

1 :: Scheduling of Dynamic In-Game Advertising

John Turner, Alan Scheller-Wolf, Sridhar Tayur

Dynamic in-game advertising uses the Internet to serve ads to video game consoles. Each time a player accesses a new level of a game, new ads are selected, downloaded, and integrated into the game. A billboard on the side of the road in a racing game may show an ad for Nike to one gamer, while showing an ad for Pepsi to another. To enable this technology, an intermediary known as the network provider contracts with advertisers and game publishers to provide the service of scheduling and serving in-game ads. We present a model for the scheduling problem faced by the leading network provider of dynamic in-game advertising in the United States. We solve our model using a two-step algorithm: 1) a large linear program is solved to establish service rates and 2) a real-time packing heuristic tracks these service rates while taking other side constraints into account. Backtesting is done over a 14-week time period, and the service rates are updated periodically by resolving the LP. Sensitivity to forecast errors is investigated and the distribution of the age of impressions is identified as a useful proxy for the revenue generated by the sales process. In our tests, we observe a make-good cost reduction of 80-87% as well as a shift in the age distribution that should lead to higher sales revenues. Due to these substantial benefits, our algorithm will be implemented in practice.

2 :: On Bounds in Network Revenue Management

Kalyan Talluri

Network Revenue Management can be formulated as a stochastic dynamic programming problem ("optimal" solution) whose exact solution is computationally intractable to solve. Consequently, a number of heuristics have been proposed in the literature, the most popular of which are the deterministic linear programming (DLP) model, and the randomized linear programming (RLP) model, both of which give upper bounds on the optimal solution value. These bounds are used to provide control values that can be used in practice to make accept/deny decisions for booking requests. Recently Adelman (2007) and Topaloglu (2007) have proposed alternate upper bounds and showed that their bounds are tighter than the DLP bound. Tight bounds are of great interest as it appears from empirical studies and practical experience that models that give tighter bounds also lead to better controls (better in the sense that they lead to more revenue). In this paper we prove relationships between all these bounds. Specifically, we show that Adelman's bound is weaker than the bound given by the RLP model, and that the RLP bound is weaker than Topaloglu's bound. Finally, we strengthen Topaloglu's bound further and show how it can be solved using column generation. To summarize our paper ranks all known bounds as follows: $DLP \geq \text{Adelman's bound} \geq PHS_{IP}$ and $PHS_{LP} \geq \text{Topaloglu's Lagrangian bound} \geq \text{Strong Lagrangian bound} \geq \text{Optimal Solution}$.

3 :: Near-Optimal Protection Level Policies for the Airline Capacity Allocation Problem

Woonghee Tim Huh, Thiam Lee, Retsef Levi

In this paper, we develop a non-parametric data-driven approach to the revenue management capacity allocation problem under monotonic fare arrivals. We introduce a new class of solutions and prove that policies from this class are near-optimal. We then describe a simple sample-driven algorithm that computes an instance from this class of policies with high probability. Furthermore, we give finite sampling 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. Furthermore, these finite sampling bounds hold for any demand distribution.

1 :: Queuing Network Approximations for Mass Dispensing and Vaccination Clinics

Ali Pilehvar, Jeffrey Herrmann

To respond to bioterrorism events or to curb outbreaks of contagious diseases, county health departments must set up and operate clinics to dispense medications or vaccines. Planning these clinics in advance of such an event requires determining clinic capacity and estimating the queueing that may occur in such facilities. We construct a queueing network model for mass dispensing and vaccination clinics and estimate the time that residents will spend at each workstation in such facilities. A key contribution is the development of useful approximations for queueing systems that have batch arrival, multiple-server batch processes and self-service stations. We compared the model's estimates to those from simulation experiments of realistic clinics using data collected from emergency preparedness exercises. Although this research was motivated by this specific application, the model should be applicable to the design and analysis of other similar queueing networks, including manufacturing systems with batch processes.

2 :: Capacity Management in Hospitals: Semi-Open Queueing Networks in the QED Regime

Galit Yom-Tov, Avishai Mandelbaum

In this research, we study capacity management problems in Hospitals. As a first stage, we analyze a Medical Unit with s nurses and n beds, which are partly/fully occupied by patients. This is a service system in which the patients need medical care, given to them by the nurses, assuming that there is a bed available in the medical unit for hospitalization. We model this medical unit by a semi-open queueing network with multiple statistically identical customers and servers. The questions we address are: How many servers (nurses) are required (staffing), and how many fixed resources (beds) are needed (allocation) in order to minimize costs while sustaining a certain service level? We answer this by developing policies that are asymptotically optimal. As the number of patients entering the system, the number of beds, and the number of servers grow to infinity together. These approximations turn out accurate for parameter values that are realistic in a hospital setting. Our asymptotic regime carefully balances quality of service against resources efficiency. As such, it has been called the QED regime (QED = Quality- and Efficiency-Driven). To date, we have developed QED approximation for the following service-level-objectives: (i) probability of blocking, (ii) probability of waiting, and (iii) expected waiting time. We have also developed fluid and diffusion limits for a closely-related open network with time-varying parameters, in which the restriction on the number of beds was relaxed, and thus blocking phenomenon cannot occur. This open model can fit other hospital scenarios such as Emergency Departments under mass-casualty events. Lastly, we are proposing several ways to continue this research in some interesting and useful directions.

3 :: Analysis of Coordinated Stochastic Systems Based on Operational Policies of Animal Shelters

Naragain Phumchusri, Julie Swann, Serhan Duran

Today there are two main policies used in animal shelters: In a "Kill" system, animals may be euthanized if the number of incoming animals is greater than available capacity. In a "No-Kill" system, a shelter rejects animals when full (diverting to capacity in another shelter if available or leaving to the street otherwise) but does not euthanize them for space. In this paper we study queueing systems with no, partial, or "full" coordination of job allocation and evaluate their performances, primarily focusing on the rejection rates and coordination costs.

1 :: Optimal Control in a Netflix-like Closed Rental System

Achal Bassamboo, Ramandeep Randhawa

We analyze the problem of product allocation to customers following the introduction of a new product in a closed rental system, such as Netflix. We consider two types of customers who differ in their rental time distributions. All customers desire the newly introduced product, and if a customer's request for this product is denied, she receives a substitute product and requests for the new product upon return. We study the control problem of minimizing the mean delay encountered by customers before obtaining the new product in a large market setting. We show that asymptotically this problem is equivalent to solving a linear program that depends on the entire rental duration distribution as opposed to mean alone. The optimal policy turns out to be a mixed priority rule that prioritizes the slower customer class while maintaining a base allocation to the faster customer class to ensure quick return.

2 :: Managing Service Systems with an Offline Waiting Option and Customer Abandonment

Amy Ward, Vasiliki Kostami, Sriram Dasu

Many service providers offer customers the choice of either waiting in a line, or going offline and returning at a dynamically determined future time. One well known example is the FASTPASS system at Disneyland. To operate such a system, the service provider must first make an upfront decision on how to allocate service capacity between the two lines. Then, during system operation, he must dynamically provide estimates of the waiting times at both lines to each arriving customer. The estimation of offline waiting times is complicated by the fact that some offline customers do not return for service at their appointed time. We show that when demand is large and service is fast, for any fixed capacity allocation decision, the two-dimensional process tracking the number of customers waiting in line and offline collapses to one dimension, and characterize the one-dimensional limit process as a reflected diffusion with linear drift. This dimensionality reduction allows us to (1) allocate capacity between the two queues in order to minimize average cost, and (2) to dynamically estimate inline and offline wait times using a simple scheme based on Little's law.

3 :: A Multi-Server Ticket Queue with Customer Abandonment

Kaan Kuzu, Susan H. Xu

1 :: Inventory Rationing for a System with Heterogeneous Customer Classes

Paul Enders, Ivo Adan, Alan Scheller-Wolf, Geert-Jan Van Houtum

Consider a single stockpoint where a single item is kept on stock. Poisson demands originate from two customer classes; demand from class 1 is served if any inventory is present, and lost if no inventory is available. Demand from class 2 is served if inventory is above a specified critical level and backordered otherwise. This policy allows for differentiation between customers in terms of service. For replenishment a base stock policy is used and replenishment lead times are assumed to be exponentially distributed. We develop an exact and efficient procedure to determine the steady state distribution of the resulting Markov process given policy parameters. Furthermore, we develop an efficient algorithm to determine the optimal parameters of this critical level policy. In a numerical study we compare the cost of the optimal critical level policy to the globally optimal cost and to two alternative, more naive, policies. The average loss by using a critical level policy is under 6% whereas the alternative policies are almost 11% and 31% from optimal. We also study the sensitivity of our policy to replenishment lead time distributions with smaller and greater variability and show that the optimal critical level policy is relatively insensitive to the coefficient of variation or the replenishment lead time distribution.

2 :: Setting Inventory Targets in a Multi-Product, Joint Service-Level Model with Correlated Stochastic Demands

Bahar Biller, Sridhar Tayur

Many practical inventory control problems involve multiple products whose demands are interdependent not only among each other, but also over time. Often the solution of these inventory problems requires the satisfaction of joint service-level criteria based on multiple products and purchase-orders over time. The determination of good inventory planning parameters satisfying such joint service-level criteria is a challenging problem, further complicated by the availability of only finite amounts of historical multivariate time-series demand data. We introduce a practical simulation-based framework that (i) provides flexible input-modeling support for the representation of multivariate and temporal dependencies in product demands; (ii) develops mixed integer linear programs specifically designed for high-dimensional settings in which joint service-level criteria naturally arise; and (iii) implements a simulation-optimization technique to determine inventory-target estimators with good statistical properties. A computational study illustrates the importance of correctly managing multi-product inventory subject to joint service-level criteria under multivariate time-series demand uncertainty.

3 :: Optimal Ordering/Rationing Policy for a Periodic Review System with Two Demand Classes and Backordering

Shaoxiang Chen, Jianjun Xu, Youyi Feng

Stock Rationing problems have been extensively studied in literature. For a finite horizon periodic review inventory system with a fixed setup cost and lost-sales, and with two priority demand classes – one stationary deterministic (constant) that must be satisfied immediately in each period and the other stochastic, the optimal policy structure is characterized in a previous research as a state dependent (s, k, S) , where s and S (ordering policy) determine when and how much to order, while k (rationing policy) specifies how much of the lower-priority (i.e., stochastic) demand to satisfy. This paper extends the results to a general setting of the problem that consists of two prioritized stochastic demand classes and allows backorders. We show that the optimality structure (s, k, S) continues to hold, though state dependent and more complex. Based on the super-modularity of the cost function, the rationing structure is partially obtained.

TA05 :: Thursday, 8:00-9:30am

Supply Chain Management

Room 1528

1 :: Contracting in a Repeated Supply Chain Relationship: The Role of Asymmetric Cost Information

Ching Chyi Lee, Justin Ren, Deming Zhou

We consider the problem of procurement contracting in a one-supplier-one-buyer relationship, where the buyer's cost information is private, but they have opportunities to transact repeatedly. The supplier moves first by offering a whole-sale price contract, after which the buyer can either accept by ordering from the supplier, or reject the offer and seek his outside option. After profits for both are collected, the game re-starts again, altogether for N times. The key interest of this paper is to see if the buyer's true cost information can be revealed during the repeated interactions (without a screening contract). We find that the buyer's 'information advantage' persists in the initial phase of the repeated game. However, as the game goes on, such an information advantage gradually dissipates, and the buyer starts to randomize between accept and reject the contract. Therefore, his private information is revealed in a probabilistic fashion. The supplier's long-run equilibrium strategy is of a threshold type: She starts with offering a low selling price and keeps doing so until her belief about the buyer's type is over a threshold. These results indicate that repeated supply chain relationships have significantly different dynamics than those observed in a one-shot supply chain relationship. We discuss insights that are helpful to firms in shaping their strategies in their long-time supplier relationship.

2 :: Supply Chain Coordination Using Buyback Contracts with Risk-Averse Retailers

Qindong Liu, Jan Stallaert

We study the problem of supply chain coordination under a buyback contract involving a risk-neutral supplier and a risk-averse retailer facing a single-period stochastic demand. We define "risk" for the retailer as the failure to reach a target profit, known as "downside risk." Consistent with financial theory, we measure this downside risk by semi-variance, i.e., the expected

squared deviation from the expected profit. We assume that the supplier offers a buyback contract to induce the retailer to order a quantity that coordinates the supply chain. We design an optimal buyback contract where the buyback price is a function of the retailer's degree of risk aversion. The supply chain can be coordinated when the retailer's risk type is observable and discrimination of buyback price is also allowed. If the retailer's type is unobservable and the supplier offers a standard contract to the retailer regardless of her risk profile, we show that the optimal buyback price from the supplier's viewpoint is the contract designed for a risk neutral retailer. For a strict risk-averse retailer, such standard buyback contract cannot coordinate the supply chain. Finally, we design a scheme where the supplier can set a menu of buyback contracts as put options where the buyback prices are equivalent to the exercise prices and the retailer acquires the proper put option corresponding to her risk profile. We show how the supplier should compute the option's ask and exercise prices such that (i) the supply chain is coordinated; (ii) the retailer honestly reveals her type and (iii) the outcome is a "win-win" situation for both parties.

3 :: Supply Chain Coordination: Can Insights be Misleading?

Valery Pavlov

By now, Supply Chain Coordination is a thoroughly studied and a well understood area. Many discovered insights are fairly intuitive, very convincing and, at this point, seem to belong to textbooks. Consider, for example, a starting point of Supply Chain Coordination, a bilateral monopoly. Without going into details, one could say that, in general, inflexible contracts, such as wholesale price contract, are inefficient because they lead to Double Marginalization, that information sharing is beneficial and that under complete information parties can always design a contract that results in a Pareto-efficient outcome. However, we argue that these insights can be easily misleading and, sometimes, may be even irrelevant from the practical perspective. The main reason of why and how insights can be potentially misleading comes from the use of game-theoretic models. One of the insights of game theory itself is that equilibrium outcomes are highly sensitive to even slight modifications of assumptions or a structure of the game. In what follows, we show that the aforementioned well-established findings, that is Double Marginalization, desirability of information sharing and Pareto-efficiency under complete information are, in fact, questionable. Our observation is that similarly looking contracting models may come to so different conclusions, that whatever insights could be possibly derived are not likely to be robust.

TA06 :: Thursday, 8:00-9:30am

Supply Chain Management

Room 1518

1 :: Value of Postponement for a Manufacturer that Produce a Line of Two Products under Uncertain Component Yields

Qi Feng, Yusen Xia

In this paper, we consider the value of postponement for a manufacturer that produces two products using the same raw materials. The production process of the manufacturer includes two stages. The first stage is from the raw material to the semi finished products. The second stage provides the finished products. Both stages incur uncertain yields. We compare the difference between pooled and separate first stage production process and provide managerial insights into when the manufacturer prefers one system to the other. Furthermore, we characterize the equilibrium for two manufacturers in which one uses pooled first stage planning and the other use separate first stage production.

2 :: The Impact of Quick Response in Inventory-Based Competition

Felipe Caro, Victor Martínez-de-Albéniz

We propose a multi-period extension of the competitive newsvendor model of Lippman and McCardle (1997) to investigate the impact of quick response under competition. For this purpose, we consider two retailers that compete in terms of inventory: customers that face a stockout at their first-choice store will look for the product at the other store. Consequently, the total demand that each retailer faces depends on the competitor's inventory level. We allow for asymmetric reordering capabilities, and we are particularly interested in the case when one of the firms has a lower ordering cost but can only produce at the beginning of the selling season, whereas the second firm has higher costs but can replenish stock in a quick response manner taking advantage of any incremental knowledge about demand (if it is available). We visualize this problem as the competition

between a traditional make-to-stock retailer that builds up inventory before the season starts versus a retailer with a responsive supply chain that can react to early demand information. We provide conditions for this game to have a unique pure-strategy subgame-perfect equilibrium, which then allows us to perform numerical comparative statics. Our results confirm in a competitive setting the intuitive fact that quick response is more beneficial when demand uncertainty is higher, or exhibits a higher correlation over time. Finally, we find that part of the competitive advantage from quick response arises from the asymmetry in response capabilities.

