OPLOG Division Seminars
From September 2018 – Onwards

Date: Wednesday, April 24th, 2019

Speaker: Ho-Yin Mak, Oxford University, UK
Title: "Peer-to-Peer Crowdshipping as an Omnichannel Retail Strategy"
Time: 12:30 PM - 1:30 PM
Place: DL 125

Abstract: Crowdshipping, a novel practice built upon the sharing economy, has been adopted by some retailers to tackle the notorious last-mile delivery problem in urban omnichannel retail. In this work, we study the potential of P2P crowdshipping, i.e., enlisting in-store shoppers to deliver online orders in their vicinity, and its impact on the marketing-operations interface for an omnichannel retailer. Specifically, P2P crowdshipping could potentially help the retailer improve delivery efficiency, and gain an additional lever for price discrimination. When these two effects interact, we find that the favorability of crowdshipping heavily depends on product and market characteristics, as well as how shopper-deliverers are reimbursed. For necessity goods, a cost-based reimbursement scheme could lead to a win-win outcome in both the retailer's profit and consumer surplus; for higher-end products, reimbursing deliverers a premium on top of their delivery costs would be favorable.

Date: Wednesday, April 17th, 2019

Speaker: Shouqiang Wang, University of Texas at Dallas
Title: "Warning Against Recurring Risks: An Information Design Approach"
Time: 2:30 PM - 3:30 PM
Place: HA 969

Abstract:

The World Health Organization seeks effective ways to alert its member states about global pandemics. Motivated by this challenge, we study a public agency's problem of designing warning policies to mitigate potential disasters that occur with advance notice. The agency privately receives early information about recurring harmful events and issues warnings to induce an uninformed stakeholder to take preemptive actions.

The agency's decision to issue a warning critically depends on its reputation, which we define as the stake-holder's belief regarding the accuracy of the agency's information. The agency faces then a trade-off between eliciting a proper response today and maintaining its reputation in order to elicit responses to future events.

We formulate this problem as a dynamic Bayesian persuasion game, which we solve in closed form. We find that the agency sometimes strategically misrepresents its advance information about a current threat in order to cultivate its future reputation. When its reputation is sufficiently low, the agency downplays the risk and actually downplays more as its reputation improves. By contrast, when its reputation is high, the agency sometimes exaggerates the threat and exaggerates more as its reputation deteriorates. Only when its reputation is moderate does the agency send warning messages that fully disclose its private information.

Our study suggests a plausible and novel rationale for some of the false alarms or omissions observed in practice. We further test the robustness of our findings to imperfect advance information, disasters without advance notice, and heterogenous receivers.

Date: Monday, April 15th, 2019

Speaker: Jing Dong, Columbia University
Title: "Optimal scheduling of proactive care with patient degradation"
Time: 2:30 PM - 3:30 PM
Place: HA 968

Abstract:

Healthcare is a limited resource environment where scarce capacity is often reserved for the most severe patients. However, there has been a growing interest in the use of preventive care to treat patients early on before they deteriorate. On one hand, providing care for patients when they are less critical could mean that fewer resources are needed to return them to a healthy, stable state. On the other hand, utilizing limited capacity for patients who may never need care in the future takes the capacity away from other more critical patients who need it now. To understand this tension, we propose a multi-server queueing model with two patient classes: moderate and urgent. A moderate patient who does not receive treatment may recover and leave or may deteriorate and become an urgent patient. In this setting, we characterize how moderate and urgent patients should be prioritized for care when proactive care for moderate patients is an option. The analysis replies on several interesting applications of optimal control theory.

This is joint work with Yue Hu and Carri Chan.

Date: Monday, April 8th, 2019

Speaker: Tamar Cohen, MIT
Title: "High-Low Promotion Policies for Peak End Demand Models"
Time: 2:30 PM - 2:30 PM
Place: HA 969

Abstract: In-store promotions are a highly effective marketing tool that can have a significant impact on profit. Incorporating important consumer behavioral effects in the demand model is crucial in order to better explain the relationship between price and consumer demand. In this research, we propose a new demand model that relies on behavioral effects (such as the minimum price within a bounded memory). We show that this new demand model predicts actual sales more accurately than current methods. In order to determine promotion strategies, subsequently, we establish that in the presence of the proposed demand model, High-Low pricing policy is optimal under diagonal dominance conditions (so that the current period price dominates both past period price effects and competitive product price effects on the demand). For promotion planning for a single item, we propose a compact Dynamic Programming (DP) approach that can find the optimal High-Low promotion plan in polynomial time. In the cases where these diagonal dominance conditions do not hold, we provide an analytical guarantee and illustrate that the proposed DP yields near-optimal solutions fast even in the absence of optimality. For the general case of promotion planning for multiple items, we illustrate that theoretically, the problem is NP-hard in the strong sense. Using the proposed DP formulation as a subroutine, we propose a Polynomial Time Approximation Scheme (PTAS) for the general case. Finally, we test our approach on data from large retailers and demonstrate an average of 5.1−15.6% increase in profit relative to the retailer's current practices.

