

OPLOG DIVISION AND COE SEMINARS

September 12th, 2016 – July 20th, 2017

Date: Thursday, July 20th, 2017

Speaker: Alberto Todeschini, University of California - Berkley

Topic: "Artificial Intelligence and Disruption: What moves fast, what slowly, and why"

Time: 1:30PM - 2:30PM

Place: Henry Angus 968

Abstract: In the last 5 years deep neural network approaches in artificial intelligence have progressed greatly and have achieved state-of-the-art performance in many tasks. This has allowed unprecedented results in endeavors as different as machine translation, image recognition, self-driving vehicles, and generation of original visual art and music. While the rate of advance is increasing across the board and widespread disruption to many industries is a safe bet, there are technical, social, and economic reasons why the speed of this disruption is distributed very unevenly and will continue to remain so for the foreseeable future. In this talk Alberto Todeschini, lecturer in artificial intelligence at UC Berkeley, explores the main factors at play along these three dimensions (technical, social, economic) and shows which industries will be disrupted quickly, which won't, and why.

Date: Friday June 2nd, 2017

Speaker: Chung Piau Teo, National University of Singapore

Topic: "On the Design of Sparse but Efficient Structures in Operations"

Time: 11:45 AM - 1:00 PM

Place: Henry Angus 969

Abstract: The belief that “a little flexibility added at the right place can reap significant benefits for operations” has been widely accepted in the community. Unfortunately, despite the extensive literature on operational flexibility, to date no known methodology exists to guide managers to design (i.e., slightly flexible) and yet efficient operations.

We address this issue using a distributional robust approach to model the performance of a stochastic system under different process structures. We use the dual prices obtained from a related conic program to guide managers in the design process. This leads to a general solution methodology for the construction of efficient sparse structures for several classes of operational problems.

Our approach can be used to design simple yet efficient structures for workforce deployment and for any level of sparsity requirement to respond to deviations and disruptions in the operational environment. Furthermore, in the case of the classical process flexibility problem, our methodology can recover the k-chain structures that are known to be extremely efficient for this type of problem when the system is balanced and symmetric.

Date: Wednesday May 3rd, 2017

Speaker: Beril Toktay, Scheller College of Business, Georgia Institute of Technology

Topic: "Design Incentives under Collective Extended Producer Responsibility: A Network Perspective"

Time: 3:00 PM - 4:00 PM

Place: Henry Angus 969

Abstract: A key goal of Extended Producer Responsibility (EPR) legislation is to provide incentives for producers to design their products for recyclability. EPR is typically implemented in a collective system, where a network of recycling resources are coordinated to fulfill the EPR obligations of a set of producers, and the resulting system cost is allocated among these producers. Collective EPR is prevalent because of its cost efficiency advantages. However, it is considered to provide inferior design incentives compared to an individual implementation (where producers fulfill their EPR obligations individually). In this paper, we revisit this assertion and investigate its fundamental underpinnings in a network setting. To this end, we develop a new biform game framework that captures producers' independent design choices (non-cooperative stage) and recognizes the need to maintain the voluntary participation of producers for the collective system to be stable (cooperative stage). This biform game subsumes the network-based operations of a collective system and captures the interdependence between producers' product design and participation decisions. We then characterize the manner in which design improvement may compromise stability and vice versa. We establish that a stable collective EPR implementation can match and even surpass an individual implementation with respect to product design outcomes. In particular, we show that when the processing technology efficiency and product recyclability are substitutes (complements), a recycling network where processor capacity pooling leads to sufficiently low (high) cost reduction will lead to superior designs in the collective system and maintain its stability, and we propose cost allocation mechanisms to achieve this dual purpose.

Keywords : Extended Producer Responsibility, Design for Environment, Biform Games, Recycling Network Operations

Date: Monday April 24th, 2017

Speaker: Tinglong Dai, Carey Business School, Johns Hopkins University

Topic: "Clinical Ambiguity and Conflicts of Interest in Interventional Cardiology Decision-Making"

Time: 2:30 PM - 3:30 PM

Place: Henry Angus 969

Abstract: Conventional wisdom suggests more-precise diagnostic testing will help reduce overtreatment. However, the literature rarely recognizes that the decision about whether to perform these tests is itself endogenous. In this paper, we study the decision to perform advanced tests in a catheterization laboratory. One key driver of inappropriate usage of percutaneous coronary intervention (PCI) procedures is physician subjectivity in “eyeballing” a coronary angiogram. Advanced intracoronary tests such as fractional flow reserve (FFR) provide more precise and objective measures of PCI appropriateness, yet the decision to perform them is endogenous and not immune to clinical ambiguity and conflicts of interest associated with the angiogram- guided decision-making.