3 :: Private Label Products: When Structural Inefficiency Increases Supply Chain Profits

Liwen Chen, Stephen Gilbert, Yusen Xia

We examine how a retailer's ability to introduce an internally produced low-end version of a national brand affects both the product assortment that is offered to consumers as well as the total profit that is generated by the supply chain. Our analysis allows for the possibility that the retailer may have some influence over the extent to which consumers are exposed to the product category. Under the assumption that consumers perceive the quality of the private label to be lower than that of the national brand, we show that there exists a wide range of conditions under which the total decentralized supply chain profits increase when a low-end private label is introduced, even though such a product would not be included in the optimal product line for a vertically integrated channel. Thus, even though the introduction of the private label is structurally inefficient, it can nevertheless result in an increase in the overall efficiency of the decentralized supply chain. We obtain our results with a relatively general family of consumer-valuation distribution functions that allows us to explore how different distributions of valuations affect the direction and magnitude of impact that the private label has upon the profits of the manufacturer, the retailer, and their combined profits. Interestingly, for product categories for which the retailer may have significant influence over their exposure to consumers, we are able to show conditions under which both the manufacturer and the retailer benefit from the introduction of the private label.

TB01 :: Thursday, 10:00-11:30am

Revenue Management

Room 1505

1 :: Revenue Management and Delay Tactics under Competition and Customer Choice

Philipp Afeche

Consider the problem of two providers who compete for heterogeneous and time-sensitive customers with private information about their preferences. How should they design price-delay menus and schedule their resources? We characterize the price-delay equilibria and discuss how they depend on capacity levels, competition, and the value of market segmentation.

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

Qian Liu, Stephen Shum

When customers make choices among alternatives that involve risk, their evaluations and decisions are affected not only by economic payoff but also emotional consequences of ex-post outcomes such as psychological disappointment. When outcome is observed, customers compare actual outcomes with prior expectations. When the actual outcome falls short of expectations, a sense of disappointment is caused. Thus, potential psychological disappointment is taken into account while customers make purchase decisions under uncertainty. Customers' aversion to disappointment creates a behavior motivation for us to revisit a firm's pricing and capacity decisions. We apply Bell's model of disappointment (1985) to customer's purchase decision problem and explore the impacts of such psychological satisfaction and disappointment on firms' pricing and capacity decisions. We study a two-period model with limited capacity in the first period and no capacity constraints in the second period. Customers are heterogeneous in both valuations of the product and degrees of aversion to disappointment. Each customer decides when to purchase based on his perceived total utility which includes both economic payoff and psychological satisfaction. We analyze the optimal pricing policy and the initial rationing level to maximize revenue. We show that when aversion to disappointment increases in valuations, that is, the higher the valuation, the larger degree of aversion to disappointment, a mark-up pricing policy with appropriate level of rationing can be used to segment the market. Moreover, when customer valuations decline do

not fast over time, mark-up policy can generate higher revenue than any mark down or uniform pricing policy. This result is a significant contrast to that under expected utility framework. We also find that when customers' levels of aversion to disappointment is independent of valuations, psychological satisfaction and disappoint doesn't affect the firm's optimal pricing and capacity decisions. Thus, in this case, there is no value in gaining additional information on psychological disappointment.

3 :: Saving Seats for Strategic Customers

Eren Cil, Martin Lariviere

We consider a service provider in a market with two segments. Members of the first request a reservation ahead of service and will not patronize the firm without one. Members of the second walk in and demand service immediately. These customers have a fixed cost of reaching the firm and may behave strategically. In equilibrium, they randomize between walking in and staying home. The service provider must decide how much of a limited capacity to make available to reservation customers. When the reservation segment offers a higher per customer margin, the firm may opt to decline some reservation requests in order to bolster walk-in demand. When walk-in customers are more valuable, we have a variation of Littlewood (1972). Where Littlewood would always save some capacity for valuable late arrivals, here it is possible that the optimal policy saves no capacity for walk-ins. Thus, it may be better to ignore rather than pamper walk-in customers.

TB02 :: Thursday, 10:00-11:30am

Healthcare / Non-profit

Room 1524

1 :: Setting Pediatric Vaccine Stockpile Levels using Decision Analysis

Ruben Proano, Sheldon Jacobson, Edward Sewell, Janet Jokela

The effectiveness of routine vaccination programs depends on the availability of a stable vaccine supply. Public health officials established the Pediatric Vaccine Stockpile Program (PVSP) as the best strategy to mitigate the impact of vaccine supply interruptions that frequently affect such supply. This paper uses the relative importance of each pediatric vaccine to define a multi-attribute utility function that models the many conflicting interests of PVSP decision-makers when ordering vaccines for the next fiscal year. This paper explores the implications of optimizing a proposed expected utility function under budget constraints for hypothetical (albeit likely) scenarios of vaccine supply, and the potential behavior of health care decision-makers. The results obtained indicate that spending the entire allocated budget does not necessarily provide the maximum utility, and that the best protection against supply interruptions is to use vaccine stockpiles to aggressively increase population immunity.

2 :: Optimal Policies for Breast Biopsy Decision-Making

Jagpreet Chhatwal, Oguzhan Alagoz, Elizabeth Burnside

Breast cancer is the most common non-skin cancer affecting women in the U.S. Every year, more than 20 million mammograms are performed in the U.S. Breast biopsy is commonly performed on the suspicious findings on mammograms to confirm the presence of cancer. Currently, 700,000 biopsies are performed annually in the U.S; 55-85% of these biopsies ultimately are found to be benign breast lesions, resulting in unnecessary treatments, patient-anxiety and expenditures. This paper addresses the decision problem faced by radiologists: when should a woman be sent for biopsy based on her mammographic features and demographic factors? This problem is formulated as a finite-horizon discrete-time Markov decision process. The optimal policy of our model shows that the decision to biopsy should take the patient's age into account; particularly, an older patient's risk threshold for biopsy should be higher than that of a younger patient. When applied to the clinical data, our model outperforms radiologists in the biopsy decision-making problem. This study also derives structural properties of the model, including sufficiency conditions that ensure the existence of a control-limit type policy and non-decreasing control-limits with age.

3 :: Influenza Vaccine Supply Chain with Multiple Agencies

Hamed Mamani, Stephen Chick, David Simchi-Levi

Billions of dollars are being allocated for influenza pandemic preparedness and vaccination is a primary weapon for fighting influenza outbreaks. We analyze a model with multiple countries purchasing vaccines. We show that selfish behavior by individual governments result in a sub-optimal vaccine allocation to different countries. i.e., vaccine stockpiles are diverted from regions where they are needed to most (e.g., countries in Southeast Asia) to the places that don't have a significant influence on the spread of the disease (e.g., US or western Europe). We finally design contractual agreements between different governments that guarantee optimal vaccine levels for all countries. These methods increase the probability of containing a pandemic at its source.

TB03 :: Thursday, 10:00-11:30am

Service Operations

Room 1511

1 :: Knowledge Management in Call Centers: How Routing Rules Influence Expertise and Service Quality

Christoph Heitz, Geoffrey Ryder, Kevin Ross

In a call center, customers are assigned to service agents by routing policies that seek to balance several objectives. Usually, these policies follow myopic rules in order to minimize the waiting time or maximize the quality experienced by the next customer. However, there is a secondary effect of the routing assignment: by learning-on-the-job, the development of the agents' expertise depends on the calls they take. In this paper, we address the effect that agent learning has on the service level experienced by customers. A dynamical model of learning and forgetting links the routing policy with knowledge acquisition, treating expertise as an endogenous rather an exogenous variable. Our results indicate that the routing may have a big impact on the knowledge level of a firm, and that classical routing policies may have a negative impact on the distribution of that knowledge.

2 :: A Routing Policy for Call Centers Designed to Respond to Unexpected Overloads

Ohad Perry, Ward Whitt

We consider a stochastic model of a call center with two customer classes and two associated service pools, with agents in each pool primarily dedicated to the designated customer class, but with all agents cross-trained and allowed to serve the other class. The service rates may depend on both the customer class and the agent type. Our goal is to exploit the cross-training to develop a routing policy that responds to unexpected overloads, while still providing good performance under normal loads. Under ordinary circumstances, we want class- i customers to be served by type- i agents. The sharing is activated automatically when there is congestion. We propose a fixed-queue-ratio assignment rule with thresholds for newly available agents. We develop approximations for the performance and show that they are accurate by comparing with simulations.

3 :: Staffing Call Centers with Differentiated Levels of Service: Approximate Solutions via Lagrangian Approach

Achal Bassamboo, Assaf Zeevi

Modern call center operations aim to provide differentiated levels of service to various classes of customers. These are typically specified by a Service Level Agreement (SLA) between the provider and the customers (or stipulated by a contract between the call center and firm outsourcing its service to it). The goal of the call center manager is to minimize staffing costs subject to meeting these pre-specified quality of service constraints. The two key decision variables in such settings are: staffing levels, i.e., number of agents of various skill sets to be staffed in a work shift; and routing control that determines the real-time assignment of incoming calls to agents. In this paper, we propose a tractable solution method for a class of such constrained optimization problems. At the heart of the approach is a transition from a constrained optimization problem, to an unconstrained cost minimization problem that, roughly speaking, is a Lagrangian formulation of the former. We show that the staffing level prescribed by the method performs as well as the optimal staffing levels of the original constrained problem.

1 :: Optimal Product Assortment under Commercial Returns

Alex Grasas, Aydin Alptekinoglu, Elif Akcali

Ever increasing product variety makes dealing with product returns more difficult for retailers. Motivated with the question of whether retailers should consider product returns when merchandising (as they compose their product assortments), we explore the interactions between product assortment and return policy decisions of a firm under two basic operational modes, make-to-order (MTO) and make-to-stock (MTS). We have a demand model grounded on individual consumer behavior. Consumers decide which product out of a given assortment to buy in the first stage of a nested multinomial logit model, and then decide to keep or return the item in the second stage. The firm, an expected profit maximizer, first makes an assortment decision from an exogenous set of potential products. We call products with high (low) expected valuation popular (eccentric), because they are more (less) likely to be purchased by a typical consumer. Second, the firm sets a refund fraction, the percentage of price refunded upon return. Finally, in the MTS case, the firm also makes inventory decisions for each product offered. Our main finding is that the optimal assortment has a counterintuitive structure for relatively strict return policies (with a sufficiently low refund fraction) and this structure differs between the two operational modes. In the MTO case, it is composed of most eccentric products; whereas, in the MTS case, a combination of most popular and most eccentric products. Therefore, not only that the retailers should carefully consider their return policy in merchandising, they should also take their basic operational mode (MTO versus MTS) into account.

2 :: Demand Management and Inventory Control for Substitutable Products

Jing-Sheng Song, Zhengliang Xue

This paper studies dynamic inventory and pricing decisions for a set of substitutable products over a finite planning horizon. We present a general stochastic, price-dependent demand model that unifies many commonly used demand models in the literature. Unsatisfied demands are backlogged. There are linear purchasing, inventory-holding, and backordering costs. The objective is to maximize the total expected discounted profit. The original formulation is not jointly concave in the decision variables and is therefore intractable. One key observation here is that the problem becomes jointly concave if we work with the inverse of the price vector -- the market share vector. We characterize the optimal policy and develop algorithms to compute it. We establish conditions under which the optimal policy demonstrates certain monotonicity property, which, in turn, can greatly enhance computation. We also analyze the myopic policy and its optimality, and present a numerical study to illustrate the interplay of the pricing and inventory decisions.

3 :: On a Make-to-Stock Queue with Advance Orders and Time Window Deliveries

Zied Jemai, Fikri Karaesmen

We consider a supplier operating a capacitated production system that receives orders in advance. The customers accept deliveries within a time window that specifies the earliest date at which deliveries are acceptable and the latest date beyond which penalties incur. We report results on the optimal base stock levels and the optimal costs for such a system as well as some basic properties in comparison to other due-date mechanisms.

1 :: Supplier Quality and Relationship Choice Under Competition

Anupam Agrawal

In this paper, we explore the quality-competition landscape. We detail out conditions under which suppliers, with the same technology, supplying similar components to different buyers, price the components differently. Using a multi-principal multi-

agent setting, we show that buyer profits from shared sources can depend on three factors : the capability of suppliers in improving quality, the existing quality levels at the suppliers and the accuracy of the buyers' in-house inspection technology. Once critical thresholds on these quality parameters are crossed, buyers can get early leads in market share by a policy of sustained investment in quality at the suppliers.

2 :: Capacity Reservation Contracts Using Common Agency

Feryal Erhun, Basak Kalkanci

We examine how the suppliers must choose their contracts in an assembly system where the components are perfectly complementary and when the manufacturer has better information on mean demand. We identify and compare optimal complex contracts and wholesale price contracts. We observe that although the complex contracts lead to a higher efficiency than wholesale price contracts, using wholesale price contracts can be a better option for some players in the system under certain conditions. We also show that many of the results in the existing literature are robust with respect to the information structure and are driven by the fact that the components are perfectly complementary in an assembly system.

3 :: Strategic Single vs. Dual Sourcing - Under Conditions of Catastrophic Disruptions and Stochastic Demand and Capacity

Wenbin Wang, Xinxin Hu, F. Robert Jacobs

This paper provides a generalized multi-period decision model that characterizes sourcing strategies in risky environments. We intend to provide a model that is more general than those found in the current literature. We consider the uncertainties caused by catastrophic disruption, stochastic demand and stochastic capacity. We assume a manufacturer sources a key component from one or two suppliers with a reliable but expensive market as backup. Each supplier is exposed to various factors that might impact its yield capacity. Some only make the capacity fluctuate, other severe factors will cause catastrophic shutdown. A supplier that experiences a shutdown is modeled to recover in a future period by certain probabilities. The generalized multi-period decision model characterizes the best procurement strategies for a manufacturer relative to single or dual sourcing. Finally our paper shed lights on how decision rules can be developed to guide the single and dual sourcing decision. To our knowledge existing research has not captured sourcing uncertainty in as general a manner as the model we propose.

TB06 :: Thursday, 10:00-11:30am

Miscellaneous OM

Room 1518

1 :: A Comparison of Price and Quantity Competition in Oligopolies with Differentiated Products

Amr Farahat, Georgia Perakis

We compare equilibrium profits under price (Bertrand) and quantity (Cournot) competition in oligopolies with an arbitrary number of firms offering gross substitute products. For the affine demand case, we provide a precise characterization of the profit relationship in terms of 1) the number of firms, 2) their relative quality and cost differences, and 3) the competition intensity, defined by the maximum absolute value of total change in competitors' demand over change in own demand due to a unit change in own price. We first examine the case where firms have the same demand sensitivity to own price and the same demand sensitivity to competitor prices. For this case, we prove that, for each firm, the profit achieved under quantity competition is at least as high as the profit achieved under price competition if the number of firms is less than 8 or if the competition intensity is less than 0.739. We also prove that the total profit of the industry under quantity competition is at least as high as the total profit under price competition if the number of firms is less than 28 or if the competition intensity is less than 0.909 or if the differences in quality and cost competitiveness between firms are small. These results are then shown numerically to qualitatively hold for general affine demand functions with variable demand sensitivities to prices. For the multinomial logit demand case, we prove that Cournot profits are at least as high as Bertrand profits for each firm when firms are symmetric. The result holds numerically for non-symmetric firms.

2 :: The Effect of Reliability of a Repairable Component on Life Cycle Costs

Kurtuluş B. Öner, Gudrun P. Kiesmüller, Geert-Jan Van Houtum

Complex technical systems such as computer networks, defense systems, material handling systems and medical systems are crucial for the operations of their users. The users require high availability because their primary processes may halt when these systems are down. In general, a system includes critical components and each failure of such a component definitely leads to a system breakdown. The availability of the system is mainly determined by these components. We investigate a situation in which an Original Equipment Manufacturer (OEM) sells systems together with service contracts that cover the whole life time of the systems. The downtime costs of the systems are high. A target system availability level is set in the contracts. The OEM provides this target level by attaining a reliability level of each critical component during system design and by managing a spare parts inventory of the component during the exploitation phase of the systems. We introduce a quantitative model for analyzing the effect of the reliability level of a critical repairable component on the Life Cycle Costs of a system. We generate insights for the effect of various input parameters on the optimal reliability level and optimal costs, and in particular for the effect of downtime cost parameters.

