

# INFORMS Revenue Management and Pricing Section Conference

---

*Scotiabank Centre*

*40 King St West, Second Floor*

*Toronto Ontario*

---

# 2018

Thank you to our Sponsors:



Scotiabank Centre  
for Customer Analytics



## Organizing Committee



**Dr. Yuri Levin**  
*Professor & Stephen J.R. Smith Chair of Analytics;  
Director, Queen's Master of Management Analytics  
program; and Director, Scotiabank Centre for  
Customer Analytics*

*yuri.levin@queensu.ca*



**Dr. Mikhail Nediak**  
*Associate Professor & Distinguished Professor of  
Operations Management  
Smith School of Business  
mikhail.nediak@queensu.ca*



**Dr. Anton Ovchinnikov**  
*Associate Professor & Distinguished Professor of  
Management Science & Operations Management  
Smith School of Business  
anton.ovchinnikov@queensu.ca*

## Program Committee



**Yossi Avivi**  
*Professor of Management  
Recanati Business School  
yaviv@post.tau.ac.il*



**Fredrik Odegaard**  
*Associate Professor, Management Science  
Ivey Business School  
fodegaard@ivey.ca*



**Georgia Perakis**  
*William F. Pounds Professor of Management  
MIT Sloan School of Management  
georgiap@mit.edu*



**Ming Hu**  
*Professor of Operations Management  
Rotman School of Management  
ming.hu@rotman.utoronto.ca*



**Sumit Kunnumkal**  
*Associate Professor of Operations Management  
Smith School of Business  
sk162@queensu.ca*



**Philipp Afèche**  
*Operations Management  
Rotman School of Management  
philipp.afeche@rotman.utoronto.ca*



**William Cooper**  
*Professor, Industrial and Systems Engineering  
University of Minnesota  
billcoop@umn.edu*



Helping  
retail banks  
improve service  
by leverage pricing  
analytics and data  
science to become  
more customer centric.



Intelligent Pricing for Better Banks

Learn how you can help banks transform customer engagement at

[www.nomissolutions.com/careers](http://www.nomissolutions.com/careers)

Schedule Day One – Thursday June 21, 2018

	Room A (MIW)	Room B	Room C	Room D
8:00 to 8:15	<b>WELCOME – Room B</b>			
8:15 to 9:45	<b>Ridesharing</b>	<b>Behavioural RM</b>	<b>RM in Online Platforms 1</b>	<b>Demand Learning</b>
	#88: Terry Taylor; Ridesharing Platforms: Competition and Automated Vehicles	#90: Gonzalo Romero; Revenue Management with Repeated Customer Interactions	#173: Wenchang Zhang; Learning, Welfare, and Profits in Service Platforms	#70: Arnoud den Boer; Dynamic Pricing with Demand Learning and Reference Effects
	#141: Pengyu Qian; The Value of State Dependent Control in Ridesharing Systems	#144: Mohamed Mostagir; Information Design in Dynamic Contests: A Lab Experiment	#60: Francisco Castro; Surge Pricing and its Spatial Supply Response	#125: Theja Tulabandhula; A Realistic Online Algorithm to Learn Buyer Behavior
	#137: Xuanming Su; Centralized Location Control in Ridesharing Networks	#149: Yang Chen; Quantifying mileage-runs	#158: Daniela Saban; Optimal commissions and subscriptions in networked markets	#171: Victor Araman; Crowdvoting New Product Introduction
	No Speaker	#56: So Yeon Chun; Should I Pay for This Purchase or Redeem Points? Effects of Loyalty Program Design on Consumer Decisions to Redeem Points	#129: Vahideh Manshadi; Two-stage Pandora's Box for Product Ranking	#104: Yifan Feng; Learning Customer Preferences From Crowdvoting
9:45 to 10:00	<b>BREAK</b>			
10:00 to 11:00	<b>Keynote: Ivan Sergienko, Element AI – Rooms A, B, C</b>			
11:00 to 11:15	<b>BREAK</b>			
11:15 to 12:45	<b>Empirics of Marketplaces</b>	<b>RM in Industry</b>	<b>RM and Economics/Finance</b>	<b>Airline RM</b>
	#115: Bo Cowgill; Competition and Specificity in Market Design: Evidence from Geotargeted Advertising	#150: Warren Lieberman; Pricing for Profit	#44: Diego Escobari; Price Discrimination and Focal Points for Tacit Collusion: Evidence from the Airline Industry	#101: Hideaki Takagi; Extension of Littlewood's rule to a multi-period model with possible standby customers
	#102: Chiara Farronato; Consumer Reviews and Regulation: Evidence from NYC Restaurants	#55: Xiao Huang; Price to Gain or Price to Retain? An Empirical Study of Hotel Pricing and Customer Cancellation Behavior	#106: Yaarit Even; Efficiency and Information Aggregation in Heterogeneous Markets	#61: Fan You; Finite-Horizon Approximate Linear Programs for an Infinite-Horizon Revenue Management Problem
	#123: Jiding Zhang; Flexibility and Relationships in Online Marketplaces: Prescriptions for Pricing and Matching	#65: Simon Hohberger; Ideas and Challenges for a State-of-the-Art Revenue Management System at the German Railway Company "Deutsche Bahn"	#97: Rongbing Huang; Flexible multi-period risk-sharing supply chain contracts under price and demand uncertainties	#80: Tomasz Sliwinski; Scenario based stochastic dynamic price optimization
	#118: Diwakar Gupta; Referral Timing and Fundraising Success in Crowdfunding	#159: Brenda Barnes; Bridging the Gap Between Academia and Practice for Mutual Benefit	#50: Robert Sanders; Reducing Retailer Food Waste through Revenue Management	#93: Darius Walczak; A Product-Focused Approach to Dynamic Programming Decomposition for the Network Revenue Management and Pricing Problems
12:45 to 1:30	<b>LUNCH</b>			
1:30 to 2:30	<b>Keynote: MIW/Fintech Panel – Rooms A, B, C</b>			
2:30 to 2:45	<b>BREAK</b>			
2:45 to 4:15	<b>Matching markets</b>	<b>Assortment Optimization</b>	<b>RM in Online Advertising</b>	<b>Dynamic Pricing</b>
	#131: Yash Kanoria; Dynamic Matching in School Choice: Efficient Seat	#47: Ruxian Wang; When Prospect Theory Meets Consumer Choice Models: Assortment and Pricing	#59: Ilan Lobel; Auction Design for ROI-Constrained Buyers	#155: Roozbeh Yousefi; Dynamic pricing of the fixed-term subscription contracts or one-time purchase offered to the strategic

If you have any questions or concerns during the conference, please notify Dean McKeown at (613)449-4592 or [dean.mckeown@queensu.ca](mailto:dean.mckeown@queensu.ca)

	Reallocation After Late Cancellations	Management with Reference Prices		customers in a heterogeneous market
	#151: Itai Ashlagi; Market Failure in Kidney Exchange	#66: Venus Lo; Assortment Optimization for a Multi-Channel Retailer with Features-based Value Boosts and Discounts	#139: Hamsa Bastani; Predicting with Proxies	#113: Jinglong Zhao; Dynamic Pricing under a Static Calendar
	#108: Nick Arnosti; How (Not) to Allocate Affordable Housing	#46: Jacob Feldman; Revenue Management versus Machine Learning; Finding optimal product displays at Alibaba	#54: Francisco Castro; The Scope of Sequential Screening with Ex-Post Participation Constraints	#142: Kate Ashley; An empirical analysis of hotel dynamic pricing practices
	No speaker	#94: Taner Bilgiç; Product Assortment Games with Effort in Online Marketplaces	#147: Antoine Desir; Design of Futures Contracts for Risk-averse Online Advertisers	#127: Qingchen Wang; Revenue Management for Parking with Advanced Reservations
<b>4:15 to 4:30</b>	<b>BREAK</b>			
<b>4:30 to 6:00</b>	<b>Incentives in online markets</b>	<b>RM Under Competition</b>	<b>Demand Estimation</b>	<b>RM in Services</b>
	#148: Ankur Mani; The Value of Price Discrimination in Large Random Networks	#117: Ruben van de Geer; Dynamic Pricing and Learning with Competition: Insights from the Dynamic Pricing Challenge	#89: Pavithra Harsha; A machine learning approach to the price-setting newsvendor problem	#152: Christopher Amaral; The Impact of Discriminatory Pricing Based on Customer Risk: An Empirical Investigation using Indirect Lending
	#133: Ashish Kabra; Effect Of Incentives In Marketplaces	#34: Ningyuan Chen; Duopoly Competition with Network Effects in Discrete Choice Models	#103: Pu He; Demand Estimation and Pricing in Long-tail Markets	#110: Xiangyu Zhang; Optimal Bidding for Highly-Valued IT Service Contracts: Theoretical Results and Practical Implications
	#74: Maxime Cohen; Frustration-based Promotions: Field Experiments in Ride-Sharing	#84: Santiago Balseiro; Learning in Repeated Auctions with Budgets: Regret Minimization and Equilibrium	#99: Ashwin Venkataraman; A Model-based Embedding Technique for Segmenting Customers	#135: Kyle Maclean; Stacking the House: Revenue Management in Live Entertainment
	#77: Amine Allouah; Prior-Independent Optimal Auctions	No speaker	#124: Divya Singhvi; Leveraging Comparables for New Product Sales Forecasting	#81: Pnina Feldman; Controlling Congestion when Consumers Choose Their Service Time
<b>6:00 to 6:15</b>	<b>CLOSING – Room B</b>			

Schedule Day Two – Friday June 22, 2018

	Room A	Room B	Room C
8:00 to 8:15	<b>WELCOME – Room B</b>		
8:15 to 9:45	<b>RM in Retailing</b>	<b>Network RM &amp; ADP</b>	<b>New Products &amp; RM</b>
	#85: Yannis Stamatopoulos; The Effects of Menu Costs on Retail Performance: Evidence from Adoption of the Electronic Shelf Label Technology	#109: Cornelia Schön, Simon Hohberger; Continuous Pricing in a Capacitated Network under Mixed Multinomial Logit Demand	#107: Amir Ahmadi-Javid; Integrating Pricing and Routing Decisions
	#82: Lennart Baardman; Increasing Profits: Leveraging Consumer Behavior to Optimize Promotions; from Nonsocial to Social	#126: Saied Samiedaluie; An Approximate Dynamic Programming Approach to Queueing Control Admission Problems	#122: Yash Kanoria; Centralized admissions for engineering colleges in India
	#35: Ram Dhurkari; Strategic Pricing Decision using Analytic Hierarchy Process	#154: Mika Sumida; A Constant-Factor Approximation Algorithm for Network Revenue Management	#153: Aaron Dukes; An Algorithm for Estimating the Benefit of an Automated Revenue Management System for Groups.
	#83: Tamar Cohen; Peak-End Demand Models and their Impact on Promotion Optimization Problems	#51: Gustavo Vulcano; Managing Uncertain Capacities for Revenue Optimization: Algorithm and Insights	#134: Pelin Pekgun; Empirically Investigating Strategic Behavior for Hotel Standby Upgrades
9:45 to 10:00	<b>BREAK</b>		
10:15 to 11:15	<b>Keynote: Scotiabank Special Session</b> 10:15 to 10:35 Akram Khaleghei, Data Scientist 10:35 to 10:55 Rogan Vleming, Senior Director Decisions Sciences, Digital Bank 10:55 to 11:15 Anshul Nanda, Director, Analytics and Business Planning		
	<b>BREAK</b>		
11:15 to 12:45	<b>RM in Online Platforms 2</b>	<b>Choice Models</b>	<b>RM and e-Commerce</b>
	#96: Zhe Liu; A Tale of Timescales: Surge Pricing and Dynamic Matching for Hotspot Demand Shock in Ride-Hailing Networks	#130: Nikos Trichakis; Network Revenue Management in the Face of Forward-looking Customers with Choice	#78: Yijun Zheng; Dynamic Pricing of Flexible Time Slots for Attended Home Delivery Management
	#76: Chiara Farronato; Consumer Protection in an Online World: When Does Occupational Licensing Matter?	#48: Gerardo Berbeglia; The generalized stochastic preference choice model	#41: Stanislav Mamonov; Time is money: Faster delivery premiums, demand elasticity and customer outcomes in online diamond r
	#164: Elaheh Fata; Online Type Matching for Revenue Maximization without Monge Conditions	#143: Fabian Bastin; Dynamic Discrete Choice Model for Railway Ticket Cancellation and Exchange Decisions	#162: Senthil Veeraraghavan; Do Ratings Cut Both Ways? Impact of Bilateral Ratings on Platforms.
	#145: Chinmoy Dutta; Locality Sensitive Hashing Based Matching for Ride-Sharing	#111: William L. Cooper; Optimal Worst-Case Pricing for a Logit Demand Model with Network Effects	#: ;
12:45 to 1:45	<b>LUNCH – RMP Board Meeting 1:00 to 1:45 Room B</b>		
1:45 to 3:15	<b>Business-to-Business</b>	<b>Dynamic Assortments</b>	<b>RM in Retailing 2</b>
	#62: Renyu (Philip) Zhang; Coopetition and Profit Sharing for Ride-sharing Platforms	#63: Yannik A. Peeters; Continuous Assortment Optimization: The extended MNL-model	#91: Mehmet Sekip Altug; Optimal Dynamic Allocation of Sales and Rental Inventory at a Retailer
	#157: Bahareh Mansouri; Suppliers' Pricing in Iterative Procurement Auctions	#92: Yufeng Cao; Assortment Optimization for Parallel Flights under an MNL Choice Model with Cheapest Fare Spikes	#163: Krishnamurthy Iyer; Signaling in Online Retail: Efficacy of Public Signals
	#68: Gregory Macnamara; Sequential Procurement through Contractual and Observational Learning	#86: Kris Ferreira; Learning to Rank an Assortment of Products	#73: A. Serdar Simsek; Value of Promotions with Delayed Incentives: An Empirical Investigation of Gift Card Promotions
	#45: Wen Jiao; Allocation and Nonlinear Pricing for Capacitated Stochastic Container Leasing System with Dynamic Arrivals	#69: Huseyin Topaloglu; Dynamic Assortment Optimization for Reusable Products with Random Usage Durations	#146: Tolga Aydinliyim; Managing Opportunistic Returns In Retail Operations
3:15 to 3:30	<b>BREAK</b>		
3:30 to 5:00	<b>Game Theory Models</b>	<b>RM with incomplete information</b>	<b>Marketplace Analytics</b>

	#71: Amir Farshbaf Geranmayeh; Coordinating advertising and pricing in a manufacturer - retailer channel in presence of strategic customer	#95: Stefano Vaccari; Information Aggregation via Consumer Reviews in Online Marketplaces	#112: Ashwin Venkataraman; A Conditional Gradient Approach for Nonparametric Estimation of Mixing Distributions
	#172: Vijay Kamble; A Truth Serum for Large Scale Evaluations	#72: Abhishek Gupta; Selling Renewable Generation with a Penalty for Shortfall	#43: Joseph Lai; Implementing A Pricing Model In R Shiny
	#167: Ming Hu; Social Pricing in the Sharing Economy: Theory and Empirical Evidence from Airbnb	#105: Adam Elmachtoub; The Value of Personalized Pricing	#160: Jue Wang; Dynamic Campaign Optimization
	No Speaker	#52: Xavier S Warnes; The Cost of Income Guarantees in Centralized Marketplaces	#140: Hamsa Bastani; Interpreting Predictive Models for Human-in-the-Loop Analytics
5:00 to 5:15	<b>CLOSING – Room B</b>		

## List of Speakers

### [A. Serdar Simsek](#)

Value of Promotions with Delayed Incentives: An Empirical Investigation of Gift Card Promotions

### [Aaron Dukes](#)

An Algorithm for Estimating the Benefit of an Automated Revenue Management System for Groups.

