Welcome / Bienvenue

Dear participants to RMP2008,

It is with great pleasure that we welcome you to this *«VIIIth INFORMS' Annual Conference on Revenue Management and Pricing»*, which reflects the multidisciplinary nature of this field, ranging from cutting edge research to industrial implementations. Slightly departing from tradition, the current edition runs over three days (two and a half, to be precise), in contrast with previous two-day meetings. This format can accomodate the increased number of presentations, while keeping with a relaxed schedule that allows for informal discussions between participants. We wish you a pleasant stay in Montreal!

Patrice Marcotte and Gilles Savard, Organizers

Chers participants à RMP2008,

Il nous fait grand plaisir de vous accueillir au *«VIIIth INFORMS' Annual Conference on Revenue Management and Pricing»*, une édition qui reflète la nature multidisciplinaire de ce champ en constante évolution, à la frontière de la recherche opérationnelle, la science de la gestion, l'économie et les mathématiques. La présente édition s'éloigne quelque peu de la tradition en s'étalant sur trois jours (deux et demi pour être précis) plutôt que deux. Ce format permet l'inclusion d'un plus grand nombre de présentations, tout en conservant un horaire léger qui laisse la part belle à des discussions informelles entre participants. Nous vous souhaitons un agréable séjour à Montréal!

Patrice Marcotte et Gilles Savard, organisateurs

Committees / Comités

Organizing Committee / Comité organisateur

Patrice Marcotte, Université de Montréal
Gilles Savard, École Polytechnique de Montréal

Scientific Committee / Comité scientifique

Christopher Anderson (Cornell)
Luce Brotcorne (Valenciennes)
Guillermo Gallego (Columbia)
Soulaymane Kachani (Columbia)
Yuri Levin (Queen's)
Costis Maglaras (Columbia)
Jeff McGill (Queen's)
Jean-François Pagé (Air Canada)
Georgia Perakis (MIT)
Ioana Popescu (INSEAD)
Mariane Riss (SNCF)
Garrett van Ryzin (Columbia)
Gustavo Vulcano (NYU)

Sponsors / Commanditaires

The organizers of the conference are grateful to our sponsors, who have contributed to the banquet (PROS), the lunches (SNCF, Air Canada), the breakfasts (ExPretio, Thalys) and the welcome cocktail (Sabre).

We also acknowledge the generous contributions of CIRRELT, GERAD, École Polytechnique and MITACS, the latter having sponsorized the five student grants.

Les organisateurs de la conférence sont reconnaissants aux commanditaires qui ont contribué au banquet (PROS), aux dîners (SNCF, Air Canada), aux déjeuners (ExPretio, Thalys) et au cocktail de bienvenue (Sabre).

Nous remercions également le CIRRELT, le GERAD, l'École Polytechnique de Montréal et tout particulièrement MITACS pour avoir commandité les cinq bourses d'étudiants.

Program Outline / Résumé du programme

	WEDNESDAY	THURSDAY	FRIDAY
7:30		Breakfast (*Agora*)	
8:00			Breakfast (*Agora*)
8:30		TRACK II II.1 Optimization Models (*S1-125*) II.2 Industry I (*S1-131*) II.3 Dynamic Pricing I (*S1-139*)	
9:00			TRACK V V.1 Mathematical Methods (*S1-111*) V.2 Pricing (*S1-131*) V.3 Customer Choice II (*S1-139*)
9:30			
10:00	Registration (*Agora*)	Break (*Agora*)	
10:30		TRACK III III.1 Game Theory II (*S1-125*) III.2 Industry II (*S1-131*) III.3 Dynamic Pricing II (*S1-139*)	Break (*Agora*)
11:00			TRACK VI VI.1 Statistics I (*S1-111*) VI.2 Retail (*S1-131*) VI.3 Customer Choice III (*S1-139*)
11:30			
12:00		Lunch (*Agora*)	
12:30			Lunch and address from the Board (*Agora*)
13:00			
13:30	Welcome Address (*S1-151*)	PLENARY SESSION (*S1-151*) M. Leboeuf	
14:00	PLENARY SESSION (*S1-151*) M. Boyer		TRACK VII VII.1 Statistics II (*S1-111*) VII.2 Applications (*S1-131*) VII.3 Internet (*S1-139*)
14:30		Break (*Agora*)	
15:00	Break (*Agora*)	TRACK IV IV.1 Congested Systems (*S1-125*) IV.2 Cargo and Freight (*S1-131*) IV.3 Dynamic Pricing III (*S1-139*)	
15:30	TRACK I I.1 Game Theory I (*S1-125*) I.2 Airlines I (*S1-131*) I.3 Customer Choice I (*S1-139*)		Closing Remarks (*S1-151*)
16:00			
16:30		BUSES LEAVE AT 16:45 BANQUET Musée Pointe-à-Callière 350, Place Royale Cocktail (17:30) Museum Visit (18:30) Banquet (19:30)	
17:00			
17:30	Cocktail (*Agora*)		
18:00			
18:30			
19:00			
19:30			
20:00			
20:30			
21:00			
21:30			
22:00			

Getting to and from Pointe-à-Callière Museum

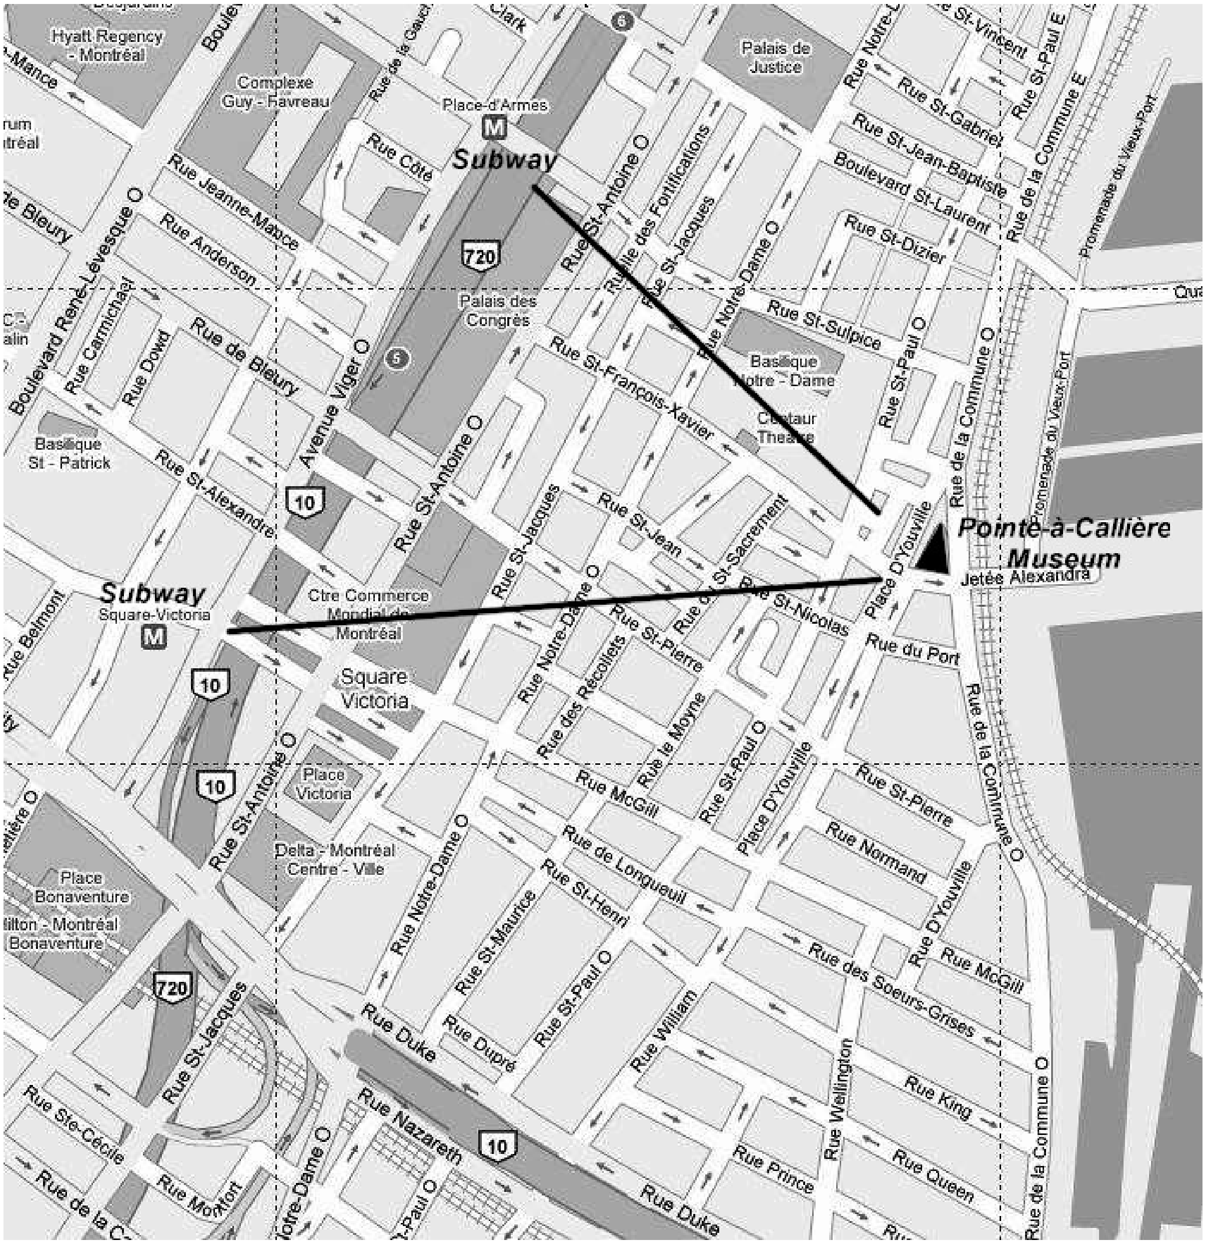

PROGRAM

WEDNESDAY, JUNE 18, 2008

Room: *Agora*

10:00 *Registration*

Room: *S1-151*

13:30 *Welcome Address*

PLENARY SESSION (Academic)

Room: *S1-151* **Chairperson:** *Gilles Savard*

14:00 *Revenue Management and Pricing under Significant Common Costs*

BOYER, MARCEL (Université de Montréal et Institut économique de Montréal)
marcel.boyer@umontreal.ca

I discuss in this paper the revenue management challenge in presence of significant common costs. Many different situations fall within such framework, the most notable case being the selling of time-related and node-related access to common infrastructures.

The common infrastructures constitute, in several ways, the skeleton of the economic and social body of our societies. They are the virtual or physical networks, on which depend the performance and competitiveness of societies and institutions and consequently the well-being of all citizens. They include among other examples those networks ensuring the proper information gathering and processing as well as communication and transmission capacity (internet, telecommunications of all types, broadband capacity, and high performance software), the transportation and distribution of energy (electricity grid, pipelines), the transportation of persons and goods (road system, public transit systems, railways, airports, seaports), as well as the provision of drinking water and the treatment of used water.

Indeed, infrastructures play two major roles in our economies: on one hand, they provide essential services that are the basis not only of productivity gains but also of poverty reduction (drinking water, electricity, mobility), and they trigger important positive externalities in all sectors of economic activity by facilitating the link between various individuals and various activities and markets, favouring social cohesion and inclusion. These positive externalities are distributed across all sectors of society and the economy through various channels related to the dynamics of demand and supply. High quality infrastructures are fundamental factors of productivity gains and increases in social wellbeing as they reduce transaction costs, shrink distances, and facilitate cultural and economic exchange between individuals, as well as trade between regions and countries. They also allow different economic actors to satisfy new demands in new places and favour the transformation of non-profitable activities into lucrative ones while increasing profit margins of existing activities. Common infrastructures are the genitors of the global village.

It is within this context that revenue management can be applied to common infrastructures. In a sense, revenue management exemplifies efficient pricing of infrastructures. I will show how revenue management of infrastructure capacity (connectivity, flexibility, safety, dependability, accessibility, speed, and user-friendliness) becomes a tool to ensure an efficient, full but with optimal congestion, use and development of the infrastructures: investment, cost allocation, and revenue management and pricing.

Room: Agora

15:00 *Coffee Break*

TRACK I

Session I.1: GAME THEORY I (15:30-17:30)

Room: *S1-125* **Chairperson:** *Dan Zhang*

Competition and Fare Class Allocation: Analysis of the Multi-Airline Problem

GAO, HUINA (University of Maryland) hgao@rhsmith.umd.edu
BALL, MICHAEL (University of Maryland) mball@rhsmith.umd.edu
KARAESMEN, ITIR (University of Maryland) ikaraes@rhsmith.umd.edu

The single-leg, fare-class allocation problem is one of the most widely studied problems in revenue management. However, a vast majority of research assumes a single decision-maker, who is risk neutral. Furthermore, an airline's fare class allocation decisions are studied in isolation from its competitors. In this paper, we consider the multi-airline problem in a competitive setting. What is unique in our model is that (i) the decision makers are risk-averse and (ii) demand is not characterized by probability distributions. We use competitive analysis of online algorithms to determine the booking control policies for the airlines. This approach determines the performance of each airline with respect to an offline optimal solution, obtained with perfect hindsight. The policies derived using competitive analysis come with worst-case performance guarantees. Assuming only the range of the fare-class demand is known, we characterize the best response-functions of the airlines using reduction on the number of possible scenarios. We show the existence and uniqueness of Nash equilibrium for a symmetric game with two airlines having only two fares. We obtain the Nash equilibrium solution in closed-form. In this particular game, competition leads to lower booking limit for the lower-fare class. We extend the analysis to the asymmetric case, and also to multiple classes and multiple airlines. Numerical experiments are used to show the impact of competition on the airlines' performance. We also use numerical examples to show the effect of ignoring competition.

Joint Memory-Dependent Pricing, Innovation, and Product Introduction Strategies

ARSLAN, HASAN (Suffolk University) harslan@suffolk.edu
KACHANI, SOULAYMANE (Columbia University) kachani@ieor.columbia.edu
SHMATOV, KYRYLO (Columbia University) kis2101@columbia.edu

We present two models that analyze how pricing, innovation, and product introduction decisions can be optimized when consumer memory impacts demand. In our first study, we consider a problem of joint product innovation and pricing for firms that produce and sell short life-cycle products in a competitive environment. We assume that the market demand for a newly introduced short life-cycle product responds to two main factors: deviation of the product's own price from a market-wide reference price, which is essentially an exponentially-weighted moving average of past prices, and deviation of the product's innovation (fashion) level from the current market average. For the first time, we incorporate in our research the joint memory effects both in pricing and the product innovation level. In terms of market positioning, we consider three types of firm strategies: (i) price discounts and promotional sales, (ii) high product turnover rate to gain a competitive edge by developing new products and restocking faster than others, and (iii) a hybrid strategy that embeds both approaches.

We study the competition among firms under a variety of market conditions through a differential game. This differential game-theoretic model accounts for different types of customer segments and different types of pricing and product innovation strategies of competing firms. Firms producing short life-cycle products may not have enough time to update their measures of market reference price and reference product innovation level. Therefore, we argue that open-loop strategies in practice fit well to such firms. We first derive open-loop Nash equilibria for an oligopoly game and show that the solution is a linear combination of exponentials. We then extend our analysis to the corresponding closed-loop Nash equilibria. We expected that open-loop and closed-loop Nash equilibrium solutions would not differ much. We confirm this intuition and quantify, both analytically and numerically, how the difference between the two equilibria

depends on the demand sensitivity parameters, and how this difference is small for reasonably low values of these parameters.

We focus on characterization of distinct classes of firms in the market in terms of their pricing and product innovation strategies. Specifically, we consider two main classes: discounter firms that control their pricing policy while keeping their product innovation level constant (no improvement of offered products and no new generation development), and high product turnover rate firms that assume a non-trivial product innovation policy while executing a constant pricing policy. We analyze specific competition among different firms that adhere to different types of pricing and product innovation strategies, and derive the corresponding Nash equilibria in these settings.

To derive insights on competition among different market positioning strategies, we conduct a case study to compare actual profit figures, achievable under different competition scenarios in different types of market, against the historically realized figures. This shows the level of improvement that may be achieved if more informed pricing and product innovation strategies are adopted. The case study is based on the example of competing retailers in the apparel industry - one that gains a competitive edge through high inventory turnover (e.g., *Inditex*) and its competitor (e.g., *H&M*), which competes largely through promoting price discounts.

In our second study, we consider joint introduction timing and pricing of successive product generations, when demand is impacted by the reference price formed by consumers. A reference price can be defined as the internal price, to which consumers compare the observed price. As customers visit the store, they develop price expectations in the form of a reference price that becomes the benchmark, against which customers compare the current price. Namely, the reference price effect implies that differences between the reference price and the current shelf price affect the demand for the product. Since reference price is formed through past price exposures of consumers, the problem of optimal price strategy for the firm results in an optimal control problem in a monopolistic environment, and a differential game in a competitive setting. Given the impact of the formed reference price on the demand for product generations, we develop optimization models for the firm to determine (i) the optimal pricing strategies for the newly introduced product generation and the incumbent product generations and (ii) the optimal introduction times for the newly developed product generations. We analyze the firm both in a monopoly and a duopoly environment.

In the monopolistic case, we explicitly characterize the optimal product introduction timing and pricing strategies for two successive generations, and propose a general quasi-analytic solution, based on the structure of the memory-dependent problem, for an arbitrary number of generations. We show that the general model can be solved much more efficiently compared to the equivalent dynamic programming model. We also extend the analysis to a duopoly environment. We deduce the equilibrium pricing and introduction time strategies in a duopoly game between two firms. We analyze both symmetric and asymmetric firms in terms of the impact of reference price on demand.

Generalized Nash Games and their Applications in Pricing

JIANG, HOUYUAN (University of Cambridge) h.jiang@jbs.cam.ac.uk
LEVIN, YURI (Queen's University) ylevin@business.queensu.ca
LEVINA, TATSIANA (Queen's University) tlevin@business.queensu.ca
MCGILL, JEFF (Queen's University) jmcgill@business.queensu.ca
NEDIAK, MIKHAIL (Queen's University) mnediak@business.queensu.ca

We present a survey of generalized Nash games that includes both theoretical developments and possible applications in pricing. The key distinction of generalized Nash games from standard ones is that the strategy sets of players depend on the strategy choices of their opponents. A simple example is capacitated pricing competition among several firms that have the option to outsource their production/service. Suppose there are m firms with capacities y_i, unit production/service costs c_i, and unit outsourcing costs $\bar{c}_i$, $i = 1, \ldots, m$. The firms simultaneously set

the prices p_i, $i = 1, \ldots, m$ which result in demands $d_i(\mathbf{p})$, $i = 1, \ldots, m$ where $\mathbf{p} = (p_1, \ldots, p_m)$. Each firm i must satisfy all of its demand and has to decide which part of this demand x_i to satisfy from its own capacity. The problem of the firm i is then

$$\begin{aligned} \max_{p_i, x_i} \quad & p_i d_i(\mathbf{p}) - c_i x_i - \bar{c}_i(d_i(\mathbf{p}) - x_i) \\ \text{s.t.} \quad & 0 \leq x_i \leq d_i(\mathbf{p}), \\ & 0 \leq x_i \leq y_i. \end{aligned}$$

This is a generalized Nash game because the prices of other firms influence the set of capacity allocation decisions of firm i.

It is interesting that while there is an extensive literature on duality, submodularity, monotone comparative statics, learning, and evolutionary approaches in standard Nash games, similar developments are missing in the analysis of generalized Nash games. We discuss these topics for the case of generalized Nash games in the context of pricing applications.

Dynamic Revenue Management in Competitive Environments

ZHANG, DAN (McGill University) dan.zhang@mcgill.ca

Oligopoly is the prevailing competitive situation in revenue management (RM) applications. A classic example is the airline industry, where RM finds its most widespread applications. It is frequently the case that a few airlines compete for customers on a given route. While price has frequently been used as the strategic decision variables in competitive revenue management models, it is also reasonable to consider quantity as the strategic decision variable. That is, we can either assume competitors respond to each other's prices or engage in a game to implement quantity strategies that jointly determine the market price.

We consider a competitive environment with *n* sellers each offering one product. Customers choose among the products following a discrete choice model, where the probability of choosing a product depends on the prices of all products. We first consider static versions of the problem where each seller has unlimited capacity and demand is deterministic. We consider two versions of the game, the price competition game and the quantity competition game. The strategic variable in the price game is the price for each seller while the strategic variable in the quantity game is the sales quantity and the price of each seller is jointly determined by the sales quantities of all sellers. We characterize and compare the Nash equilibria of the two games and show that they can be significantly different. When demand is specified by multinomial logit demand function, we show that the Nash equilibrium of the quantity competition game induces a fixed price policy for each seller which does not change as the number of players in the game increases, which is significantly different from that of the price competition game. We also provide comparative statics results. Our analysis is extended to the dynamic setting with predetermined capacity for each seller and stochastic sequential customer arrivals. Throughout, we emphasize the distinction between price competition and quantity competition and systematically analyze and compare the two. Numerical examples are given to illustrate the results.

Session I.2: AIRLINES I (15:30-17:30)

Room: *S1-131* **Chairperson:** *Huseyin Topaloglu*

Applying Support Vector Machine to Passenger Name Record Based Cancellation Forecasting for Revenue Management

ROMERO MORALES, DOLORES (University of Oxford) dolores.romero-morales@sbs.ox.ac.uk
WANG, JINGBO (University of Oxford) jingbo.wang@sbs.ox.ac.uk

A revenue management system must take into account the possibility that a booking may be cancelled, or that a booked customer may fail to show up at the time of service (no-show), which is a special case of cancellation that happens at the time of service. Forecasted rates of cancellations are needed to estimate the net demand of the service, as well as for defining overbooking levels. Accuracy of these rates is crucial to a revenue management system especially at the later part of the booking horizon. It has already been acknowledged in the literature that data mining models using Passenger Name Record (PNR) information outperform traditional seasonal average models [2, 4, 5, 7]. In our presentation at the 2007 INFORMS Pricing and Revenue Management Section [6], we reviewed the existing PNR data mining models, which mainly address the no-show case, and proposed a framework for PNR data mining based forecasting of cancellations happening at any time during the booking horizon.

However, there is still a scope for fundamental research in this area, since the data mining methods that have been applied to the cancellation forecasting problem, such as decision tree, logistic regression and naive Bayes, are not the only competitive ones, see [1]. Support Vector Machine (SVM) was originally proposed as a classification method [3]. For two classes, it maps input data points to a higher dimensional space and tries to find a separating hyperplane for the two classes, which maximizes the margin. Platt [8] used logistic regression to link a data point's distance to the hyperplane with its class probability and made it possible to use SVM for probability estimation. A recent study [1] has shown that SVM is among the most competitive methods in terms of probability estimation. In this talk, we will show how SVM can be applied to the cancellation forecasting problem. To the best of our knowledge, this is the first time SVM is applied to this problem.

The biggest challenge of applying SVM to PNR based cancellation forecasting is the computational complexity. We have obtained two real-world databases of reservation records. One is from a major hub-based European airline with around 2,000,000 entries each containing about 100 attributes. The other one is from a major UK hotel chain with nearly 500,000 entries each containing over 400 attributes. With such magnitude of the dataset, it would be extremely costly, if not impossible, to train SVMs. As a result, the challenge of refining and speeding-up SVM but at the same time preserving its excellent performance arises.

In this study, we propose an algorithm that accelerates the SVM training process by transforming the huge original dataset into a much more compact form. The "accelerated" SVM is tested on the two above mentioned datasets and outperforms decision tree and logistic regression based models. Compared with SVM built on the huge original data, the accelerated SVM leads to very similar forecasting accuracy but reduces the training time by 10 to 100 times. In addition, the compact representation of information produced by the transformation algorithm could help practitioners better understand the cancellation forecasting problem.

References

[1] R. Caruana and A. Niculescu-Mizil. An empirical comparison of supervised learning algorithms. In Proceedings of the Twenty-Third International Conference on Machine Learning, pages 161-168, 2006.

[2] W. C. Chiang, J. C. H. Chen, and X. Xu. An overview of research on revenue management: current issues and future research. International Journal of Revenue Management, 1(1):97-128, 2007.

[3] N. Cristianini and J. Shawe-Taylor. An Introduction to Support Vector Machines and Other Kernel-Based Learning Methods. Cambridge University Press, Cambridge, UK, 2000.

[4] C. Hueglin and F. Vannotti. Data mining techniques to improve forecast accuracy in airline business. In Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pages 438-442, 2001.
[5] R. D. Lawrence, S. J. Hong, and J. Cherrier. Passenger-based predictive modeling of airline no-show rates. In Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pages 397-406, 2003.
[6] D. R. Morales and J. Wang. Data mining based no-show and cancellation forecasting. Presentation at the Seventh Annual INFORMS Revenue Management and Pricing Section Conference, 2007.
[7] R. Neuling, S. Riedel, and K.-U. Kalka. New approaches to origin and destination and no-show forecasting: Excavating the passenger name records treasure. Journal of Revenue and Pricing Management, 3(1):62{72, 2003.
[8] J. C. Platt. Probabilistic outputs for support vector machines and comparison to regularized likelihood methods. In A. J. Smola, P. Bartlett, B. Schäolkopf, and D. Schuurmans, editors, Advances in Large Margin Classifiers, Cambridge, MA, 2000. MIT Press.

Stochastic Approximation for Capacity Allocation on a Single Flight Leg

TOPALOGLU, HUSEYIN (Cornell University) ht88@cornell.edu
KUNNUMKAL, SUMIT (Indian School of Business) sumit_kunnumkal@isb.edu

In the setting we consider, we use c seats available on a single flight leg to satisfy the demands from n fare classes. We index the fare classes such that the demand from fare class 1 arrives first and the demand from fare class n arrives last. If we sell a seat to fare class j, then we generate a revenue of r_j. We employ the standard assumption that the revenues satisfy $0 < r_1 \leq r_2 \leq \ldots \leq r_n$ so that the demands from the cheaper fare classes arrive earlier. We let D_j be the random demand from fare class j, which takes integer values.