Date: Monday, April 1st, 2019

Speaker: Van-Anh Truong, Columbia University
Title: "Dynamic Optimization of Mobile Push Advertising Campaigns"
Time: 2:30 PM - 3:30 PM
Place: HA 968

Abstract: We study a novel resource-allocation problem faced by Alibaba Group. In this problem, mobile "push messages" must be sent over the course of a day to hundreds of millions of users. Each message can be sent to any number of users, and yields a reward when it generates a clickthrough, subject to a budget constraint on the total reward over all users for the message. This budget represents the maximum amount that an advertiser is willing to pay for clickthroughs for the message on a given day. Given users' diverse preferences, the problem aims to deliver the "right messages" to the "right users" to maximize ad revenues without overwhelming each user with too many messages. Due to the large size of the real application, we analyze algorithms for the above problem in an asymptotic regime. We consider a novel scaling of the problem "size," called big-data scaling. In this scaling, as the problem size grows, the number of users, as well as their diversity, grow. The scaling captures the fact that individual

user information remains highly granular and distinctive even as the size of the user base increases. We prove that solving the problem as a static assignment problem results in a regret of O(√ t), where t is the parameter scaling the problem. Furthermore, adding a single recourse opportunity, by sending push messages in two cycles over the course of a day and making use of information observed in the first cycle to adapt decisions in the second cycle, can reduce the regret to O(t 1/4 log t). Finally, the difference in regret between the static and dynamic strategy can be Ω(√ t). Numerical experiments on three real data sets, each containing several hundred million users, show that the latter strategy improves the regret of the former by at least 10%-50%.

This is joint work with Xinshang Wang, Shenghuo Zhu, and Qiong Zhang.

Date: Monday, March 18th, 2019

Speaker: Opher Baron, Rotman School of Business, University of Toronto
Title: "Data Driven Forecasting and Revenue Management with Word of Mouth Dependent Reference Prices"
Time: 2:30 PM - 3:30 PM
Place: DL 125

Abstract: We use data on the retail business of TMall to consider revenue management in the presence of reference prices. We consider questions as: Does reference price improves demand forecasts? And, does considering word of mouth (WOM) in the formation of reference price forecasts benefit revenue management? Revenue management considers the impact of prices on current and future sales via forecasts. The related literature uses reference price to study this impact empirically and theoretically, mostly focusing on sales of specific items to repeat customers. It ignores the effect of WOM on reference price. We (I) develop scalable, data driven methodologies to compare the effectiveness of different forecasts, (II) demonstrate these methodologies, (III) introduce models capturing the effect of WOM on reference price, and (IV) formulate and study (theoretically and numerically) the revenue management problem when forecasts are reference price-dependent. We provide a foundation for systematic implementation of revenue management in the presence of reference price effects. We (I) formulate the impact of WOM on reference price and investigate the effectiveness of different forecasts; and (II) demonstrate that revenue management could benefit from WOM-dependent reference price models. The improved accuracy and revenue management performance of our forecasting models with reference price affected by WOM support their usage in practice.

Joint work with: Deng Chang, Southern University of Science and Technology and He Simai, Yuan Hongsong, Shanghai University of Finance and Economics.

Date: Monday, March 11th, 2019
Speaker: Nikos Trichakis, MIT
Title: "On the Efficacy of Static Prices for Revenue Management in the Face of Strategic Customers"
Time: 2:30 PM - 3:30 PM
Place: **HA 254**

Abstract: We consider canonical network revenue management (RM) problems wherein a monopolist seller posts prices for multiple products that are for sale over a fixed horizon so as to maximize expected revenues. Products are differentiated and subject to joint capacity constraints. For different types of customer strategic behavior, we derive performance guarantees for the efficacy of static pricing. In particular, when customers strategize only on when to buy, i.e., they are "forward-looking," we show that static prices are asymptotically optimal for a broad class of customer utility models. When customers strategize on both when and what to buy, i.e., they are both forward-looking and choose what product to buy, we derive a constant-factor guarantee.

Date: Monday, March 4th, 2019
Speaker: Kevin Shang, Duke University
Title: "Inventory Management for Multidivisional Firms with Cash Pooling"
Time: 2;30 pm - 3:30 pm
Place: **DL 125**

Abstract: Cash pooling is a powerful management tool that allows each division's cash balance to be transferred to a single account managed by the corporate treasury in the headquarter. While the reported benefits of cash pooling are associated with the reduction of transaction and financing costs, the value of cash pooling is not clear from a perspective of improving operational efficiency. In this talk, we examine the benefit of cash pooling on inventory replenishment for multidivisional firms through two models. The first model considers a series supply chain in which demand occurs at the most downstream division, and each division orders from its upstream division. The second model considers a distribution supply chain in which each division replenishes from an outside supplier to meet its local demand. The corporate treasury receives cash payments from customers and determines how much to invest externally for a positive return in each period. There are holding costs for the on-hand inventory and unfilled demands incur backorder costs. The objective is to obtain the optimal joint cash retention and inventory replenishment policy that maximizes the expected net worth (equity). For the series model, we show that the optimal policy has a simple structure—each division implements a base-stock policy for inventory replenishment; the corporate treasury monitors the system working capital and implements a two-threshold policy for cash retention. For the distribution model, we provide a simple and effective heuristic derived from the construction of a lower bound to the optimal value function. Our lower bound improves the so-called Lagrangian-relaxation bound and the induced-penalty bound in the multi-echelon literature. We quantify the value of cash pooling for both models. Our research leads to a new online simulation game, called Cash Beer Game, which incorporates cash flows into the standard Beer Game. I shall demonstrate this online game in this talk.

