

INVENTORY COMPETITION AND INCENTIVES TO BACK-ORDER*

Serguei Netessine
The Wharton School of Business
University of Pennsylvania
netessin@wharton.upenn.edu

Nils Rudi
INSEAD
Fontainebleau, France
nils.rudi@insead.edu

Yunzeng Wang
Weatherhead School of Management
Case Western Reserve University
yxw36@po.cwru.edu

Forthcoming in *IIE Transactions*

July 2002, Revised July 2003, November 2004 and August 2005

Abstract

In this paper we consider the issue of inventory control in a multi-period environment with competition on product availability. Specifically, when a product is out of stock, the customer often must choose between placing a back order or turning to a competitor selling a similar product. We consider a competition in which customers may switch between two retailers (substitute) in the case of a stock-out at the retailer of their first choice. In a multi-period setting, the following four situations may arise if the product is out of stock: sales may be lost; customers may back-order the product with their first-choice retailer; customers may back-order the product with their second-choice retailer; or customers may attempt to acquire the product according to some other more complex rule. The question we address is: how do the equilibrium stocking quantities and profits of the retailers depend on the customers' back-ordering behaviors?

In this work we consider the four alternative back-ordering scenarios and formulate each problem as a stochastic multi-period game. Under appropriate conditions, we show that a stationary base-stock inventory policy is a Nash equilibrium of the game that can be found by considering an appropriate static game. We derive conditions for the existence and uniqueness of such a policy and conduct comparative statics analysis. Analytical expressions for the optimality conditions facilitate managerial insights into the effects of various back-ordering mechanisms.

*The authors would like to thank Ton de Kok, three anonymous referees and Matthew Sobel for helpful comments. This paper was previously titled "Dynamic inventory competition and customer retention."

Further, we recognize that often a retailer is willing to offer a monetary incentive to induce a customer to back-order instead of going to the competitor. Therefore, it is necessary to coordinate incentive decisions with operational decisions about inventory control. We analyze the impact of incentives to back-order the product on the optimal stocking policies under competition and determine the conditions that guarantee monotonicity of the equilibrium inventory in the amount of the incentive offered. Our analysis also suggests that, counterintuitively, companies might benefit from making their inventories “visible” to competitors’ customers, since doing so reduces the level of competition, decreases optimal inventories and simultaneously increases profits for both players.

1 Introduction

Recently one of the authors of this paper purchased a new Volkswagen Passat. Living in a small city, he was restricted to buying a car from one of two local Volkswagen dealers. Unfortunately, the first dealer he visited was out of stock on the sought-after Passat configuration, but offered to back-order the car and give an additional discount to make up for the delay. Despite the offer, the author decided to take his chances at another dealer, where he found the configuration of his choice and made a purchase.

Situations like this occur quite often in various retail and industrial settings: a customer who does not find a certain product at the first-choice retailer might decide to switch to another retailer selling the same product or a close substitute. A wealth of research literature addresses the problem of optimally stocking substitutable products under competition. Traditionally, though, this problem is analyzed in a single-period, newsvendor-like setting, and hence the standard assumptions include the risk of lost sales and the salvage of remaining products at a loss at the end of the period. Another feature of a single-period model that is not preserved in a more general multi-period setting is the fact that demand for each retailer depends on the competitor’s but not the retailer’s own inventory.

In some situations customers are willing to back-order the product in the case of a stock-out. For example, a car-buying trip rarely results in an immediate purchase since there are many variations to choose from. Often, the desired car is not available, and the customer faces the choice of back-ordering the car with the first dealer or continuing the search at another dealer. Furthermore, the second dealer may be out of stock too, and the customer faces the dilemma of back-ordering the car with this second dealer or perhaps returning to the first dealer and back-ordering the car there. In such a situation, the total demand faced by each retailer generally depends on the retailer’s own inventory level as well as the competitor’s inventory, and thus retailers compete for customers by setting the stocking quantities of the product. A recent survey of retailers has found that “of the customers that do not find what they want on the shelf, 40% either defer the purchase or go to another store to find the item” (Andraski and Haedicke [2]). Clearly, operational decisions about inventory control that must be made in connection with customer switching and back-ordering

behavior differ from those that arise in a single-period setting. We seek to better understand the influence of customers' decisions to back-order a product on the optimal stocking policies and the resulting profits of the competing retailers, since this influence is key to a conceptual understanding as well as to generating rules for managerial decisions.

Another major issue arising under multi-period competition is giving customers incentives to back-order: it might be profitable for the retailer to offer a monetary enticement (as in the Volkswagen example at the beginning) to induce more customers to back-order the product rather than go to the competitor. In practice, customer incentives are often handled by the marketing department of a company, while stocking decisions are independently set by the operations department. Hence, it is important to understand how the marketing decision to offer a monetary incentive to back-order the product affects the operational decisions involved in selecting an optimal inventory replenishment policy under competition.

1.1 Summary of the main results

In this paper, we analyze situations in which retailers compete for customers by setting the stocking quantities of a single product with exogenously given prices. Specifically, in a multi-period setting we consider two retailers that simultaneously make inventory replenishment decisions at the beginning of each period using a periodic review base-stock policy. Each retailer's demand is a function of the retailer's own inventory as well as the competitor's inventory in the current period, but neither demand depends on any past decisions by either of the two firms. Leftover inventory at the end of the period is carried over to the next period, incurring an inventory holding cost. We begin the analysis by formulating the multiple-period problem in a quite general setting and proving that under appropriate regularity conditions an infinite horizon policy under which both retailers employ stationary base-stock inventory levels is a Nash equilibrium, i.e., a competitive equilibrium can be found by solving an appropriately defined single-period static game.

With respect to customer back-ordering behavior, we formulate four models. Demand that is unsatisfied by both retailers is either completely lost (Model I), or the product is back-ordered (Models II-IV), with retailers incurring penalty charges for backlogging customers. For the case of backlogging we further consider the following scenarios. In Model II we assume that in the case of a stock-out, at, say, retailer i , those customers who are willing to switch to retailer j do so and are backlogged with retailer j in the case that retailer j cannot satisfy them in the same period. In Model III we assume that in the case that demand is not filled initially by retailer i , those customers who are willing to switch do so *only if* retailer j has inventory to satisfy them in the same period; otherwise, they stay and are backlogged with retailer i . As we demonstrate, in the first three models we analyze situations in which the total (effective) demand that a retailer faces in each period (from the first-choice and the second-choice customers) is a piece-wise linear function of the inventory levels of the two retailers. In Model IV we analyze a backlogging problem

in which the mean of the total demand that each retailer faces is an arbitrary function of the stocking quantities of the two retailers. This last model may account for effects other than demand substitution. Examples include the stimulating effect of inventory on demand. (Wolfe [35] provides extensive empirical data to show that weekly sales of some merchandise are strongly correlated with the weekly *beginning* inventory. See also Balakrishnan et al. [5].)

For all four models, we derive tractable analytical solutions, and, whenever we are able to, determine conditions that guarantee the existence/uniqueness of a stationary equilibrium. We show that Model I results in higher inventories and lower profits than Model II. We also conduct sensitivity analysis of equilibrium solutions to changes in the problem parameters. Numerical experiments suggest that different customer back-ordering behavior may result in drastically different inventory decisions and profits. In addition to making an analytical comparison between Model I and II, we demonstrate numerically that Model II results in higher inventories and lower profits than Model III. Therefore, for certain problem parameters it might be in retailers' interest to invest in customer service so as to transition from Model I to Model II. To transition from Model II to Model III, it might be worthwhile to invest in an information system that makes the competitor's inventory visible to customers (if doing so is practical). This result is somewhat counterintuitive: while in practice competitors often tend to limit the information exchange, we find that inventory visibility mitigates competitive overstocking by reducing customer switching, which in turn results in lower inventories and higher profits.

As we noted above, it is sometimes reasonable to expect that the number of customers willing to back-order a product is a function of a monetary incentive that accompanies a back order. We therefore analyze the impact of offering a monetary incentive on the optimal inventory policy by introducing an appropriate relationship between the proportion of backlogging customers and the incentive to back-order. Our main result is that, under some technical assumptions, in Models II and IV the competitors' optimal inventory policies are monotone in the amount of the incentive offered. Specifically, an increase in the incentive offered by retailer i leads to an increase in retailer i 's inventory and a decrease in retailer j 's inventory. Numerical experiments show that, if these technical assumptions are not satisfied, this monotonicity is not necessarily preserved. Moreover, if retailers' inventories are visible to competitors' customers, as in Model III, then offering any incentive at all may be detrimental.

The contributions of this paper are twofold. First, we analyze the inventory policies of two competing retailers and make progress in considering four different nonlinear back-ordering scenarios that arise as a result of competition and customer switching behavior. As such, our paper extends the stream of research on static inventory competition with lost sales by considering a multi-period duopolistic environment and analyzing the impact of customer backlogging behavior, phenomena that previous papers have not studied. As we show, these issues have important implications for firms' profitability. Second, we address the issue of giving customers an incentive to back-order the

product and provide conditions that guarantee monotonicity of equilibrium inventory levels in such incentive.