3 :: Process Investment Decisions of a Startup under Uncertainty

Fehmi Tanrisever, Sinan Erzurumlu, Nitin Joglekar

TC01 :: Thursday, 2:00-3:30pm

Revenue Management

Room 1505

1 :: Competition and Contracting in Service Industries

Gabriel Weintraub, Ramesh Johari

Dynamic in-game advertising uses the Internet to serve ads to video game consoles. Each time a player accesses a new level of a game, new ads are selected, downloaded, and integrated into the game. A billboard on the side of the road in a racing game may show an ad for Nike to one gamer, while showing an ad for Pepsi to another. To enable this technology, an intermediary known as the network provider contracts with advertisers and game publishers to provide the service of scheduling and serving in-game ads. We present a model for the scheduling problem faced by the leading network provider of dynamic in-game advertising in the United States. We solve our model using a two-step algorithm: 1) a large linear program is solved to establish service rates and 2) a real-time packing heuristic tracks these service rates while taking other side constraints into account. Backtesting is done over a 14-week time period, and the service rates are updated periodically by resolving the LP. Sensitivity to forecast errors is investigated and the distribution of the age of impressions is identified as a useful proxy for the revenue generated by the sales process. In our tests, we observe a make-good cost reduction of 80-87% as well as a shift in the age distribution that should lead to higher sales revenues. Due to these substantial benefits, our algorithm will be implemented in practice.

2 :: Towards a Unified Theory of Procurement Contract Design: Production Flexibility, Spot Market Trading and the Structure of Option Contracts

Pamela Pei, Tunca Tunay, David Simchi-Levi

3 :: Revenue and Capacity Management for a Multi-class Stochastic Service System

Dan Zhang

We study a stochastic service system modeled as an $M/M/c/c$ queue accessed by several customer classes. We first consider a dynamic pricing problem where the prices for all customer classes are controlled over time to maximize the total discounted reward for the service provider. The problem is formulated in a continuous-time Markov decision process (MDP) framework. We study the structural properties of the problem. Furthermore, we study a deterministic fluid formulation as an optimal control problem that admits a simple solution. The resulting policies can be adapted to control the corresponding stochastic system. We build on the fluid model to analyze a capacity management problem that takes into account the contingent

revenue management decisions. Throughout the paper, we cast an admission control problem where the prices are fixed and customer service requests are either accepted or rejected upon arrival as a special case.

TC02 :: Thursday, 2:00-3:30pm

Healthcare / Non-profit

Room 1524

1 :: Reducing Boarding in a Post-Anesthesia Care Unit

Carter Price, Bruce Golden, Michael Harrington, Edward Wasil

Boarding occurs when a surgical patient spends the night in the post-anesthesia care unit (PACU) because a post-operative bed is unavailable. As hospitals operate closer to capacity, boarding becomes increasingly more common. We use integer programming to develop a surgical schedule designed to minimize the number of boarders by better matching the flow of patients into the hospital with those being discharged.

2 :: When Should Hospitals Stagger Nurse Schedules?

Wen-Ya Wang, Diwakar Gupta, Sandra Potthoff

Hospitals routinely develop nurse schedules that cover a period of 2 to 8 weeks and are posted 4 to 8 weeks in advance. Once posted, changes to the schedule require voluntary participation by the nurses. We propose a strategy where nurses are divided into a small number of cohorts and schedules are phase shifted for different cohorts. This allows the nursing unit managers to benefit from more frequent demand updating without having to renegotiate scheduling rules with the nurses' unions. We show that staggering schedules by cohort can be beneficial when forecast accuracy improves upon postponing the time epoch at which the forecast is developed.

3 :: Dynamic Scheduling of Outpatient Appointments under Patient No-shows and Cancellations

Nan Liu, Serhan Ziya, Vidyadhar Kulkarni

We propose a Markov Decision Process formulation for the outpatient scheduling problem under patient no-shows and cancellations, and develop a heuristic policy based on the idea of using one step of the policy-improvement algorithm on a state-independent policy. The heuristic assigns an appointment date to each patient upon his/her request, depending on the system state, i.e., the current appointment schedule. Using the data from a Family Medicine Center, we estimate some of the model parameters, in particular cancellation and no-show probabilities as a function of the appointment delay of the patient. Our simulation study strongly suggests that our heuristic method can bring significant improvements over some of the policies that have been proposed and are being used in practice especially when the clinic capacity is small with respect to the patient demand.

TC03 :: Thursday, 2:00-3:30pm

Service Operations

Room 1511

1 :: Competition in Large Scale Service Systems: Do Waiting Time Standards Matter?

Itay Gurvich, Gad Allon

We consider a market with several large-scale service-providers that compete on both prices and service levels as measured by the response time to a service request. The literature on the many-server heavy-traffic regime provides significant simplifications towards the optimal capacity sizing of large-scale providers operating as monopolists in the market but falls short of providing similar simplifications for a competition setting in which each firm's decision is affected by its competitors' actions. In this talk, we introduce a framework that combines many-server heavy-traffic analysis and the concept of epsilon-Nash equilibrium and apply it to the study of equilibria in a market with multiple large-scale service providers. In analogy to fluid and diffusion approximations for queueing systems, the proposed framework allows us to provide first-order and second-order characterization results for the equilibria in these markets, thus going beyond the tractability boundaries imposed by the regular Nash equilibrium analysis.

2 :: Quality-Speed Tradeoff in Care Intensive Services

Mehmet Fazil Pac, Senthil Veeraraghavan

In this paper we analyze the trade-off between the service speed and the service quality for firms maximizing revenues. The quality of the service declines, either abruptly or gradually, with increasing service speed in several service environments. In this study we model a service system where the expected service quality varies with service speed. An arriving customer decides whether to join a service based on its value, speed of service and his waiting costs. Under such a circumstance increasing the service rate reduces the waiting costs, but also reduces the expected service quality. Our aim is to analyze the trade-off between service speed and service quality for the consumer and the service provider. We derive the equilibrium system characteristics for a monopoly firm, and show its optimal service speed. In a competitive setting, where firms serve a common queue, we find that in equilibrium both the firms, might prefer to compete on service speeds instead of service value. This provides an explanation for the presence of high service speeds and low service values even in care-intensive service processes.

3 :: How Heavy Should Traffic Be? Revenue Optimization with Non-linear Delay Costs

Sunil Kumar, Ramandeep Randhawa

We study the impact of non-linear, convex delay costs on the pricing and capacity choice decisions of a monopolistic firm. We model the firm as either a single-server or a multiserver queue with i.i.d exponential service times. For the most part, we consider an arrival rates of the form $n\lambda(p + EW^r)$, where n is the market size, λ is the "demand curve", p is the price charged by the firm, W is the steady-state waiting time in the queue, and $r \geq 1$ is a parameter fixed up front. In particular, customer behavior as captured by $\lambda(\cdot)$ is assumed to be independent of the market size. Analyzing the properties of the exact solution to the firm's profit maximization problem, we demonstrate that in a large market setting, i.e., as n grows without bound, the firm chooses to approach heavy traffic in the traditional $1 - O(n^{-1})$ manner only if $r = 1$. In general, the traffic intensity of the optimal solution satisfies the relation $1 - C_1 n^{-\frac{1}{r+1}} \leq \rho_n \leq 1 - C_2 n^{-\frac{1}{r+1}}$ for some constants $C_1, C_2 > 0$. That is, the nature of the delay cost affects the rate at which the firm chooses to approach heavy traffic, and the firm operates in a "heavier" traffic regime than the conventional regimes chosen for asymptotic analysis. In particular, our results suggest that the firm will choose not to operate in the so-called Halfin-Whitt regime in the multiple server case. We obtain similar results when studying the problems of capacity sizing at a fixed price, which is equivalent to selecting capacity to manage customer delay costs, and pricing at fixed capacity levels. We briefly discuss a two-class setting and show that if the parameter r is different for the two classes, the firm operates in a regime dictated by the class that has the higher delay sensitivity. We extend this analysis to the case when $r < 1$, i.e., concave delay costs. We show that the optimal regime is "lighter" than conventional and the way in which heavy traffic is approached depends both on the delay costs as well as whether we are in a single or multi server setting.

TC04 :: Thursday, 2:00-3:30pm

Inventory Theory / Multi-echelon

Room 1520

1 :: A Customer-Item Decomposition Approach to Stochastic Inventory Systems with Correlation

Yimin Yu, Saif Benjaafar

We consider an inventory system with a single stage, periodic review, correlated, non-stationary stochastic demand and correlated, non-stationary stochastic and sequential leadtimes. We use the customer-item decomposition approach to decompose the problem into sub-problems, each involving a single customer-item pair. We then formulate each subproblem as an optimal stopping problem. We use properties that arise from this formulation to show that the optimal policy is a state-dependent base-stock policy and to show, for several cases, that the optimal policy can be obtained via a polynomial time algorithm. We also use the formulation to construct a myopic heuristic which leads to an explicit solution for the optimal policy in the form of a critical fractile. We characterize conditions under which the myopic heuristic is optimal.

2 :: Asymptotic Optimality of Order-up-to Policies in Lost Sales Inventory Systems

Woonghee Tim Huh, Ganesh Janakiraman, John Muckstadt, Paat Rusmevichientong

We study a single-product single-location inventory system under periodic review, where excess demand is lost and the replenishment lead time is positive. The performance measure of interest is the long run average holding cost and lost sales penalty. For a large class of demand distributions, we show that when the lost sales penalty becomes large compared to the holding cost, the relative difference between the cost of the optimal policy and the best order-up-to policy converges to zero. For any given cost parameters, we establish a bound on this relative difference. Numerical experiments show that the best order-up-to policy performs well, yielding an average cost that is within 1.5% of the optimal cost when the ratio between the lost sales penalty and the holding cost is 100.

3 :: Average Optimal Policies for Inventory Models with Convex Ordering Cost

Frederick Zahm

If an inventory system faces its ordering decision in a short-term operational context, it is plausible that an increased order quantity would see increased marginal cost for additional units---and thus a convex ordering cost function may be appropriate. We investigate a periodic-review inventory model under a long-run average expected cost criterion, formulated as a Markov decision process with complete observations. The ordering cost for any period is a convex, piecewise linear function of the (real-valued) order quantity of a single product. Orders are delivered immediately. All parameters are stationary. Excess demand is fully backlogged. The inventory holding (or backlogging penalty) cost for any period is a convex function of the (real-valued) inventory level. The random demands are independent and identically distributed across periods; the (known) distribution function is subject to mild assumptions. After reviewing relevant literature, we outline our argument that there exists an optimal policy of the finite generalized base stock type; in closing, we indicate potential extensions of this result.

TC05 :: Thursday, 2:00-3:30pm

Supply Chain Management

Room 1528

1 :: Disruption-Risk Management in Global Supply Chains

Robin Dillon, Joseph Mazzola

Global supply chains are an integral part of the Twenty-First century economy. A disruption occurring within a supply chain, whether it be attributable to a natural disaster or a human-induced event, presents substantial risk to companies within the supply chain and the markets that it serves. We discuss a risk-based modeling approach to managing GSC disruption. We introduce the concept of a supply-risk network to capture potential disruptions and the attendant risks they impose on the flow of goods and information. Using this framework we then formulate a GSC disruption-risk model that allows for organizations within the supply chain to plan strategically for the sourcing and flow of goods throughout the supply chain in a manner that directly incorporates the risk of disruption. The GSC disruption-risk model is formulated as a two-stage stochastic integer programming problem with fixed recourse. We illustrate the supply-risk network modeling approach on an example and discuss managerial implications.

2 :: Disruption Management with Cross Pricing and Cross Production

Tieming Liu, Ho-Yin Mak, Max Shen

We study the optimal capacity investment portfolio of a firm by considering responsive pricing and partial manufacturing flexibility. The firm produces two substitutable products and invests in resources that can be used for both products but with different efficiencies. We characterize the structure of the optimal pricing and allocation policies and derive managerial insights on the benefits of production and pricing flexibilities.

3 :: Supply Contracts in the Face of Production Disruptions

Suleyman Demirel, Roman Kapuscinski, Ching-hua Chen-Ritzo

We consider a single product with deterministic demand and exogenous price in a multi-period setting. A single buyer has an opportunity to engage in a contract with a reliable supplier (he), an unreliable one (she), or both. The unreliable supplier faces random disruptions, disabling her to produce anything for a number of consecutive periods. The reliable supplier can either serve as a primary (regular) supplier, or as a backup supplier. When she serves as the primary supplier, there is no need for any additional supplier and she builds sufficient capacity to serve the retailer. When she does not serve as the primary supplier, she may or may not serve as a backup supplier. We consider these two cases separately. When the reliable supplier is limited to serving as a primary supplier, she may charge a premium for being reliable, while the unreliable supplier may end up paying an unreliability penalty. When the reliable supplier can serve either as a primary supplier or a backup supplier, there exist no pure-strategy Nash equilibria in most realistic settings if suppliers are allowed to quote single prices. If, however, suppliers can quote prices contingent on the retailer's sourcing strategy, i.e. single sourcing vs. dual sourcing, the pricing game has unique pure-strategy Nash equilibrium. In equilibrium, the downstream buyer may use both suppliers or just one. In the case of single-sourcing from the unreliable supplier, the unreliable supplier offers rebates for any items not delivered on time. Reliable supplier may be used as a backup supplier. In the case of single-sourcing from reliable supplier, she charges reliability premium.

TC06 :: Thursday, 2:00-3:30pm

Miscellaneous OM

Room 1518

1 :: Improving Valuation under Consumer Search: Implications for Pricing and Profits

Olga Perdikaki, Jayashankar Swaminathan

In this paper, we study the impact of influencing consumer valuation under monopoly and competitive settings. We first develop a stylized model for changes in consumer valuation and how it affects the eventual purchase decision. Using the above, we analyze the price and profits realized by the firm. Under monopoly, we show there exists a threshold cost of investment beyond which the firm will not strive to improve consumer valuation. We also show that a stochastically larger consumer valuation leads to higher prices and profits. Surprisingly, we find that in certain regimes the optimal price charged by the firm may indeed decrease with the purchase cost per unit. In a duopoly setting, we consider situations where a consumer's valuation may be increased by one retailer; however, the consumer may decide to buy the product from the other retailer. We characterize the Nash equilibrium in the above situation and study it under different regimes related to market expansion, retailers' proximity, direction of consumer flow, and degree of changes in consumer valuation. Among other results, we find that under certain competitive regimes, a retailer who improves consumer valuation beyond a certain level may get hurt. We also show the existence of regimes where a retailer can overcome cannibalization effects due to competition by improving consumer valuation. Our results indicate that before investing to improve consumer valuation firms should pay close attention to presence of a close competitor, direction of consumer flow, and market characteristics.

2 :: A Thorough Economic Evaluation of the ISO 2859:2006 Acceptance Sampling Schemes

George Nenes, Yiannis Nikolaidis

In order to carry out Acceptance Sampling, companies often use sampling schemes that are determined by easy-to-use quality standards. However, these standards do not take quality costs directly into account. This paper aims at evaluating economically the single-sampling plans recommended by the ISO 2859:2006. It has been initially motivated by the case of CRYSTAL S.A., a Greek company specialized in commercial refrigerators, which uses the ELOT 398.0 and 398.1 standards - essentially the Greek equivalent to the ISO 2859:2006 - in order to control the quality of the incoming raw materials.

3 :: Sales Force Compensation in Settings with Servitization

Zeynep Aksin, Taha Dogru, Evrim Didem Gunes

In this paper we focus on settings where customers are offered a bundle of products and services. For firms that offer such consumption systems, we model the sales force incentive design problem that explicitly accounts for customer's relationship with the firm and how the effort of the sales agent affects this relationship.