Date: Monday, March 4th, 2019

Speaker: Kevin Shang, Duke University
Title: "Inventory Management for Multidivisional Firms with Cash Pooling"
Time: 2;30 PM - 3:30 pm
Place: **DL 125**

Abstract: Cash pooling is a powerful management tool that allows each division's cash balance to be transferred to a single account managed by the corporate treasury in the headquarter. While the reported benefits of cash pooling are associated with the reduction of transaction and financing costs, the value of cash pooling is not clear from a perspective of improving operational efficiency. In this talk, we examine the benefit of cash pooling on inventory replenishment for multidivisional firms through two models. The first model considers a series supply chain in which demand occurs at the most downstream division, and each division orders from its upstream division. The second model considers a distribution supply chain in which each division replenishes from an outside supplier to meet its local demand. The corporate treasury receives cash payments from customers and determines how much to invest externally for a positive return in each period. There are holding costs for the on-hand inventory and unfilled demands incur backorder costs. The objective is to obtain the optimal joint cash retention and inventory replenishment policy that maximizes the expected net worth (equity). For the series model, we show that the optimal policy has a simple structure—each division implements a base-stock policy for inventory replenishment; the corporate treasury monitors the system working capital and implements a two-threshold policy for cash retention. For the distribution model, we provide a simple and effective heuristic derived from the construction of a lower bound to the optimal value function. Our lower bound improves the so-called Lagrangian-relaxation bound and the induced-penalty bound in the multi-echelon literature. We quantify the value of cash pooling for both models. Our research leads to a new online simulation game, called Cash Beer Game, which incorporates cash flows into the standard Beer Game. I shall demonstrate this online game in this talk.

Date: Monday, February 25th, 2019

Speaker: Kostas Bimpikis, Stanford University
Title: "Spatial Pricing in Ride-Sharing Networks"
Time: 2:30 PM - 3:30 PM
Place: HA 968

Abstract: We explore spatial price discrimination in the context of a ride-sharing platform that serves a network of locations. Riders are heterogeneous in terms of their destination preferences and their willingness to pay for receiving service. Drivers decide whether, when, and where to provide service so as to maximize their expected earnings, given the platform's prices. Our findings highlight the impact of the demand pattern on the platform's prices, profits, and the induced consumer surplus. In particular, we establish that profits and consumer surplus are maximized when the demand pattern is "balanced" across the network's locations. In addition, we show that they both increase monotonically with the balancedness of the demand pattern (as formalized by its structural properties). Furthermore, if the demand pattern is not balanced, the platform can benefit substantially from pricing rides differently depending on the location they originate from. Finally, we consider a number of alternative pricing and compensation schemes that are commonly used in practice and explore their performance for the platform.

Joint work with Ozan Candogan and Daniela Saban

Link to the paper: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2868080

Date: Monday, February 11th, 2019

Speaker: David Kirkpatrick, Professor Emeritus, UBC Computer Science Department
Title: "Minimizing Interference Potential Among Moving Entities"
Time: 2:30 PM - 3:30 PM
Place: HA 969

Abstract: We consider the problem of monitoring the interference among a collection of entities moving with bounded speed in d-dimensional Euclidean space. Uncertainty in entity locations due to unmonitored and unpredictable motion gives rise to a space of possible entity configurations at each moment in time, with possibly very different interference properties. We define different measures of what we call the interference potential of such spaces to describe the interference that might actually occur.

We study the extent to which restricted monitoring frequency impacts interference potential, through the analysis of a clairvoyant scheme (one that knows the trajectories of all entities) subject to the same monitoring frequency restriction.

This forms a benchmark for the analysis of uninformed schemes. In this framework, we describe and analyse an adaptive monitoring scheme for minimizing interference potential over time that is competitive (to within a constant factor) with any other scheme (in particular, a clairvoyant scheme) over modest sized time intervals.

As a natural application, imagine that the entities are transmission sources, with associated broadcast ranges, moving in three dimensions. Two such entities, transmitting on the same channel, interfere if their broadcast ranges intersect. Uncertainty in the location of a transmission source effectively expands its broadcast range to a potential broadcast range. The chromatic number of the intersection graph of these potential broadcast ranges, one of our interference potential measures, gives the minimum number of broadcast channels required to avoid interference. Our scheme provides an adaptive, locally updated, channel assignment scheme that is competitive over time, in terms of the number of broadcast channels used with a fixed monitoring frequency, with any other such scheme.

(Based on joint work with Daniel Busto and Will Evans.)