### [Abhishek Gupta](#)

Selling Renewable Generation with a Penalty for Shortfall

### [Adam Elmachtoub](#)

The Value of Personalized Pricing

### [Amine Allouah](#)

Prior-Independent Optimal Auctions

### [Amir Ahmadi-Javid, Maryam Daryalal](#)

Integrating Pricing and Routing Decisions

### [Amir Farshbaf Geranmayeh](#)

Coordinating advertising and pricing in a manufacturer - retailer channel in presence of strategic customer

### [Ankur Mani](#)

The Value of Price Discrimination in Large Random Networks

### [Antoine Desir](#)



Design of Futures Contracts for Risk-averse Online Advertisers

[Arnoud den Boer](#)

Dynamic Pricing with Demand Learning and Reference Effects

[Ashish Kabra](#)

Effect Of Incentives In Marketplaces

[Ashwin Venkataraman](#)

A Conditional Gradient Approach for Nonparametric Estimation of Mixing Distributions

[Ashwin Venkataraman](#)

A Conditional Gradient Approach for Nonparametric Estimation of Mixing Distributions

[Ashwin Venkataraman](#)

A Model-based Embedding Technique for Segmenting Customers

[Bahareh Mansouri](#)

Suppliers' Pricing in Iterative Procurement Auctions

[Bo Cowgill](#)

Competition and Specificity in Market Design: Evidence from Geotargeted Advertising

[Brenda Barnes](#)

Bridging the Gap Between Academia and Practice for Mutual Benefit

[Chiara Farronato](#)

Consumer Reviews and Regulation: Evidence from NYC Restaurants

[Chiara Farronato](#)

Consumer Protection in an Online World: When Does Occupational Licensing Matter?

[Chinmoy Dutta](#)

Locality Sensitive Hashing Based Matching for Ride-Sharing

[Christopher Amaral](#)

The Impact of Discriminatory Pricing Based on Customer Risk: An Empirical Investigation using Indirect Lending

[Cornelia Schön, Simon Hohberger](#)

Continuous Pricing in a Capacitated Network under Mixed Multinomial Logit Demand

[Daniela Saban](#)

Optimal commissions and subscriptions in networked markets

[Darius Walczak](#)

A Product-Focused Approach to Dynamic Programming Decomposition for the Network Revenue Management and Pricing Problems

[Diego Escobari](#)

Price Discrimination and Focal Points for Tacit Collusion: Evidence from the Airline Industry

[Divya Singhvi](#)

Leveraging Comparables for New Product Sales Forecasting

[Diwakar Gupta](#)

Referral Timing and Fundraising Success in Crowdfunding

[Elaheh Fata](#)

Online Type Matching for Revenue Maximization without Monge Conditions

[Fabian Bastin](#)

Dynamic Discrete Choice Model for Railway Ticket Cancellation and Exchange Decisions

[Fan You](#)

Finite-Horizon Approximate Linear Programs for an Infinite-Horizon Revenue Management Problem

[Francisco Castro](#)

Surge Pricing and its Spatial Supply Response

[Francisco Castro](#)

The Scope of Sequential Screening with Ex-Post Participation Constraints

[Gerardo Berbeglia](#)

The generalized stochastic preference choice model

[Gonzalo Romero](#)

Revenue Management with Repeated Customer Interactions

[Gregory Macnamara](#)

Sequential Procurement through Contractual and Observational Learning

[Gustavo Vulcano](#)

Managing Uncertain Capacities for Revenue Optimization: Algorithm and Insights

[Hamsa Bastani](#)

Predicting with Proxies

[Hamsa Bastani](#)

Interpreting Predictive Models for Human-in-the-Loop Analytics

[Hideaki Takagi](#)

Extension of Littlewood's rule to a multi-period model with possible standby customers

[Huseyin Topaloglu](#)

Dynamic Assortment Optimization for Reusable Products with Random Usage Durations

[Ilan Lobel](#)

Auction Design for ROI-Constrained Buyers

[Yash Kanoria](#)

Dynamic Matching in School Choice: Efficient Seat Reallocation After Late Cancellations

[Itai Ashlagi](#)

Market Failure in Kidney Exchange

[Jacob Feldman](#)

Revenue Management versus Machine Learning: Finding optimal product displays at Alibaba

[Jiding Zhang](#)

Flexibility and Relationships in Online Marketplaces: Prescriptions for Pricing and Matching

[Jinglong Zhao](#)

Dynamic Pricing under a Static Calendar

[Joseph Lai](#)

Implementing A Pricing Model In R Shiny

[Jue Wang](#)

Dynamic Campaign Optimization

[Kate Ashley](#)

An empirical analysis of hotel dynamic pricing practices

[Kris Ferreira](#)

Learning to Rank an Assortment of Products

[Krishnamurthy Iyer](#)

Signaling in Online Retail: Efficacy of Public Signals

[Kyle Maclean](#)

Stacking the House: Revenue Management in Live Entertainment

[Lennart Baardman](#)

Increasing Profits: Leveraging Consumer Behavior to Optimize Promotions; from Nonsocial to Social

[Maxime Cohen](#)

Frustration-based Promotions: Field Experiments in Ride-Sharing

[Mehmet Sekip Altug](#)

Optimal Dynamic Allocation of Sales and Rental Inventory at a Retailer

[Mika Sumida](#)

A Constant-Factor Approximation Algorithm for Network Revenue Management

[Ming Hu](#)

Social Pricing in the Sharing Economy: Theory and Empirical Evidence from Airbnb

[Mohamed Mostagir](#)

Information Design in Dynamic Contests: A Lab Experiment

[Nick Arnosti](#)

How (Not) to Allocate Affordable Housing

[Nikos Trichakis](#)

Network Revenue Management in the Face of Forward-looking Customers with Choice

[Ningyuan Chen](#)

Duopoly Competition with Network Effects in Discrete Choice Models

[Pavithra Harsha](#)

A machine learning approach to the price-setting newsvendor problem

[Pelin Pekgun](#)

Empirically Investigating Strategic Behavior for Hotel Standby Upgrades

[Pengyu Qian](#)

The Value of State Dependent Control in Ridesharing Systems

[Pnina Feldman](#)

Controlling Congestion when Consumers Choose Their Service Time

[Pu He](#)

Demand Estimation and Pricing in Long-tail Markets

[Qingchen Wang](#)

Revenue Management for Parking with Advanced Reservations

[Ram Dhurkari](#)

Strategic Pricing Decision using Analytic Hierarchy Process

[Renyu \(Philip\) Zhang](#)

Coopetition and Profit Sharing for Ride-sharing Platforms

[Robert Sanders](#)

## Reducing Retailer Food Waste through Revenue Management

### [Rongbing Huang](#)

Flexible multi-period risk-sharing supply chain contracts under price and demand uncertainties

### [Roosbeh Yousefi](#)

Dynamic pricing of the fixed-term subscription contracts or one-time purchase offered to the strategic customers in a heterogeneous market

### [Ruben van de Geer](#)

Dynamic Pricing and Learning with Competition: Insights from the Dynamic Pricing Challenge at the 2017 INFORMS RM & Pricing Conference

### [Ruxian Wang](#)

When Prospect Theory Meets Consumer Choice Models: Assortment and Pricing Management with Reference Prices

### [Saied Samiedaluie](#)

An Approximate Dynamic Programming Approach to Queueing Control Admission Problems

### [Santiago Balseiro](#)

Learning in Repeated Auctions with Budgets: Regret Minimization and Equilibrium

### [Senthil Veeraraghavan](#)

Do Ratings Cut Both Ways? Impact of Bilateral Ratings on Platforms.

### [Simon Hohberger](#)



Ideas and Challenges for a State-of-the-Art Revenue Management System at the German Railway Company “Deutsche Bahn”

[So Yeon Chun](#)

Should I Pay for This Purchase or Redeem Points? Effects of Loyalty Program Design on Consumer Decisions to Redeem Points

[Stanislav Mamonov](#)

Time is money: Faster delivery premiums, demand elasticity and customer outcomes in online diamond r

[Stefano Vaccari](#)

Information Aggregation via Consumer Reviews in Online Marketplaces

[Tamar Cohen](#)

Peak-End Demand Models and their Impact on Promotion Optimization Problems

[Taner Bilgiç](#)

Product Assortment Games with Effort in Online Marketplaces

[Terry Taylor](#)

Ridesharing Platforms: Competition and Automated Vehicles

[Theja Tulabandhula](#)

A Realistic Online Algorithm to Learn Buyer Behavior

[Tolga Aydinliyim](#)

## Managing Opportunistic Returns In Retail Operations

### [Tomasz Sliwinski](#)

Scenario based stochastic dynamic price optimization

### [Vahideh Manshadi](#)

Two-stage Pandora's Box for Product Ranking

### [Venus Lo](#)

Assortment Optimization for a Multi-Channel Retailer with Features-based Value Boosts and Discounts

### [Victor Araman](#)

Crowdvoting New Product Introduction

### [Vijay Kamble](#)

A Truth Serum for Large Scale Evaluations

### [Warren Lieberman](#)

Pricing for Profit

### [Wen Jiao](#)

Allocation and Nonlinear Pricing for Capacitated Stochastic Container Leasing System with Dynamic Arrivals

### [Wenchang Zhang](#)

Learning, Welfare, and Profits in Service Platforms

[William L. Cooper](#)

Optimal Worst-Case Pricing for a Logit Demand Model with Network Effects

[Xavier S Warnes](#)

The Cost of Income Guarantees in Centralized Marketplaces

[Xiangyu Zhang](#)

Optimal Bidding for Highly-Valued IT Service Contracts: Theoretical Results and Practical Implications

[Xiao Huang](#)

Price to Gain or Price to Retain? An Empirical Study of Hotel Pricing and Customer Cancellation Behavior

[Xuanming Su](#)

Centralized Location Control in Ridesharing Networks

[Yaarit Even](#)

Efficiency and Information Aggregation in Heterogeneous Markets

[Yang Chen](#)

Quantifying mileage-runs

[Yannik A. Peeters](#)

Continuous Assortment Optimization: The extended MNL-model

[Yannis Stamatopoulos](#)

The Effects of Menu Costs on Retail Performance: Evidence from Adoption of the Electronic Shelf Label Technology

[Yash Kanoria](#)

Centralized admissions for engineering colleges in India

[Yifan Feng](#)

Learning Customer Preferences From Crowdvoting

[Yijun Zheng](#)

Dynamic Pricing of Flexible Time Slots for Attended Home Delivery Management

[Yufeng Cao](#)

Assortment Optimization for Parallel Flights under an MNL Choice Model with Cheapest Fare Spikes

[Zhe Liu](#)

A Tale of Timescales: Surge Pricing and Dynamic Matching for Hotspot Demand Shock in Ride-Hailing Networks

## [Gonzalo Romero](#)

### **Revenue Management with Repeated Customer Interactions**

We model and analyze a revenue management problem where a platform interacts with a set of customers over a number of periods. Unlike traditional network revenue management, we consider customers who can dynamically change state between periods. A customer's state depends on the quality of their past service and determines the amount of budget that they allocate to the platform. These dynamics create a trade-off between the platform myopically maximizing short-term revenues, versus maximizing the long-term goodwill of its customers. We (i) show that, in general, finite look-ahead policies can perform arbitrarily poorly in this repeated setting and (ii) identify a set of reasonable conditions under which myopic policies that ignore the budget dynamics are either optimal or near-optimal.

## [Wenchang Zhang](#)

### **Learning, Welfare, and Profits in Service Platforms**

In the paper, we investigate the revenue impact of informational policies in online service platforms. In particular, we show that by delaying the revelation of quality information of service providers, the platform can improve its revenue by experimenting more newcomers and maintaining more high-type providers in steady state.

## [Arnoud den Boer](#)

### **Dynamic Pricing with Demand Learning and Reference Effects**

We study dynamic pricing problems with incomplete information about the demand function, and with the presence of reference price effects. That means that the demand is not only affected by the current selling price, but also by consumers' price expectations. We show how to handle these reference effects, in order to build asymptotically optimal pricing policies.

## [Yannis Stamatopoulos](#)

### **The Effects of Menu Costs on Retail Performance: Evidence from Adoption of the Electronic Shelf Label Technology**

We use the adoption of electronic shelf labels (ESLs) by a major international grocery retailer in 2015 in the United Kingdom to identify the effects of reducing physical menu costs (operational costs of price adjustment) on retail performance. The ESL technology essentially eliminates the physical costs associated with price adjustment (e.g., costs of printing and distributing price tags). We find that the elimination of physical menu costs increased revenues, decreased the average price per unit sold, and increased sales volumes. We also find that ESL adoption increased price-adjustment volume, decreased the average size of a price adjustment, and

decreased the batching of price changes across different products. Finally, we find that ESL adoption had a statistically significant effect on the volume of downward price changes, but not on the volume of upward price changes.

## Cornelia Schön, Simon Hohberger

### **Continuous Pricing in a Capacitated Network under Mixed Multinomial Logit Demand**

In this paper, we consider the deterministic multi-product multi-resource dynamic pricing (DMMDP) problem with continuous prices under the mixed multinomial logit (MMNL) choice model. The DMMDP problem arises in many applications where pricing decisions for multiple products should be optimized jointly, in particular when products have cross-effects on demands of other products and/or different products share the same resources (Talluri and van Ryzin 2004, Chen and Chen 2015). Applications are manifold, such as revenue management for airlines, railways, and hotels, assortment pricing in retailing, or product line pricing in consumer goods industries. The corresponding problem instances in practice are typically of large-scale size such that efficient solution techniques are required.

The MMNL model is considered to be a powerful choice model that captures heterogeneous cross-effects in demand and that can approximate any random utility choice model arbitrarily closely (McFadden and Train 2000). It has received increasing attention in the related field of product assortment (PA) optimization, involving a seller's discrete decisions about the selection of products and their prices (see, e.g., Feldman and Topaloglu 2015, Kunnumkal 2015, Méndez-Díaz 2015). Since the PA problem is NP-hard under the MMNL choice model (Rusmevichientong et al. 2014, Désir et al. 2014), much work has been focused on deriving upper bounds and efficient approximations, with the recent exception of Sen et al. (2017) who propose an exact conic MIP approach.

On the other hand, the MMNL model and its incorporation into the DMMDP problem has only received scant attention in the dynamic pricing literature (e.g., Keller et al. 2014), despite its theoretical and practical relevance. The more common approach so far has been to incorporate the standard single-segment MNL choice model into the DMMDP problem (see, e.g., Dong et al. 2009, Zhang and Lu 2013, Keller et al. 2014). The logit profit function is known to be concave with respect to demand (Hanson and Martin 1996, Song and Xue 2007, Dong et al. 2009, Li and Huh 2011). In case of a single customer segment, there is a linearizable one-to-one mapping between product prices and MNL choice probabilities such that the demand model satisfies some regularity conditions, and the resulting optimization problem DMMDP is a convex optimization (minimization) problem in demand. In case the logit choice model encompasses multiple customer segments with heterogeneous price sensitivity parameters, the convex problem structure can still be maintained if it is feasible to simultaneously quote individual prices for the same product to each segment according to a first- or third-degree price discrimination (Schön 2010a, b). However, this requires the capability to identify a priori which customer segment an incoming sales request belongs to. In the more common case considered here, where the same product is offered at a uniform price to all customer requests occurring at the same time, the convexity property is lost, since non-convex constraints need to be introduced to ensure price consistency across segments with overlapping consideration sets. Accordingly, the DMMDP problem under MMNL choice is non-linear non-convex and thus difficult to solve in general.