We find that under a high conflict-of-interest level, the physician would perform the advanced test only when the visually assessed angiographic stenosis is sufficiently large, because of a financial disincentive: performing the advanced test may lower PCI revenue if the test results argue against the procedure. Surprisingly, despite this disincentive, a more revenue-driven physician can be more inclined to perform the advanced test. In addition, under an intermediate conflict-of-interest level, the physician may either perform the advanced test for all cases meriting consideration of a PCI procedure or never perform it. These results help to inform policy initiatives in several ways: (1) direct intervention in the conflict-of-interest level may backfire and lead to more PCI overuse; (2) reducing the risk of the advanced test has a behavior-inducing effect, yet a modest risk reduction may lower patient welfare; and (3) using parameter values based on the U.S. interventional cardiology practice, we numerically illustrate the impact of increasing the FFR reimbursement rate on reducing overstenting and improving patient welfare.

Much effort in reducing overtreatment has been expended on “the first mile” (developing new diagnostic tools and therapies), yet “last mile problems” in this domain, such as an interventional cardiologist’s decision to perform the FFR test, have been little explored. Our research highlights the central importance of endogeneity of advanced intracoronary testing in policymaking surrounding interventional cardiology.

Date: Monday April 10th, 2017

Speaker: Soroush Saghaian, Harvard Kennedy School, Harvard University

Topic: "Workload Management in Telemedical Physician Triage and Other Knowledge-Based Service Systems"

Time: 2:30 PM - 3:30 PM

Place: Henry Angus 969

Abstract: Telemedical Physician Triage (TPT) is an example of a Hierarchical Knowledge-Based Service System (HKBSS), in which a second level of decision agent (telemedical physician) renders a decision on cases referred to him/her by the primary level agents (triage nurses). Managing the speed-versus-quality tradeoff in such systems presents a unique challenge because of the interplay between agent knowledge and flow of work between the two levels. We develop a novel model of agent knowledge, based on the beta distribution, and deploy it in a Partially Observable Markov Decision Process (POMDP) model to describe the optimal policy for deciding which cases (patients) to refer to the second level for further evaluation. We show that this policy has a monotone control-limit structure that reduces the fraction of decisions made at the upper level as workload increases. Because the optimal policy is complex, we use structural insights from it to design two practical heuristics. These heuristics enable an HKBSS to respond efficiently to workload shifts by adjusting the criteria for referring decisions to the upper level based on partial real-time queue length information. Finally, we conduct analytic and numerical analyses to derive insights into the management of a TPT system. We find that (1) the telemedical physician should evaluate more patients as congestion in the emergency room waiting area increases, (2) training that improves accuracy of the physician and/or nurses can be effective even if it only does so for a single patient type, but training that improves consistency must do so for all patient types to be effective, and (3) patient classification in triage should consider environmental and operational conditions in addition to the patient's medical condition.

Date: Monday, April 3rd, 2017

Speaker: Donald Lee, Yale School of Management

Topic: “Can arrival processes be modeled by simple waves?”

Time: 2:00 PM - 3:00 PM

Place: David Lam - DL 125

Abstract: Many real world arrival processes exhibit strong cyclic behaviour. For example, emergency department (ED) demand often peaks around midday and drops off at night. In this talk we investigate whether arrival rates can be approximated by a simple sum of sinusoids. Such models can potentially bridge the gap between analytic modelling and high resolution simulations: They are simple enough to be tractable for the former, but may also be accurate enough to use for the latter for prescriptive capacity planning.

We discuss some unexpected estimation pitfalls that can arise when estimating the frequency components in the arrival rate using intuitive heuristics. Motivated by these concerns, we propose a simple graphical solution that has surprisingly good performance guarantees. To aid capacity planning, we apply the procedure to patient arrivals data from the Yale ED.

Joint work with Ningyuan Chen (HKUST) and Sahand Negahban (Yale).