Date: Monday, February 4th, 2019

Speaker: Juan Serpa, McGill University
Title: "Oversight and Efficiency in Public Projects"
Time: 2:30 PM - 3:30 PM
Place: HA 969

Abstract: In the U.S., four in ten public infrastructure projects report delays or cost overruns. To tackle this problem, regulators often scrutinize the project contractor's operations. We investigate the causal effect of government oversight on project efficiency by gleaning 262,857 projects that span seventy-one U.S. federal agencies and 54,739 contractors. Our identification strategy exploits a regulatory bylaw: if a project's anticipated budget exceeds a threshold value, the contractor's operations are subject to surveillance from independent procurement officers; otherwise, these operational checks are waived. Using a regression discontinuity design, we find that oversight is obstructive to the project's operations, especially when the contractor (i) has no prior experience in public projects, (ii) is paid with a fixed-price contract that includes performance-based incentives, and (iii) performs a labor-intensive task. In contrast, oversight is least obstructive — or beneficial — when the contractor (i) is experienced, (ii) is paid with a time-and-materials contract, and (iii) performs a machine-intensive task.

Date: Monday, January 28th, 2019

Speaker: Vibhuti Dhingra, PhD Student, Operations and Logistics Division, UBC Sauder School of Business
Title: "Managing Reputation Risk in Supply Chains: The Role of a Risk-sharing Contract"
Time: 2:30 PM - 3:30 PM
Place: HA 968

Abstract: When a supplier fails to comply with social and environmental standards, the buyer's reputation suffers. We study the role of a risk-sharing contract in managing such reputation concerns. We find that risk-sharing can both decrease and increase supplier violations, and we characterize conditions where each occurs. When risk-sharing reduces violations, the buyer clearly benefits from a lower reputation risk. But if the reputation costs are low, we show that the buyer can benefit by using risk-sharing to squeeze the supplier's margin even if this increases violations.

*Joint work with Harish Krishnan

Date: Monday, January 21st, 2019

Speaker: Maurice Queyranne, Operations and Logistics Division, UBC Sauder School of Business
Title: "Modeling Convex Subsets and Related Shape Requirements with Mixed Integer Programming"
Time: 2:30 PM - 3:30 PM
Place: HA 968

Abstract: Many planning and location problems in forestry, mining, political districting, farmland consolidation, as well as in data mining, entail the selection of a subset, or a partition into subsets, that should satisfy some, often ill-specified, shape constraints. A typical such shape constraint is that the subsets be convex, or "approximately convex". In general terms, a subset S of a given (finite) set of points in a convexity structure is convex if every given point that is in the convex hull of S is itself in S. We are interested in modeling such convexity restrictions when the given set of points is finite, i.e., in obtaining polyhedral descriptions of the characteristic vectors of the convex subsets, and extended formulations that are compact (i.e., polynomial-sized) and if possible ideal (i.e., with integer extreme solutions). This leads us to consider the associated optimization problem, where each point has a given weight (of arbitrary sign) and we seek a convex subset with maximum (or minimum) total weight. In many applications, these restrictions and optimization problems arise in a low-dimensional space and are subject to

additional constraints.

Modeling convex subset restrictions is well understood for the one-dimensional case. On the other hand, the optimization problem is NP-hard for the standard (vector space) convexity in dimension three and higher, and hard to approximate when the dimension is part of the input. For the two-dimensional (planar) case, convexity can be enforced by a polynomial (quartic) number of linear inequalities in the natural binary variables, but the resulting formulation is very weak. We present a compact ideal extended formulation, derived from the cubic-time dynamic programming optimization algorithm of Eppstein et al. (1992) and Bautista-Santiago et al. (2011). We point out open questions on finding more compact and/or tighter formulations, as well as exact or approximate combinatorial separation and optimization algorithms. We also consider related notions of convexity in partially ordered sets and in networks and metric spaces. (This talk presents joint work with numerous co-authors.)

Date: Monday, January 14th, 2019

Speaker: Shawn Mankad, Cornell University
Title: "A for Effort? Using the Crowd to Identify Moral Hazard in NYC Restaurant Hygiene Inspections"
Time: 2:30 PM - 3:30 PM
Place: HA 968

Abstract: From an upset stomach to a life-threatening foodborne illness, getting sick is all too common after eating in restaurants. While health inspection programs are designed to protect consumers, such inspections typically occur at wide intervals of time, allowing restaurant hygiene to remain unmonitored in the interim periods. Information provided in online reviews may be effectively used in these interim periods to gauge restaurant hygiene. In this paper, we use textual information from online reviews of restaurants to effectively identify cases of hygiene violations in restaurants, even after the restaurant has been inspected and certified, thereby identifying moral hazard. Using a dataset of restaurant hygiene inspections in New York City from 2010 through 2016, and the associated set of online reviews for the same set of restaurants from Yelp, we utilize supervised machine learning techniques to develop a hygiene dictionary specifically crafted to identify hygiene-related concerns. The use of this dictionary, and the related word counts in online reviews, allows us to identify systematic instances of moral hazard, wherein restaurants with positive hygiene inspection scores are seen to regress in their hygiene maintenance within 90 days of receiving the inspection scores. To the extent that social media provides some visibility into the hygiene practices of restaurants, we argue that the effects of information asymmetry that lead to moral hazard may be partially mitigated in this context. Based on our work, we also provide strategies for how cities and policy-makers may design effective restaurant inspection programs, through a combination of traditional inspections and the appropriate use of social media.