1.2 Literature survey

A large body of operations literature studies the common phenomenon whereby customers substitute one product with another or switch from one retailer to another when their first-choice product or retailer is stocked out. The stream of literature most relevant to our work is the one that considers substitution under competition, i.e., when substitutable products are sold by different companies that compete for customers. In a single-period (newsvendor) setting, Parlar [26] models the inventory decisions of two competing retailers selling substitute products and shows the existence and uniqueness of the Nash equilibrium. Wang and Parlar [34] extend the model to three retailers. Karjalainen [13], Lippman and McCardle [18], Mahajan and van Ryzin [19, 20], Netessine and Rudi [24] and Netessine and Zhang [25] further study this problem for an arbitrary number of retailers. Anupindi and Bassok [3], Avsar and Baykal-Gursoy [4] and Nagarajan and Rajagopalan [23] analyze the impact of substitution in a multi-period setting with lost sales. To the best of our knowledge, this line of research has thus far been constrained within the single-period framework (or a multi-period framework with an assumption of lost sales), where the modeling of demand backlogging is not an issue and hence differs from our multiple-period problem. Papers by Parlar [26], Wang and Parlar [34], Karjalainen [13], Netessine and Rudi [24] and Anupindi and Bassok [3] model customer switching behavior similarly to our Model I, where we assume that there is no back-ordering. Lippman and McCardle [18] have a more general model with several rules for allocating demand to competing retailers. Mahajan and van Ryzin [19], [20] model demand as a stochastic sequence of heterogeneous customers who choose dynamically among available products based on utility maximization criteria. The closest to our work is a recent paper by Li and Ha [16] in which the authors consider a two-period variant of the inventory competition problem and allow back-ordering with the first-choice retailer only. However, a common feature of all of these papers is that the effective demand for each retailer depends only on *competitors'* inventory. As we show in Models II-IV, in a more general case of multi-period competition with backlogging, effective demand should also depend on the retailer's *own* inventory (due to back-ordering) so that additional complexity is introduced into the analysis and optimality conditions. Hence, single-period inventory competition papers do not capture some of the effects that we analyze. A large portion of research on demand substitution focuses on centralized inventory management decisions. We refer Interested readers to Mahajan and van Ryzin [19] for a comprehensive review of this stream of literature.

Our work fits within the stream of research on stochastic multi-period games that Shapley [30] initiated with his seminal paper. While a number of papers model single-period inventory competition, an analysis of multi-period stochastic games involving inventory decisions by competing retailers is scarce: except for work that includes Kirman and Sobel [14] almost 30 years ago and recent work by Bernstein and Federgruen [6], the literature has been rather silent on the issues specific

to multi-period oligopolies with inventories. Kirman and Sobel [14] consider an oligopoly in which retailers set prices and inventory levels but compete on price only (that is, the demand faced by each of two retailers is a function of both retailers' prices but not their inventories). They show that the stationary mixed pricing policy in which firms randomize their prices is a Nash equilibrium. Bernstein and Federgruen [6] analyze a similar model. They recognize that randomized policies are undesirable in practice, determine the conditions for the existence of stationary pure strategy equilibrium policies, and further analyze the game under more specific assumptions about the nature of competition. The major difference between these two papers and our work is that in their models retailers compete on *price* (even though inventory decisions are made as well), whereas we take prices as exogenous (a rather standard approach in operations literature – the same assumption is made in all the related single-period competition papers cited above) and focus on competition for *inventory* (product availability). One standard justification for taking prices as exogenous is that in many situations prices are fixed for long periods of time, whereas inventory replenishment decisions are made much more frequently.

Competition for inventory among firms located in different echelons of the supply chain has attracted significant attention among researchers. Representative publications in this stream include Cachon and Zipkin [7], Lee and Whang [15], Chen [8] and Porteus [27]. In all of these papers the supplier and the buyer in a two-stage serial supply chain independently choose base-stock policies resulting in suboptimal decisions from a supply chain perspective. Clearly, the setting for such a problem differs greatly from ours, where competition among retailers takes place within the same supply chain echelon. With regard to customer back-ordering behavior under competition, we are aware only of previous work that considers forms of back-ordering that are no different from noncompetitive back-ordering, i.e., the customer either back-orders the product or leaves without making a purchase (as in Kirman and Sobel [14], Bernstein and Federgruen [6], and Cachon and Zipkin [7]). To the best of our knowledge there is no previous work that considers situations where the customer can switch to a competitor and back-order the product there. A large body of literature in marketing extensively studies the customer choice process (see, for example, Chapter 2 in Lilien et. al. [17]) and hence is related to our models of different back-ordering behavior. However, the marketing literature rarely accounts for inventory issues and does not explicitly model back-ordering.

With regard to incentives to back-order, the closest work is DeCroix and Arreola-Risa [10], who also assume that the number of customers willing to back-order the product can be influenced by monetary incentives. They, however, consider only a single monopolistic company. Furthermore, our modeling technique differs from theirs and, unlike DeCroix and Arreola-Risa [10], we do not address the optimal incentive that has to be offered. Another closely related paper is Moinzadeh and Ingene [21], which considers a company simultaneously setting inventory levels for two different but substitutable products: one for immediate and one for delayed delivery. This work is close to ours in that the authors assume that the price for the product with delayed delivery affects the

number of customers who are willing to stay with the retailer rather than go elsewhere. The difference, however, is that we consider a competitive problem setting and substitution occurs between the companies rather than within the company. Cheung [9] considers a continuous review model in which a discount can be offered to customers willing to accept the back-ordering option even before the inventory is depleted, but the proportion of customers back-ordering the product is not a function of a monetary incentive. The work by Gans [11] is also related to incentives to back-order, since it provides insight into the effect of switching behavior on the service levels offered by competing suppliers. However, Gans focuses on customer loyalty as a result of past experience with the company, whereas our paper discusses the immediate impact of stock-outs. As such, the model of Gans is more dynamic than ours.

The remainder of the paper is organized as follows. Section 2 contains the multiple-period model formulation for two competing retailers, and Section 3 presents our results for the different backlogging scenarios. We demonstrate how offering customers a monetary incentive to back-order affects the equilibrium inventory policy in Section 4. We report numerical experiments in Section 5 and make concluding remarks in Section 6.

2 Multi-period model formulation

We consider a competitive duopoly. For simplicity, we assume that there are infinitely many time periods (a finite-horizon model can be similarly analyzed). At the beginning of each period $t = 1, 2, \dots$, two retailers review their inventories and simultaneously make replenishment decisions. We let x_i^t denote the initial inventory of retailer $i = 1, 2$ at the onset of period t . We let Q_i^t denote the order quantity chosen by retailer i in period t . We assume that the inventory replenishment is instantaneous so that $y_i^t = x_i^t + Q_i^t$, the order-up-to inventory level, is the total inventory available at retailer i at the beginning of period t . Then, the constraint $Q_i^t \geq 0$ is equivalent to $y_i^t \geq x_i^t$, and the decision of choosing an order quantity Q_i^t is equivalent to choosing an order-up-to level y_i^t for a given initial inventory x_i^t .

We let D_i^t denote the exogenously given random demand for the product of retailer i in period t from the customers for whom retailer i is a first choice. We assume that *the first-choice demand does not depend on any past decisions made by the two competing firms*. This is a strong assumption because in practice customers may condition their decision to buy from the firm on their past experience with the firm and its competitor. Relaxing this assumption, however, greatly complicates the analysis because the players in this case might condition their behavior on their past decisions. Hence, many equilibrium outcomes arise. We assume such customer behavior away, which is plausible in situations in which each particular customer purchases only once (or very rarely, as is the case with cars) and his experience is not communicated to other customers (i.e., word-of-mouth advertising does not have a strong effect). Bernstein and Federgruen [6] and Kirman and Sobel [14] make similar assumptions. We further assume that the demand distribution for each retailer is

estimated using the exogenously given prices, which accounts for the existing price differential (if any). We assume that D_i^t is a nonnegative continuous random variable. Continuity of demand is a common abstraction that is used to simplify the exposition since in practice demand is discrete. Further, we let \bar{D}_i^t denote the total (effective) demand for the product of retailer i from both the first-choice customers and the second-choice customers, i.e., customers who prefer retailer j but switch to retailer i because retailer j is out of stock. As becomes apparent shortly, in the most general case \bar{D}_i^t depends on the beginning inventory of retailer j in period t , namely y_j^t , as well as on the beginning inventory of retailer i , namely, y_i^t , where $i, j = 1, 2$ and $i \neq j$ ($i, j = 1, 2$ hereafter). That is, the demand realization \bar{D}_i^t is a function of y_i^t and y_j^t , which is denoted by $\bar{D}_i^t(y_i^t, y_j^t)$. In the subsequent sections we present several models of $\bar{D}_i^t(y_i^t, y_j^t)$. For convenience we often omit the arguments and simply write \bar{D}_i^t . Note that this very general definition of the total demand that retailer i faces allows for an arbitrary dependence of demand on starting inventory levels. For example, $\bar{D}_i^t(y_i^t, y_j^t)$ may include the dependence of demand on inventory levels due not only to substitution, but also to the stimulating effects of inventory on demand.

We let retailer i have the following stationary cost and revenue parameters: unit cost of the product c_i , unit revenue r_i , unit inventory holding cost per period h_i , unit cost of backlogging demand per period p_i^1 , and discount factor per period β_i . Also, f_X denotes the density function of random variable X . We first consider the problem with the assumption of lost sales. The inventory balance equations are:

$$x_i^{t+1} = (y_i^t - \bar{D}_i^t)^+, \quad i = 1, 2, \quad t = 1, 2, \dots$$

When the order-up-to levels (y_i^t, y_j^t) are chosen by the two retailers in period t , retailer i 's single-period expected net profit under the lost sales assumption is given by

$$E \left[r_i \min(y_i^t, \bar{D}_i^t) - h_i (y_i^t - \bar{D}_i^t)^+ - c_i Q_i^t \right], \quad i = 1, 2, \quad t = 1, 2, \dots$$

A retailer's total profit is the expectation of the sum of his discounted intra-period profit. Starting with the initial inventories $x^1 \equiv (x_1^1, x_2^1)$ for the retailers in period 1, let $\pi_i(x^1)$ denote the total infinite-horizon profit of retailer i . When the two retailers follow an arbitrary feasible ordering policy $\{(y_1^t, y_2^t), t = 1, 2, \dots\}$, we can write retailer i 's total profit for the lost-sales case as

$$\begin{aligned} \pi_i(x^1) &= E \sum_{t=1}^{\infty} \beta_i^{t-1} \left[r_i \min(y_i^t, \bar{D}_i^t) - h_i (y_i^t - \bar{D}_i^t)^+ - c_i Q_i^t \right] \\ &= E \left\{ \sum_{t=2}^{\infty} \beta_i^{t-1} \left[r_i \min(y_i^t, \bar{D}_i^t) - h_i (y_i^t - \bar{D}_i^t)^+ - c_i \left(y_i^t - (y_i^{t-1} - \bar{D}_i^{t-1})^+ \right) \right] \right. \\ &\quad \left. + r_i \min(y_i^1, \bar{D}_i^1) - h_i (y_i^1 - \bar{D}_i^1)^+ - c_i (y_i^1 - x_i^1) \right\} \end{aligned}$$