How to efficiently solve the continuous pricing problem with multiple segments to optimality is still an open problem, and we want to contribute to narrow this gap. Our contributions are as follows:

- First, we analyze the DMMDP problem with continuous prices and price consistency constraints under the MMNL choice model in detail with regard to its mathematical structure.
- We present an approximate optimization problem to derive an upper bound on the optimal profit and to determine heuristic solutions. The approximate problem is convex, and can therefore be solved efficiently even for large problem instances. An experimental study shows that the approach is very promising with regard to run time performance and solution quality.
- We present a convex mixed-integer programming approach that allows to tighten the upper bound arbitrarily close-to-optimum and to determine provably near-optimal solutions of the original problem; to our knowledge, this is the first approach to approximately tackle the problem under the MMNL choice model with a performance guarantee; in our experiments, we are able to approximately solve medium-sized problem instances in reasonable time. Furthermore, we discuss the potential benefits we gain by allowing prices to be continuous rather than restricting them to discrete values with regard to solution quality and run time performance.
- The suggested dynamic pricing approach is applied to a real-world revenue management case study of the German long-distance railway network.

## [Amir Ahmadi-Javid, Maryam Daryalal](#)

### **Integrating Pricing and Routing Decisions**

Integrating Pricing and Routing Decisions: Application in Last-Mile Delivery

## [Terry Taylor](#)

### **Ridesharing Platforms: Competition and Automated Vehicles**

Ridesharing platforms simultaneously compete for a common pool of supply (independent driver-workers) and demand (rider-consumers). Ridesharing platforms are aggressively pursuing the development of autonomous vehicles, and have publicly stated that they anticipate autonomous vehicles will be an important component of their operating models. What are the implications of competition and access to autonomous vehicle technology for the management of ridesharing platforms? How does a platform's access to supply-side (namely, autonomous vehicle) technology change prescriptions for its demand-side (namely, pricing) decisions? How does the presence of demand-side competition change prescriptions for a platform's supply-side (namely, autonomous vehicle fleet size) decisions?

## [Mohamed Mostagir](#)

### **Information Design in Dynamic Contests: A Lab Experiment**

If you have any questions or concerns during the conference, please notify Dean McKeown at (613)449-4592 or [dean.mckeown@queensu.ca](mailto:dean.mckeown@queensu.ca)

Contests are a common mechanism for extracting effort from participants. Their use is widespread in a variety of settings like workplace promotions, crowdsourcing innovation, healthcare quality, etc. One of the pivotal aspects of contest design is the contest's information structure: what information should the contest designer provide to participants and when should this information be revealed? The answers to these questions have important implications to how players behave and the overall outcome of the contest. We design a novel experiment to evaluate how different information disclosure policies perform in a lab setting and compare our findings with various theoretical predictions.

## [Francisco Castro](#)

### **The Scope of Sequential Screening with Ex-Post Participation Constraints**

We consider the pricing problem faced by a platform matching price sensitive customers to flexible supply units within a geographic area. This can be interpreted as the problem faced in the short-term by a ride-hailing platform matching supply and demand within a city. We propose a framework in which a platform selects prices for the different locations, and drivers respond by choosing where to relocate based on prices, travel costs and driver congestion levels. Our contributions are along two dimensions. We first derive general results on the structure of optimal pricing policies. In particular, we elucidate structural properties of supply equilibria and the corresponding utilities that emerge and establish a form of spatial decomposition, which allows us to localize the analysis. In turn, uncovering an appropriate knapsack structure to the platform's problem, we establish a crisp local characterization of an optimal solution and the corresponding supply response. We then specialize the analysis to a family of models that isolates the impact of supply demand imbalances by introducing a demand shock. We derive in quasi-closed form the optimal solution across the city, highlighting the implications of the strategic nature of supply units. In particular, we show that the platform will use prices to create damaged regions where demand is shut-down or driver congestion is artificially high, incentivizing a suitable number of drivers to relocate towards the demand shock. Furthermore, the optimal solution, while better balancing supply and demand around the shock, also ends up inducing movement away from it.

## [Theja Tulabandhula](#)

### **A Realistic Online Algorithm to Learn Buyer Behavior**

We propose a new efficient online algorithm to learn the parameters governing the purchasing behavior of a utility maximizing buyer, who responds to prices, in a repeated interaction setting. The key feature of our algorithm is that it can learn even non-linear buyer utility while working with arbitrary price constraints that the seller may impose. This overcomes a major shortcoming of previous approaches, which use unrealistic prices to learn these parameters making them unsuitable in practice.

## [Lennart Baardman](#)

### **Increasing Profits: Leveraging Consumer Behavior to Optimize Promotions; from Nonsocial to Social**

If you have any questions or concerns during the conference, please notify Dean McKeown at (613)449-4592 or [dean.mckeown@queensu.ca](mailto:dean.mckeown@queensu.ca)



Detecting trends in fashion can help retailers determine effective personalized promotion plans more easily. Access to social media data can be important in order to be able to understand these trends. Unfortunately, social data is usually unavailable. We introduce a personalized demand model that captures customer trends and show an efficient method for estimating it from transaction data. Furthermore, we illustrate how this personalized demand model can allow us to do targeted promotion strategies. We show that the promotion targeting problem is NP-hard to solve exactly, and develop an efficient greedy approach that scales to large datasets and finds provably near-optimal solutions.

## [Saied Samiedaluie](#)

### **An Approximate Dynamic Programming Approach to Queueing Control Admission Problems**

We study a classical queueing control problem with multiple classes of customers. The queue is a loss system; i.e., arriving customers are rejected if all servers are busy. When a server is available, the decision is whether to admit an arriving customer and collect a lump-sum revenue. We model this problem as a continuous-time infinite-horizon dynamic program and solve it using approximate linear programming (ALP). We study several alternative approximation architectures and numerically investigate their policy performance. Our approach is potentially useful for a wide variety of queueing control problems.

## [Yash Kanoria](#)

### **Centralized admissions for engineering colleges in India**

Centralized admissions for engineering colleges in India: We designed and implemented a new joint seat allocation process for undergraduate admissions to over 500 programs spread across 80 Centrally Funded Technical Institutes (CFTIs) in India, including the Indian Institutes of Technology (IITs). Our process is based on the well known Deferred Acceptance algorithm, but complex rules regarding seat reservations led us to make a number of algorithmic innovations, including (i) a practical heuristic for incorporating a non-nested common quota, (ii) a method to “dereserve” seats with no modifications to the core software, and (iii) considering reserved category candidates for reserved seats before unreserved seats so as to reduce variability in the number of reserved category candidates admitted. Our new seat allocation process went live in 2015, and based on its success, including significant reduction in vacancies, it has been used again in 2016 and 2017, with continuing efforts to identify and address pending issues. While vacancies at the IITs have reduced by over 70%, a significant number of vacancies persist in the non-IIT CFTIs due to students not taking up the offered seats. We suggest some approaches to reduce these vacancies.

## [Pengyu Qian](#)

### **The Value of State Dependent Control in Ridesharing Systems**

We study the state dependent control for a closed queueing network model inspired by ridesharing systems. The platform can choose which supply unit to assign to an incoming customer; the unit becomes available again at customer's destination. The platform aims to minimize proportion of dropped demand.

We propose a family of simple policies called Scaled MaxWeight (SMW) and prove that under complete resource pooling (analogous to Hall's condition), any SMW policy induces an exponential decay of demand-drop probability as the number of supply units goes to infinity. We further show that there is an SMW policy that achieves the \*optimal\* exponent among all assignment policies and specify it analytically

## [Yang Chen](#)

### **Quantifying mileage-runs**

“Mileage-runs” refer to the strategic behavior when consumers purposefully fly/spend more with a primary goal of “gaming” the loyalty program. This work investigates the extent of mileage-running based on a proprietary dataset obtained from a large loyalty program.

## [Daniela Saban](#)

### **Optimal commissions and subscriptions in networked markets**

We consider a platform providing a marketplace for a group of sellers and buyers to trade with each other. Inexchange for providing service, the platform charges commissions and subscription fees to sellers and buyers.

Potential buyers and sellers are divided into types; not all buyer and seller types are compatible with each other. We encode the compatibility among different types using a bipartite network. The platform cannot directly control the transaction prices, i.e., traders determine the prices at which the goods/services are exchanged. We establish that, in order to maximize its revenues, the platform may need to charge different commissions/subscriptions to different types, depending on their network position. In fact, we show that if the same commissions/subscriptions are employed for all agents on the same side, the revenue loss can be unbounded. We complement this worst-case result by providing a bound on the revenue loss in terms of the supply/demand imbalance across the network under homogeneous value distributions. Surprisingly, we also show that, in general, charging commissions/subscriptions to only one side of the market (i.e., only to buyers or only to sellers) leads to lower revenues than optimal, even when different types on the same side can be charged different fees. Furthermore, we characterize the impact of the network structure on the revenues of the platform. Finally, we investigate how the commissions/subscriptions chosen by the platform impact social welfare. We establish that under some convexity assumptions on the value distributions, the revenue-maximizing commissions/subscriptions induce at least 2/3 of the maximum achievable social welfare.

## [Victor Araman](#)

### **Crowdvoting New Product Introduction**

Launching new products into the marketplace is a complex and risky endeavor that companies must continuously undertake. In this paper, we consider a seller who has the ability to first test the market and gather demand information before deciding whether or not to launch a new product. In particular, we consider the case in which the seller sets up an online voting system and offers multiple versions of the product - differentiated through their quality levels and prices - for potential customers to vote on. We investigate the optimal design of such a crowdvoting system in order for it to provide an effective demand forecast and allow the seller to identify which version if any to commercialize.

## [Mika Sumida](#)

### **A Constant-Factor Approximation Algorithm for Network Revenue Management**

We provide a constant-factor approximation algorithm for network revenue management problems. In our approximation algorithm, we construct an approximate policy using value function approximations that are expressed as linear combinations of basis functions. We use a backward recursion to compute the coefficients of the basis functions in the linear combinations. If each product uses at most  $L$  resources, then the total expected revenue obtained by our approximate policy is at least  $1/(1 + L)$  of the optimal total expected revenue. In many network revenue management settings, although the number of resources and products can become large, the number of resources used by a product remains bounded. In this case, our approximate policy provides a constant-factor performance guarantee. To our knowledge, our approximate policy is the first constant-factor approximation algorithm for network revenue management problems. Our approach can incorporate the customer choice behavior among the products and allows the products to use multiple units of a resource, while still maintaining the performance guarantee. In our computational experiments, we demonstrate that our approximate policy performs quite well, providing total expected revenues that are substantially better than its theoretical performance guarantee.

## [Aaron Dukes](#)

### **An Algorithm for Estimating the Benefit of an Automated Revenue Management System for Groups.**

Issues involved in Estimating Benefit provided by an automated Hotel Revenue Management System for Group business, and an algorithm designed to circumvent those issues.

## [Xuanming Su](#)

### **Centralized Location Control in Ridesharing Networks**

We build a spatial equilibrium model to study traffic flows in ridesharing networks. We consider two regimes: drivers follow a central routing plan (i.e., centralized) or choose their own directions (i.e., decentralized). Our results measure the value of centralized routing.

## [So Yeon Chun](#)

### **Should I Pay for This Purchase or Redeem Points? Effects of Loyalty Program Design on Consumer Decisions to Redeem Points**

In this research, we study consumers' choices between redeeming loyalty points or spending money. Specifically, we propose that design characteristics that vary across loyalty programs lead to systematic differences in the way consumers spend loyalty points compared with money.

## [Vahideh Manshadi](#)

### **Two-stage Pandora's Box for Product Ranking**

On online platforms, consumers face an abundance of options that are displayed in the form of a position ranking. Only products placed in the first few positions are readily accessible to the consumer, and she needs to exert effort to access more options. We study how platforms with different business models should rank products to maximize their profit.

## [Yifan Feng](#)

### **Learning Customer Preferences From Crowdvoting**

We study a seller introducing a new product with multiple potential product versions into the marketplace. In order to pick the version that is most likely to be preferred by customers, the seller uses an online system that allows potential buyers to vote for their preferred versions. We study how to dynamically customize each individual voter's choice set, in order to most efficiently learn overall customer preferences. We propose an algorithm that balances breadth of choice and accuracy in determining the best product. We show this algorithm is asymptotically optimal in speed of learning.

## [Tamar Cohen](#)

### **Peak-End Demand Models and their Impact on Promotion Optimization Problems**

In-store promotions are a highly effective marketing tool that can have a significant impact on revenue. 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) as well as sophisticated machine learning tools (a 2-stage version of the Gaussian Process). We show that this new demand model predicts actual sales more accurately than current methods. In order to determine promotion strategies, subsequently, we suggest a compact Dynamic Programming (DP) approach that uses the proposed demand model. We characterize structural properties of the optimal policy that allow us to solve the proposed model

tractably. Furthermore, using these properties, we establish conditions under which the proposed DP solves the promotion planning optimization problem exactly. In the cases where these 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. Finally, we test our approach on data from large retailers and demonstrate 5.9-19.1% relative improvement in the precision of the demand prediction as well as an average of 9.1-11.6% increase in profit relative to the retailer's current practices.

## Gustavo Vulcano

### **Managing Uncertain Capacities for Revenue Optimization: Algorithm and Insights**

We study the problem of managing uncertain capacities for revenue optimization over a network of resources, accounting for two variants according to the source of the uncertainty. In the first one, the initial allocation of capacities is subject to changes during the selling horizon. In the second one, the availability of physical capacities by the time of service execution is random.

For both variants of the problem, during some initial, pre-specified time of the selling horizon, physical capacities remain uncertain.

The control policy analyzed is aligned with the current industry practice, where the seller sets a virtual capacity and a bid-price for each resource, and collects revenues from an arriving stream of customers. Admitted requests that cannot be accommodated within the final, effective capacities, incur a penalty cost. The firm's objective is to maximize the total cumulative net revenue (sales revenue minus penalty cost).

The problem arises in practice for instance when airlines are subject to last-minute change of aircrafts, and in cargo revenue management where the capacity left by the passengers' load is used for freight. We develop a stochastic gradient algorithm for this problem whose limit points are stationary points of the expected net revenue function. Then, through an exhaustive numerical study, we show that our controls are computed efficiently and deliver revenues that are almost consistently higher than the ones obtained from benchmarks based on the widely-adopted deterministic linear programming model. Finally, we obtain managerial insights about the inconvenience of delaying the allocation of resources and about the increasing losses in the degree of capacity heterogeneity.

## Pelin Pekgun

### **Empirically Investigating Strategic Behavior for Hotel Standby Upgrades**

In this research, we investigate the existence and extent of strategic customer behavior, utilizing a major hotel chain's 16-month booking and standby upgrades data with two research goals: First, we seek to understand customer decision-making dynamics in the presence of standby upgrades, and estimate the percentage of customers who exhibit strategic behavior. Second, we seek to identify hotel and upgrade characteristics that are more or less likely to result in strategic behavior. Following the empirical analysis, we propose a new pricing strategy where loyalty members and non-loyalty customers are offered different prices for premium rooms and standby upgrades. Through a numerical study, we compare the performance of such a pricing policy against two benchmark policies, reflective of what a hotel chain could do if they did not know the extent of strategic customers.

