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

ANALYTICS FOR X 2024

NEXT-GEN SCHOLAR'S SYMPOSIUM

28 - 29 August, 2024

A conference promoting the enhancement and dissemination of knowledge related to Operations Research, Operations Management and Analytics, as well as efficient industrial practices.

DAVID SIMCHI-LEVI
Massachusetts Institute of Technology

MING HU
University of Toronto

YING-JU CHEN
Hong Kong University of Science and Technology

28 - 29 August, 2024

09.00 AM - 06.00 PM

NUS Innovation 4.0

Register here

More information: https://iora.nus.edu.sg/afx2024/

FOREWORD

Dear Participants,

Welcome to the Analytics for X 2024 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 and the Department of Analytics and Operations at the National University of Singapore, this conference returns to Singapore on 28-29 August 2024, following a successful event in Chongqing, China.

This year's event introduces the Next-Gen Scholars Symposium, celebrating outstanding PhD students in Analytics, Operations Research, and Operations Management. These scholars will present their groundbreaking research and engage with the broader analytics community.

The conference features an impressive lineup of keynote speakers, including Professor Ming Hu from the University of Toronto, Professor David Simchi-Levi from MIT, and Professor Ying-Ju Chen from Hong Kong University of Science and Technology. Their insights set the stage for diverse presentations on machine learning, optimization, differential privacy, and more.

Key sessions include discussions on data-driven operations management, robust optimization, and innovations in pricing and resource allocation. These contributions reflect the conference's commitment to advancing knowledge in analytics and industrial practices.

Following the conference, the IORA Industry Day on 30 August 2024 offers additional opportunities to network with local industries.

We extend our gratitude to all contributors and organizers whose efforts have made this event possible. We hope the insights shared here inspire ongoing exploration and collaboration in the analytics landscape.

Thank you for joining us in this exciting journey.

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

Ng Szu Hui

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

Teo Chung Piaw

Executive Director, Institute of Operations Research and Analytics and Provost's Chair Professor, Department of Analytics and Operations

TABLE OF CONTENTS

CONFERENCE SCHEDULE

DAY 1 - 28 AUGUST 2024

8:30-9:00	Registration
8:55-9:00	Opening Address
9:00-10:00	**Keynote Address: Big, Small, and Small+ Data-Driven OM** Ming Hu Professor Joseph L. Rotman School of Management, University of Toronto
10:00-10:30	Tea Break
10:30-11:00	**When Should you Offer an Upgrade: Online Upgrading Mechanisms for Resource Allocation** Zijie Zhou Massachusetts Institute of Technology
11:00-11:30	**Machine Learning-Augmented Optimization of Large Bilevel and Two-Stage Stochastic Programs: Application to Cycling Network Design** Bo Lin University of Toronto
11:30-12:00	**Differential Privacy via Distributionally Robust Optimization** Aras Selvi Imperial College London
12:00-13:15	Lunch Break
13:15-13:45	**LEGO: Optimal Online Learning under Sequential Price Competition** Shukai Li Northwestern University
13:45-14:15	**Statistical Properties of Robust Satisficing** Zhiyi Li Hong Kong University of Science and Technology
14:15-14:45	**Sinkhorn Distributionally Robust Optimization** Jie Wang Georgia Institute of Technology
14:45-15:15	Tea Break

15:15-15:45	**Randomized Assortment Optimization** Zhengchao Wang Imperial College London
15:45-16:15	**A Behavioral Model for Multi-armed Bandits: Theoretical Framework and Experimental Evidence** Jingying Ding National University of Singapore
16:15-17:15	**Keynote Address: Postgraduate program applications: simultaneous search, sequential outcomes, and reservation fees** Ying-Ju Chen Hong Kong University of Science and Technology
18:30-21:00	**Dinner Reception (by invitation only)** The Scholar Chinese Restaurant Kent Ridge Guild House 9 Kent Ridge Drive, Singapore 119241

DAY 2 - 29 AUGUST 2024

9:00-10:00	**Keynote Address: Reinventing Operations Management's Research and Practice with Data Science** David Simchi-Levi Professor of Civil and Environmental Engineering; Director, MIT Data Science Lab, Massachusetts Institute of Technology
10:00-10:30	Tea Break
10:30-11:00	**Offline Pricing with Shape-Restricted Demands** Jingren Liu National University of Singapore
11:00-11:30	**Bayesian Online Multiple Testing: A Resource Allocation Approach** Feng Zhu Massachusetts Institute of Technology
11:30-12:00	**Optimal Performance Feedback in Gamified Services** Lin Chen INSEAD
12:00-13:15	Lunch Break