Date: Friday, January 4th, 2019

Speaker: Ashwin Venkataraman, Harvard Institute for Quantitative Social Sciences
Title: "Fitting Large-scale Mixture of Logit Models: A Convex Optimization Approach"
Time: 3:00 PM - 4:00 PM
Place: HA 966

Abstract: The wealth of data that firms are able to collect nowadays provides new opportunities and capabilities for firms to better solve classical problems within operations. At the same time, the data imposes new and unique challenges. My research aims at building novel methodologies to address these challenges, drawing upon ideas and techniques proposed in the machine learning literature.

Problem: In this talk, I will present one of the problems I have worked on: learning customer demand. This is an important problem for firms since demand predictions serve as key inputs to many critical decisions. Mixture of logit models are popular tools for modeling customer demand within operations, since they account for heterogeneity in customer preferences as well as capture rich substitution patterns in customers' choice behavior. The standard setup involves fitting the mixture model to historical sales transactions and product availability data, and then using the fitted model to make predictions.

Challenges: The main challenge in estimating the mixture model is that the mixing distribution is often unknown, and imposing a priori parametric assumptions can lead to model misspecification issues, resulting in poor predictions and suboptimal decisions. In addition, existing methods typically estimate the mixture model by solving non-convex optimization problems, and therefore lack strong convergence guarantees. These methods also suffer from many computational issues which are exacerbated when applied to the large-scale transaction data that firms have now access to.

Contributions: We propose a new methodology for nonparametric estimation of large-scale mixture of logit models. We formulate the estimation problem as a constrained convex program, and apply the conditional gradient (a.k.a. Frank-Wolfe) algorithm to solve this convex program. Theoretically, we show that our estimator has a sublinear convergence rate to the global optimal solution, and characterize the structure of the mixture types recovered by our method. We also empirically validate our methodology: (1) using synthetic data, we show that our nonparametric estimator is able to recover different ground-truth mixing distributions, despite having no knowledge of their underlying structure; (2) on real data, we show that our method outperforms the standard expectation-maximization (EM) benchmark not only on speed and predictive accuracy, but also decision accuracy: it is able to extract around 23% more revenue from customers. This shows that the method used to fit the mixture model matters greatly in practice.

This talk is based on joint work with Prof. Srikanth Jagabathula (NYU Stern and Harvard Business School) and Prof. Lakshminarayanan Subramanian (NYU Courant).

Date: Wednesday, December 12th, 2018

Speaker: Yining Wang, Carnegie Mellon University
Title: "Dynamic Assortment Planning under Plain and Contextual Discrete Choice Models"
Time: 3:00 PM - 4:00 PM
Place: HA 966

Abstract: In this talk, I will present my works (joint with Xi Chen and Yuan Zhou) on dynamic assortment planning with various discrete choice models, including the plain multinomial logit (MNL) model and the linear contextual MNL model. For each arriving customer, the seller offers an assortment of substitutable products, and then the customer makes purchases according to pre-specified choice models. Since all utility parameters of customers are unknown, the seller needs to simultaneously learn customers' choice behavior and make dynamic decisions on assortments based on the current knowledge. For both the plain MNL and the contextual MNL discrete choice models, we develop efficient dynamic policy and establish near-optimal upper bounds on their corresponding regret, an important evaluation measure of the expected revenue of the planned assortments. Numerical results of our proposed policies will also be presented.

Date: Monday, December 10th, 2018

Speaker: Christopher Ryan, The University of Chicago School of Business
Title: "Mixed-integer bilevel representability"
Time: 2:30 PM - 3:30 PM
Place: Henry Angus 968

Abstract: We study the representability of sets that admit extended formulations using mixed-integer bilevel programs. We show that feasible regions modeled by continuous bilevel constraints (with no integer variables), complementarity constraints, and polyhedral reverse convex constraints are all finite unions of polyhedra. Conversely, any finite union of polyhedra can be represented using any one of these three paradigms. We then prove that the feasible region of bilevel problems with integer constraints exclusively in the upper level is a finite union of sets representable by mixed-integer programs and vice versa.

Further, we prove that, up to topological closures, we do not get additional modeling power by allowing integer variables in the lower level as well. To establish the last statement, we prove that the family of sets that are finite unions of mixed-integer representable sets forms an algebra of sets (up to topological closures).

Date: Friday, December 7th, 2018

Speaker: Lennart Baardman, MIT Operations Research Centre
Title: "Learning to Optimize in Offline and Online Advertising"
Time: 3:00 PM - 4:00 PM
Place: HA 968