## Warren Lieberman

### **Pricing for Profit**

Traditional pricing approaches, such as "everyday low pricing," may be easy to implement, but typically result in less revenue, and far less profit, than data-driven strategies. Furthermore, many companies underestimate the willingness of potential customers to pay higher prices for a product that they perceive to be better (as opposed to what the company perceives to be better). This presentation highlights a new dynamic approach to pricing and product segmentation that has led to revenue increases of 4 – 6 percent for a number of companies in the self-storage industry.

In addition, for industries where customers subscribe to a product and customer subscription rates can be increased periodically, we will show how adopting a differentiated subscription rate increase policy has led to revenue increases that are 1.5 – 2.5 percentage points greater than more common approaches where an across-the-board subscription rate increase is implemented.

## Diego Escobari

### **Price Discrimination and Focal Points for Tacit Collusion: Evidence from the Airline Industry**

We use unique data sets with round-the-clock posted fares and a regression discontinuity design to identify price discrimination in advance-purchase discounts. Price discrimination increases fares by 14% between two and one week before departure, and by 7.6% between three and two weeks to departure. While competition reduces price discrimination, it is unaffected by product variety for a multiproduct monopolist. The results show that the arbitrary thresholds of 7 and 14 days-in-advance serve as focal points for tacit collusion and to implement price discrimination in competitive markets. For round-trip tickets price discrimination depends on the days-in-advance for both the outbound and inbound flights.

## Hideaki Takagi

### **Extension of Littlewood's rule to a multi-period model with possible standby customers**

Explicit formulas in terms of multiple integrals and numerical examples for the direct extension of Littlewood's classical rule (1972) to a static multi-period model, including the relation to the Brumelle-McGill theorem (1993) and Robinson's model (1995) with standby customers.

## Zhe Liu

### **A Tale of Timescales: Surge Pricing and Dynamic Matching for Hotspot Demand Shock in Ride-Hailing Networks**

We study an online platform that operates a ride-hailing network with price and delay sensitive riders and strategic drivers. Our model jointly considers surge pricing (rider price and driver wage) and dynamic spatial matching in the platform's profit maximization problem, responding to a demand shock at a hotspot with uncertain magnitude and/or duration. Surge pricing is meant to a) moderate demand and b) incentivize supply to proactively reposition toward the hotspot. Dynamic matching trades off non-hotspot local matches for more lucrative hotspot matches, within service level constraints.

Our results show the interplay between important timescales, e.g., rider patience, demand shock duration, and drivers' travel delay to the hotspot, and their impact on system performance. The distinctive features of this work lie on the focus of system transient under non-stationary demand, the network setting, and drivers' strategic responsiveness to surge signals given delayed incentives.

## Nikos Trichakis

### **Network Revenue Management in the Face of Forward-looking Customers with Choice**

We consider a canonical network revenue management model in which a seller offers multiple products, which consume capacitated resources, for sale to utility-maximizing customers who choose both (1) when to buy, and (2) which product to buy. We use a novel decomposition approach for a suitable dynamic mechanism design problem to derive for all non-anticipating dynamic pricing policies an upper bound to expected revenues. We employ the bound to obtain a constant-factor guarantee for the performance of static pricing in a fluid-type regime.

## Yijun Zheng

### **Dynamic Pricing of Flexible Time Slots for Attended Home Delivery Management**

We demonstrate the dynamic pricing of regular and flexible time slots offered for the attended home delivery management. A flexible slot is introduced as a combination of several regular time slots in which customer's orders need to be delivered. Numerical results based on realistically-sized scenarios are also presented to show that expected profit increases when adding flexible time slots rather than using only regular time slots.

## [Chiara Farronato](#)

### **Consumer Reviews and Regulation: Evidence from NYC Restaurants**

We investigate how two signals of restaurant quality, hygiene grade cards and online reviews, affect consumer choice and restaurant hygiene. Unlike hygiene cards, online reviews contain information about multiple dimensions of restaurant quality. To extract signals of hygiene from online reviews, we exploit the fact that health inspectors look for different types of violations and we apply machine learning methods to predict the occurrence of individual violations from review text. Using out-of-sample prediction accuracy as a measure of signal informativeness, we find substantial heterogeneity in how informative reviews are about different violations. Reviews are more informative about food handling and pest violations than facilities and maintenance violations. Next, we estimate the effect of hygiene information contained in online reviews on consumer demand and restaurant hygiene choices. We find that consumer demand is more sensitive to more informative signals of hygiene. In addition, restaurants that are reviewed online are more likely to comply with hygiene standards for which online reviews provide a more informative signal. Our results have implications for the allocation of limited regulator resources when consumers rate service providers online.

## [Xiao Huang](#)

### **Price to Gain or Price to Retain? An Empirical Study of Hotel Pricing and Customer Cancellation Behavior**

We use hotel transaction data to investigate the customer cancellation behavior. We show that cancellation rates are highly correlated with booking rates and the posted rates at the time of cancellation. We discuss the implications of this finding on hotel overbooking. Our results complement the empirical research on strategic customer behavior in revenue management and highlight the importance of accounting for the behavioral impact of pricing.

## [Yaarit Even](#)

### **Efficiency and Information Aggregation in Heterogeneous Markets**

This paper studies how heterogeneous preferences shape the informational and allocative efficiency of centralized markets with asymmetric information. We start by showing that introducing agent-level heterogeneity (e.g., in terms of trading costs) to the standard rational expectations equilibrium models reduces price informativeness by creating a bias towards the private information of agents with smaller trading costs. We then establish that this reduction in price informativeness in turn manifests itself as an informational externality: in the presence of heterogeneity, agents do not internalize the impact of their trading decisions on the information revealed to others via prices, even in competitive markets. We conclude the paper by investigating the welfare implications of market segmentation in the presence of the heterogeneity-induced informational externality.



## [Fan You](#)

### **Finite-Horizon Approximate Linear Programs for an Infinite-Horizon Revenue Management Problem**

We consider a rolling-horizon revenue management problem that can be formulated as an infinite horizon discounted cost Markov Decision Process. We consider affine and finite-horizon approximations, and show these admit compact representations that can be solved efficiently. The resulting approximations can be used to construct control policies that probabilistically allocate demand, and we use a numerical study to evaluate their performance

## [Chiara Farronato](#)

### **Consumer Protection in an Online World: When Does Occupational Licensing Matter?**

This paper examines how consumers' choices of service professionals are related to the occupational licensing status of service providers and how state-level outcomes are related to licensing regulations. The setting is a large, US online labor market platform, where professional service providers bid on individual consumers' projects. We outline an econometric approach to estimate the causal effect of a consumer being informed about the professional's licensing status and the effect of the professional's online reputation on the consumer's choice of whom to hire. Our findings suggest that a professional's reputation matters more than the professional's licensing status. We also discuss how outcomes at the state level are related to variation in licensing stringency across states.

## [Gerardo Berbeglia](#)

### **The generalized stochastic preference choice model**

We propose a new discrete choice model, which we called the Generalized Stochastic Preference (GSP) that generalizes the random utility model (RUM). The GSP is constructed by extending the stochastic preference model (another representation of RUM) in which additional consumer types are "non-rational".

As such, the (GSP) model is a non-parametric model that has the flexibility to fit complex choice behavior and can be estimated using the similar techniques as those used to estimate a RUM. The model is also amenable to sparse solutions. For example, simple GSP's can easily (and also exactly), replicate well-known choice social experiments that possess strong regularity violations which cannot be explained by any choice model contained in the RUM class.

## [Stanislav Mamonov](#)

## **Time is money: Faster delivery premiums, demand elasticity and customer outcomes in online diamond r**

We examine price premiums and demand elasticity in the context of an online diamond retailer. We find that the retailer employs tiered delivery premium pricing. We also find that ~7.5% price premium for faster delivery translates into a nearly 50% reduction in the likelihood that a diamond will be sold.

### Jiding Zhang

## **Flexibility and Relationships in Online Marketplaces: Prescriptions for Pricing and Matching**

Online marketplaces have grown and diversified as intermediaries for services, extending from the flexible (Uber and Airbnb) to those in which trust and relationships are valued (\$800B home services market). Using data from a leading online labor marketplace, we empirically relate the operational performance of the platform's pricing and matching regimes to a simple market characteristic: how clients derive comparative value from flexibility and/or relationships with service providers. Extending recent theoretical modeling in dynamic mechanism design, our structural empirical methods use large-scale market data to prescribe pricing and matching policies.

### Simon Hohberger

## **Ideas and Challenges for a State-of-the-Art Revenue Management System at the German Railway Company “Deutsche Bahn”**

Although Railway Revenue Management offers great opportunities to increase revenue, it has received little attention in terms of research during the past decades. Deutsche Bahn (DB), the German railway provider, started using revenue management techniques in 2002 to control the availability of discounted tickets. Since most of the sold tickets use more than one leg, and a leg is used by many itineraries, the RM problem has a strong network (O&D) structure. Furthermore, customers typically have several possible train connections between two cities within a day (e.g. 37 between Frankfurt and Berlin). As a traditional independent demand model cannot cover the switching effects between the product alternatives, a discrete choice model should be considered to represent the customers' behavior more precisely.

In our talk, we first briefly show the special structure of the railway network in Germany and give an overview of the current RM system. In the main part of the presentation, we will give insights into the challenges that would arise in practice for a railway company with an O&D and price-based revenue management model that includes customer choice behavior.

### Rongbing Huang

## **Flexible multi-period risk-sharing supply chain contracts under price and demand uncertainties**

We study a supply chain contract for an international supply chain. As exchange rates between two countries fluctuate, the prices paid by the buyer also fluctuate. The research is to investigate the value of the quantity flexibility supply contract when the buyer and the seller enter a risk-sharing agreement.

### [Tomasz Sliwinski](#)

#### **Scenario based stochastic dynamic price optimization**

We consider a dynamic pricing problem modeled as a multistage, scenario based stochastic optimization, with the booking horizon divided into a number of the booking periods. The scenarios are generated based on the historical data and include competition prices and price-independent numbers of incoming reservations over the remaining booking horizon. The demand response is defined over the relative price difference to the competition, and is modeled by the parametric logistic function. We apply sophisticated model extensions to aggregate multiple scenario objectives into typical two criteria or maximum regret model. We show the results of the optimization performed in the simulated environment utilizing real-live testing data sets.

### [Elaheh Fata](#)

#### **Online Type Matching for Revenue Maximization without Monge Conditions**

In today's online marketplaces, the problem of how to dynamically match agents in real-time is becoming increasingly prevalent. It captures important tradeoffs faced by the emerging sharing economy, as well as classical inventory and production decisions faced by the travel and transportation industries. The dynamic matching of agents is an inherently challenging task due to the uncertainty in the agents to arrive in the future.

We consider this problem from the perspective of a central platform who, through historical data, has already established the types of agents that could potentially arrive and the relationships between them. This information is represented in a weighted undirected graph, where the vertices represent a finite list of potential agent types, and the weight of the edge between two vertices represents the reward gained from matching two agents with those types (and the lack of an edge represents a reward of zero). Two agent types are compatible if their corresponding vertices share an edge. The platform's goal is to maximize the total reward it earns from matching compatible agents over the course of the time horizon.

Motivated by [Hu and Zhou, 2016], the goal of this paper is to leverage prior knowledge of this weighted graph to make better dynamic matching decisions. However, unlike [Hu and Zhou, 2016], who focus on necessary and sufficient conditions (called modified Monge conditions) required of the graph for a dynamic matching policy to be optimal, we focus on bounding the sub-optimality when these conditions are not satisfied. Specifically, we ask: for a given weighted graph, what is the best dynamic matching policy and performance guarantee possible,

If you have any questions or concerns during the conference, please notify Dean McKeown at (613)449-4592 or [dean.mckeown@queensu.ca](mailto:dean.mckeown@queensu.ca)

which holds under any sequence of agent arrivals? We take the same analysis approach as [Ball and Queyranne, 2009], who consider an airline seat allocation problem, and derive the best-possible reservation policy and performance guarantee, based on the potential fare prices of the flight. Our results are instead dependent on the entire weighted graph, and generalizes the results of [Ball and Queyranne, 2009].

## [Fabian Bastin](#)

### **Dynamic Discrete Choice Model for Railway Ticket Cancellation and Exchange Decisions**

The increasing use of internet as a major ticket distribution channel has resulted in passengers becoming more strategic to fare policy. This potentially induces passengers to book the ticket well in advance in order to obtain a lower fare ticket, and later adjust their ticket when they are sure about trip scheduling. This is especially true in flexible refund markets where ticket cancellation and exchange behavior has been recognized as having major impacts on revenues. In this paper, we propose an inter-temporal choice model of ticket cancellation and exchange for railway passengers where customers are assumed to be forward looking agents. A dynamic discrete choice model (DDCM) is applied to predict the timing in which ticket exchange or cancellation occurs in response to fare and trip schedule uncertainty. The problem is formulated as an optimal stopping problem, and a two steps look-ahead policy is adopted to approximate the dynamic programming problem. The approach is applied to real ticket reservation data for intercity railway trips. Estimations results indicate that the DDCM provides more intuitive results when compared to multinomial logit (MNL) models. In addition, validation results show that DDCM has better prediction capability than MNL. The approach developed here in the context of exchange and refund policies for railway revenue management can be extended and applied to other industries that operate under flexible refund policies.

## [Diwakar Gupta](#)

### **Referral Timing and Fundraising Success in Crowdfunding**

Entrepreneurs frequently leverage social contacts to generate traffic to a crowdfunding campaign. Considering campaign fundraising progress, we address the question of exactly when an entrepreneur should involve his or her social connections. A visitor, upon arriving at the campaign, may be more likely to contribute in the presence of prior contributions from others, leading to a rich-get-richer scenario; however, prior contributions may also dissuade altruistic contributors, crowding them out, resulting in diminishing marginal returns to prior capital accumulation. With these possibilities in mind, we investigate our question via a combination of empirical and analytical methods.

## [Brenda Barnes](#)

### **Bridging the Gap Between Academia and Practice for Mutual Benefit**

In this talk, I'll share my ideas for bringing practitioners and academics in our field closer together based on 30 years of combined industry/consulting experience, work with top business schools, research on industry/academic collaboration and 17 years as a Section member and former Chair. In addition, I'll present ideas and facilitate discussion on ways to increase practitioner involvement in the Section.

## [Robert Sanders](#)

### **Reducing Retailer Food Waste through Revenue Management**

I study the incentives of retailers that create food waste using a novel data set that matches a large supermarket chain's loyalty-card database, including the time stamp for customer trips, with the firm's product-level information on marginal cost, shelf life (perishability), and daily production-processes. In a descriptive analysis, I first show the firm's production of waste correlates positively with the degree of demand uncertainty and with its market power. This relationship is consistent with the classic newsvendor problem and suggests positive waste levels are an endogenously-determined outcome. I then conduct a structural analysis of the artisanal bread category to assess the supermarket chain's incentives to mitigate waste. Currently, grocery retailers have not yet adopted revenue management (intraday markdowns) due to insufficient scale. A hypothetical policy that reduces the costs of revenue management sufficiently to induce adoption would reduce total planned waste for the chain by 13%. Surprisingly, the incentive to reduce waste varies across the individual stores: in several stores, the firm endogenously increases planned waste, suggesting revenue management alone may be insufficient to curb grocery waste. I then simulate the effects of a ban recently enacted in California. The ban increases firms' waste-disposal-costs by mandating composting. For the supermarket chain, this policy increases the returns to revenue management adoption from 7% to 11% and induces the firm to use intraday markdowns, resulting in a combined waste reduction of 31%.