Dynamic Programming Formulation We let x_j be the remaining capacity just before making the decisions for fare class j and y_j be the remaining capacity just after making the decisions for fare class j. Using x_j as the state variable and y_j as the decision variable, the optimal policy can be found by solving the optimality equation

$$v_j(x_j, D_j) = \max_{[x_j - D_j]^+ \leq y_j \leq x_j} \Big\{ r_j \, [x_j - y_j] + \mathbb{E}\{v_{j+1}(y_j, D_{j+1})\} \Big\}. \tag{1}$$

It is possible to show that $\{v_j(\cdot, D_j) : j = 1, \ldots, n\}$ are piecewise-linear concave functions with points of nondifferentiability being a subset of integers. In this case, it is easy to see that the optimal policy is characterized by a set of protection levels $\{y_j^* : j = 1, \ldots, n\}$, where y_j^* can be computed as a maximizer of the piecewise-linear concave function

$$f_j(y_j) = -r_j \, y_j + \mathbb{E}\{v_{j+1}(y_j, D_{j+1})\} \tag{2}$$

over the interval $[0, c]$. This is to say that if the remaining capacity just before making the decisions for fare class j is x_j and the demand from fare class j is D_j, then it is optimal to sell $\min\{[x_j - y_j^*]^+, D_j\}$ seats to fare class j. It is also possible to show that the optimal protection levels are nested. In other words, there exist a set of optimal protection levels $\{y_j^* : j = 1, \ldots, n\}$ that satisfy $y_1^* \geq y_2^* \geq \ldots \geq y_n^*$.

Stochastic Approximation Method Our goal is to compute the optimal protection levels through sampled trajectories of the system. By (2), we can compute a stochastic subgradient of $f_j(\cdot)$ at y_j through

$$\Delta_j(y_j, D_{j+1}) = -r_j + \dot{v}_{j+1}(y_j, D_{j+1}), \tag{3}$$

where we use $\dot{v}_{j+1}(y_j, D_{j+1})$ to denote a stochastic subgradient of $\mathbb{E}\{v_{j+1}(\cdot, D_{j+1})\}$ at y_j. In this case, letting $\{y_j^k : j = 1, \ldots, n\}$ be the estimates of the optimal protection levels at iteration k, $\{D_j^k : k = 1, \ldots, n\}$ be the demand random variables at iteration k and $\{\alpha_j^k : j = 1, \ldots, n\}_k$ be a sequence of step size parameters, we can update our estimates of the optimal protection levels by

$$y_j^{k+1} = \min\left\{ \left[y_j^k + \alpha_j^k \, \Delta_j(y_j^k, D_{j+1}^k) \right]^+, c \right\},$$

where the operator $\min\{[\cdot]^+, c\}$ ensures that the estimates of the optimal protection levels always lie in the interval $[0, c]$. However, this approach is not realistic because the computation in (3) requires the knowledge of $\{v_j(\cdot, \cdot) : j = 1, \ldots, n\}$. The method we propose builds tractable approximations to the stochastic subgradients of $\{f_j(\cdot) : j = 1, \ldots, n\}$.

Since $f_j(\cdot)$ is concave and the optimal protection level y_j^* is a maximizer of this function over the interval $[0, c]$, we can use a standard argument to write (1) as

$$v_j(x_j, D_j) = \begin{cases} r_j \, D_j + \mathbb{E}\{v_{j+1}(x_j - D_j, D_{j+1})\} & \text{if } y_j^* < x_j - D_j \\ r_j \, [x_j - y_j^*] + \mathbb{E}\{v_{j+1}(y_j^*, D_{j+1})\} & \text{if } x_j - D_j \leq y_j^* \leq x_j \\ \mathbb{E}\{v_{j+1}(x_j, D_{j+1})\} & \text{if } x_j < y_j^*. \end{cases}$$

Therefore, we can compute a stochastic subgradient of $\mathbb{E}\{v_j(\cdot, D_j)\}$ at x_j through the recursion

$$\dot{v}_j(x_j, D_j) = \begin{cases} \mathbb{E}\{\dot{v}_{j+1}(x_j - D_j, D_{j+1})\} & \text{if } y_j^* < x_j - D_j \\ r_j & \text{if } x_j - D_j \leq y_j^* \leq x_j \\ \mathbb{E}\{\dot{v}_{j+1}(x_j, D_{j+1})\} & \text{if } x_j < y_j^*. \end{cases} \tag{4}$$

To construct tractable approximations to the stochastic subgradients of $\{f_j(\cdot) : j = 1, \ldots, n\}$, we *mimic* the computation in (4) by using the estimates of the optimal protection levels. In particular, letting $\{y_j^k : j = 1, \ldots, n\}$ be the estimates of the optimal protection levels at iteration k and using $\mathcal{O}(\cdot)$ to denote the operator that rounds a scalar to a nearest integer by breaking ties arbitrarily, we recursively define

$$\rho_j^k(x_j, D_j, D_{j+1}, \ldots, D_n) = \begin{cases} \rho_{j+1}^k(x_j - D_j, D_{j+1}, \ldots, D_n) & \text{if } \mathcal{O}(y_j^k) < x_j - D_j \\ r_j & \text{if } x_j - D_j \leq \mathcal{O}(y_j^k) \leq x_j \\ \rho_{j+1}^k(x_j, D_{j+1}, \ldots, D_n) & \text{if } x_j < \mathcal{O}(y_j^k). \end{cases}$$

We use $\rho_j^k(x_j, D_j, D_{j+1}, \ldots, D_n)$ to approximate $\dot{v}_j(x_j, D_j)$. Specifically, at iteration k, we replace $\dot{v}_{j+1}(y_j, D_{j+1})$ in (3) with $\rho_{j+1}^k(y_j, D_{j+1}, \ldots, D_n)$ and use

$$s_j^k(y_j, D_{j+1}, \ldots, D_n) = -r_j + \rho_{j+1}^k(y_j, D_{j+1}, \ldots, D_n)$$

to approximate a stochastic subgradient of $f_j(\cdot)$ at y_j. Therefore, we use the algorithm

$$y_j^{k+1} = \max\left\{ \min\left\{ \left[y_j^k + \alpha_j^k \, s_j^k(y_j^k, D_{j+1}^k, \ldots, D_n^k) \right]^+, c \right\}, \mathcal{O}(y_{j+1}^{k+1}) \right\} \tag{5}$$

to compute the optimal protection levels.

Main Results We have the following results. 1) We show that the iterates of the stochastic approximation method in (5) converge to a set of optimal protection levels with probability 1. 2) The way we update our estimates of the optimal protection levels in (5) ensures that $c \geq \mathcal{O}(y_1^k) \geq \mathcal{O}(y_2^k) \geq \ldots \geq \mathcal{O}(y_n^k) \geq 0$ for all $k = 1, 2, \ldots$. Therefore, our approach provides integer and nested protection levels to use "on the fly" as we search for the optimal protection levels. 3) We show that only knowing whether the demand strictly exceeds the number of seats that we make available for sale to each fare class is adequate to compute $\rho_j^k(x_j, D_j^k, D_{j+1}^k, \ldots, D_n^k)$. Thus, our approach is applicable when the demand information is censored by the seat availability. 4) Our approach provides stochastic subgradients of the value function, which may be useful to choose the capacity of the flight leg. 5) In contrast to some of the earlier literature, our approach does not require the assumption that the demand random variables are continuous.

Internet Penetration and Capacity Utilization in the US Airline Industry

DANA, JAMES (Northeastern University) j.dana@neu.edu
ORLOV, EUGENE (Compass Lexecon) eorlov@compasslexecon.com

Airline capacity utilization, or load factors, increased dramatically between 1993 and 2007, after staying fairly level for the first 15 years following deregulation. Improvements in demand forecasting, capacity management, and revenue management are potential explanations, but revenue management systems were widely adopted in the 1980's, significantly before the increase in load factor. We argue that consumers' adoption of the Internet, and their use of the Internet to investigate and purchase airline tickets, explains recent in creases in airlines' load factors. Using metropolitan area measures of Internet penetration, we find strong evidence that differences in the rate of change of Internet penetration explain differences in the rate of change of airline airport-pair load factors. We argue that these increases, and a significant part of the associated $1 billion reduction in airlines' annual costs, represent a previously unmeasured social welfare benefit of the Internet.

Session I.3: CUSTOMER CHOICE I (15:30-17:30)

Room: *S1-139* **Chairperson:** *Gustavo Vulcano*

Revenue and Response Time Management for Queueing Systems with Customer Choice

AFECHE, PHILIPP (University of Toronto) afeche@rotman.utoronto.ca
PAVLIN, MICHAEL (University of Toronto) michael.pavlin06@rotman.utoronto.ca

Motivation and Positioning

A range of service and manufacturing firms, such as Amazon, Dell or Federal Express, operate in markets with time-sensitive customers whose willingness to buy and pay for a product or service also depends on the response time between order placement and delivery. Taking advantage of significant heterogeneity in customer preferences – some value speedy service more than others – firms typically offer a price-response time menu, giving impatient customers the option to pay more for faster delivery while charging less for longer lead times. Such response time-based price and service differentiation can serve as a valuable revenue management tool. However, the joint problem of determining the revenue-maximizing price-response time menu and the corresponding capacity scheduling policy is challenging, particularly if the provider only has aggregate customer information, e.g., based on market research, but cannot tell apart individual customers. In this common scenario, all customers can choose among all available price response time menu options and do so in line with their own self-interest. The provider must take into account this customer choice behavior which significantly complicates her decisions.

While this joint problem of pricing and response-time management under customer choice has recently received considerable interest in the revenue management literature, significant gaps remain in understanding the relationship between the underlying market attributes, the key features of the optimal price-response time menu, and the resulting customer selection and segmentation structure. The work we hope to present aims to contribute to closing these gaps.

Methodological Framework

Our study focuses on the case of a make-to-order service or manufacturing operation which we model as an M/G/1 queueing system. The capacity level may be fixed or endogenously determined by the provider’s optimization. The market consists of a heterogeneous pool of small customers. The preferences of each customer are summarized by two attributes, a value or willingness to pay for immediate delivery and a delay cost parameter which specifies how delay reduces her willingness to pay. A distinctive feature of this work is that it considers how alternative specifications of the value-delay cost distribution affect a firm’s price-response time decisions. The analysis builds on a mechanism design approach and work conservation laws to formulate the price-response time menu design problem as a constrained nonlinear optimization problem, the solution of which is analyzed for selected cases of interest.

Research Questions and Findings

Existing studies typically restrict the relationship between values and delay costs in a way that implies a certain fixed ranking among customers – some are always more preferred than others. This work relaxes some of these restrictions on the value-delay cost distribution and investigates the resulting impact on the optimal price-response time menu and the resulting customer selection and segmentation structure. In particular, it investigates the relationship between the three following elements. 1. Market attributes, including the value distribution, the delay cost distribution and the correlation between the two. 2. Structure of the optimal menu, including the number of service options and the price and delay differentiation between different options. 3. Structure of customer selection and segmentation: which customers are served and which excluded? Which customers are pooled into a common service option?

This work is currently in progress. Existing results suggest that this is a fruitful project that may provide a range of unconventional results and place the results of existing studies in a more general context. For example, preliminary results suggest that it may be optimal to exclude the "middle" market under some conditions.

On "Buy Up" as a Model of Customer Choice

LI, LE (University of Minnesota) leli@me.umn.edu
COOPER, WILLIAM (University of Minnesota) billcoop@me.umn.edu

Historically, many revenue management systems have relied on quantitative models founded on the notion of exogenous demands for fare classes. In such models, customers are assumed to "belong" to particular fare classes. Before the widespread use of the internet to book tickets, such models were perhaps a close representation of reality. Recently, it has been recognized that it is important to consider developing and using models that explicitly incorporate customer choice behavior. One of the first approaches for modeling choice in revenue management, apparently dating back at least to the late 1980s, was to incorporate "buy-up" into models based upon exogenous demands.

In a "buy-up" model, customers are still assumed to belong to fare classes. However, those customers who belong to a low fare class but arrive to find the class closed, are assumed to buy up to a higher fare with a certain fixed probability. If an airline decides to use a model with buy up, it must estimate demand distributions for the fare classes, and it also must estimate buy-up probabilities.

If customers do indeed make choices among fare classes, but an airline uses a model that does not explicitly incorporate such choice behavior and that instead relies on the notion of exogenous demand (without buy up), then, not surprisingly, poor revenue management performance may result. More interestingly, if the airline subsequently uses data in an apparently reasonable way to attempt to refine estimates of the model's parameters, then the booking controls produced by such a model may in fact systematically deteriorate — rather than improve — over time. This so-called spiral down can be viewed as an effect of a modeling error — in this case, the use of a model for which no setting of the parameters can correctly represent customer behavior. In this talk we will address the extent to which models that use buy-up can mitigate such effects and the extent to which they are potentially susceptible. We consider the performance of buy-up models when they are applied to settings where customer choice behavior is actually more complex. We analyze the evolution of booking limits (produced by the model with buy up) over a sequence of flights, assuming that an airline uses a seemingly reasonable method to translate booking records into estimates of demand distributions and buy-up probabilities, which are in turn used in the determination of booking limits. The talk will focus on long-run behavior of the booking limits.

Saving Seats for Strategic Customers

LARIVIERE, MARTIN (Northwestern University) m-lariviere@northwestern.edu
CIL, EREN (Northwestern University) e-cil@kellogg.northwestern.edu

When do customers ask for service? That is a question facing many service providers. Some customers may want to reserve a spot weeks before service is needed. Others may simply walk-in. Obviously, commitments made to the former limit the provider's ability to respond to the latter and misjudging demand of either type can be costly for the firm. How to split capacity between customers that arrive at different points in time has consequently been one of the bedrock problems of revenue management.

Here we consider responding to early and late service requests when late-arriving customers must exert effort to request service, say by having to visit the provider. Further, they must request service without knowing how many seats are available or how many other walk-in customers have ventured in. From the customer's perspective, walking-in requires weighing the chance of getting a seat against the cost of asking for service. Customers must think strategically because getting a seat depends on how others act. From the provider's perspective, giving out reservations not only limits how much can be sold to walk ins, it also impacts the actual number of walk ins who request service.

While service providers ranging from hair dressers to golf courses face this challenge, we focus on restaurants and develop a simple model. A restaurant serves two segments who desire service at the same time but request service at different times. The first segment contacts the firm ahead of the service occasion while the second requires immediate service. Customers in the latter, walk-in segment incur a cost to request service and do not know how many seats are available relative to how many walk-in customers are in the market. In equilibrium, some walk-in customers may choose not to request service.

We consider two cases. In the first, early demand is more profitable but uncertain. Here, it may be optimal to set aside seats for walk-in customers. If a large number of seats are available via reservations, walk-in traffic may be insufficient to compensate for poor reservation demand realizations. The restaurateur commits to potentially turning away some valuable early service requests in order to assure a high level of walk-in traffic. The restaurateur is not guaranteed a full dining room. The optimal policy targets a probability of empty seats that increases with the profitability of the reservation segment and the walk-in segment's cost to request service.

Saving seats for strategic walk-in customers is intuitively appealing. Remarkably, this result does not necessarily carry over to our second setting which follows the assumption of Littlewood's classical model (Littlewood, 1972). Here it is the late-arriving walk-in segment that is more profitable. We show that the firm only restricts reservations when the margin on reservation customers and the walk-in segment's cost to request service are low. Otherwise, the restaurateur makes all capacity available for reservations. That is, it may be better to ignore the more valuable segment than to save seats for them when those customers behave strategically.

On the Value of Psychological Elation and Disappointment Effects on Capacity and Pricing Decisions

LIU, QIAN (Hong Kong University of Science and Technology) qianliu@ust.hk
SHUM, STEPHEN (Hong Kong University of Science and Technology) sshum@ust.hk

Incorporating customer behavior into pricing and capacity decisions has been gaining ascending attention in the academic operation management (OM) community. The traditional OM literature assumes that demand is generated exogenously, that is, customers are passive takers of market environments and operations decisions. However, more realistically, customers are involved in decision making process and behave actively in response to firms' pricing and capacity decisions. For example, faced with dynamic pricing, customers game around as adversaries and they may postpone or accelerate their purchases to maximize gains. Some recent research on this topic has employed expected utility theory, which assumes that customers are perfect maximizers of economic utility when they make decisions under uncertainty. But do customers really behave like this in practice? Experience tells us that customers are far from being perfectly rational subject to cognitive limitations and psychological biases. When customers make choices among alternatives that involve risk, utility maximization may not be a complete description of customer behavior. Faced with decisions under uncertainty, customers compare the actual outcome with prior expectations. When the actual outcome falls below expectations, disappointment results; when the actual outcome exceeds expectations, a sense of elation can result. This sense of elation or disappointment depends on the disparity between prior expectations and ex-post outcomes: the larger the disparity, the stronger the elation or disappointment. Therefore, customer decisions and evaluations are affected not only by economic payoff, but also by psychological satisfaction or disappointment. When an outcome does not match up to prior expectations, customers' perceived utility is reduced by disappointment. This concept of disappointment should be embedded in expected utility theory in a prescriptive model.

In this work, we are particularly interested in such psychological satisfaction effects on firms' pricing and capacity decisions. We consider a monopoly firm whose selling season is divided into two periods. In the first period, the firm has limited capacity and sells each unit of product at a high price. In the second period, we assume the firm has enough capacity and sells at a low price. A case in point can be new product launch. The firm has limited inventory in the introductory period and more inventory is built up in the later period. The other example is a "first-100-customers-get-30%-off" event in retailing business. We consider a large market with a deterministic size. Customers have heterogeneous valuations and unit demand for the good. Each potential customer decides when to purchase. However, due to limited capacity, a customer who attempts to purchase in period 1 may not be able to get one product. We assume that every buyer in period 2 is guaranteed to obtain one product. There is no capacity constraint in the second period.

We apply Bell's model ([1]) to each customer's purchase decision problem. The total utility perceived by each customer is determined by:

Total Utility = Economic Payoff + Psychological Satisfaction

Each customer decides when to purchase based on the perceived total utility. Given a fill-rate and pricing policies, we can characterize the optimal purchase behavior of a particular customer as a function of valuation v and psychological factor k, which is the difference in the degrees of psychological disappointment and elation effects. Our focus is to investigate the impact of this psychological elation and disappointment on the firm's capacity and pricing decisions. In particular, what is the optimal amount of rationing created in the first period? Which pricing policy should a firm adopt? The results depend on the relations between customers' valuations and their psychological factors.

We first look at the case in which psychological factors depend directly on valuations; that is, the psychological impact factor, k, is a function of the valuation, v, of the customers. We show that when k(v) is increasing, that is, the larger the valuation, the more disappointed the customer if he fails in obtaining a product early, the firm should create rationing in the first period and employ a mark-up policy. Specifically, the firm stocks such that the optimal threshold valuation is at least as large as the second period price. Under the mark-up policy, the ideal case for the firm is to sell only to customers with valuations less than the second period price only in the first period while leaving all the affordable customers to buy late at a higher price. In fact, we can derive the sufficient and necessary conditions to support the second period price as the optimal threshold value. This result is somewhat sharp. Recall that when customer purchase behavior is determined by the expected utility (economic payoff) only, under the perfect information and deterministic demand setting, the firm should either use a uniform price to serve the entire market, or at the best, employ the markdown policy to segment the market. Under the framework of economic payoff, the mark-up policy is never optimal for the firm. However, when psychological satisfaction of customers also affects their purchase decisions besides the economic payoff, this may help the firm in the sense that the mark-up policy can also be used to segment the market and generate higher revenue than the mark-down policy does.

We also study the case in which psychological factor indirectly depends on valuation. That is, the psychological factor of a customer depends on the relationship of his valuation and the price choice in the second period. We find that when high-value customers have larger psychological factor, the firm should create rationing in period 1; again, a mark-up policy can be beneficial to the firm under certain conditions, especially when high-value customers have a higher degree of disappointment and (or) customers' valuations do not decline too much over time. Otherwise, no rationing is involved.

When the valuations of customers are independent of their psychological factors, regardless of the distribution functions, it is always optimal for the firm to employ a mark-down policy. This result is somewhat surprising. It suggests that the information on psychological factors does not provide any value when making the firm's capacity and pricing decisions as long as the psychological factors are not related to valuations of customers.

References

[1] D. Bell. Disappointment in decision making under uncertainty. Operations Research, 33:1-27.

Selling with Reservations in the Presence of Strategic Consumers

Osadchiy, Nikolay (New York University) nosadchi@stern.nyu.edu
Vulcano, Gustavo (New York University) gvulcano@stern.nyu.edu

We analyze a revenue management problem in which a seller facing an arriving stream of strategic consumers operates a *pricing with reservations* scheme. Upon arrival, each consumer, trying to maximize his own surplus, must decide either to buy at the full price and get the item immediately, or to place a non-withdrawable reservation at a discount price and wait until the end of the sales season where the leftover units are allocated according to some priority: Earlier reservations first, or later reservations first.

As a benchmark, we use a two-period dynamic pricing scheme that models the usual markdown practice: In the first period, units are sold at the full price, and in the second period the price is discounted. Consumers arriving early in the sales horizon choose between buying now or waiting for the clearance season, where leftover units are allocated randomly among the consumers who decided to wait. The remaining units are depleted among new consumers that visit the store during the clearance season.

Our approach consists of two stages. First, we study structural properties of the problem, and show that the equilibrium strategy for the three versions of this game is of the threshold type, meaning that a consumer will buy at the full price only if his valuation is above a function of his arrival time. This consumer's strategy can be computed using an iterative algorithm in a function space, provably convergent under some conditions. Unfortunately, this procedure is computationally intensive. To overcome this, in the second stage we formulate an asymptotic version of each of the problems, in which the demand rate and the initial number of units grow proportionally large. We get a simple closed form for the equilibrium strategies in this regime, which are then used as approximate solutions for the original problems. Our computations show that this heuristic is accurate for mid- to big-size problems. Finally, through an extensive numerical study, we identify cases where the use of pricing with reservations leads to revenue improvements over the standard markdown practice.

Room: *Agora*

17:30 *Cocktail*

THURSDAY, JUNE 19, 2008

Room: *Agora*

7:30 *Breakfast*

TRACK II

Session II.1: OPTIMIZATION MODELS (8:30-10:00)

Room: *S1-125* **Chairperson:** *Géraldine Heilporn*

Dynamic Programming Decomposition Methods for Joint Capacity Allocation and Overbooking in Network Revenue Management

ERDELYI, ALEXANDER (Cornell University) ae69@cornell.edu
TOPALOGLU, HUSEYIN (Cornell University) ht88@cornell.edu

We propose a dynamic programming-based method to jointly make capacity allocation and overbooking decisions in network revenue management. In the setting we consider, we have a set of flight legs that can be used to satisfy the itinerary requests that arrive randomly over time. At each time period, an itinerary request arrives and we need to decide whether to accept or reject the itinerary request. An accepted itinerary request generates a revenue and becomes a reservation. At the departure time of the flight legs, each reservation shows up with a certain probability and we need to decide which reservations should be allowed boarding.

The itinerary requests arrive over the time periods $\mathcal{T} = \{\tau, \ldots, 1\}$. The flight legs depart at time period 0. The set of flight legs is $\mathcal{L}$ and the set of itineraries is $\mathcal{J}$. The capacity on flight leg i is c_i. The probability that there is a request for itinerary j at time period t is p_{jt}. If we accept a request for itinerary j, then we generate a revenue of f_j and this reservation shows up at the departure time with probability q_j. If we allow boarding to a reservation for itinerary j, then we consume a_{ij} units of capacity on flight leg i. The penalty cost of denying boarding to a reservation for itinerary j is γ_j. For brevity of the presentation, we assume that the reservations are not canceled over the time periods in $\mathcal{T}$ and there are no refunds.

Deterministic Linear Program Assuming that the itinerary requests and the reservations that show up at the departure time take on their expected values, we can formulate the problem as a linear program. We let z_j be the number of requests for itinerary j that we plan to accept and y_j be the number of reservations for itinerary j that we plan to deny boarding. In this case, the network revenue management problem can be formulated as

$$
\begin{aligned}
\max \quad & \sum_{j\in\mathcal{J}} f_j\, z_j - \sum_{j\in\mathcal{J}} \gamma_j\, y_j & & (1)\\
\text{subject to} \quad & \sum_{j\in\mathcal{J}} a_{ij}\,[q_j\, z_j - y_j] \le c_i & i \in \mathcal{L} \qquad & (2)\\
& z_j \le \sum_{t\in\mathcal{T}} p_{jt} & j \in \mathcal{J} \qquad & (3)\\
& y_j - q_j\, z_j \le 0 & j \in \mathcal{J} \qquad & (4)\\
& z_j,\, y_j \ge 0 & j \in \mathcal{J}. \qquad & (5)
\end{aligned}
$$

The problem above assumes that if we accept z_j requests for itinerary j, then $q_j\, z_j$ reservations for itinerary j show up at the departure time. There are two uses of Problem (1)-(5). First, this problem can be used to derive a decision rule to accept or reject the itinerary requests. This is essentially an extension of the widely known bid price decision rule that compares the revenue from an itinerary request with the opportunity cost of the capacities that are consumed. Second, it is possible to show that the optimal objective value of problem (1)-(5) provides an upper bound on the total expected profit of the optimal policy.

Dynamic Programming Decomposition Our goal is to build tractable *approximations to the dynamic programming formulation* of the network revenue management problem. We exploit the deterministic linear program for this purpose. Letting $\{\mu_i^* : i \in \mathcal{L}\}$ be the optimal values of the dual variables associated with constraints (2), we choose an arbitrary flight leg i and relax constraints (2) for all other flight legs by associating the dual multipliers $\{\mu_k^* : k \in \mathcal{L} \setminus \{i\}\}$. This yields the linear program

$$\max \quad \sum_{j \in \mathcal{J}} \Big[f_j - \sum_{k \in \mathcal{L} \setminus \{i\}} q_j \, a_{kj} \, \mu_k^* \Big] z_j - \sum_{j \in \mathcal{J}} \Big[\gamma_j - \sum_{k \in \mathcal{L} \setminus \{i\}} a_{kj} \, \mu_k^* \Big] y_j + \sum_{k \in \mathcal{L} \setminus \{i\}} c_k \, \mu_k^* \tag{6}$$

$$\text{subject to} \quad \sum_{j \in \mathcal{J}} a_{ij} \, [q_j \, z_j - y_j] \leq c_i \tag{7}$$

$$(3), (4), (5). \tag{8}$$

By duality theory, the optimal objective value of the problem above is equal to the optimal objective value of problem (1)-(5).