13:15-13:45	**Cross- vs. In-Region Courier Routing in On-Demand Delivery** Arseniy Gorbushin University Of Toronto, Rotman School of Management
13:45-14:15	**Optimal probability measure decomposition** Xin Tong National University of Singapore
14:15-14:45	**Unraveling the Implications of Silent Labor Time (SLT) in the Gig Economy** Srishti Arora INSEAD
14:45-15:15	Tea Break
15:15-15:45	**Robust Auction Design with Support Information** Jerry Anunrojwong Columbia Business School
15:45-16:15	**Neural-Network Mixed Logit Choice Model: Statistical and Optimality Guarantees** Zhi Wang University of Texas at Austin
16:15-16:45	**A Nonparametric Approach with Marginals for Modeling Consumer Choice** Yanqiu Ruan Singapore University of Technology and Design
16:45-17:00	Closing Address

KEYNOTE SPEAKERS

KEYNOTE ADDRESS

David Simchi-Levi
Professor of Engineering Systems
Director, MIT Data Science Lab
Massachusetts Institute of Technology

Reinventing Operations Management's Research and Practice with Data Science

ABSTRACT

In this talk we show how data-driven research fosters the development of new engineering and scientific methods that explain, predict, and change behavior. We report on a few projects with online and brick-and-mortar retailers where we combine machine learning, optimization and econometrics techniques to improve business performance.

KEYNOTE ADDRESS

Ming Hu
Distinguished Professor of Business Operations and Analytics
University of Toronto

Big, Small, and Small+ Data-Driven OM

ABSTRACT

We discuss three data-driven decision-making scenarios in classical operations management (inventory or pricing) settings, using big or small data or small data but with one chance of experimentation. First, with big data, we study a contextual-based newsvendor problem using deep neural networks (DNN). Empirical process theory is pivotal in ensuring that the asymptotic behavior of observed data converges to the true underlying distribution as the sample size increases. We provide theoretical guarantees in terms of excess risk bounds for the DNN solution, characterized by the network structure and sample size, validating the applicability of DNNs in relevant OM contexts. These excess risk bounds exhibit polynomial growth in the feature dimension and attain the minimax convergence rate (with respect to the sample size) in expectation. Second, with small data, traditional frequentist methods may be ineffective, and we propose using the empirical Bayes (EB) method to achieve transfer learning in estimating unknown parameters using data across many products and subsequently making decisions based on these estimates. We illustrate this approach with a multi-product pricing problem, employing a hierarchical feature-based demand model and a nonparametric maximum likelihood method to derive the prior from the data. The effectiveness of the EB method is demonstrated by characterizing the regret bound using an oracle benchmark that presumes prior knowledge of the underlying distribution. Third, we study the benefits of one-shot price experimentation for a seller in setting a price who only knows the exact purchase probability associated with a single historical price and aims to maximize the worst-case revenue ratio compared to an oracle with complete knowledge of the value distribution. We analytically characterize the optimal distributionally robust experimental and final price points, obtain their tight performance guarantee for any historical purchase probability, and then evaluate the value of experimentation, which exhibits a two-modal behavior with respect to the historical purchase probability.

KEYNOTE ADDRESS

Ying-Ju Chen
Crown Worldwide Professor of Business
Hong Kong University of Science and Technology

Postgraduate program applications: simultaneous search, sequential outcomes, and reservation fees

ABSTRACT

This paper studies a simultaneous-search problem in which a player observes the outcomes sequentially, and must pay reservation fees to maintain eligibility for recalling the earlier offers. We use postgraduate program applications to illustrate the key ingredients of this family of problems. We develop a parsimonious model with two categories of schools: reach schools, which the player feels very happy upon joining, but the chance of getting into one is low; and safety schools, which are a safer choice but not as exciting. The player first decides on the application portfolio, and then the outcomes from the schools applied to arrive randomly over time. We start with the extreme case wherein the safety schools always admit the player. We show that it suffices to focus on the last safety school, which allows us to conveniently represent the player's value function by a product form of the probability of entering the last safety period and the expected payoff from then on.

We show that the player's payoff after applications is increasing and discrete concave in both the numbers of reach and safety schools. We also develop a recursive dynamic programming algorithm when admissions to safety schools are no longer guaranteed. We demonstrate instances in which the player applies to more safety schools when either the reservation fee gets higher or the admission probability drops lower, and articulate how these arise from the portfolio optimization consideration. This has strong managerial implications for service providers in devising their reservation fees and admission rates, especially for institutions that are not universally favored by prospective applicants.