1 :: A Sequential Auction-Bargaining Procurement Model

Kun Soo Park, Woonghee Tim Huh

One of the most important decisions that a firm faces in managing its supply chain is a procurement decision: selecting suitable suppliers among many potential competing sellers and reducing the purchase cost. While both auction and bargaining have been extensively studied in the literature, the research that combines both auction and bargaining is limited. In this paper, we consider a combined auction-bargaining model in a setting where a single buyer procures an indivisible good from one of many competing sellers. The procurement model that we analyze is a sequential model consisting of the auction phase followed by the bargaining phase. In the auction phase, the sellers submit bids, and the seller with the lowest bid is selected as the winning bidder. In the bargaining phase, the buyer audits the cost of the winning seller and then negotiates with him to determine the final price. For this auction-bargaining model, we find a symmetric bidding strategy equilibrium for the sellers in a closed form, which is simple to understand and closely related to the classical results in the auction and bargaining literature. We also show that the auction-bargaining model generates at least as much profit to the buyer as the standard auction or sequential bargaining model.

2 :: Revenue Management with Price Bargaining

Nicola Secomandi, Atul Bhandari

We consider the revenue management problem of a seller who owns a finite inventory of a single product that can be sold during an infinite horizon when demand is Bernoulli. The extant literature assumes that the seller has all the negotiation ability. For example, this implies that the seller can enforce a posted, take it or leave it, price. We also consider the cases when this ability resides entirely with the buyers, is equally shared by the seller and the buyers, and negotiations are structured. We develop and analyze a Markov decision process that computes the seller's total discounted expected revenue under a given price bargaining mechanism. We apply this model and compare the seller's performance in each of these cases, both analytically and numerically. We also numerically investigate robustness issues related to misspecifying the seller's negotiation ability.

3 :: Generalized Cournot Competition for Multiple Products and Price of Anarchy

Jonathan Kluberg, Georgia Perakis

In this research we study a generalized model for Cournot oligopolistic competition. The main goal in this paper is to understand Cournot competition when firms are producing multiple products and are faced with a variety of constraints. We first study existence and uniqueness of Cournot equilibria under general constraints. We then compare the total profit under Cournot to the one obtained under a centralized setting, (i.e., where a single firm controls all the products and maximizes the total profit). Our goal is to understand how the presence of competition affects the overall profit in the system and what the key drivers of the inefficiencies that arise due to competition are.

1 :: Optimal Control of Inventory Systems with Multiple Types of Remanufacturable Products

Sean Zhou, Zhijie Tao, Xiuli Chao

Product returns have become a significant part for many manufacturing, distribution, and retail systems. Since products are returned under different operational conditions, they usually require different remanufacturing effort/costs. In this paper, we study a single-product periodic-review inventory system with multiple types of returns. The serviceable products used to fulfill stochastic customer demand can be either manufactured/ordered, or remanufactured from the returned products, and the objective is to minimize the expected total discounted cost over the planning horizon. We are interested in the optimal

manufacturing/remanufacturing policies for systems without or with possible disposals of returned products. We show that, under some circumstances but not all, the optimal policy has a simple form and can be completely characterized by a sequence of constant control parameters. However, in other scenarios, the optimal policy is quite complicated and its control parameters are state-dependent. As the state-dependent policy is difficult to compute and implement, we develop simple heuristic policies for managing the system. Numerical studies show that the heuristics are effective. This work is motivated by a project with a major energy company that manages its inventory through the options of ordering and remanufacturing returned products of various conditions.

2 :: Forecasting Trade-In Returns Based on Return Merchandise Authorization

Jiayi Li, Duncan Fong, Susan Xu

3 :: Transient vs. Steady State Profitability of Remanufacturing

Yue Jin, Yihao Lu, Ana Muriel

We investigate the profitability of offering remanufactured products for a monopoly firm in both transient and steady state settings, extending the work in Jin, Muriel and Lu (2007). We characterize a threshold for the remanufacturing cost below which it is optimal to offer remanufactured products, and focus on analyzing its dependence on the consumer profile. This analysis allows us to identify when it is customer segmentation, and not the cost difference, that drives the remanufacturing decision. That is, we find conditions under which it is optimal to offer remanufactured products even if they are as costly to produce as new products. Consequently, it becomes cost effective to substitute them and offer new products under the remanufactured label when the stream of returned, remanufacturable products is low.

TD03 :: Thursday, 4:00-5:30pm

Service Operations

Room 1511

1 :: Block Time Estimation and Robust Airline Scheduling

Yu-Ching Lee, Diego Klabjan, Milind Sohoni

Airline schedule planning divisions determine flight frequencies and preliminary departure times. We develop a robust optimization model that computes exact departure and flight block-times based on the provided stochastic estimates. Operating cost and passenger revenue are captured in the model. Service level based on connecting traffic and the DOT measure of on-time performance is also taken into account.

2 :: Airline Scheduled Time Competition and On-Time Performance

Yan Dong, Kefeng Xu

Airline service performance has become a focal point of the industry-wide competition. This study focuses on airline competition on price and scheduled flight time, as well as the competition's impact on airline on-time performance. It establishes a sequential game between airlines with respect to price and scheduled flight time, and discusses multiple equilibria of scheduled times under conditions of competition intensities and customer values of on-time performance. Hypotheses generated through the game theoretic model are tested empirically using airline data. This research contributes to literature in the area of customer relationship management. The findings provide implications to how customer acquisition and customer retention interact in airline competition.

3 :: "Information Boosting" Methodologies for Multiple Data Sources in Workforce Management

Jianying Hu, Yingdong Lu, Aleksandra Mojsilovic, Ana Radovanovic, Mark Squillante

We consider problems that arise due to the lack of compatibility and linkages among workforce data in enterprise resource planning (ERP) systems and propose an "information boosting" methodology that exploits various relationships in the data to detect errors and improve the quality of information for workforce forecasting and planning. We present a three-step

methodology to infer the missing information and correct errors in the existing data. We test the methodology using real-life ERP data from the IBM Corporation, and demonstrate the benefits of our information boosting approach.

TD04 :: Thursday, 4:00-5:30pm

Inventory Theory / Multi-echelon

Room 1520

1 :: New Policies for the Stochastic Inventory Control Problem with Two Supply Sources

Anshul Sheopuri, Ganesh Janakiraman, Sridhar Seshadri

We study an inventory system under periodic review in the presence of two suppliers (or delivery modes). The emergency supplier has a shorter lead-time than the regular supplier, but the unit price he offers is higher. Excess demand is backlogged. We show that the classical "Lost Sales inventory problem" is a special case of this problem. Then, we generalize the recently studied class of Dual Index policies (Veeraraghavan and Scheller-Wolf (2007)) by proposing two classes of policies. The first class consists of policies that have an order-up-to structure for the emergency supplier. We provide analytical results that are useful for determining optimal or near-optimal policies within this class. This analysis and the policies that we propose leverage the connections we make between our problem and the lost sales problem. The second class consists of policies that have an order-up-to structure for the combined orders of the two suppliers. Here, we derive bounds on the optimal order quantity from the emergency supplier, in any period, and use these bounds for finding effective policies within this class. Finally, we undertake an elaborate computational investigation to compare the performance of the policies we propose with that of Dual Index policies. One of our policies provides an average cost-saving of 1.1 % over the Best Dual Index policy and has the same computational requirements. Another policy that we propose has a cost performance similar to the Best Dual Index policy but its computational requirements are lower.

2 :: Inventories with Multiple Supply Sources and Networks of Queues with Overflow Bypasses

Jing-Sheng Song, Paul Zipkin

Consider an inventory system with multiple supply sources. The replenishment leadtimes from each source are stochastic, representing congestion and disruption. We develop performance evaluation tools for a family of reasonable order policies. These policies take into account real-time supply information, which can be obtained through tracking technologies such as GPS and RFID. Performance evaluation of such state-dependent policies is generally hard. The main thrust of the paper is to show that, under these policies, the supply system becomes a network of queues with a special routing mechanism called an overflow bypass. The solution has a simple product form. Thus, we obtain closed-form performance measures. These results reinterpret and extend Moinzadeh and Schmidt's analysis of a system with two sources having deterministic leadtimes.

3 :: Adaptive Data-Driven Inventory Control Policies Based on the Kaplan-Meier Estimator

Tim Huh, Retsef Levi, James B. Orlin, Paat Rusmevichientong

Using the well-known Kaplan-Meier estimator from statistics, we propose a new class of nonparametric adaptive data-driven policies for stochastic inventory control problems. We focus on the distribution-free newsvendor model with censored demands. The assumption is that the demand distribution is not known and there is only past sales data available. We study the theoretical performance of the new policies and show that for discrete demand distributions they converge almost surely to the set of optimal solutions. Extensive computational experiments suggest that the new policies converge for general demand distributions, and demonstrate that they outperform previously known policies. As a byproduct of the theoretical analysis, we obtain new results on the asymptotic consistency of the Kaplan-Meier estimator that significantly extend existing work in statistics. To the best of our knowledge, this is the first paper to apply the Kaplan-Meier estimator within an adaptive optimization algorithm, and it is the first application to stochastic inventory control models. We believe that this work will lead to additional applications in other domains.

1 :: Inventory Management for an Assembly System Subject to Component Supplier Disruptions

Gregory DeCroix

We analyze a multi-component assembly system where one or more of the components is subject to random supply disruptions. We establish structural properties of the optimal ordering policy, and fully characterize the policy in a special case. For the general case we propose a heuristic policy that performs well in numerical trials. We use that policy and/or the optimal policy (when it can be computed) to numerically explore how various aspects of system structure affect the cost impact of supply disruptions.

2 :: Bargaining Power and Supply Base Diversification

Zhixi Wan, Damian Beil

We examine a supply base diversification problem faced by a buyer who periodically holds auctions to allocate short-term supply contracts among a cohort of suppliers. To mitigate significant cost shocks to procurement, the buyer can diversify her supply base by selecting suppliers from different regions. We find that the optimal degree of supply base diversification depends on the buyer's bargaining power in the supply chain. At one extreme, when the buyer has full bargaining power and thus can dictatorially implement the optimal mechanism, she prefers to fully diversify. At the other extreme, when the buyer uses an open-descending auction due to her lack of bargaining power, she must consider protecting herself against potential predatory price escalation from powerful suppliers; surprisingly, she does so precisely when the likelihood and scale of cost shocks are large.

3 :: Multi-period Procurement Strategies in the Presence of Unreliable Suppliers

Awi Federgruen, Nan Yang

We propose and analyze a general periodic review model, in which the firm has access to a set of potential suppliers, each with specific yield and price characteristics. Assuming unsatisfied demand is backlogged, the firm incurs, as in standard inventory models, three types of costs: (i) procurement costs; (ii) inventory carrying costs for units carried over from one period to the next and (iii) backlogging costs. However, in contrast to standard inventory models, a procurement strategy requires the specification, in each period, of (i) the set of suppliers to be retained, (ii) their respective shares in this period's replenishments as well as (iii) the traditional aggregate order placed (among the various suppliers). We develop an efficient algorithm to determine, in each period, as a function of the prevailing inventory, the (unique) optimal set of suppliers, the aggregate order as well as the suppliers' shares in this order. While our initial analysis pertains to a finite planning horizon, we show that our results and algorithm can be extended to an infinite horizon setting: a stationary policy is optimal and its structure is identical to that obtained in the finite horizon case. Most importantly, we derive many structural results which generate important managerial insights.

1 :: Optimal Structural Policies for Ambiguity and Risk Averse Inventory and Pricing Models

Xin Chen, Peng Sun

This paper discusses multi-period stochastic joint inventory and pricing models when the decision maker is risk and ambiguity averse. We present infinite horizon models with discounted and long run average optimization criteria. The main result of this paper is to establish the optimality of stationary $(s; S; p)$ policies for the infinite horizon inventory and pricing models, for which the existing proof techniques for the risk neutral counterparts may not be extended.

2 :: Optimal Market-Making with Risk Aversion

Miao Song, Kan Huang, David Simchi-Levi

Market-makers have the obligation to trade fixed amounts of securities at quoted bid or ask prices, and their inventories are exposed to the potential loss when the market price moves in an undesirable direction. One approach to reduce the risk brought by price uncertainty is to adjust the inventory at the price of losing potential spread gain. Using stochastic dynamic programming, we show that a threshold inventory control policy is optimal with respect to an exponential utility criterion, and more general results are obtained for mean-variance analysis.

3 :: Invest or Exit? Optimal Decisions in the Face of a Declining Profit Stream

Dharma Kwon

Even in the face of deteriorating and highly volatile demand, firms often invest in rather than discard aging technologies. In order to study this phenomenon, we model the firm's profit stream as a Brownian motion with negative drift. At each point in time, the firm can continue operations, or it can stop and exit the project. In addition, there is a one-time option to make an investment which boosts the project's profit rate. Using stochastic calculus, we show that the optimal policy is characterized by three thresholds. There are investment and exit thresholds before investment, and there is a threshold for exit after investment. We also affect a comparative statics analysis of the thresholds with respect to the drift and the volatility of the Brownian motion. When the profit boost upon investment is sufficiently large, we find a counter-intuitive result: an increase in volatility induces the firm to invest earlier.

FA01 :: Friday, 8:00-9:30am

Revenue Management

Room 1505

1 :: Pricing as a Hedge against Demand Uncertainty

Guillaume Roels

In this paper, we investigate the role of pricing in the newsvendor model as a way to cope with demand uncertainty, as well as the relative efficiency of ignoring demand uncertainty in the pricing decision. In contrast to previous research, we find that charging a positive risk premium is generally optimal, at least when the demand coefficient of variation is large, when demand is restricted to be nonnegative. Moreover, we demonstrate that, under an additive (multiplicative) demand model, ignoring demand uncertainty in the pricing decision can be very inefficient (is relatively efficient), even though the riskless price is close to (can be very different than) the optimal price.

2 :: Product Design with Threshold Preferences

Matulya Bansal, Costis Maglaras

We study the product design problem of a revenue-maximizing monopolist firm that serves a market where customers are heterogeneous with respect to their valuations and desire for a quality attribute, and are characterized by a perhaps novel model of customer choice behavior. Specifically, instead of optimizing the net utility that results from an appropriate combination of prices and quality levels, customers are "satisficers" in that they seek to buy the cheapest product with quality above a certain customer-specific threshold. This model dates back to Simon's work in the 1950's (Simon 1955) and can be thought of as a model of bounded rationality for customer choice. We characterize the structural properties & identify the optimal product menu for this model.

3 :: "We Will be Right with You": Managing Customers with Vague Promises

Gad Allon, Achal Bassamboo, Itai Gurvich

Delay announcements informing customers about anticipated service delays are prevalent in service-oriented systems. Which delay announcements a service provider should use is a complex question, and its answer depends on both the dynamics of the underlying queueing system and on the customer behavior. We examine this problem of information communication by

considering a model in which both the firm and the customers act strategically: the firm in choosing its delay announcement, and the customers in interpreting these announcements and in making the decision about when to join the system and when to balk. We characterize the equilibrium language that emerges between the service provider and her customers. The analysis of the emerging equilibria provides new and interesting insights into customer-firm information sharing. We show that even though the information provided to customers is non-verifiable and non-credible, it improves the profits of the firm and the expected utility of the customers. Further, the information could be as simple as “High Congestion”/“Low Congestion” announcements, or could be as detailed as the true state of the system. We also show that firms may choose to shade some of the truth by using intentional vagueness to lure customers.

FA02 :: Friday, 8:00-9:30am

Sustainable Operations

Room 1524

1 :: Product Introduction and Retail Pricing Strategies for Durable Products in Decentralized Channels

Shuya Yin, Saibal Ray, Haresh Gurnani

This paper studies a decentralized manufacturer-retailer channel dealing with a durable product and renewable customers. We pinpoint the roles played by secondary and tertiary markets as well as advance information sharing (individually and jointly) on the manufacturer's new model introduction strategy, retail prices and profits in such a channel. While our model captures different types of durable products, the focus is specifically on the textbook industry. In that context, we are able to analyze the underlying reasons behind the increased frequency of new edition releases and rising prices for new textbooks.

