M&SOM
MANUFACTURING & SERVICE OPERATIONS MANAGEMENT

Winter 2016 Special Issue on Practice-Focused Research

OM Forum—Practice-Based Research in Operations Management: What It Is, Why Do It, Related Challenges, and How to Overcome Them

Jérémie Gallien
London Business School, London NW1 4SA, United Kingdom, jgallien@london.edu

Stephen C. Graves
Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, sgraves@mit.edu

Alan Scheller-Wolf
Tepper School of Business, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, awolf@andrew.cmu.edu

Practice-based research—research performed with the intent of improving the operation of a collaborating practitioner—is an important endeavor for our field: such work may reveal new problems, interesting phenomena, and may also generate data, educational material, and solutions to important practical problems. We argue that the practical relevance of any operations management (OM) research is driven by the two dimensions of generalizability and validity, which together offer a framework for contrasting the potential strengths and weaknesses of theory-based and practice-based research. We review challenges and strategies for successfully engaging in practice-based research, including: choosing a good problem; establishing and managing a relationship with a practitioner; validation; and impact estimation. Finally, we discuss possible ways to encourage more practice-based research in OM. In particular, we argue that our field should, in general, put more emphasis on research validity.

Optimization in Online Content Recommendation Services: Beyond Click-Through Rates

Omar Besbes
Columbia University, New York, New York 10027, ob2105@gsb.columbia.edu

Yonatan Gur
Stanford University, Stanford, California 94305, ygur@stanford.edu

Assaf Zeevi
Columbia University, New York, New York 10027, assaf@gsb.columbia.edu

A new class of online services allows Internet media sites to direct users from articles they are currently reading to other content they may be interested in. This process creates a "browsing path" along which there is potential for repeated interaction between the user and the provider, giving rise to a dynamic optimization problem. A key metric that often underlies this recommendation process is the click-through rate (CTR) of candidate articles. Whereas CTR is a measure of instantaneous click likelihood, we analyze the performance improvement that one may achieve by some lookahead that accounts for the potential future path of users. To that end, by using some data of user path history at major media sites, we introduce and derive a representation of content along two key dimensions: clickability, the likelihood to click to an article when it is recommended; and engageability, the likelihood to click from an article when it hosts a recommendation. We then propose a class of heuristics that leverage both clickability and engageability, and provide theoretical support for favoring such path-focused heuristics over myopic heuristics that focus only on clickability (no lookahead). We conduct a live pilot experiment that measures the performance of a practical proxy of our proposed class, when integrated into the operating system of a worldwide leading provider of content recommendations, allowing us to estimate the aggregate improvement in clicks per visit relative to the CTR-driven current practice. The documented improvement highlights the importance and the practicality of efficiently incorporating the future path of users in real time.

Proactive Customer Education, Customer Retention, and Demand for Technology Support: Evidence from a Field Experiment

German F. Retana
INCAE Business School, Alajuela, Costa Rica, gf@incae.edu

Chris Forman, D. J. Wu
Georgia Institute of Business, Atlanta, Georgia 30308
{chris.forman@scheller.gatech.edu, dj.wu@scheller.gatech.edu}

Do service provider efforts to educate customers influence customer outcomes? We analyze the outcome of a field experiment executed by a major public cloud infrastructure services provider in 2011. Out of 2,673 customers who adopted the service during the experiment, 366 received a service intervention: an engagement through which the provider offered initial guidance on how to use basic features of the service. Before execution, it was unclear if this proactive customer education would have positive or negative effects on customer retention and demand for technology support. We show the treatment reduces by half the number of customers who churn from the service during the first week. Further, treated customers ask 19.55% fewer questions during the first week of their tenure than the controls. Although the treatment's effects decay within one week, we show that such proactive customer education can have significant economic benefits for the provider. In particular, we find that treated customers increase their accumulated usage of the service by 46.57% in the eight months after sign-up. Finally, we provide evidence that the effects of the treatment are strongest among customers who have less experience with the provider.

Pricing Personalized Bundles: A New Approach and an Empirical Study

Zhengliang Xue
IBM T. J. Watson Research Center, Yorktown Heights, New York 10598, zxue@us.ibm.com

Zizhuo Wang
Department of Industrial and Systems Engineering, University of Minnesota, Minneapolis, Minnesota 55414, zwang@umn.edu

Markus Ettl
IBM T. J. Watson Research Center, Yorktown Heights, New York 10598, msettl@us.ibm.com