ABSTRACTS FOR CONFERENCE PAPERS

DAY 1

When Should you Offer an Upgrade: Online Upgrading Mechanisms for Resource Allocation

Patrick Jaillet
(Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology)

Chara Podimata
(Sloan School of Management, Massachusetts Institute of Technology)

Andrew Vakhutinsky
(Oracle Lab)

Zijie Zhou
(Operations Research Center, Massachusetts Institute of Technology)

Contact Email: Zijie Zhou (zhou98@mit.edu)

Abstract:
In this work, we study an upgrading scheme for online resource allocation problems. We work in a sequential setting, where at each round a request for a resource arrives and the decision-maker has to decide whether to accept it (and thus, offer the resource) or reject it. The resources are ordered in terms of their value. If the decision-maker decides to accept the request, they can offer an *upgrade-for-a-fee* to the next more valuable resource. This fee is dynamically decided based on the currently available resources. After the upgrade-fora-fee option is presented to the requester, they can either accept it, get upgraded, and pay the additional fee, or reject it and maintain their originally allocated resource.

We take the perspective of the decision-maker and wish to design upgrading mechanisms in a way that simultaneously maximizes revenue and minimizes underutilization of resources. Both of these desiderata are encapsulated in a notion of *regret* that we define, and according to which we measure our algorithms' performance. We present a fast algorithm that achieves *O(logT)* regret.

Finally, we implemented our algorithm utilizing data from a real hotel and estimated our upgrading mechanism would increase the annual revenue by over 17%.

Machine Learning-Augmented Optimization of Large Bilevel and Two-Stage Stochastic Programs: Application to Cycling Network Design

Timothy C. Y. Chan
(University of Toronto)

Bo Lin
(University of Toronto)

Shoshanna Saxe
(University of Toronto)

Contact Email: Bo Lin (imbo.lin@mail.utoronto.ca)

Abstract:
Motivated by a cycling infrastructure planning application, we present a machine learning approach to solving bilevel programs with a large number of independent followers, which as a special case includes two-stage stochastic programming. We propose an optimization model that explicitly considers a sampled subset of followers and exploits a machine learning model to estimate the objective values of unsampled followers. Unlike existing approaches, we embed machine learning model training into the optimization problem, which allows us to employ follower features that cannot be represented using leader decisions. We prove bounds on the optimality gap of the generated leader decision as measured by the original objective that considers the full follower set. We develop follower sampling algorithms to tighten the bounds and a representation learning approach to learn follower features, which are used as inputs to our machine learning model. Through numerical studies, we show that our approach generates leader decisions of higher quality compared to baselines. Finally, we perform a real-world case study in Toronto, Canada, where we solve a cycling network design problem with over one million followers. Compared to the current practice, our approach improves a transportation metric by 19.2% and can lead to a potential cost saving of $18M.

Differential Privacy via Distributionally Robust Optimization

Aras Selvi
(Imperial College Business School)

Huikang Liu
(Research Institute for Interdisciplinary Sciences, School of Information Management and Engineering, Shanghai University of Finance and Economics)

Wolfram Wiesemann
(Imperial College Business School)

Contact Email: Aras Selvi (a.selvi19@imperial.ac.uk)

Abstract:
In recent years, differential privacy has emerged as the *de facto* standard for sharing statistics of datasets while limiting the disclosure of private information about the involved individuals. This is achieved by randomly perturbing the statistics to be published, which in turn leads to a privacy-accuracy trade-off: larger perturbations provide stronger privacy guarantees, but they result in less accurate statistics that offer lower utility to the recipients. Of particular interest are therefore optimal mechanisms that provide the highest accuracy for a pre-selected level of privacy. To date, work in this area has focused on specifying families of perturbations a priori and subsequently proving their asymptotic and/or best-in-class optimality.

In this paper, we develop a class of mechanisms that enjoy non-asymptotic and unconditional optimality guarantees. To this end, we formulate the mechanism design problem as an infinite-dimensional distributionally robust optimization problem. We show that the problem affords a strong dual, and we exploit this duality to develop converging hierarchies of finite-dimensional upper and lower bounding problems. Our upper (primal) bounds correspond to implementable perturbations whose suboptimality can be bounded by our lower (dual) bounds. Both bounding problems can be solved within seconds via cutting plane techniques that exploit the inherent problem structure. Our numerical experiments demonstrate that our perturbations can outperform the previously best results from the literature on artificial as well as standard benchmark problems.

LEGO: Optimal Online Learning under Sequential Price Competition

Shukai Li
(Department of Industrial Engineering & Management Sciences, Northwestern University)

Cong Shi
(Herbert Business School, University of Miami)

Sanjay Mehrotra
(Department of Industrial Engineering & Management Sciences, Northwestern University)

Contact Email: Shukai Li (shukaili2024@u.northwestern.edu)

Abstract:
We consider price competition among multiple sellers over a selling horizon of T periods. In each period, sellers simultaneously offer their prices and subsequently observe their respective demand that is unobservable to competitors. The realized demand of each seller depends on the prices of all sellers following a private unknown linear model. We propose a least-squares estimation then gradient optimization (LEGO) policy, which does not require sellers to communicate demand information or coordinate price experiments throughout the selling horizon. We show that our policy, when employed by all sellers, leads at a fast convergence rate $O\left(\frac{1}{\sqrt{T}}\right)$ to the Nash equilibrium prices that sellers would reach if they were fully informed. Meanwhile, each seller achieves an optimal order-of-$\sqrt{T}$ regret relative to a dynamic benchmark policy. Our analysis further shows that the *unknown individual price sensitivity contributes to the major difficulty of dynamic pricing in sequential competition* and forces regret to the order-of-$\sqrt{T}$ in the worst case. If each seller knows their individual price sensitivity coefficient, then a gradient optimization policy can achieve an optimal order-of-$\frac{1}{\sqrt{T}}$ convergence rate to Nash equilibrium as well as an optimal order-of-$\log T$ regret.