Comparing problem (1)-(5) with problem (6)-(8) and noting that we have constraint (7) only for flight leg i in problem (6)-(8), problem (6)-(8) is the deterministic linear program for the single-leg revenue management problem that takes place over flight leg i only. In this single-leg revenue management problem, if we accept a request for itinerary j, then we generate a revenue of $\tilde{f}_j^i = f_j - \sum_{k \in \mathcal{L} \setminus \{i\}} q_j \, a_{kj} \, \mu_k^*$ and if we deny boarding to a reservation for itinerary j, then we incur a penalty cost of $\tilde{\gamma}_j^i = \gamma_j - \sum_{k \in \mathcal{L} \setminus \{i\}} a_{kj} \, \mu_k^*$. On the other hand, letting x_{jt} be the total number of reservations for itinerary j at the beginning of time period t and using $x_t = \{x_{jt} : j \in \mathcal{J}\}$ to capture the state of the reservations, we can find the optimal policy for this single-leg revenue management problem by solving the optimality equation

$$u_t^i(x_t) = \sum_{j \in \mathcal{J}} p_{jt} \max\{\tilde{f}_j^i + u_{t-1}^i(x_t + e_j), u_{t-1}^i(x_t)\} + \Big[1 - \sum_{j \in \mathcal{J}} p_{jt}\Big] \, u_{t-1}^i(x_t). \tag{9}$$

As a function of x_0, if the reservations that show up at the departure time are given by the random variable $S_0(x_0) = \{S_{j0}(x_{j0}) : j \in \mathcal{J}\}$, then the cost of denying the reservations is

$$V^i(S_0(x_0)) = \min \quad \sum_{j \in \mathcal{J}} \tilde{\gamma}_j^i \, y_j \tag{10}$$

$$\text{subject to} \quad \sum_{j \in \mathcal{J}} a_{ij} \, [S_{j0}(x_{j0}) - y_j] \leq c_i \tag{11}$$

$$0 \leq y_j \leq S_{j0}(x_{j0}) \qquad j \in \mathcal{J}. \tag{12}$$

Thus, the boundary condition in the optimality equation in (9) is $u_0^i(x_0) = -\mathbf{E}\{V^i(S_0(x_0))\}$.

Main Results We have the following results. 1) Letting $\bar{0}$ be the $|\mathcal{J}|$-dimensional vector of zeros, it is possible to show that $u_\tau^i(\bar{0}) + \sum_{k \in \mathcal{L} \setminus \{i\}} c_k \, \mu_k^*$ is an upper bound on the total expected profit of the optimal policy. Furthermore, this upper bound is tighter than the one provided by the optimal objective value of problem (1)-(5). 2) The state variable in the optimality equation in (9) is $|\mathcal{J}|$-dimensional. Letting $\mathcal{J}_i = \{j \in \mathcal{J} : a_{ij} \neq 0\}$, the number of dimensions of this state variable can be reduced to $|\mathcal{J}_i|$. 3) We can use state aggregation ideas to obtain very high-quality solutions to the optimality equation in (9). 4) We can use $\{\sum_{i \in \mathcal{L}} u_t^i(x_t) : t \in \mathcal{T}\}$ as approximations to the value functions. The profit obtained by using these value function approximations is 3-5% better than the profit obtained by the decision rule prescribed by the deterministic linear program in (1)-(5).

Optimal Dynamic Promotion and the Knapsack Problem for Perishable Items

JACKO, PETER (Universidad Carlos III de Madrid) peter.jacko@gmail.com

Senior managers in retail industry make important decisions upon assortment planning, product pricing, and product promotion. While product assortment is a strategic decision taken over a long-term planning period (Kök et al., 2006), the latter two are both strategic and tactical: they can be used in day-to-day marketing decisions to dynamically adjust to demand variations. Within the food retail industry, the necessity, frequency, and complexity of pricing and promotion decisions are further magnified by *perishability* of food products.

There is a strong need by retail managers for a "soft" marketing tool, which would dynamically allow them to improve sales and revenues, yet not altering product prices. For, dynamic pricing models may prescribe to change prices too often or in an "unsystematic" fashion, which contradicts discrete-time decision making, implementation costs, and retail brand image strategy. In addition, the price reduction must usually be done over all units of the product, thus losing possible profit from customers willing to pay the original, higher price. Furthermore, dynamic pricing results in a trade-off between markdowns and stockouts, since markdowns may damage producers, while stockouts may damage retailers.

We design a revenue management model, in which demand is stimulated by moving a number of product units to a promotion space, rather than by price changes. Thus, we address the problem of filling a promotion location with limited space to maximize the expected revenue, which we have termed the *Knapsack Problem for Perishable Items* (KPPI). Examples of the promotion space include shelves close to the cash register, promotion kiosks, or a depot used for selling via the Internet.

We solve the KPPI using a problem decomposition into single product unit subproblems. A natural mathematical setting for the KPPI subproblems is the restless bandit model (cf. Whittle, 1988; Niño-Mora, 2002), a fundamental stochastic model for resolving a trade-off between exploration and exploitation in an optimal fashion. In our model the bandits (perishable items) are *restless*, because they can get sold regardless of being in the knapsack or not, the time horizon is *finite*, and we are to select *more than one* item for the knapsack, which is allowed to be filled partially, due to the *heterogeneity* of the items.

Each product unit is assigned a *promotion priority index*, which captures the opportunity cost of promotion, as a function of its price, lifetime, expected demand, and expected promotion power. These indices are then used for each item as objective-function value coefficients in a (classic) knapsack problem, whose solution yields a well-performing heuristic for the KPPI. We thus mix up two models: the restless bandit problem and the knapsack problem.

1 Optimal Dynamic Promotion under Time-Homogeneous Demand

Suppose that the item perishes in T time periods, implying a final cost $c > 0$ if not sold before. Let $1 - p$ be the probability that a promoted item is sold in one period, and $1 - q$ that of a non-promoted item ($q > p$). Future costs are discounted by the one-period discount factor β. The next proposition asserts that, optimally, an item with lower probabilities of being sold when not promoted is assigned a higher promotion priority index, and that the index is nondecreasing over time. Hence, once the item is optimally chosen for promotion, then it should remain promoted until it perishes.

Proposition 1. *The perishable item is indexable with its index*

$$\nu^* = \frac{c(\beta q - \beta p)(\beta p)^{T-1}}{1 - (\beta q - \beta p)\frac{1-(\beta p)^{T-1}}{1-\beta p}}. \tag{1}$$

The results can be extended to the long-run average criterion by taking the limit $\beta \to 1$. Further, if $q = 1$, then the index reduces to $\nu^* = c(1 - p)$, the well-known $c\mu$-rule. The details and extensions, including the case with non-homogeneous demand and product inventories, are given in the full-length paper (Jacko, 2007), where we prove that, under certain demand regularity conditions, the necessity of marketing actions increases with larger inventory and with shorter product lifetime.

2 Knapsack Problem for Perishable Items

Since the dynamic programming formulation of the KPPI is most likely to be intractable, we relax the problem, which allows us to decompose it into single-item cases solved above.

Whittle (1988) proposed for restless bandits what has become known as the *Whittle's relaxation*: replace the infinite set of resource constraints by one constraint requiring to use the full resource only *in expectation*, resulting in a sort of *perfect market assumption*. This is solved by the Lagrangian method: for a given penalizing parameter, the relaxation can be decomposed and analyzed separately for each item, so we can apply the results obtained above. The penalizing parameter can be interpreted as the competitive market price of space, the resource provided by the knapsack.

Since the indices measure the opportunity cost (the true economic value) of promoting an item, we propose the following heuristic construction for the (non-relaxed) KPPI: "Promote the items that are given by an optimal solution to a knapsack problem with indices as the objective function value coefficients and item volumes as the knapsack constraint weights". Our simulation studies suggest that the proposed heuristic outperforms the existing ones of Whittle (1988) and Niño-Mora (2002).

3 Conclusions

We have developed a dynamic and stochastic model for promotion of perishable products, and proposed a solution that has a natural economic interpretation and suggests itself to be easily implementable in practice. These advantages come at the cost of possible suboptimality of such a dynamic solution, which is, however, negligible and smaller than the cost of implementing naive marketing solutions or other existing proposals. The model has an appealing property of being extensible to a variety of ad-hoc requirements that managers or certain circumstances may impose.

More generally, KPPI offers a comprehensive modeling framework that may be used in other applications, since the items considered in knapsack problems are often perishable, either naturally or due to special restrictions. Apart from product promotion in other retail industries, applications arise transplantation medicine or task management.

References

Jacko, P. (2007). Optimal dynamic promotion and the knapsack problem for perishable items. submitted.

Kök, A. G., Fisher, M., and Vaidyanathan, R. (2006). *Retail Supply Chain Management*, chapter Assortment Planning: Review of Literature and Industry Practice. Kluwer Publishers. (to appear).

Niño-Mora, J. (2002). Dynamic allocation indices for restless projects and queueing admission control: A polyhedral approach. *Mathematical Programming, Series A*, 93(3):361–413.

Whittle, P. (1988). Restless bandits: Activity allocation in a changing world. *A Celebration of Applied Probability, J. Gani (Ed.), Journal of Applied Probability*, 25A:287–298.

Near-Optimal Protection Level Policies for the Capacity Allocation Problem

LEE, THIAM (Columbia University) thl2102@columbia.edu
HUH, TIM (Columbia University) th2113@columbia.edu
RETSET, LEVI (Massachusetts Institute of Technology) retsef@mit.edu

In this paper, we develop a non-parametric data-driven approach to a classical revenue management problems — the capacity allocation problem under monotonic fare arrivals [1, 3]. We introduce a new class of solutions called ε-*backwards accurate policies*, and prove that these policies are near-optimal. That is, for each accuracy level α, if ε is sufficiently small, the expected revenue of each ε-backwards accurate policy is at least $(1-\alpha)$ times the optimal one. We then describe a simple algorithm that uses uncensored independent demand samples to compute an instance from this class of solutions with high probability. Furthermore, we give polynomial uniform bounds (finite sample bounds) on the number of demand samples needed to ensure the expected revenue of the computed solution is, with a given confidence probability, within some specified ratio of that of an optimal solution, where by uniform bounds we mean bounds that hold for *any* demand distribution.

In the capacity allocation problem, we sell airplane seats on a flight leg. We suppose that a flight has total capacity x_0 to be sold to $M+1$ different classes of demand indexed by $1,\ldots,M+1$, each willing to pay a different fare $f_1,\ldots,f_{M+1}$ and whose quantities are $D_1,\ldots,D_{M+1}$ respectively. We assume that the demand classes (also termed fare classes) are monotonic in their fares, $f_1 \geq f_2 \geq \ldots \geq f_{M+1}$, and that the demands for the fare classes $\mathbf{D} = (D_1,\ldots,D_{M+1})$ are random variables that are independent and have finite expectations and whose cumulative distribution functions are $F_1,\ldots,F_{M+1}$. Furthermore, we assume that demand for fare class $k+1$ arrives completely before demand for fare class k, for $k = 1,\ldots,M$. From Brumelle and McGill [1] and Robinson [3], an optimal policy for this problem is of the form of protection levels. A protection level p_k is the amount of capacity that is reserved exclusively for sales to customers of fare classes 1 through k inclusive, or equivalently that once the remaining amount of capacity is p_k or less we stop selling to customers of fare classes $k+1,\ldots,M$. The expected revenue-to-go at the beginning of fare class i sales using a policy $(p_1,\ldots,p_{i-1})$ can then be expressed recursively as $W_i(x) = \int_0^{x-p_{i-1}} r_i f_i + W_{i-1}(x-r_i)dF_i(r_i) + \bar{F}_i(x-p_{i-1})\left((x-p_{i-1})f_i + W_{i-1}(p_{i-1})\right)$ for $x \geq p_{i-1}$, and $W_i(x) = W_{i-1}(x)$ for $x < p_{i-1}$, with the boundary condition $W_0(\cdot) = 0$. The optimal protection level vector can be characterized as follows: if $\mathbf{p}^{*,M} = (p_1^*,\ldots,p_M^*)$ is a protection level vector satisfying $(f_{k+1}/f_1) = \Pr\left(D_1 > p_1^*,\ldots,\sum_{i=1}^k D_i > p_k^*\right)$ for all $k = 1,\ldots,M$, then $\mathbf{p}^{*,M}$ is an optimal protection level vector.

We define and analyze what we term ε-backwards accurate policies, whose construction we will outline further on in this abstract. An ε-backwards accurate policy $\hat{\mathbf{p}}^M = (\hat{p}_1,\ldots,\hat{p}_M)$ satisfies the optimality condition above to within some small margin of error, in other words it satisfies $(f_{k+1}/f_1) - \varepsilon \leq \Pr\left(D_1 > \hat{p}_1,\ldots,\sum_{i=1}^k D_i > \hat{p}_k\right) \leq (f_{k+1}/f_1) + \varepsilon$ for some $\varepsilon > 0$ for all k. We find that using an ε-backwards accurate policy leads to an unusual structure in the expected revenue-to-go function with respect to remaining capacity. In particular, we find that the function is almost concave in terms of its gradient, more explicitly if $\hat{W}_i(x)$ is the revenue-to-go with remaining capacity x and using an ε-backwards accurate policy at the stage where the i-th highest paying customers are arriving then for all $0 \leq a \leq b$, $(\partial\hat{W}_i/\partial x)(b) \leq (\partial\hat{W}_i/\partial x)(a) + 2\varepsilon f_1$. Exploiting this structure allows us to relate the performance of an ε-backwards accurate protection level policy with that of an optimal one, and we find that the relative gap of the two, letting $\hat{W}$ be the overall expected revenue of the ε-backwards accurate policy and W^* be that of an optimal one, can be upper bound as follows: $(W^* - \hat{W})/W^* \leq (2\varepsilon f_1^2 M)/(f_M(f_{M+1} - \varepsilon f_1))$. Note that this bound is independent of the demand distributions as well as the starting capacities.

We also give an algorithm that computes from sufficient demand samples some ε-backwards accurate policy with high probability. An outline of this algorithm is as follows: perform the proceeding iteratively for $k = 1$ through M in sequence. Given that we have computed the first through $(k-1)$-th protection levels $\hat{\mathbf{p}}^{k-1} = (\hat{p}_1,\ldots,\hat{p}_{k-1})$, we compute the k-th protection level, $\hat{p}_k$, in the following manner: use some number of uncensored independent demand samples to form an empirical distribution of the demands. Then, using protection levels $\hat{\mathbf{p}}^{k-1}$ for the first through $(k-1)$-th protection level, find the quantity for the k-th protection level such that the optimality conditions are satisfied under the empirical distribution and set $\hat{p}_k$ as that quantity. We show that our algorithm is able to output a policy that is $(1-\alpha)$-optimal (its expected revenue is at least $1-\alpha$ that of an optimal policy) with probability $1-\gamma$ when we use $N(\alpha,\gamma)$ demand samples, where

$$N(\alpha,\gamma) = \frac{M\left(2f_1^2 M + \alpha f_1 f_M\right)^2 \cdot (\ln 2M - \ln\gamma)}{2\left(\alpha f_M f_{M+1}\right)^2} = \mathcal{O}\left(\left(\frac{f_1}{f_{M+1}}\right)^4 \cdot \frac{M^3(\log M - \log\gamma)}{\alpha^2}\right).$$

Our paper addresses a gap in the literature with respect to addressing the capacity allocation problem under limited or no known demand information, but instead with access to demand samples. Before us, Van Ryzin and McGill [4] and Huh and Rusmevichientong [2] developed stochastic gradient ascent algorithms that update the protection levels from selling season to selling season based on realized sales, converging to an optimal solution at a shown asymptotic rate. In contrast to their papers, our work takes a different approach that allows us to say something much stronger than an asymptotic rate of convergence for our algorithm: we find uniform sampling bounds, in fact polynomial ones, explicitly linking desired performance to the number of samples required to guarantee that. To do this we had to assume more — the availability of uncensored demand samples. However, we argue that this assumption is fair as several modern sales channels for flight reservation systems possess the ability to collect on customers' requests, beyond transacted sales data, which gives information on the underlying demand distribution. For example, travel booking websites such as Travelocity or Orbitz that sell flight tickets on the Internet can easily collate data on their customers' search and browsing history that allows for uncensoring of sales data.

Additionally, we note that while in this abstract we assumed the fare classes arrive monotonically, in the full paper we generalize this result to fare classes that could arrive non-monotonically.

A Polyhedral Study of the Network Pricing Problem with Connected Toll Arcs

HEILPORN, GÉRALDINE (Université de Montréal, Université Libre de Bruxelles) gheilpor@ulb.ac.be
LABBÉ, MARTINE (Université Libre de Bruxelles) mlabbe@ulb.ac.be
MARCOTTE, PATRICE (Université de Montréal) marcotte@iro.umontreal.ca
SAVARD, GILLES (École Polytechnique de Montréal) gilles.savard@polymtl.ca

Consider the tarification problem of maximizing the revenue generated by tolls set on a subset of arcs of a transportation network, where origin-destination flows (commodities) are assigned to shortest paths with respect to the sum of tolls and initial costs.

This talk is concerned with a particular case of the above problem, in which all toll arcs are connected and constitute a path like in a real highway topology. In order to allow more flexibility, a complete toll subgraph is considered, where each toll arc corresponds to a toll subpath. Two variants of this problem are studied, with or without specific constraints linking together tolls on the arcs.

First described by a bilevel formulation, the problem can be modelled as a linear mixed integer program. It is proved to be *NP*-complete by a reduction to 3-*SAT*. Several families of valid inequalities are proposed for this problem, which exploit the underlying network structure and strengthen some important constraints of the initial model.

Their efficiency is first shown theoretically. Focusing on instances involving one or two commodities, we prove that most of the valid inequalities, as well as several initial constraints, are facet defining for the convex hull of feasible solutions for those restricted problems. A complete description of the convex hull of feasible solutions for a single commodity problem is also given.

Numerical tests have also been conducted, and highlight the real efficiency of the valid inequalities for the multi-commodity case. Several of the valid inequalities proposed seem very efficient, at least to decrease the gap or number of nodes in the Branch and Cut algorithm. Some of those also allow to decrease the computing time for one variants of the problem.

Finally, we point out the links between the problem studied in this talk and a more classical design and pricing problem in economics.

References

[1] G. Heilporn, M. Labbé, P. Marcotte, and G. Savard. A polyhedral study of the network pricing problem with connected toll arcs. Technical Report 579, ULB, 2007.

Session II.2: INDUSTRY I (8:30-10:00)

Room: *S1-131* **Chairperson:** *Gilles Savard*

Improving Pricing Performance

OOSTEN, MAARTEN (PROS) moosten@prospricing.com

In this presentation we will approach pricing optimization as the process of improving pricing performance. The objective of the optimization process most often includes multiple success criteria. While profits and revenues are typically an important part of pricing, objectives can also include setting prices that achieve certain levels of market share, or prices that don't impinge upon customer satisfaction.

The process of improving pricing performance encompasses many different scientific disciplines including data mining, statistics, econometric modeling, and operations research. However the hallmark of the optimization process is not the level of sophistication of the methods that are applied but the fact that the process relies on a feedback loop. Decisions are made and the effects are measured. Based on these measurements, a decision process is then updated.

There are many different types of pricing including list pricing, deal pricing, contract pricing, promotional pricing, and others. Yet, despite the differences, these various price optimizations appear to share a common process structure. The ultimate goal of this process is to allow for quick, appropriate, and manageable pricing responses to changing market conditions.

A New Revenue Optimization Tool for High-Speed Railway Finding the Right Equilibrium between Revenue Growth and Commercial Objectives

RISS, MARIANE (SNCF, Paris) mariane.riss@sncf.fr
CÔTÉ, JEAN-PHILIPPE (ExPretio Technologies)
ROBBE, NICOLAS (Thalys, Bruxelles)
SAVARD, GILLES (École Polytechnique de Montréal) gilles.savard@polymtl.ca

In order to better tackle airline competition and to prepare for the upcoming deregulation of European rail transportation over the coming years, SNCF (the French national railway operator) is working in partnership with ExPretio Technologies to develop a new-generation Revenue Optimization tool. The tool, built on a "bilevel" mathematical optimization model, allows to take explicitly into account customer purchase behaviour and the offers of competitors. This new approach, initially developed at the Center for Research on Transportation (CRT) of the University of Montréal and École Polytechnique Montréal to fulfill the needs of airlines, is now being adapted to SNCF's and its affiliates' business context (fare structure, nature of the customer base, inventory management policies and practices, etc.).

This new tool is now integrated in SNCF's information technology and data-processing environment and consists of a combinatorial optimization model, a demand forecasting model, and a customer behaviour model. This innovative approach places customer purchase behaviour at the centre of the modelling process. Indeed, contrary to traditional RM approaches, we do not suppose that demand forecasts are independent and that customers are captive of one given type of product. In our model, the customer has access to complete information about the transportation services being offered on the market and makes purchase decision according to his own preferences and choice criteria.

In this talk, we will present this new Revenue Optimization approach and contrast it with more classical methods currently applied in this field. We also underline the benefits brought by our model and discuss some results obtained through large-scale and live benchmarks on the network of THALYS. We conclude by listing some areas where we can improve the model as we go through the industrialization and deployment of the tool.

Session II.3: DYNAMIC PRICING I (8:30-10:00)

Room: *S1-139* **Chairperson:** *Mikhail Nediak*

Dynamic Pricing in a Market with Strategic Customers

BANSAL, MATULYA (Columbia Business School) mbansal11@gsb.columbia.edu
MAGLARAS, COSTIS (Columbia Business School) c.maglaras@gsb.columbia.edu

We consider a revenue-maximizing monopolist firm, the seller, that sells a homogeneous good to a market of heterogeneous strategic customers that differ in their valuations and risk preferences. The firm seeks to discriminate customers by selling the product over a period of time at different prices while controlling its availability, i.e. by creating rationing risk. Customers observe or anticipate the price and associated product availability at different points in time and decide when, if at all, to attempt to purchase the product in a way that maximizes their expected net utility from their purchase that accounts for the rationing risk associated with each time period. The seller's problem is to choose its dynamic pricing and product availability strategies to maximize its profitability taking into account the strategic customer choice behavior. Though some of these aspects have been studied elsewhere, a common feature of the past analysis is the stylized nature of the problem with respect to customer heterogeneity, e.g. by assuming two discrete types of customers, uniformly distributed valuations, no price control, and/or focus on two product prices. Our goal is to suggest an approach to this problem that is intuitive and allows one to relax these restrictive market assumptions.

Using a discrete (but arbitrary) customer valuation distribution, we first show that the above dynamic pricing problem can be reformulated as a mechanism design problem. While rationing, a key quality attribute in our setting, affects capacity consumption and differentiates our problem from the standard product design problem, the "standard machinery" applied in mechanism design is still useful in our setting.

Second, we show that prices are monotonic in fill-rates, i.e. periods with higher prices are also characterized by higher availability rates. It suffices to offer as many products as there are discrete valuations (types). If the number of offered products is strictly less than the number of customer types, then the offered products partition the customer types in accordance with their strategic choice behavior into contiguous classes (sets of customer types), so that the quality (and therefore price) of the product purchased by customers varies monotonically with their valuations. These properties allow us to reformulate the problem as a non-linear optimization problem in terms of the price decisions above. While the resulting problem is not easy to solve in general, it does provide a formulation that is amenable to a brute force computational solution. In many special cases, the problem simplifies considerably. Given a market with multiple types and a general valuation distribution, we analyze the two product problem, and the problem with risk-neutral customers. The first problem is solved in closed-form, while the second problem is reformulated as a LP, and hence can be easily solved. The revenue resulting from this LP is shown to be a lower bound of the revenue achievable in the problem with risk-averse customers. Its solution consists of offering at most two products, irrespective of the size and composition of the market, with the two product solution being offered only if the capacity constraint is tight.

Third, we show that when risk-aversion is low - which is practically relevant in many application settings - the general problem can be addressed hierarchically through two LPs.

1. The first LP corresponds to the risk-neutral problem that uses the various market primitives but treats all consumer types as being risk-neutral.
2. A second LP that solves for price and rationing risk perturbations around the risk-neutral solution and takes into account the risk-aversion of the various customer types.

The above decomposition is justified asymptotically as the risk-aversion coefficients of all market participants' approach 1 (i.e. the risk-neutral case). The perturbation in item 2 above is such that the two-product nature of the risk-neutral solution is preserved even for the problem with low risk aversion. This lends credibility to a practical heuristic that would focus on optimizing over a two product offering, which as mentioned earlier can be solved very efficiently even without this asymptotic decomposition.

Our approach unifies and extends several previously established results under a common and intuitive framework. The above problem can also be extended to include, for example, the possibility for some of the customers to be myopic (and to always purchase in a given time period), customer utility to be time-dependent, and customers to also differ in terms of their best outside opportunity which affects their no-purchase utility threshold, to name but a few.

On Pricing in a Non-Stationary Environment: Complexity Characterization and Optimal Policies

BESBES, OMAR (The Wharton School) obesbes@wharton.upenn.edu
ZEEVI, ASSAF (Columbia University) assaf@gsb.columbia.edu

We consider a dynamic pricing problem in an environment where the customers' willingness to-pay (WTP) distribution may change at some point over the selling horizon. Customers arrive sequentially and make purchase decisions based on a quoted price and their private reservation price. The seller knows the WTP distribution pre- and post- change, but the time of this change is unbeknown to him. The performance of a pricing policy is measured in terms of regret: the loss in revenues relative to a clairvoyant oracle which knows the time of change prior to the start of the selling season. We derive universal lower bounds on the regret and develop pricing strategies that achieve the order of these bounds; in that sense, the proposed prescriptions cannot be improved upon. We show that price experimentation is in general necessary for optimal performance, hence highlighting the role of dynamic pricing in changing environments. We then show that when market response can be monitored without sacrificing revenues, the pricing problem becomes simpler: passive monitoring schemes suffice to achieve optimal performance; and the regret achieved is of significantly smaller order than in the general case. Our formulation allows for almost arbitrary consumer WTP distributions, and purchase request patterns.

