some rivals' prices, correlated experimentation induces an omitted-variable bias in demand estimation. We show that this bias determines the conjectures that govern the long-run equilibrium. Notably, when this learning bias vanishes, for example under full price feedback or independent experimentation of unobserved rivals, the learning dynamics converge to the standard Nash equilibrium. We provide simple sufficient conditions on demand for convergence in standard models and establish a finite-sample guarantee, showing that the mean squared price error decays at a rate of $\mathcal{O}(T^{-1/2})$.

Service-Induced Congestion in Memory-Constrained LLM Serving

Ruicheng Ao (Massachusetts Institute of Technology)

Jing Dong (Columbia University)

Gan Luo (Peking University)

David Simchi-Levi (Massachusetts Institute of Technology)

Contact email: Ao Ruicheng (aorc@mit.edu)

Abstract:

In large language model (LLM) serving, each request accumulates persistent graphics processing unit (GPU) memory during service as its key-value cache grows with every generated token. Under high concurrency, aggregate memory usage therefore increases endogenously over time: the service process itself creates future capacity pressure. When memory capacity is exceeded, systems evict active requests, discarding cached state and restarting them later, which wastes computation and reduces throughput. We develop a discrete-time dynamical model of memory-constrained LLM inference that captures admission, memory growth, and eviction under continuous batching. In the saturated-input regime, the system admits both eviction-free fixed points and limit cycles with evictions. For homogeneous workloads, we show that the eviction-free equilibrium is unstable and that, except for a Lebesgue-measure-zero exact-capture set, the system converges to a unique worst-case limit cycle that is asymptotically stable outside this exceptional set, with throughput losses as large as 50%. For heterogeneous workloads, we prove a stability criterion in the two-class common-input setting and explain how the survival-polynomial mechanism generalizes to multiple classes and heterogeneous-input lengths. Under an input-dominated scaling regime, coprime decoding lengths stabilize the eviction-free equilibrium, while non-coprime lengths create synchronized modes that drive instability. These results characterize when workload heterogeneity desynchronizes completions and helps stabilize memory-constrained serving. More broadly, we identify service-induced congestion as a structural instability mechanism and derive scheduling design principles for sustaining high throughput.

The Value of Advance Perishability Information in Dynamic Inventory Control

Luanxi Mao (National University of Singapore)

Mabel Chou (National University of Singapore)

Xiaobo Li (National University of Singapore)

Contact email: Mao Luanxi (luanxi@u.nus.edu)

Abstract:

Fresh perishables often have uncertain remaining shelf lives, making it difficult for grocery retailers to know how much of their on-hand inventory will spoil. Advances in sensing technologies offer a way to reduce this uncertainty by providing perishability information about inventory condition before replenishment. In this paper, we study a finite-horizon periodic-review inventory model under the baseline system and the advance perishability information (API) system. The baseline retailer places replenishment orders under perishability uncertainty and follows a state-dependent ordering policy. With information, the API retailer observes the realized perishability before ordering, and the optimal policy is a standard base-stock policy. We further quantify the benefit of API through the value of advance perishability information (VAPI) and analytically characterize its structural properties. We also find that although API reduces expected inventory cost, it does not necessarily reduce food waste. Overall, the waste reduction benefit of API is stronger when demand is more predictable and perishability is more variable.

Maximum Load Assortment Optimization: Approximation Algorithms and Adaptivity Gaps

Omar El Housni (Cornell University)

Marouane Ibn Brahim (Cornell University)

Danny Segev (Tel Aviv University)

Contact email: Marouane Ibn Brahim (mi262@cornell.edu)

Abstract:

Motivated by modern-day applications such as attended home delivery and preference-based group scheduling, where decision makers wish to steer a large number of customers toward choosing the exact same alternative, we introduce a novel class of assortment optimization problems, referred to as maximum load assortment optimization. In such settings, given a universe of substitutable products, we are facing a stream of customers, each choosing between either selecting a product out of an offered assortment or opting to leave without making a selection. Assuming that these decisions are governed by the multinomial logit choice model, we define the random load of any underlying product as the total number of customers who select it. Our objective is to offer an assortment of products to each customer so that the expected maximum load across all products is maximized. We consider both static and dynamic formulations of the maximum load assortment optimization problem. In the static setting, a single offer set is carried throughout the entire process of customer arrivals, whereas in the dynamic setting, the decision maker offers a personalized assortment to each customer, based on the entire information available at that time. As can only be expected, both formulations present a wide range of computational challenges and analytical questions. The main contribution of this paper resides in proposing efficient algorithmic approaches for computing near-optimal static and dynamic assortment policies. In particular, we develop a polynomial time approximation scheme for the static problem formulation. Additionally, we demonstrate that an elegant policy utilizing weight-ordered assortments yields a $1/2$ approximation. Concurrently, we prove that such policies are sufficiently strong to provide a $1/4$ approximation with respect to the dynamic formulation, establishing a constant factor bound on its adaptivity gap. Finally, we design an adaptive policy whose expected maximum load is within factor $1-\epsilon$ of optimal, admitting a quasi-polynomial time implementation.