2 :: An Empirical Investigation of Environmental Announcements and the Market Value of the Firm

Brian Jacobs, Vinod Singhal, Ravi Subramanian

Researchers have studied the relationship between firm environmental performance and financial performance using a variety of measures and methods, often with inconclusive results. We use event study methodology to estimate the average abnormal return of firm stock price after print media announcements relating to good environmental performance. After an extensive review of recent announcements, we categorize them into five major types -- awards, certifications, initiatives, resource conservation efforts, and products -- and examine the market reaction to each type. We find that the market does not react to announcements of environmental performance in such broad categorizations. Rather, only announcements relating to objective third party recognitions at the national or international level result in a detectable market reaction. Specifically, announcements of awards for environmental performance granted by a federal agency and reports of ISO 14001 certification result in a statistically significant 0.80% average increase in stock price.

3 :: Adoption of Voluntary Environmental Standards

Charles Corbett, Suresh Muthulingam

We examine the role of signaling and of intrinsic benefits in the adoption of the individual elements of the voluntary LEED (Leadership in Energy and Environmental Design) standards for green buildings. We use goodness-of-fit tests on data for all 442 LEED certified buildings and find that neither signaling nor pursuit of intrinsic benefits can independently explain the observed adoption pattern, but that a combination of the two factors can. We also find tentative evidence that the adoption decision is made sequentially: organizations first choose a level of certification (consistent with signaling), and then choose how many LEED elements to adopt given their chosen level of certification (consistent with pursuing intrinsic benefits). We relate our findings to some open questions in the literature on diffusion of technology and draw implications for the design and the future development of similar voluntary standards and eco-labels.

1 :: Solving Continuous Replenishment Inventory Routing Problems

Samuel Fomundam, Jeffrey Herrmann

This research is motivated by work with public health officials who must plan the logistics for resupplying points of dispensing (PODs), which will dispense medications to the public in case of a bioterrorist attack such as anthrax. After receiving an initial but limited supply of medication, the PODs will operate continuously. Vehicles will resupply the PODs continuously from a central depot that has a stockpile of medication. Each vehicle repeatedly follows the same route, starting out as soon as it can after returning to the depot. At each site, the vehicle delivers enough medication to replace what was consumed since the last visit. This paper formulates this problem, which is related to vehicle routing, identifies a special case, and presents the results of testing heuristics for the general case.

2 :: The Consistent Vehicle Routing Problem

Chris Groer, Bruce Golden, Edward Wasil

In the small package shipping industry, companies try to differentiate themselves by providing high levels of customer service. This can be accomplished in several ways including on-line tracking of packages, ensuring on-time delivery, and offering residential pick ups. Some companies want their drivers to develop relationships with customers on a route and would like the same drivers to visit the same customers at roughly the same time on each day that the customers need service. These service requirements together with traditional constraints on vehicle capacity and route length define a variant of the classical capacitated vehicle routing problem (VRP) that we call the Consistent VRP. In this paper, we formulate the problem as a mixed integer program (MIP) and develop a metaheuristic algorithm to solve the Consistent VRP. We compare the performance of our algorithm to the optimal MIP solutions for a set of small problems and then apply our algorithm to several large simulated data sets and a real-world data set with more than 3,700 customers. The solutions produced by our algorithm on all problem sets do a very good job of meeting customer service objectives with routes that have a low total travel time.

3 :: A Savings-based Shipment Consolidation Algorithm

Donald Warsing, Michael Kay

We present a method for constructing shipment routes that minimize the sum of inventory holding and freight transportation costs. Considering transportation cost as a joint ordering cost, this problem can be viewed as a joint replenishment problem (JRP). We formulate and solve a logistically-realistic version of the JRP, allocating a set of items efficiently across transportation vehicles, including the possibility that some items might be better left in singleton sets, shipping by themselves on container-load or less-than-container-load modes. We formulate transportation costs in a flexible, realistic, and accurate way, and we utilize a savings-based algorithm to build routes that reduce the joint costs of transportation across the set of items to be shipped, while also seeking to reduce system inventory holding costs. Computational analysis demonstrates that the consolidation solution can generate dramatic improvements in both transportation and inventory holding costs versus shipping the items independently and versus more traditional distance-only approaches to finding good shipment routes. Our algorithm can also be applied directly to a more general problem, where the shipments have arbitrary origin and destination points, thus demonstrating the applicability of our solution algorithm beyond vehicle-routing-type problems.

1 :: Inventory Management of Perishables: New Results

Itir Karaesmen, Alan Scheller-Wolf, Borga Deniz, Ismail Civelek

We consider a discrete-time inventory model for perishable goods having separate demand streams for items of different ages. We study the effect of substitution under two practical replenishment policies: replenishing inventory according to order-up-to level policies based on either (i) total inventory in system (TIS) or (ii) new items only (NIS). Given these policies for an item with two-periods of lifetime, we distinguish four different levels of substitution: (1) demand for an item can only be satisfied by an item of that age (No-Substitution); (2) demand for new items can only be satisfied by new ones, but excess demand for old items can be satisfied by new (Downward-Substitution); (3) demand for old items can only be satisfied by old, but excess demand for new items can be satisfied by old (Upward-Substitution); (4) both downward and upward substitution are employed (Full-Substitution). These four models are sufficient to study inventory issuance rules such as FIFO and LIFO, and other practical prioritization rules to satisfy the demand. We compare the four substitution models analytically in terms of their infinite horizon time-average costs, providing conditions on cost parameters, including substitution costs, that determine when (if at all) one substitution model is more profitable than the others, given that only a proportion of customers accept substitutes. We show that substitution can be economically beneficial even when substitution costs are very high, or economically harmful even when substitution costs are zero when inventory is replenished using TIS. In contrast, our results on NIS are more intuitive. We complement our results with computational experiments exploring when a particular substitution-model replenishment-policy pair performs well. Our computations indicate that, contrary to earlier literature, NIS is in general superior to TIS, when new items are more valuable than old and the demand for new items is not negligible.

2 :: Inventory Management under Price Protection

Karthik Sourirajan, Roman Kapuscinski, Markus Ettl

In high-tech industries where technological obsolescence is high, price protection has become a standard element of contracts between manufacturers and distributors. While price protection was intended to give distributors sufficient incentives to stock enough inventory, empirical evidence suggests that price protection often leads to overstocking. We analyze acceptable mechanisms to reduce excess inventory caused by price protection.

3 :: Optimization of Stock Levels for Service Tool Inventory

Ingrid Vliegen, Ana Busic, Alan Scheller-Wolf, Geert-Jan van Houtum

Many Original Equipment Manufacturers not only sell machines to their customers, but also service contracts in which a certain system availability level is agreed upon. To ensure this availability level is met, a sufficiently fast provisioning of spare parts, service engineers, and service tools is required when a machine failure occurs. In this talk, we focus on the planning problem for service tools, and base ourselves on an actual service tools planning problem of an OEM in the semiconductor supplier industry. Their investment in service tools is in the order of tens of millions of euros, while currently no planning method is available that takes into account the specific characteristics of service tools. Therefore, we consider a single-location, multi-item inventory model for service tools. Multiple service tools are kept on stock at a local warehouse. Independent Poisson demand streams arrive requesting different sets of tools. Those tools from a requested set that are in stock are immediately released; they are "in use" for an exponential amount of time, after which they are returned together. Requested tools that are not in stock are delivered via an emergency channel (e.g., from a central warehouse); they may be considered as lost sales for the warehouse under consideration. Thus, our model is a partial order service model featuring coupled demands and coupled returns - sets of tools are released and returned together. Our objective is to minimize the investment in service tools, subject to meeting the service level agreed upon for the aggregate order fill rate, i.e., the percentage of demands for which all requested tools are delivered from stock. The aim of this paper is to provide a heuristic to determine near-optimal stock levels for this optimization problem. We develop two heuristics, and compare our heuristics with a Lagrangian lower bound on the costs, as well as with two other heuristics, including the one currently in use at the OEM, on the basis of costs, service level and running time of the heuristics. Our heuristics outperform the others on cost, while the service levels are met, but the running times are high. The other heuristics have lower running times, but lead to inferior solutions, either because the solutions are not feasible or because the costs are too high.

1 :: Reverse Auction Procurement with Flexible Noncompetitive Contracts

Sean Zhou, Zhijie Tao, Nianbing Zhang, Gangshu Cai

More and more companies nowadays are using reverse auctions to procure products and services. We consider the case of a buyer who wants to procure multiple units of a certain product from a group of potential suppliers. The buyer adopts a hybrid procurement mechanism, in which she first offers some flexible noncompetitive contracts to suppliers followed by a reverse auction. Each supplier given the contract can accept the offer, so he does not need to participate in the auction but commits to selling a unit to the buyer. The contract is flexible in the sense that the selected supplier is allowed to join the auction if he turns down the offer. This model is an extension of the model proposed by Engelbrecht-Wiggans and Katok (2006), in which they prohibit the selected suppliers from joining the auction if they turn down the offer. Our extension yields some interesting results and insights. We show that, in some cases, the suppliers prefer accepting the flexible noncompetitive offers to declining them and joining the auction. Also, the buyer can benefit from offering such kind of contracts to suppliers as opposed to just conducting a regular reverse auction or offering the noncompetitive contracts proposed by Engelbrecht-Wiggans and Katok (2006). More importantly, when the buyer makes multiple offers, we analyze the game behavior of the selected suppliers and characterize the Nash equilibrium. These analyses enable us to further provide some guidelines on the optimal number of contracts that the buyer should offer.

2 :: When is It Cheaper to Outsource Integration

Qing Ye, Izak Duenyas, Roman Kapuscinski

In this paper we model a system in which manufacturer purchases multiple components and then integrates them into an end product. For each component, there are multiple suppliers capable of providing it. The manufacturer considers outsourcing the assembly to one of his suppliers, rather than integrating the final product in-house. Since the production and integration costs are not publicly known, auctions are used to reward those contracts. In the first stage, the manufacturer uses first-price sealed-bid reverse auction to determine which supplier will become the integrator. Then, the winner of the first auction, i.e., the integrator, uses reverse auctions to procure the required components from the subset of the suppliers in the first auction. In order to evaluate the manufacturer's performance, we analyze the suppliers' bidding strategies in both stages. An explicit solution is derived for a canonical case where the suppliers' cost distributions are uniform and the integration costs are linear. We compare the manufacturer's total cost when manufacturer procures components by himself and integrates them in-house versus when integration is auctioned. Managerial insights are provided to help the manufacturer choose his auctioning strategy.

3 :: A Differential Pricing Mechanism for Vertical Information Sharing under Horizontal Competition

Aditya Jain, Sridhar Seshadri, Milind Sohoni

In this paper we consider a two-echelon supply chain in which downstream retailers engage in Cournot competition. Each retailer is endowed with private information about the uncertain demand intercept. The manufacturer is the Stackelberg leader who sets the contract terms with the retailers, and benefits from retailers sharing their private information. Previous work has shown that a common price set by the manufacturer leaks the shared information to all retailer, thus hurting retailers under competition. We show that a retailer can benefit from such information leakage by manipulating demand beliefs of other retailers, and thus has an incentive to understate his demand information in any information sharing agreement. We then derive a wholesale pricing scheme that alleviates this problem, i.e., enforces truthful revelation. This pricing scheme rewards (penalizes) a retailer for providing optimistic (pessimistic) information by allowing the manufacturer to price discriminate among retailers based on their shared information. We show that under such a pricing scheme a retailer is better off by sharing information in all circumstances, i.e., information sharing is a dominant strategy for all retailers. Furthermore, the manufacturer's profit, as well as the supply chain surplus, increase when all the retailers share information. Finally, we

argue that the dominant symmetric Nash equilibrium is characterized by all retailers revealing their information under the proposed linear contracting mechanism.

FA06 :: Friday, 8:00-9:30am

Operations / Finance Interface

Room 1518

1 :: Inventory Control and Risk Management of Natural Gas Storage Assets

Owen Wu, Derek Wang, Jacky Qin

Managing a natural gas storage asset involves injection and withdrawal of natural gas and risk management via trading on spot and futures markets. The objective is to shape the probability distribution of end-of-winter profit, so as to balance the down-side risk and up-side profit. The firm takes profit not only from the winter-to-summer futures price differentials but also from the spot-futures price differentials due to higher spot market volatility. Physical constraints are also present: injection and withdrawal of natural gas are subject to the storage capacity constraint, injection/withdrawal rate constraint, and the delivery schedule constraint. In this paper, we analyze a model that captures all the above essential features. We study the structure of the optimal inventory control and trading strategy, and also construct heuristic policies that is numerically shown to be near-optimal.

2 :: Integrated Optimization of Procurement, Processing and Trade of Commodities in a Network Environment

Sripad Devalkar, Ravi Anupindi, Amitabh Sinha

We consider an integrated optimization problem for a firm involved in procurement, processing and trade of commodities in a network environment. We derive optimal policies for a risk-neutral firm, when the processed commodity(ies) are sold using futures instruments. When the firm has unlimited processing capacity, we find that an 'all or nothing' procurement policy is optimal and it is optimal to postpone all processing till the last possible period. With finite processing capacity, the optimal procurement quantity is dependent on the starting inventory levels and it might be optimal for the firm to process earlier in the horizon. We also develop heuristics for computing near optimal policies and develop an upper bound for the value function using the technique of Lagrangian relaxation. We conduct numerical studies to quantify the performance of the heuristics with respect to the upper bound.

3 :: Analysis of Energy Bridge Technology in the Liquefied-Natural-Gas Industry

Erkut Sonmez, Nicola Secomandi, Alan Scheller-Wolf, Sunder Kekre

Liquefied natural gas (LNG) is natural gas that has been converted to liquid form at -260F for efficient storage and transport over long distances. According to the U.S. Energy Information Administration's forecasts (EIA 2006), local gas production will be soon unable to meet demand in industrialized countries; imported LNG will play an important role to close this gap. Thus, increasing LNG imports will require additional re-gasification capacity. There are two types of LNG import (i.e. re-gasification) terminals: onshore and offshore. Energy Bridge (EB) is a recently developed offshore LNG re-gasification and delivery technology developed by Excelebrate Energy. The major advantage of an EB is that it does not require costly onshore re-gasification terminals. Despite EB's advantages, it has yet to be widely adopted and used by the LNG industry due to its low LNG throughput rate. In order to cope with this drawback, Excelebrate Energy has developed and tested LNG transshipment technology for ship-to-ship LNG transfer. The use of transshipping could enable greater utilization of an EB Port, enhancing overall efficiency and throughput while reducing costs. In this extended abstract, we analyze the leverage effect of the transshipment business model on EB technology.

FB01 :: Friday, 10:00-11:30am

Operations / Marketing Interface

Room 1505

1 :: The Impact of Counterfeits on Brand Name Products

Jie Zhang, Jeff Hong, Rachel Zhang

This paper investigates the impact of counterfeits on the price, market share and profitability of brand name products and the strategies for brand name companies to fight counterfeiting. We adopt the multinomial logit model which takes into consideration uncertain consumer purchasing behavior and allows consumers the choices of purchasing an authentic product, purchasing a fake, or not buying. We first study how non-deceptive counterfeiting products, which customers know at time of purchase that the products are counterfeits with certainty, affect a brand name product in monopoly markets. We then define the "competency level" of a product, authentic or counterfeiting, which directs us to different strategies for fighting counterfeiting in monopoly markets, competitive markets where there are multiple brand name companies with competing products that are all victims of counterfeiters, and nonhomogeneous markets where there are several consumer groups each with different utility towards an authentic product. We also discuss fighting strategies for brand name companies facing deceptive counterfeits.