Statistical Properties of Robust Satisficing

Zhiyi Li
(The Hong Kong University of Science and Technology)

Yunbei Xu
(Laboratory for Information and Decision Systems, Massachusetts Institute of Technology)

Ruohan Zhan
(Department of Industrial Engineering and Decision Analytics, The Hong Kong University of Science and Technology)

Contact email: Zhiyi Li (lizhiyi11745@gmail.com)

Abstract:
The Robust Satisficing (RS) model is an emerging approach to robust optimization, offering streamlined procedures and robust generalization across various applications. However, the statistical theory of RS remains unexplored in the literature. This paper fills in the gap by comprehensively analyzing the theoretical properties of the RS model.

Notably, the RS structure offers a more straightforward path to deriving statistical guarantees compared to the seminal Distributionally Robust Optimization (DRO), resulting in a richer set of results. In particular, under the Wasserstein distance, we establish two-sided confidence intervals for the optimal loss without the need to solve a minimax optimization problem explicitly. We further provide finite-sample generalization error bounds for the RS optimizer. Importantly, our results extend to scenarios involving distribution shifts, where discrepancies exist between the sampling and target distributions. We further extend the RS model to the case of f-divergence and provide an asymptotic upper bound on the generalization error. Our numerical experiments show that the RS model consistently outperforms the baseline empirical risk minimization in small-sample regimes and under distribution shifts. Furthermore, compared to the DRO model, the RS model exhibits lower sensitivity to hyperparameter tuning, highlighting its practicability for robustness considerations.

Sinkhorn Distributionally Robust Optimization

Jie Wang
(School of Industrial and Systems Engineering, Georgia Institute of Technology)

Rui Gao
(Department of Information, Risk, and Operations Management, University of Texas at Austin)

Yao Xie
(School of Industrial and Systems Engineering, Georgia Institute of Technology)

Contact email: Jie Wang (jwang3163@gatech.edu)

Abstract:
We study distributionally robust optimization (DRO) with Sinkhorn distance—a variant of Wasserstein distance based on entropic regularization. We derive convex programming dual reformulation for general nominal distributions, transport costs, and loss functions. Compared with Wasserstein DRO, our proposed approach offers enhanced computational tractability for a broader class of loss functions, and the worst-case distribution exhibits greater plausibility in practical scenarios. To solve the dual reformulation, we develop a stochastic mirror descent algorithm with biased gradient oracles. Remarkably, this algorithm achieves near-optimal sample complexity for both smooth and nonsmooth loss functions, nearly matching the sample complexity of the Empirical Risk Minimization counterpart. Finally, we provide numerical examples using synthetic and real data to demonstrate its superior performance.

Randomized Assortment Optimization

Zhengchao Wang
(Imperial College Business School, Imperial College London)

Heikki Peura
(Aalto University School of Business)

Wolfram Wiesemanna
(Imperial College Business School, Imperial College London)

Contact email: Zhengchao Wang (z.wang19@imperial.ac.uk)

Abstract:
When a firm selects an assortment of products to offer to customers, it uses a choice model to anticipate their probability of purchasing each product. In practice, the estimation of these models is subject to statistical errors, which may lead to significantly suboptimal assortment decisions. Recent work has addressed this issue using robust optimization, where the true parameter values are assumed unknown and the firm chooses an assortment that maximizes its worst-case expected revenues over an uncertainty set of likely parameter values, thus mitigating estimation errors. In this paper, we introduce the concept of *randomization* into the robust assortment optimization literature. We show that the standard approach of deterministically selecting a single assortment to offer is not always optimal in the robust assortment optimization problem. Instead, the firm can improve its worst-case expected revenues by selecting an assortment randomly according to a prudently designed probability distribution. We demonstrate this potential benefit of randomization both theoretically in an abstract problem formulation as well as empirically across three popular choice models: the multinomial logit model, the Markov chain model, and the preference ranking model. We show how an optimal randomization strategy can be determined exactly and heuristically. Besides the superior in-sample performance of randomized assortments, we demonstrate improved out-of-sample performance in a datadriven setting that combines estimation with optimization. Our results suggest that more general versions of the assortment optimization problem—incorporating business constraints, more flexible choice models and/or more general uncertainty sets—tend to be more receptive to the benefits of randomization.