DAY 2

Speed of Intervention in Algorithmic Markets: Controlling Collusion and Stability

Tong Xie (University of Chicago, Chicago)

Rene Caldentey (University of Chicago, Chicago)

Martin Haugh (Imperial Business School, Imperial College London)

Contact email: Xie Tong (tong.xie@chicagobooth.edu)

Abstract:

Algorithmic pricing is prevalent in online marketplaces, raising concerns that learning algorithms may sustain supra-competitive outcomes. We study two sellers in a repeated Prisoner's Dilemma who delegate decisions to memoryless value-based learning algorithms. A platform steers demand toward sellers choosing the more competitive action. Using a fluid approximation, we characterize how intervention intensity, learning architecture, and implementation speed determine market dynamics. Stronger intervention lowers cooperation levels at the fixed point. Algorithmic responsiveness determines the location of stationary outcomes, while action-dependent learning rates determine their local stability. As intervention intensifies, the fixed point can lose stability through a Hopf bifurcation and its local basin of attraction shrinks as the stability threshold is approached. Consequently, an abrupt intervention may leave inherited Q -values outside the target's local basin of attraction and lead to persistent oscillations, whereas gradual intervention leads to a stable market outcome. We also derive local bounds on implementation speed and rollout time. Our results show that platform regulation of algorithmic collusion may require choosing not only how much to intervene, but also how quickly to implement the intervention relative to the algorithms' learning dynamics.

Multi-Agent Queueing Network: Coordination with Learning

Andrew E. B. Lim (National University of Singapore)

ZhaoXuan Wei (National University of Singapore)

Hanqin Zhang (National University of Singapore)

Contact Email: Wei Zhaoxuan (zhaoxuanwei@u.nus.edu)

Abstract:

We study decentralized control of a Jackson network consisting of a pricing agent who controls the exogenous arrival process, and service agents who control the service rate at each station. The system operates under the collection of decentralized control policies determined within informational silos. The control policy of each agent is obtained by optimizing a stochastic model of her own operations that accurately captures her impact on the system but may misspecify the impact of other agents. We show that agents can be coordinated with a transfer price mechanism to achieve optimal centralized efficiency, even when local models are misspecified, and propose a multi-agent reinforcement learning (MARL) algorithm for computing the optimal transfers in a fully decentralized fashion. Our RL algorithm is shown to converge to the optimal transfer mechanism and the sample complexity for computing the ϵ -optimal transfer is given.

Satisficing Regret Minimization in Bandits: Constant Bound and Light-Tailed Distribution

Qing Feng (Cornell University)

Tianyi Ma (Cornell University)

Ruihao Zhu (Cornell University)

Contact Email: Feng Qing (qf48@cornell.edu)

Abstract:

Motivated by the concept of satisficing in decision-making, we consider the problem of satisficing regret minimization in bandit optimization. In this setting, the learner aims at selecting satisficing arms (arms with mean reward exceeding a certain threshold value) as frequently as possible. 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 expected satisficing regret for a wide variety of bandit optimization problems in the realizable case (i.e., a satisficing arm exists). 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 on both synthetic datasets and a real-world dynamic pricing case study.

Queueing Causal Models: Comparative Analytics in Queueing Systems

Opher Baron (University of Toronto)

Dmitry Krass Rotman (University of Toronto)

Arik Senderovich (York University)

Mark van der Laan (University of California, Berkeley)

Zhenghang Xu (University of Toronto)

Contact Email: Xu Zhenghang(zhenghang.xu@rotman.utoronto.ca)

Abstract:

How can we evaluate alternative queueing policies from historical data without fully specifying the underlying queueing model? We develop Structural Causal Queueing Models (SCQM), a data-driven framework combining causal modeling, machine learning, and simulation. SCQM represents queueing dynamics through causal dependencies and structural assignment functions. Using event order and a sufficient history bound, we construct default parent sets that provide a starting point for learning these mechanisms from observational data. The estimated model then supports policy evaluation by replacing the relevant policy mechanism and simulating the resulting queueing trajectories. Numerical experiments demonstrate accuracy comparable to queueing-theoretic benchmarks, including settings where analytical solutions are available. SCQM thus complements queueing theory by enabling data-driven comparative analytics when the underlying system dynamics are not fully known.