¹ p_i represents the compensation paid to the customer who is willing to backlog the product. We assume that all backlogged demand is immediately satisfied in the next period.

$$\begin{aligned}
&= c_i x_i^1 + E \sum_{t=1}^{\infty} \beta_i^{t-1} \left[r_i \min(y_i^t, \bar{D}_i^t) - h_i (y_i^t - \bar{D}_i^t)^+ - c_i y_i^t + \beta_i c_i (y_i^t - \bar{D}_i^t)^+ \right] \\
&= c_i x_i^1 + E \sum_{t=1}^{\infty} \beta_i^{t-1} \left[(r_i - c_i) y_i^t - (r_i + h_i - \beta_i c_i) (y_i^t - \bar{D}_i^t)^+ \right], \quad i = 1, 2,
\end{aligned}$$

where the second equality holds, since $Q_i^t = y_i^t - x_i^t$ and $x_i^t = (y_i^{t-1} - \bar{D}_i^{t-1})^+$ for $t \geq 2$; the last equality uses the fact that $\min\{a, b\} = a - (a - b)^+$. Thus, we can write

$$\pi_i(x^1) = c_i x_i^1 + E \sum_{t=1}^{\infty} \beta_i^{t-1} G_i^t(y_i^t, y_j^t), \quad i, j = 1, 2,$$

where

$$G_i^t(y_i^t, y_j^t) = E \left[(r_i - c_i) y_i^t - (r_i + h_i - \beta_i c_i) (y_i^t - \bar{D}_i^t)^+ \right], \quad i, j = 1, 2, t = 1, 2, \dots$$

This expression can be rewritten as

$$G_i^t(y_i^t, y_j^t) = (r_i - c_i) E \bar{D}_i^t - E \left[(r_i - c_i) (\bar{D}_i^t - y_i^t)^+ + (h_i + (1 - \beta_i) c_i) (y_i^t - \bar{D}_i^t)^+ \right], \quad i, j = 1, 2, t = 1, 2, \dots$$

Using the notation $m_i = r_i - c_i$ for unit margin, $u_i^L = r_i - c_i$ for unit underage cost in the lost-sales model (although in this case $m_i = u_i^L$, the reason to use a different notation for the same quantity becomes clear shortly), and $o_i = h_i + c_i(1 - \beta_i)$ for unit overage cost, we rewrite the single-period objective function in a final form:

$$G_i^t(y_i^t, y_j^t) = E \left[m_i \bar{D}_i^t - u_i^L (\bar{D}_i^t - y_i^t)^+ - o_i (y_i^t - \bar{D}_i^t)^+ \right], \quad i, j = 1, 2, t = 1, 2, \dots \quad (1)$$

Note that in the case of backlogging, the inventory balance equations are $x_i^{t+1} = y_i^t - \bar{D}_i^t$, $i = 1, 2$ and the single-period expected net profit is determined by

$$E \left[r_i \bar{D}_i^t - h_i (y_i^t - \bar{D}_i^t)^+ - p_i (\bar{D}_i^t - y_i^t)^+ - c_i Q_i^t \right], \quad i = 1, 2, t = 1, 2, \dots$$

It is readily verified that in the case of backlogging we arrive at the same expression for a single-period objective function (1), with the only difference being that u_i^L is replaced by $u_i^B = p_i - c_i(1 - \beta_i)$ (there is an explicit backlogging penalty). Hence, for both the lost-sales and backlogging cases we have a generic expression for a single-period objective function (1). When deriving results that are common to both models, we write u_i for the unit-underage cost. We also assume throughout the paper that $u_i^B \leq m_i$.

2.1 Optimality of the stationary inventory policy

We suppose that D_i^t , $t = 1, 2, \dots$ are *i.i.d.* random variables for $i = 1, 2$, that is, that demand is independent across periods but not necessarily independent between retailers. We assume that

D_i^t , $t = 1, 2, \dots$ are nonnegative and possess a continuous differentiable distribution function that is stationary over time. If that is the case, then $G_i^t(y_i^t, y_j^t) = G_i(y_i^t, y_j^t)$ since the order-up-to levels are achievable. We let Δ denote the two-person, noncooperative, static (single-period) game in which player i ($i = 1, 2$) chooses y_i and her payoff function is $G_i(y_i, y_j)$ as defined in (1). A pure-strategy Nash equilibrium of this game is such a pair of base-stock policies that no player wants to deviate unilaterally from it. A pure-strategy Nash equilibrium of the multi-period game is a sequence of such pairs for every period. We let (\bar{y}_1, \bar{y}_2) denote a pure-strategy equilibrium point of the static game Δ , provided such a point exists. Kirman and Sobel [14] demonstrate that there is a randomized stationary policy in a price game where stocking decisions of retailers affect their own (but not their competitor's) profit. We now establish that in our inventory game the *stationary, nonrandomized* inventory policy is a Nash equilibrium.

Proposition 1 *Suppose that demands D_i^t in each period are i.i.d. random variables, functions $\bar{D}_i^t(y_i^t, y_j^t)$ have stationary dependence on (y_i^t, y_j^t) , $\beta_i < 1$ and $(x_1^1, x_2^1) \leq (\bar{y}_1, \bar{y}_2)$. Then the stationary base-stock inventory policy such that $(y_1^t, y_2^t) = (\bar{y}_1, \bar{y}_2)$ for all $t = 1, 2, \dots$, provided it exists, is a pure-strategy Nash equilibrium in the multi-period game.*

Proof: The sequential policy $(y_1^t, y_2^t) = (\bar{y}_1, \bar{y}_2)$ for all $t = 1, 2, \dots$ is *term-by-term* optimal for each player i to maximize its profit $\pi_i(x^1)$. Therefore, to prove the proposition it is sufficient to prove feasibility, i.e., $(x_1^t, x_2^t) \leq (y_1^t, y_2^t) = (\bar{y}_1, \bar{y}_2)$ for all t . By assumption in the proposition, $(x_1^1, x_2^1) \leq (\bar{y}_1, \bar{y}_2)$, so $(y_1^1, y_2^1) = (\bar{y}_1, \bar{y}_2)$ is feasible. Then, for the case of back-ordering (and similarly for the lost-sales case),

$$(x_1^{t+1}, x_2^{t+1}) = (y_1^t, y_2^t) - (\bar{D}_1^t, \bar{D}_2^t) = (\bar{y}_1, \bar{y}_2) - (\bar{D}_1^t, \bar{D}_2^t) \leq (\bar{y}_1, \bar{y}_2), \quad i = 1, 2, \quad t = 1, 2, \dots$$

since $(\bar{D}_1^t, \bar{D}_2^t) \geq 0$. So $(y_1^{t+1}, y_2^{t+1}) = (\bar{y}_1, \bar{y}_2)$ is again feasible, and the solution is stationary according to Chapter 9, "Sequential Games," in Heyman and Sobel [12]. \square

It is worth noting that, even if the equilibrium in a single-period game Δ is unique, this fact alone does not guarantee that the multi-period equilibrium is unique as well, since other more complex nonstationary strategies may arise as a Nash equilibrium. However, since stationary base-stock inventory policies are intuitively appealing, simple to implement in practice and standard in the operations literature, it is particularly important to know that such a strategy is a Nash equilibrium. From now on we focus on stationary equilibria in pure strategies, which means that it suffices to characterize equilibria in a static game (1). For the rest of the paper we drop time superscript t whenever appropriate and analyze a static single-period game with appropriate cost/revenue parameters and demand distributions.

3 Models of customer backlogging behavior

First, it is important to ensure that a Nash equilibrium exists in a single-period game. Lippman and McCardle [18] prove the existence of a Nash equilibrium in a single-period game in a model where each retailer's demand depends on a competitor's (but not the retailer's own) inventory. Since our model is more general in this respect, their proof is not directly applicable to our problem setting. The next two propositions provide quite general conditions for the existence of a pure-strategy Nash equilibrium in the static game with demand substitution.

Proposition 2 *A pure-strategy Nash equilibrium exists in a static game Δ if each realization of the single-period demand \bar{D}_i faced by each retailer is concave in the retailer's own inventory y_i .*

Proof: The sufficient conditions for the existence of a pure-strategy Nash equilibrium are (1) compact, convex strategy sets, (2) the continuity of the players' payoffs and (3) the concavity of each player's objective function (Moulin [22]). Condition (1) is satisfied by choosing a large-enough closed set $[0, M] \times [0, M]$ containing the players' strategies. Condition (2) is satisfied since we assume continuity of distribution functions. Finally, to show condition (3) we rewrite the objective function (1) as follows:

$$\begin{aligned} G_i(y_i, y_j) &= E \left[m_i \bar{D}_i - u_i (\bar{D}_i - y_i)^+ - o_i (y_i - \bar{D}_i)^+ \right] \\ &= E \left[m_i \bar{D}_i - u_i \bar{D}_i + u_i \min(\bar{D}_i, y_i) - o_i y_i + o_i \min(\bar{D}_i, y_i) \right] \\ &= E \left[(m_i - u_i) \bar{D}_i + (u_i + o_i) \min(\bar{D}_i, y_i) - o_i y_i \right], \quad i, j = 1, 2. \end{aligned} \quad (2)$$

A standard result from Rockafellar [28] is that a point-wise minimum of an arbitrary collection of concave functions is a concave function. Hence, $\min(\bar{D}_i, y_i)$ is a concave function, and so $G_i(y_i, y_j)$ is an expectation of a sum of concave functions that is itself a concave function (recall that $m_i \geq u_i$). \square

Sometimes requiring concavity, as in Proposition 2, is too restrictive. Alternatively, one may employ a technique similar to that of Lippman and McCardle [18] and show the existence of an equilibrium through supermodularity. The following proposition makes precise the regularity condition that is required in this case.