Competitive Dynamic Pricing with Guarantees in the Presence of Strategic Consumers

AVIV, YOSSI (Washington University) aviv@olin.wustl.edu
LEVIN, YURI (Queen's University) ylevin@business.queensu.ca
NEDIAK, MIKHAIL (Queen's University) mnediak@business.queensu.ca

We consider a two-period model of duopolistic capacitated dynamic pricing competition when each firm's strategy can include price matching, and consumer response to current market conditions is strategic. The products offered by the firms are differentiated, and consumer choice is described using a random utility model that takes expected product availability into account. We study the properties of equilibrium strategies, and whether price matching can be an effective tool in countering strategic consumer behavior.

Room: *Agora*

10:00 *Coffee Break*

TRACK III

Session III.1: GAME THEORY II (10:30-12:00)

Room: *S1-125* **Chairperson:** *Robert Shumsky*

The Anarchy of Price Competition

FARAHAT, AMR (Cornell University) aaf33@cornell.edu
PERAKIS, GEORGIA (Massachusetts Institute of Technology) georgiap@mit.edu

We study price competition in oligopolies of differentiated products where each firm offers one or more products. Most literature on price competition makes at least one of the following assumptions: A1) the products are perfect substitutes, A2) the industry consists of two firms only, A3) the industry consists of multiple firms but they are all symmetric, A4) each firm offers one product only. In this paper, we relax *all* four assumptions. The products are gross, but not perfect, substitutes; the number of firms is arbitrary; own and cross price effects are firm dependent; marginal production costs are constant but vary by product; and the demand model is an extension of the popular affine function that avoids various negativity and existence problems (see 1 below). We investigate a number of questions regarding the effects of competition, compared to centralized pricing, on industry profits and total welfare. Our main contributions are:

1. We formulate a demand system where a unit price increase of a product leads to a fixed nonnegative change in the demand of other products as long as this product has positive demand. This is a more realistic extension of the popular affine demand system that avoids i) negative demands, and ii) nonexistence problems associated with alternative extensions of affine functions to regions with negative demand that would occur in the absence of A4. We motivate the formulation using a concave quadratic utility function of a representative consumer and show that the demand system reduces to solving a linear program.
2. We show that a unique Nash equilibrium exists, and we characterize it in closed form. The existence and uniqueness results are significant because in our model, the strategy space is not compact and the payoff functions are not differentiable, not quasi-concave, and not available in closed form.
3. We provide sharp lower and upper bounds on the ratio of decentralized to centralized profits that are *independent of marginal costs*. The lower bound is based on the minimum eigenvalue of the demand sensitivity matrix and the upper bound is based on comparing diagonal elements of two derived matrices. We then develop simpler bounds, with more intuitive economic interpretation, based on a quantification of the degree of competition intensity.
4. We investigate the effect of competition on total surplus and on the tradeoff between consumers' surplus and industry profits (producers' surplus). We find that: i) the loss of efficiency due to competition is no more than 25%, ii) price equilibria are, in general, not Pareto efficient, making significant improvements in industry profits possible without sacrificing consumers' surplus, and iii) price fixing can improve total surplus, but by no more than 10%.

Dynamic Oligopoly: Cournot Equilibrium in Exhaustible Resource Markets

PANG, ZHAN (James) (University of Calgary) Silfamars@Gmail.com

This paper studies an oligopoly problem in an exhaustible natural resource market with a continuous-time Cournot-Nash framework. With limited resource reserve capacity, the oligopolist must optimally allocate the resource over time. The aggregate supply of all the players is cleared at a market price according to an inverse demand curve which is impacted by a random shock. By exploiting the inter-temporal nature of the supply process of the players, we can characterize the dynamics of subgame perfect Nash equilibrium market price explicitly. With iso-elastic demand model, we show that (1) equilibrium market clearing price follows a Geometric Brownian Motion and the discounted market clearing price is a martingale, which verifies the *Hotelling rule* (inter-temporal arbitrage conditions) that the price of exhaustible natural resource must grow over time exponentially with the rate of interest along the efficient resource exploitation in a competitive equilibrium; (2) with GBM random shock, each player's supply rate and aggregate supply rate decline overtime; (3) Cournot competition produces competitive equilibrium of the aggregate supply rate and spot prices; (4) Cournot competition leads to over-explore behavior.

Revenue Management with Price Bargaining

BHANDARI, ATUL (University of Pittsburgh) atulb@pitt.edu
SECOMANDI, NICOLA (Carnegie Mellon University) ns7@andrew.cmu.edu

It is well known that in business, as opposed to consumer, markets commercial transactions are typically carried out through buyer-seller bilateral negotiations rather than arm's length mechanisms. But the price based revenue management literature has only focused on the posted price and auction cases (Talluri and van Ryzin 2004 [3, Chapters 5-6]) that do not fully capture the bilateral nature of business negotiations. We focus on modeling bilateral price bargaining in a revenue management setting.

We consider a risk neutral seller who owns a finite inventory of a single product that can be sold during an infinite horizon divided into discrete time periods. At the beginning of each time period there is a positive probability of at most one buyer's purchase request for a single unit of the seller's inventory, that is, demand is Bernoulli. Buyers are risk neutral and do not strategize on when they place their purchase requests. If a sale is made, the seller's inventory decreases by one unit. The seller's and each requesting buyer's valuations for the marginal unit that is up for sale are their private information, which we model using independent probability distributions that are common knowledge to the seller and all potential buyers.

Differently from the existing literature, we consider the case when the seller and each arriving buyer negotiate the sale price. We take a game-theoretic mechanism design approach to model the *outcome* of the Bayesian price bargaining problem (Myerson 1991 [2, Chapters 6 and 10]) that ensues between the seller and a buyer who places a purchase request. That is, we do not model the *process* of price bargaining, we only focus on the outcome of this process. Thus, the seller faces a sequence of static Bayesian price bargaining problems *linked* by his *inventory* availability that changes dynamically as sales are made over time.

By the well known revelation principle, we restrict attention to direct and feasible mechanisms, which are roughly mediated plans that satisfy incentive compatibility and individual rationality constraints, that is, telling the truth is a Bayesian Nash equilibrium for both the seller and each arriving buyer and these players cannot be forced to trade unless it is advantageous to them, respectively. We show how to embed a direct and feasible mechanism for static bilateral price bargaining within a dynamic Markov decision process (MDP) that computes the seller's total discounted expected revenue for a given initial inventory availability during an infinite horizon. Under an informational assumption and a regularity condition, we exploit the structure of the problem and the incentive compatibility constraints to simplify the MDP formulation by showing that optimally solving the resulting MDP reduces to solving a finite sequence of univariate fixed point problems, one for each inventory level.

We investigate the following structural issue. Suppose that one mechanism is known to outperform another mechanism in terms of interim expected utility for the seller in a static setting. That is, the seller's expected utility given that he knows his type (marginal valuation) is no smaller under the former mechanism than the latter one. Does this imply that his optimal value functions when these mechanisms are employed to resolve each price negotiation in our dynamic setting satisfy the same relationship? This question is nontrivial, because the seller's marginal valuation for the same inventory availability, in any time period, is not necessarily the same under the two mechanisms. We show that the answer to this question is positive in the context of our model.

We apply this result to compare four price bargaining mechanisms analyzed by Myerson (1985 [1]) in the context of the (static) symmetric uniform trading problem (SUTP), where it is common knowledge to both the seller and the buyer that their beliefs about each other's marginal valuation are independently and uniformly distributed on the unit interval. Three of these mechanisms differ in terms of whether the seller has all the bargaining ability, this ability resides entirely with the buyer, or the seller and the buyer equally share this ability. (Myerson (1985 [1]) defines this ability as the ability "to argue articulately and persuasively in the negotiation process.") The fourth mechanism corresponds to a structure negotiation policy. We show that in our dynamic setting the seller is better off by having more, rather than less, bargaining ability for all inventory levels. The seller is also better off in the structured negotiation environment compared to the case when all the negotiation ability resides with each buyer.

We numerically measure the differences between the seller's optimal value functions in each of these cases, that is, we quantify whether the highlighted performance differences are significant or negligible, and explaining their nature. For the range of parameters considered, there is not a significant advantage for the seller to having all the bargaining ability than sharing it equally with each buyer, but in both these cases he is significantly better off than when this ability resides entirely with each buyer. Compared to the latter case, the seller performs significantly better in the structured negotiation environment, but not as well as in the two former cases for reasonable relative inventory availability (days of demand assuming that a sale always occurs upon a buyer's arrival). However, we observe that the seller can be better off in the structured negotiation case for sufficiently high relative inventory availability. We leverage our structural results to explain these findings. We also investigate robustness issues related to model misspecification.

References

[1] R. B. Myerson. Analysis of two bargaining problems with incomplete information. In A. E. Roth, editor, *Game-theoretic Models of Bargaining*, pages 115-147. Cambridge University Press, Cambridge, UK, 1985.

[2] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, MA, USA, 1991.

[3] K. T. Talluri and G. J. van Ryzin. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers, Boston, MA, USA, 2004.

Managing Revenues in Alliances Under Incomplete Information

WRIGHT, CHRISTOPHER (University of Rochester) wrightc1@simon.rochester.edu
GROENEVELT, HARRY (University of Rochester) groenevelt@simon.rochester.edu
SHUMSKY, ROBERT (Tuck School at Dartmouth) robert.shumsky@dartmouth.edu

The revenue management problem for airlines utilizing interline itineraries, such as those commonly used within alliances, is considerably more complex than that faced by individual airlines managing isolated networks. While alliances seek to coordinate the partners' actions, the partners within the alliance may be unwilling or unable to share complete information about their operations because of competitive, technical and legal barriers. For example, they may not provide complete information about demand forecasts to each other because this would require sharing information on all flights and itineraries that they operate – including those itineraries on which they compete.

Wright et al. (2008) formulate a model for a two-partner alliance in which one partner can sell an interline itinerary only after paying a posted transfer price to its partner for the relevant subitinerary. They derive the equilibrium decision rules for static proration scheme that are prevalent in practice, as well as for several dynamic schemes based on suggestions from practitioners. Their formulation assumes a game of *complete information*: within each time period, each airline knows its partner's inventory level, and both have identical forecasts of future arrival probabilities and revenue distributions over the entire alliance network. The authors leave open the research question of how such policies could be managed when the partners are faced with incomplete information at the time they make decisions. We begin to address that problem here.

For dynamic transfer price (proration) schemes to function, each partner must post a list of current transfer prices for each interline itinerary. In this paper we assume an extreme case: that these transfer prices are the *only* information that each airline has about the state of its partner's network. Specifically, its partner's forecasts (revenue and arrival processes) and inventory levels are not shared. Therefore, each partner's decisions are based solely on the transfer prices, their own forecasts for their own networks, and their own inventory levels. Under this new game with *incomplete information,* we examine the decision rules of each partner. We explore the airlines' equilibrium policies under this game and compare the policies, as well as the overall performance of the alliance, to both the centralized control (first-best) alliance and an alliance that shares complete information.

Reference

Wright, C., H. Groenevelt and R. Shumsky. 2008. "Dynamic Revenue Management in Airline Alliances." Working paper, Available at SSRN: http://ssrn.com/abstract=1105135.

Session III.2: INDUSTRY II (10:30-12:00)

Room: *S1-131* **Chairperson:** *Darius Walczak*

Introduction to Data Normalization

QUILLINAN, JOHN (Walt Disney World) John.Quillinan@disney.com

We will introduce the concept of data normalization in pricing. RASPOS is a priceable model where the primary goal is to determine optimal prices that maximize revenue and profitability. The fundamental first step in identifying optimal pricing is to identify the price demand relationship. The demand of a product is influenced by many explanatory variables such as attendance, demographics of guests, rain, temperature, product availability. In order to extract the price-demand relationship, we must isolate the effect of price on demand from the effects of other explanatory variables. Normalization process is employed to isolate the impact of the significant explanatory variables on demand.

Normalization occurs on a weekly level, and within demand zones. All the data is at a weekly level -- Daily demand of a product aggregated weekly and daily park attendance aggregated weekly. Demand zones are created by grouping merchandise locations for the purpose of using only applicable explanatory variables. Explanatory variables are identified through external analysis. We use Lookups in PPSS to input the explanatory variables into PPSS, and assign explanatory variables to specific products at the demand zone level. We will conclude by discussing what explanatory variables might apply to other businesses.

Choice Models and the Marginal Revenue Data Transformation

WALCZAK, DARIUS (PROS Revenue Management) dwalczak@prosrm.com
MARDAN, SETAREH (PROS Revenue Management) smardan@prosrm.com
KALLESEN, ROYCE (PROS Revenue Management) rkallesen@prosrm.com

We start by briefly reviewing popular choice models appearing in optimization problems, both in airline RM and in Pricing. We then show how to reduce a single-resource optimization to the traditional yield optimization by means of the marginal revenue data transformation. We provide an interpretation of the EMSR heuristic for the case with demand that buys down, and show that this interpretation is consistent with the data transformation.

We indicate how to extend the transformation to O&D network for some classes of the problem and report on experiences from airline industry.

We provide a high-level overview of how the technique facilitates implementing in real-life inventory control system new models of demand that go beyond the traditional yieldable assumption.

Some simulation results and comparisons to other relevant control methods are provided as.

Session III.3: DYNAMIC PRICING II (10:30-12:00)

Room: *S1-139* **Chairperson:** *Georgia Perakis*

Integrated Inventory Allocation and Markdown Pricing at Multiple Stores

CHEN, MING (University of Maryland, College Park) chenming@rhsmith.umd.edu
CHEN, ZHI-LONG (University of Maryland, College Park) zchen@rhsmith.umd.edu
PUNDOOR, GURUPRASAD (ILOG) gpundoor@ilog.com

Over the last two decades, significant research efforts have been made in the area of dynamic pricing. In a typical dynamic pricing problem, one attempts to determine the optimal pricing scheme for selling a given amount of inventory over a finite horizon. The basic problem originates from retail industry where firms are either selling seasonal products, fashion goods or products with limited life time such as consumer electronics. Most of the existing models in this area consider a single store and a single product and very few papers have examined the situation involving multiple products or multiple stores. Bitran, Caldentey, and Mondschein (1998) is the only paper we are aware of that investigates a multiple-store dynamic pricing problem. In practice, however, many retailers (e.g., Best Buy, Circuit City, Jcrew) own and operate many stores and the pricing and inventory allocation decisions must be coordinated among these stores, and therefore, solving the pricing problem for each individual store independently may not be adequate and can result in inferior pricing decisions in this case. In addition, most of the existing papers focus on stylized models where some special assumptions are made with respect to demand distributions and allowable prices. The managerial insights or structural properties obtained by such models will not work for more complex practical problems with complex business rules that have to be followed.

We study a real-world problem faced by a large retailer that owns multiple (e.g., 50) stores in the context of clearance sales. These stores are served by a central warehouse. At the beginning of the clearance sales, the total inventory in the central warehouse is given and the retailer has to determine jointly how much inventory to allocate to each store and what markdown prices to implement over time at each store in order to achieve the maximum possible revenue. There are a number of business rules that the retailer has to follow when making the inventory allocation and determining the pricing decisions. For example, each store has to be allocated at least a certain amount of inventory (minimum inventory allocation), prices can only be non-increasing (markdown) over time, the number of markdowns at each store can not exceed a given limit, retail stores in the same geographical region have to employ similar pricing decisions (price difference can only be within a certain range), and the magnitude of markdowns has to be within a given range in each period. These complex business constraints make all available insights in the pricing literature inapplicable. We show that the problem with even a single store, a subset of these business constraints, and deterministic demand functions is NP-complete. Furthermore, the retailer's knowledge about the probability distribution of the demand at a given price may improve over time once new information becomes available. Hence, the retailer needs to employ a rolling horizon approach where the problem is resolved at the beginning of each period by utilizing new information about demand distributions.

We model stochastic demand using discrete demand scenarios based on the retailer's latest knowledge about the demand distribution. The advantage of using scenarios is that it allows us to model demand correlation across different time periods. Most existing markdown pricing literature assumes that demands at different time periods are independent. However, in practice, due to the word of mouth effect or due to the finite population, the demand in one period is usually dependent on the sales in previous periods. We formulate our problem at the beginning of each period as a MIP (mixed integer programming) formulation with demand scenarios. Due to the large scale of the formulation, it is impractical to solve it directly by an MIP solver. Therefore, we propose a Lagrangian relaxation approach and decompose the problem by store clusters. We find that the most commonly-used approaches for finding the optimal Lagrangian multipliers are very sensitive to the initial values as well the step size chosen and very difficult to implement. However, by looking at the economic interpretation of the Lagrangian multipliers for our problem, we show that the possible values of the Lagrangian multipliers must fall within a certain interval. We thus employ a search procedure to find good multipliers to use. This search procedure works very well and the procedure converges only after a few iterations for most cases. We demonstrate through numerical experiments that the Lagrangian relaxation approach achieves high-quality solutions (98% of the optimal for most cases) within a reasonably amount of time (within 2 hours for most cases) for problems of practical size.

We implement our approach on a rolling horizon basis in which the problem is re-solved in each period after demand realization from the previous period has been observed and more accurate knowledge about demand distribution has been obtained. We show how the demand learning techniques (e.g., Bayesian update) discussed by Bitran and Wadhwa (1996) can be incorporated in our model. We compare our approach with the deterministic solutions as well as several simple heuristic solutions that are commonly used in practice and demonstrate by numerical examples how our approach outperforms others under various circumstances. We conclude that our approach can be easily applied in practice for solving large scale markdown pricing optimization problems such as in the case of large retail chain with many stores.

References

Bitran G., Caldentey R., and Mondschein S. 1998. “Coordinating Clearance Markdown Sales of Seasonal Products in Retail Chains”. Operations Research, 46(5), 609-624

Bitran G., and Wadhwa H. 1996. “A Methodology for Demand Learning with an Application to the Optimal Pricing of Seasonal Products”. Working paper, MIT Sloan School of Management.

A Model of Dynamic Pricing with Strategic Customers and an Information Provider

YANG, SONG (University of Chicago Graduate School of Business) syang1@chicagogsb.edu
BIRGE, JOHN (University of Chicago Graduate School of Business) john.birge@chicagogsb.edu

Recently, dynamic pricing problems with strategic customer behavior have drawn increasing attention from academia. Within this work, many papers focus on equilibrium results in the dynamic game between the seller and customers. Due to the complexity of the dynamic game, common assumptions are that the seller and customers have symmetric information; and customers are capable of solving highly complicated quantitative models.

However, in reality, the seller usually owns some private information, such as real-time inventory level, demand forecasts, etc., which is unavailable to customers. Overestimating customers' knowledge may lead to an overly conservative pricing policy. Moreover, while customers may have varying computational capabilities, the surge of information technology provides a simple solution. Recently, many websites and software have been developed to help customers compare prices of the same product from different sources, show historical price patterns, and even provide price forecasts. Designed properly, these information providers allow customers to increase their consumer surplus, free of the burden of heavy computation.

To shed light on the above problems, we impose a different information structure in this study. We assume customers have no information other than the current price and a signal on the future price evolution from an information provider, who possesses a subset of the seller's information (historical prices, certain demand parameters, etc.). Customers then employ certain simple rules to decide whether to make an immediate purchase or wait for a possible markdown in the future. If they decide to wait, they will revisit the store in the future, which is governed by a certain stochastic process.

Other than this assumption, our model carries most common assumptions in the literature, including finite stock and sale horizon, customer heterogeneity, and stochastic demand. As the number of waiting customers is unobservable to the seller, the seller solves a finite-horizon partially observed Markov decision process. Upper bounds and structural results are derived. While the high dimension of the state space prevents us from solving the problem exactly, different heuristic and approximation algorithms are employed. Numerical results are presented.

Dynamic Pricing Using Scenario Based Optimization

RUBEN LOBEL (Massachusetts Institute of Technology) rlobel@mit.edu
GEORGIA PERAKIS (Massachusetts Institute of Technology) georgiap@mit.edu

In many industries, managers are faced with the challenge of selling a given amount of a certain product within a fixed time horizon. This could be the case of an airline selling tickets for a flight, a hotel trying to book rooms, or a retail store selling articles within the current fashion season. In any of these cases, we have the two common traits: a fixed initial inventory that cannot be replenished during the selling horizon; a small or zero salvage value for the product after the end of the season. In the setting of this paper, we assume there is only one seller, selling a single product, with salvage value of zero at the end of the selling horizon. We also assume that customers are myopic and demand is observed by the firm as an external aggregate demand process.

There is an extensive literature on this problem using a stochastic optimization approach, where we assume the demand uncertainty is given by a certain probability distribution and we would like to maximize our expected revenue. On the other hand, we consider an optimization approach to the dynamic pricing problem, where our goal is to determine the pricing strategy that maximizes revenue robustly for all possible scenarios of demand, without making distributional assumptions.

While previous robust optimization models can be tractable, most of them have a static structure. The main contribution of this paper is the exploration a closed-loop robust approach to pricing, *i.e.*, instead of finding a single static pricing decision, we are looking for pricing policies that are functions of the previous realized demands. The drawback of introducing adaptability to the model is that the problem becomes non-convex, loosing tractability. To tackle this issue, we introduce a scenario based optimization approach that can solve this problem, with a confidence level based on the number of samples used. We will show how this methodology can be used for data-driven pricing or adapted for a random sampling optimization approach.

Furthermore, we will demonstrate the methodology using numerical simulations. We first consider the open-loop robust problem, where the original problem is tractable, comparing the exact solution with the scenario based approach. We also compare the simulated performance of the closed-loop solution with the static robust, showing the gains of introducing adaptability. We also look at the performance of the adaptable robust solution compared with a stochastic optimization approach and observe how their performance can be affected when we are mistaken about the true demand distribution. Finally, we will propose an iterative algorithm, which can be used to refine the scenario pool and has shown to improve the performance of the model in the numerical tests.

Room: *Agora*

12:00 *Lunch*

PLENARY SESSION (Industry)

Room: *S1-151* **Chairperson:** *Mariane Riss*

13:30 *A French Touch In High Speed Rail Revenue Management*

MICHEL LEBOEUF (SNCF) michel.leboeuf@sncf.fr

Room: *Agora*

14:30 *Coffee Break*

TRACK IV

Session IV.1: CONGESTED SYSTEMS (15:00-16:30)

Room: *S1-125* **Chairperson:** *Nicolas Stier-Moses*

Pricing Services Subject to Congestion: Sell Subscriptions or Charge a Per-Use Fee?

CACHON, GERARD (The Wharton School) cachon@wharton.upenn.edu
FELDMAN, PNINA (The Wharton School) pninaf@wharton.upenn.edu

How should a firm price its service when congestion is an unavoidable reality? Customers dislike congestion, so a firm has an incentive to ensure it provides reasonably fast service. At the same time, the firm needs to earn an economic profit, so the firm's pricing scheme must generate a sufficient amount of revenue. Furthermore, these issues are closely linked: the chosen pricing scheme influences how frequently customers use a service, which dictates the level of congestion; congestion correlates with the customers' perceived value for the service, and that determines the amount of revenue the firm can generate.

A natural option is to charge customers a per-use fee or toll. Naor (1969) began this line of research and there have been many subsequent extensions of his basic model, but nearly always with a focus on per-use fees. (See Hassin and Haviv 2003 for a broad survey of this literature.)

Although the emphasis in the queuing literature has been placed on per-use pricing, other pricing schemes are observed in practice. Most notably, some firms sell subscriptions for the use of their service: a health club may charge an annual membership that allows a customer to use the facility without additional charge for each visit; American Online, an Internet service provider, initially charged customers per-use access fees but later switched to subscription pricing (a monthly access fee with no usage limitation); Netflix, a retailer that provides movie DVDs for rental, also uses subscription pricing (a monthly fee for an unlimited number of rentals); Disney charges an entry fee for its theme park without charging per ride on the attractions; etcetera.

Despite the existence of subscriptions in practice, a subscription pricing strategy has a clear limitation in the presence of congestion effects: subscribers are not charged per use, so it is intuitive that they seek service too frequently (e.g., use the health club too often), thereby increasing congestion and decreasing the value all subscribers receive from the service. As a result, in a setting with clear congestion costs (e.g., in a queuing model) one might assume that subscription pricing would be inferior to per-use pricing. However, in this paper we demonstrate that subscription pricing may indeed be a firm's better pricing strategy despite its limitations with respect to congestion. We do so in three different capacity management scenarios: (i) the firm's service capacity is exogenously fixed; (ii) the firm's service capacity adjusts to meet an industry standard for congestion; and (iii) the firm endogenously chooses its service capacity in addition to its pricing policy.

Using a queuing framework, we find that a firm may prefer subscription pricing over per-use pricing even if consumers dislike congestion. Furthermore, subscription pricing may be preferable in situations that would a priori suggest a preference for per-use pricing: when the industry has a high standard for congestion (i.e., a standard for having customers spend little time in the service process); or when customers' waiting costs are high (so that congestion is costly) or when capacity is expensive (so that it is costly to build sufficient capacity to reduce congestion). Subscription pricing can dominate in these situations because it is capable of generating higher revenue than per-use pricing. This is particularly important when the firm must incur a large capacity cost to operate with little congestion. Next, we find that the absolute advantage of subscription pricing can be considerable whereas the absolute advantage of per-use pricing is generally modest - per-use pricing generates higher revenue or earns higher profit only when revenue or profit is reasonably low. Overall, we conclude that the emphasis on per-use pricing in the queuing literature is somewhat misplaced - we provide evidence that subscription pricing can indeed be the preferable pricing strategy even in services that experience congestion.

References

Hassin, R., M. Haviv. 2003. To queue or not to queue: equilibrium behavior in queueing systems. Kluwer Academic Publishers, Boston, Massachusetts.

Naor, P. 1969. The regulation of queue size by levying tolls. Econometrica 37(1) 15-24.