AI-Generated Evaluations for Identification and Inference with Missing Data: A Weak Shadow Variable Approach

Hongyu Chen (Massachusetts Institute of Technology)

David Simchi-Levi (Massachusetts Institute of Technology)

Ruoxuan Xiong (Emory University)

Contact Email: Chen Hongyu (chenhy@mit.edu)

Abstract:

Across business and social science applications, outcomes are often missing in ways that depend on the unobserved outcomes themselves. In service systems, for example, whether a customer submits a rating depends on the rating they would have provided. Such missing-not-at-random (MNAR) mechanisms make population quantities difficult to identify without strong assumptions on the observation process. Meanwhile, rich unstructured data, such as customer interaction histories, are increasingly available and can be used to construct structured measurements using tools such as large language models (LLMs). In this work, we develop an assumption-lean partial identification framework that uses such measurements as weak shadow variables, defined as outcome-informative proxies that are conditionally independent of missingness given the true outcome and observed covariates. Importantly, they need not accurately predict missing outcomes or satisfy the completeness requirement in the classical shadow variable literature. For identification, we characterize sharp bounds on population quantities through a pair of linear programs. For estimation and inference, we propose a localized penalized estimator that remains feasible under sampling error, and a subsampling algorithm for constructing confidence intervals. In semi-synthetic experiments using real customer-service dialogues, weak-shadow-variable intervals are about 89% narrower than those without auxiliary information, while their midpoints have around 41% lower estimation error than classical MNAR methods.

Learning to Rank under Strategic Corruption: A Nonatomic Approach

Qinzhen Li (National University of Singapore)

Yifan Feng (National University of Singapore)

Hongfan (Kevin) Chen (The Chinese University of Hong Kong)

Contact Email: Li Qinzhen (qinzhen@u.nus.edu)

Abstract:

We study a dynamic learning-to-rank problem in which N agents strategically inflate the reward signals that are used to rank them, with corruption budgets chosen endogenously. The learner observes noisy manipulated outcomes and seeks to maximize total true reward over T periods. Under a general supermodular ranking-reward model, we show that efficient learning remains achievable despite strategic manipulation. Specifically, a simple Experiment-Then-Commit policy attains regret bounded in T , under a strategy profile that forms an $o(1)$ -equilibrium among agents as N and T grow.

Our analysis proceeds in three parts. First, we show that agents' strategic interactions converge to a static nonatomic game. Second, we find a nonatomic equilibrium in this limit, where agents' manipulation intensity increases with their types. Third, we transport the limit equilibrium to the finite-agent, finite-horizon setting, which yields non-asymptotic approximate-equilibrium and regret guarantees.

Optimizing Service Operations via LLM-Powered Multi-Agent Simulation

Yanyuan Wang (The Hong Kong University of Science and Technology)

Xiaowei Zhang (The Hong Kong University of Science and Technology)

Contact email: Wang Yanyuan (yanyuan.wang@connect.ust.hk)

Abstract:

Service system performance depends on how participants respond to design choices, but modeling these responses is hard due to the complexity of human behavior. We introduce an LLM-powered multi-agent simulation (LLM-MAS) framework for optimizing service operations. We cast the problem as stochastic optimization with decision-dependent uncertainty: design choices are embedded in prompts and shape the distribution of outcomes from interacting LLM-powered agents. By embedding key numerical information in prompts and extracting it from LLM-generated text, we model this uncertainty as a controlled Markov chain. We develop an on-trajectory learning algorithm that, on a single simulation run, simultaneously constructs zeroth-order gradient estimates and updates design parameters to optimize steady-state performance. We also incorporate variance reduction techniques. In a sustainable supply chain application, our method outperforms several benchmarks, including black-box optimization and using LLMs as numerical solvers or as role-playing system designers. A case study on optimal contest design with real behavioral data shows that LLM-MAS is both a cost-effective evaluator of known designs and an exploratory tool that can uncover viable designs overlooked by traditional approaches.

ABOUT THE ORGANIZERS

Department of Analytics and Operations, NUS Business School

The **Department of Analytics & Operations (DAO)** produces cutting-edge research and disseminates knowledge in Business Analytics and Operations Management that impact academia, industry and society. DAO is well-known internationally as a research centre of excellence in supply chain and logistics. It has been responsible for the quantitative dimension of business school education, particularly in Analytics, Operations Research and Supply Chain Management.

The department attracts top students from around the world, who are taught by its world-leading faculty. The department's rigorous, industry-relevant, and impactful research is also infused into the education of future generations.

Proudly fostering a robust research environment, DAO supports and organizes international conferences, with its seminar series regularly drawing academics and researchers from across NUS and other Singapore universities. In its efforts to engage the wider community, DAO is actively involved with three NUS cross-faculty research institutes and centers: the Business Analytics Centre, The Logistics Institute–Asia Pacific, and the Institute for Operations Research and Analytics.