Proposition 3 *A pure-strategy Nash equilibrium exists in a static game Δ if each realization of the single-period demand \bar{D}_i faced by each retailer is submodular in variables (y_i, y_j) and $(\bar{D}_i - y_i)$ is a decreasing function of y_i, y_j .*

Proof: The sufficient conditions for the existence of a pure-strategy Nash equilibrium are (1) the compactness of the strategy space, (2) the continuity of the players' payoffs and (3) the supermodularity of each player's objective function (Topkis [32]). Conditions (1) and (2) are shown similarly to Proposition 2. To show (3) for a two-player game we can redefine variables and prove supermodularity in $(y_i, -y_j)$, which is equivalent to proving submodularity in (y_i, y_j) . Consider

expression (2). The first and third terms are clearly submodular, so it remains to show that the second term is submodular as well. Rewrite the second term as follows:

$$(u_i + o_i) \min(\bar{D}_i, y_i) = (u_i + o_i) \min(\bar{D}_i - y_i, 0) + (u_i + o_i) y_i.$$

In this expression, $\min(\bar{D}_i - y_i, 0)$ is a composition of the concave increasing function $\min(\cdot, 0)$ with submodular function $\bar{D}_i - y_i$ that is monotone. According to Table 1 in Topkis [31] such a composition is itself a submodular function. Hence, $(u_i + o_i) \min(\bar{D}_i, y_i)$ is a submodular function as well. Finally, $G_i(y_i, y_j)$ is an expectation of a sum of submodular functions that is itself a submodular function (see Topkis [32]). \square

Proposition 3 generalizes the existence result (Theorem 1) of Lippman and McCardle [18] to the situation in which the retailers' inventories affect both their own and their competitor's demand. Notice that the requirement that $(\bar{D}_i - y_i)$ is decreasing in y_i, y_j is quite intuitive. We expect that \bar{D}_i is increasing in retailer's own inventory y_i but not too fast: a unit increase in inventory should not generate more than one extra unit of demand. Also, it is quite natural that under competition \bar{D}_i is decreasing in y_j .

It is hard to obtain any further analytical results without restricting ourselves to a somewhat narrower class of effective demand functions $\bar{D}_i(y_i, y_j)$. We now consider four specific models of backlogging under competition. In Models I-III we assume that the products of the two retailers are partially substitutable in the customers' eyes. Customers arrive at the store with a product preference that does not depend on current or past inventory levels. Each customer has a first-choice product. If this product is out of stock, only a certain percentage of customers are willing to purchase the other product. More specifically, let α_{ij} , $0 \leq \alpha_{ij} \leq 1$, denote the deterministic proportion of retailer i 's customers who are *willing* to buy from (switch to) retailer j in the case that retailer i experiences a stock-out during this period. We assume that this proportion accounts for the price differential between the two firms. This approach to modeling substitution is a common abstraction that has been used extensively in the literature (see Netessine and Rudi [23] for references). Whenever it is necessary to differentiate among the three models, we use superscripts *I*, *II* and *III*.

3.1 Model I: lost sales

The first model is for *lost sales*. After accounting for the effect of product substitution, we have the following equation for inventory transition:

$$x_i^{t+1} = \left(y_i^t - D_i^t - \alpha_{ji} (D_j^t - y_j^t)^+ \right)^+ = \left(y_i^t - \bar{D}_i^t \right)^+, \quad i, j = 1, 2, t = 1, 2, \dots,$$

where $\bar{D}_i^t = D_i^t + \alpha_{ji} (D_j^t - y_j^t)^+$ is the effective single-period demand for retailer i . Note that the single-period static model with lost sales is essentially identical to the one analyzed by Parlar [26] and Netessine and Rudi [24]. To complete the exposition, we reproduce the optimality condition here. Netessine and Rudi [24] demonstrate that there exists a unique, globally stable (see Moulin [22] page 129 for the definition of stability) pure-strategy Nash equilibrium in this game, and it is characterized by the following first-order conditions:

$$\Pr(\bar{D}_i < \bar{y}_i) = \frac{u_i^L}{u_i^L + o_i}, \quad i = 1, 2. \quad (3)$$

Note that in this case the total demand faced by the retailer is a piecewise linear function of the competitor's inventory (which is captured by \bar{D}_i) but not a function of the retailer's own inventory.

3.2 Model II: back-ordering with the second-choice retailer

In this model, we assume that in the instance where retailer i is out of stock, those customers who are willing to switch from retailer i to retailer j , i.e., $\alpha_{ij} (D_i^t - y_i^t)^+$, do so and are backlogged with retailer j if not filled by retailer j in the same period. Customers that do not switch, i.e., $(1 - \alpha_{ij}) (D_i^t - y_i^t)^+$, are backlogged with retailer i . This scenario represents a situation in which the customer can observe the inventory of only one retailer at a time. The inventory transition equation can be written as

$$x_i^{t+1} = y_i^t - D_i^t + \alpha_{ij} (D_i^t - y_i^t)^+ - \alpha_{ji} (D_j^t - y_j^t)^+ = y_i^t - \bar{D}_i^t, \quad i, j = 1, 2, \quad t = 1, 2, \dots$$

where $\bar{D}_i^t = D_i^t - \alpha_{ij} (D_i^t - y_i^t)^+ + \alpha_{ji} (D_j^t - y_j^t)^+$. Clearly, each retailer's demand is a function of the competitor's as well as the retailer's own inventory, which makes this model quite different from the previous literature on single-period inventory competition. The next proposition establishes the uniqueness of the competitive equilibrium without any additional assumptions.

Proposition 4 *There exists a unique, globally stable pure-strategy Nash equilibrium in a static, single-period game in Model II. It is characterized by the following set of optimality conditions:*

$$\Pr(\bar{D}_i < \bar{y}_i) = \frac{u_i^B}{u_i^B + o_i} + \alpha_{ij} \frac{m_i - u_i^B}{u_i^B + o_i} \Pr(D_i > \bar{y}_i), \quad i, j = 1, 2. \quad (4)$$

Proof: The expression for the effective demand can be rewritten as follows:

$$\begin{aligned} \bar{D}_i &= D_i - \alpha_{ij} (D_i - \min(D_i, y_i)) + \alpha_{ji} (D_j - y_j)^+ \\ &= D_i (1 - \alpha_{ij}) + \alpha_{ij} \min(D_i, y_i) + \alpha_{ji} (D_j - y_j)^+, \quad i, j = 1, 2, \end{aligned}$$

where $\min(D_i, y_i)$ is a concave function (Rockafellar [28]). Hence, $\bar{D}_i(y_i, y_j)$ is concave in y_i for any realization of demand, and by Proposition 2 there exists at least one pure-strategy Nash equilibrium.

Furthermore, the objective function for Model II can be expanded as follows:

$$\begin{aligned}
G_i(y_i, y_j) &= E[(m_i - u_i^B)\bar{D}_i + (u_i^B + o_i) \min(\bar{D}_i, y_i) - o_i y_i] \\
&= E[(m_i - u_i^B) (D_i - \alpha_{ij} (D_i - y_i)^+ + \alpha_{ji} (D_j - y_j)^+) \\
&\quad + (u_i^B + o_i) \min(D_i - \alpha_{ij} (D_i - y_i)^+ + \alpha_{ji} (D_j - y_j)^+, y_i) - o_i y_i], \quad i, j = 1, 2.
\end{aligned}$$

Using the technique for taking derivatives described in Rudi [29], the first derivatives are found as follows:

$$\begin{aligned}
\frac{\partial G_i(y_i, y_j)}{\partial y_i} &= \alpha_{ij}(m_i - u_i^B) \Pr(D_i > y_i) + (u_i^B + o_i) \Pr(\bar{D}_i > y_i) \\
&\quad + \alpha_{ij}(u_i^B + o_i) \Pr(\bar{D}_i < y_i, D_i > y_i) - o_i, \quad i, j = 1, 2.
\end{aligned}$$

It is readily verified that $\Pr(\bar{D}_i < y_i, D_i > y_i) = 0$. Hence, the first derivative is

$$\frac{\partial G_i(y_i, y_j)}{\partial y_i} = \alpha_{ij}(m_i - u_i^B) \Pr(D_i > y_i) + (u_i^B + o_i) \Pr(\bar{D}_i > y_i) - o_i, \quad i, j = 1, 2,$$

and the optimality conditions follow. Furthermore, a sufficient condition for the uniqueness of the Nash equilibrium is that the slopes of the retailers' best-response functions never exceed 1 in an absolute value (see Moulin [22]), which is equivalent to the following condition:

$$\left| \frac{\partial^2 G_i(y_i, y_j)}{\partial y_i \partial y_j} \right| < \left| \frac{\partial^2 G_i(y_i, y_j)}{\partial y_i^2} \right|, \quad i, j = 1, 2. \tag{5}$$

The second-order derivatives needed to verify these conditions are

$$\frac{\partial^2 G_i(y_i, y_j)}{\partial y_i \partial y_j} = -\alpha_{ji}(u_i^B + o_i) f_{\bar{D}_i | D_j > y_j}(y_i) \Pr(D_j > y_j).$$

Furthermore,

$$\begin{aligned}
\frac{\partial^2 G_i(y_i, y_j)}{\partial y_i^2} &= -\alpha_{ij}(m_i - u_i^B) f_{D_i}(y_i) - (u_i^B + o_i) f_{\bar{D}_i}(y_i) \\
&\quad + \alpha_{ij}(u_i^B + o_i) f_{\bar{D}_i | D_i > y_i}(y_i) \Pr(D_i > y_i) \\
&= -\alpha_{ij}(m_i - u_i^B) f_{D_i}(y_i) - (u_i^B + o_i) f_{\bar{D}_i}(y_i)
\end{aligned}$$

because $f_{\bar{D}_i | D_i > y_i}(y_i) \Pr(D_i > y_i)$ vanishes. To see that the required inequality holds, notice that

$$f_{\bar{D}_i}(y_i) \geq f_{\bar{D}_i | D_j > y_j}(y_i) \Pr(D_j > y_j).$$

This completes the proof. \square

Note that in this model the total demand faced by the retailer is a piecewise linear function of the competitor's inventory as well as of the retailer's own inventory. The first-order conditions (4) we obtained can be interpreted as follows. First, without the second term on the right, the solution becomes the same as in Model I where the effective demand that retailer i faces depends only on the competitor's inventory. The extra term appears because effective demand also depends on the retailer's *own* inventory. That is, if demand in the current period exceeds inventory ($D_i > y_i$), then customers switch to the competitor, resulting in the expected relative loss $\alpha_{ij} \frac{m_i - u_i^B}{u_i^B + o_i} \Pr(D_i > \bar{y}_i)$ for player i . Hence, the newsvendor ratio on the right-hand side is adjusted up to account for this effect.