Date: Monday March 27th, 2017

Speaker: Pengyi Shi, Krannert School of Management, Purdue University

Topic: "Inpatient Bed Overflow: An Approximate Dynamic Programming Approach"

Time: 2:30 PM - 3:30 PM

Place: Henry Angus 969

Abstract: Hospital inpatient beds are usually grouped into several wards, with each ward being designated to serve patients from certain “primary” specialty. When a patient waits excessively long in the emergency department before a primary bed becomes available, hospital managers may have to overflow the patient to a non-primary bed though it is undesirable. To aid the decision making on when and where to overflow patients, we model hospital inpatient flow as a multi-class, multi-pool parallel-server queueing system and formulate a discrete-time, infinite-horizon average cost Markov decision process (MDP). This MDP incorporates many realistic and important features such as the patient arrival and discharge patterns depending on the times of the day.

To overcome the curse-of-dimensionality of this formulated MDP, we resort to the approximate dynamic programming (ADP) technique. A critical part in the algorithm is to choose appropriate basis functions to approximate the relative value function. Using a novel combination of fluid control and single-pool approximation, we develop an analytical form to approximate the relative value function at the midnight, which then guides the choice of the basis functions for different times of the day. We demonstrate, via numerical experiments in realistic hospital settings, that our proposed ADP algorithm is remarkably effective in finding good overflow policies that can significantly improve various system performance over some commonly used overflow strategies.

This is a joint work with Jim Dai at Cornell University.

Date: Monday, March 20th, 2017

Speaker: Philipp Afeche, Rotman School of Management, University of Toronto

Topic: "Ride-Hailing Networks with Strategic Drivers: The Impact of Platform Control Capabilities on Performance"

Time: 2:30 PM - 3:30 PM

Place: Henry Angus 969

Abstract: This work is motivated by the emergence of ride-hailing platforms such as Uber, Lyft and Gett that match demand (passengers) with service capacity (drivers) over a geographically dispersed network. This matching problem is complicated by two challenges. (i) There are significant demand imbalances in the network. (ii) Drivers are self-interested and behave strategically in deciding whether to join, and if so, how to reposition (route) themselves when not transporting passengers.

To address these challenges we study the value of two operational controls, demand-side admission control and supply-side repositioning control, on the performance of a revenue-maximizing ride-hailing platform. Considering a fluid model of a two-location network in a game-theoretic framework, we characterize the system equilibrium under three operating regimes, ranging from minimal control to centralized admission and repositioning control. These results contribute novel insights on the interplay between the platform's admission control and the drivers' strategic routing decisions.

We also quantify the impact of control capabilities on the platform revenue, the capacity and the per-driver profits. The value of control is largest at moderate utilization and increases with demand imbalances.

(Joint work with Costis Maglaras and Zhe Liu.)

Centre for Operations Research and Analytics in Health Seminar

Date : Tuesday, March 14th, 2017

Speaker: Ingeborg Bikker, University of Twente, The Netherlands

Topic: "Advance Multi-Appointment Scheduling with Resource Compatibility Restrictions"

Time: Noon - 1.00 PM

Place: David Lam - DL 414

Abstract: To provide radiotherapy patients with timely access to care, we consider an advance scheduling problem in which patients require a series of radiation appointments on one of several treatment units. Unit requirements vary per patient type. In practice, appointments are typically planned in the first available slot of a suitable unit, leaving no space for urgent patients.

We formulate a Markov decision process to identify good policies for scheduling appointments on units, while reducing access times in a cost-effective manner. We use value function approximation and column generation to derive approximate solutions, and simulate their performance for a case of the British Columbia Cancer Agency.

Joint work with: A. Sauré (University of Ottawa), X. Ma (British Columbia Cancer Agency), N. Horvath (British Columbia Cancer Agency), S. Tyldesley (British Columbia Cancer Agency), M.L. Puterman (University of British Columbia).

Date: Monday March 6th, 2017

Speaker: Yichuan (Daniel) Ding, OPLOG Division, UBC Sauder School of Business

Topic: "A Fluid Model for an Overloaded Bipartite Queueing System with Heterogeneous Matching Utilities"

Time: Noon - 1:30 PM

Place: Henry Angus 233

Abstract: We consider a bipartite queueing system (BQS) with multiple types of servers and customers, where different customer-server combinations may generate different utilities. Whenever a server is available, it serves the customer with the highest index, which is the sum of a customer's waiting index and the matching index. We assume that the waiting index is an increasing function of a customer's waiting time and the matching index depends on both the customer and the server's types. We develop a fluid model to approximate the behavior of such a BQS system, and show that the fluid limit process can be computed over any finite horizon. We prove that the fluid limit process converges to a unique steady state, which can be efficiently computed. These results enable the policy designer to predict the behavior of a BQS when an indexing priority rule takes the form as we have defined, and choose an indexing formula that optimizes a given set of performance metrics. We illustrate the use of our machinery by analyzing the public housing assignment in the city of Pittsburgh using a real data set.