DAO faculty members frequently publish in top journals in Analytics and Operations, including Operations Research, Management Science, MSOM, and Production and Operations Management. Several faculty members also serve as editors for these prestigious journals.



Department of Industrial Systems Engineering & Management, NUS College of Design and Engineering

The **Industrial Systems Engineering and Management (ISEM)** Department has come a long way since its formation in the early 1970s under the sponsorship of the Ford Foundation. Its MSc program (the first in Engineering in Singapore) has been the flagship program of the department for the last 30 years and is still vital and relevant to Singapore's economy. Its BEng (ISE) graduates are much sought after by industries and the program has been accredited by the Engineering Accreditation Board (EAB).

Moving in tandem with the national thrust in transforming NUS to a leading research-based university, the research program in the ISEM Department has also received world recognition in the area of Quality and Reliability Engineering with its faculty members represented in editorial boards of many international journals in this field. Research in Logistics has been well-funded and is gaining momentum and international visibility.

The ISEM Department's vision is to be the department of choice, which, in turn, will put NUS on the world map as the university of choice. While the department will continue to play a role in reinforcing Singapore's key areas of excellence in ports, aviation, infra-structure and high-tech industries, we are also conducting research inspired by regional/global needs. Aligned with the vision of producing engineering leaders, we hope that our alumni will make a difference in our community as well as in the international arena.



Institute of Operations Research and Analytics

Established in November 2016, the **Institute of Operations Research and Analytics (IORA)** conducts cutting-edge research on the optimization, analysis and management of service systems, including model formulation, algorithm design, analysis of service strategies, and software development. As part of the NUS's Smart Nation Research Cluster, IORA works with various partners to develop innovative measures and inventive solutions for real-world issues.

Since its inception, IORA has been steadfast in its core missions:

- To develop world-class research programs to modernize the practice of Operations Research in this new data-intensive environment.
- To modernize the teaching and nurture a new generation of PhD students, well versed in the tools and theories in the integration of Data Analytics into Operations Research.
- To leverage on the combined expertise of NUS faculty in related areas to establish a centre of excellence in the field of Operations Research and Analytics.

The IORA faculty is made up of strong Operations Research and Analytics expertise housed in various departments including computer science, economics, engineering, mathematics, statistics, analytics and operations. Taking a multidisciplinary approach to research, our students learn from faculty members at the forefront of their respective fields. Graduates are trained to be conversant in the new science of data driven analytics, with expertise in Operations Research, Machine Learning and Computational Techniques, capable of discovering new solutions and models to make smart inference from high dimensional and voluminous data.




Operational Research Society of Singapore

Operational Research Society of Singapore (ORSS) was registered as a professional association in Singapore on 21st November 1975. It was formed in recognition of the need for such an association to promote the field of Operational Research or Operations Research (OR) in Singapore.

As we are well aware, OR has had profound impacts addressing complex, real-world problems by employing advanced analytical methods. This critical field influences a diverse range of industries, including logistics, supply chain management, healthcare, finance, and public policy, and many others. Over the almost 5 decades since its formation, ORSS has played a vital role in promoting OR and related fields in Singapore. It offers a platform for industry professionals, academics, and students to explore and engage with the latest OR methodologies and applications. As a central hub, it disseminates innovative research findings and best practices, driving the field's growth in Singapore.

Today, ORSS continues to foster knowledge exchange and collaboration, ensuring that Operations Research remains a driving force behind Singapore's efficiency, resilience, and innovation.

ORGANIZING TEAM



Rajiv K

Sr. Research Fellow
IORA



Mao Luanxi

Student
IORA



Cheung Wang Chi

Assistant Professor
ISEM



Yuan Yuxin

Student
ISEM



Qin Hanzhang

Assistant Professor
ISEM



Nigel Soh

Student
ISEM



Zhang Guoyun


Student
ISEM



Zhao Long

Assistant Professor
DAO

GETTING THERE



To get to Innovation 4.0, NUS, scan this code or visit <https://bit.ly/4cciXHL> for directions.



By MRT/Bus:

- Kent Ridge MRT Station (Circle Line)
- Clementi MRT Station (East-West Line)
 - Transfer to Service 96 at the Clementi Bus Interchange.
- Buona Vista MRT Station (East-West Line, Circle Line)
 - Transfer to Service 95 at Bus Stop (3) outside Exit D/across North Buona Vista Road, in front of the Ministry of Education building).

Once in the campus, take the campus shuttle and alight at TCOMS, Opp TCOMS or COM3 and walk to Innovation 4.0.

By Taxi/Private Hire:

- Input Innovation 4.0, NUS into your GPS or ride-hailing app.

By Car:

- Note that parking during term time might be limited, so we recommend alternative modes of travel. Consult the NUS mobility services for latest parking info.

Notes