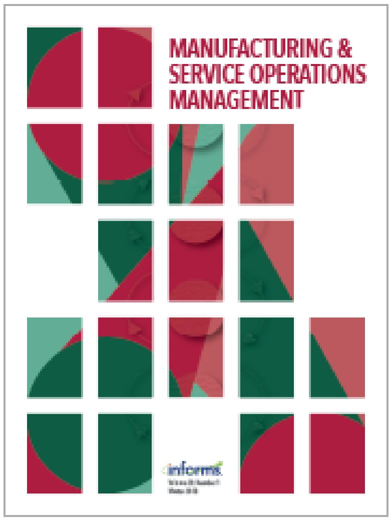

Manufacturing & Service Operations Management

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

MSOM Society Student Paper Competition: Abstracts of 2021 Winners

MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 24, No. 3, May–June 2022, pp. 1872–1874
http://pubsonline.informs.org/journal/msom
ISSN 1523-4614 (print), ISSN 1526-5498 (online)

MSOM Society Student Paper Competition: Abstracts of 2021 Winners

https://doi.org/10.1287/msom.2022.1097

Abstract. The journal is pleased to publish the abstracts of the six finalists of the 2021 Manufacturing and Service Operations Management Society's student paper competition. The 2021 prize committee was chaired by Vishal Agrawal (Georgetown University), Florin Ciocan (INSEAD), and Yanchong Zheng (Massachusetts Institute of Technology). The judges were Adam Elmachtoub, Adem Orsdemir, Amrita Kundu, Antoine Desir, Anyan Qi, Arian Aflaki, Arzum Akkas, Ashish Kabra, Bin Hu, Bora Keskin, Brent Moritz, Can Zhang, Chloe Kim Glaeser, Dan Iancu, Daniel Freund, Daniel Lin, Daniela Saban, David F. Drake, Dawson Kaaua, Divya Singhvi, Ekaterina Astashkina, Elena Belavina, Elodie Adida, Enis Kayis, Ersin Korpeoglu, Evgeny Kagan, Fabian Sting, Fanyin Zheng, Fei Gao, Fernanda Bravo, Francisco Castro, Georgina Hall, Gonzalo Romero, Guangwen Kong, Guoming Lai, Hamsa Bastani, Hessam Bavafa, Hummy Song, Ioannis (Yannis) Stamatopoulos, Ioannis Bellos, Iris Wang, Jake Feldman, Jason Acimovic, Jiankun Sun, Jiaru Bai, John Silberholz, Joline Uichanco, Jonas Oddur Jonasson, Jose Guajardo, Kaitlin Daniels, Kenan Arifoglu, Lennart Baardman, Leon Valdes, Lesley Meng, Luyi Gui, Luyi Yang, Mary Parkinson, Mazhar Arikan, Mehmet Ayvaci, Miao Bai, Michael Freeman, Ming Hu, Morvarid Rahmani, Mumin Kurtulus, Nan Yang, Nektarios Oraiopoulos, Nikhil Garg, Nil Karacaoglu, Nitin Bakshi, Nur Sunar, Olga Perdikaki, Ovunc Yilmaz, Ozan Candogan, Ozge Sahin, Panos Markou, Pascale Crama, Pengyi Shi, Pnina Feldman, Qiuping Yu, Renyu Zhang, Ruslan Momot, Ruth Beer, Ruxian Wang, Saed Alizamir, Safak Yucel, Samantha Keppler, Sanjith Gopalakrishnan, Santiago Gallino, Sarah Yini Gao, Sebastien Martin, Serdar Simsek, Seyed Emadi, Shima Nassiri, Shouqiang Wang, Siddharth Singh, Simone Marinesi, So Yeon Chun, Somya Singhvi, Song-Hee Kim, Soo-Haeng Cho, Soroush Saghafian, Sriram Dasu, Stefanus Jasin, Stephen Leider, Suresh Muthulingam, Suvrat Dhanorkar, Tian Chan, Tim Kraft, Tom TAN, Tugce Martagan, Velibor Misic, Vishal Gupta, Weiming Zhu, Xiajun Amy Pan, Xiaoshan Peng, Xiaoyang Long, Yangfang (Helen) Zhou, Yehua Wei, Yiangos Papanastasiou, Ying-Ju Chen, Yinghao Zhang, Yoni Gur, Yuqian Xu, Zhaohui (Zoey) Jiang, Zumbul Atan.