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

Jingying Ding
(Department of Analytics and Operations, National University of Singapore)

Yifan Feng
(Department of Analytics and Operations, National University of Singapore)

Ying Rong
(Antai College of Economics and Management, Shanghai Jiao Tong University)

Contact email: Jingying Ding (jyding@nus.edu.sg)

Abstract:
The exploration-exploitation trade-off is a fundamental concept arising when the decision maker needs to make repeated choices, whose rewards are unknown in priori. This paper delves into how humans navigate this trade-off through the lens of the multi-armed bandit (MAB) problem. Inspired by both behavioral economics and online learning literature, we introduce a novel family of behavioral policies called Myopic Quantal Choice (MQC). It is a dynamic adaptation of quantal choice models, deriving anticipated utilities directly from past rewards. MQC offers a simple method to describe the arm selection process, and yet is rich enough to quantify the exploration-exploitation trade-off through a "shrinkage rate of exploration."

Through both non-asymptotic and asymptotic analysis, we show that MQC admits intuitive properties that match the qualitative patterns of the laboratory experiment data. Particularly, MQC always converges to the optimal arm, thus capturing the "learning" effect. In addition, we characterize MQC's regret in the asymptotic regime and demonstrate the effects of "over-" and "under-" exploration; Over-exploration with a too small shrinkage rate parameter results in gradual deterioration of the lower bound of regret, while under-exploration with a too large rate parameter leads to sudden deterioration. Analysis of laboratory experiment data reveals that the MQC model excels in predictive power compared to other behavioral models. Insights from the asymptotic regime also extend to the finite horizon experiments. Particularly, when we fit the MQC model to the data, a prevalent tendency toward over-exploration becomes evident.

DAY 2

Offline Pricing with Shape-Restricted Demands

Jingren Liu
(Institute of Operations Research and Analytics, National University of Singapore)

Mabel C. Chou
(Institute of Operations Research and Analytics, National University of Singapore)

Hanzhang Qin
(Institute of Operations Research and Analytics & Department of Industrial Systems Engineering and Management, National University of Singapore)

Contact Email: Jingren Liu (jingren.liu@u.nus.edu)

Abstract:
We consider a fundamental problem in revenue management: feature-based offline pricing, where a firm needs to optimize a "one-shot" pricing decision for a single product based on feature information. The availability of historical data with price, features, and uncensored demands facilitates the demand learning process. However, evaluating a potentially more profitable pricing policy in an offline setting poses challenges due to the lack of counterfactual outcomes in observational data: We only observe how demand responds to the prices that were charged under a specific logging policy, but not to any other prices that could have been charged. This issue motivates the use of off-policy evaluation and off-policy optimization. In off-policy evaluation, our objective is to evaluate the revenue generated by a new pricing policy. Meanwhile, off-policy optimization aims to determine a price policy that generates near-optimal revenue. To achieve these objectives, we propose a doubly robust off-policy estimate for the revenue function. We show that this estimate minimizes the mean squared error in off-policy evaluation and that off-policy optimization based on this estimate is near-optimal with theoretical convergence guarantees. We emphasize that some favorable properties in estimation and decision-making arise from imposing a shape restriction on the demand estimation procedure.

Online Local False Discovery Rate Control: A Resource Allocation Approach

Ruicheng Ao
(Institute for Data, Systems, and Society, Massachusetts Institute of Technology)

Hongyu Chen
(Institute for Data, Systems, and Society, Massachusetts Institute of Technology)

David Simchi-Levi
(Institute for Data, Systems, and Society & Department of Civil and Environmental Engineering & Operations Research Center, Massachusetts Institute of Technology)

Feng Zhu
(Institute for Data, Systems, and Society, Massachusetts Institute of Technology)

Contact Email: Feng Zhu (fengzhu@mit.edu)

Abstract:
We consider the problem of sequentially conducting multiple experiments where each experiment corresponds to a hypothesis testing task. At each time point, the experimenter must make an irrevocable decision of whether to reject the null hypothesis (or equivalently claim a discovery) before the next experimental result arrives. The goal is to maximize the number of discoveries while maintaining a low error rate at all time points measured by local False Discovery Rate (FDR). We formulate the problem as an online knapsack problem with exogenous random budget replenishment. We start with general arrival distributions and show that a simple policy achieves a $O(\sqrt{T})$ regret. We complement the result by showing that such regret rate is in general not improvable. We then shift our focus to discrete arrival distributions. We find that many existing re-solving heuristics in the online resource allocation literature, albeit achieve bounded loss in canonical settings, may incur a $\Omega(\sqrt{T})$ or even a $\Omega(T)$ regret. With the observation that canonical policies tend to be too optimistic and over claim discoveries, we propose a novel policy that incorporates budget safety buffers. It turns out that a little more safety can greatly enhance efficiency — small additional logarithmic buffers suffice to reduce the regret from $\Omega(\sqrt{T})$ or even $\Omega(T)$ to $O(ln^2 T)$. From a practical perspective, we extend the policy to the scenario with continuous arrival distributions, time-dependent information structures, as well as unknown T. We conduct both synthetic experiments and empirical applications on a time series data from New York City taxi passengers to validate the performance of our proposed policies. Our results emphasize how effective policies should be designed to reach a balance between circumventing wrong accept and reducing wrong reject in online resource allocation problems with exogenous budget replenishment.