Date: Monday, February 6th, 2017

Speaker: Michael Jong Kim, University of Toronto

Topic: "Deconstructing the Exploration-Exploitation Trade-off for Bayesian Bandits"

Time: 2:30 PM - 4:30 PM

Place: Henry Angus 969

Abstract: Bayesian bandit problems study the fundamental exploration vs. exploitation trade-off that arises in many business applications. That is, how to choose between choices that seem profitable now, with less explored choices that may yield higher profits later. It is known that the optimal policy for Bayesian bandit problems is the Gittins index policy. Computing the Gittins index for these problems however is intractable. This paper makes two contributions. First, we introduce the concept of an approximate learning trajectory as a new approach to approximating the dynamics of future learning, which is motivated by the Bernstein-von Mises Theorem. Using this approximation, we show how the expected cost-to-go terms in the dynamic programming equations for Bayesian bandits can be easily obtained, which allows for an efficient computation of an approximate Gittins index. We prove that the approximate Gittins index is asymptotically optimal in that it approaches the true Gittins index in certain limiting signal and prior standard

deviation regimes. Second, we discuss how easy computation of the approximate Gittins index allows us to decompose the Gittins index into tractable terms that characterize how the values of exploration and exploitation depend on primitives of the problem. In particular, we show that the Gittins index can be expressed as the sum of the estimate of its expected reward under the posterior (value of exploitation) and a ``boost" that we interpret as the value of exploration. The ``exploration boost" is equal to a measure of upside potential of a decision net a penalty that is decreasing in the quality of the signal. Intuitively, this tells us that a decision is only worth exploring if the upside is high and the signals are of sufficiently high quality that we can learn the true value of the expected reward within the effective time horizon.

Date: Thursday, February 2nd, 2017

Speaker: Xiaoshan Peng, Booth School of Business, University of Chicago

Topic: "Managing the Callback Option under Arrival Rate Uncertainty"

Time: 3:30 PM - 5:00 PM

Place: Henry Angus 966

Abstract: We study how to manage the callback option effectively to mitigate congestion due to temporary surges in the arrivals to a call center. The call arrival process can be an arbitrary point process, allowing uncertainty and temporary surges in the arrival rate, provided that the system is stable. However, particular attention will be paid to the Poisson process with the Cox-Ingersoll-Ross (CIR) process as its stochastic intensity both in our model development and numerical results because of its practical importance, although our theoretical results hold for any arbitrary point process. When a customer arrives, the call center manager reviews the system state and decides whether to keep him in the online queue or to offer the callback option. For each customer in the online queue, she incurs a waiting cost of h per time unit. Similarly, whenever she routes a customer to an offline queue (for a callback later), she incurs a one-time penalty of p . Initially, we allow complete foresight policies that look into the entire future. We first study the case where all customers are willing to accept a callback offer. A simple lookahead policy that looks into the future for the next p/h time units is pathwise optimal among the complete foresight policies. Next, we consider the setting where some customers may reject the callback offer. We show that a modified lookahead policy that looks into the future arrivals and service completion times for the next p/h time units and uses the current number of customers in the system who previously rejected a callback offer (but does not look into the accept/reject decisions of future customers) is pathwise optimal among the complete foresight policies. Building on the insights gleaned from the optimal lookahead policies, we also propose a non-anticipating policy, referred to as the line policy, to decide when to offer the callback option. Lastly, we conduct a simulation study using a dataset from a US bank call center which shows that the line policy has excellent performance.

Date: Wednesday, February 1st, 2017

Speaker: Vijay Kamble, Stanford University

Topic: "Matching while Learning"

Time: 2:30 PM - 4:00 PM

Place: Henry Angus 968

Abstract: We consider the problem faced by a service platform that needs to match supply with demand, but also to learn attributes of new arrivals in order to match them better in future. We introduce a benchmark model with heterogeneous workers and jobs that arrive over time. Job types are known to the platform, but worker types are unknown and must be learned by observing match outcomes. Workers depart after performing a certain number of jobs. The payoff from a match depends on the pair of types and the goal is to maximize the steady state rate of accumulation of payoff.