This paper studies the pricing strategies for personalized product bundles. In such problems, a seller provides a variety of products for which customers can construct a personalized bundle and send a request for quote (RFQ) to the seller. The seller, after reviewing the RFQ, has to determine a price based on which the customer either purchases the whole bundle or nothing. Such problems are faced by many companies in practice, and they are very difficult because of the potential unlimited possible configurations of the bundle and the correlations among the individual products. In this paper, we propose a novel top-down and bottom-up approach to solve this problem. In the top-down step, we decompose the bundle into each component and calibrate a value score for each component. In the bottom-up step, we aggregate the components back to the bundle, define important features of the bundle, and segment different RFQs by those bundle features as well as customer attributes. Then we estimate a utility function for each segment based on historical sales data and derive an optimal price for each incoming RFQ. We show that such a model overcomes the aforementioned difficulties and can be implemented efficiently. We test our approach using empirical data from a major information technology service provider and the test result shows that the proposed approach can improve the effectiveness of pricing significantly.

Analytics for an Online Retailer: Demand Forecasting and Price Optimization

Kris Johnson Ferreira
Technology and Operations Management Unit, Harvard Business School, Boston, Massachusetts 02163, kferreira@hbs.edu

Bin Hong Alex Lee
Engineering Systems Division, Massachusetts Institute of Technology, Cambridge, Massachusetts 02193, binhong@mit.edu

David Simchi-Levi
Engineering Systems Division Department of Civil and Environmental Engineerineering, Institute for Data, Systems, and Society, and the Operations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts 02193, dslevi@mit.edu

We present our work with an online retailer, Rue La La, as an example of how a retailer can use its wealth of data to optimize pricing decisions on a daily basis. Rue La La is in the online fashion sample sales industry, where they offer extremely limited-time discounts on designer apparel and accessories. One of the retailer's main challenges is pricing and predicting demand for products that it has never sold before, which account for the majority of sales and revenue. To tackle this challenge, we use machine learning techniques to estimate historical lost sales and predict future demand of new products. The nonparametric structure of our demand prediction model, along with the dependence of a product's demand on the price of competing products, pose new challenges on translating the demand forecasts into a pricing policy. We develop an algorithm to efficiently solve the subsequent multiproduct price optimization that incorporates reference price effects, and we create and implement this algorithm into a pricing decision support tool for Rue La La's daily use. We conduct a field experiment and find that sales does not decrease because of implementing tool recommended price increases for medium and high price point products. Finally, we estimate an increase in revenue of the test group by approximately 9.7% with an associated 90% confidence interval of [2.3%, 17.8%].

ᔕ ᔕ

Improving Store Liquidation

Nathan C. Craig
The Ohio State University, Columbus, Ohio 43210, craig.186@osu.edu

Ananth Raman
Harvard Business School, Boston, Massachusetts 02163, araman@hbs.edu

Store liquidation is the time-constrained divestment of retail outlets through an in-store sale of inventory. The retail industry depends extensively on store liquidation, both to allow managers of going concerns to divest stores in efforts to enhance performance and as a means for investors to recover capital from failed ventures. Retailers sell billions of dollars of inventory annually during store liquidations. This paper introduces the store liquidation problem to the literature and presents a technique for optimizing key store liquidation decisions, including markdowns, inventory transfers, and the timing of store closings. We propose a heuristic for solving the store liquidation problem and evaluate the performance of this method. Through applications, we show that our approach could improve net recovery on cost (i.e., the profit obtained during a liquidation stated as a percentage of the cost value of liquidated inventory) by two to five percentage points in the cases we examined. Further, we discuss ways in which current practice in store liquidation differs from the decisions identified by our method, and we trace the consequences of these differences.

ᔕ ᔕ

Stockout-Based Substitution and Inventory Planning in Textbook Retailing

Joonkyum Lee
Sogang Business School, Sogang University, Seoul 121-742, Korea, jklee@sogang.ac.kr

Vishal Gaur
Johnson Graduate School of Management, Cornell University, Ithaca, New York 14853, vg77@cornell.edu

Suresh Muthulingam
Smeal College of Business, The Pennsylvania State University, University Park, Pennsylvania 16802, suresh@psu.edu

Gary F. Swisher
The Cornell Store, Cornell University, Ithaca, New York 14853, garyswisher@cornell.edu

We demonstrate the value of utility-based choice models to estimate demand and plan inventory for new and used textbooks in the presence of consumer choice and stockout-based substitution at a university textbook retailer. Demand information is censored, the exact time of stockout is not observed, and the short selling season often does not allow for replenishment. Using data for 26,749 book titles from 2007 to 2011 and a simulation experiment calibrated on real data, we show that an attribute-based choice model generates accurate demand estimates (mean absolute percentage error less than 1%) even when nearly 90% of the textbooks in the fit sample experience stockout. This performance is driven by the heterogeneity of product attributes and is robust to the occurrence of product returns. We implement this model at the bookstore in a controlled field experiment and obtain over 10% increase in profit. The results show that accounting for asymmetric and stockout-based substitution in demand estimation and inventory planning enables us to make systematic corrections in inventory mix and inventory level compared to the existing process.