Abstract: Big data and the internet are shifting the advertising paradigm. The amount of readily available data from both point-of-sale systems and cookies has grown, enabling a shift from qualitative design to quantitative tools for advertising. At the same time, the internet has created new marketing channels that allow personalized advertising next to traditional mass marketing. In this work, we address how a retailer can maximize profits both in their offline retail store as well as their online search/social channel. The different channels share common goals, but at the same time, they are characterized by distinguishing features such as the length of the planning horizon and the ability to personalize advertisements. In the online setting, we take an exploration-exploitation perspective to model how an advertiser should periodically select a portfolio of target audiences to bid on in online advertising auctions. Current methods often consider the revenue and cost of bidding on a target to be known. However, many target audiences have never been shown an ad before, and hence, learning the value of each target while optimizing revenue is important. Borrowing ideas from upper confidence bound algorithms and robust optimization, we develop an optimistic-robust learning algorithm with provably-bounded regret. Computations, based on real advertising campaign data, show that this algorithm outperforms other passive and active learning approaches. In the offline setting, we model how to schedule promotion vehicles (e.g., endcap displays, printed flyers, tv commercials) to maximize profits under capacity constraints as a non-linear integer optimization problem. Even though this problem is APX-hard, we demonstrate that it can be solved efficiently through fast and practically near-optimal algorithms. One of these algorithms is based on solving an innovative approximate integer optimization model of polynomial size. Due to its polynomial size, the computation time of this model is limited, while being able to optimize to any level of accuracy. An added benefit of this model is that it can incorporate numerous new constraints that retailers might want to impose on the problem. Using supermarket data, we show our models lead to a profit increase of 2-9%.

Date: Monday, December 3rd, 2018

Speaker: Karla Hoffman, George Mason University
Title: "The use of hybrid optimization algorithms for problems that are either unsolvable by state-of-the-art software packages or for problems where such codes are too slow for the real-time application"
Time: 2:30 PM - 3:30 PM
Place: Henry Angus 968

Abstract: Optimization algorithms have increasingly played a role in improving the operations of many standard corporate activities such as the supply chain, real-time scheduling and routing, and determining the allocation and pricing of goods and services through auctions. In this talk, we examine problems that arose in a high-profile government auction where the Federal Communications Commission (FCC) bought back spectrum from TV stations, packed the remaining broadcasters into a smaller swath of spectrum and sold the acquired spectrum to the wireless industry. Optimization is used before, during and after this auction to assure that multiple governmental goals are met.

Our second example considers how the military can use similar optimization strategies to allocate limited spectrum during combat situations when spectrum is scarce and communication vital. In both instances, the problems require the solution to extremely difficult optimization problems. In each case, our approach is to use a combination of heuristics, decompositions, and constraint programming to create an overall algorithm that is capable of solving to global or near global optimality problems with millions of variables and hundreds of thousands of constraints.

We finally present an example of a real-time routing and scheduling problem where one needs near-optimal solutions in less than a second. In this case, we ask the question: Under what conditions are the following algorithms the best choice: enumeration, constraint programming, heuristics, decision diagrams or global optimization techniques best?

In all applications, we use realistic data sets and make the data sets available to the research community.

Date: Monday, November 19th, 2018

Speaker: Martin Puterman, UBC Sauder School of Business
Title: "Points Gained: a New Metric for Assessing National Football League (NFL) Team Performance"
Time: 2:30 PM - 3:30 PM
Place: HA 968

Abstract: In this talk I describe our revisit to Carter and Machol's 1971 paper "Operations Research on Football" with new data, a theoretical foundation and new ways to use the results. Our research uses rather obscure Markov Decision Process theory to develop a rigorous framework for deriving league value functions. It shows that the obvious value function equation has a unique solution that equals the bias of the underlying Markov reward process thus providing a precise interpretation for the value function. Our research proceeds by extending Brodie's recent work on "shots gained" in golf, by showing how to use the Bellman equation to derive a "points gained" metric for football. Data analysis from the 2013-2016 NFL seasons shows how the points gained metric identifies and provides new insights into specific factors that distinguish team performance.

This talk is based on joint work with Tim Chan and Craig Fernandes of the University of Toronto.

Date: Friday, November 9th, 2018

Speaker: Miguel Anjos, Polytechnique Montreal
Title: "Mathematical optimization approaches for facility layout problems: The state-of-the-art and future research directions"
Time: Noon - 1:00 PM
Place: Henry Angus **333**

Abstract: Facility layout problems are an important class of operations research problems that has been studied for several decades. Most variants of facility layout are NP-hard, therefore global optimal solutions are difficult or impossible to compute in reasonable time. Mathematical optimization approaches that guarantee global optimality of solutions or tight bounds on the global optimal value have nevertheless been successfully applied to several variants of facility layout. This review covers three classes of layout problems, namely row layout, unequal-areas layout, and multifloor layout. We summarize the main contributions to the area made using mathematical optimization, mostly mixed integer linear optimization and conic optimization. For each class of problems, we also briefly discuss directions that remain open for future research.

Date: Monday, October 29th, 2018

Speaker: Jussi Keppo, National University of Singapore Business School
Title: "Investment Decisions and Falling Cost of Data Analytics"
Time: 2:30-3:30 PM
Place: Henry Angus 968

Abstract: We model a risk-averse decision maker who optimizes the size of an investment, leverage used in the investment, and the level of information on the investment through costly data analytics. We show that borrowing-constrained or highly risk-averse investors have low demand for data analytics. We also show that the demand of data analytics is highest for investment opportunities with high expected returns, and the demand of data analytics is lumpy for opportunities with low expected returns even without analytics fixed cost. Furthermore, the falling cost of data analytics raises investors' leverage, which leads to higher losses during crises.