Our main contribution is a complete characterization of the structure of the optimal policy in the limit that each worker performs many jobs. The platform faces a tradeoff for each worker between myopically maximizing payoffs (exploitation) and learning the type of the worker (exploration). This creates a multitude of multiarmed bandit problems, one for each worker, coupled together by the constraint on the availability of jobs of different types (capacity constraints). We find that the platform should estimate a shadow price for each job type, and use payoffs adjusted by these prices, first to determine its learning goals, and then, for each worker, (i) to balance learning with payoffs during the “learning phase”, and (ii) to myopically match after it has achieved its learning goals during the “exploitation phase”.

Date: Friday, January 27th, 2017

Speaker: Antoine Desir, Columbia University

Topic: "A markovian approach for choice modeling and assortment optimization"

Time: 2:30 PM - 4:00 PM

Place: Henry Angus 968

Abstract: Which set of products should be offered to arriving consumers to maximize expected revenue? This is a core revenue management problem known as assortment optimization which applies to a wide variety of settings (from retail to e-commerces). Discrete choice model theory offers a way to mathematically model the substitution behavior of consumers and provide a key ingredient for this problem. Many choice models have been proposed in the literature, introducing a fundamental tradeoff between model expressiveness and computational complexity. In particular, the assortment optimization problem is notoriously hard for general choice models. In this talk, we look at a new framework which tries to strike a good balance between expressiveness and tractability. In particular, the substitution behavior of consumers is

modeled as transitions in a Markov chain. By doing so, this new model helps alleviate the Independence of Irrelevant Alternatives (IIA) property, a well-known limitation of the popular multinomial logit model. Moreover, it provides a good approximation to the class of random utility models. We show that not only this model has a great predictive power, it is also tractable from a computational perspective. In particular, we give a algorithm framework to derive efficient algorithm for different variants of the assortment optimization problem.

Date: Monday, January 9th, 2017

Speaker: Michael Freeman, Judge School of Business, University of Cambridge

Topic: "Economies of Scale and Scope in Hospitals: An Empirical Study of Volume Spillovers"

Time: 2:30 PM - 4:00 PM

Place: Henry Angus 969

Abstract: General hospitals across the world are becoming larger (i.e. admitting more patients each year) and more complex (i.e. offering wider portfolios of services to higher acuity patients with more diverse care needs). Although prior work has shown that increased volume is positively associated with patient outcomes, it is less clear how volume affects costs in these complex organizations. This paper investigates this relationship using panel data for 14 service lines comprising both elective and emergency admissions across 130 hospitals in England over a period of nine years. Although we find significant economies of scale for both elective and emergency admissions, we also find evidence of negative spillovers across the two admission types, with increased elective volume at a hospital being associated with an increase in the cost of emergency care. Furthermore, for emergency admissions we find evidence of positive spillovers across service lines – increased emergency activity in one service line is associated with lower costs of emergency care in other service lines. By contrast, we find no evidence of such spillovers across service lines for elective admissions. Our findings have implications for individual hospitals and for the organization of regional hospital systems. Specifically, at the hospital level our findings suggest that growth strategies that target elective patients may have unintended negative productivity implications for emergency services. At the regional level, our findings offer support for the reorganization of regional hospital systems toward general hospitals that focus on the provision of emergency care across a full range of services, complemented by high-volume specialist hospitals that focus on elective services in a single service line.

Date: Thursday January 5th, 2017

Speaker: Ali Aouad, MIT

Topic: "Revenue Management in Face of Choice Heterogeneity"

Time: 12:30 PM - 2:00 PM

Place: Henry Angus 969

Abstract: Modern-day applications in e-commerce and brick-and-mortar retailing involve complex customer choice behaviors. Modeling this choice heterogeneity strikes a delicate balance between explaining large-scale data and prescribing efficient operational policies. At the strategic level, for the assortment selection problem, we propose a "consider-then-choose" modeling approach, borne out by the marketing literature. Experiments on a large purchase panel dataset demonstrate the strong predictive power of our models against common benchmarks. We develop a dynamic programming framework and show that many empirically vetted assumptions on how customers consider and then choose lead to tractable optimization models. Our algorithm dominates state-of-the-art commercial solvers in several regimes. Further, at the operational level, we study joint assortment and inventory management where customers show a dynamic substitution behavior. We derive the first provably good policies by revealing hidden submodular-like structure. Our approach is an order of magnitude faster than existing heuristics and increases revenue by 6% to 12% in experiments.

This work is based on several papers jointly with Profs. Vivek Farias, Retsef Levi and Danny Segev.