We now compare Models I and II in terms of the level of service offered to customers as well as inventory policies. To this end, we define the equilibrium in-stock probability for firm i as $\Pr(\bar{D}_i < \bar{y}_i)$.

Proposition 5 *Equilibrium in-stock probabilities are higher in Model I than in Model II, $\Pr(\bar{D}_i^I < \bar{y}_i^I) \geq \Pr(\bar{D}_i^{II} < \bar{y}_i^{II})$, $i = 1, 2$. Furthermore, if the problem is symmetric, then equilibrium order-up-to levels are higher in Model I than in Model II, $\bar{y}_i^I \geq \bar{y}_i^{II}$, $i = 1, 2$, and equilibrium profits are lower in Model I than in Model II, $G_i^I(\bar{y}_i^I, \bar{y}_j^I) \leq G_i^{II}(\bar{y}_i^{II}, \bar{y}_j^{II})$, $i = 1, 2$.*

Proof: For Model II, we notice that

$$\begin{aligned}
\Pr(\bar{D}_i < \bar{y}_i) &= \Pr(D_i - \alpha_{ij} (D_i - \bar{y}_i)^+ + \alpha_{ji} (D_j - \bar{y}_j)^+ < \bar{y}_i) \\
&= \Pr((1 - \alpha_{ij}) (D_i - \bar{y}_i) + \alpha_{ji} (D_j - \bar{y}_j)^+ < 0, D_i > \bar{y}_i) \\
&\quad + \Pr(D_i + \alpha_{ji} (D_j - \bar{y}_j)^+ < \bar{y}_i, D_i < \bar{y}_i) \\
&= \Pr(D_i + \alpha_{ji} (D_j - \bar{y}_j)^+ < \bar{y}_i, D_i < \bar{y}_i) \\
&= \Pr(D_i + \alpha_{ji} (D_j - \bar{y}_j)^+ < \bar{y}_i).
\end{aligned}$$

Furthermore, from (4) we have

$$\begin{aligned}
\Pr(D_i + \alpha_{ji} (D_j - \bar{y}_j)^+ < \bar{y}_i) &= \alpha_{ij} \frac{m_i - u_i^B}{u_i^B + o_i} \Pr(D_i > \bar{y}_i) + \frac{u_i^B}{u_i^B + o_i} \\
&\leq \alpha_{ij} \frac{m_i - u_i^B}{u_i^B + o_i} \Pr(D_i + \alpha_{ji} (D_j - \bar{y}_j)^+ > \bar{y}_i) + \frac{u_i^B}{u_i^B + o_i}.
\end{aligned}$$

After simplifying, we obtain

$$\Pr(\bar{D}_i < \bar{y}_i) = \Pr(D_i + \alpha_{ji} (D_j - \bar{y}_j)^+ < \bar{y}_i) \leq \frac{\alpha_{ij} (m_i - u_i^B) + u_i^B}{\alpha_{ij} (m_i - u_i^B) + u_i^B + o_i}.$$

In order to show that the in-stock probability is higher in Model I, it remains to show that

$$\Pr(\bar{D}_i < \bar{y}_i) = \frac{u_i^L}{u_i^L + o_i} \geq \frac{\alpha_{ij} (m_i - u_i^B) + u_i^B}{\alpha_{ij} (m_i - u_i^B) + u_i^B + o_i} \Leftrightarrow u_i^L \geq \alpha_{ij} (m_i - u_i^B) + u_i^B.$$

This is clearly the case because $u_i^I = m_i \geq u_i^B$ and $\alpha_{ij} \leq 1$. To compare order-up-to levels, we note that

$$\Pr(\bar{D}_i^I < \bar{y}_i^I) = \Pr(D_i + \alpha_{ji} (D_j - \bar{y}_j^I)^+ < \bar{y}_i^I) \geq \Pr(D_i + \alpha_{ji} (D_j - \bar{y}_j^{II})^+ < \bar{y}_i^{II}) = \Pr(\bar{D}_i^{II} < \bar{y}_i^I).$$

Since the problem is symmetric, we know that one of the equilibria must be symmetric. Moreover, since the equilibrium is unique, it follows that this unique equilibrium is symmetric. If we restrict our attention to the symmetric equilibria, only two situations need to be analyzed: $\bar{y}_i^I \leq y_i^{II}$, $\bar{y}_j^I \leq y_j^{II}$ and $\bar{y}_i^I \geq y_i^{II}$, $\bar{y}_j^I \geq y_j^{II}$. The first case clearly does not satisfy the above inequality thus we are left with $\bar{y}_i^I \geq y_i^{II}$, $\bar{y}_j^I \geq y_j^{II}$. Finally, to compare profits we note that $G_i^I(\bar{y}_i^I, \bar{y}_j^I) \leq G_i^I(\bar{y}_i^I, \bar{y}_j^{II})$ because more demand is available for player i , and furthermore $G_i^I(\bar{y}_i^I, \bar{y}_j^{II}) \leq G_i^{II}(\bar{y}_i^I, \bar{y}_j^{II})$, because orders are backlogged in Model II. Finally, it is quite clear that $G_i^{II}(\bar{y}_i^I, \bar{y}_j^{II}) \leq G_i^{II}(\bar{y}_i^{II}, \bar{y}_j^{II})$, and the result follows. \square

It would seem that in Model I retailers face a total demand that is smaller than in Model II, and therefore Model II should result in higher inventories. But on the contrary, we see that, if unsatisfied customers back-order with the second-choice retailer, competition is reduced and both retailers tend to establish lower levels of service as well as lower order-up-to levels because the danger of losing customers to a competitor is reduced, leading to the reduction in competitive overstocking. As a result, single-period models with inventory competition are likely to overstate the effect of overstocking because the possibility of back-ordering dampens the competition through reduction in the number of customers who switch between competitors.

3.3 Model III: back-ordering with the first-choice retailer

The second backlogging case is to assume that, if their orders are not filled by one retailer, those customers who are willing to switch do so only if the other retailer has inventory to satisfy them in the same period; otherwise they stay and are backlogged with the original retailer. Notice that, intuitively, this model is in some sense less competitive than Model II since customers switch less often. The number of customers that actually switch from i to j is limited by the number of customers willing to switch from i to j and the number of customers that can be served by retailer j immediately, which is given by

$$\min \left\{ \alpha_{ij} (D_i^t - y_i^t)^+, (y_j^t - D_j^t)^+ \right\}.$$

Thus, we have the following inventory transition equation:

$$\begin{aligned} x_i^{t+1} &= y_i^t - D_i^t + \min \left\{ \alpha_{ij} (D_i^t - y_i^t)^+, (y_j^t - D_j^t)^+ \right\} - \min \left\{ \alpha_{ji} (D_j^t - y_j^t)^+, (y_i^t - D_i^t)^+ \right\} \\ &= y_i^t - \bar{D}_i^t, \quad i, j = 1, 2, \quad t = 1, 2, \dots \end{aligned}$$

where $\bar{D}_i^t = D_i^t - \min \left\{ \alpha_{ij} (D_i^t - y_i^t)^+, (y_j^t - D_j^t)^+ \right\} + \min \left\{ \alpha_{ji} (D_j^t - y_j^t)^+, (y_i^t - D_i^t)^+ \right\}$. Due to the complexity of the expression for effective demand, certain difficulties arise with the analysis of this model. As a result, we are able only to obtain optimality conditions but not to prove the existence of an equilibrium. To see the complexity of this model, note that the effective demand in Model III is generally not concave in y_i since the term $-\min \left\{ \alpha_{ij} (D_i - y_i)^+, (y_j - D_j)^+ \right\}$ is, in general, neither concave nor convex in y_i . Hence, we cannot employ Proposition 2. Moreover, one can also verify that demand in this model is generally not submodular in (y_i, y_j) since the same term is supermodular in (y_i, y_j) , so we cannot employ Proposition 3 either. However, the optimality conditions that we obtain shortly allow us to find the equilibrium numerically. In our numerical experiments the equilibrium in Model III always exists and is always unique.