## [Darius Walczak](#)

### **A Product-Focused Approach to Dynamic Programming Decomposition for the Network Revenue Management and Pricing Problems**

We consider a dynamic programming formulation of the network revenue management problem. Due to the curse of dimensionality, various approximate dynamic programming methods based on the classical resource-based decomposition approach are used in practice. We revisit our earlier alternative idea of the product-based network decomposition and develop a novel dynamic model for the problem that relies on simultaneous solving of sub-problems. We present numerical analysis to compare the revenue and computation performance of the new method with the more classical ones.

## [Chinmoy Dutta](#)

### **Locality Sensitive Hashing Based Matching for Ride-Sharing**

We consider the problem of matching riders for sharing rides in a ride sharing platform. A standard approach to the ride-matching problem is to first generate all feasible matchings of rides in the match pool constrained by the occupancy limit of a car. The generated matchings are then filtered to ensure user experience constraints. Lastly, a constraint optimization problem, typically a Mixed Integer program (MIP), is solved to choose the optimal set of matchings.

A major bottleneck with the above approach is the inefficiency of the first step of generating the feasible matchings. A typical way to generate the feasible matchings is to iteratively generate the matchings of increasing sizes. For a large ride-sharing platform, it is common to have tens of thousands of ride requests in the match pool at peak times in dense regions. To be able to generate all feasible matchings thus becomes a daunting combinatorial problem which is neither scalable nor practical.

In this work, we propose a novel method to circumvent this combinatorial challenge. Our approach employs techniques from locality sensitive hashing to find similar items in a high dimensional space.

## [William L. Cooper](#)

### **Optimal Worst-Case Pricing for a Logit Demand Model with Network Effects**

Pricing for a Logit Demand Model with Network Effects

## [Ruxian Wang](#)

### **When Prospect Theory Meets Consumer Choice Models: Assortment and Pricing Management with Reference Prices**

Reference prices are the points against which consumers evaluate their willingness to pay for the products they are considering. Reference prices may arise in several scenarios and can be influenced by time or context. Repeated consumers may remember the prices encountered on the past purchase occasions; for example, historical prices or future price trajectory may also be available online for certain products. A temporal reference price is formed by past prices; consumers adapt to or assimilate past prices of the product they purchased or considered. In these cases, when evaluating a product, the consumers may take its historical prices into account, in addition to the product's attributes and current price. The literature refers to this effect as internal or temporal reference price. However, it is documented that consumers have consistently demonstrated a limited ability to recall past prices they paid before, so the accuracy of the internal reference prices may be in doubt; see, e.g., Dickson and Sawyer (1990). We consider other definitions of what constitutes a reference price. Unlike time in the past prices, context refers to the prices of other products in the same or similar category in a store at the point of purchase. Some consumers evaluate the price of a product by comparing with the prices of other products on the shelf. When the consumers make a choice among multiple similar products, the prices of other products may also be taken into account in the evaluation of a particular product. The external information on other products in this purchase environment determines what consumers

think they should pay for a particular product, which impacts their choice behavior. Similarly, this kind of effect is called external or contextual reference price. Because the formation of an external reference price involves a comparison between the price of a product and the current prices of other products, it captures the effects of this reference price across products; see, e.g., Mazumdar and Papatla (2000). How do consumers form their external reference prices? Are these reference prices mostly related to the lowest, the highest, the median, the average market prices or the price of a particular product? The lowest price may play an important role in product evaluation, because it is often featured in local newspapers or displayed in the best shelf location. The price of a particular product/brand may form another reference price in some consumer choice scenarios. In addition, assortment variety may also affect the evaluation of certain products. In this paper, we incorporate reference prices into consumer choice models, validate their importance by empirical study and investigate the impact of various external reference prices on consumer choice behavior, as well as on assortment planning and pricing. The multinomial logit model has been widely used as a model of choice behavior (e.g., Luce (1959) and McFadden (1974)), but it exhibits a restrictive substitution pattern, known as the independence of irrelevant alternatives (IIA) property, a special case of the so-called simple scalability; see Tversky (1972). Simple scalability implies the order of the choice probabilities for two alternatives – not necessarily the ratio – is independent of the offer set or the attributes of other alternatives. In particular, if one alternative is preferred to another in one scenario, it is preferred in any other scenario. However, the order of choice probabilities does not necessarily hold if a reference price exists. If the changes of the offer set or prices of existing products lead to a dramatically different reference price, the preference order may reverse, because the changes of utility or disutility due to the reference price may overwhelm others. Under many other consumer choice models including the multinomial logit and preference-based choice models, consumer surplus is always higher as more products are offered or the attractiveness of some existing products is higher by, for instance, increasing feature values or decreasing prices. Different from these choice models, the “more is not always better” phenomenon may happen under the reference-dependent choice models; that is consumer surplus may be lower due to the effects of different reference prices by changing offer set or prices of existing products. The choice models with reference prices allow more flexible substitution patterns. Our empirical study on a data set about consumer choices for ketchup brands shows that incorporating reference prices into choice models can significantly improve the goodness-of-fit and prediction accuracy of consumer choice behavior. Moreover, the reference prices (e.g., the lowest price and the store brand price) play an important role in consumer choices for ketchup brands. In addition, we also show that failure to account for reference prices may lead to substantial losses in some scenarios, where the reference price effects indeed exist. Therefore, incorporating reference prices into consumer choice models and investigating the associated assortment planning and pricing problems is necessary. In particular, if the reference price is defined by the lowest price, we show that a quasi-markup-ordered assortment, which includes products following the markup-decreasing order plus at most one additional item, is optimal. For the pricing problem, we show that the same-markup/same-price policy, which charges the same markup for high-cost products and the same price for low-cost products, is optimal. For the reference price defined by assortment variety, the same-markup/same-price/same markup policy is

optimal: high-cost products charge the same markup; medium-cost products charge the same price that is equal to the reference price; low-cost products charge another same markup. In the above discussions, consumers form their reference prices based on the offer set and product prices on a particular purchase occasion. In some other scenarios, firms may be able to choose the reference price or may have more direct influence on the reference price to some extent by, for instance, advertising or salesforce efforts. In particular, firms may advertise their products through multiple media or exert extra salesforce efforts to influence the reference price. As a result, the reference price may increase and consequently consumers' willingness to pay may also increase. However, it may be costly to change the reference price. We also characterize the optimal reference price level if firms can directly choose it (perhaps) at a cost. Reference price is a relatively new concept in the field of operations management, especially for revenue management and pricing under consumer choice models. We take the widely used multinomial logit model as a showcase to examine the effects of reference prices. Incorporating reference prices into other consumer choice models and investigating the effects analytically or empirically on choice behavior, assortment planning and pricing under consumer heterogeneity would be another interesting and useful future research topic.

## [Ilan Lobel](#)

### **Auction Design for ROI-Constrained Buyers**

We combine theory and empirics to (i) show that some buyers in online advertising markets are financially constrained and (ii) demonstrate how to design auctions that take into account such financial constraints. We use data from a field experiment where reserve prices were randomized on Google's advertising exchange. We find that, contrary to the predictions of classical auction theory, a significant set of buyers lowers their bids when reserve prices go up. We show that this behavior can be explained if we assume buyers have constraints on their minimum return on investment (ROI). We proceed to design auctions for ROI-constrained buyers. We show that optimal auctions for symmetric ROI-constrained buyers are either second-price auctions with reduced reserve prices or subsidized second-price auctions. For asymmetric buyers, the optimal auction involves a modification of virtual values. Going back to the data, we show that using ROI-aware optimal auctions can lead to large revenue gains as well as large welfare gains for buyers.

## [Renyu \(Philip\) Zhang](#)

### **Cooperation and Profit Sharing for Ride-sharing Platforms**

We propose to model the cooperation between different ride-sharing platforms using the Multinomial Logit (MNL) choice model to capture the fact that riders face several alternatives. The drivers are self-scheduled so that they can choose when and how long to work. Different platforms engage in a price competition to maximize their own profits. We also explicitly model the congestion of both platforms, and capture the practice that a platform may incur extra labor costs (i.e., additional driver wage to induce enough drivers that can cover the excessive demand) when the system is congested. The platforms engage in cooperation by introducing a new joint service (e.g., the taxi-sharing service offered by Curb and Via). We base our model on current practices in the ride-hailing industry and study the impact of introducing the cooperation partnership on different



stakeholders of the market: platforms, riders, and drivers. In particular, our model enables us to characterize conditions under which a well-designed profit sharing contract could be beneficial for platforms, riders and drivers.

## [Yannik A. Peeters](#)

### **Continuous Assortment Optimization: The extended MNL-model**

Traditionally, assortments are thought of as being a combination of discrete items. One can easily think of a situation where the choice set is actually continuous. To handle this different setting, we will consider an extension of the multinomial logit choice model. By sequentially offering assortments, we will explore the customer's preference and exploit that information in a policy that balances the exploration - exploitation trade-off.

## [Mehmet Sekip Altug](#)

### **Optimal Dynamic Allocation of Sales and Rental Inventory at a Retailer**

We consider a retailer that simultaneously sells and rents its product over a given horizon. In every period, the retailer faces uncertain demand that splits between renters and buyers based on their utility. We characterize the optimal dynamic rental allocation policy and study its properties. We propose an implementable heuristic. We then extend our results to the duopoly case and characterize the equilibrium. We derive the conditions that lead to “pure sales”, “pure rental” or “mixed” strategy equilibrium and discuss their implications.

## [Venus Lo](#)

### **Assortment Optimization for a Multi-Channel Retailer with Features-based Value Boosts and Discounts**

We consider a retailer who sells online and in a physical store offline. Customers have initial product valuations, but boost/discount values if they observe under/over-hyped product features in-store. The retailer selects an assortment to offer in-store to maximize expected revenue while managing customers' valuation. Our model organizes products on a tree so that a leaf is a product and a non-leaf vertex is a feature common to leaves in its subtree. Online and offline customers have different consideration sets but both adjust valuations based on the offline assortment. This problem is NP-hard and we present a FPTAS, which performs much better than its theoretical guarantee under computation.

## [Hamsa Bastani](#)

### **Predicting with Proxies**

An increasingly common problem with machine learning is that true outcomes of interest are costly to measure. Instead, practitioners use proxy outcomes for prediction and subsequent decision-making. For example, a content recommender may make critical decisions based on predicted click-through rates (easily available) rather than user experience (costly to measure). However, the proxy outcome may be biased and result in poor decisions. We introduce a novel estimator that leverages many proxy outcomes and a few true outcomes to provably detect and reduce bias. We evaluate our method on a real dataset on Expedia personalized hotel recommendations, and find significant improvements in predictive accuracy.

## [Jinglong Zhao](#)

### **Dynamic Pricing under a Static Calendar**

This paper pertains to the following dynamic pricing problem. A firm is endowed with a finite, discrete, unreplaceable inventory of a single item to sell over a finite time horizon. Demand is stochastic but stationary, and the demand function, which specifies the rate of demand as a function of the price offered, is given in advance. The goal is to dynamically control prices offered to maximize the firm's revenue before the time horizon or inventory runs out.

From our collaborations with a large Consumer Packaged Goods company, we have found that while they appreciate the advantages of dynamic pricing, it is operationally beneficial for them to plan out a deterministic price trajectory, i.e. a calendar with a price for every point in time, in advance. It is possible to deviate from this calendar as demand is observed, but there is a significant overhead in doing so, and thus deviation should be reserved for situations where the realized demand was drastically higher or lower than expected. Motivated by this, we formulate the dynamic pricing problem under static calendar constraints.

First, we consider the problem where zero deviation from the initial calendar is allowed (i.e. the deviation penalty is infinity), and derive structural results for the optimal calendar similar to those of Gallego and van Ryzin (1994) for the optimal pricing policy. Second, we consider this problem's adaptivity gap---the fraction of optimal revenue lost when the firm must set an initial calendar and pay a penalty for deviating from it. We show that the best-possible revenue ratio can be obtained. Another feature of our approach is that its analysis can be naturally adapted to the discrete-time setting, leading to a ratio bound which is tight with respect to both the starting inventory level and the total number of time steps. Finally, we have numerical experiments showing that the optimal calendar, and even its basic structure, is surprisingly sensitive to the exact value of the deviation penalty.

## [Bahareh Mansouri](#)

### **Suppliers' Pricing in Iterative Procurement Auctions**

Procurement combinatorial auctions are increasingly drawing attentions as tools to drive down procurement costs for the auctioneer who is in fact the buyer in a procurement setting. The majority of literature in this regards focuses on various aspects of the auctioneer's problem who in fact implements and controls the whole process. Nonetheless, studies have shown that this disproportionate focus on the auctioneer's interest leads to the

suppliers' suspicion that the buyer acts opportunistically towards the suppliers. This concern is significant since it affects the suppliers' subsequent actions to safeguard their desired returns: from reducing their product quality and service to terminating their buyer-seller relationship with the auctioneer. The main intention behind this paper is to help bridge this gap and balance the suppliers' perception of the auction unfairness by providing a novice pricing structure that slightly cuts on the auctioneer's revenue with the promise that the suppliers are not worse off as they continue to bid in auction rounds . In an iterative setting we analytically prove the convergence of the proposed auction. Additional experimental results are proposed on the price/profit dynamics of both auction parties.

## [Yufeng Cao](#)

### **Assortment Optimization for Parallel Flights under an MNL Choice Model with Cheapest Fare Spikes**

The classical multinomial logit (MNL) choice model does not capture the phenomenon that airline customers tend to choose the cheapest fare classes. We study an assortment optimization problem for parallel flights under an extended spiked-MNL model, which introduces a separate attractiveness parameter for the cheapest available fare class on each flight. We show that the corresponding optimal assortment policy selects revenue-ordered assortments. We also propose static booking limit heuristics based on deterministic approximations of the problem. We evaluate different assortment policies in numerical experiments using both synthetic and real-world data provided by an airline partner.

## [Krishnamurthy Iyer](#)

### **Signaling in Online Retail: Efficacy of Public Signals**

We study optimal signaling of inventory and demand information to customers in a two-period retail setting. With homogeneous customers, we show that the revenue-optimal signaling mechanism is public.

## [Nick Arnosti](#)

### **How (Not) to Allocate Affordable Housing**

We study several common methods for allocating affordable housing, and show three main results. First, several methods with very different descriptions yield identical outcomes. Second, there is often a tradeoff between targeting agents with the highest need, and matching agents to suitable developments. Third, match quality can be improved by selecting a small subset of winners and giving them choices of where to live.

## [Jacob Feldman](#)

### **Revenue Management versus Machine Learning: Finding optimal product displays at Alibaba**

We develop and implement a featurized and scalable multinomial logit (MNL) model to predict purchase probabilities for each arriving customer based on customer and product features. We also developed a  $O(N^2)$  algorithm to solve the corresponding assortment optimization problem with cardinality constraints. We conducted a large-scale field experiments over 40 million customers by implementing our assortment approach on Alibaba's in-store recommendation system and comparing it with Alibaba's current machine learning based approach.

## Francisco Castro

### **Surge Pricing and its Spatial Supply Response**

We study the classic sequential screening problem in the presence of ex-post participation constraints for the buyer. A leading example is the online display advertising market, in which publishers frequently cannot use up-front fees and instead use transaction-contingent fees. We establish when the optimal selling mechanism is static and buyers are not screened with respect to their interim type, or dynamic and the buyers are screened with respect to their interim type. Our main result establishes a necessary and sufficient condition under which the static contract is optimal for general distributions of ex-post values.

We completely characterize the optimal contract with binary interim types and continuum of ex-post values. If the means of the interim types are sufficiently close, then no screening is optimal. If they are sufficiently apart, then a dynamic contract becomes optimal. Importantly, the latter contract randomizes the low type buyer while giving a deterministic allocation to the high type.