Date: Monday, October 22nd, 2018

Speaker: Henry Wolkowicz, University of Waterloo
Title: "Facial Reduction in Cone Optimization with Applications to Matrix Completions"
Time: 2:30 PM - 3:30 PM
Place: HA 968

Abstract: Strict feasibility is at the heart of convex optimization. This is needed for optimality conditions, stability, and algorithmic development. New optimization modelling techniques and convex relaxations for hard nonconvex problems have shown that the loss of strict feasibility is a much more pronounced phenomenon than previously realized. These new developments suggest a reappraisal. We describe the various reasons for the loss of strict feasibility, whether due to poor modelling choices or (more interestingly) rich underlying structure, and describe ways to cope with it and, in particular, "take advantage of it".

Date: Monday, October 15th, 2018

Speaker: Fernando Bernstein, Duke University
Title: "Competition between Two-Sided Platforms under Demand and Supply Congestion Effects"
Time: 2:30 PM - 3:30 PM
Place: HA 233

Abstract: This paper explores the impact of competition between platforms in the sharing economy. Examples include the cases of Uber and Lyft in the context of ride-sharing platforms. In particular, we consider competition between two platforms that offer a common service (e.g., rides) through a set of independent drivers to a market of customers. Each platform sets a price that is charged to customers for obtaining service provided by a driver. A portion of that price is paid to the driver that delivers the service. Both customers' and drivers' utilities are sensitive to the payment terms set by the platform and are also sensitive to congestion in the system (given by the relative number of customers and drivers in the market). We consider two possible scenarios. The first one, termed "single-homing," assumes that drivers work through a single platform. In the second scenario, termed "multi-homing" (or "multi-apping" as it is known in practice), drivers deliver service through both platforms. In both the single-homing and multi-homing scenarios, we study the equilibrium prices that emerge from the competitive interaction between the platforms and explore the supply and demand outcomes that can arise at equilibrium. We leverage the model to study some practical questions that have received significant press attention (and stirred some controversies) in the ride-sharing industry. The first involves the issue of surge pricing. The second involves the increasingly common practice of drivers choosing to operate on multiple platforms (multi-homing). We find that raising prices in response to a surge in demand makes platforms, drivers and customers better off than if platforms were constrained to charge the same prices that would arise under normal demand levels. We also compare the platforms', drivers', and customers' performance when all drivers either single-home or all multi-home. We find that, while individual drivers may have an incentive to multi-home, all players are worse off when all drivers multi-home.

Date: Thursday, October 11th, 2018

Speaker: Ilbin Lee, University of Alberta, School of Business
Title: "Is Modeling Heterogeneous Transition Patterns Beneficial for Sequential Decision-Making?"
Time: 12:30 PM - 1:30 PM
Place: HA 454

Abstract: In recent applications of Markov decision process (MDP), transition probabilities and rewards are often estimated from large-scale sequential data. In cases where sequences are obtained by simulating a single system, one can safely assume that all sequences follow the same

model. However, in health applications, sequential data is often collected from a population where each sequence corresponds to a person. Thus, there may be sub-populations that exhibit heterogeneous transition patterns. For example, in a large cohort with a certain disease, there may be patients whose disease status progresses faster than other patients. For such a group, estimating a separate transition probability matrix and applying the corresponding optimal treatment plan can improve their outcomes. In this talk, I formally define the benefit of modeling heterogeneity and derive a probabilistic bound on the benefit. The theoretical bound gives us intuition that as the transition models of sub-populations become more "similar" to each other, modeling heterogeneity becomes less beneficial. I present empirical analysis illustrating the theoretical result and show that we need big enough samples in order to identify the benefit of modeling heterogeneity. I also suggest an empirical method to estimate the benefit of modeling heterogeneity based on bootstrapping and numerically illustrate it.

Date: Monday, October 1st, 2018

Speaker: Saibal Ray, McGill University
Title: "Short-term Housing Rentals and Corporatization of Platform Pricing"
Time: 2:30 PM - 3:30 PM
Place: Henry Angus 968

Abstract: In this paper, we model an online platform for short-term rentals of housing assets and analyze how a change in the pricing strategy due to the growing influx of corporate players to the platform affects the performance and payoffs of stakeholders. The focus on housing assets also enables us to understand the implications for two socially important and related channels---hotels and long-term rentals. Many short-term online house rental platforms like Airbnb started off by facilitating Peer-to-Peer (P2P) matching between individual house owners and visitors with the aim of creating surplus for both supply and demand sides. However, recently, concerns have been raised about the growing prevalence of corporate players who control a significant number of housing assets on such platforms as well as about the increasing scarcity in the long-term rental market due to capacity being diverted to sharing platforms---a phenomenon we term as ``crowding-out''. In this paper, we show that corporatization potentially has: (i) a direct effect on the stakeholders of such platforms by changing its objective from maximizing the number of matches to maximizing profit; and (ii) that this change in the objective and the resulting impact on performance might be one of the driving factors behind the capacity scarcity seen in the long-term rental market.