Transparent or Not? Optimal Performance Feedback in Gamified Services

Lin Chen
(INSEAD)

Guillaume Roels
(INSEAD)

Contact Email: Lin Chen (lin.chen@insead.edu)

Abstract:
Problem Definition: In many games or gamified services (e.g., fitness), users often receive feedback upon service completion. Such performance feedback, sometimes presented together with a goal and/or other users' scores, shapes their perception of individual performance (through prospect theory) and relative status (through social comparison). How transparent should service providers be in their disclosure of individual performance feedback to enhance user utility?

Methodology/Results: We employ a Bayesian persuasion framework to determine the optimal information disclosure policy, taking into account whether a goal is specified and whether the other users' scores are communicated. We find that when no goal is specified but the other users' scores are communicated, full (resp., no) information is optimal when users are ahead-seeking (resp., behind-averse); that is, the optimal information policy is either fully transparent or fully opaque. When a goal is specified but the other users' scores are not communicated, an upper censorship policy is optimal: the provider should reveal the exact scores to the low-performing users and only tell the high-performing users that they lie in the top range. Moreover, the higher the goal, the higher the cutoff between the two subsets of users. When a goal is specified and the other scores are communicated, the optimal information policy is again only fully transparent for a subset of users. Which subset depends on the nature of social comparison: When users are ahead-seeking, it is optimal to inform users in a range containing the goal that they are in that range; whereas, when they are behind-averse, it is optimal to additionally tell high-performing users their performance is in the top range.

Managerial Implications: Our paper offers guidelines to providers of games or gamified services to enhance user utility, and thus increase the value of their service, by engineering the design of their relative performance feedback.

Cross- vs. In-Region Courier Routing in On-Demand Delivery

Arseniy Gorbushin
(Rotman School of Management, University of Toronto)

Ming Hu
(Rotman School of Management, University of Toronto)

Xiaotang Yang
(Rotman School of Management, University of Toronto)

Yun Zhou
(Degroot School of Business, MacMaster University)

Contact Email: Arseniy Gorbushin (arseniy.gorbushin@rotman.utoronto.ca)

Abstract:
The on-demand delivery firms for food and grocery have grown rapidly. One claimed benefit of those firms is the potential to optimize courier routing by sharing couriers across the city and among many vendors, such as restaurants or grocery stores. However, it is puzzling to observe that the largest food-delivery platforms do allow restaurants to employ their own delivery units and grocery delivery pair couriers with single warehouses to deliver in their local neighborhood. To study whether and how sharing couriers across the city benefits a firm, we consider a spatial queuing model in which couriers are servers, and it takes a different amount of traveling time to serve customers depending on their vendor of choice, their own location, and the dispatch policy. Surprisingly, we find that in many cases, the in-region policy with a dedicated courier to each vendor can outperform the cross-region policy with a shared courier fleet among vendors. This result is attributed to the randomness and (potential) imbalance in the courier allocation that cross-region routing creates. Under the growth target strategy of achieving an exogenous demand rate, if the market is sufficiently large, the cross-region policy achieves a higher profit than the in-region one, and otherwise, if the market is small, the in-region policy is more profitable. However, under the profit maximization strategy of endogenizing the demand rate, we find that it is even more likely the in-region policy is optimal: in addition to a market size condition, a high enough service value is required for the cross-region policy to be optimal. This is because, under profit maximization, the firm may want to increase the delivery fee and limit distant customers, which reduces the benefit of cross-region routing. In addition, we show that in those markets where restaurants have a more distinctive cuisine, the cross-region policy tends to do better, and otherwise, in those markets with similar restaurants, the in-region policy has an edge. Lastly, we extend the results from the base model to account for multiple couriers and impatient customers.

Optimal probability measure decomposition

Jiangze Han
(University of British Columbia)

Christopher Thomas Ryan
(University of British Columbia)

Xin T. Tong
(National University of Singapore)

Contact Email: Xin T. Tong (xin.t.tong@nus.edu.sg)

Abstract:
We examine the infinite-dimensional optimization problem of finding a decomposition of a probability measure into K probability sub-measures to minimize specific loss functions inspired by applications in clustering and user grouping. We analytically explore the structures of the support of optimal sub-measures and introduce algorithms based on Wasserstein gradient flow, demonstrating their convergence. Numerical results illustrate the implementability of our algorithms and provide further insights.