## Kate Ashley

### **An empirical analysis of hotel dynamic pricing practices**

This paper examines the dynamic pricing trends observed in practice for hotels in several U.S. metropolitan areas. We characterize price variation across online booking channels, as well as intertemporal price variation across both booking date and check-in date. Our analysis suggests relatively low adoption of sophisticated intertemporal pricing schemes, despite increased industry attention to dynamic pricing strategies. We develop a simulation model to test the potential impact of price control mechanisms that allow for finer control of the rates offered across different dates of stay; preliminary results indicate that there is significant room for revenue management techniques to improve the hotelier's bottom line.

## Gregory Macnamara

### **Sequential Procurement through Contractual and Observational Learning**

We study a dynamic game of incomplete information that models the interactions between a Principal ("Buyer"), who demands the same good or service repeatedly over time, and an Agent ("Seller"), who can produce the good at a marginal cost and an average quality that are his private information. The quality of the delivered

good is stochastic in each time period and unknown in advance. Moreover, the delivered quality is not (objectively) ex post verifiable.

We characterize the Buyer-optimal mechanism and identify two forms of learning that may take place when there are two types of Sellers: high and low. We identify key parametric regimes which are characterized by different structures of the Buyer's optimal learning dynamics.

## [Kris Ferreira](#)

### **Learning to Rank an Assortment of Products**

We consider the product ranking challenge that e-tailers face when their customers typically do not have a good idea of the product assortment offered. These customers form an impression of the assortment after looking only at products ranked in the initial positions, and then decide whether they want to continue browsing all products or leave the site. We propose an online algorithm that learns consumer preferences and converges to the optimal full-information ranking. Finally, we test our algorithm at a large e-tailer.

## [A. Serdar Simsek](#)

### **Value of Promotions with Delayed Incentives: An Empirical Investigation of Gift Card Promotions**

Gift cards have become a popular vehicle for promotional campaigns run by many departmental, consumer electronic, and online retail stores. Using a proprietary data set from a large department store, we investigate how customers respond to these promotions as well as its effectiveness as a promotional vehicle for retailers.

## [Taner Bilgic](#)

### **Product Assortment Games with Effort in Online Marketplaces**

We study a duopoly where sellers are competing in product assortment and (advertising) effort decisions in an online marketplace. Sellers choose their product assortment and decide on a level of costly (advertising) effort which jointly effect their demand. Selling prices are taken as exogenous to isolate the effect of joint assortment and effort decisions. Since the online marketplace owner has full information about the sales and assortments of all sellers, we wonder whether the platform can provide assortment and advertising effort advice to its sellers as a value-added service. We show that an equilibrium in product assortment and effort exists in a duoplistic setting with non-identical sellers. Then, we build a computational setting where we give the effort opportunity exclusively to one of the sellers as if this opportunity is "won" by this seller as a result of an auction process which is not modeled. Using the computational framework, we observe that offering effort opportunity to one seller causes the marketplace to increase marketplace expected profits in both decentralized and centralized settings. The seller who has the effort opportunity has a substantial advantage against the competitor. This substantial advantage decreases if the marketplace centrally suggests (and somehow enforces) assortment and effort decisions and/or if the other seller has an exclusive product in her assortment. We also observe that competition always leads to more effort.

## [Antoine Desir](#)

### **Design of Futures Contracts for Risk-averse Online Advertisers**

Buying display ad impressions via auctions in Internet advertising exchanges comes with significant allocation and price uncertainties. We consider the problem of designing a contract to mitigate this risk. In particular, we propose augmenting the traditional auctions with the option of buying at a premium a Market-Maker contract that removes uncertainties in outcome. We rigorously analyze the equilibrium outcome in the presence of a Market-Maker contract and show how to design it to yield an improvement both in the seller's revenue and in the buyer's utilities, therefore improving the total welfare.

## [Qingchen Wang](#)

### **Revenue Management for Parking with Advanced Reservations**

We investigate dynamic pricing of reservations for parking garages. Parking spots are reusable throughout the day and customers have varying length-of-stay. We present a simulation study using real data from a parking garage operator.

## [Wen Jiao](#)

### **Allocation and Nonlinear Pricing for Capacitated Stochastic Container Leasing System with Dynamic Arrivals**

Nonlinear Pricing for Capacitated Stochastic Container Leasing System with Dynamic Arrivals

## [Huseyin Topaloglu](#)

### **Dynamic Assortment Optimization for Reusable Products with Random Usage Durations**

We consider multi-product dynamic assortment problems with reusable products, in which each arriving customer chooses a product within an offered assortment, uses the product for a random duration of time, and returns the product back to the firm to be used by other customers. The goal is to find a policy for deciding on the assortment to offer to each customer so that the total expected revenue over a finite selling horizon is maximized. The dynamic programming formulation of this problem requires a high-dimensional state variable that keeps track of the on-hand product inventories, as well as the products that are currently in use. We present a tractable approach to compute a policy that is guaranteed to obtain at least 50% of the optimal total expected revenue. This policy is based on constructing linear approximations to the optimal value functions. The approximations are computed through an efficient backward recursion over the time periods in the selling horizon. When the usage duration is infinite or follows a negative binomial distribution, we also discuss how to efficiently perform rollout on a simple static policy. Performing rollout corresponds to using separable and

nonlinear value function approximations. The resulting policy is also guaranteed to obtain at least 50% of the optimal total expected revenue. The special case of our model with infinite usage durations captures the revenue management problem under customer choice over parallel flight legs operating between the same origin-destination pair. We provide computational experiments based on simulated data for parallel flights and real parking transaction data for the city of Seattle. Our computational experiments demonstrate that the practical performance of our policies is substantially better than their performance guarantees and performing rollout yields noticeable improvements.

## [Tolga Aydinliyim](#)

### **Managing Opportunistic Returns In Retail Operations**

Considering “honest-customers” versus “renters,” we assess the merits “targeted-refunds” and “menu-of-refunds” proposals retailers use to mitigate opportunistic consumer returns. Relative to benchmarks with “no-renters” (Su, 2009) and “uniform-refunds,” we find that the menu-of-refunds proposal with a low restocking-fee can separate customer types when leftover/returned units must be significantly marked down.

## [Ruben van de Geer](#)

### **Dynamic Pricing and Learning with Competition: Insights from the Dynamic Pricing Challenge at the.**

In this talk the results of the Dynamic Pricing Challenge are presented, held on the occasion of the 17th INFORMS Revenue Management and Pricing Section Conference on June 29-30, 2017 in Amsterdam, The Netherlands.

## [Pavithra Harsha](#)

### **A machine learning approach to the price-setting newsvendor problem**

We develop a practical framework for modeling of the price-setting newsvendor problem, which includes statistical estimation and price optimization methods for estimating the optimal solutions and associated confidence intervals. The specific novelty of the framework is that the relevant statistical estimation methods are carried out in close conjunction with the requirements of the optimization problem, which we show in this context requires the estimation of three distinct aspects of the demand distribution, namely the mean, quantile and superquantile (also known as conditional value-at-risk, CVaR). We seek to estimate these quantities in practical settings which may involve large-scale data sets and where the relationship between demand and multiple covariates may exhibit conditional heteroskedasticity, with no further assumptions on the demand distribution. This motivates us to investigate different distribution-free statistical estimators, which are broadly based on generalized linear regression (GLR), mixed-quantile regression (MQR), and superquantile regression (SQR). We extend the statistical estimation methods, notably to relax the implicit assumption of homoskedasticity in earlier versions of MQR, and to obtain a novel and exact large-scale decomposition method

that is computationally efficient for SQR (these extensions are of independent interest, besides the newsvendor application discussed here). Our computational experiments indicate the importance of methods that model heteroskedasticity, the use of regularization techniques in big-data settings and suggest that quantile-based methods such as MQR and SQR provide better solutions for a wide range of demand distributions, although for certain location-scale demand distributions that are similar to the Normal distribution, GLR may be preferable.

## Christopher Amaral

### **The Impact of Discriminatory Pricing Based on Customer Risk: An Empirical Investigation using Indirect Lending**

There is strong evidence in the marketing literature that suggests that consumers often differ in their willingness to pay (Besanko, Dubé, and Gupta 2003), which highlights the potential for firms to profit by varying prices by consumer segment. While segments can be formed using various factors, one potential factor that has received less academic attention in the price discrimination literature is risk. To the extent that customers with different risk profiles differ in their willingness to pay, pricing by risk class may be advantageous. The concept of risk is especially applicable in the context of consumer credit. Before the late 1990s the majority of lenders offered approved loan applicants a single rate and as a result rejected applications from higher-risk customers (Johnson 1992). As costs associated with data storage declined and technology associated with underwriting improved, lenders started incorporating customer credit risk into loan pricing decisions (Edelberg 2006). This practice is a form of price discrimination and is referred to as “risk-based pricing”, which involves the classification of borrowers into risk segments that are priced differently. This pricing strategy allows the lender to charge higher rates to riskier borrowers as compensation for increased default risk. Prior to risk-based pricing, higher credit risk customers were rationed out of the market, since lending to them at the “house rate” would not be profitable (Bostic 2002). To the extent that risk-based pricing increases consumer welfare through a reduction in credit rationing and improves financial institution profitability and thus economic stability, price discrimination based on risk may be advantageous.

Given the importance of consumer credit pricing and the potential benefits associated with price discrimination based on risk, our primary objective is to optimize loan prices across risk segments and geographic regions over time, incorporating the three key stakeholders involved in the loan process: the lender (principal), the agent, and the consumer. Using individual-level loan application data from a North American Financial Institution, we build a four stage structural model of choice, which enables us to account for the underlying behavioural process of these three parties each of whom act with a goal to maximize their respective utilities (Chintagunta et al. 2006). Since the structural approach results in parameter estimates that are invariant to policy (Bronnenberg, Rossi, and Vilcassim 2005; Erdem and Keane 1996; Mazzeo 2006) it fits our model framework well and allows us to better understand the implications of switching from a uniform pricing policy to a discriminatory pricing policy based on risk.

Our model framework has the following structure: First, we model the financial institution’s decision to approve a loan application based on customer, product/brand, and loan characteristics. Second, we model the agent’s decision to select a loan rate to offer the customer from the combinations of discrete rates and agent incentives that are available. Third, we model the customer’s decision to accept a loan offer and purchase the product, which allows us to determine price sensitivities across risk classes. Fourth, we model the customer’s choice to



default on the loan as a function of customer rate and other loan characteristics, which is critical in accounting for price dependent risk. Failing to take price dependent risk into consideration often leads to optimized rates that are too high (Phillips, Raffard, and Solutions 2011). Finally, given the estimation results for the financial institution, agent, and consumer decisions, we optimize the loan rate and agent incentive using dynamic programming. Thus, our modeling approach integrates the decisions of the principal (i.e. financial institution), the agent, and the consumer, allowing us to optimize the incentive structure and the customer rate simultaneously.

The current research surrounding risk-based pricing is scant and empirical investigations exploring the topic are even more limited (Edelberg 2006). This is largely because credit history information is proprietary and inaccessible to academics (Getter 2006). While a few studies explore the effects of risk-based pricing (see for example Einav, Jenkins, and Levin (2013) and Phillips, Şimşek, and Van Ryzin (2015)), to the best of our knowledge, our study is the first to simultaneously account for the three key stakeholders in the process. The results suggest that consumers' price sensitivities decrease with higher levels of customer risk. This dispersion in price sensitivities across risk classes provides the financial institution with suitable conditions to implement risk-based pricing. We examine the validity of the model and the level of confidence in the associated predictions across the potential pricing policy space and use conservative bounds on predicted demand to assess the lift from price optimization. The implementation of risk-based pricing is expected to lead to double-digit increases in the financial institution's loan profits. Additionally, consumer welfare is expected to improve since lower-credit quality customers are expected to experience increases in access to credit.

## [Amir Farshbaf Geranmayeh](#)

### **Coordinating advertising and pricing in a manufacturer - retailer channel in presence of strategic customer**

Pricing and Advertising in a Supply Chain in Presence of Strategic Consumers

## [Stefano Vaccari](#)

### **Information Aggregation via Consumer Reviews in Online Marketplaces**

We analyze of a model of social learning from online consumer reviews, when consumers can choose between multiple competing alternatives. We study how the presence of choice impacts learning outcomes and learning transients. Finally, we address different versions of the control problem that the platform can consider in order to speed up learning and purchases.

## [Ashwin Venkataraman](#)

### **A Conditional Gradient Approach for Nonparametric Estimation of Mixing Distributions**

Mixture models are very versatile tools and have been used extensively in diverse fields such as operations, marketing and econometrics. A key challenge in estimating mixture models is that the mixing distribution is

often unknown and imposing apriori parametric assumptions can lead to model misspecification issues. In this paper, we propose a new methodology for nonparametric estimation of the mixing distribution. We formulate the likelihood-based estimation problem as a constrained convex program and our key contribution is applying the conditional gradient (aka Frank-Wolfe) algorithm to solve this convex program, showing that it iteratively generates the support of the mixing distribution so that the algorithm may be terminated at the desired number of mixture components. For mixtures of logit models, we establish sublinear convergence rate of our estimator and characterize the structure of the estimated mixing distribution, showing that it corresponds closely to the notion of consideration sets in existing literature. We test our approach using two case studies on real data. We show that it outperforms the standard expectation-maximization (EM) benchmark on speed (16x faster), in-sample fit (up to 24% improvement in the log-likelihood loss), and predictive (average 27% improvement in the RMSE metric for predicting market shares) and decision accuracies (extracts ~23% more revenue). We also demonstrate, using synthetic data, that our estimator is robust to different ground-truth mixing distributions; and provide an extension that can account for endogeneity in product features.

## [Ashish Kabra](#)

### **Effect Of Incentives In Marketplaces**

Marketplace operators run aggressive incentive schemes to achieve scale, that is key to the efficacy, survival and eventual domination of a marketplace. This study quantifies and compares the effect of incentives given to the “buyer” side and “seller” side using data from a leading ride-hailing market. We build a structural model to accurately capture the driver and passenger response to incentives, and the nature of incentives. Driver effort on the platform is unobserved, for which we devise a novel local matching model based imputation method. We find that in short-term (current week) passenger incentives are more effective while the opposite is true in the long-term (next 3 months).

## [Ningyuan Chen](#)

### **Duopoly Competition with Network Effects in Discrete Choice Models**

We consider two firms that sell substitutable products to a market of network connected customers. The customers make purchases based on the multinomial logit model and the firms compete for their purchasing probabilities.

We characterize possible Nash equilibria: when the network effect is weak, there is a single symmetric equilibrium and the two firms split the market; when the network effect is strong, there exist asymmetric equilibria and one firm emerges as the dominant player; when the product quality is low and the network effect is neither too weak nor too strong, the resulting market equilibrium is never symmetric although the firms are ex ante symmetric.

## [Pu He](#)

## **Demand Estimation and Pricing in Long-tail Markets**

Long tail distributions in sales or market share data have long been an issue in empirical studies in areas such as economics, operations and marketing, and it is increasingly common nowadays with more detailed level of data available and many more products being offered in places like online retailers. In this work we first show that in long tail markets, there will be significant profit loss if biased demand estimates from traditional demand estimation are used as an input to pricing. Then we introduce a new two-step procedure to solve the problem utilizing machine learning algorithm.