The 2021 prize winners are as follows:

First Prize
Searching for the Best Yardstick: Cost of Quality Improvements in the U.S. Hospital Industry
Jong Myeong Lim, University of Pennsylvania

Second Prizes (in alphabetical order according to the author's last name):
Optimal Pricing with a Single Point
Achraf Bahamou, Columbia University
Online Policies for Efficient Volunteer Crowdsourcing
Scott Rodilitz, Yale University

Finalists (in alphabetical order according to the author's last name):
Contextual Learning with Online Convex Optimization
Esmaeil Keyvanshokooh, University of Michigan

How Does Telemedicine Shape Physician's Practice in Mental Health?
Manqi Li, University of Michigan

Distributionally Robust Batch Contextual Bandits
Nian Si, Stanford University

Searching for the Best Yardstick: Cost of Quality Improvements in the U.S. Hospital Industry

Jong Myeong Lim
University of Pennsylvania, jongmlim@wharton.upenn.edu
Advisors: Ken Moon and Sergei Savin, University of Pennsylvania

The Hospital Value-Based Purchasing (VBP) program is Medicare's implementation of yardstick incentives applied to hospitals in the United States. Under the VBP program, 2% of all Medicare payments to hospitals, estimated to be US$1.9B in fiscal year 2021, are withheld and redistributed based on their relative performance in the quality of delivered care. We develop a dynamic equilibrium model in which hospitals are engaged in a repeated competition under yardstick incentives. Using structural estimation methods, we recover key parameters that govern hospitals' decisions to invest in quality improvement, including the financial and nonfinancial costs and uncertain outcomes of investment. By dynamically solving for hospitals' individually optimal investment policies, we estimate the trajectory of quality improvements for each hospital, including its investment decisions and quality levels throughout the implementation of the VBP program. Our counterfactual analyses explore the benefits, on the one hand, of modifying the overall size of the yardstick incentives and, on the other hand, of implementing a more focused program tailored to hospital type. We find that increasing the size of the incentives from 2% to 4% would result in an additional quality investment of US$1.2B from 2011 to 2018, leading to a 3.3% reduction in the average rate of central line–associated bloodstream infections (CLABSI). Applying yardstick incentives to the tailored hospital peer groups, even without changing the size of the incentives, can lead to an average reduction of 1.4% in the rate of CLABSI among groups of hospitals associated with the highest costs of quality investment.

Optimal Pricing with a Single Point

Achraf Bahamou
Columbia University, achraf.bahamou@columbia.edu
Advisor: Omar Besbes, Columbia University

We study the following fundamental data-driven pricing problem: how can/should a decision maker price its product based on observations at a single historic price? The decision maker optimizes over (potentially randomized) pricing policies to maximize the worst-case ratio of the revenue the decision maker can garner compared with an oracle with full knowledge of the distribution of values when the latter is only assumed to belong to a broad nonparametric set. In particular, our framework applies to the widely used regular and monotone nondecreasing hazard rate (mhr) classes of distributions. For settings in which the seller knows the exact probability of sale associated with one historic price or only a confidence interval for it, we fully characterize optimal performance and near-optimal pricing algorithms that adjust to the information at hand. The framework we develop, in general, allows us to characterize optimal performance for deterministic or more general randomized mechanisms and leads to fundamental novel insights on the value of information for pricing. As examples, against mhr distributions, we show that it is possible to guarantee 85% of oracle performance if one knows that half of the customers have bought at the historic price, and if only 1% of the customers bought, it is still possible to guarantee 51% of oracle performance.

Online Policies for Efficient Volunteer Crowdsourcing

Scott Rodilitz
University of California, Los Angeles, scott.rodilitz@gmail.com
Advisor: Vahideh Manshadi, Yale University