Date: Wednesday January 4th, 2017

Speaker: Anton Braverman, Cornell University

Topic: "Steady-state diffusion approximations in service systems: engineering solutions and error bounds"

Time: 2:30 PM - 4:00 PM

Place: Henry Angus 969

Abstract: Steady-state diffusion approximations are commonly used to approximate models of large scale service systems. In this talk I will introduce a framework based on Stein's method that can be used a) as an engineering solution for generating good steady-state approximations and b) as a mathematical tool for establishing error bounds for these approximations. These approximations are often universally accurate in multiple parameter regions, from underloaded, to critically loaded, to overloaded (when customers abandon).

As a running example, I will use the many server queue with customer abandonment and phase-type service time distribution, which is a fundamental building block in service system models.

CENTRE FOR TRANSPORTATION STUDIES

Date: Friday December 16th, 2016

Speaker: Stefan K. Sløk-Madsen, Copenhagen Business School

Topic: "Commercial Capabilities and Entrepreneurial Value Capturing in Dynamic Maritime Markets: The Case of North Sea Oil Service Companies"

Time: Noon - 1:30 PM

Place: Henry Angus 969

Abstract: This paper develops and tests a theory of entrepreneurial value capturing in maritime markets. The framework is argued to be applicable in all maritime fields and other fields with similar attributes, but is specifically tested on Oil Service companies operating in the North Sea region. This specific empirical application, however, mirrors general maritime concerns of derived demand and high capital intensity and knowledge specificity. The paper introduces the construct of Commercial Capabilities and models the relevance of such a subset of firm capabilities from an entrepreneurial dynamic market process view; building on Kirznerian alertness, Hayekian capital heterogeneity, and Knightian uncertainty. The theory helps explain value capturing from a firm perspective but also subsequent new firm creation or value loss. The model is tested and relevant managerial implications, as well as reflections on further research, presented. The paper is written so as to be relevant for maritime economists, in moving an emergent maritime entrepreneurship agenda forward, but also mainstream entrepreneurship research, and applied Austrian economics, in attempting to test theory in an industry setting only sparingly engaged by non-maritime researchers.

Date: Wednesday December 14th, 2016

Speaker: Anke van Zuylen, College of William and Mary

Topic: "Improved Tours and Paths for the Traveling Salesman: Classical Tools and Recent Advances"

Time: 3:30 PM - 5:00 PM

Place: Henry Angus 968

Abstract: The traveling salesman problem (TSP) is perhaps the most famous problem in combinatorial optimization: given a set of cities and the distances for traveling between each pair of cities, the goal of the problem is to find the shortest tour that visits each city once and returns to its starting point. Besides many applications in logistics, algorithms for the TSP have been used in areas as diverse as finding the most efficient way to drill circuit boards to mapping genomes. Finding an optimal tour is known to be hard; nevertheless, the past decades have seen great progress in solving real-world instances of the TSP, and new techniques for the TSP have been driving progress for other combinatorial optimization problems as well. In this talk, we will describe a classical algorithm for the TSP by Christofides, which finds tours that are close to optimal. We then discuss two key new ideas for getting better solutions if the salesman has a different starting and ending point and wants to visit all other cities on the way. These ideas lead to new algorithms that find solutions that are closer to optimal, which can be seen as steps towards resolving a famous conjecture.

Date: Friday, November 25th, 2016

Speaker: Leon Valdes, MIT Sloan School of Business

Topic: "Supply Chain Visibility and Social Responsibility: Investigating Consumers' Behaviors and Motives"

Time: Noon - 1:00 PM

Place: Henry Angus 966

Abstract: Consumers today want to know more about where and how the products they purchase are being made. To create transparency requires a company to both gain visibility into its supply chain and disclose information to consumers. We focus on the dimension of visibility and investigate when companies can benefit from greater supply chain visibility. To do so, we design an incentivized human-subject experiment to study two key questions: (i) When does supply chain visibility impact consumers' valuations of a company's social responsibility practices in its upstream supply chain? (ii) What roles do consumers' reciprocal motives and prosocial orientations play in affecting their valuations? Our results show that consumers value greater visibility either when workers in the upstream supply chain are disadvantaged or when consumers use uncertainty as an excuse not to pay for social responsibility (i.e., they exhibit a self-serving bias). In addition, we observe that high prosocial consumers do not exhibit strong reciprocal motives. Conversely, reciprocal motives significantly increase low prosocial consumers' valuations under high visibility. Our work adds to the experimental literature on transparency in social responsibility (which has primarily studied disclosure) by examining the equally important but understudied dimension of visibility.