### [Xiangyu Zhang](#)

## **Optimal Bidding for Highly-Valued IT Service Contracts: Theoretical Results and Practical Implications**

Information technology (IT) service providers compete to win highly-valued IT service contracts, worth tens to hundreds of millions of dollars, in a tender process. The process begins with a request for proposals (RFP) from a client detailing the required services. Each service provider then prepares a solution and price for review by the client, and the client selects a provider to win the contract. Prior literature shows that features other than price, including the service provider's relationship with the client, contribute to the client's selection decision and may even exceed price in importance. Indeed, a good relationship with the client increases the provider's chance of winning the deal even if that provider's bidding price is not the lowest. Typically, the more deals the provider wins with a client, the better their relationship becomes. Additionally, when a deal is won, the provider installs their software and hardware infrastructure at the client's facility. This reduces the provider's costs for future deals with the same client, because some of the infrastructure setup cost will have already been paid. It also makes it more expensive for other providers to deliver services to that client in the future because it would require switching the IT infrastructure for previous installations to one that is compatible with the new provider's technology.

Thus, it might be beneficial for the provider to lower their price with the objective of improving the client relationship and installing their infrastructure, increasing their potential profits from future contracts. This might be particularly true for new clients with which the provider has not yet done business. Given all this, one would expect the myopic bidding policy, which tries to maximize the expected profit of the current deal, to set prices that are too high.

In this work, we provide theoretical and numerical results that give insight into the provider's optimal bidding strategy, and that demonstrate that this intuition is true: the optimal price is indeed smaller than the myopic price. Our approach considers the trade-off between long-term client relationships and life-time profitability. Our work contributes to and extends the literature on bidding/auctions with switching costs in two ways. First, we study a unique combination of features specific to IT service contracts. One of these unique features is the dependence of the provider cost on the prior bids with the same client. In prior literature, the provider's cost is assumed to be independent of the history with the client. Second, prior literature on the most closely related problem setting, Cosguner et al. 2016, leaves open theoretical analysis of the optimal bidding strategy, focusing exclusively on empirical results.

We model the provider-client interaction over multiple deals as a Markov Decision Process (MDP). The provider-client relationship is modeled by a score that becomes higher, in a stochastic sense, with each deal won by the provider and lower with each deal lost. The provider's cost for delivering services is determined by this relationship score and an i.i.d. set of requirements in the RFP. The probability that the provider wins a deal is also determined by this relationship, the RFP's requirements, and the price. Our analysis assumes that a deal's expected profit is unimodal in the price. Although our analysis is otherwise general, we leverage in our numerical experiments a highly accurate (>90% accuracy) classifier that provides this probability that was trained and evaluated using real data from one of the world's largest IT service providers.

In this setting, we investigate the provider's client-specific bidding strategy that maximizes the total discounted expected lifetime profit. We prove a structural result: that the bidding price under the optimal strategy is lower than the price under the myopic policy, in line with the intuition described above. We also provide empirical results using the same dataset from the large IT service provider and a solution to the aforementioned MDP that leverages a low-dimensional representation of the value function. Our numerical results verify the aforementioned theorems and quantify the value of discounting the myopically optimal price to create good relationships.

## [Vijay Kamble](#)

### **A Truth Serum for Large Scale Evaluations**

A major challenge in obtaining large-scale evaluations of products or services in reputation systems is that of eliciting honest responses from agents in the absence of verifiability. We propose a new reward mechanism with strong incentive properties applicable in a wide variety of such settings. This mechanism has a simple and intuitive output agreement structure: an agent gets a reward only if her response for an evaluation matches that of her peer. But instead of the reward being the same across different answers, it is inversely proportional to a popularity index of each answer. This index is a second order population statistic that captures how frequently two agents performing the same evaluation agree on the particular answer. Rare agreements thus earn a higher reward than agreements that are relatively more common.

In the regime where there are a large number of evaluation tasks, we show that truthful behavior is a strict Bayes-Nash equilibrium of the game induced by the mechanism. Further, we show that the truthful equilibrium is approximately optimal in terms of expected payoffs to the agents across all symmetric equilibria, where the approximation error vanishes in the number of evaluation tasks. Moreover, under a mild condition on strategy space, we show that any symmetric equilibrium that gives a higher expected payoff than the truthful equilibrium must be close to being fully informative if the number of evaluations is large. These last two results are driven by a new notion of an agreement measure that is shown to be monotonic in information loss. This notion and its properties are of independent interest.

## [Abhishek Gupta](#)

### **Selling Renewable Generation with a Penalty for Shortfall**

We consider the problem of selling renewable electricity in a two-stage market to a number of load serving entities (LSEs). Since the generation is random, there is a possibility of shortfall, that is, some of the loads will not receive the amount of electricity promised in the first stage of the market. In case of shortfall, the renewable generator promises to pay a penalty, which is linear in the shortfall, to each LSE. We derive allocation and pricing rules that induces all load serving entities to bid truthfully (dominant strategy incentive compatible) their willingness to pay per unit electricity.

## [Joseph Lai](#)

### **Implementing A Pricing Model In R Shiny**

A real-world example demonstrating the simplification of a black-box pricing model into a glass-box pricing model, and how it was easily implemented into R Shiny.

## [Maxime Cohen](#)

### **Frustration-based Promotions: Field Experiments in Ride-Sharing**

In this talk, we examine whether a firm should proactively send compensation to users who have experienced a frustration (i.e., a poor service quality). In collaboration with one of the leading ride-sharing platforms, Via, we designed and ran three field experiments to investigate how different compensation types affect the engagement of riders who experienced a frustration.

## [Ashwin Venkataraman](#)

### **A Conditional Gradient Approach for Nonparametric Estimation of Mixing Distributions**

We consider the problem of segmenting a large population of customers into non-overlapping groups with similar preferences, using diverse preference observations such as purchases, ratings, clicks, etc. over subsets of items. We focus on the setting where the universe of items is large (ranging from thousands to millions) and unstructured (lacking well-defined attributes) and each customer provides observations for only a few items. These data characteristics limit the applicability of existing techniques in marketing and machine learning. To overcome these limitations, we propose a model-based embedding technique which takes the customer observations and a probabilistic model class generating the observations as inputs, and outputs an embedding--a low-dimensional representation in Euclidean space---for each customer. We then cluster the embeddings to obtain the segments. Theoretically, we derive precise necessary and sufficient conditions that guarantee asymptotic recovery of the true segments. Empirically, we demonstrate the speed and performance of our method in two real-world case studies: (a) upto 84% improvement in accuracy of new movie recommendations on the MovieLens dataset and (b) upto 8% improvement in the performance of similar product

recommendations algorithm on an offline dataset at eBay. We show that our method outperforms standard latent class, empirical bayesian and demographic-based techniques.

## Ming Hu

### **Social Pricing in the Sharing Economy: Theory and Empirical Evidence from Airbnb**

The sharing economy blurs the line between economic and social exchanges. In this paper, we study the influence of the social utility on pricing in a sharing economy. We first offer empirical evidence that social utility impacts transaction prices in Airbnb. Given this empirical evidence, we then theoretically investigate the implications of social utility for the sharing economy participants (i.e., service user and service provider), the platform, and the social planner.

## Adam Elmachtoub

### **The Value of Personalized Pricing**

In this work, we study the value of personalized pricing under the fundamental setting of a single-product monopolist. Using only simple metrics regarding the customer valuation distribution, we provide tight, closed-form upper bounds on the value of personalized pricing. We use these bounds to study the two key assumptions underlying personalized pricing: (i) the firm can charge a distinct price to each customer and (ii) the firm can perfectly predict customer valuations.

## Jue Wang

### **Dynamic Campaign Optimization**

While the effect of advertising on various mind-set metrics have received considerable attention in the marketing literature, the insights on how advertising plays a role on facilitating individual progression through the purchase funnel and what the best ad strategy is when targeting individuals with varying degrees of engagement with brands are still relatively scant. Advertising works differently across individuals and this heterogeneity in ad elasticity is expected to be prominent across customers at different stages in the customer journey.

In this paper, we study the dynamic endogenous relationship between the advertising strategy and the prospect's latent state on the path to conversion, and optimize campaign design for each individual and over time. We use individual-level data on online and offline channel-specific advertising exposure over a range of ad campaigns from a large bank in the South American market. Furthermore, we control for prospective customers' financial situation and borrowing relationship with competitors through data from a credit-reporting agency which tracks individuals' debt borrowing and repayment activities. Competitive controls have not been considered in previous literature on the topic, and help us understand how campaign effectiveness is influenced by the prospect's incumbent relationship with other financial institutions.

Our model framework follows two stages. We, first, estimate a Hidden Markov Model that can help the firm dynamically infer its relationship with the prospect from their responses to the campaigns. Given the results, we identify campaign characteristics and financial behaviors that are effective on conversion for prospects at different stages of the purchase funnel. We, then, leverage this relationship inference model to build a dynamic campaign optimization model using a Partially Observable Markov Decision Process (POMDP), which helps us select the optimal strategy to maximize the expected reward of conversion under realistic business rules.

## [Amine Allouah](#)

### **Prior-Independent Optimal Auctions**

Auctions are widely used in practice. While also extensively studied in the literature, most of the developments rely on significant assumptions about the seller's knowledge. In this work, we study the design of optimal prior-independent selling mechanisms. In particular, the seller faces buyers whose values are drawn from an unknown distribution, and only knows that the distribution belongs to a particular class of distributions. We analyze a maximin setting in which the seller attempts to optimize the worst-case fraction of revenues compared to those of an oracle with knowledge of the distribution. Our results are along two dimensions. We first characterize the structure of optimal mechanisms. Leveraging such structure, we then establish tight lower and upper bounds on performance, leading to a crisp characterization of optimal performance for a spectrum of families of distributions. In particular, we show that a second price auction is an optimal mechanism, in the maximin sense, when the seller knows that distribution of buyers belongs the monotone increasing hazard rate class of distributions and is near optimal when this distribution belongs to the regular class of distributions.

## [Santiago Balseiro](#)

### **Learning in Repeated Auctions with Budgets: Regret Minimization and Equilibrium**

In online advertising markets, advertisers often purchase ad placements through bidding in repeated auctions based on realized viewer information. We study how budget-constrained advertisers may compete in such sequential auctions in the presence of uncertainty about future bidding opportunities and competition. We formulate this problem as a sequential game of incomplete information, where bidders know neither their own valuation distribution, nor the budgets and valuation distributions of their competitors. We introduce a family of practical bidding strategies we refer to as "adaptive pacing" strategies, in which advertisers adjust their bids according to the sample path of expenditures they exhibit, and analyze the performance of these strategies in different competitive settings. We establish the asymptotic optimality of these strategies when competitors' bids are independent and identically distributed over auctions, but also when competing bids are arbitrary. When all the bidders adopt these strategies, we establish the convergence of the induced dynamics and characterize a regime (well motivated in the context of online advertising markets) under which these strategies constitute an approximate Nash equilibrium in dynamic strategies: the benefit from unilaterally deviating to other strategies, including ones with access to complete information, becomes negligible as the number of auctions and competitors grows large. This establishes a connection between regret minimization and market stability, by which advertisers can essentially follow approximate equilibrium bidding strategies that also ensure the best performance that can be guaranteed off equilibrium.

## [Divya Singhvi](#)

### **Leveraging Comparables for New Product Sales Forecasting**

We introduce a simple and intuitive approach for forecasting demand for new products. The approach identifies clusters of comparable products and accurately predicts sales for new products in these product clusters. The approach we propose is both fast, scalable, and intuitive as it mimics the current operational practice of retailers. In collaboration with two big industry partners, we provide prediction improvements of the order of 40%-50% (WMAPE) that are robust across various product categories. In addition, we also present analytical guarantees on running time and prediction accuracy.

## [Pnina Feldman](#)

### **Controlling Congestion when Consumers Choose Their Service Time**

In many service environments consumers have a need for service and choose how long they want to use it. We propose alternative mechanisms, such as imposing a limit for the time spent in service and charging a price per unit of time, for controlling congestion in such settings. We investigate their optimality for maximizing firm's profits as well as consumer and social welfare.

## [Xavier S Warnes](#)

### **The Cost of Income Guarantees in Centralized Marketplaces**

In many online systems where consumers are assigned to providers of services, a centralized planner aims to maximize instantaneous revenue while accounting for various considerations that may impact long run sustainability.

In this paper we have three main objectives. First, to quantify the magnitude of revenue losses incurred by a centralized platform when its allocations are restricted by these long run considerations. Second, to understand the instances that generate the largest losses. Third, to conduct an empirical study aimed at understanding the magnitude of these losses in real-world problems.

## [Hamsa Bastani](#)

### **Interpreting Predictive Models for Human-in-the-Loop Analytics**

Machine learning is increasingly used to make consequential decisions from real-world observational datasets, which may have significant bias or confounding. Interpretability provides a means for domain experts to understand and validate the resulting model. We propose an approach for interpreting a blackbox model by extracting a decision tree that approximates the model. Our algorithm provably avoids overfitting by actively sampling new training points using the blackbox model. We evaluate our algorithm on a random forest to predict diabetes risk on a real electronic medical record dataset. Domain experts (physicians) used our interpretation to discover an unexpected causal issue; we were able to then verify that this endogeneity indeed existed in our data, underscoring the value of interpretability.



## Bo Cowgill

### **Competition and Specificity in Market Design: Evidence from Geotargeted Advertising**

How should market designers tradeoff liquidity and specificity? We study a natural experiment in the release of a new ad targeting feature into select geographic markets using a regression discontinuity. The experiment affects the specificity advertising assets in the markets (ie, the availability of targeting a city or a zip code). We find evidence that additional specificity reduces the total number of ad impressions delivered by the platform, as advertisers concentrate bidding into fewer, targeted markets. Despite this, we find an overall positive effects on revenue grow thin the treated areas. This appears to be driven mainly by increases in click through rates, not through ad prices (which were set by auction) and by entry of new advertisers to offset the fragmentation of demand.