Nonprofit crowdsourcing platforms, such as food recovery organizations, rely on volunteers to perform time-sensitive tasks. To encourage volunteers to complete a task, platforms use nudging mechanisms to notify a subset of volunteers with the hope that at least one of them responds positively. However, because excessive notifications may reduce volunteer engagement, the platform faces a trade-off between notifying more volunteers for the current task and saving them for future ones. Motivated by these applications, we introduce the online volunteer notification problem, a generalization of online stochastic bipartite matching in which tasks arrive following a known time-varying distribution over task types. Upon arrival of a task, the platform notifies a subset of volunteers with the objective of minimizing the number of missed tasks. To capture each volunteer's adverse reaction to excessive notifications, we assume that a notification triggers a random period of inactivity, during which a volunteer ignores all notifications. However, if a volunteer is active and notified, the volunteer performs the task with a given pair-specific match probability that captures the volunteer's preference for the task. We develop an online randomized policy that achieves a constant-factor guarantee close to the upper bound we establish for the performance of any online policy. The design of our policy relies on sparsifying an ex ante feasible solution by solving a sequence of dynamic programs. Further, in collaboration with Food Rescue U.S., a volunteer-based food recovery platform, we demonstrate the effectiveness of our policy by testing it on the platform's data from various locations across the United States.

Contextual Learning with Online Convex Optimization

Esmaeil Keyvanshokooh
University of Michigan, keyvan@tamu.edu
Advisors: Mark P. Van Oyen and Cong Shi, University of Michigan

Optimizing the treatment regimen is a crucial medical decision-making problem. This can be thought of as a two-dimensional decision-making problem with a nested structure as it involves determining the optimal medication along with its optimal dose. In many cases, it is a challenge to determine the most effective medication for an individual. Even when a suitable medication is identified, dosing it optimally remains a significant challenge. Making these two nested decisions involves adaptively learning a personalized disease-progression control model. We formulate this as a new, contextual, multiarmed bandit model under a two-dimensional control with a nested structure. We develop a joint contextual learning and optimization algorithm for this model, which sequentially selects for a patient (i) the best medication based on the patient's contextual information and (ii) the corresponding dose that is optimized over the prior history of all patients who received the same medication. This algorithm synergizes the strength of the contextual bandit with online convex optimization techniques in a seamless fashion. We prove that it admits a sublinear regret, which is tight up to a logarithmic factor. We demonstrate the practicality of our learning-optimization methodology using clinical trial data on high–blood pressure patients with type 2 diabetes at high risk of cardiovascular diseases. Our analysis suggests that our method has the potential to outperform the current practice. We benchmark several policies to show the advantage of our method and provide several insights. Our framework could be widely used in many applications with two-dimensional nested decision making.

How Does Telemedicine Shape Physician's Practice in Mental Health?

Manqi Li
University of Michigan, manqili@umich.edu
Advisor: Shima Nassiri, Amazon.com

Fueled by the prevalence of digital devices and internet access, telemedicine (visiting physicians through real-time telecommunication) is becoming an important mode of service. In this work, we study whether the adoption of telemedicine has an impact on physicians' behavior in terms of scheduling related follow-up visits. To answer this question, we use a changes-in-changes model to estimate the effect of adopting telemedicine on the length of the interval between two related visits, namely, the related visit interval (RVI). Our results show that physicians schedule related visits with shorter RVIs in the short term after adopting telemedicine. As a result, physicians can admit more patients to their panel as they adopt telemedicine for a longer time. Thus, in the long run, adoption of telemedicine results in experiencing a heavier workload and scheduling related visits with longer RVIs. The adoption effect is also spilled over to the scheduling decision made during in-office visits with a decrease in RVI length in the short term and an increase in the long term. Furthermore, we show that physicians tend to schedule more frequent follow-up visits after a telemedicine visit because of the uncertainty in patients' health status in a remote visit. This study sheds light on the benefits and unintended consequences of adopting telemedicine as this mode of service is more widely utilized.

Distributionally Robust Batch Contextual Bandits

Nian Si
Stanford University, niansi@stanford.edu
Advisor: Jose Blanchet, Stanford University

Policy learning using historic observational data are an important problem that finds widespread applications. Examples include selecting offers, prices, and advertisements to send to customers as well as selecting which medication to prescribe to a patient. However, existing literature rests on the crucial assumption that the future environment in which the learned policy will be deployed is the same as the past environment that has generated the data (an assumption that is often false or too coarse an approximation). In this paper, we lift this assumption and aim to learn a distributionally robust policy with incomplete (bandit) observational data. We first present a policy evaluation procedure in the ambiguous environment. We further propose a novel learning algorithm that is able to learn a robust policy to adversarial perturbations and unknown covariate shifts with a performance guarantee based on the theory of uniform convergence. Finally, we test empirically the quality of our results in synthetic data sets and provide a comprehensive application of our methods in the context of a real-world voting data set.