2 :: Pricing and Production Planning under Supply and Demand Uncertainty

Burak Kazaz

This paper provides a comprehensive review of the modeling approaches that can be used in order to study the interactions between operational and marketing decisions. It addresses the problem of determining the sale price and the production quantity under supply and demand uncertainty. The problem is commonly observed in the agricultural industry. We consider a firm that initially leases farm space in order to grow fruit. However, the realized amount of fruit supply fluctuates due to weather conditions, diseases, etc. At the end of the growing season, the firm has an opportunity to purchase an additional supply from other growers at a unit cost that changes with the realized supply. Specifically, the lower the realized supply, the higher the unit purchasing cost. We term this as the yield-dependent purchasing cost. The firm then makes two production decisions under demand uncertainty, and determines the optimal amount of the realized supply to be converted to finished product, and the amount of additional supplies to be obtained from the spot market. Two modeling approaches are presented for the price-setting behavior. In the Early Pricing model, the sale price is determined when the leasing agreement is made, and in the Postponed Pricing model, it is set after observing the realized supply. The yield-dependent purchasing cost differentiates this work from studies that feature price-setting behavior under both supply and demand uncertainty. In traditional models where the unit purchasing cost is a constant, the firm experiences either the case of "backlogging" (all demand is satisfied) or "lost sales" (some of the demand is not satisfied), but not both. The yield-dependent cost, however, forces the firm to experience both cases simultaneously. The paper identifies how the results regarding the optimal sale price and production decisions change with backlogging and lost sales independently, and later with the yield-dependent cost. It proves that the optimal stocking level is a constant in the case of backlogging, and decreases with the sale price in the case of lost sales. The optimal sale price, on the other hand, increases in the stocking factor in the case of backlogging, and decreases in the case of lost sales. Under the presence of the yield-dependent purchasing cost, the optimal stocking factor is greater than or equal to that of the complete backlogging and lost sales, and the optimal sale is in between the values for the cases of complete backlogging and lost sales.

3 :: Combined Replenishment and Pricing Policies for a Perishable Product with Age-Dependent Demand

Li-Ming Chen, Amar Sapra

In this paper, we consider combined replenishment and pricing strategies for a perishable product whose demand depends on both price and age. We develop three periodic review models with increasing flexibility in decision making to identify the optimal replenishment frequency, order quantity and pricing strategy. In the first model, we fix the replenishment interval and determine price in each period. In the second model, we allow the possibility of order being placed every period. In the third model, we further allow partial salvage of inventory if no order is being placed and if on-hand inventory is excessive. In addition to characterizing the form of replenishment and pricing strategies, we conduct numerical experiments to determine the relative benefit of flexibility in order placement decisions and partial salvage of inventory.

1 :: Flexibility in Dynamic Processing Systems: Optimality of Tailored Chaining without Complete Resource Pooling

Achal Bassamboo, Ramandeep Randhawa, Jan van Mieghem

In this paper, we study the role of flexibility in dynamic settings that arise in processing or service systems. Such systems typically serve or process different jobs (or customers) that require different modes of service. Deciding the amount of capacity to install is a key decision that the system manager needs to make. The manager typically has at his disposal different types of capacity. In particular, there is dedicated capacity that consists of resources which have the ability to serve/process only one kind of jobs; and flexible capacity which can serve different kinds of jobs. We show that an optimal solution does not exhibit the complete resource pooling phenomenon. We further demonstrate the optimality of tailored flexibility solutions that use dedicated capacity to serve the base or fairly certain demand, and flexible resources to help with the variable demand.

2 :: Is Inventory's Fiscal Year End Effect Caused by Sales Timing? A Test Using a Natural Experiment from Germany

Richard Lai

We first document inventory's fiscal year end (FYE) effect, in which inventory levels are especially low not just at calendar year end, but also at fiscal year end. Using a panel of all inventory-holding U.S. firms in COMPUSTAT, we find inventory is 10% lower in the fourth fiscal quarter than other fiscal quarters, controlling for calendar quarters. In aggregate, the 10% is a staggering \$47 billion, valued at cost of goods. One cause of the FYE effect might be sales timing, in which managerial incentives lead firms to pull some next-quarter sales into the fourth quarters of the fiscal year. This is not an obvious cause; the literature suggests three alternative explanations. To test for sales timing, we employ a novel natural experiment based on Germany's tax code change in 2000, when some firms change their fiscal year ends in a way that is plausibly exogenous to inventory patterns. We find that these German firms also have lower inventory at FYE. This result is robust to corrections for possible treatment selection using the Heckit procedure and propensity scoring. We also examine mediator and moderator effects. Taken together, the evidence is consistent with sales timing, but is not explained by the alternative hypotheses. Finally, we estimate that 1 percentage point lower inventory at FYE is associated with an economically significant 1.7% lower Tobin's q, which could be due to lower gross profits and higher operating costs like inventory holdings expenses and capacity investments. We conclude with a discussion of limitations, next steps, and some intriguing implications for research and practice.

3 :: The Impact of International Trade on Inventory Levels - Evidence from U.S. Manufacturers

Chaodong Han, Yan Dong, Martin Dresner

With increased global operations and the extended supply chains that result, inventory theory predicts that inventory levels should rise due to a greater need for safety and cycle stock. However, there appears to be an "inventory paradox" as empirical research has documented an overall decline in inventory levels for most U.S. manufacturing sectors. Based on location theory, we argue that international trade provides a viable channel for firms to efficiently allocate inventories along global supply chains, and may be contributing to the decline in domestic U.S. inventories. Using a panel of 19 three-digit NAICS industry categories, including manufacturing plant-based data for 1997-2005, our research provides supporting evidence for the impact of international trade on the declining trend in U.S. domestic manufacturing inventories.

1 :: A Multi-stage Decomposition Heuristic for the Container Stowage Problem

Mehmet Gumus, Philip Kaminsky, Erik Tiemroth, Mehmet Ayik

A significant and growing volume of global trade travels on container ships, and each year larger ships are introduced. For this reason, the Container Stowage Problem, a fundamental problem in marine cargo transportation involving the optimal

assignment of shipping containers of various types to specific storage locations in container ship at each port in order to maximize loading and unloading efficiency and minimize shipping costs, is growing in importance. In spite of this, the problem is currently solved in an ad-hoc manner based on experience and rules-of thumb. In this paper, we develop a multi-stage decomposition heuristic for this problem that accounts for the many complexities of real-world versions of the problem. This approach is performing well in testing on real-life problems, and is currently being developed for inclusion in commercial software.

2 :: Dynamic Supply Routing for a Make-To-Order Manufacturing Network

Jeremie Gallien, John Foreman, Julie Alspaugh, Chee Chong Teo, Rohit Bhatnagar

Dell's supply chain for desktops involves Asian vendors shipping components by sea to several U.S. plants. While suppliers are responsible for shipping enough inventory, Dell can re-route and expedite their shipments while in transit to balance inventory across sites and dynamically track geographic demand patterns. We present optimization models developed with Dell to address this routing challenge and some insights gained from their implementation.

3 :: Integrated Production and Outbound Distribution Scheduling in a Supply Chain: Research Overview and Future Directions

Zhi-Long Chen

In many applications involving make-to-order or time-sensitive (e.g., perishable, seasonal) products, finished orders are often delivered to customers immediately or shortly after the production. Consequently, there is little or no finished product inventory in the supply chain such that production and outbound distribution are very intimately linked and must be scheduled jointly in order to achieve a desired on-time delivery performance at minimum total cost. Research on integrated scheduling models of production and outbound distribution is relatively recent, but is growing very rapidly. In this paper, we provide a survey of such existing models. We present a unified model representation scheme, classify existing models into several different classes, and for each class of the models give an overview of the optimality properties, computational tractability, and solution algorithms for the various problems studied in the literature. We clarify the tractability of some open problems left in the literature and some new problems by providing intractability proofs or polynomial-time exact algorithms. We also identify several problem areas and issues for future research.

FB04 :: Friday, 10:00-11:30am

Inventory Theory / Multi-echelon

Room 1520

1 :: Optimal Expediting Policies for an Inventory System with Stochastic Lead Time under Radio Frequency Identification

Chiwon Kim, Diego Klabjan, David Simchi-Levi

Radio frequency identification is a nascent technology that allows more efficient tracking of goods and thus visibility into the supply chain. Many discussions revolve around the deployment costs and the resulting return on investment. New business processes hinging on radio frequency identification can improve the value proposition of this technology. We consider a serial supply chain with a radio frequency deployment at each installation and an expediting option from an installation to the manufacturing facility. The goods move stochastically among the installations and the system faces a stochastic demand at the manufacturer. Radio frequency technology allows to capture the state of the system, i.e., the time and location of goods, at any point in time. We identify systems that yield simple and tractable optimal policies. Both regular ordering and expediting follow a variant of the base stock policy. We provide some sufficient and necessary conditions to identify such systems. In addition, important managerial insights linking expediting costs and base stock levels are also provided.

2 :: Evaluation of Supply Chains with Inventory Record Inaccuracy and Implications on IT Investments

Gurhan Kok, Kevin Shang

Inventory record inaccuracy affects supply chain performance because it leads to ineffective replenishment decisions. Conducting cycle counts is a common approach to correct inventory record. It is not clear, however, how inaccuracy at

different locations impacts the supply-chain performance and how a cycle-count program should be designed. This paper aims to answer these questions by considering two basic supply chains: series and distribution systems. Each location implements a local base-stock policy. We assume that a random loss occurs on the physical inventory at each location in each period, such as shrinkage. The loss is not observed by the information system that carries inventory records. The discrepancy between the physical inventory and the inventory record accumulates until a location performs a cycle count. We first provide a recursion to evaluate the average supply chain cost per period. For fixed cycle-count policies, we construct an algorithm to compute the best local base-stock levels. These analytical results are used to examine the system performance under different count cycles. Our study suggests that it is more effective to conduct a cycle count at downstream stages in general. It also suggests that cycle counts should be conducted at a stage with short lead time, high holding cost rate, and low backorder cost rate. Our research provides insights on choosing locations to install RFID systems, or the choke point selection, from the perspective of reducing inventory record inaccuracy.

3 :: Information-Sensitive Replenishment when Inventory Records are Inaccurate

Adam Mersereau

In a lost sales inventory system with record inaccuracy, we consider replenishment policies based on a “Bayesian Inventory Record” (BIR). We formulate the problem of optimal BIR-based replenishment as a partially observed Markov decision process (POMDP). We analyze simple one- and two-period versions of the problem, and we implement an approximate POMDP algorithm for longer horizons. We identify an “information effect” that incentivizes a forward-looking inventory manager to stock less than a myopic one so as to improve information content for future decisions. Our research connects with research on inventory systems with unknown demand distributions, where the analogous result is to stock more.

FB05 :: Friday, 10:00-11:30am

Supply Chain Management

Room 1528

1 :: Procurement Strategies in a Three-Tier Supply Chain With Private Demand Information

Dimitris Kostamis, Izak Duenyas

We consider a three-tier supply chain with one OEM, one tier-1 and one tier-2 supplier. The OEM has private demand information, but no market power over the suppliers, and can choose between contracting with both suppliers and contracting with the tier-1 supplier only. We find that the OEM is better off contracting only with the tier-1 supplier if the suppliers employ full price-quantity schedules. On the other hand, we show that the OEM sometimes is better off contracting with both suppliers if they restrict themselves to direct mechanisms (i.e., they only offer as many contracts as the number of possible OEM demand types).

2 :: Delegation vs. Control of Component Procurement

Enis Kayis, Feryal Erhun, Erica Plambeck

A manufacturer must choose to delegate component procurement to its tier-1 supplier, or to control component procurement by contracting with both the tier-1 supplier and the tier-2 component supplier. Both suppliers have private cost information and the manufacturer has an alternative source of supply with cost known only to herself. This paper validates that if the firms may use arbitrarily complex contracts, then the manufacturer has the same expected profit with delegation as with control of component procurement. If, however, the firms use simple price-only contracts, then the manufacturer may achieve strictly greater expected profit with either delegation or control, depending upon the price of the alternative source, the selling price of the end product and, most importantly, what the manufacturer knows about her suppliers' costs. A numerical study shows that over a wide range of conditions, if the manufacturer chooses delegation versus control correctly, then she achieves nearly as much expected profit with price-only contracts as with the complex optimal contracts. However, the study identifies conditions under which price-only contracts perform poorly, and also shows that the loss may be extremely high when the manufacturer chooses delegation versus control incorrectly. Other factors, such as different information as well as supply chain structures (serial versus assembly), also affect the manufacturer's delegation versus control decisions under price-only contracts.

3 :: Buying from the Babbling Newsvendor: Availability Information and Cheap Talk

Gad Allon, Achal Bassamboo

Provision of real-time information by a firm to its customers has become prevalent in recent years in both the service and retail sectors. In this paper, we study a retail operations model where customers are strategic in both their actions and in the way they interpret information, while the retailer is strategic in the way it provides information. This paper focuses on the ability (or the lack thereof) to communicate credibly unverifiable information. We develop a game-theoretic framework to study this type of communication and discuss the equilibrium language emerging between the retailer and its customers. We show that for a single-retailer setting, the equilibrium language that emerges carries no information. In this sense, a single-retailer providing information on its own cannot create any credibility with the customers. We explore several remedies so that the firm can credibly disclose availability information to its customers. While in these remedies we show that the firm may be able to reveal complete information, the firm would prefer to shade some information and use intentional vagueness.

FB06 :: Friday, 10:00-11:30am

Inventory Theory / Multi-echelon

Room 1518

1 :: On Market Demand Mis-specification in a Two-level Supply Chain

Takamichi Hosoda, Stephen Disney

We investigate the impact of market demand process mis-specification on a serially linked two-level supply chain. Box-Jenkins models are used to represent both the true market demand process and a mis-specified model of the market demand process. Our analysis suggests that mis-specification does not always result in additional costs. Indeed by exploiting mis-specification, the supply chain can reduce both the total supply chain inventory and the production costs simultaneously. A managerial insight is revealed: Poor forecast accuracy is not always bad for the supply chain costs, from a total cost viewpoint. In other words, employing more accurate forecasting methods may actually result in higher total supply chain costs.

2 :: The Role of Special Orders in Retail Supply Chains

Hao-Wei Chen, Diwakar Gupta, Haresh Gurnani

We present a mathematical model of a two-player retail supply chain, consisting of a manufacturer and a retailer, in the presence of special (expedited) orders. We characterize the retailer's and the manufacturer's optimal stocking and production policies and their responses to changes in supply chain parameters. Our analysis shows that in the presence of special orders, the manufacturer's profit increases more than the retailer's, especially when the customer participation rate is high. In contrast, if more customers are willing to wait for the out-of-stock item (without the need for expedited delivery), then this reduces the manufacturer's power and the retailer's profit increases with customer participation rate.

3 :: To Negotiate or Not: The Effect of Capacity

Chia-wei Kuo, Hyun-soo Ahn, Goker Aydin

We consider a supply chain where the manufacturer has limited capacity and the retailer chooses one of two pricing regimes, posting a fixed price or negotiating, when selling to customers who are heterogeneous in their willingness-to-pay. We analyze the retailer's quantity and pricing regime decisions as well as the manufacturer's inducement of a pricing regime via the wholesale price. We show how capacity influences the manufacturer's wholesale price decision and the resultant pricing regime.

FC01 :: Friday, 1:30-3:00pm

Operations / Marketing Interface

Room 1505

1 :: Pricing Services Subject to Congestion: Sell Subscriptions or Charge on a Per-Use Basis?

Gerard Cachon, Pnina Feldman

In the queueing literature, price mechanisms used to control congestion have primarily been of the pay-per-use type. Subscription or two-part-tariff received almost no attention in this line of research. We compare the three price mechanisms in a simple queueing setting for both the short (fixed capacity) and the long (the firm decides on both the capacity level and the price) term. While on a pay-per-use basis, customers buy on the spot (and thus are aware of their valuation at the time of purchase), we view subscription pricing as a form of advance selling in which the acts of purchase and consumption are separated in time. We find that social optimal is achieved by a two-part tariff scheme in which the firm is able to extract all rent from the consumers. We also find that compared to the social optimal congestion, the system is under-congested under pay-per-use and over-congested when selling subscriptions. When capacity is fixed, we show that when the system is lightly congested, the firm is always better-off selling subscriptions and that subscription can be better than per-use even if the former results in considerably more congestion. When capacity is not fixed, we observe that pay-per-use results in under-investment in capacity, while subscription induces over-investment. Using a numerical study, we conclude that the firm should generally avoid a per-use pricing scheme – it is better-off establishing either a subscription scheme (when capacity cost is not too high) or a two-part tariff.