## Yash Kanoria

### **Dynamic Matching in School Choice: Efficient Seat Reallocation After Late Cancellations**

In many public school assignment systems throughout the United States, students are required to submit preferences over schools at an early stage, when they do not know their options outside of the public school system. Consequently, a significant fraction of the students do not use their allotted public school seat, leading to significant inefficiency. In the NYC public high school system, over 80,000 students are assigned a seat each year in March, and about 10% of these students choose not to attend a public school in September.<sup>1</sup> Moreover, schools find out about many of these vacated seats only when students do not show up when classes begin, and then reassign students via decentralized mechanisms that run months into the school year. A well-designed reassignment process, run after students learn about their outside options, could lead to significant gains in overall welfare. Yet no systematic way of reassigning students to unused seats has been proposed in the literature. Our goal is to design an explicit reassignment mechanism run at a late stage of the matching process that reassigns students efficiently. We consider a two-stage model of school assignment with finitely many schools, where students learn their outside option after the first round assignment and need to be reassigned. Student first round assignments are obtained via the standard Deferred Acceptance mechanism (DA), with school preferences given by weak priorities with a single tie-breaking lottery across all schools.<sup>2</sup> Afterwards, some students may be presented with better outside options (such as admission to a private school) and may no longer be interested in the seat allotted to them. In the second round, students are invited to re-submit their (new) ordinal preferences over schools. The goal is to reassign the seats so that the resulting assignment is efficient, and so that the overall (two-stage) mechanism is strategy-proof and does not penalize students for participating in the second round. As a significant fraction of vacated seats are reassigned only after the start of the school year, an additional goal is to ensure that the reassignment process minimizes the number of students who need to be reassigned. We suggest a class of mechanisms with desirable efficiency and incentive properties --- the Permuted Lottery Deferred Acceptance mechanisms (PLDA) --- in which the assignments in both stages are given by DA, the initial assignment serves as a guarantee in the second round, and the lottery numbers across the two rounds are correlated. The PLDA mechanisms are natural extensions of existing assignment mechanisms and retain many of their desirable properties. In addition, the PLDA

mechanisms can be implemented by either running Deferred Acceptance again in a centralized second round with updated preferences, or via a decentralized second-round waitlist system that closely mirrors current reassignment processes. Our key insight is that within the class of PLDA mechanisms, the mechanism designer can leverage the correlation between tie-breaking lotteries to achieve different operational goals. In particular, we show that reversing the lottery between the two rounds minimizes the number of reassigned students. Our main theoretical result is that under an intuitive condition, which we term the order condition, all PLDAs produce the same distribution over the final assignment, and the reverse lottery DA minimizes reassignment. The order condition can be interpreted as all school shaving the same order of over demand in the two rounds, despite changes in student preferences, and is satisfied when dropouts are uniform across first round student preferences, or when students have common preferences. In the latter setting, if in addition there are no school priorities then our results are very intuitive: by reversing the lottery we move a few students many schools up their preference lists rather than moving many students a few schools up their preference lists, thereby eliminating unnecessary cascades of reassignment. Surprisingly, however, our theoretical results hold in a general setting with heterogeneous student preferences and arbitrary priorities at schools. We also give an axiomatic justification for the class of PLDA mechanisms: in the case with no school priorities, they are equivalent to the class of mechanisms that are two-round strategy-proof while satisfying natural efficiency and symmetry requirements. Finally, we conduct empirical investigations based on data from the New York City high school admissions system. We investigate a class of PLDAs that includes RLDA, a proxy for current practice where we rerun DA using the original lottery order (termed Forward Lottery Deferred Acceptance or FLDA), and rerunning DA using an independent random lottery. We find that all these mechanisms perform similarly in terms of allocative efficiency, but RLDA reduces the number of reassigned students significantly. For instance, in the NYC public school system data set from 2004-2005, we find that current practice (FLDA) results in about 7,200 reassignments out of a total of about 75,000 students who remained in the public school system. Switching to our proposed RLDA mechanism would more than halve the number of reassigned students to fewer than 3,100. Our findings support our theoretical results and suggest they have implications for practice. Based on our results, we have proposed the Permuted Lottery Deferred Acceptance mechanisms for reassignment in school choice to the NYC Department of Education and Boston Public Schools. Our findings have been positively received and we are in discussions about how to tailor our suggestions to suit the particular constraints of these markets.

## [Kyle Maclean](#)

### **Stacking the House: Revenue Management in Live Entertainment**

The first generation of live entertainment (e.g. theatrical, sporting, and concert events) Revenue Management (RM) consisted of applying variable pricing based on location, colloquially referred to as scaling the house. Comparatively little attention has been given to the inventory management aspects of RM. This is interesting given the capacity issues that arise in practice because of all-or-nothing group buying behavior. For instance, a large number of single seats that remain but are difficult to sell.

The purpose of this paper is to bridge this gap and propose a dynamic method for the optimal capacity control. Specifically, our objective is to analyze how venues should dynamically update the set of seats made available based on remaining time, capacity, and customer group size - which we refer to as stacking the house.

We consider the sales process of ticket sales for a specific performance. We assume a customer arrival process over a discrete number of time periods where in each period there is a known probability of a specific group size arriving. The venue's seating map is divided into price differentiated segments that have a fixed price per seat. Upon a customer arrival requesting, the revenue manager decides which segment(s) to make available, with the feasible segments containing sufficient remaining capacity for the group size. Customers, upon seeing the segment(s) made available, decide between the offered segments as well as a no-buy option. From a manager's point of view, there exists a choice function that provides the probability of a customer choosing a specific segment. This choice function depends upon the segments offered. We assume that within a given segment, customers are seated from left-to-right (or right-to-left). This is not as restrictive as it may seem, as segments need not correspond to entire physical rows. We formulate the problem as a Markov Decision Process (MDP) and provide general structural properties of the problem. To illustrate how the parameters in our model can be parameterized, we obtained the 2016 transactional sales data for an annual North American professional sporting event.<sup>1</sup>

The event, held in a large metropolitan city, annually attracts thousands of domestic and international spectators. The data spanned 32,000 sales transactions for 13 separate "shows", and included price paid, seat locations, as well as the date and time of transaction. We show how choice and group arrival probabilities can be determined from the data using straightforward techniques. The model and corresponding solution suffers from the "curse of dimensionality". To make the problem tractable, we propose and analyze a set of heuristics that work by ignoring a combination of (i): the ability for customers to make a choice, (ii) the uncertain future customer demand and (iii) the full capacity configuration. The heuristics studied include a "No Single Seat" allocation method, currently in place on the Ticket Master platform, that works by offering all seats which would not leave a single seat open. We also propose and describe a "single-segment" version of the full MDP. Finally, we propose a heuristic which pre-allocates expected demands to the seating map, opening any seats which have been pre-allocated to a specific group size. We then conduct a set of computational experiments to find the average revenue enhancement above an "offer all" policy. We find that managerially significant revenue enhancements are possible and that these occur over a wide range of demand levels and group demand distributions.

## [Roozbeh Yousefi](#)

### **Dynamic pricing of the fixed-term subscription contracts or one-time purchase offered to the strategic customers in a heterogeneous market**

Subscriptions are contracts that a company makes with its customers to deliver a service regularly, or provides access to the service for the customer and this access can be unlimited or capped per the subscription contract. The emergence of subscription services has been more pronounced as technological constraints are resolved. For example, unlimited usage on cell phone services, internet services, and digital content are now more common than before. The revenue management in subscription services must consider the rational behavior of customers. Faced with a set of different possible substitutable services, a customer carefully chooses her purchase to maximize her utility. We develop a discrete choice demand model utilizing the Generalized Extreme Value distribution of the utility to capture the strategic customer behavior. The market is divided to several segments and in each segments the customers faces the choice among waiting, purchasing a contract or a single use. The

market size is changing dynamically as new subscriptions begin and old subscriptions expire. The heterogeneity of the market is modeled by the difference in the level of strategic behavior in each segment, difference in service evaluation, purchase intensity, and correlation between the discrete choices. The service provider (the monopoly) maximizes its time discounted revenue by setting the selling prices for a contract or a single purchase at any point in time (i.e., the dynamic pricing problem). The decision problem is formulated as a continuous-time optimal control model capturing all the aspects mentioned above. The model explains the dynamics of the state variables, which include the market size of each segment and the expected utility of a customer in each segment. We derived necessary optimality conditions of the model by extending the minimum principle for multiple delayed optimal control problems presented in Göllmann and Maurer (2014). The co-states differential equations were determined and the resultant system of equations should be satisfied considering the optimal pricing policy. We demonstrate the optimal pricing results in numerical experiments. We had several observations in our experiments. It is confirmed that a two-part tariff is superior to toll based pricing or pure subscription pricing schemes. We observed that with offering short subscriptions, the dominated element of the revenue comes from single-use purchases. But it is reversed when the length of the subscription contracts increases. The sensitivity of the revenue is also analyzed upon variation of each group of customers' evaluation of the service and upon changes in the degree of being strategic in each segment of the market. Our numerical experiments also showed that more myopic customers generates more revenue. Cited Work Göllmann, L., & Maurer, H. (2014). Theory and applications of optimal control problems with multiple time-delays. *Journal of Industrial & Management Optimization*, 10(2), 413-441.

## [Ram Dhurkari](#)

### **Strategic Pricing Decision using Analytic Hierarchy Process**

Price is the only "P" of marketing mix that generates revenue and therefore how firms should set and adjust their prices in order to maximize profitability is very important. During last few decades, many quantitative techniques were developed to dynamically adjust prices so that the right prices are in the right place at the right time, through the right channel and to the right customer. However, most of these quantitative techniques are based upon various inherent assumptions and use limited number of parameters and variables. In addition, these quantitative techniques often use precise historical data but provide little scope to incorporate the preferences of the decision maker (DM) in the process of making pricing related decisions. Studies have shown that the judgments are always inferior to forecasting models because forecasting models use precise historical data. However, the quantitative model lack in capturing all the relevant realities of the environment. In certain situations, the judgments of the decision maker can give superior results. Further, quantitative models are well suited to estimate an exact price offering but judgments based models can very well be used to estimate price brackets and positioning. Setting an exact price offering often depends upon the price brackets and positioning. This is because the ultimate price figure should always align with the marketing/pricing strategy of the firm. Deciding marketing/pricing strategy of the firm require processing various quantitative and qualitative indicators of the market as well as of the firm. This requires understanding of the market through the expertise, knowledge, and preference of the pricing expert or the decision maker. However, there is no decision framework available in the literature to take a decision on the pricing strategy of a firm using the preference or the judgment of the pricing expert or the marketing manager. Using the findings of Rao and Kartono (2009), this research work proposes a decision hierarchy

that can be used to take a decision on the pricing strategy of a firm. The decision hierarchy begins with the overall objective of setting or selecting a best pricing strategy for the new/existing product or services of the firm. Down further, the decision hierarchy contains two criteria, six sub-criteria, and three possible strategic alternatives as shown in the figure 1. The two broad criteria considered here are the objectives of the firm and the market determinants. To select a best pricing strategy for the firm, the decision maker will define the relative importance between the objectives of the firm and the market determinants. Each of these two criteria is measured with the help of three sub-criteria. The decision maker will then define the relative importance of these three sub-criteria with respect to the criteria under which they are placed in the decision hierarchy. This is required because a firm may have several objectives to achieve but all will have different priorities with respect to the overall objective of the firm. Similarly, several market determinants are required to be assessed and their relative priorities are to be defined in terms of their importance in shaping the pricing strategy of the firm. Leaves of the decision hierarchy are the three possible alternatives. The decision maker will then define the relative importance of the three alternatives separately and independently with respect to each of the sub-criteria placed immediately above the alternatives in the decision hierarchy. Since the assumptions of the problem structure closely match with the axiomatic assumptions of the Analytic Hierarchy Process(AHP), a popular multi-criteria decision making method, this research also demonstrates how the AHP method can be used to solve the strategic pricing decision problem.

## [Senthil Veeraraghavan](#)

### **Do Ratings Cut Both Ways? Impact of Bilateral Ratings on Platforms.**

Traditional online platforms (e.g., Amazon Marketplace) use Unilateral Rating System (URS), in which customers rate sellers. However, sharing economy platforms (e.g., Uber, Lyft, Airbnb) have adopted Bilateral Rating System (BRS) that also allows service providers to rate customers, and even select customers based on their ratings. BRS is often purported to be better than URS, as BRS unlocks the hidden information for one-time interactions, by revealing ratings of customers to service providers before they make accept/reject decisions. We compare URS and BRS in the context of a ride-sharing service to study their impact on the decisions as well as revenue/welfare of all stakeholders. We find that BRS changes drivers' effort structure significantly, and influences the pricing policy of the platform, as well as the welfare of drivers and customers. Bilateral Ratings, despite containing more information and choice, are not necessarily better for platform revenues, and they may reduce driver revenues and consumer welfare.

## [Ankur Mani](#)

### **The Value of Price Discrimination in Large Random Networks**

The use of social network information for marketing has become prevalent in the recent years. Firms such as Gilt Groupe use information (from sites like Klout.com) for targeted advertising, steering, identifying influencers in the social network for product promotion and viral marketing, as well as price discrimination for products with externalities. The use of such marketing strategies has raised concerns from consumer agencies and government because they inherently cause inequity. In the meantime, there has been significant research on efficient marketing strategies using network effects, pricing network effects, and it seems that

for small deterministic networks there is value in using network information. However, most real networks are large and noisy; it is not clear how valuable such techniques are for large random networks and whether it is worth using them given the concerns of inequity. In this paper, we focus on the issue of price discrimination based upon position of individuals in the social network. The problem of price discrimination based upon position in the social network has received significant attention in recent years. However, the question of interest in the research has mostly been to characterize and identify optimal prices and profits in deterministic networks and the complexity of computing optimal prices. For example, these works typically suggest the firm give discounts to more central people in the network to stimulate demand from others at a marked up price, which is consistent with what many firms have been doing. However, in reality, firms are very protective about pricing decisions even though they offer different prices to different people. This lack of transparency can reduce customer satisfaction and create mistrust. In view of this dilemma, our main goal is to quantify the value of price discrimination as the difference in a monopolist's profit under two prices: (i) optimal discriminative prices, and (ii) optimal uniform prices and thus shed light on the tradeoffs of price discrimination. To be specific, we study the value of price discrimination in large random graphs, in particular Erdos-Renyi graphs. To model the scenario, we consider a sequence of problems indexed by the network size. In each problem, a monopolist sells a divisible good with positive network externalities to the consumers at individualized prices and the consumers choose their optimal consumption level based upon the price and the consumption of their peers in the social network. The social network is a Erdos-Renyi graph with certain average degree. In our work, we identify tight upper and lower bounds on asymptotic profit rates and the value of price discrimination in the size of the network. These rates depend upon the average degree of the network. Our main results suggest: 1. In all cases, the value of price discrimination is sub linear in the size of the network for all network densities suggesting that the added profits from price discrimination as a fraction of overall profits is asymptotically disappearing. 2. The value of price discrimination is the highest when the average degree of the network is independent of the size of the network. In this case, the network has a giant component with sufficient imbalance that a few stars influence many in the giant component. In this case price discrimination may be valuable for decently large networks. However, such networks are very sparse as compared to real world social networks.

3. If the average degree decreases with the size of the network then the network is very fragmented and so the influence of all nodes is limited within their components, leading to small gains from price discrimination.  
4. If the average degree increases with the size of the network (at least as the logarithm of the network size) then the network tends to be very balanced and all positions have roughly equal influence. Therefore, the value of price discrimination decreases and is asymptotically zero. To obtain these results we build upon the existing literature in random graph theory, while at the same time, we develop new insights about the count of walks of different lengths in random graphs and spectra of random graphs which could be independently interesting on its own. We also study the value of price discrimination in large scale-free networks and find that the value of price discrimination is high only if the average degree is small. As the average degree of the network increases, the value of price discrimination decreases. We also show that when the network information is noisy, simple pricing policies that use only local network information for computing individual prices are robust and provide higher profits than the optimal pricing policies that use the whole network information for computing individual prices. Our results have important managerial implications. Firms need to be more careful about using price discrimination because the value of such discriminative pricing policies under many cases may not be substantial while the inequity and the lack of transparency in pricing can lead to lower customer satisfaction and mistrust.

## **Market Failure in Kidney Exchange**

Kidney exchange markets enable approximately 800 transplants per year. Patients in this market need a kidney transplant, and most come with a living, but incompatible, donor. Transplants are organized using patient-donor swaps and chains initiated by altruistic donors. The importance of this market has grown because of a shortage of organs from deceased donors, and because current law prohibits using monetary compensation to acquire organs.

This paper evaluates the kidney exchange market in the US (after more than a decade long) and shows that, despite significant success in organizing transplants, the market suffer from two serious and fixable market failures that together result in hundreds of lost transplants per year. Indeed, our descriptive evidence suggests that the market is fragmented and operates inefficiently. We use price-theoretic arguments to explain that this inefficiency arises from two standard market failures. Each failure corresponds to a specific problem with current institutions. Both problems yield suboptimal incentives for hospitals, the key decision makers. First, large national kidney exchange platforms use inefficient mechanisms to match patients. Due to biological compatibility, there is substantial variation in the additional transplants that result from different patient and donor types being submitted to the national platform. But current mechanisms ignore these differences in social value. This limitation in the design pushes hospitals to match patients with high social value outside the platform, even when it is inefficient. Second, there is scope for agency problems: hospitals face most of the costs of participating in the national platforms but receive only a fraction of the benefits. We show that fixing these two problems can result in hundreds more transplants per year.