Unraveling the Implications of Silent Labor Time (SLT) in the Gig Economy

Srishti Arora
(INSEAD)

Vivek Choudhary
(Nanyang Technological University)

Sameer Hasija
(INSEAD)

Ivan Gorban
(Careem)

Selim Turki
(Careem)

Ahmed Shekhani
(Careem)

Contact email: Srishti Arora (srishti.arora@insead.edu)

Abstract:
The gig economy has become integral to the global economy, driven by the labor flexibility it affords companies and the self-scheduling options available to workers. However, it also introduces an additional discretion for workers; they need to search for tasks during downtime, i.e., intervals when the platform does not assign them tasks. This period of uncompensated task-seeking, which we call "Silent Labor Time (SLT)," necessitates balancing effort between searching and executing tasks, impacting execution time. Our research aims to establish how the effort allocated during SLT affects workers' performance and earnings and identify factors that moderate this relationship. In food delivery, the distance drivers travel to find the next order, called "relocation distance," represents the effort allocated during SLT. Collaborating with a food delivery platform, we find that, on average, drivers relocate 2.6 km before each order, and a km increase in relocation distance reduces order allocation by 5.4%, order speed by 2.7%, and earnings by 14.8% in the subsequent hour. The primary reason for this decline is drivers allocating significant effort to searching for tasks during SLT, subsequently conserving energy when executing tasks. Relocations that are not towards familiar clusters and reduce supply-demand balance are most detrimental to workers' performance and earnings. Our findings suggest that relocation adversely affects drivers' earnings and operational performance in subsequent orders, ultimately impacting the platform's efficiency. We offer actionable insights by suggesting strategies for the management of effort allocation during SLT by different driver groups.

Robust Auction Design with Support Information

Jerry Anunrojwong
(Columbia University, Graduate School of Business)

Santiago R. Balseiro
(Columbia University, Graduate School of Business)

Omar Besbes
(Columbia University, Graduate School of Business)

Contact Email: Jerry Anunrojwong (jerryanunroj@gmail.com)

Abstract:
A seller wants to sell an indivisible item to n buyers. The buyer valuations are drawn i.i.d. from a distribution, but the seller does not know this distribution; the seller only knows the support [a, b]. To be robust against the lack of knowledge of the environment and buyers' behavior, the seller optimizes over dominant strategy incentive compatible (DSIC) mechanisms, and measures the worst-case performance relative to an oracle with complete knowledge of buyers' valuations. Our analysis encompasses both the regret and the approximation ratio objectives.

For these objectives, we derive an optimal mechanism in quasi-closed form, and the associated performance, as a function of the support and the number of buyers n. Our analysis reveals three regimes of support information and a new class of robust mechanisms. i.) With "low" support information, the optimal mechanism is a second-price auction (SPA) with a random reserve, a focal class in the earlier literature. ii.) With "high" support information, we show that second-price auctions are strictly suboptimal, and we establish that an optimal mechanism belongs to a novel class of mechanisms we introduce, which we call pooling auctions (POOL); whenever the highest value is above a threshold, the mechanism still allocates to the highest bidder, but otherwise the mechanism allocates to a uniformly random buyer, i.e., pools low types. iii.) With "moderate" support information, we establish that a randomization between SPA (with a random reserve price) and POOL (with a random threshold) is optimal.

We also characterize optimal mechanisms within nested central subclasses of mechanisms: standard mechanisms that only allocate to the maximum value bidder, SPA with random reserve, and SPA with no reserve. We show strict separations in terms of performance across classes, implying that deviating from standard mechanisms is necessary for robustness. Lastly, we show that the same results hold under other distribution classes that capture "positive dependence", namely: i.i.d., mixture of i.i.d., and exchangeable and affiliated distributions, as well as i.i.d. regular distributions.

Neural-Network Mixed Logit Choice Model: Statistical and Optimality Guarantees

Zhi Wang
(Department of Information, Risk and Operations Management, McCombs School of Business, The University of Texas at Austin)

Rui Gao
(Department of Information, Risk and Operations Management, McCombs School of Business, The University of Texas at Austin)

Shuang Li
(School of Data Science The Chinese University of Hong Kong (Shenzhen))

Contact Email: Zhi Wang (zhi.wang@mccombs.utexas.edu)

Abstract:
The mixed logit choice model, widely used in operations, marketing, and econometrics, represents choice probabilities as mixtures of multinomial logits. This study explores the effectiveness of representing the mixed logit model using a single-hidden-layer neural network, which approximates the mixture distribution as an equally weighted distribution on finite consumer types. From the statistical perspective, we show that the approximation error of the neural network does not suffer from the curse of dimensionality, and that overparameterization does not lead to overfitting. From the optimization perspective, we prove the noisy gradient descent algorithm can find the global optimizer of the entropic regularized non-convex parameter learning problem, up to an error inversely proportional to the number of consumer types. Experiments on synthetic and real datasets validate the algorithm's superior in-sample and out-of-sample performance compared to existing benchmarks. These findings underscore the potential of even shallow yet over-parameterized neural network representations, coupled with efficient training algorithms, to effectively learn complex choice models with strong statistical and optimality guarantees.