Enabling State Dependent Priority Service by Using Pricing Mechanisms That Encourage Customers to Jump the Queue

GILONI, AVI (Yeshiva University) agiloni@yu.edu
TROY, PHIL (Les Entreprises TROYWARE) PhilTroy@Mail.com

Given the importance of maximizing the benefits less waiting costs that organizational service facilities generate serving internal customers, there has been a considerable amount of research on using pricing mechanisms to optimally select the customers processed by these facilities. This research includes analyses of the use of state dependent tolls, i.e. tolls that are a function of the number of customers in the facility; it also includes analyses on providing priority service to customers incurring waiting costs at higher rates, where customers are induced to pick the appropriate priority via tolls. However, to the best of our knowledge, there is no research that shows how to use pricing mechanisms to optimally enable state dependent priority service in a practical manner.

One possible policy for optimally enabling state dependent priority service via tolls is to extend the use of state dependent tolls to an ordered set of queues, where customers in higher ordered queues are served before customers in lower ordered queues. Unfortunately, this policy raises several concerns, one of which is that the complexity of computing the optimal tolls, typically via policy iteration, grows very quickly in the number of queues and in the number of customers allowed in each queue.

Perhaps even more important are the remaining concerns. The first of these is the difficulty of determining expected waiting costs, and thus deciding to which queue, if any, customers should be admitted to. This difficulty arises because the expected waiting costs of customers joining lower priority queues are a function of the queues into which subsequent customers are accepted, which are in turn a function of those subsequent customers' expected waiting costs.

A second concern is that the variability in the possible benefits less waiting costs generated for individual customers when arriving to facilities in a particular state can be very large, which can result in dissatisfaction when benefits less tolls are smaller than expected. That variability is a function of the processing times of customers in the queue to which they will be accepted, as well as the variability of the arrivals of future customers who might join higher priority queues. To see this, consider a customer who is accepted to a lower priority queue. When no other customers are accepted into a higher priority queue before this customer is served, this customer's waiting cost will be relatively low. In contrast, when other customers are accepted to higher priority queues before this customer is served, this customer's waiting costs may be significantly higher.

The third of these concerns is the seeming disability of this policy to handle non-linear waiting costs. This issue arises when the rate at which a customer incurs waiting costs changes, either positively or negatively, at which time it might be desirable to move the customer into a different queue, so as to continue to minimize the waiting costs of all customers in the facility. Unfortunately this is not possible with this policy, as to allow it would further complicate the computation of the optimal tools. It is also not possible because it would entail tracking the state at all times of all customers waiting for service, which would in turn require an infinite state space, which would in turn preclude the computation of the optimal tolls.

To address these concerns, we construct a new policy in which a state dependent pricing mechanism is used to select customers that are admitted to a single queue. To minimize expected waiting costs, admitted customers are allowed to jump the queue, both when they arrive and when the rate at which they incur waiting costs changes, provided they pay customers, or are paid by customers, for changes in their expected waiting costs.

While this new policy is similar to the policy of applying state dependent admission control to an ordered set of queues, we show that the differences make it possible to address the concerns raised by the former policy. We start by showing that this new policy makes it easy to compute customers' expected net benefits. This is because the compensation mechanism inherent to this policy delays the need for computing the costs of prioritization until customers of higher priority arrive to the facility, so that each customer only need compute their expected waiting costs based on the customers already in the facility.

We next show that the new policy does not increase the variance of customers' expected benefits less waiting costs. This is because prioritization is implemented via jumping, and because customers already waiting for service are compensated for any changes in these amounts caused by the jumping.

To address the concerns related to non-linear waiting costs, we recall that the new policy allows customers, whose rate of incurring waiting costs change, to change their position in the queue, provided they compensate or are compensated for changes in waiting costs they cause by that jumping. To preclude the need to track the detailed state of each customer, we approximate customers' non-linear waiting costs as a series of piecewise linear waiting cost functions that customers of a particular group transition between at a random rate, which limits the state space to a large, but nonetheless, finite state space.

To address the concerns related to computing the tolls for this policy, we observe that the transition equations are relatively sparse, and suggest the use of value iteration to compute those tolls. We further observe that the state dependent values associated with the optimal policy computed by policy iteration of the first come first served version of this problem suggest a good starting point for the value iteration process. We also observe that it may not be necessary to fully compute via value iteration the values of all possible states, but that it may instead be possible to compute needed values in a just in time manner.

In addition to addressing the above concerns, this policy has a further advantage in that there are conditions under which it appears to be in each customer's best interest, at least in the short term, to state their gross benefits and waiting cost functions accurately.

Pricing with Markups in Competitive Markets with Congestion

CORREA, JOSE (Universidad Adolfo Ibanez) correa@uai.cl
STIER-MOSES, NICOLAS (Columbia University) stier@gsb.columbia.edu

Congestion games model competitive situations in which consumers must select resources in a network to minimize their individual costs (Rosenthal 1973). These games assume that the cost a consumer pays to each resource depends on the total utilization of that resource. For example, this structure readily models situations in which there is congestion or increasing marginal costs. A common interpretation is that these resources are exogenous to the model and respond to the laws of supply and demand. This means that when the demand for the resource is high, the price consumers need to pay to use the resource are going to be high too. For a set of cost functions for resources, Beckmann, McGuire, and Winsten (1956) show how an equilibrium for the consumers can be characterized and computed. Recent research in this area focused on measuring the inefficiency that arises from this competition, measured with respect to the solution that minimizes the total cost among all consumers.

In contrast to earlier models, this paper looks at making the actions of the resources endogenous. We add their pricing decisions as a first stage of the game, deferring the resource selection process to a second stage. Indeed, resources have a production function which maps the quantity they produce to the total production cost. Substitutable resources may have different production costs because of technological, investment or capacity differences. Since resources produce to make profit, they advertise to consumers a cost function that is higher than the production function. In this work, we assume that the producers’ decision is the (multiplicative) markup to apply to production costs. We highlight that producers do not advertise a fixed price but the cost function itself. Production functions are private information, known only to the corresponding producer.

Summarizing, the considered game is as follows: producers decide a markup and announce a cost function that is proportional to the production function, consumers learn those functions and place orders as in a standard congestion game. The market structure in the consumer market (the second stage) may range from one monopolist, to many oligopolists, to the case of perfect competition. In the first stage, producers play a game between them to fix markups, predicting what consumers are going to do in the second stage. Producers face the tradeoff of charging more to increase the revenue per unit but lose part of the market, or charging less to get more of the market but reduce the revenue per unit. Our approach relates to supply function equilibrium models, which are commonly used for modeling electricity networks (see, e.g., Rudkevich 2005). An important difference, though, is that we do not necessarily impose a market clearing price because consumers can buy from different producers at different prices, a situation that arises from the existence of congestion externalities and market power.

In some instances of the model, markups are large because some producers exploit their market power and their superior efficiency. In addition, when there is not sufficient competition in the producers' market, a producer may be able to extract more profit by increasing its markup, which may force other producers to react by also increasing their markups. Part of the problem of markups being possibly too large arises because our model considers that there is a fixed demand, which forces consumers to buy regardless of the price. Although the assumption of inelastic demand is natural in some applications, for others it is not. For that reason, we impose a condition to guarantee that the market is competitive enough. This condition assumes that no producer is significantly better that the average producer, meaning that nobody has a significant competitive advantage compared to the rest.

Under the assumption, we can prove that all producer markups are bounded. Using this fact and Brower's Fixed Point Theorem, one can prove that the game always has a Nash equilibrium. In addition, we measure the inefficiency generated by markups by looking at the total production cost under a Nash equilibrium and comparing it to that under the hypothetical solution if consumers based their decisions on the real production costs (which they do not have access to). It turns out that under the previously-mentioned assumption, and for the cases of substitutable and complementary resources, the cost increase is not significant. Without the assumption, though, the inefficiency introduced by the markups can be significant because markups themselves can be significant.

Another interesting phenomenon that arises from the study of this model is that increased cooperation does not lead to increased efficiency. For example, we compare the situations when the consumer market consists of a single monopolistic consumer and when it consists of infinite consumers without market power. Solving this model, all agents are worse-off when selling to a monopolist than when selling to the infinite consumers.

References

Beckmann, M. J., C. B. McGuire, and C. B. Winsten (1956). Studies in the Economics of Transportation. Yale University Press, New Haven, CT.

Rosenthal, R. W. (1973). A class of games possessing pure-strategy Nash equilibria. International Journal of Game Theory 2, 65–67.

Rudkevich, A. (2005). On the supply function equilibrium and its applications in electricity markets. Decision Support Systems 40, 409-425.

Session IV.2: CARGO AND FREIGHT (15:00-16:30)

Room: *S1-131* **Chairperson:** *Mikhail Nediak*

Carrier-Forwarder Contracts in the Air Cargo Industry

GUPTA, DIWAKAR (University of Minnesota) guptad@me.umn.edu

The air-cargo business, a 40 billion dollar industry worldwide, is increasingly a significant source of revenue for passenger airlines who carry nearly two thirds of worldwide airfreight either on passenger flights or on freight-only flights. This industry is forecasted to grow at an average annual rate of 6.4% over the next two decades (see Boeing Company 2005). Factors such as increasing international trade, leaner operations, and time-to-market based competition are driving rapid expansion of the air-cargo business (Dahl 2005, and West 2005).

The air-cargo service chain consists of shippers, consolidators (also called freight forwarders), carriers, integrators and consignees. Shipments originate from shippers who wish to have them delivered to the designated consignees. Shippers can use the services of an integrator, send the package directly through the carrier, or use a consolidator. Integrators are vertically integrated firms, having both transportation equipment (trucks, planes, etc.) and warehouses. In contrast, the carrier–forwarder combinations represent decentralized service chains. Carriers own planes whereas forwarders add value by performing services such as freight booking, pick-up (from shipper location), storage, import–export documentation (e.g. customs' clearance), consolidation of cargo into unit load devices at origin, breaking of bulk shipments at destination, and local delivery to the consignee. In addition, forwarders have marketing expertise and ties to local industry, which helps to generate demand that is not directly accessible to carriers.

Carriers (C) typically preallocate capacity on specific flights to forwarders (F) at a discount price; see Bazaraa et al. (2001) and DeLain and O'Meara (2004) for additional details on carrier–forwarder interactions. Such preallocations may cover a period of a few months to a year at a time (Reuters 2001). Freight charges depend on the volume generated by the forwarder, the type of goods shipped, the strength of demand in that market segment, and competition (see Thuermer 2005 and Chiger 2005 for a discussion of air-cargo pricing). The amount of preallocated capacity, determined by the carrier, may depend on the market share of the forwarder, her historical record in filling up the allocated space, and the amount of anticipated ad-hoc demand.

Preallocations are sometimes guaranteed, but it is not unusual for the forwarder's cargo to be bumped to later flights during high-demand seasons. Moreover, even when allocations are guaranteed, carriers at a pre-specified time before the flight's departure, customarily take back any unused capacity from forwarders and attempt to sell this capacity to ad-hoc customers. The timing of this exchange can be up to 48 hours before the flight's departure. Ad-hoc prices are usually higher. This should be anticipated by the forwarders in capacity-constrained markets where capacity buyback would matter to them. However, ad-hoc prices may be lower in some cases, especially during slow periods. Note that ad-hoc demand is usually not large enough to use all available cargo space, which explains why carriers sell to forwarders in high-demand markets.

This talk will examine the efficiency of two variants of C–F contracts, one involving allocations that are not guaranteed and the other involving guaranteed allocations. These contracts are compared to the option of selling without recourse in two different problem settings. In the first instance, we consider a situation where the forwarder exerts a non-verifiable effort on promoting sales upon receiving its allotment from the carrier. This effort is increasing in the size of the allotment. However, since the marginal benefit to the forwarder is smaller than that of the service chain's, it may not exert the optimal amount of effort. Preallocations are not firm commitments by the carrier, but any attempt to repossess capacity must not adversely affect the forwarder ex-post. We examine the extent to which contracts with the provision to buyback capacity result in service chain efficiency. Full details can be found in Gupta (2007).

In the second instance, we consider a situation where forwarder's type, which determines its demand, contribution margin, and carrier's cost of serving the forwarder, is its private information. Here, preallocations are guaranteed, but unused capacity may be bought back at a pre-announced price. In addition to giving carrier an opportunity to sell unused capacity on an ad-hoc basis, capacity buyback serves as an ex-post signal to the carrier about the forwarder's type. This talk will present models that show that capacity buybacks are beneficial to the carrier. We also quantify the impact of using price-only contracts on optimal capacity allocations and the amount of informational rents paid by the carrier. This part of the talk extends results obtained in a recent paper by Amaruchkul et al. (2007).

The main take-away of the proposed presentation is that contrary to intuition, preallocations without a firm purchase commitment from the forwarder, may be better for the carrier than firm sales. Depending on the situation, the flexibility inherent in contracts involving guaranteed and non-guaranteed allocations permit the carrier to respond effectively to either informational asymmetry or non-verifiable effort.

References

Amaruchkul, K., W. L. Cooper, D. Gupta. 2007. Air-cargo capacity contracts under asymmetric information. Working Paper, University of Minnesota.

Bazaraa, M., J. D. Hurley, E. L. Johnson, G. L. Nemhauser, J. S. Sokol, E.-P. Chew, C.-H. Huang, I. Mok, K.-C. Tan, C.-P. Teo. 2001. The Asia Pacific air cargo system Research Paper No: TLI-AP/00/01. The Logistics Institute-Asia Pacific. National University of Singapore and Georgia Institute of Technology. Retrieved on January 06, 2007 from http://www.tliap.nus.edu.sg/Library/default.aspx.

Boeing Company. 2005. World air cargo forecast 2002–2003. Retrieved March 30, 2005 from http://www. boeing.com/commercial/cargo/exec_summary.html.

Chiger, S. 2005. Negotiating a better parcel carrier contract. Catalog Age 38.

Dahl, R. V. 2005. Airfreight rebound: Growth in shipments, need for replacement aircraft dictate doubling of fleet in next 20 years. Aviation Week & Space Technology 162(3) 57–59.

DeLain, L., E. O'Meara. 2004. Building a business case for revenue management. Journal of Revenue and Pricing Management 2(4) 368–377.

Gupta, D. 2007. Flexible Carrier-Forwarder contracts for air cargo business. To appear in the Journal of Revenue and Pricing Management.

Reuters. 2001. Lufthansa cargo confident on new risk-sharing contract Retrieved November 08, 2005 from http://www.tiaca.org/articles/2001/10/12/8FE5CB0F2C5E4A7580D200E810C9AC2D.asp.

Thuermer, K. E. 2005. Air cargo pricing: No simple matter. Logistics Management 44(11) 36–42.

West, A. 2005. Transportation: Commercial. Standard & Poor's Industry Surveys Retrieved January 05, 2007 from http://www.netadvantage.standardandpoors.com/NASApp/NetAdvantage/ showIndustrySurvey.do?code=trc&date=/trc_0805/trc_0805.htm.

Freight Transportation Contracting Under Uncertainty: A Real Options Approach

TSAI, MEI-TING (University of California, Irvine) meitingt@uci.edu
REGAN, AMELIA (University of California, Irvine) aregan@uci.edu
SAPHORES, JEAN-DANIEL (University of California, Irvine) saphores@uci.edu

In lean and demand-responsive logistics systems, orders need to be delivered rapidly, accurately and reliably, even under demand uncertainty. Increasing burdens on the industry motivate the introduction of new methods to manage transportation service contracts. One way to hedge transportation capacity and cost volatility would be to use concepts from the theory of real options to craft derivative contracts. Therefore, the purpose of this paper is to investigate the adoption of derivative contracts in trucking, which is the dominant mode of freight transportation. We present (1) necessary conditions for the emergence of a market for truckload (TL) options; (2) derivation of an option pricing model; and (3) numerical results showing the potential benefits for shippers and carriers of adopting derivative contracts in trucking.

1. Necessary conditions

To-date, ocean transportation is the only mode of transportation for which derivatives contract has been applied. Thus, we start by reviewing the development of derivatives markets in the maritime industry. Freight derivatives contracts are popular and effective tools for hedging freight rates in the shipping industry. Their introduction can be traced back to 1985 when the Baltic International Freight Futures Exchange (BIFFEX) was established. At the beginning, BIFFEX worked well. However, trading volumes began to fall in 1989 and unfortunately, contracts ceased to exist in 2002. The termination of BIFFEX was mainly due to low liquidity and poor hedging performance. The underlying asset of BIFFEX contracts, the Baltic Freight Index (BFI), was a weighted average of spot prices from 11 shipping routes. The weighting and composition of the index changed over the years. Moreover, the BFI was diverse in terms of cargoes and vessel sizes, whereas most market agents work on specific routes, so they demand contracts tailored to their specific needs. Consequently, BIFFEX contracts did not perform well as hedging instruments. In spite of these problems, the hedging function of freight derivatives contracts is regarded positively by many in this industry. Different types of contracts have been launched since 1995, and forward freight agreements (FFAs), Futures, and Options are currently available for trading. Either FFAs or Futures contracts obligate the seller and the buyer with a predetermined price at a specified future time for an agreed route. They are agreements between two parties where the payoff is the difference between the spot price and the contract price. However, FFAs are traded in over-the-counter markets while Futures are traded on an exchange. Options provide the holder with the right, but not the obligation, to buy or sell an asset at a predetermined price on or before a predetermined date. They are all financial settlement which means no physical shipping capacity is involved in the contracts. The Baltic Exchange, NYMEX, and IMAREX are the three main markets for trading these contracts. Each has specific products and trading rules, but their common characteristic is increasing trading volumes.

Based on our findings and on the experience accumulated in maritime industry, we draw conclusions for necessary conditions for the emergence of a market for truckload (TL) options.

2. TL options pricing model

As in the maritime industry, the demand for derivatives contracts for hedging uncertainty in trucking has appeared. Unlike the maritime industry, however, trucking needs derivatives not only for hedging price uncertainty but also for capacity uncertainty. In other words, the current financial-settlement derivatives used in the maritime industry are not sufficient for hedging demand uncertainty in the trucking industry because those derivatives are operated separately from the physical shipping markets. As a result, the price risk may be fixed but the capacity for delivering cargos is not guaranteed. Hence, we propose trucking capacity option contracts. As a starting point, this paper focuses on truckload services in spot markets with European style options.

Definition 1. A European TL call option is a derivative contract that names the buyer a right, not an obligation, to buy one truckload at the predetermined (strike) price K on predetermined (expiration) date T. The payoff of the European TL call option depends on the strike price and the spot price on expiration date S_T, which is given by

$$h_c(S_T) = \max[(S_T - K), 0]. \qquad (1)$$

A European TL put option is a derivative contract that gives the buyer the right (it is not an obligation) to sell one truckload at a predetermined (strike) price K on a predetermined (expiration) date T. The payoff of the European TL put option depends on the strike price and the spot price on expiration date S_T. It is given by

$$h_P(S_T) = \max[(K - S_T), 0]. \qquad (2)$$

The price of an option is the present value such that the expectation of future payoff equal to zero, which is known as free-arbitrage pricing. Specifically, the values of call and put options are given by equation (3) and (4), respectively.

$$C_t = E[e^{-r(T-t)} h_c(S_T) | F_t] \qquad (3)$$

$$P_t = E[e^{-r(T-t)} h_p(S_T) | F_t] \qquad (4)$$

where $h(S_T)$ is the payoff function defined as equation (1) and (2); S_T is the TL spot rate; r is the interest rate; t and T denote the present and maturity, respectively; F_t is the information filtration at time t which describes the information available about the spot rate up to time t. Obviously, in addition to defining the payoff functions, we need to identify the TL spot rate function S_t so that the prices of options can be obtained. As argued by Dixit and Pindyck (1994), we can expect current prices to be related to long-run marginal costs. While prices move up and down in the short-run, they will eventually revert back to long run marginal costs in a competitive market. Following this argument, we assume that the logarithm of the logarithm of spot prices follows a mean-reverting process such as the Ornstein-Uhlenbeck process (Karlin and Taylor, 1981).

Given these assumptions, we first develop the partial differential equation (PDE) for a TL call option. A well-known analytic formula of TL call option price is derived by solving the resulting PDE. Next using the put-call parity, we derive a formula for the price of a TL put option.

3. A numerical example

Obtaining reliable truck load spot price data for this type of research is not a simple task because to-date there does not appear to be a unique spot market for truck loads on various routes. It is therefore not surprising that there are still no published papers on optimal contracting in the trucking industry (for example the many papers on the use of variants of combinatorial auctions) based on real data from the industry. Delivery service prices for specific lanes are not available on a daily basis; instead, electronic marketplaces provide monthly statistics (such as minimum, maximum, and average). We use these statistics to infer the parameters of the underlying spot price for truck loads. We then rely on this information and our option pricing formula to present a numerical example of TL option pricing.

Cargo Revenue Management with Allotments and Spot Market Demand

LEVIN, YURI (Queen's University) ylevin@business.queensu.ca
NEDIAK, MIKHAIL (Queen's University) mnediak@business.queensu.ca
TOPALOGLU, HUSEYIN (Cornell University) ht88@cornell.edu

We consider a problem of an airline that operates a single flight leg to transport cargo between a particular origin destination pair. The airline can sell its cargo capacity either through an allotment contract or on the spot market. The goal is to simultaneously choose the allotment size and find a booking control policy for the booking requests on the spot market so as to maximize the total expected profit. We formulate the booking control problem on the spot market as a dynamic program, and construct approximations to its value function to estimate the total expected profit from the spot market. This allows us to choose the allotment size by maximizing the sum of the revenue from the allotments and the total expected profit from the spot market. We show that our value function approximations provide upper bounds on the total expected revenue from the spot market and our booking control policy satisfies several desirable properties. We also examine trade-offs between the spot market and allotment sales.

Session IV.3: DYNAMIC PRICING III (15:00-16:00)

Room: *S1-139* **Chairperson:** *Houyuan Jiang*

Optimal Prices and Advertising Expenditure with Inventory Constraints

RASMUSSEN, HENNING (University of Western Ontario) hrasmuss@uwo.ca
MACDONALD, LEO (Kennesaw State University) lmacdon4@kennesaw.edu

One of the most important problems in modern business practices is the determination of the optimum price for a product and the optimum expenditure on advertisement. Most published papers deal with either of these procedures in isolation instead of considering the joint pricing/advertising problem. In this paper we consider the joint problem for both deterministic and stochastic models for both a monopoly situation (new product introduction) in the presence of inventory constraints.

During the last few decades, a large number of papers has been published on the topic of dynamic pricing of new products in both monopolistic and oligopolistic settings. The Bass model[1] of course started a stream of research in new product diffusion models, with numerous extensions adding control variables such as price and advertising. The early work in this area is summarized in [4]. At first these models did not consider demand uncertainty, but since then a limited amount of research has been done on incorporating uncertainty into the models, within monopolistic as well as oligopolistic markets, such as Eliashberg and Jeuland [3], and Raman and Chatterjee [6]. Similarly a large number of papers have appeared which treat the problem of optimal advertisement for both monopolistic and oligopolistic situations. Among more recent work we can mention [5] who use a stochastic differential games approach to study the oligopolistic case. A combined study of the optimal pricing and advertising policies seems to only have been considered by [7] who studied a deterministic model for the oligopoly case.

In this paper, we analyze a dynamic pricing and advertising strategy of a new differentiated product for a single firm over a finite planning horizon and the main new feature of our research is that we incorporate the effects of inventory constraints and salvage. The sales dynamic of this product is modeled as a linear stochastic process which incorporates diffusion effect (especially word of mouth effect) and demand uncertainty, and we focus on the uncertainty in terms of the inability of the firm to predict demand with complete certainty, and specifically, we are considering the uncertainty pattern that increases linearly in cumulative sales, i.e. as cumulative sales increase, the market becomes diffuse and less certain. We formulate this stochastic disturbance via a Wiener process. Since we are not considering the saturation effects as in the Bass model, our model is primarily designed for non-durable products or durable products in their very early stage of their life cycles when diffusion effects are dominate. We set up a limit on inventory and a 'penalty' for selling more than the limit at the end of planning time and also we have a salvage value associated with the inventory state at the end of time horizon.

In the monopolistic setting, we are using dynamic programming in order to maximize the total expected profit and hence we obtain Bellman's equation to get the optimal price and advertising strategies. We analyze the implications of demand uncertainty on the price strategy by comparing stochastic and corresponding deterministic demand situations. We get closed-form solutions of the optimal strategies for the deterministic model, but we can only solve numerically the problem for the stochastic model, which we use to demonstrate how the solutions and optimal policies are impacted by the parameters, particularly, the demand uncertainty.

References

[1] Bass, M.F. (1969), "A new product growth for model consumer durables", Management Science, Vol.15, pp. 215-227.
[2] Dockner, E.J. and Jorgensen S. (1988), "Optimal pricing strategies for new products in dynamic oligopolies", Marketing Science, Vol. 87, pp. 315-334
[3] Eliashberg, J. and Jeuland, A. (1986), "The impact of competitive entry in a developing market upon dynamic pricing strategies", Vol. 5, pp. 20-36
[4] Kalish, S. (1983), "Monopolist pricing with dynamic demand and production cost", Marketing Science, Vol. 2, pp. 135-159

[5] Prasad, A. and Sethi, S. P. (2004), "Competitive advertising under uncertainty: A stochastic differential game approach", J. Optimization and Applications, Vol. 123, pp. 163-185.
[6] Raman, K. and Chatterjee, R. (1995), "Optimal Monopolist pricing under demand uncertainty in dynamic markets", Management Science, Vol. 41, pp. 144-162
[7] Thompson, G. L. and Teng, J. T. (1984) "Optimal pricing and advertising policies for new product oligopoly Models", Marketing Science, Vol. 3, pp.148-168.

Dynamic Pricing of Perishable Goods for a Monopoly and a Duopoly in the Presence of Strategic Consumers

MANTIN, BENNY (University of Waterloo) bmantin@uwaterloo.ca
GRANOT, DANIEL (University of British Columbia) daniel.granot@sauder.ubc.ca

We introduce and analyze monopoly and duopoly dynamic pricing models in the presence of both myopic and strategic consumers. Considering a two-period horizon, we assume that in a monopoly, active consumers visit the same retailer in both periods, and in a duopoly, each of them visits one of the retailers in the first period and switches to the competing retailer in the second period. While the myopic consumers are impatient and purchase the good as soon as they encounter a price which is below their valuation, the strategic consumers, who consider the prices along the path of their visit pattern and their discounting factor, purchase the good, if at all, when their surplus is maximized, as long as it is positive.