Proposition 6 *Suppose there exist interior equilibria (possibly multiple) in Model III. Then all equilibria satisfy the following set of optimality conditions:*

$$\begin{aligned} \Pr(\bar{D}_i < \bar{y}_i) &= \frac{u_i^B}{u_i^B + o_i} + \frac{m_i - u_i^B}{u_i^B + o_i} \alpha_{ij} \Pr \left(\bar{y}_i < D_i < \bar{y}_i + \frac{\bar{y}_j - D_j}{\alpha_{ij}} \right) \\ &+ \frac{m_i - u_i^B}{u_i^B + o_i} \Pr(\bar{y}_i - \alpha_{ji} (D_j - \bar{y}_j) < D_i < \bar{y}_i), \quad i, j = 1, 2. \end{aligned} \quad (6)$$

Proof: The objective function for Model III can be expanded as follows:

$$\begin{aligned} G_i(y_i, y_j) &= E[(m_i - u_i^B)\bar{D}_i + (u_i^B + o_i) \min(\bar{D}_i, y_i) - o_i y_i] \\ &= E[(m_i - u_i^B) \left(D_i - \min \left\{ \alpha_{ij} (D_i - y_i)^+, (y_j - D_j)^+ \right\} \right. \\ &\quad \left. + \min \left\{ \alpha_{ji} (D_j - y_j)^+, (y_i - D_i)^+ \right\} \right) \\ &\quad + (u_i^B + o_i) \min \left(D_i - \min \left\{ \alpha_{ij} (D_i - y_i)^+, (y_j - D_j)^+ \right\} \right. \\ &\quad \left. + \min \left\{ \alpha_{ji} (D_j - y_j)^+, (y_i - D_i)^+ \right\}, y_i \right) - o_i y_i], \quad i, j = 1, 2. \end{aligned}$$

The first derivatives are:

$$\begin{aligned} \frac{\partial G_i(y_i, y_j)}{\partial y_i} &= (m_i - u_i^B) \alpha_{ij} \Pr(\alpha_{ij} (D_i - y_i) < y_j - D_j, D_i > y_i) \\ &+ (m_i - u_i^B) \Pr(\alpha_{ji} (D_j - y_j) > y_i - D_i, D_i < y_i) \\ &+ (u_i^B + o_i) \Pr(\bar{D}_i > y_i) \\ &+ (u_i^B + o_i) \alpha_{ij} \Pr(\bar{D}_i < y_i, \alpha_{ij} (D_i - y_i) < y_j - D_j, D_i > y_i) \\ &+ (u_i^B + o_i) \Pr(\bar{D}_i < y_i, \alpha_{ji} (D_j - y_j) > y_i - D_i, D_i < y_i) - o_i, \quad i, j = 1, 2. \end{aligned}$$

We consider first the term

$$\Pr(\bar{D}_i < y_i, \alpha_{ij} (D_i - y_i) < y_j - D_j, D_i > y_i).$$

First, we observe that since $D_i > y_i$ and $\alpha_{ij}(D_i - y_i) < y_j - D_j$, we can conclude that $\bar{D}_i = D_i - \alpha_{ij}(D_i - y_i)$. Hence, the first inequality inside the bracket, $\bar{D}_i < y_i$, expands to $D_i - \alpha_{ij}(D_i - y_i) < y_i$ or similarly $D_i < y_i$, which contradicts the last inequality inside the bracket, $D_i > y_i$. Hence, this term is always zero. Similarly, in the term

$$\Pr(\bar{D}_i < y_i, \alpha_{ji}(D_j - y_j) > y_i - D_i, D_i < y_i)$$

under the conditions $D_i < y_i$ and $\alpha_{ji}(D_j - y_j) > y_i - D_i$, we can conclude that $\bar{D}_i = D_i + (y_i - D_i) = y_i$, resulting in this term being always zero. Hence, the expression for the derivative simplifies to

$$\begin{aligned} \frac{\partial G_i(y_i, y_j)}{\partial y_i} &= (m_i - u_i^B)\alpha_{ij} \Pr\left(y_i < D_i < y_i + \frac{y_j - D_j}{\alpha_{ij}}\right) \\ &\quad + (m_i - u_i^B) \Pr(y_i - \alpha_{ji}(D_j - y_j) < D_i < y_i) \\ &\quad + (u_i^B + o_i) \Pr(\bar{D}_i > y_i) - o_i, \quad i, j = 1, 2, \end{aligned}$$

and the resulting optimality conditions follow. This completes the proof. \square

As in the previous model, the retailer's effective demand in Model III is a piecewise linear function of own and the competitor's inventories. The optimality condition can be interpreted as follows. The fact that the retailer's effective demand depends on own inventory level is captured by the second and third terms on the right-hand side of (6) that adjust the otherwise standard newsvendor ratio up. Seemingly, this should lead to a higher level of service in Model III than in Model I, but this assertion is hard to verify analytically. The second term on the right-hand side can be interpreted as follows: if this period's demand for the products of player i , D_i , exceeds current inventory y_i , then there is a chance that customers may switch to player j if she has inventory to satisfy this demand immediately, $D_j < y_j$. The probability of a simultaneous shortage at player i and excess at player j is $\Pr(\bar{y}_i < D_i < \bar{y}_i + \frac{\bar{y}_j - D_j}{\alpha_{ij}})$, and the relative expected cost of losing a customer to the competitor becomes $\frac{m_i - u_i^B}{u_i^B + o_i} \alpha_{ij} \Pr(\bar{y}_i < D_i < \bar{y}_i + \frac{\bar{y}_j - D_j}{\alpha_{ij}})$. This term then adjusts the right-hand side of the equation up to increase the standard newsvendor ratio to reflect that losing customers to competition is costly. The third term on the right $\frac{m_i - u_i^B}{u_i^B + o_i} \Pr(\bar{y}_i - \alpha_{ji}(D_j - \bar{y}_j) < D_i < \bar{y}_i)$ has a different interpretation. Even though player i might have enough inventory to satisfy his own demand ($D_i < y_i$), he has a potential to capture additional demand from player j . However, this demand only materializes if there is inventory to satisfy it immediately. The probability of both excess demand at player j and sufficient inventory at player i is $\Pr(\bar{y}_i - \alpha_{ji}(D_j - \bar{y}_j) < D_i < \bar{y}_i)$, and the relative expected cost of not being able to capture customers switching from player j is $\frac{m_i - u_i^B}{u_i^B + o_i} \Pr(\bar{y}_i - \alpha_{ji}(D_j - \bar{y}_j) < D_i < \bar{y}_i)$. This term adjusts the right-hand side of the equation up to increase the standard newsvendor ratio to reflect that failure to capture overflow customers from the competitor is costly. Finally, the fact that the retailer's demand depends on the competitor's inventory is captured in \bar{D} . Note that in Model III conditions sufficient for customers to switch

between the two firms depend on inventories and demands at both players, whereas in Model II the condition is simply $D_i > y_i$. One way to interpret this observation is that on average there is less customer switching in Model III than in Model II, so in some sense Model III is less competitive than Model II. The implications of this result are discussed further in later sections.

3.4 Model IV: nonlinear back-ordering rule

The previous three back-ordering models focus on customer behavior due to substitution and essentially assume that the relationship between inventory and demand is piecewise linear. There are, however, other situations in which substitution due to stock-outs either does not occur or has a small impact. Instead, other effects may be present. Large inventory by itself might have a stimulating effect on the customers. Car dealers, for example, usually place all their inventory in a parking lot in front of the dealership to attract customers' attention (see Wang and Gerchak [33] for other examples and numerous references).

In light of these considerations, we analyze an alternative model that allows us to tackle situations in which demand for each retailer depends on each retailer's inventory in a nonlinear way, i.e., $\bar{D}_i(y_i, y_j)$ is a nonlinear function of y_i, y_j . A completely general analysis in such a case is complex, so we make additional technical assumptions, including the assumption that only the mean of the total demand that the retailer faces depends on the initial inventory levels of the two retailers. Such an assumption does not hold for Models I-III, so it is not appropriate when the main effect is the substitution when product is out of stock. However, it captures the essence of the problem in other situations (e.g., stimulating effects of inventory) when demand depends on competitors' inventory policies. Hence, this model can be used to gain insights into the issues involved, since it allows for analytical tractability that cannot be achieved using previously described models. This assumption is also frequently encountered in the operations, marketing, and economics literature.

Previous analysis of Models I-III and common sense lead us to believe that the demand faced by each retailer should be increasing in the retailer's own inventory and decreasing in the competitor's inventory. Taken together, these arguments can be summarized as follows:

Assumption 1a. $\bar{D}_i(y_i, y_j) = \eta_i(y_i, y_j) + \varepsilon_i$, where ε_i is an arbitrarily distributed random variable, density $f_{\varepsilon_i}(\cdot)$, distribution $F_{\varepsilon_i}(\cdot)$, and $\eta_i(y_i, y_j)$ is a positive real-valued function such that $\bar{D}_i > 0$, $1 \geq \frac{\partial \eta_i}{\partial y_i} \geq 0$, $\frac{\partial \eta_i}{\partial y_j} \leq 0$, $i, j = 1, 2$.

To provide a more specific example of when Model IV might be a good approximation of real-life situations, we suppose that the total market size is known with near certainty and is equal to η . The part of the population attracted by each of the two competitors depends on the firms' respective inventory policies so that each firm gets $\eta_1(y_1, y_2)$, $\eta_2(y_1, y_2)$ customers with $\eta_1 + \eta_2 = \eta$. There is, however, some uncertainty with respect to whether each particular customer makes a purchase

or not. For example, a customer may have some preliminary idea about the product, but his preference may change when he sees it. This uncertainty is firm-specific (i.e., once the customer visits the firm, his decision to purchase is not affected by the competitor's inventory policy but is only affected by the firm-specific or product-specific characteristics, so that random shocks ε_i could be correlated and asymmetric). Hence, the resulting demand is $\eta_i(y_i, y_j) + \varepsilon_i$ such that $\eta_i(y_i, y_j)$ can be adjusted up (e.g., customers tend to buy more than one unit of the product) or down (some customers may not buy at all).

Throughout the analysis it is understood that η_i and its derivatives are functions of y_i, y_j and are evaluated at appropriately selected y_i and y_j . To ensure the existence of an equilibrium in the game, an additional assumption is needed regarding the second-order effects. We present two alternative and quite intuitive technical assumptions about the demand distributions $\bar{D}_i(y_i, y_j)$. First, it is reasonable to believe that in a majority of situations, $\eta_i(y_i, y_j)$ should exhibit a decreasing marginal rate of return in y_i . In other words, increasing y_i increases demand \bar{D}_i at a decreasing rate:

Assumption 2a. $\frac{\partial^2 \eta_i}{\partial y_i^2} < 0, i = 1, 2.$

Such an assumption is clearly sufficient to prove the existence of an equilibrium, due to Proposition 2. An alternative assumption could be made about the second-order cross-effect. Since the products of retailers i and j are physical substitutes, it is reasonable to expect that they are also substitutes in an economic sense, that is, increasing the stocking quantity of one retailer reduces the marginal benefit of increasing the other retailer's stocking quantity.

Assumption 2b. $\frac{\partial^2 \eta_i}{\partial y_i \partial y_j} < 0, i, j = 1, 2.$

Clearly, Assumption 2b suffices to guarantee the existence of an equilibrium as well, due to Proposition 3. Moreover, for two players, Assumption 2b is satisfied by a number of standard demand functions including the linear, Logit, Cobb-Douglas and CES demand functions (see Bernstein and Federgruen [6]).