A Nonparametric Approach with Marginals for Modeling Consumer Choice

Yanqiu Ruan
(Singapore University of Technology and Design)

Xiaobo Li
(National University of Singapore)

Karthyek Murthy
(Singapore University of Technology and Design)

Karthik Natarajan
(Singapore University of Technology and Design)

Contact Email: Yanqiu Ruan (yanqiu_ruan@mymail.sutd.edu.sg)

Abstract:
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 utility 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 polynomialsized 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 convenient to be 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 bestfitting MDM choice probabilities reduces to solving a mixed integer convex program. The results extend naturally to the case where the alternatives can be grouped based on the similarity of the marginal distributions of the utilities. Numerical experiments show that MDM provides better representational power and prediction accuracy than multinominal logit and significantly better computational performance than RUM.

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 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.

ORGANIZING TEAM

Hou Xinyu
Student
IORA

Li Jingying
Student
IORA

Wu Lan
Student
DAO

Yang Wenjie
Student
NUSRI-CQ

Chen Zhi
Assistant Professor
DAO

Cheung Wang Chi
Assistant Professor
ISEM

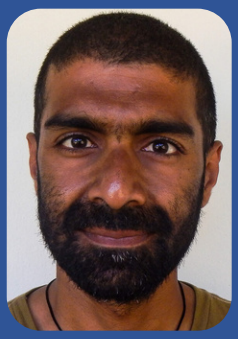

Rajiv K
Sr. Research Fellow
IORA

GETTING HERE

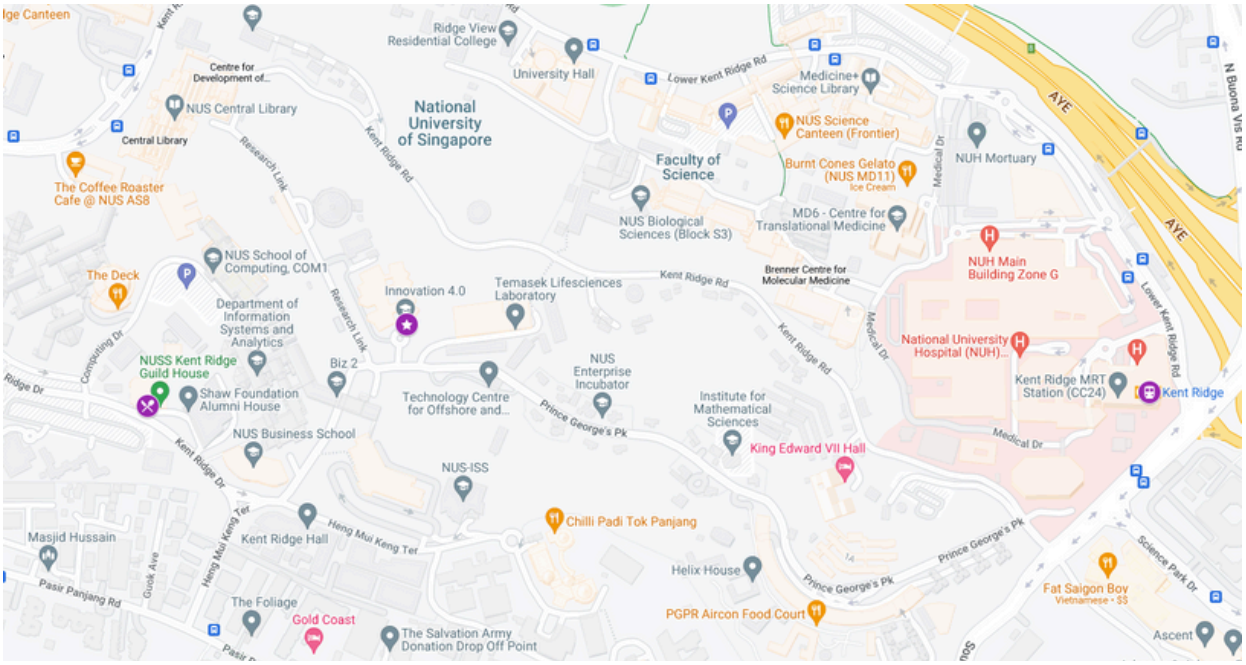

To get to Innovation 4.0, NUS, scan this code or visit https://bit.ly/4cciXHL for directions.

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 alternate modes of travel. Consult the NUS mobility services for latest parking info.

Notes