The pricing policy at equilibrium prescribes similar actions in both a monopoly and a duopoly - to skim higher valuation consumers in the first period. More precisely, the pricing policy employed by retailers depends on the market composition and consumers' discounting factor. When there are sufficiently many strategic consumers, or when the discounting factor of the strategic consumers is sufficiently low, the retailers skim high-valuation consumers of both types in the first period (Pricing Policy i, or PP(i)). Otherwise, they skim only myopic high-valuation consumers in the first period and none of the strategic consumers purchases the good in the first period (PP(ii)). When retailers switch from PP(i) to PP(ii), prices at both periods spike upwards, as to address demand diverted from the first to the second period.

We further show that strategic consumer behavior inflicts losses to retailers, and under competition the relative loss is larger. However, while a monopoly is always better off with a lower consumers' discounting factor, this is not necessarily true for competing retailers. That is, a monopolistic retailer should investigate ways to mitigate the strategic consumers' patience, e.g., by enhancing the "virtual" value of the good, while, on the other hand, such patience mitigation may result, in some situations, with a lower duopoly profit.

Ignorance of strategic consumer behavior is not a wise strategy for a monopoly. Yet, this may be beneficial for competing retailers. In other words, by being oblivious to strategic consumer behavior the monopoly is always worse off, while the duopoly may be better off. Moreover, when retailers treat all consumers as myopic, the duopoly profit may exceed the monopoly profit.

Extensions of the analysis are studied for (*i*) the case of perfect markets, wherein both retailers and strategic consumers share a common discounting factor, (*ii*) capacitated settings, and (*iii*) longer horizons, wherein we characterize the equilibrium pricing policy and resulting expected profit when all consumers are strategic. The analysis of the capacitated settings reveals that with a limited inventory, retailers update their prices such that, in expectation, all inventory is depleted over the course of the two-period selling horizon by matching the corresponding demand. If inventory is sufficiently limited, then retailers resort to a fixed price policy and all inventory is depleted in the first period, in which case strategic consumer behavior is eliminated.

Dynamic Capacity Allocation and Pricing under Competition

JIANG, HOUYUAN (University of Cambridge) h.jiang@jbs.cam.ac.uk
LEVIN, YURI (Queen's University) YLevin@business.queensu.ca
NEDIAK, MIKHAIL (Queen's University) MNediak@business.queensu.ca

A monopolistic company provides certain services/products to customers over a planning horizon $[0, T-1]$, which is divided into T equally-spaced time periods indexed by t. The company repeats the same services over the planning horizon, with a fixed capacity in each period, which may be a function of t. Service capacity may have more than one attribute such as weight and volume.

Customer services are decomposed into a number of non-overlapping types, indexed by k. Customers make bookings of services in advance. Each booking is completely specified by its type k, booking time t', check-in time t (time when the item is ready for service), and due time t'' (nominal time for service to be completed) such that $t' \leq t \leq t''$. The total demand for each class of bookings $kt'tt''$ is assumed to be a deterministic non-negative real number, and is a function of prices for all possible values of t and t''. This demand function says that the demand for product $kt'tt''$ is completely determined by the prices, announced at period t', of all products of type k. We are particulary interested in the linear demand function.

A booking of class $kt'tt''$ may be serviced between check-in time t and due time t''. Because of a capacity limit at time t, some bookings due at t may not be serviced before or at t. Those unserviced bookings will be outsourced to a third party. At the beginning of period t, the company announces its prices, and, at the end of period t, the company decides which checked-in bookings are serviced, and which checked-in bookings due at t are outsourced.

We make the following contributions to the literature. First, we prove that a fixed price policy for a monopolist is asymptotically optimal to a dynamic price policy. Second, we study monotone properties of prices for a monopolist. Third, we show that the strategy profile based on the fixed price policy in a competitive environment is an ε-Nash equilibrium of the dynamic pricing game. Fourth, we study comparative statics results for both monopolistic and oligopolistic pricing problems. Lastly, we provide insights on capacity allocation policies.

Musée Pointe-à-Callière

17:30 *Banquet*

Two buses leave in front of the Pavillon Jean-Coutu at 16:45.

FRIDAY, JUNE 20, 2008

Room: *Agora*

8:00 *Breakfast*

TRACK V

Session V.1: MATHEMATICAL METHODS (9:00-10:30)

Room: *S1-111* **Chairperson:** *Guillaume Roels*

Revenue Management of the Risk-Averse Newsboy with Atemporal Untity

LI, MICHAEL Z.F. (Nanyang Technological University) zfli@ntu.edu.sg
ZHUANG, WEIFEN (Nanyang Technological University) U050023@ntu.edu.sg

This paper studies the impact of risk-aversion and risk on the static two-class newsboy model in revenue management. We compare the optimal booking policy between risk-neutral and risk-averse and we have shown that the optimal booking limit under risk-aversion is monotone on absolute risk-averse factor for any atemporal utility function. We have also examined the impact of an increase in demand risk in terms of three risk measurements – first-order, second-order and third-order stochastic dominance on the optimal booking control policies. Numerical examples will be used to illustrate and support our theoretical results.

Irrevocable Dynamic Assortment Strategies

FARIAS, VIVEK (Massachusetts Institute of Technology) vivekf@mit.edu
MADAN, RITESH (Qualcomm-Flarion Technologies) rkmadan@stanfordalumni.org

- We introduce a new 'irrevocable' heuristic for the multi-armed bandit problem with multiple simultaneous arm pulls. Computational experiments for a generative family of bandit problems with up to 500 arms indicate losses of up to about 8-10% percent relative to an optimal strategy with *no restriction* on exploration (i.e. an optimal strategy that is *allowed recourse* to arms that were discarded from the assortment in the past). In addition, we are able to demonstrate a *uniform* bound on performance loss relative to such an optimal strategy for a general class of bandits - our heuristic earns expected rewards that are always within a factor of 8 of an optimal, potentially non-irrevocable policy. Such a bound represents the first non-asymptotic performance guarantee for problems of this type and employs a novel combinatorial analysis that draws on recent developments on the 'adaptivity gap' for stochastic packing problems.
- Assuming that an individual arm has $O(\Sigma)$ states and a time horizon of T steps, our heuristic requires a total $O(n\Sigma^2 \log T)$ computations per time step amortized over the time horizon. In comparison, the simplest theoretically sound heuristics in existence for this problem (such as Whittle's heristic) require $O(n\Sigma^2 T)$ computations per time step. We consequently establish that our heuristic is computationally attractive. Further, the heuristic has an appealing index interpretation.

The computational simplicity and performance of our heuristic in combination with its 'irrevocable' property, makes it a potentially useful solution to practical dynamic assortment problems.

Consider the operations of a 'fast-fashion' retailer such as Zara. As a consequence of their highly flexible procurement capabilities, such retailers are able to adjust the assortment of products offered on sale at their stores to quickly adapt to popular fashion trends. This operational model relies crucially on an effective technology to learn from purchase data, and adjust product assortments based on such data.

The multi-armed bandit is a convenient mathematical model for the design of dynamic assortment heuristics. Such heuristics typically call for a fair amount of 'exploration'. For instance, in the fast-fashion context, such heuristics may choose to discard from the assortment a product presently being offered for sale in favor of a new product whose popularity is not known precisely but later choose to re-introduce the discontinued product. While such exploration may appear necessary if one is to discover profitable bandit arms, enabling such a heuristic in practice will typically call for a great number of adjustments to the product assortment - a requirement that is both expensive and undesirable. This begs the following question: Is it possible to design a heuristic for the multi-armed bandit problem that comes close to being optimal with a minimal number of adjustments to the set of arms pulled over time?

With this question in mind, we consider a multi-armed bandit with n arms, $k \geq 1$ of which may be pulled simultaneously at any given time. We wish to maximize expected rewards earned over a finite horizon of T time periods. As an example, one may think of each of these arms as coins of an unknown bias. Flipping a coin garners potential rewards and further allows us to refine our knowledge of the coins bias. We introduce a new heuristic for the multi-armed bandit problem that possesses the following 'irrevocable' property: should an arm presently in the 'assortment' of arms pulled be removed from the assortment at some point in time, it will never again be re-introduced into the assortment. It is clear that such a heuristic makes the minimal possible number of changes to the assortment of arms pulled over T time steps. What is perhaps surprising, is that the restriction to an irrevocable policy is for many relevant classes of bandits far less expensive than one might expect. More specifically, we offer the following contributions:

Approximate Dynamic Programming Approach to Network Revenue Management with Customer Choice and Concave Approximation of the Value Function

STRAUSS, ARNE KARSTEN (Lancaster University Management School) a.strauss@lancaster.ac.uk
MEISSNER, JOERN (Lancaster University Management School) joe@meiss.com

Network capacity models attempt to maximize some reward function when customers buy bundles of multiple resources. The interdependence of the resources, commonly referred to as network effects, provides difficulty in solving the problem. A notable, now classical approach was taken by Glover et al. (1982) by proposing a deterministic LP (DLP) solution to the network capacity problem in an airline context which triggered considerable research in possible reformulations and extensions, and became a standard in many industrial applications. While easy to implement and solve, it is well-known that DLP has many disadvantages because of its simple structure and restrictive assumptions. A significant limitation on the applicability of classical models is the assumption of independent demand; customers do not usually arrive in the system with only one particular product in mind but rather are faced with choices among a variety of alternatives. In response to this, interest has arisen in recent years to incorporate customer choice into these models and to make better use of stochastic information, further increasing the complexity.

This development drives current efforts to design powerful and practical heuristics that still can manage problems of this scope. Gallego et al. (2004), for example, proposed a method that translates the DLP model into the choice-based context (called CDLP). Because of its analytical elegance, the linear programming approach to ADP as presented for example in de Farias and Van Roy (2003) has seen considerable attention, most notably in the papers of Bront et al. (2007), Zhang and Adelman (2007), Kunnumkal and Topaloglu (2007) and Farias and Van Roy (2007). In particular, the latter show that the linear programming approach to ADP can outperform the widely used classical DLP method.

We improve on the network revenue management models with customer choice of Zhang and Adelman (2007) by approximating the value function of the Markov decision process formulation with a concave function which is separable across resource inventory levels as introduced by Farias and Van Roy (2007) in the no-choice context. This approximation reflects the intuitive interpretation of diminishing marginal utility of resource levels and allows for significantly improved accuracy compared to currently available approaches. The resulting approximation yields provable tighter bounds on the exact optimum than the approaches of van Ryzin and Liu (2007), Zhang and Adelman (2007) and Kunnumkal and Topaloglu (2007). As a consequence of the obtained bound being tighter than the one of CDLP, we have asymptotic optimality of our proposed approximation as demand, capacity and time horizon are linearly scaled up. A seemingly new upper bound relationship between the approaches of Zhang and Adelman (2007) and Kunnumkal and Topaloglu (2007) is shown, namely that the former provides a tighter upper bound on the objective value than the latter.

The increased accuracy of our approach comes at the expense of a more complex system that needs to be solved. However, efficient computation is achieved by column generation techniques for the multinomial logit choice model with disjoint consideration sets. For this case, we show that the column generation subproblem can be reduced to a linear mixed integer program. Similar to the results of Zhang and Adelman (2007), our approximation uses only efficient sets in the optimal solution, where efficiency is the most beneficial trade-off between the expected total revenue and expected total resource consumption. Computational experiments provide insights on the relative performance of policies that are based on our approach.

A Riskless Price Heuristic for the Newsvendor

ROELS, GUILLAUME (University of California) groels@anderson.ucla.edu

To cope with demand uncertainty, firms can choose to hold safety stocks and/or charge risk premiums. The role of safety stocks as a hedge against demand uncertainty has been extensively studied with the newsvendor model, and the insights are generally robust: Independently of the demand distribution, a firm needs to hold positive safety stock when the underage costs dominate, and negative safety stock when the overage costs dominate.

In contrast, the role of risk premiums is ambiguous, as different models reach different conclusions. In particular, it has been shown that a *positive* risk premium should be charged when the impact of uncertainty on the demand curve is *multiplicative*, while a *negative* risk premium should be charged when the impact of uncertainty is *additive*. This divergence in insights is somewhat disturbing, especially because most companies do not know their demand curve, let alone the impact of uncertainty on their demand curve.

In this talk, we offer new perspectives on this problem. We first show that the riskless price (i.e., the price that ignores demand uncertainty) leads to positive safety stock if and only if the demand elasticity is smaller than 2, independently of the demand model (e.g., additive, multiplicative, Poisson).

Using worst-case bounds over all demand probability distributions, we then demonstrate that it is relatively robust to price at the riskless price. In particular, we show that, for the additive demand model, the riskless price is always larger than 2/3 the optimal price. We then show that the riskless price can sometimes be larger than the optimal price when demand is restricted to be nonnegative, in contrast to the classical result by Mills (1959), but that it is never larger than twice the optimal price. For the multiplicative demand model, we prove that, even if the riskless price can be very different from the optimal price, the loss of optimality from choosing the riskless price is generally small, and can be bounded by a constant in special cases. We then show through numerical examples that the loss of optimality is generally within a few percents from the optimal profit.

Combining our results, we conclude that charging the riskless price is a robust decision, leading to positive (negative) safety stock whenever the demand elasticity is smaller (larger) than 2.

Session V.2: PRICING (9:00-10:30)

Room: *S1-131* **Chairperson:** *Guillermo Gallego*

Optimal Timing of Price-Quote Revisions

SOUSA LOBO, MIGUEL (Duke University) mlobo@duke.edu

The practice of offering individual discounts to buyers is more common in some industries than others, and some sellers are quicker than others to offer a discount as an incentive for a specific buyer to complete a transaction. How can the seller determine the optimal timing of individual discounts, and how deep should the eventual discount be?

While it is diffcult to obtain a precise estimate for the number and total value of transactions in which the price is negotiated, this seems to be the dominant mode in business-to-business transactions which, given the length of modern value chains, comprise the majority of economic transactions by value. While there exists an extensive literature relating to this problem, particularly in the fields of industrial organization and bargaining, it has received little attention in the revenue management literature which has mostly focused on posted-price transactions.

There is strong support, both empirical and from theory, to believe that it can be of benefit to a seller to negotiate prices individually with each buyer, and that the equilibirum strategy can result in a negotiation delay in transactions. The main goal of this article is to take a step in operationalizing the problem from the seller's side. We discuss how the seller can estimate the aggregate buyers' strategy from sales history, and determine the optimal response under the practical constraint of discrete price revisions.

A key feature of our model is that the buyer's evaluation and decision process is not instantaneous, nor is the search for alternative sellers. High-valuation buyers may either find an alternative or accept the seller's offer. Low-valuation buyers can only terminate their search if a lower-cost alternative is found. It can then be expected that high-valuation buyers will terminate the search at a higher rate than low-valuation buyers. If this is the case, conditional on a buyer not having yet made a purchase, the probability that the buyer is low-valuation increases over time, and the seller has an increasing incentive to offer a lower price. A partial form of third-degree price discrimination is achieved by making low-valuation buyers wait for a discount. Our analysis shows, with plausible model parameters, significant gains in expected revenue for the seller from such a policy (see Table 1) and, which is crucial to justify the widespread use of negotiated prices and individual discounts, that these gains can be robust to model uncertainty.

	(i) $q_1 = 0.05$ $q_2 = 0.55$	(ii) $q_1 = 0.10$ $q_2 = 0.50$	(iii) $q_1 = 0.20$ $q_2 = 0.40$
Full price discrimination (upper bound)	41.7%	83.3%	33.3%
$\beta = 0.2$	12.5%	48.5%	16.2%
$\beta = 0.5$	2.3%	32.1%	8.1%
$\beta = 1$	0.0%	20.8%	3.3%
$\beta = 2$	0.0%	12.3%	0.8%
$\beta = 5$	0.0%	5.6%	0.0%

We derive a number of structural and technical results of interest. We provide a detailed characterization of the two-price model, including general conditions for uniqueness of the optimal policy and a closed-form solution under the assumption of exponential distributions on the random search and purchase-delay times. The model also provides support for a result which may be counter-intuitive, but is consistent with anecdotal evidence: if alternative suppliers are harder to find and the seller has more monopoly power, then the seller is more likely to offer substantial individual discounts to some buyers. For the multiple-price problem, we provide a dynamic-programming formulation and establish detailed structural properties. We are able to propagate quasi-concavity of the maximand in the decision variable by propagating a number of properties of the value function including, somewhat surprisingly, convexity in one of the parameters.

Models used in the revenue management literature generally presume the availability of at least approximate estimates (or, more commonly, exact knowledge) of the demand function. However, if prices are set on an individual basis, the price-elasticity of demand cannot be estimated from historical data. The number of sales at each price do not necessarily reect the distribution of valuations but, rather, the outcomes of the negotiation processes, which is to say, how often salespeople settled on different price discounts. However, given a sufficient number of sales, much information is available that makes the problem amenable to a quantitative revenue management approach. If information about all customer contacts is extant, the seller can estimate the probability distribution of the delay in buyer purchases, and estimate the effect of an earlier or later timing of the offer of a discount. This can be interpreted as estimating a time-elasticity of demand. We discuss implementation and statistical issues of identifiability and estimation, including Bayesian maximum-likelihood estimates and full estimation of the posterior distribution by Markov-chain Monte Carlo integration. While we find that there are identifiability concerns in the problem under study, with the appropriate estimation procedures, and with decisions that take the uncertainty into account, the seller can nevertheless derive substantial gains in expected revenue.

Pricing Access Services and Warranties

GALLEGO, GUILLERMO (Columbia University) gmg2@columbia.edu
SAHIN, OZGE (University of Michigan) ozge@umich.edu

This paper is motivated by recent empirical evidence on time inconsistency and overestimation behavior by consumers. We investigate the structure of the contracts that a firm should offer in industries that charge access fees to consumers. Typical examples are extended warranties sold with consumer electronics, club memberships, printing services, and on-demand financial services. In our model the provider offers a menu of contracts in the form of three part tariffs and the consumer self selects the contract that maximizes their utility. After the purchase of the contract, consumption takes place over time. We explicitly model the consumer's utility as a function of her optimal consumption over time. The goal of this paper is to provide guidelines to managers on pricing access services and warranties if the consumers are facing time inconsistency and have different risk preferences.

Upgrades, Upsells and Pricing in Revenue Management

GALLEGO, GUILLERMO (Columbia University) gmg2@columbia.edu
STEFANESCU, CATALINA (London Business School) cstefanescu@london.edu

Capacity providers often experience a mismatch between supply and demand that can be partially alleviated while improving revenues by allowing for product upgrades. When prices are fixed and demands are independent, the problem is to decide which customer demands to upgrade to which products and when. We show that a fairness constraint can be imposed without loss of optimality under mild conditions. We also investigate a model that limits upgrades to the next higher quality product and provide necessary and sufficient conditions for its revenues to be as high as that of any less restricted upgrade model. Resellers of capacity also have an incentive to use upgrades as a mechanism to entice customers to higher quality products with higher commission margins. We show that this practice can be very profitable and that the profits can be much larger than direct commissions from sales would indicate. This suggests that primary providers may be in a more powerful position to negotiate commissions or impose sales volume constraints. We then investigate the case where sellers have pricing flexibility and customer demand is driven by a choice model. We derive pricing formulas under the assumption that demand for products follows a multinomial logit model. For this model we show that neither upgrades nor upsells can improve profits when margins are homogenous and there is complete freedom in selecting prices. However, upgrades can improve revenues significantly when margins are heterogeneous and sensible business constraints on prices are imposed.

Session V.3: CUSTOMER CHOICE II (9:00-10:30)

Room: *S1-139* **Chairperson:** *James Dana*

Dynamic Assortment Optimization with a Multinomial Logit Choice Model and Capacity Constraint

RUSMEVICHIENTONG, PAAT (Cornell University) paatrus@cornell.edu
SHEN, MAX (University of California-Berkeley) shen@ieor.berkeley.edu
SHMOYS, DAVID (Cornell University) shmoys@cs.cornell.edu

Given limited resources such as shelf capacity and budget, a retailer must determine which assortment of products to offer to maximize profits. In some cases, the retailer does not know the demand *a priori*, and in fact, the probability that a customer will purchase a product depends on the assortment offered. The retailer can learn the demand by offering different assortments of products and estimating the demand from the resulting sales and assortment decisions.

We formulate a stylized model that captures some of the issues commonly present in assortment optimization problems, including the capacity constraint, the uncertainty in the demand distribution, and the dependence of the purchase probability on the assortment decision. Assume that we have N products indexed by $1, 2, \ldots, N$. For each $i = 1, \ldots, N$, let $w_i > 0$ denote the marginal profit of product i. Due to a capacity constraint, we can offer at most C products to the customers. The goal is to determine the subset of at most C products that will yield the maximum profit.

To account for dependence of the selection probability on the assortment decision, we use a multinomial logit (MNL) choice model, which is one of the most commonly used choice models in economics, marketing, and operations management (see, for example, McFadden (1974), Ben-Akiva and Lerman (1985), Mahajan and van Ryzin (1998), and the references therein). Under the MNL choice model, it well known that, for each assortment $S \subseteq \{1, \ldots, N\}$, the probability $\theta_j(S)$ that a customer chooses product j is given by:

$$\theta_j(S) = \begin{cases} e^{\mu_j} / \left(1 + \sum_{k \in S} e^{\mu_k}\right), & \text{if } j \in S, \\ 1 / \left(1 + \sum_{k \in S} e^{\mu_k}\right), & \text{if } j = 0, \\ 0, & \text{otherwise.} \end{cases}$$

where for each j, μ_j denote the mean utility that a customer assigns to product j. We denote the option of no purchase by 0 and set $\mu_0 = 0$. We define $v_j = e^{\mu_j}$ for all j. Then, the expected profit $f(S)$ associated with the assortment S is given by $f(S) = \sum_{j \in S} w_j \theta_j(S) = \frac{\sum_{j \in S} w_j v_j}{1 + \sum_{j \in S} v_j}$. The assortment optimization problem, which we refer to as the CAPACITATED MULTINOMIAL LOGIT ASSORTMENT PROBLEM **(CMNL)**, is defined by: $Z^* = \max_{S \subseteq \{1,\ldots,N\}: |S| \le C} \; f(S)$.

We are interested in the setting where the mean utilities μ_i's are *unknown*, and we have to infer these values by offering different assortments and estimating the parameters based on the resulting purchases. In this case, we aim to develop an adaptive policy that generates a sequence of assortments $S_1, S_2, \ldots,$ such that $|S_t| \le C$ for each t, and S_t depends only on the customer selections and assortment decisions in the previous $t-1$ periods. As a performance measure, we will consider the running average expected profit $\frac{1}{T} \sum_{t=1}^{T} E\left[f(S_t)\right]$, and we want this average to converge to Z^* – the maximum expected profit that we can achieve had we known the true mean utilities – as the number of time periods T increases to infinity.

Our approach consists of two parts. In the first part, which we refer to as *Static Optimization*, we assume that the mean utilities $\mu_1, \ldots, \mu_N$ are known in advance, and focus on finding a solution to the **CMNL** problem. We present a polynomial-time algorithm for finding the optimal assortment and establish structural properties of the solution. In the second part, which we refer to as *Dynamic Optimization*, we leverage the insights from the static problem and extend the algorithm to the setting when the mean utilities are unknown and must be estimated from historical data. We now describe the contributions and key insights of the paper.

When $C = N$, Gallego et al. (2004) and Liu and Van Ryzin (2004) show that the optimal solution to the **CMNL** problem can be found among the following N assortments: $\Lambda_1 = \{\lambda_1\}$, $\Lambda_2 = \{\lambda_1, \lambda_2\}, \ldots, \Lambda_N = \{\lambda_1, \ldots, \lambda_N\}$, where $(\lambda_1, \ldots, \lambda_N)$ represents the ordering of the products based

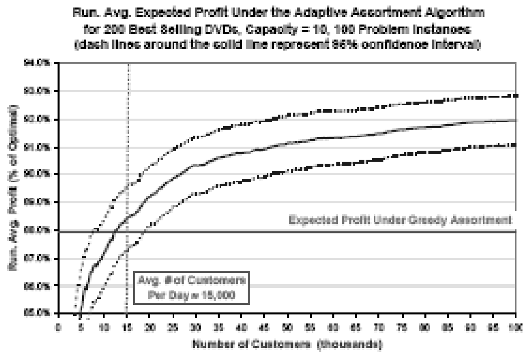

Figure 1: The running average expected profit under our proposed adaptive policy (solid blue line). The dash lines above and below the solid line represent 95% confidence intervals. The greedy assortment corresponds to the top 10 most profitable products.

on their marginal profits, that is, $w_{\lambda_1} \geq w_{\lambda_2} \cdots \geq w_{\lambda_N}$. For $C < N$, however, we show that restricting our attention to the top C most profitable products can lead to suboptimal solutions. Using the idea pioneered by Megiddo (1979), we then develop a polynomial-time algorithm for solving the **CMNL** problem that has a running time of $O\left(N^2\right)$. The algorithm generates a sequence $\langle A_0, A_1, \ldots, A_K\rangle$ of assortments that guarantees to contain an optimal solution with $K = O(N^2)$. Given the widespread use of the MNL choice model, we believe it is important to bring the community's attention to this method, and we have provided a detailed description of the algorithm in the full paper as it applies to our setting.

By exploiting the specific structure of our multinomial logit choice model, we derive a new upper bound associated with the sequence $\langle A_0, A_1, \ldots, A_K\rangle$ generated by Megiddo's method. We show that the number of distinct subsets among $A_0, A_1, \ldots, A_K$ is at most $C(N + C - 1)$. Furthermore, we show that the expected profit associated with the sequence $\langle A_0, A_1, \ldots, A_K\rangle$ is *unimodal* in the sense that there exists an index m^* such that $f\left(A_1\right) \leq f(A_2) \leq \cdots \leq f\left(A_{m*}\right)$ and $f\left(A_{m*}\right) \geq f\left(A_{m*+1}\right) \geq \cdots \geq f\left(A_K\right)$. This result enables us to identify the optimal assortment efficiently. To our knowledge, these represent the first structural results associated with the **CMNL** problem.