Proposition 7 *Under Assumption 1a and either Assumption 2a or 2b, there exists at least one pure strategy Nash equilibrium, characterized by the following set of optimality conditions:*

$$\Pr(\bar{D}_i < \bar{y}_i) = \frac{u_i^B}{u_i^B + o_i} + \frac{m_i}{u_i^B + o_i} \frac{\frac{\partial \eta_i}{\partial y_i}}{1 - \frac{\partial \eta_i}{\partial y_i}}, i = 1, 2. \quad (7)$$

Proof: Existence trivially follows from Propositions 2 and 3. Furthermore, the first derivatives of the objective function are found as follows:

$$\frac{\partial G_i(y_i, y_j)}{\partial y_i} = \left(m_i - u_i^B\right) \frac{\partial \eta_i}{\partial y_i} - \left(u_i^B + o_i\right) \Pr(\bar{D}_i < y_i) \left(1 - \frac{\partial \eta_i}{\partial y_i}\right) + u_i^B, i = 1, 2.$$

The set of optimality conditions follows after equating to zero and rearranging. \square

Note that the optimality conditions we obtained have nice interpretable properties. First, the optimal fractile (right-hand side of (7)) has two distinct parts. Part one is a standard newsvendor fractile. Part two is an extra term that accounts for the fact that the effective demand depends on the retailer's own inventory (the fact that demand depends on the competitor's inventory is captured in \bar{D}_i). This term is similar to the probability terms encountered in Models II and III. If $\frac{\partial \eta_i}{\partial y_i} = 0$ (demand does not depend on the retailer's own stocking quantity but perhaps depends on the competitor's stocking quantity), we arrive at a solution identical to the solution for Model I. Moreover, if $\bar{D}_i(y_i, y_j) = D_i$, then we arrive at the classic newsvendor solution. The optimal in-stock probability is higher for higher values of $\frac{\partial \eta_i}{\partial y_i}$.

We need some further assumptions to guarantee the uniqueness of the equilibrium. To this end, we assume that the sum of the absolute changes due to a unit increase in the retailer's own inventory and the competitor's inventory does not exceed 1. This is a reasonable assumption in most situations, since it is hard to expect that change in one unit of inventory can cause an effect significant enough to change demand by more than one unit. For Models I-III, for example, $\partial E\bar{D}_i(y_i, y_j)/\partial y_i < \alpha_{ij}$ and $\partial E\bar{D}_i(y_i, y_j)/\partial y_j < \alpha_{ji}$, and therefore Assumption 3 takes the form of $\alpha_{ij} + \alpha_{ji} < 1$, which is a reasonable condition for most practical situations. Formally, we have

Assumption 3. $\left| \frac{\partial \eta_i}{\partial y_i} \right| + \left| \frac{\partial \eta_i}{\partial y_j} \right| \leq 1, \quad i, j = 1, 2.$

Finally, we assume that the marginal value of a retailer's own inventory y_i is more sensitive to y_i than to y_j , an assumption that is rather standard in economics and that holds for a number of standard demand functions (see Bernstein and Federgruen [6]).

Assumption 4. $\left| \frac{\partial^2 \eta_i}{\partial y_i \partial y_j} \right| < \left| \frac{\partial^2 \eta_i}{\partial y_i^2} \right|, \quad i, j = 1, 2.$

The additional Assumptions 3 and 4 suffice to show the uniqueness of the equilibrium.

Proposition 8 *Suppose that one of the pairs of conditions of Proposition 7 hold and moreover Assumptions 3 and 4 hold as well. Then there exists a unique, globally stable Nash equilibrium in the static game of Model IV. It is characterized by optimality conditions (7).*

Proof: To demonstrate the uniqueness, we again employ the sufficient condition (5) from Moulin [22] that was used earlier. The second derivatives are:

$$\begin{aligned} \frac{\partial^2 G_i(y_i, y_j)}{\partial y_i^2} &= (m_i - u_i^B) \frac{\partial^2 \eta_i}{\partial y_i^2} - (u_i^B + o_i) f_{\bar{D}_i}(y_i) \left(1 - \frac{\partial \eta_i}{\partial y_i}\right)^2 \\ &\quad + (u_i^B + o_i) \Pr(\bar{D}_i < y_i) \frac{\partial^2 \eta_i}{\partial y_i^2}, \quad i, j = 1, 2, \end{aligned}$$

$$\begin{aligned} \frac{\partial^2 G_i(y_i, y_j)}{\partial y_i \partial y_j} &= (m_i - u_i^B) \frac{\partial^2 \eta_i}{\partial y_i \partial y_j} + (u_i^B + o_i) f_{\bar{D}_i}(y_i) \left(1 - \frac{\partial \eta_i}{\partial y_i}\right) \frac{\partial \eta_i}{\partial y_j} \\ &\quad + (u_i^B + o_i) \Pr(\bar{D}_i < y_i) \frac{\partial^2 \eta_i}{\partial y_i \partial y_j}, \quad i, j = 1, 2. \end{aligned}$$

After term-by-term comparison, we see that $\left| \frac{\partial^2 \eta_i}{\partial y_i^2} \right| > \left| \frac{\partial^2 \eta_i}{\partial y_i \partial y_j} \right|$ by Assumption 4 and that $\left(1 - \frac{\partial \eta_i}{\partial y_i}\right) > \left| \frac{\partial \eta_i}{\partial y_j} \right|$ by Assumption 3, sufficient for the proof. \square

3.4.1 Example: linear demand function

To give a more specific example for Model IV, we consider the linear form of dependence between the expected demand and inventory. Specifically, we assume that retailers face the following demand distributions,

$$\bar{D}_i = A_i + b_i^1 y_i - b_i^2 y_j + \varepsilon_i, \quad i, j = 1, 2,$$

where $b_i^1 + b_i^2 < 1$, $b_i^1 > b_i^2$, and $b_i^1, b_i^2 > 0$, $\varepsilon_i \sim N(0, \sigma_i)$ and A_i s are large enough so that the probability of negative demand is negligible. We denote by $\Phi(\cdot)$ the standard Normal distribution function. Then the optimality conditions (7) yield

$$\begin{aligned} \bar{y}_1 &= A_1 + b_1^1 \bar{y}_1 - b_1^2 \bar{y}_2 + \sigma_1 z_1, \\ \bar{y}_2 &= A_2 + b_2^1 \bar{y}_2 - b_2^2 \bar{y}_1 + \sigma_2 z_2, \end{aligned}$$

where $z_i = \Phi^{-1}\left(\frac{u_i^B}{u_i^B + o_i} + \frac{m_i}{u_i^B + o_i} \frac{b_i^1}{1 - b_i^1}\right)$. The solution in a closed form is:

$$\bar{y}_i = \frac{(A_i + \sigma_i z_i) (1 - b_j^1) - (A_j + \sigma_j z_j) b_i^2}{(1 - b_i^1) (1 - b_j^1) - b_i^2 b_j^2}, \quad i, j = 1, 2.$$

From this solution, sensitivity analysis to all problem parameters is rather straightforward. In case the retailers are symmetric, the solution becomes

$$\bar{y}_1 = \bar{y}_2 = \frac{A + \sigma z}{1 - (b^1 - b^2)}. \quad (8)$$

3.4.2 Comparative statics

Assumption 2b is particularly natural in this problem setting and useful when obtaining comparative statics of the game. Namely, using this assumption we are able to characterize the shift in equilibrium base-stock levels as a response to changes in cost and revenue parameters. The next proposition makes this statement precise.²

²One can also verify that the same characterizations hold for Models I and II because these games can be shown to be submodular.

Proposition 9 *Suppose Assumptions 1a, 2b, 3 and 4 hold and further suppose that (\bar{y}_1, \bar{y}_2) is an equilibrium of the game. Then an increase in $r_i, p_i, -c_i, -h_i, \beta_i$ (alternatively, a decrease in $r_j, p_j, -c_j, -h_j, \beta_j$) leads to a new equilibrium (\hat{y}_1, \hat{y}_2) such that $\bar{y}_1 \geq \hat{y}_1$ and $\bar{y}_2 \leq \hat{y}_2$.*

Proof: Since the game is submodular by Assumption 2b, we redefine $\tilde{y}_2 = -y_2$ to obtain a supermodular game in (y_1, \tilde{y}_2) . In supermodular games, the sufficient condition for parametric monotonicity of equilibrium in parameter θ is the property of the increasing differences³ of the players' objective functions in decision variables and parameter θ (see Topkis [32], Theorem 4.2.2). Furthermore, the sufficient condition for increasing differences in (y_1, θ) and (\tilde{y}_2, θ) is the non-negativity of the second-order cross-partial derivatives. We shall consider sensitivity to player i 's parameters (sensitivity to player j 's parameters is derived similarly):

$$\begin{aligned} \frac{\partial^2 G_i(y_i, y_j)}{\partial y_i \partial r_i} &= \frac{\partial \eta_i}{\partial y_i} \geq 0, \\ \frac{\partial^2 G_i(y_i, y_j)}{\partial y_i \partial p_i} &= (1 - \Pr(\bar{D}_i < y_i)) \left(1 - \frac{\partial \eta_i}{\partial y_i}\right) \geq 0, \\ \frac{\partial^2 G_i(y_i, y_j)}{\partial y_i \partial (-c_i)} &= -\beta_i \left(1 - \frac{\partial \eta_i}{\partial y_i}\right) + 1 \geq 0, \\ \frac{\partial^2 G_i(y_i, y_j)}{\partial y_i \partial (-h_i)} &= \Pr(\bar{D}_i < y_i) \left(1 - \frac{\partial \eta_i}{\partial y_i}\right) \geq 0, \\ \frac{\partial^2 G_i(y_i, y_j)}{\partial y_i \partial \beta_i} &= c_i \left(1 - \frac{\partial \eta_i}{\partial y_i}\right) \geq 0. \end{aligned}$$

Clearly, all corresponding derivatives for player j with respect to player i 's parameters are zero, and the proof is complete. \square

Proposition 9 extends Theorem 4 of Lippman and McCardle [18] into multiple periods and also tests sensitivity to parameters pertaining to the multiple-period models. However, since customer back-ordering behavior is at the heart of the game, one may wonder if our comparative statics analysis with respect to p (back-order penalty) is too simplified: it is likely that changing p affects demand distribution for each retailer as well. This issue is thoroughly analyzed in the next section.