We develop two game-theoretical models to capture the differences between P2P and corporate pricing strategies and employ the rational expectation concept to fully characterize the equilibrium decisions. We then compare the two to evaluate the impact of pricing strategies on the equilibrium performance of the platform (in terms of price, capacity availability, proportion of capacity occupied) and on the equilibrium payoffs of the stakeholders. We identify two factors that determine whether or not asset owners wish to join the sharing platform: (i) a short-term supply-demand mismatch between the two sides of the platform (i.e., owners and visitors) that creates a risk for the asset being unoccupied; and (ii) the relative profit potential of the platform

with respect to the long-term rental channel. We show that low occupancy risk and/or high price premium potential cause owners to join the short-term sharing platform, leaving fewer assets available to the long-term rental market. This is the underlying cause of the empirically observed phenomenon of crowding-out. Moreover, the change in the pricing strategy due to the entry of corporate players (i) increases the likelihood of crowding-out; (ii) leads to higher asset prices and availability and lower asset utilization; and (iii) increases profits for asset owners and decreases surpluses for consumers. We further extend our model to examine the impact of endogenous external channels and show that our results are relatively robust. Our results not only analytically justify the concerns expressed by policy makers regarding the potentially detrimental effects of the corporatization of online sharing platforms on the long-term rental market but also provide industry regulators/managers with means to cope with them, e.g., through a more competitive secondary channel.

*Joint work with Han Zhu and Mehmet Gumus

Date: Monday, September 24th, 2018

Speaker: Xuanming Su, Wharton School, University of Pennsylvania
Title: "Centralized Routing in Ride-hailing Networks"
Time: 2:30 PM - 3:30 PM
Place: Henry Angus 968

Abstract: In ride-hailing networks, where should drivers go after dropping off passengers? This paper studies first best and second best routing plans. The first best is the set of routes that minimizes the number of cars needed to satisfy a given demand pattern, and the second best incorporates additional incentive compatibility constraints for drivers to voluntarily follow the plan. The first best routing plan is not incentive compatible because drivers may prefer to wait around drop-off locations. Only after enough cars have accumulated (i.e., it takes sufficiently long to get the next ride) will they drive empty to other locations. Applying our model to ride-hailing data in San Francisco, we show that centralized routing can reduce the number of cars by about 30% while maintaining the same level of service

---

Date: Monday, September 17th, 2018

Speaker: George Shanthikumar, Purdue University
Title: "A Framework for Data Integrated Prescriptive Operations Management"
Time: 2:30 PM - 3:30 PM
Place: Henry Angus 968

Abstract: We provide a framework for data integrated modeling for prescriptive operations Management. Specific attention is shown to overcoming structural and statistical errors. This is achieved through operational statistics and objective operational learning which are built on the

basis of data integration and cross validation. We will illustrate how regularization in sample approximation approaches and data driven robust optimization with cross validation relates to operational objectives and operational statistics. Applications in pricing and revenue management, inventory control and staffing in service systems will be demonstrated.

* Joint work with Qi Feng, Krannert School of Management, Purdue University

---

Date: Friday, September 14th, 2018

Speaker: Shoshana Anily, Tel Aviv University
Title: "Full characterization of the non-negative core of some cooperative games"
Time: 12:30 PM - 1:30 PM
Place: Henry Angus 968

Abstract: We propose a method to design cost allocation contracts that help maintain the stability of strategic alliances among firms by using cooperative game theory. The partners of the alliance increase their efficiency by sharing their assets. We introduce a new sufficient condition for total balancedness of regular games, and a full characterization of their non-negative core. A regular game is defined by a finite number of resources owned by the players. The initial cost of a player is a function of the vector of quantities of the resources that the player owns. The characteristic function value of a coalition is a symmetric real function of the vectors of its members. Within this class we focus on centralizing aggregation games, meaning that the formation of a coalition is equivalent to aggregating its players into one artificial player whose cost is an intermediate value of the costs of the aggregated players. We prove that under a certain decreasing variation condition, a centralizing aggregation game is totally balanced and its non-negative core is fully characterized.

We present a few non-concave games in operations management that their non-negative core is fully characterized, by showing that they satisfy the conditions presented in this article.
The paper has been accepted for publication in *Naval Research Logistics*

Date: Monday, September 10th, 2018

Speaker: Rajesh Vijayaraghavan, AIS Division, UBC Sauder School of Business
Title: "Recognizing Loan Losses in Banks: An Examination of Alternative Approaches"
Time: 2:30 PM - 3:30 PM
Place: Henry Angus 968
Abstract: I investigate the accounting rules for loan loss recognition in banks. In June 2016 the FASB issued a new rule, effective in December 2019, that will replace current GAAP with a model that allows banks to use broader information to estimate loan loss allowances.
To empirically examine current GAAP and the new model, I exploit differences in the information sets allowed under the old and the new rules. Using a methodology that combines

micro data and machine learning techniques, I provide evidence that it is possible to construct a loan loss recognition model that outperforms the current GAAP without expanding the information set beyond that permitted under the current rule. I find that expanding this model's information set does not significantly improve its performance. My model's predicted allowances would have been materially larger at the outset of the financial crisis than actual reported bank estimates. The differences are due to that my model consistently assigns larger weights to certain input variables relative to current GAAP. I also find that weakly capitalized banks under provision relative to well capitalized banks. My results provide a novel method to examine aspects of the new accounting rule before it comes into effect. The findings suggest that the way information is used, rather than the use of broader information set improves the estimates of loan loss allowance.