This is joint work with Professors Tim Kraft (UVA) and Karen Zheng (MIT).

Date: Monday, November 21st, 2016

Speaker: Chris Ryan, University of Chicago, Booth School of Business

Topic: "Monotonicity of optimal contracts without the first-order approach"

Time: 2:30 - 3:30 PM

Place: David Lam - DL 125

Abstract: We develop a simple sufficient condition for an optimal contract of a moral hazard problem to be monotone in the output signal. Existing results on monotonicity require conditions on the output distribution (namely, the monotone likelihood ratio property (MLRP)) and additional conditions to ensure that agent's decision is approachable via the first-order approach of replacing that problem with its first-order conditions. We know of no positive monotonicity results in the setting where the first-order approach does not apply. Indeed, it is well-documented that when there are finitely many possible outputs, and the first-order approach does not apply, the MLRP alone is insufficient to guarantee monotonicity. However, we show that when there is an interval of possible output signals, the MLRP does suffice to establish monotonicity under additional technical assumptions that do not ensure the validity of the first-order approach. To establish this result we examine necessary optimality conditions for moral hazard problems using a penalty function approach. We then manipulate these conditions and provide sufficient conditions for when they coincide with a simple version of the moral hazard problem with only two constraints. In this two-constraint problem, monotonicity is established directly via a strong characterization of its optimal solutions.

This is joint work with Rongzhu Ke.

Date: Monday, November 7th, 2016

Speaker: So Yeon Chun, McDonough School of Business, Georgetown University

Topic: "Setting the Optimal Value of Loyalty Points"

Time: 2:00 PM - 3:00 PM

Place: Henry Angus 969

Abstract: A loyalty program introduces a new currency---the points---through which customers transact with a firm. We study the problem of optimally setting the monetary value of points, i.e., pricing in this new currency, in a multi-period setting. We first show that point pricing is different from cash pricing primarily due to the way points are accounted for, as liabilities on the firm's balance sheet. This introduces subtle channels through which the firm's decisions affect its financial performance, and exacerbates the importance of certain managerial considerations such as taxation or earnings smoothing incentives.

We characterize the optimal cash and point pricing policies, and find that they mimic ``base-

stock, list price" policies in inventory management. In particular, point prices/values are always set so that the total value of points reaches a "base-stock" target, and cash prices are charged so as to maximize the firm's cash flows under the optimal loyalty point values. Under a profit-maximizing policy, the total value of loyalty points is set independently of the firm's realized financial performance. In contrast, we find that under the aforementioned managerial considerations, the optimal value of points becomes state-dependent, and is increasing (decreasing) under strong (weak) operating performance. In this sense, our work shows that loyalty points can act as a hedging tool against uncertainty in future performance, providing a new rationale for their existence, even in the absence of competition.

Date: Monday, October 24th, 2016

Speaker: Hubert Pun, Ivey Business School, Western University

Topic: "The Competing with Copycats when Customers Are Strategic"

Time: 2:30 PM - 3:30 PM

Place: Henry Angus 969

Abstract: In this paper, we use a two-period game theoretical model to examine the decisions of a manufacturer and a copycat firm who are competing for strategic customers. The manufacturer decides on the amount of its market expansion advertising investment in the first period and on its pricing strategy in both periods. Advertising increases the "size of the pie," but eventually the manufacturer may end up inadvertently sharing the benefits with the copycat. After the first period, the copycat makes a market-entry decision, and, if it opts to enter, it also decides on a pricing strategy. The customers are forward-looking strategic, and they decide whether or not to buy, when to buy, and which product to buy. We find that, interestingly, lower quality levels of the manufacturer's product may increase the manufacturer's prices and profit. Moreover, the manufacturer may be worse off when customers are more likely to purchase its product immediately rather than wait for a price reduction or for the copycat's product. Finally, the copycat may be worse off when customers withhold their purchases in the first period in anticipation of the possibility of copycat product becoming available in a later period.

Date: Monday, October 17th, 2016

Speaker: Bruce Shepherd, McGill University

Topic: "Confluent Flows, Rooted Clustering and Stable Matchings"

Time: 2:30 PM - 3:30 PM

Place: Henry Angus 969

Abstract: We describe a model for network flows which only use edges of a tree. We examine the motivations and then develop new formulations which are needed to present a Maximum Flow Theorem. We will see that our algorithmic solution has connections to the classical stable matching problem.