2 :: Capacity Investment and Pricing Strategy under Product Evolution: An Analytical Model

Xiao Huang, Greys Sosic

Consider a market with two products, in which the new product is along some attributes inferior to the old one. Our model consists of an entrant who can make only the new product, and an incumbent who can manufacture both. A segment of the customers is loyal to the new product, while the rest can purchase either product. We study how the firms strategically determine their capacities and prices under technology evolution, and determine equilibria under different assumptions about the capacity cost.

3 :: Customer Rebates and Retailer Incentives in the Presence of Competition

Ozgun Caliskan Demirag, Pinar Keskinocak, Julie Swann

We analyze the effect of competition on the firms' promotional decisions. We study two-stage supply chains where the promotions are offered by the manufacturers and can be directed to their exclusive retailers (retailer incentives) or to the end customers (customer rebates). We model price discrimination by the retailers and use a game theoretical framework to investigate how the effectiveness of the two promotions is impacted by the competition. We find that the manufacturers can benefit from offering retailer incentives while the retailers may be hurt in the existence of competition. In the case of customer rebates, an equilibrium outcome exists where one manufacturer is better off and the other is worse off compared to the no-promotion case.

FC02 :: Friday, 1:30-3:00pm

Empirical Research

Room 1524

1 :: Estimating the Effect of Stock-out on Sales using Store-level Data

Marcelo Olivares, Andres Musalem, Christian Terwiesch, Daniel Corsten, Eric Bradlow

We develop an econometric model to estimate the effect of stock-outs on sales using store-level data. We use a random utility model to describe consumer choice among the products available at a store. In many applications, product availability is not perfectly known at every point in time (e.g. because of inventory record inaccuracy), which complicates the estimation. We use data augmentation techniques and Bayesian estimation methods to estimate the model with this type of data. We test our methods with real data from 12 supermarket stores on 25 SKUs of shampoo. We use our model to estimate the cost of shortage and to evaluate heuristics to mitigate the impact of stock-outs.

2 :: Inside the Blackbox of Collaborative Planning, Forecasting, and Replenishment (CPFR): An Empirical Analysis of the Learning Process in a Dynamic Business Environment

Yuliang Yao, Rajiv Kohli, Susan Sherer, Jerold Cederlund

Collaborative Planning, Forecasting, and Replenishment (CPFR) is an information systems enabled set of business processes deployed between supply chain partners to insure that the right product is delivered to the right retail space at the right time to meet consumer needs – an increasingly critical goal for competitive supply chains. In order to derive value from CPFR via improved forecast accuracy and higher inventory turns, firms must learn a new set of practices that facilitate collaboration, planning, forecasting, and replenishment. Organizational learning theory proposes that value realization is a learning process and the learning curves may be uniquely dependent on the nature of the program, the tradeoff between exploitation and exploration, as well as the potential spillovers between program components. We construct an econometric model to examine the learning process between a leading telecommunication equipment manufacturer and its partner, a national electronics retailer, in a high velocity business environment characterized by constant new product innovation coupled with rapidly changing customer demands. Using a dataset for 9 product lines over 2-1/2 years, we examine the role of learning in CPFR and its influence on operational performance. We find that two key components of CPFR, collaborative forecasting and collaborative replenishment, follow two distinct learning curves. Consistent with learning theory, forecast accuracy improved immediately after CPFR implementation but at a decreasing rate whereas inventory turns decreased prior to eventually increasing. We find learning spillover from earlier products to later products with later products having higher initial forecast accuracy. Further, we find that as the demand for new products changes, the success of CPFR is dependent upon the firms' learning to recalibrate inventory levels in the channel.

3 :: Testing the Validity of a Demand Model: An Operations Perspective

Omar Besbes, Robert Phillips, Assaf Zeevi

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).

FC03 :: Friday, 1:30-3:00pm

Supply Chain Management

Room 1511

1 :: An Available-to-Promise Production-Inventory System with Pseudo Orders

Long Gao, Susan Xu

We study an order promising problem in the presence of pseudo order information in an Available-to-Promise (ATP) production-inventory system. We develop a Markov framework that captures lumpiness, non-stationarity and volatility, three key characteristics of pseudo order information. We then characterize the structure of the optimal order acceptance policy for both capacity and inventory constrained cases by a sequence of thresholds. Our characterization has the features of a simple policy form and computational efficiency. We also characterize the effects of lead time and demand variability on the system performance and the policy parameters. We numerically compare the performance of three commonly used policies with that of the optimal policy. Our results show that ignoring pseudo order information and resource rationing can severely jeopardize the performance and profitability of the system.

2 :: Integrating Inventory Planning with Project Management in Project-Driven Supply Chains

Ching-Yu Chen, Yao Zhao

We present a new modeling paradigm to help firms make inventory decisions and project decisions jointly. We model a project-driven supply chains by a combined network of material supply and project activities facing random demand and develop mathematic programming model to optimize inventory levels, activity durations and project schedule simultaneously, so as to strike the balance between inventory cost and project cost.

3 :: Real-Time Available to Promise: A Stochastic Programming Model with Decomposition and Sampling Techniques Chien-Yu Chen

A two-stage stochastic programming model is developed to pre-allocate available-to-promise (ATP) into distinct buckets. With pre-allocated ATP quantities, a firm can promise customer orders from different demand channels and time periods in real time. We apply Benders decomposition and importance sampling techniques to reduce solving time and obtain a close-optimal solution over a huge number of stochastic demand scenarios. We also examine solution quality by investigating the statistical optimality gap.

FC04 :: Friday, 1:30-3:00pm

Inventory Theory / Multi-echelon

Room 1520

1 :: Comparison of (s,S) and (R,T) Policies in a Serial Supply Chain with Fixed Costs and Information Sharing Jin Kyung Kwak, Srinagesh Gavirneni

We investigate how the retailer's inventory policy affects the total cost of a two-stage serial supply chain with a fixed cost at the retailer. When the retailer uses the locally optimal (s,S) policy, there is randomness in both order time and order quantity to the supplier. When the retailer uses the suboptimal (but arguably more practical) (R,T) policy, the supplier sees randomness only in order quantity. After formulating and solving the non-stationary inventory control problems faced by the supplier when the retailer uses an (R,T) policy, we perform an extensive computational study to evaluate the performance of the supply chain under (R,T) policy. We observed that the retailer using an (R,T) policy instead of its locally optimal (s,S) policy is better for the entire supply chain in 88.77% of the cases when the supplier has knowledge of the end-customer demand information. This behavior was magnified when the retailer costs were low, the supplier costs were high, and when the end-customer demands had a medium level of variability.

2 :: On the Structure of Serial Inventory Systems with Lost Sales Woonghee Tim Huh, Ganesh Janakiraman

For single-item periodic-review serial inventory systems with lost sales, we establish a structural result on the vector of optimal order quantities, which is analogous to the classical single-stage result of Karlin and Scarf and Morton. This result represents, to our knowledge, the first general structural result for the multi-echelon lost sales inventory systems. Our analysis is based on the recent L-natural-convexity approach of Zipkin.

3 :: Ordering and Pricing a Fashion Product to Minimize Regret Charles Wang, Scott Webster, Sidong Zhang

We consider the problem of ordering and pricing a leading-edge fashion product. The retailer cannot re-order during the brief period of time when the product is targeted to a segment of high-paying consumers. Product that remains after the "season" will be marked down for broader market appeal or sold through a liquidator. The retailer does not know the probability distribution of consumer valuation, but is able to accurately estimate the market size as well as lower and upper limits on valuation. The objective is to minimize the maximum loss in expected profit due to market ignorance, or minimax regret. We derive closed-form expressions for optimal solutions to four types of single-product minimax regret ordering and/or pricing problems under two models of market demand, and we apply these results to solve multi-product problems with budget constraints. The expressions elucidate the interplay among key factors that influence optimal decisions and the value of investments to reduce market ignorance. Interestingly, the optimal ordering and pricing rules are quite simple and are unaffected by the model of market demand.

1 :: The Curse of Reliability: Outsourcing Restoration Services for Infrequent, High-Impact Equipment Failures

Sang-Hyun Kim, Morris Cohen, Serguei Netessine, Senthil Veeraraghavan

Firms in industries such as aerospace and defense, high-tech manufacturing, and telecommunications rely on functioning mission-critical equipment and cannot afford significant operational downtime due to equipment failures. Although resolving problems quickly is essential to minimizing the impact of disruptions, the additional challenge in a decentralized supply chain is incentivizing the suppliers to provide prompt restoration services. A key obstacle is that equipment failures may occur rarely, making it too expensive for a supplier to commit the necessary resources since they will be idle most of the time. A widely adopted incentive mechanism is performance-based contracting (PBC), in which suppliers receive compensation based on realized equipment downtime. It is believed to be an effective means to incentivize speedy restoration services. Using a model that builds upon the principal-agent framework, in this paper we show that designing a successful PBC creates nontrivial challenges that are unique to this environment. Namely, due to the infrequent and random nature of equipment failures, a seemingly innocuous choice of performance measures used in contracts may create unexpected incentives, resulting in counterintuitive optimal contract structures. We find that the definition of the performance metric becomes critical, i.e., in general a contract based on sample-average downtime is superior to a contract based on cumulative downtime. We also find that, paradoxically, a firm that values a consistent level of service delivery would find PBC most costly to implement when a product is most reliable. In such cases, high equipment reliability may be accompanied by excessive contracting costs.

2 :: Centralization vs. Competition in Subcontracting Operations

Tolga Aydinliyim, George Vairaktarakis

Subcontracting and outsourcing have recently become prominent business practices across many industries. Subcontracting of industrial production is generally based on the short-term need for additional production capacity. We model such a business relationship where a group of manufacturers subcontract part of their workload to a third-party, each with the objective of completing his total workload the soonest possible. In such a setting, competition arises due to the sequencing of the subcontracted workloads at the thirdparty facility. Neither the current business practice of the first-come-first-serve (FCFS) processing of the subcontracted workloads, nor the competitive equilibrium schedules developed in the earlier studies result to the maximum utilization of the third-party capacity. For this reason, the third-party finds an optimal centralized schedule and devises a savings sharing scheme that makes it more profitable for all manufacturers to follow the centralized schedule rather than the original. We develop optimal and/or approximate algorithms that produce centralized schedules and propose incentive schemes to promote coordination. We also perform a computational experiment that quantifies the value of coordination in this production chain first against the original schedule and then against Nash equilibrium schedules under complete or incomplete information.

3 :: Managing Suppliers' Delivery Performance with Service Level Agreements

Liping Liang, Derek Atkins

Service-level agreements (SLA) are often observed in practice to manage suppliers' delivery performance. Under an SLA, a supplier agrees to meet a pre-specified service level over a review period. We investigate the optimal SLA between a supplier and a buyer when the supplier can invest both in inventory and in leadtime reduction to meet the service level target. We compare the performance of a ready-rate contract and a ready-rate-with-window contract.

1 :: Dynamic Inventory Competition with Stockout-Based Substitution

Tava Olsen, Rodney Parker

We consider a duopoly where retailers compete by providing inventory under the circumstances where unsatisfied customers may seek satisfaction elsewhere, backlog, or leave. A very general framework is formulated to address a variety of customer avenues when stock is unavailable. The extreme possibilities of fully backlogging (at either retailer) or fully lost sales have dominated the inventory literature so far; we incorporate both of these possibilities along with a blend of these situations, in a competitive context. We find a base-stock inventory policy is the sustainable equilibrium policy in the infinite horizon under several conditions. Further, we give conditions under which the open loop equilibrium is also closed loop. We demonstrate how the policy changes with respect to the inventory parameters and compare the competitive and integrated solutions. We see it is not difficult to observe situations where competition raises the total inventory held but an individual retailer may end up understocking.

2 :: Inventory Centralization Games with Price-Dependent Demand

Xin Chen

Consider a distribution system consisting of a set of retailers facing a single period price dependent demand of a single product. By taking advantage of the risk pooling effect and the quantity/volume discount provided by suppliers or third party carriers, the retailers may place joint orders and keep inventory at central warehouses before demand realization, and allocate inventory among themselves after demand realization to reduce their operating costs. Under rather general assumptions, we prove that there is a stable allocation of profits among the retailers in the sense that the resulting inventory centralization game has a nonempty core. We also show how to compute an allocation in the core.

3 :: Multi-Echelon Inventory System with Intermediate Product Demand: A Perspective on Allocation Policies

Suman Niranjana, Frank Ciarallo

In this paper we develop mathematical model for a capacitated three-echelon inventory system with intermediate, external product demand in one of the upper echelons under different inventory allocation policies. The components are procured from the external suppliers, are assembled into intermediate products and a final product. The inventory system considered in this paper consists of a combination of assembly and serial stages. The demand and capacity are considered as stochastic in nature. A fixed lead time is used between the echelons. Optimal base stock levels for the components, intermediate and final product based on a required customer service level at each stage is determined. A simulation based inventory optimization approach using Infinitesimal Perturbation Analysis (IPA) is employed to determine the optimal base stock levels. We describe some useful initial insights derived from the solutions to these allocation problems.

1 :: Optimal Markdown Pricing: Implications of Inventory Display Formats in the Presence of Strategic Customers

Rui Yin, Yossi Aviv, Amit Pazgal, Christopher Tang

Consider a retailer who sells a limited inventory of a product to strategic customers over a selling season. The retailer attempts to maximize his profit by making three decisions: the initial order quantity, the regular price, and the post-season clearance price. By pre-announcing the regular and post-season clearance prices, customers with heterogeneous valuations and stochastic arrivals can make their strategic purchasing decision by comparing two surpluses: the (known) current surplus obtained from purchasing the product at the regular price; and the (expected) surplus obtained from waiting for the post-season clearance price. Clearly, the expected surplus depends on the likelihood of getting the product at the end of the season.

In this paper we consider two inventory display formats: Display All (DA) and Display One (DO). Under the DA format, the retailer displays all available units so that each arriving customer has perfect information about the actual inventory level. Under the DO format, the retailer displays only one unit at a time so that each customer knows about product availability but not the actual inventory level. Given the retailer's decisions, we prove that, for each display format, the rational expectations' equilibrium purchasing rule for the customers is of a threshold-type. Anticipating customers' purchasing behavior in equilibrium, we determine the retailer's expected profit and develop theoretical bounds on the optimal prices, quantities, and profits. We conducted an extensive numerical study and found that, for any given initial level of inventory, the optimal regular price, as well as the retailer's expected profit, are higher under the DO format. This price premium suggests that, relative to the DA format, the DO format creates a sense of scarcity in the consumers' mind. When the retailer can optimally set the initial level of inventory, we found that a change from DA to DO never leads to a simultaneous decline in inventory and pricing and consequentially never lowers the expected retailer profits. We observed that while a change in display format can be beneficial to the retailer, it is far from totally eliminating the adverse impact of strategic consumer behavior. Our numerical study also enables us to characterize the conditions under which the DO format performs best relative to the DA format.

2 :: On the Value of Inventory Information and Availability When Selling to Strategic Consumers

Xuanming Su, Fuqiang Zhang

This paper studies the role of product availability in attracting consumer demand. We start with a newsvendor model, but additionally assume that consumers must incur some search cost in order to visit the seller. The seller sets an observable price and an unobservable stocking quantity. Consumers anticipate the likelihood of stockout and determine whether to visit the seller. We characterize the rational expectations (RE) equilibrium in this game. We show that the seller can improve profits by providing inventory information (i.e., the seller makes the stocking quantity publicly observable) or availability guarantees (i.e., the seller promises to compensate consumers in the event of stockout). Interestingly, the seller has an incentive to over-compensate consumers during stockouts, relative to the first-best benchmark under which social welfare is maximized. We find that first-best outcomes do not arise in equilibrium, but can be supported when the seller uses a combination of inventory information and availability guarantees. Finally, we examine the robustness of these conclusions by extending our analysis to accommodate heterogeneous consumer valuations.