As the first step toward solving the dynamic optimization problem, we introduce a parameter estimation technique based on maximum likelihood estimation (MLE). We develop error bounds that relate the quality of our parameter estimates (based on MLE) as a function of the number of customers. We show that if we offer our assortments to T independent customers, then for any $\epsilon > 0$, the probability that our estimates differ from their true values by more than ϵ is $O\left(e^{-\epsilon^4 T}\right)$.

Combining the results from the static optimization problem and the error bound in parameter estimation, we develop an adaptive algorithm for the setting when the mean utilities μ_i's are unknown and must be estimated from historical sales and assortment decisions. The proposed policy generates a sequence of assortments $S_1, S_2, \ldots,$ such that the T-period average expected profit $\frac{1}{T}\sum_{t=1}^{T} E\left[f\left(S_t\right)\right]$ converges to Z^* at the rate of $O\left((\log T)^2/T\right)$.

We then test our proposed policy using actual DVD sales data from an online retailer. Figure 1 shows the running average expected profit over time under our proposed adaptive policy, when averaged over 100 problem instances each with a capacity of 10. We observe that the average expected profit approaches 90% of the optimal profit Z^* after approximately two days of sales.

References

Ben-Akiva, M., and S. Lerman. 1985. *Discrete choice analysis: Theory and application to travel demand.* MIT Press.

Gallego, G., G. Iyengar, R. Phillips, and A. Dubey. 2004. Managing Flexible Products on a Network. Working Paper.

Liu, Q., and G. Van Ryzin. 2004. On the Choice-Based Linear Programming Model for Network Revenue Management. To appear in MSOM.

Mahajan, S., and G. van Ryzin. 1998. Retail inventories and consumer choice. In *Quantitative Models for Supply Chain Management*, ed. S. Tayur, R. Ganeshan, and M. Magazine, 491–551: Kluwer Academic Publishers.

McFadden, D. 1974. Conditional Logit Analysis of Qualitative Choice Behavior. In *Frontiers in Econometrics*, 105–142: Academic Press.

Megiddo, N. 1979. Combinatorial Optimization with Rational Objective Functions. *MOR* 4 (4): 414–424.

Network Design under a Discrete Choice: a Bilevel Programming Approach

GILBERT, FRANÇOIS (Université de Montréal) gilbert@iro.umontreal.ca
MARCOTTE, PATRICE (Université de Montréal) marcotte@iro.umontreal.ca
SAVARD, GILLES (École polytechnique de Montréal) gilles.savard@polymtl.ca

We consider a class of network design problem where demand is modeled according to discrete choice theory. This probabilistic approach, much like deterministic multiclass models, accounts for population heterogeneity .

The basic problem consists in setting taxes on a subset of the network links and constitutes a nonlinear extension of the problem adressed in [4]. When commuter choices are limited by finite capacities of the network links, the route choice model, derived from information theory, generalizes the Logit model [2]. The resulting revenue maximization problem is differentiable but non-concave and does not directly lend itself to global resolution tools such as the combinatorial techniques used in [1]. We will show that, by properly approximating the continuous Logit flow distribution, one can globally solve the problem through an equivalent mixed integer formulation similar to the one used in the deterministic case.

Network design problems fit the Stackelberg game framework. Two players are involved : a leader and a follower. The leader plays first and has perfect knowledge of the follower strategy, while the follower only responds to whatever action the leader has taken. In our case the leader's goal is to maximize his revenue by setting of tolls, and possibly of capacities, on some of the network links. The follower's response is described by a discrete choice model parametrised by the leader's toll policy. This scheme is most naturally formulated as a bilevel mathematical program.

More precisely let $\mathcal{G} = (\mathcal{A}, \mathcal{N})$ an oriented graph, $\mathcal{A}^m \subset \mathcal{A}$ the subset of toll links, and Q a set of origin-destination (OD) couple characterized by demand d_q for each $q \in \mathcal{Q}$. Let $\mathcal{R}^q$ the set of paths associated with OD couples q and $\mathcal{R}^q_m \subset \mathcal{R}^q$ the subset of paths associated with OD couple q that includes at least one toll link. Let us note $\mathcal{T}^q \subset \mathbb{R}^{|\mathcal{R}^q_m|}$ the set of toll vectors defined in the path space that are supported by an equivalent link space toll policy. Each path $r \in \mathcal{R}^q$ is characterized by its fix cost c^q_r and a toll t^q_r, summed along it's links. Let us also define Λ^q the arc-path incidence matrix associated with OD couple q and $\Omega^q : \mathcal{R}^q_m \to \mathcal{R}^q$ the mapping between toll path space and path space for OD couple q. The bilevel program associated with the network design problem is formulated as :

$$
\begin{aligned}
\max_{x,t} \quad & \sum_{q\in\mathcal{Q}} x^q\, \Omega^q\, t^q \\
\text{s.t} \quad & t^q \in \mathcal{T}^q \subset \mathbb{R}^{|\mathcal{R}^q_m|} && q \in \mathcal{Q} \\
& x \in \arg\min_y \quad \sum_{q\in\mathcal{Q}} (c^q + \Omega^q\, t^q) y^q + \frac{1}{\theta} y^q \log y^q \\
& \text{s.t.} \quad \sum_{r\in\mathcal{R}^q} y^q_r = d_q && q \in \mathcal{Q} \\
& \qquad \sum_{r\in\mathcal{R}^q} \Lambda^q y^q \le u && q \in \mathcal{Q} \\
& \qquad y^q \ge 0 && q \in \mathcal{Q}
\end{aligned}
$$

where the vectors x^q et y^q represent the demand allocation on paths R^q, and the logarithm is taken componentwise. The constant θ is a scale parameter that reflects the variance of the underlying statistical process.

The following scenarios are considered in increasing order of complexity :

- unbounded capacities with tolls as decision variables ;
- bounded capacities with tolls as decision variables ;
- bounded capacities with tolls and capacities as decision variables ;

Different resolution schemes are presented and tested on instances involving several hundred paths. Note that such problems cannot be tackled by traditional global optimization tolls such as interval arithmetics [3]. The following resolution methods and approximation schemes are considered :

- trust region method adapted to bilevel programs
- piecewise linear approximation of the lower level objective
- piecewise quadratic approximation of the lower level objective

In all cases we obtain an approximating mixed integer formulation that is solved to optimality.

Note :
We have adopted a path formulation for the sake of simplicity. This greatly simplifies the presentation and will alow generalization to more elaborate route choice models. Yet, in order to tackle larger instances, a link formulation is required and is available for most of our models.

Références

1. L. Brotcorne, M. Labbé, P. Marcotte and G. Savard, "A bilevel model for toll optimization on a multicommodity transportation network" Transportation Science 35 (2001) 1–14.
2. S. Erlander, "Efficiency and the logit model" Annals of Operations Research 82 (1998) 203–218.
3. C. A. Floudas, "Deterministic global optimization : theory, algorithms and applications" Kluwer, Dordrecht 1999.
4. M. Labbé, P. Marcotte and G. Savard, "A bilevel model of taxation and its application to optimal highway pricing" Management Science 44 (1998) 1595–1607.
5. P. Marcotte, S. Nguyen and A. Schoeb, "A strategic model of traffic assignment in static capacitated networks" Operations Research 52 (2004) 191–212.

Seat Value Measures for Theaters for Theaters and Stadiums

VEERARAGHAVAN, SENTHIL (University of Pennsylvania) senthilv@wharton.upenn.edu
, RAMNATH (University of Pennsylvania) ramnathv@wharton.upenn.edu

Theaters and sporting events have several characteristics that are well suited to Revenue Management (RM) methods. There are many different customer segments (e.g. corporations, season ticket holders, families, students) each with varying usage patterns and willingness to pay. Ticket prices depend on several factors such as the location of the seat, the popularity of the event, the expected demand, group affiliations of customers, seasonalities, bulk-sales terms and advance-purchase restriction (Talluri and van Ryzin 2004).

The simplest revenue management lever used by theaters and stadiums is price discrimination based on seat quality, which largely depends on the location of the seat relative to the stage or playing field. For example, theaters "scale the house" by pricing the front rows at high prices and reducing prices all the way to the nosebleed section (Courty 2000). In contrast, seat value in a given travel class is less sensitive to seat location for airline seats (as they primarily serve as a conduit for transporting a person from an origin to a destination).

While there is extensive research studying ticket pricing and selling strategies to maximize season revenues, there is little empirical research on seat value as perceived by consumers. According to Talluri and van Ryzin (2004) "*fear of negative customer reactions and consequent loss of customer goodwill are the main reasons firms seem to be avoiding bolder demand management strategies*". Hence, it is imperative to understand ex-post seat valuations in order to be able to better implement revenue management strategies.

The dependence of perceived seat value on location can be fairly sophisticated. For example, in theaters, seats in the middle of a row might be preferred over seats toward the outside of a row further forward, and seats at the front of second-level sections are sometimes preferred to seats at the back of first-level sections (Leslie 2000). While this ordering of seat value by location is subjectively understood by theaters and stadiums, there is no objective measure quantifying it.

Our research complements traditional RM literature for theater and sports businesses by providing (1) an objective measure of seat value adjusted for price (Seat Value Index), (2) a methodology to study the relationship between the ex-post seat value perceived by consumers and seat location, and (3) pricing recommendations to achieve a *"desired level of value"*.[1] We apply this methodology to a proprietary dataset collected by a professional baseball franchise in Japan, from a survey of its patrons during a weeknight game.

We capture the ex-post (latent) value realized by patrons attending an event through a reported ordinal measure (with J levels) called Seat Value Index (SVI). We assume that a respondent i derives her SVI by categorizing her ex-post (latent) valuation, V_i^*, into buckets defined by the boundaries (thresholds) $\{\tau_i^0, \tau_i^1, ..., \tau_i^J\}$. Mathematically, $SVI_i = j$, if and only if $\tau_i^{j-1} < V_i^* \leq \tau_i^j$, where seat value is increasing in j. The model specification is completed by assuming that $V_i^* = \mathbf{x}_i^T\beta + \epsilon_i$, where $\mathbf{x}_i$ is a vector of consumer and seat characteristics (excluding a constant), β is the associated vector of parameters, and ϵ_i is an idiosyncratic experience term. This leads us to the ordinal regression model. In addition to the standard model, we also incorporate differences in reporting thresholds and heterogeneity in the distribution of valuations, to obtain unbiased parameter estimates. The detailed analysis is available in the full version of the paper. We report some key insights obtained from applying this methodology to the baseball dataset:

1. **Experience Goods:** Repeated visits to the ballpark reduce uncertainty in valuations and help consumers "learn" the "true value" of seats. This is consistent with the notion that theater and stadium events are experience goods with residual uncertainty. In particular, we find that a customer who has visited the ball park five times has 29% lower variance in his realized valuation as compared to a first timer.

2. **Seats are Asymmetric:** Seat Value Indices reported by consumers seated in the ball park are *asymmetric*. Patrons seated on the third base side (left field) are less likely to have extreme valuations, as compared to those seated on the first base side (right field). This asymmetry is intriguing and appears counter-intuitive. However, there are several possible explanations including higher incidence of foul balls/home runs, and the location of the home-team dugout on the left field side.

3. **Valuable Seat Locations:** Customers seated at the upper deck enjoy higher mean net-valuations and also have more heterogeneity in realized valuations. This has two possible implications that are important from the franchise's perspective. First, consumers might be responding positively to the availability of lower priced upper deck seats. Second, even though the upper deck seats are located far off, they could still be providing comparatively higher ex-post valuations to consumers who sat there.

Our framework and methodology can be easily employed in any theater/stadium where seat value is influenced by seat location characteristics and customer attributes. We also provide specific pricing recommendations to achieve a "service-level" objective in the full version of the paper.

[1] This notion is analogous to "fill-rate"/"service-level" measures employed in operations management. While focusing on a desired fill-rate/service-level might be sub-optimal for short-run profit maximization, it improves availability, leading to long-run benefits. Quantity adjustments are more difficult in theaters/stadiums, but price adjustments to "satisfice" value can be made.

Revenue Management by Sequential Screening

AKAN, MUSTAFA (Northwestern University) m-akan@kellogg.northwestern.edu
ATA, BARIS (Northwestern University) b-ata@kellogg.northwestern.edu
DANA, JAMES (Northeastern University & HBS) j.dana@neu.edu

We consider a dynamic model of revenue management with strategic, i.e. forward looking, consumers. The consumers are heterogeneous in their valuations or willingness to pay. Each consumer knows his type, which is defined as the distribution of his valuation, at time zero, but not his true valuation. Rather, they learn their true valuations sequentially over time. Consumers with higher valuations, e.g. business travelers, learn their true valuations after the ones with lower valuations, e.g. leisure travelers. In this setting, a monopolist system manager strives to maximize her profits by sequentially screening the consumers. Using a mechanism design approach, we show that the optimal mechanism is a menu of expiring refund contracts. We also identify the conditions under which the system manager can achieve the first-best solution, thereby extracting the entire (expected) consumer surplus. Under the optimal mechanism, contracting takes place after the consumers learn their types but before they learn their true valuations. Moreover, the monopolist finds it optimal to ration different types of consumers to various degrees. Finally, we discuss how the results change under different assumptions on the dynamics of learning.

Room: *Agora*

10:30 *Coffee Break*

TRACK VI

Session VI.1: STATISTICS I (11:00-12:30)

Room: *S1-111* **Chairperson:** *Garrett van Ryzin*

When to Buy a Ticket

LEE, MISUK (Georgia Institute of Technology)
CASTILLO, MARCO (Georgia Institute of Technology)
GARROW, LAURIE (Georgia Institute of Technology) Laurie.Garrow@ce.gatech.edu

A common practice in the revenue management literature is to analyze what the optimal pricing for alternative demand schedules. This paper investigate the opposite question: what is the optimal search and purchasing behavior of strategic buyers given price evolution and what does this imply for optimal pricing? This question is important not only because it adds realism to the analysis but because consumers' preferences can have surprising effects on what the optimal pricing policy is. This is a first step towards a model where demand and supply are jointly determined.

Our approach considers that buyers are strategic and choose among available alternatives in each period but also consider the potential opportunities available to them in the future. The basic decision problem for the buyer is then to decide how intensely search and what to do given the available options cognizant that other (better) opportunities might appear in the future. Our model combines two strands of literature: random utility models and optimal search models. Random utility models are useful because they allow us to model taste heterogeneity and market uncertainty. Search models allow us to determine not only when to buy but how intensely search for alternatives. Importantly, our models are flexible enough as to allow us to analyze the effect of cognitive biases on search and purchasing behavior.

We show that consumers follow a simple cut-off rule that depends on future prices and own optimal behavior. Buyers' likelihood of purchasing increases as the deadline approaches and expected future prices increase. Consumers are shown to delay purchases as more alternatives are made available and with mean preserving spreads in prices. Importantly, this implies that demand cannot be assumed independent of future prices and therefore optimal pricing cannot be reactive to demand. For consumers, goods available today and tomorrow are substitutes. Because our model allows for randomness in tastes or information, our model allows for different degrees of substitution. That is, buyers are thought as uncertain of their own future preferences making the same good sold today or tomorrow imperfect substitutes.

The fact that buyers are strategic has consequences for the optimal pricing of a monopolist. We show that with imperfect temporal substitution the optimal policy for a monopolist is mark-down pricing. This holds for a large class of models allowing for preferences heterogeneity. Intriguingly, we also show that mark-up pricing can be the optimal strategy if consumers time preferences are inconsistent and unsophisticated (Strotz, 1956; Phelps and Pollak, 1968). That is, we show that the strategic behavior and distribution of preferences of buyers can have large consequences. Consumers that procrastinate might end up buying at higher prices than optimal. Prices can increase even without capacity constraints. We consider that this result is important empirically. Indeed, there is growing evidence that people posses time inconsistent preferences. This might imply that buyers facing search problems with a deadline are willing to buy services or products that reduce the potential costs of suboptimal behavior.

While our model has several applications, we concentrate on the case of airline industry. The growth of search engines has increased the level of competition among carriers and also reduced the search costs of individuals. The fact that customers have to consider an enlarged and changing choice set justify our modeling assumption that customers might be uncertain about their tastes on the future. More importantly, the growth of internet sites giving recommendations about *when to buy a ticket* makes the analysis of optimal search and purchase behavior most pressing.

Estimating Primary Demand for Substitutable Products from Sales Transaction Data

VAN RYZIN, GARRETT (Columbia University) gjv1@columbia.edu
VULCANO, GUSTAVO (NYU Stern School) gvulcano@stern.nyu.edu
RATLIFF, RICHARD (Sabre Holdings) Richard.Ratli®@sabre-holdings.com

Two important problems in retail demand forecasting are: 1) estimating turned away demand when items are sold out, and 2) properly accounting for substitution effects among related items. For simplicity, most retail demand forecasting approaches in practice rely on time-series models of observed sales data, which treat each separate stock keeping unit (SKU) as independent. However, if demand turned away when a customer's first choice is unavailable (referred to as *spilled* demand) is ignored, the resulting demand forecasts are downwardly biased; this underestimation can be severe if products are unavailable for long periods of time. Additionally, if some products are unavailable it will increase observed sales in substitute products which are available (referred to as *recapture*); ignoring recapture in demand forecasting will lead to an overestimation bias among the set of available SKU's. Correcting for both spill and recapture effects is important in order to establish a good estimate of the underlying primary, or first-choice, demand for products. This paper provides a simple and computationally efficient approach for estimating primary demand using only observed historical sales data.

While spilled demand is not observable from sales transactions, various statistical techniques have been proposed to estimate this information using data on sales during periods when the product is available. Collectively these techniques are known as *demand untruncation* or *uncensoring* methods (the terms are equivalent). One of the most popular (and accurate) methods of untruncation is the expectation maximization (EM) algorithm. EM procedures ordinarily employ iterative methods to estimate the underlying parameters of interest (in this case, demand by SKU across a set of historical data). The EM method works using alternating steps of computing conditional expected values of the parameter estimates to obtain an expected log-likelihood function (the "E-step") and maximizing this function to obtain improved estimates (the "M-step"). Traditionally, retail forecasts that employ the EM approach have been limited to untruncating sales history for individual SKU's and disregard recapture effects from substitute products.

To provide greater variety for customers, retailers typically provide assortments of goods comprised of numerous different brands of an item or items with slightly different features (e.g., grocery stores provide an entire aisle dedicated to different breakfast cereals). Classical economic theory involving substitution effects provides techniques for estimating demand shifts due to changes in prices of alternative offerings.

However, an important practical problem is how to fit such demand models based on readily available data, which in most retail settings consists of observed sales, prices and on-hand inventory quantities by SKU. In our work here, we assume these are the only data available. A convenient and widely used approach for estimating the demand of different SKU's within a set of similar items is to use customer choice models, such as the multinomial logit (MNL). Customer choice models predict the likelihood of customers purchasing a specific product from a set of related, available products based on their relative attractiveness. A convenient aspect of these models is that the likelihood of purchase can be readily recalculated if the mix of available related products changes (e.g., due to another item being sold out or restocked). By assigning preference weights across each of the available products, customer choice models are useful in correcting for recapture effects in observed historical sales.

We propose a novel method of integrating customer choice models with the EM method to untruncate demand and correct for recapture effects across an entire set of related products sales history. The key idea is to view the problem in terms of primary demand (or first-choice) demand and to treat the observed sales as incomplete observations of primary demand. We then apply the expectation-maximization (EM) method to this incomplete, primary demand model and show that it leads to a simple, highly efficient iterative procedure for estimating the model which provably converges to a stationary point of the log-likelihood function. The approach is practical and effective, as illustrated on several industry data sets.

Session VI.2: RETAIL (11:00-12:30)

Room: *S1-131* **Chairperson:** *JongWook Lim*

Multiple Supplier Competition under Consignment Inventory Programs with Competing Retailers

HU, MING (Columbia University) mh2252@columbia.edu

We investigate the equilibrium behavior of Stackelberg games where two-echelon supply chains with a set of competing suppliers serves a network of competing retailers. We assume the demand system is governed by a linear structure. This seems to be the first paper to consider both the supplier and retailer competition in a single model.

In the Stackelberg game with deterministic demand, the sequence of actions is assumed as follows: (1) the suppliers set wholesale prices simultaneously; (2) the retailers set retail prices simultaneously given the wholesale pricing information from the first-stage available. We provide sufficient conditions for the existence and uniqueness of Nash equilibrium to this Stackelberg game.

Then we take demand uncertainty into account by considering a Stackelberg game of suppliers as leaders and retailers as followers under consignment inventory program. The sequence of events is summarized in Figure 1: (1) Competing suppliers simultaneously announce wholesale prices and decides how much to procure/manufacture in advance for various retailers; (2) Competing retailers simultaneously announce retail prices; (3) Consumer market is realized and unsold consignment inventories are returned and salvaged by suppliers. We provide sufficient conditions for the existence of Nash equilibrium to this Stackelberg game.

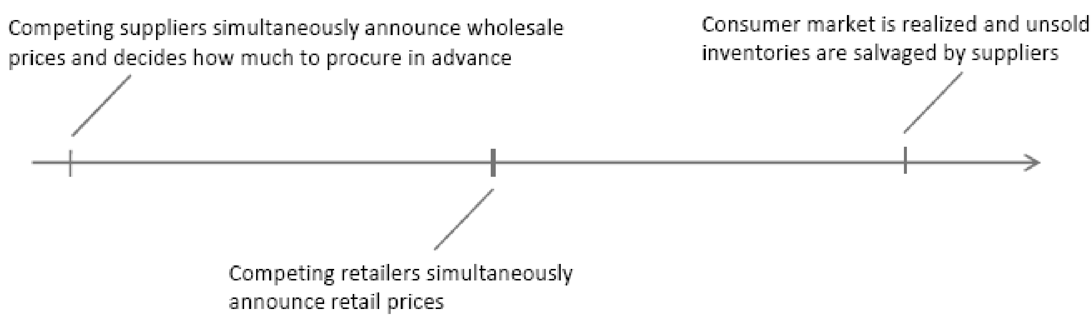

Figure 1: Stochastic Model Timeline

The Calibration of Price Response Function and the Profit Maximization Price Point for Cosmetic Market in China

LIM, JONGWOOK (Samsung OpenTide China) frontier@opentide.com.cn
LEE, YEISUN (Samsung OpenTide China) wyeth@opentide.com.cn
KWAK, YOUNGSIK (Jinju National University) pricing@benet.co.kr

Although the searching for the optimal price point for a given product is the one of the most interesting issues for marketers, the attempt to identify the profit maximization price point for a given product has been relatively rare. To estimate the optimal price point, we need the price response function and the cost function for the product or service. To measure the disaggregate price response curve in the Chinese cosmetic market, we used the conjoint analysis-based procedure including ACA (adaptive conjoint analysis) and CBC (choice-based conjoint analysis).

The curve enables us to estimate the individual choice probabilities corresponding to applied price points for BRAND-A, the one of market followers in Chinese cosmetic industry. Furthermore, we aggregate the individual price response curve into the market price response function for BRAND-A.

After we add the cost curve for each sales volume for each price point, we can estimate the profit-maximization price point for BRAND-A.

We apply mixture model to the part-worth for each observation and identify two segments. One segment is located at economic zone with domestic brands like BRAND-A and BRAND-B. The other is positioned at premium market with BRAND-C and BRAND-D. With results on hands, this multinational company plans to launch two brands for each segment with two optimal price levels for each segment. This case shows that mixture model is a useful segmentation tool to identify hidden segments and the conjoint analysis is the one of the most effective tool to calibrate for the optimal price point.

Session VI.3: CUSTOMER CHOICE III (11:00-12:30)

Room: *S1-139* **Chairperson:** *Kinshuk Jerath*

Optimal Design of a Name-Your-Own-Price Channel when Customers Behave Strategically

WILSON, JOHN (Ivey School of Business) jwilson@ivey.ca
ZHANG, GUOREN (Ivey School of Business)

A retailer places a certain product for sale on the Internet. Customers are invited to "name their own price" for the product. The retailer will accept a given bid x with probability equal to $p(x)$. It is assumed that customers know the function $p(.)$ and will place bids that maximize their individual expected profits. Knowing that customers will behave this way, the retailer wants to choose the function $p(.)$ that maximizes the retailer's expected profit. We demonstrate that there is an explicit ε-optimal solution to this problem.

For purposes of simplicity, we consider one product: an economy rental car on a specific date. A customer is allowed to bid only once. If the bid is too low, Priceline responds with specific information on what is a "good" bid and what is a "bad" bid. At this stage-before the bid is processed-the customer is asked to consider revising the bid. Once the bid is submitted, it is either accepted or rejected by Priceline and no further bidding is allowed. There are various websites where a customer can see the bidding histories of other customers. This paper considers the situation where Priceline would forestall this information gathering by customers and simply release a function that will let customers know the chances that a given bid will be accepted. We will stick to this simple rental car case and assume that the retailer has more than enough inventory to satisfy all possible demand. As mentioned in Terwiesch et al. (2005), this situation reflects the reality of many cases where rental car companies, for instance, allow their product to be discounted on Priceline. In the rental car case, a customer can easily access rental car prices from the various companies or find the cheapest rate available on one of the well known sites such as Expedia.com or Travelocity.com. We assume that this price is H and that a customer will not bid more than this on the Priceline site. We assume that Priceline has market research available to it concerning the behavior of its customers. This assumption is similar to ones made in Ding et al. (2005), Vulcano (2002) and Terwiesch et al. (2005). In particular, we assume that the reserve price of a randomly selected customer is described by the density function $f(.)$ with support on closed interval $[0,H]$. Each customer is assumed to make a bid that maximizes his or her expected return. The retailer, knowing that the customer will act strategically by finding H or finding $p(.)$ (if it is not provided) wants to select a function $p(.)$ that will maximize the retailer's expected profit. We will show, for any given $\varepsilon>0$, ε-optimal functions can be found. Each possible bidder has a maximum price (called the reserve price) that he or she is willing to pay for the item. Assume that $f(.)$ represents the density function for the distribution of such prices. Since the bidder can always obtain the item by paying full price on the market, it is assumed that the density $f(.)$ is distributed over the range $[0,H]$, where H can be interpreted as the full retail price of the item. Given that, in its most general form, this is a calculus of variations problem, it is somewhat surprising that explicit analytical solutions for ε-optimal functions can be found. The case of the retailer allowing multiple bids will also be considered.