Date: Monday, October 3rd, 2016

Speaker: Amr Farahat, Olin Business School, University of Washington at St. Louis

Topic: "The Multi-Product Newsvendor Problem with Customer Choice"

Time: 2:30 PM - 3:30 PM

Place: Henry Angus 969

Abstract: We address the multi-product newsvendor problem under a general specification of customer choice behavior. We develop a methodology that yields upper bounds on the optimal value as well as feasible inventory solutions. The methodology is based on an approximate Jordan decomposition of the state transition matrix. Two specializations of the methodology are presented: one leads to a decomposition by customer into a sequence of assortment optimization problems and the second leads to a decomposition by product into a collection of independent newsvendor problems. We conduct computational experiments and find that the proposed methodology outperforms existing bounds.

Joint work with Joonkyum Lee.

Date: Monday, September 26th, 2016

Speaker: Ilan Lobel, Stern School of Business, New York University

Topic: "Feature-based Dynamic Pricing"

Time: 2:30 PM - 3:30 PM

Place: Henry Angus 969

Abstract: We consider the problem faced by a firm that receives highly differentiated products in an online fashion and needs to price them in order to sell them to its customer base. Products are described by vectors of features and the market value of each product is linear in the values of the features. The firm does not initially know the values of the different features, but it can learn the values of the features based on whether products were sold at the posted prices in the past. This model is motivated by a question in online advertising, where impressions arrive over time and can be described by vectors of features. We first consider a multi-dimensional version of binary search over polyhedral sets, and show that it has exponential worst-case regret. We then propose a modification of the prior algorithm where uncertainty sets are replaced by their Lowner-John ellipsoids. We show that this algorithm has a worst-case regret that is quadratic in the dimensionality of the feature space and logarithmic in the time horizon.

Joint work with Maxime Cohen and Renato Paes Leme.

Date: Monday, September 19th, 2016

Speaker: Dorit Hochbaum, IEOR department, University of California, Berkeley

Topic: "Combinatorial algorithms for clustering, image segmentation and data mining"

Time: 2:30 PM - 3:30 PM

Place: Henry Angus 969

Abstract: We present a model for clustering which combines two criteria: Given a collection of objects with pairwise similarity measure, the problem is to find a cluster that is as dissimilar as possible from the complement, while having as much similarity as possible within the cluster. The two objectives are combined either as a ratio or with linear weights. The ratio problem, and its linear weighted version, are solved by a combinatorial algorithm within the complexity of a single minimum s,t-cut algorithm. This problem (HNC) is closely related to the NP-hard problem of normalized cut that is often used in image segmentation and for which heuristic solutions are generated with the eigenvector technique (spectral method).

The relationship of HNC to normalized cut is generalized to a problem we call "q-normalized cut". It is shown that the spectral method that solves for the Fiedler eigenvector of a related matrix is a continuous relaxation of the problem. In contrast, the generalization of the combinatorial algorithm solves a discrete problem resulting from a relaxation of a single sum constraint. We study the relationship between these two relaxations and demonstrate a number of advantages for the combinatorial algorithm. These advantages include a better approximation, in practice, of the normalized cut objective for image segmentation benchmark problems.

HNC can, and has been used, as a supervised or unsupervised machine learning technique. It has been used for data mining, and its comparison to leading machine learning techniques on datasets selected from the UCI data mining benchmark and other benchmarks indicates that its use of pairwise comparisons is powerful in improving quality of clustering. Time permitting, we will discuss methods that have been employed to make the HNC algorithms family scalable and applicable for large scale data sets.

Date: Monday, September 12th, 2016

Speaker: Lawrence Robinson, Johnson School of Management, Cornell University

Topic: "Appointment Scheduling with Strategic Patients"

Time: 2:30 PM - 3:30 PM

Place: Henry Angus 969

Abstract: The appointment scheduling literature generally assumes that patients will arrive promptly. In this paper, we instead model the doctor as a Stackelberg leader who sets the appointment times. Her patients act as Stackelberg followers, who then choose a target arrival time conditional on these appointment times. Patients wish to minimize their expected waiting time, but need to maintain a good working relationship with the doctor, which we model by their

incurring a psychic “guilt aversion cost” if they make her wait. We investigate the effect that strategic patient behavior has on the physician’s optimal appointment schedule.

* Joint work with Tava L. Olsen, University of Auckland, New Zealand.