3 :: Inventories and Customer Search Behavior when Product Quality is Uncertain

Laurens Debo, Garrett van Ryzin

Inventory availability can influence customers' perceptions of product quality, especially with new, unknown or innovative products. Customers who find products out of stock at several retailers may infer that many other customers bought and therefore value it; this information may induce them to buy as well. In this paper, we study this phenomenon. We analyze a model in which the equilibrium customer search process and purchasing behavior depends on the costs and noise inherent in customer search, the ex ante market consensus estimate of quality, the strength of each customer's private signal about quality and the firm's replenishment policy. We find that a single aggregate parameter of all these factors, which we call "search difficulty", can be used to characterize the equilibrium outcome. Equilibria are not unique in general and involve customer search under certain conditions and no search under other conditions; moreover, search may be "intensive" in certain equilibria and "superficial" in others. We show in which cases it is beneficial for the firm to encourage search and in which cases it is detrimental. Lastly, we discuss implications of our model for managing innovative products with uncertain quality.

FD02 :: Friday, 4:00-5:30pm

Empirical Research

Room 1524

1 :: Should Large Hospitals Teach Less? An Empirical Study of the Relations between Hospital Volume, Teaching Status and Service Quality

Carol Theokary, Justin Ren

This paper contributes to research on quality drivers in healthcare settings by examining the relationships between patient volume, the extent of teaching mission and operational service quality in U.S. hospitals. To develop a model that accurately assesses the impact of patient volume and teaching status on quality, we draw from three streams of related research – volume-quality relationship, comparative quality of care in teaching and non-teaching hospitals and quality drivers in service institutions. We propose that the impact of a hospital teaching status on its service quality can vary across hospitals of different sizes. To test our proposition, we use a new large dataset that measures the objective operations quality on the treatment for heart attack of all major US hospitals. Our findings indicate that (1) small and medium size major teaching hospitals surpass their non-teaching counterparts by 1.1% on average (2) in contrast, large major teaching hospitals trail large non-teaching hospitals by 0.5% on average (3) the marginal impact of training an additional resident on service quality is linearly decreasing with volume, and it becomes detrimental to quality for large hospitals. Our results suggest that the concentration of expertise found in teaching institutions significantly contributes to enhancing quality of service in small and medium size hospitals but is masked by the challenges associated with training residents in large-scale operations. A practical implication of this work is that hospitals should re-evaluate the extent of their teaching mission in light of their current patient volume.

2 :: Manufacturing with Chinese Characteristics: Evidence from Auto-makers

Richard Lai, Justin Ren

In this paper, we study the fast-growing Chinese auto manufacturing industry, and compare it to that of the US and Japan. Using a large dataset on major automobile manufacturers of the three countries, we evaluate the technical efficiency and production capabilities of each firm, and quantify the effect of several key efficiency drivers such as inventory management, scale economies, vertical integration, and ownership structure. Our results uncover important differences in China's fast-developing auto industry. China's automakers overall have significantly higher labor productivity, but much lower capital productivity. However, China's joint ventures are more efficient and their production functions are more similar to Japanese and US automaker than to their Chinese peers. This indicates the benefit of technology transfer. Finally, we find that inefficient inventory management among Chinese automakers accounts for their most technical inefficiency.

3 :: Quality Risk in Offshore Manufacturing: An Explanation and Empirical Test

John Gray, Aleda Roth

Does offshoring pose a quality risk? We draw from the literature in organizational theory and quality management to argue that offshoring does pose a quality risk. We empirically test this hypothesis in the drug manufacturing industry. To evaluate a plant's quality risk, we use a heuristic procedure, developed with industry experts, that is applied to Food and Drug Administration inspection data. We compare the quality risk posed by matched pairs of plants owned by the same company in the US mainland and Puerto Rico, which while technically "nearshore," is a reasonable and conservative, proxy for an "offshore" destination. Our findings that offshore plants have significantly higher quality risk complement our previous work which demonstrated that domestic outsourcing poses a quality risk in the drug manufacturing industry.

FD03 :: Friday, 4:00-5:30pm

Supply Chain Management

Room 1511

1 :: Regulatory Trade Risk and Supply Chain Strategy

Yimin Wang, Wendell Gilland, Brian Tomlin

Trade regulations are an important driver of supply chain strategy in many industries. For example, the textile, paper, chemical, and steel industries grapple with significant levels of non-tariff barriers (NTBs) such as voluntary export restraints and safeguard controls. We describe and analyze four often observed supply chain strategies in industries subject to NTBs; direct procurement, split procurement, direct outward processing (D-OPA), and indirect outward processing (I-OPA). We characterize the optimal procurement quantities for each of these four strategies, and also analytically and numerically explore how industry and country characteristics influence the firm's strategy preference. Among other results, we prove that the direct and split procurement strategies are more likely to be preferred in an uncertain NTB environment but the D-OPA strategy is more

likely to be preferred if the expected NTB price is high or the domestic production lead time is short. Also, the D-OPA strategy becomes increasingly attractive as market uncertainty increases, but only if the domestic lead time is short. From a policy perspective, we find that careful attention should be paid to industry characteristics, such as aggregate demand and demand volatility, when setting the mandated domestic production fraction associated with the D-OPA strategy.

2 :: Optimal Transfer Pricing and Offshoring Strategies of Global Supply Chains

Masha Shunko, Laurens Debo, Srinagesh Gavimani

One of the challenges facing global supply chains is the decision associated with offshoring of manufacturing. Offshoring decisions are affected by many factors including cost advantages and tax advantages in different countries of operation. Use of transfer prices allows the firms to shift income between the divisions of the firms to take advantage of favorable tax rates and also can serve as an incentive mechanism for local managements. Specific role of transfer pricing is determined by the global supply chain design. We study the use of transfer pricing in combination with offshoring decisions for different supply chain designs under perfect information and compare the resulting profits. We conclude that different design structures affect the role of transfer pricing in the supply chain and consequently lead to different offshoring and pricing strategies. Through the comparison of resulting profits, we conclude which designs are preferred for a firm with certain characteristics and explain the reasons for these preferences. We demonstrate that regardless of the supply chain design structure, the only scenario in which the firm is not able to exploit both tax and cost advantages to the full capacity is the case when the tax advantage is in the foreign country and the cost advantage is in the local country. We also show that under perfect information the only case when the firm has to deviate from "all-or-nothing" offshoring policy is when the setting of selling price is decentralized to the local management level while the offshoring and transfer pricing decision are made at the central level. These insights change in the presence of information asymmetry between divisions.

3 :: Global Facility Network Design with Responsive Pricing

Lingxiu Dong, Panos Kouvelis, Ping Su

This paper considers a facility network design problem for a global firm who sells to two markets: the domestic market and a foreign market. The firm faces uncertainty in both market demands and exchange rate. The firm has the flexibility of conducting final assembly and pricing products after uncertainties of exchange rate and demands are resolved. The key network design problem faced by the firm is in which of the two markets to locate facilities for the long leadtime component production activities, the corresponding capacity levels, and logistical linkages between the two markets. For general downward sloping demand functions we provide a complete characterization of the optimal global facility network configuration. We show that uncertainty, either in demand or exchange rate, increases the attractiveness of centralizing in one market versus establishing capacities in both markets. We then conduct the sensitivity analysis of the optimal configuration with respect to cost parameters and uncertainties. We find that ex-post pricing flexibility strengthens the substitutability of the capacities in the two markets by reducing the influence of the capacity locations on profit maximization, and hence enhances the robustness of the firm's global facility network configuration decision to changes in demand and exchange rate volatility and their correlations. The strong capacity substitutability effect also explains the counterintuitive result that the increase of market size of one market might decrease the attractiveness of centralizing capacity in that market. We further show that exchange rate uncertainty and demand uncertainty may have distinctively different effects on the optimal capacity investment. For linear market demand functions, the optimal centralized capacity may decrease in the exchange rate volatility but always increases in the demand uncertainty.

FD04 :: Friday, 4:00-5:30pm

Inventory Theory / Multi-echelon

Room 1520

1 :: Incentives for Transshipment in Decentralized Supply Chains with Competing Retailers

Jing Shao

Transshipment of product between retailers is a common practice in many supply chains. Transshipment facilitates a better match between supply and demand, and also allows firms to exploit the benefits of risk pooling. However, in a decentralized supply chain, transshipment may not benefit all the parties in the supply chain. This paper investigates the incentives for transshipment in a decentralized supply chain, where a manufacturer distributes a product through competing retailers. We study the benefit from transshipment, to the manufacturer and retailers, under three possible scenarios. Under each scenario, the firms use a linear wholesale price contract. In the first scenario, both the wholesale price and the transshipment price are exogenous; in the second scenario, we allow the manufacturer to endogenously choose the wholesale price, while the transshipment price remains exogenous. In both these scenarios, the manufacturer and retailers' profits are closely related to the transshipment price, and transshipment is not always in the favor of both parties. In the third scenario, both the wholesale price and the transshipment price are endogenous. Under this scenario, the outcome depends on the amount of control that the manufacturer and retailers have, over the transshipment price and the decision to adopt transshipment.

2 :: Inventory Sharing under Decentralized Preventive Transshipments

Ying Rong, Lawrence Snyder, Yang Sun

We consider two stores in a decentralized system with two demand subperiods. Replenishment orders are made before the first subperiod, and the stores may make preventive transshipments to one another between the subperiods. We prove that the transshipment decision has a dominant strategy, called a control-band conserving transfer policy, under which each store chooses a quantity to transship in or out that will keep its second-subperiod starting inventory level within a range called a control band. We prove that the stores' optimal order (replenishment) policy is a modified base-stock policy in which a store's order-up-to level depends on the initial inventory at the other. We prove that there exists a unique Nash equilibrium of order quantities. Moreover, we prove that there does not exist a transfer price that coordinates the decentralized supply chain.

3 :: Decentralized Inventory Sharing with Asymmetric Information

Hui Zhao, Xinghao Yan

In this paper, we study an inventory sharing game between two retailers who privately hold demand information, non-cooperatively determine their order quantities, but cooperatively share inventory with each other. Characterizing retailers' Bayesian Nash equilibrium order quantities when no demand information is shared between the retailers (NIS), we find that sharing demand information does not always benefit retailers. However, a coordination mechanism can be developed if the retailers share truthful demand information with each other before they each make ordering decisions. Unfortunately, an incentive compatibility analysis shows that the retailers have the incentives to untruthfully reveal their demand information under such coordination mechanism. A truth-inducing scheme is then designed, which, together with the coordination mechanism, achieves the coordination of the retailers and guarantees an all-win situation. To our knowledge, this paper is the first attempt to study inventory sharing systems with asymmetric information.

FD05 :: Friday, 4:00-5:30pm

Supply Chain Management

Room 1528

1 :: A Population Monotonic Allocation Scheme in Joint Replenishment Systems

Jiawei Zhang

We consider joint replenishment systems where retailers belong to different firms or are decentralized divisions of the same firm. The retailers form a coalition and operate in a centralized manner in order to reduce cost. This paper addresses the issue of how to allocate the cost of the coalition among the retailers. We propose a cost allocation scheme under which no retailer would be worse off when a new retailer joins the coalition. This cost allocation scheme provides a strong incentive for the retailers to cooperate.

2 :: Group-buying of Competing Retailers

Rachel Chen, Paolo Roma

Under group-buying, quantity discounts are offered based on the buyers' aggregated purchasing quantity, instead of individual ordering quantity. As the price decreases with the total purchasing quantity, buyers receive lower prices than they would otherwise be able to get individually. Previous studies on group buying focus on the benefit buyers receive in terms of reduced acquisition costs or enhanced bargaining power. In this paper, we show that buyers might get hurt from such cooperation. Specifically, we consider a two-level distribution channel, with a single manufacturer (seller) and two retailers (buyers) who compete for end customers. We show that under competition, retailers always prefer group-buying if they are symmetric (i.e., identical). With asymmetric retailers (i.e., differ in market base and/or efficiency), group-buying is always beneficial to the smaller (less efficient) player. However, it can be detrimental to the larger (more efficient) one. We also show that the manufacturer can benefit from retailers' group purchasing. The managerial insights from our study provide guidelines for retailers who face the decision of whether or not to cooperate in purchase.

3 :: A Study on Coalitions Formation in a Distribution Supply Chain

Adel Elomri, Asma Ghaffari, Zied Jemai, Yves Dallery

In this paper, we are concerned with the analysis of alliances (coalitions) that may arise in a one supplier multiple-retailers distribution system. These retailers have to replenish their inventory from the supplier to satisfy a deterministic and constant rate demand of final customers. Retailers may choose to cooperate by grouping their orders as a big order. Thus, coalition formation may be a key topic that improves the competitive situation where every retailer is looking for an optimal replenishment strategy according to his own objectives and parameters. The goal is to cluster the retailers into exhaustive and disjoint coalitions (alliances) that maximize the total profit. This partition is called coalition structure. The number of coalition structures is too large to allow exhaustive search for the optimal one, furthermore the problem is NP-complete. We focus on analyzing the cooperative behavior of the competing retailers. Further, we propose an effective algorithm that can find the optimal coalition structure via a partial search in the coalition structures space.

FD06 :: Friday, 4:00-5:30pm

Inventory Theory / Multi-echelon

Room 1518

1 :: Optimal Control of a Production-Inventory System with both Backorders and Lost Sales

Saif Benjaafar, Mohsen Elhafsi, Tingliang Huang

We consider the optimal control of a production inventory-system with a single product and two customer classes where items are produced one unit at a time. Upon arrival, customer orders can be fulfilled from existing inventory, if there is any, backordered, or rejected. The two classes are differentiated by their backorder and lost sales costs. At each decision epoch, we must determine whether or not to produce an item and if so, whether to use this item to increase inventory or to reduce backlog. At each decision epoch, we must also determine whether or not to satisfy demand from a particular class (should one arise), backorder it, or reject it. In doing so, we must balance inventory holding costs against the costs of backordering and lost sales. We formulate the problem as a Markov decision process and use it to characterize the structure of the optimal policy. We show that the optimal policy can be described by three state-dependent thresholds: a production base-stock level and two order-admission levels, one for each class. The production base-stock level determines when production takes place and how to allocate items that are produced. This base-stock level also determines when orders from the class with the lower shortage costs (class 2) are backordered and not fulfilled from inventory. The order-admission levels determine when orders should be rejected. We show that the threshold levels are monotonic (either non-increasing or non-decreasing) in the backorder level of class 2. Using numerical results, we compare the performance of the optimal policy against several heuristics and show that those that do not allow for the possibility of both backordering and rejecting orders can perform poorly.

2 :: Centralization versus Decentralization: Risk Pooling, Risk Diversification, and Supply Disruptions in a One-Warehouse Multiple-Retailer System

Amanda Schmitt, Lawrence Snyder, Max Shen

We investigate optimal system design in a One-Warehouse Multiple-Retailer system where supply may be disrupted, and examine the costs and cost variances of the system when either a centralized or decentralized inventory setting is used. We show that using a decentralized inventory model for this system reduces cost variance through risk diversification, and that when demand is deterministic and supply may be disrupted, a decentralized inventory setting is optimal. This is counter to the result when supply is deterministic and demand is stochastic, where risk pooling causes centralization to be optimal. When both supply may be disrupted and demand is stochastic, we demonstrate that a risk-averse firm should typically choose a decentralized setting as its optimal inventory system design.

3 :: Flexibility of Reliable Suppliers and Monitoring of Unreliable Suppliers: How Much Should a Manufacturer Invest?

Soroush Saghafian, Mark Van Oyen

We address two important remedies to increase flexibility and reliability of supply chains in facing disruptions: (1) the critical value of a flexible supplier, and (2) the importance of monitoring unreliable suppliers. We model dynamics of disruptions as discrete time Markov chains and consider different cases of information availability on the suppliers' disruption risk levels where such information may or may not be known by the manufacturing firm. By distinguishing between the firm's perception and underlying reality, we then derive quantitative measures and methods to evaluate the true value of implementing these techniques. These measures provide managers with analytical methods and insights about the amount of money that a risk neutral manufacturer should be willing to invest to implement these techniques.