Revenue Management with Strategic Customers: Last-Minute Selling and Opaque Selling

JERATH, KINSHUK (University of Pennsylvania) kinshuk @wharton.upenn.edu
NETESSINE, SERGUEI (University of Pennsylvania) netessine @wharton.upenn.edu
VEERARAGHAVAN, SENTHIL (University of Pennsylvania) senthilv@wharton.upenn.edu

Companies in a variety of industries (e.g., airlines, hotels, theatres) often use last-minute sales to dispose of unsold capacity. The consequences of such a strategy are not immediately obvious: more discounted last-minute tickets lead to more consumers anticipating the discount and delaying the purchase rather than buying at the regular (higher) prices, and hence potentially reducing revenues for the company. To mitigate such behavior, many service providers have turned to opaque intermediaries such as hotwire.com that hide many descriptive attributes of the service (e.g., departure times for airline tickets) so that the buyer cannot fully predict the ultimate service provider. Using a stylized economic model, this paper attempts to explain and compare the benefits of last-minute sales directly to consumers vs. through an opaque intermediary.

We utilize the notion of rational expectations equilibrium to model consumer purchasing decisions: consumers make early purchase decisions based on expectations regarding future availability, and these expectations are correct in equilibrium. We show that direct last-minute sales are preferred over selling through an opaque intermediary when consumer valuations for travel are high and/or there is little service differentiation between competing service providers; otherwise, opaque selling dominates. Moreover, contrary to the usual belief that such sales are purely mechanisms for disposal of unused capacity, we show that opaque selling becomes more and more preferred over direct last-minute selling as the probability of having high demand increases. At the extreme, firms will never use direct last-minute sales when there is no demand uncertainty but they may still employ the opaque sales channel.

Markov-Based Models of Airline Passengers' Searching and Purchasing Behavior

LEE, MISUK (Georgia Institute of Technology)
GARROW, LAURIE (Georgia Institute of Technology) Laurie.Garrow@ce.gatech.edu
POST, DAVID (Sigma Zen)

This study analyzes how individuals search for information about air ticket attributes and prices. The data for this study is based on actual data from an airline website that employs a unique strategy for selling tickets at low prices to travelers with a high degree of travel flexibility (e.g., the individual can travel on different days, does not require a lot of advance notice of when the trip is schedule prior to departing, etc.) As part of this study, the following questions are explored:

- When searching for information, are individuals likely to change their departure time windows, length of stay, advance notice windows?
- Do individuals search along multiple parameters simultaneously, or linearly?
- Do individuals start with an initially "large" search space, or start small and then expand?
- Are there any "screen dependencies," in the sense that what individuals have seen two or three screens earlier influences where customers are going to search next?

Markov-based methods are used to explore these questions. The state-space is decomposed into three primary components to capture reneging behavior (defined as situations in which customers leave the screen without purchasing or engaging in search), purchasing behavior (defined as situations in which customers purchase directly from the first screen) and searching behavior (defined as screen-to-screen searches that may ultimately lead to purchase or no purchase events). Empirical results show that search intensity is highly dependent on the relative discount level, i.e., customers seriously begin to engage in search when discount levels reach 30%. Purchase probabilities increase as the discount level increases, particularly above the 30% level. Most significant, there is a high degree of correlation (0.70) observed across screens, that is, there is a "stickiness factor" in the sense that customers are seen to fixate on specific attributes of the search process (e.g., are inflexible on the amount of time they want to spend at the destination). Failure to account for this high degree of correlation across screens leads to high forecast errors.

Currently, we are working to integrate the behavioral findings into a framework that can support firm pricing policies. This framework leverages the fact that conversion probabilities, or the probabilities of purchasing, are driven both by price (reflected in the discount level), in addition to well-defined attributes characterizing the search space (e.g., how far in advance the individual wants to be notified of their flight departure dates and times, the number of days on which the individual is willing to depart, etc.). Essentially, the underlying optimization algorithm seeks to maximize profits and/or conversion probabilities by enabling the airline firm to influence demand via discount levels that are set via three searching parameters: departure time windows, length of stay windows, and advance notice windows.

It is envisioned that the results of this study can be used by major carriers as a method for selling excess inventory via opaque sales channels that bypasses third-party intermediaries. In addition, it is envisioned that the results of this study can be used to support the development of additional opaque products for flexible travelers (one such case is currently occurring in Europe).

Room: *Agora*

12:30 *Lunch and address from the Board*

TRACK VII

Session VII.1: STATISTICS II (14:00-15:30)

Room: *S1-111* **Chairperson:** *Omar Besbes*

Modeling the Choice of an Airline Itinerary and Fare Product Using Booking and Seat Availability Data

CARRIER, EMMANUEL (Massachusetts Institute of Technology) ecarrier@mit.edu

Over the last ten years, the rapid growth of low-cost airlines and the development of web-based distribution of airline tickets have transformed the competitive environment in the airline industry worldwide. The relaxation of fares rules by low-cost airlines has disrupted the pricing and revenue management models of large network airlines and a better understanding of passenger choice behavior is required to support the development of new strategies to compete more effectively in the current marketplace. However, previous studies of airline passenger choice have not adequately represented the impact of pricing and revenue management or the heterogeneity of behaviour across different categories of airline travelers. This research develops a modeling framework that better reflects these major characteristics of airline markets and their impact on airline passenger choice.

In order to analyze the joint choice of an airline itinerary and fare product, a unique dataset was developed. Booking data was combined with seat availability data collected daily over a 90-day booking period to determine which fare products were available at the time of the booking. Fare rules such as advance purchase and Saturday night stay requirements were also applied to incorporate the impact of airline pricing on the choice set of each traveler. Characteristics of the trip and the traveler such as the distribution channel of the ticket or frequent flyer membership were retrieved from the bookings records and used to segment the demand. They were included as explanatory variables of a latent class choice model in which several factors can be used simultaneously to segment the market without dividing the bookings into many small sub-segments. In addition, a new formulation of a continuous function of time was proposed to model the time-of-day preferences of airline travelers in short-haul markets. Instead of being set to a full 24 hours, the duration of the daily cycle was estimated to account for the low attractiveness of some periods of the day such as nighttime.

Estimation results over a sample of 2000 bookings from three European shorthaul markets show that the latent class structure of the model and the use of a continuous function of time lead to a significant improvement in the fit of the model compared to previous specifications based on a deterministic segmentation of the demand or time-period dummies. In addition, the latent class choice model provides a more intuitive segmentation of the market between a core of time-sensitive business travelers and a mixed class of price conscious business and leisure travelers.

This research extends the scope of potential applications of passenger choice models to additional airline planning decisions such as pricing and revenue management. In particular, parameter estimates of the model were applied to forecast the sell-up behavior of airline passengers, a major input required by the newly proposed revenue management models designed to maximize revenues under less restricted fare structures.

The Impact of the Internet on Airline Fares

BRUNGER, WILLIAM (Case Western Reserve University) bbrung@coair.com

Customers who use Internet-based online travel agencies (OTAs) to purchase “Clearly Leisure” trips (defined as trips that were booked at least 14-days in advance, where the customer stayed over the weekend) pay significantly less (more than 10% less in our sample) for similar itineraries in the same markets than those who purchase through traditional travel agencies. This difference persists *even though* the fares and inventory offered by the airline through each channel are identical because of contractual obligations and airline marketing strategies. The purpose of this study is to examine this Internet Price Effect (IPE). Specifically, this study uses detailed passenger-level trip data (previously unreported in the literature) from a database maintained by Continental Airlines to examine differences in fares paid across ticket distribution channels, focusing on OTAs compared to traditional travel agencies, in order to attempt to quantify the incremental impact of Internet search on the fares customers pay.

The study is operationalized as a multivariate regression in which the dependent variable is the Fare Paid (transformed to its natural log). The independent variable of interest is the distribution channel through which the ticket was purchased. The study controls for trip characteristics (advance purchase, length of stay and group size), customer differences (frequent flyer status and trip origin in an airline hub) and market-effect dummy variables. The study introduces several explanatory variables new to the literature, including the Bid Price (or opportunity cost) of the seat, which explains as much as a third of the variation in Fare Paid. Control variables take into account precedents from previous studies in the literature. The study finds that, after controlling for this set of variables, the IPE persists at a negative 5-8% level across a wide range of timeframes and market sizes. By reformulating the dependent variable to reflect surplus revenue contribution, the study further demonstrates that, even though OTA customers pay less for their seats, they are almost as profitable as customers who use traditional travel agencies, and in some time frames, are more profitable.

Testing the Validity of a Demand Model: An Operations Perspective

ZEEVI, ASSAF (Columbia University) assaf@gsb.columbia.edu
PHILLIPS, ROBERT (Nomis Solutions) robert.phillips@nomissolutions.com
BESBES, OMAR (University of Pennsylvania) obesbes@wharton.upenn.edu

The statistics and econometrics literature have developed powerful methods for testing the validity (specification) of a model. Unlike statisticians, managers are typically more interested in the performance of their decisions rather than the validity of the model from which they are derived. Focusing on the problem of demand model specification in the context of revenue management, we propose a framework and a statistical test that captures this perspective. Theoretical properties of the test are established and its efficacy is illustrated both on synthetic examples, as well as on an empirical data set in the realm of financial services. It is shown that traditional model-based goodness-of-fit tests may consistently reject simple parametric models of consumer response (e.g., the ubiquitous logit model), while at the same time these models may "pass" the proposed performance-based test. Such situations arise when pricing decisions induced by the assumed model structure, achieve a performance (profits) that cannot be distinguished statistically from the true optimal performance (i.e., when one knows the actual underlying response function that governs realized demand).

Session VII.2: APPLICATIONS (14:00-15:30)

Room: *S1-131* **Chairperson:** *Weifen Zhuang*

Optimal Column Assignment for a Vending Machine

FUKUCHI, JUNICHIRO (Gakushuin University) junichiro.fukuchi@gakushuin.ac.jp
ITOH, HAJIME (Otaru University of Commerce)
TAKEUCHI, TOSHIKO (Gakushuin University)

We investigate a method of finding the optimal assignment of items to columns in a vending machine. There are spaces called columns in a vending machine. A column in a vending machine is a space where stock of an item are held. We formulate a mathematical programming problem for the optimal column assignment under the stochastic demand and investigate properties of a heuristic method of finding the optimal solution.

We assume that the demand processes for items are independent Poisson processes. The objective function in our formulation is the long-run profit rate under joint replenishment policy. Since replenishment cost for a vending machine is not low, some type of joint replenishment policy is used in practice. We assume that the inventory is continuously reviewed and the lead time for a replenishment is zero. We consider the following joint replenishment policy.

1. Stock is replenished at the instance that some item is sold out.
2. All the items are replenished to the initial level when an item is replenished.

The problem to be solved is a nonlinear integer programming problem with many local maxima. We use the Life Span Method, a variant of the tabu search, to obtain a near optimal solution. The results of numerical experiments will be given.

Asymmetric Pricing to Increase MMS Revenue

SAMANTA, SWADESH KUMAR (University of Essex) ssksam@essex.ac.uk
WOODS, J. (University of Essex) woodjt@essex.ac.uk
GHANBARI, M. (University of Essex) ghan@essex.ac.uk

Research into the impact of pricing on revenue has traditionally focused on the airline, hotel and rental car industries. The basic characteristics of these industries are that they have perishable inventory, relatively fixed capacity, predictable demand and high fixed costs. Networked industries such as telecommunications have similar properties but have tended to attract less research attention (Chiang et. al. 2007). Telecommunication services are key factors for the growth and development of other services. However in order to make the services popular the prices have to be acceptable to the user. This research identifies an innovative pricing strategy for the Multimedia Messaging Service (MMS) which increases the revenue to the operator whilst simultaneously benefitting the consumers.

Multimedia Messaging Service (MMS) take-up by consumers is far less than the projection when it was initially deployed (Forum Nokia, 2002). There are many reasons such as interoperability, handset compatibility, coverage, lack of content, low penetration, poor usability, time taken to compose a message, high price, and high handset cost (Pohjola and Kilkki, 2005; Portio Research, 2007). Pohjola and Kilkki (2005) showed that price and the time taken to compose an MMS are more important compared to other factors such as usability, coverage, and message dropping. The interoperability problems are now minimised due to the signing of an agreement between the operators (GSM Association, 2007; Softmedia, 2008). The cost of the handset is also substantially reduced due to technological advances and economies of scale in mass production. As an example the current (as of March 2008) average cost of an MMS capable handset (with camera) is 150 to 200 euros compared to 500 to 700 euros in 2002 (3G Press Release, 2002). Moreover the advancement of media adaptation (transcoding) has minimised interoperability and handset compatibility problems (Coulombe and Grassel, 2004). With the introduction of automatic handset configuration over the air, users are free from the burden of manual configuration (Comverse, 2007; Krol, 2007). As these issues are addressed, the service becomes more popular, but it is not seen to be increasing at the expected rate. It is our conjecture that there are other issues which researchers have not yet addressed. Although most published research indicates that high price is the major factor behind the low MMS growth (Informa, 2006), few have considered factors such as the cost incurred and the value gained (utility) by both the user and operator.

Based on a market survey in the Taiwan telecom market, Hsu et. al., (2007) showed that user's utility has a major bearing on MMS adoption. They found that impacts from other factors such as usability are dependent on whether the users are early or late adaptors to the service. The utility to a user can be quantified as a combination of the expenditure incurred to send an MMS and the value gained from it. The expenditure incurred by a sender is not only dependent on the price paid to the operator, but the additional cost to buy an MMS capable handset and the notional cost of the time spent to compose and send an MMS. A (rational) user will procure an MMS handset only when the additional value gained from sending an MMS exceeds the sum of the extra costs. The value gained from MMS depends on whether it is considered as a message or entertainment service. An MMS can consist of text, audio and video and therefore some users may consider it as a messaging service whereas others may focus on its entertainment value (Reiter Ilan, 2003). Therefore it is natural for a user to compare the cost of an MMS with other modes of communications such as SMS (Short Message Service) or a voice call before taking a decision whether to send or not.

As of April 2007, the MMS usage volume in China was 0.3 MMS messages per mobile user per month compared to 1.2 in the USA (China Mobile, 2007; AT&T, 2007). We also observed similar usage in countries where mobile penetration has gone beyond 100% (The Register, 2004; Dudley, 2007). As an example, the mobile penetration in Finland reached over 100% in 2004 but MMS usage (0.12 MMS) per user per month was much less compared to 37 SMS messages (Bouwman et. al., 2007). As of January 2008, the mobile penetration in the UK was more than 115% but the MMS usage was 0.6 message per user per month which is far less compared to 90 SMS (Mobile Data Association, 2008a). We find that more than 42% of mobile population in Japan use MMS compared to 25% in the UK and 12% in the USA (Dudley, 2007; Mobile Data Association, 2008b; Mobile Marketing Association, 2008).

Low MMS subscription and usage is a direct result of the current pricing mechanisms. The pricing strategy currently used by major European and American operators does not discriminate between the users with subsidized handsets or those who purchased their own. Users who buy their own handsets should reap the benefit of a reduced price for the service. This would not only increase the penetration of MMS capable handsets but also the use of the service. The main obstruction to the growth of MMS is the adoption of a similar pricing strategy to that of SMS (Wireless Profit, 2004). The operator generally charges at the moment the user sends the message, irrespective of whether it has been delivered intact or not delivered at all. The solution proposed in this work is to break the charging into two parts[1]: sending and receiving.

We developed a pricing model, based on the cost to provision an MMS and utility gained from it. The following factors are considered: 1) cost incurred by the operator to provision, 2) cost incurred by the user, and 3) utility gained by both the operator and user. No other research has considered these factors and their novel evaluation. Based on real data from BSNL, Indian telecom operator we show that the operator's cost to provision an SMS, an MMS (of size 30KB) and a per-minute-call has a ratio of 1:5:10. The cost incurred by a sender is not only dependent on the price paid to the operator, but the additional cost to buy an MMS capable handset and the notional cost on the time spent to compose and send an MMS. While most published research (Informa, 2006) compares the price of an MMS with that of an SMS, we showed that the comparison should be done both on an SMS and a per-minute-call. We compared the price (as of April 2007) for six countries; Japan, UK, USA, China, India and Brazil and we find that the average MMS price has a close relationship with the (operator's provisioning) cost if compared with SMS. However the average price has no relationship if compared with a per-minute-call. We show that the component cost to the operator to provision an MMS has more similarity with a per-minute-call than that of an SMS. The cost of a 30KB MMS is half that of a per-minute-call, but the price varies from 1/3 of a per-minute-call in Japan to 5 times in India, with the average being 1.5 for the six countries. These observations provide the motivation for the formulation of an ideal pricing mechanism which can be implemented in any country.

If we consider the provisioning cost, then the ideal price of an MMS should have two parts; a fixed part for the connection and a variable part for the transport (according to the size of the MMS message). We assume that a sender will derive less utility compared to the recipient of an MMS and therefore we propose a lower price to send and higher price to receive. We assume that the value gained is dependent on the number of MMS sent and received and the incremental gain decreases with an increase in total number of messages. Therefore we propose that total price paid by a user should be a concave function of the number of messages sent and received. This is in contrast to the current implementation where operator charges fixed price per message irrespective of the number of messages. The demand for MMS subscription is modeled based on cost of the handset and the price to send an MMS considering other factors such as the income, cultural background and educational level of the individual as same. We modeled the volume of MMS usage based on cost of handset and price to send and receive. The model compares the utility gained by both the operators and users against different combinations of handset cost and price.

Using a General Error Regression (Book and Young, 1997) and real billing data from BSNL we derived the subscription elasticity on the cost of the handset. We show that subscription can be increased by 200%, if additional cost to procure the handset is reduced by 50% (by providing handset subsidy) which conforms to the observation made by Dewenter et al. (2007). Using the model we show how a trade-off between price to send and receive an MMS (person-to-person) can increase revenue to the operator. We show that by splitting the charging into two parts the operator can increase net revenue from the MMS service by more than 50% compared to one part charging. As this is based only on person-to-person messages, an operator's revenue is likely to be significantly higher if messages delivered from applications such as mobile advertisements are considered. Although the validity of the model for mobile subscription and call usage is demonstrated by the authors (Samanta et. al., in press), we acknowledge that with more data it can be further tested for MMS. We do not have sufficient data to

[1] In a few countries (e.g. USA) both the sender and receiver pay for all services (voice, SMS and MMS), but both parties pay the same amount.

perform an extensive test and if we did, it would be difficult to compare with other published work as little research has been done in this area (Chiang, 2007).

Optimal Capacity Control for Multiple Medical Diagnostic Facilities

LI, MICHAEL Z.F. (Nanyang Technological University) zfli@ntu.edu.sg & u050023@ntu.edu.sg

Capacity management for expensive medical diagnostic facilities is critically important in healthcare industries. In this paper, we examine multiple-facility problem with three customer segments by studying the underlying stochastic dynamic control problem. Given the arrival processes for inpatients and emergency patients and any outpatient schedule, we have proved that the optimal rationing operator preserves the multimodularity of the value function, which leads to optimal control policies that can be fully characterized by monotone switching curves. We propose two applicable and numerically solvable scheduling for outpatient– vertical scheduling and horizontal scheduling and perform numerical experiments on the two scheduling under different cost and revenue structure. Through comparing the differences of the optimal profit under vertical scheduling and horizontal scheduling, we evaluate their relative performances and gain managerial insights.

Session VII.3: INTERNET (14:00-15:30)

Room: *S1-139* **Chairperson:** *Chris Anderson*

Service Differentiation in Grid Computing: Analysis of a Shared-Resource System

YAHAV, INBAL (University of Maryland) iyahav @rhsmith.umd.edu
KARAESMEN, ITIR (University of Maryland) ikaraes@rhsmith.umd.edu
RASCHID, LOUIQA (University of Maryland) louiqa@umiacs.umd.edu

Until recently, when a company needed more computing capacity or faster processors, it would buy more powerful, expensive computational resources. The emergence of on-demand computing, wrapped up in a computational grid architecture, opens a world of new computing opportunities with reduced costs and processing time. In essence, the computational grid enables coordinating and sharing of distributed resources in an on-demand fashion. The allocation of resources is transparent to customers, who basically experience a virtual super-computer.

Currently, majority of the grid computing providers such as IBM, HP, Sun Microsystems, Amazon Elastic Compute Cloud (EC2) offer services on a long-term basis. The services usually charge a flat, usage-based fee or use 'quantity discounts' where the fee is decreasing in usage (see also the survey in Baron, 2003). For example, EC2 (Amazon EC2, 2007) offers a virtual computing environment in return to a flat fee. Amazon offers the service to registered customers only, and the service is guaranteed on-demand: Amazon limits the number of customers by the server capacity, which enables Amazon to guarantee a minimum level of service. However, Amazon's cap on the number of registered users is regardless of potential customers' attributes and/or preferences. Such pricing and admissions policies are clearly suboptimal given the differences in customers' frequency of use, willingness-to-pay, and delay sensitivity.

In this paper, we show how service providers can benefit from service differentiation by segmenting the market on multiple attributes including service guarantees, price, and delay penalties. We analyze a system with two different service agreements which we call constant use (CU) and spot market (SM). The CU agreement is a proxy for the current system where all customers are guaranteed the service at a flat, usage-based fee. Any delays are subject to penalties in CU. The fee and delay penalties are non-negotiated and pre-determined. In the SM, the customers are allowed to 'name-their-own-fee' but the service provider can accept/reject job requests without any penalties. We focus on admissions control decisions for the SM customers. We first analyze a Markovian model where SM customers need immediate processing (i.e., no slack is allowed). We show that threshold admission policies are optimal in this case: A SM customer is accepted when his/her price is above a threshold level. This threshold is a function of the workload of CU customers in the system. For the general model, we study both scheduling and admissions decisions. We show earliest deadline-first is the optimal scheduling rule for each type of customer in isolation. Given the complexity, we develop approximations for the joint scheduling and

admissions control problem. We use computational experiments and study the benefit of service differentiation by quantifying the additional profits generated by introducing the service agreement SM. The paper ends with a discussion of how the proposed system can be used to provide information on probability of acceptance of a job to SM customers; such information influences future bids of the customers.

Revenue Management for Online Advertising: Pay-per-Impression and Pay-per-Click Pricing

FRIDGEIRSDOTTIR, KRISTIN (London Business School) kristin@london.edu
NAJAFI ASADOLAHI, SAMI (London Business School) snajafi@london.edu

The Internet is currently the fastest growing advertising medium. Online advertising brings new opportunities and has many different characteristics from advertising in traditional media that support efficient and mechanized decision making. We consider online advertising from the viewpoint of a web publisher that generates revenues by selling advertising space on its website. The advertisers approach the web publisher and request their ad to be displayed to a certain number of visitors to the website. We consider two pricing schemes: pay-per-impression and pay-per-click pricing.

The web publisher faces the problems of pricing and managing capacity of the advertising space with the objective of maximizing the revenues generated. The advertisers are assumed to be impatient and not willing to wait if an advertising slot is not available. We suggest a queueing model (a loss model) for the operation of the web publisher considering the uncertainty of both the demand (the advertisers) and the supply (the visitors) with the advertising slots acting as servers. The two pricing schemes result in two different queueing models. Both are different from known models in the literature. For both models we derive closed form solutions of the probability distribution of number of advertisers in the system. We compare these two queueuing systems to known systems in the literature.

Having characterized the operation of the web publisher, we study its revenue maximization problem and determine the optimal advertising price. We provide managerial insights on the optimal price and compare the two pricing schemes.

Setting Prices on Priceline

ANDERSON, CHRIS (Cornell University) cka9@cornell.edu

The pricing of services (rooms, rental cars, airline seats etc…) online has dramatically changed how service firms reach customers. Online travel sales are expected to exceed offline (or traditional sales channels) by 2008 with 60% of these online sales through supplier managed websites and 40% through online travel agents (OTAs) like Expedia, Orbitz and Priceline. Initial thoughts about pricing online where very positive as from a marketing standpoint, firms now had several new methods to reach customers with the opportunity to increasingly segment customers across these new channels. Over time service providers have increased efforts to move customers back to firm managed distribution channels (firm specific websites and call centers) in an effort to control costs as firms save commissions and fees and maintain direct contact with the customer to facilitate loyalty programs and other marketing efforts.

In parallel with this migration to service provider managed channels (supplier websites and call centers) has been an effort to streamline prices and create price parity (equivalent prices regardless of booking method) across all distribution channels. The belief being that price parity instills some level of trust and comfort with the customer as they don't need to shop around for the best prices (for a given service provider). From the consumer standpoint online pricing and online travel agents like Orbitz, Expedia, Travelocity etc… and meta sites like Kayak that consolidate prices across different OTAs and suppliers have greatly simplified shopping as consumers can price compare with ease. Opaque channels like Priceline and Hotwire, those channels where the consumer does not know the service provider until purchase is completed, are one form of online selling which does not afford rate shopping by consumers and as a result provide firms an opportunity to segment customers and simultaneously offer multiple prices to the market without causing dilution of higher posted rates.

Consumers after visiting Priceline.com can choose two forms of shopping/purchasing; they can use Priceline like a traditional OTA and rate shop competing service providers or choose the “Name Your Own Price” option. Priceline originated as the “Name Your Own Price” marketplace where consumers indicate what they want to pay for service via a bid which is then accepted or rejected by service providers. In the name your price model customers specify details about the service – e.g. for a hotel the location and star level, then post a price (guaranteed with a credit card) that they are willing to pay. If unsuccessful with their bid customers can not re-bid on the identical product for a specified time (24 hours for hotels). After the bid is posted Priceline then determines if there is a service provider willing to accept the price. Priceline provides detailed daily bid reports to service providers, these bid reports detail all bids placed and whether or not they were accepted by a service provider.

In this talk I provide a brief introduction to opaque pricing and outline how Priceline allocates services to bidders (and how this differs to other opaque channels like Hotwire) and the impact this has upon what prices a service provider should provide to Priceline. Then using actual data that Priceline provides to service providers I will outline a model this is currently being used to determine optimal prices for firms to post on Priceline.

Room: *S1-151* *Patrice Marcotte and Gilles Savard*

15:30 Closing Remarks