Setting Planned Leadtimes in Customer-Order-Driven Assembly Systems

Zümbül Atan, Ton de Kok, Nico P. Dellaert
Department of Industrial Engineering and Innovation Sciences, Eindhoven University of Technology, 5600 MB, Eindhoven, Nethlerlands {z.atan@tue.nl, a.g.d.kok@tue.nl, n.p.dellaert@tue.nl}
Richard van Boxel, Fred Janssen
ASML, 5504DR Veldhoven, Netherlands {richard.van.boxel@asml.com, fred.janssen@asml.com}

We study an assembly system with a number of parallel multistage processes feeding a multistage final assembly process. Each stage has a stochastic throughput time. We assume that the system is controlled by planned leadtimes at each stage. From these planned leadtimes the start and due times of all stages can be derived. If a job finishes at a particular stage and has to wait before the start of the next job(s), a holding cost proportional to the waiting time is incurred. A penalty cost proportional to the lateness is incurred when the last stage of the final assembly process finishes after its due time. The objective is to determine planned leadtimes for each individual stage, such that the expected cost of a customer order is minimized.

We derive the recursive equations for the tardiness and earliness at all stages and an exact expression for the expected cost. We discuss the similarity between these expressions and those for serial inventory systems. Based on this observation and a conjecture related to the generalized Newsvendor equations, we develop an iterative heuristic procedure. Comparison with a numerical optimization method confirms the accuracy of the heuristic. Finally, we discuss an application of the model to a real-life case, showing the added value of a system-wide optimization of planned leadtimes compared to current practice.

Accurate Emergency Department Wait Time Prediction

Erjie Ang, Sara Kwasnick, Mohsen Bayati, Erica L. Plambeck
Graduate School of Business, Stanford University, Stanford, California 94305
{erjieang@stanford.edu, kwasnick@stanford.edu, bayati@stanford.edu, elp@stanford.edu}
Michael Aratow
San Mateo Medical Center, San Mateo, California 94403, maratow@smcgov.org

This paper proposes the Q-Lasso method for wait time prediction, which combines statistical learning with fluid model estimators. In historical data from four remarkably different hospitals, Q-Lasso predicts the emergency department (ED) wait time for low-acuity patients with greater accuracy than rolling average methods (currently used by hospitals), fluid model estimators (from the service operations management literature), and quantile regression methods (from the emergency medicine literature). Q-Lasso achieves greater accuracy largely by correcting errors of underestimation in which a patient waits for longer than predicted. Implemented on the external website and in the triage room of the San Mateo Medical Center (SMMC), Q-Lasso achieves over 30% lower mean squared prediction error than would occur with the best rolling average method. The paper describes challenges and insights from the implementation at SMMC.

Optimization and Simulation of Orthopedic Spine Surgery Cases at Mayo Clinic

Asli Ozen
University of Massachusetts Amherst, Amherst, Massachusetts 01003, aslozen@gmail.com
Yariv Marmor
ORT Braude College, 2161002 Karmiel, Israel; and Mayo Clinic, Rochester, Minnesota 55905, myariv@braude.ac.il
Thomas Rohleder
3185 Rosemary Lane NE, Rochester, Minnesota 55906, trohleder@charter.net
Hari Balasubramanian
University of Massachusetts Amherst, Amherst, Massachusetts 01003, hbalasubraman@ecs.umass.edu
Jeanne Huddleston, Paul Huddleston
Mayo Clinic, Rochester, Minnesota 55905

{huddleston.jeanne@mayo.edu, huddleston.mayo@mayo.ed}

Spine surgeries tend to be lengthy (mean time of 4 hours) and highly variable (with some surgeries lasting 18 hours or more). This variability along with patient preferences driving scheduling decisions resulted in both low operating room (OR) utilization and significant overtime for surgical teams at Mayo Clinic. In this paper we discuss the development of an improved scheduling approach for spine surgeries over a rolling planning horizon. First, data mining and statistical analysis was performed using a large data set to identify categories of surgeries that could be grouped together based on surgical time distributions and could be categorized at the time of case scheduling. These surgical categories are then used in a hierarchical optimization approach with the objective of maximizing a weighted combination of OR utilization and net profit. The optimization model is explored to consider trade-offs and relationships among utilization levels, financial performance, overtime allowance, and case mix. The new scheduling approach was implemented via a custom Web-based application that allowed the surgeons and schedulers to interactively identify best surgical days with patients. A pilot implementation resulted in a utilization increase of 19% and a reduction in overtime by 10%.