4 Incentives to back-order

In practice, retailers often attempt to influence customer behavior by offering a monetary incentive that persuades the customer to back-order the out-of-stock product rather than go to another retailer. The natural operational question arises: how does offering a monetary incentive influence the optimal stocking decisions in a competitive situation? Promotional decisions (e.g., offering monetary incentives) are usually made by the marketing department, while stocking policies are

³When all functions are continuously differentiable (as in this paper), increasing (decreasing) differences are equivalent to supermodularity (submodularity). However, to simplify references to Topkis's results, we use the terminology of increasing/decreasing differences.

controlled by the operations department. In such a situation it is crucial for operations managers to understand what effect a monetary incentive to back-order the product has on the inventory replenishment policy. Proposition 9 answers this question only partially, since it does not account for the effects that incentives have on the demand distribution. In particular, we may expect that an incentive increases the total demand faced by the retailer and decreases his competitor's effective demand and therefore, perhaps, increases the retailer's own stocking quantity and decreases the competitor's stocking quantity. In this section we define the precise conditions in which this is indeed the case. Our numerical experiments later demonstrate that such a reaction is not the only possible outcome (see Section 5) .

Another interesting problem is to find out if there is an *optimal* monetary incentive that maximizes a retailer's profit. It is beyond the scope of this paper to consider the optimal setting of the monetary incentive. Such a problem deserves a separate investigation mainly due to technical difficulties: the uniqueness of the solution even without competition can be shown only under rather restrictive conditions, and it is likely that in a competitive situation even such a basic result might not be available.

In this section, our goal is to characterize an optimal operational response to the promotional decisions of the marketing department, or, more precisely, the impact of an incentive to back-order on equilibrium inventory decisions under competition. We analyze this problem for two models only: Models II and IV. In Model I such an analysis is not relevant, since there is no back-ordering. In Model III such an analysis is obscured by the complexity of the demand expression and hence only numerical analysis is provided (see Section 5). For Models II and IV, however, we introduce a dependence between the total demand faced by the retailer and the monetary incentives offered. Exploring the structural properties of supermodular games helps us to answer the question definitively. For the rest of this section, recall that p_i is a monetary incentive offered by retailer i to the customer to persuade the customer to back-order the product.

4.1 The incentive to back-order in Model II

In this model, customer behavior is characterized by the coefficients α_{12} and α_{21} . If the retailer offers a monetary incentive to the customer to increase the proportion of customers willing to back-order the product rather than switch to a competitor, these coefficients depend on the amount of compensation. Specifically, we assume that:

Assumption 5. $\alpha_{ij} = \alpha_{ij}(p_i)$, $\frac{\partial \alpha_{ij}}{\partial p_i} < 0$, $\frac{\partial \alpha_{ji}}{\partial p_i} = 0$.

The assumption is a very natural one: the proportion of customers willing to switch from retailer i to retailer j decreases as the amount of compensation retailer i offers increases. Since customers switch from retailer i to retailer j without knowing if retailer j has inventory to satisfy demand

immediately, it is reasonable to assume that customers who choose retailer i are also unaware of the incentives offered by retailer j and hence α_{ij} is not a function of p_j . We also assume that the first-choice demand for each retailer is not a function of the incentive. This is a plausible assumption if the customer learns about the incentive only after coming to the store (e.g., if the word-of-mouth effect is not too strong). The next proposition shows a condition sufficient for monotonicity of both players' inventories in the monetary incentive. To this end, we define $\tilde{y}_j = -y_j$.

Proposition 10 *Suppose that Assumption 5 holds in Model II. Then:*

- 1) *the players' objective functions are supermodular in (y_i, \tilde{y}_j) ,*
- 2) *player j 's objective function has increasing differences in (\tilde{y}_j, p_i) . Furthermore, player i 's objective function has increasing differences in (y_i, p_i) at any p_i^o satisfying*

$$(1 - \alpha_{ij}(p_i^o)) + (m_i - u_i^B) \left. \frac{\partial \alpha_{ij}}{\partial p_i} \right|_{p_i^o} \geq 0, \quad (9)$$

- 3) *unique optimal inventory policies are such that $\bar{y}_i(p_i)$ is increasing and $\bar{y}_j(p_i)$ is decreasing in p_i at any p_i^o satisfying (9).*

Proof: To prove 10.1, it is sufficient to show that the second-order cross-partial derivatives of the players' objective functions are positive (see Topkis [32])

$$\frac{\partial^2 G_i(y_i, y_j)}{\partial y_i \partial \tilde{y}_j} = \alpha_{ji}(u_i^B + o_i) f_{\bar{D}_i | D_j > -\tilde{y}_j}(y_i) \Pr(D_j > -\tilde{y}_j) > 0, \quad i, j = 1, 2.$$

To prove 10.2, it is sufficient to show supermodularity (see Theorem 2.6.1 in Topkis [32]), which is again verified through the second-order cross-partial derivatives

$$\begin{aligned} \frac{\partial^2 G_i(y_i, y_j)}{\partial y_i \partial p_i} &= \left(\frac{\partial \alpha_{ij}}{\partial p_i} (m_i - u_i^B) - \alpha_{ij} \right) \Pr(D_i > y_i) + \Pr(\bar{D}_i > y_i) \\ &\quad + (u_i^B + o_i) \frac{\partial}{\partial p_i} \Pr(\bar{D}_i > y_i), \quad i, j = 1, 2. \end{aligned}$$

Note that

$$\begin{aligned} \Pr(\bar{D}_i > y_i) &= \Pr(\bar{D}_i > y_i, D_i > y_i) + \Pr(\bar{D}_i > y_i, D_i < y_i) \\ &= \Pr(D_i(1 - \alpha_{ij}) + \alpha_{ji}(D_j - y_j)^+ > y_i(1 - \alpha_{ij}), D_i > y_i) \\ &\quad + \Pr(D_i + \alpha_{ji}(D_j - y_j)^+ > y_i, D_i < y_i) \\ &= \Pr(D_i > y_i) + \Pr(D_i + \alpha_{ji}(D_j - y_j)^+ > y_i, D_i < y_i), \end{aligned}$$

so that $\Pr(\bar{D}_i > y_i)$ is independent of α_{ij} and therefore $\frac{\partial}{\partial p_i} \Pr(\bar{D}_i > y_i) = 0$. We can rewrite the

expression for the derivative as follows:

$$\frac{\partial^2 G_i(y_i, y_j)}{\partial y_i \partial p_i} = \left(\frac{\partial \alpha_{ij}}{\partial p_i} (m_i - u_i^B) - \alpha_{ij} \right) \Pr(D_i > y_i) + \Pr(D_i > y_i) + \Pr(D_i + \alpha_{ji}(D_j - y_j)^+ > y_i, D_i < y_i).$$

Ignoring the last term since it is positive, we can obtain the sufficient condition for the positivity of this derivative:

$$(1 - \alpha_{ij}) + \frac{\partial \alpha_{ij}}{\partial p_i} (m_i - u_i^B) \geq 0.$$

The second cross-partial derivative is trivially positive:

$$\frac{\partial^2 G_j(y_i, y_j)}{\partial \tilde{y}_j \partial p_i} = -(u_j + o_j) \frac{\partial}{\partial p_i} \Pr(\bar{D}_j > -\tilde{y}_j) \geq 0, \quad i, j = 1, 2.$$

Finally, 10.3 follows from 10.1 and 10.2 and Theorem 4.2.2 in Topkis [32]. \square

Proposition 10 demonstrates the effect that offering an incentive to back-order has on the equilibrium stocking policies of the competitors under condition (9): if the marketing department of retailer i decides to offer (or similarly increase) an incentive to customers willing to back-order the product, then the operations department of retailer i should simultaneously increase the stocking quantity of the product. At the same time, retailer j should decrease its stocking quantity. Numerical experiments show that other situations are possible besides those implied by condition (9). Notice also that condition (9) in itself has a nice managerial interpretation: the marginal benefit of an additional unit of the retailer's own inventory is higher at higher values of the incentive as long as condition (9) is satisfied. On the other hand, the marginal benefit of an additional unit of the competitor's inventory is always lower at higher values of the incentive.

4.2 Incentive to back-order in Model IV

As in the previous section, we introduce additional assumptions about the dependence between demand and compensation. First, we assume that the impact of incentives on total demand follows a general functional form, but to keep the solution tractable we assume that there are no cross-effects. Clearly, a higher incentive offered by retailer i should increase the effective demand of retailer i and decrease the effective demand of retailer j . Second, we can expect that an incentive has a decreasing marginal effect. The next assumption formalizes these observations and is an alternative to Assumption 1a.

Assumption 1b. $\bar{D}_i(y_i, y_j) = \eta_i(y_i, y_j) + \xi(p_i, p_j) + \varepsilon_i$, where ε_i is an arbitrarily distributed random variable with density $f_{\varepsilon_i}(\cdot)$ and distribution $F_{\varepsilon_i}(\cdot)$, and $\eta_i(y_i, y_j), \xi(p_i, p_j)$ are positive real-valued functions such that $\bar{D}_i > 0$, $1 \geq \frac{\partial \eta_i}{\partial y_i} \geq 0$, $\frac{\partial \eta_i}{\partial y_j} \leq 0$, $\frac{\partial \xi_i(p_i, p_j)}{\partial p_i} > 0$, $\frac{\partial \xi_i(p_i, p_j)}{\partial p_j} < 0$, $i, j = 1, 2$.

Our main result for Model IV is summarized in the following proposition.

Proposition 11

- 1) Under Assumptions 1b and 2b, the players' objective functions are supermodular in (y_i, \tilde{y}_j) .
- 2) Under Assumptions 1b and 3, the players' objective functions have increasing differences in (y_i, p_i) and (\tilde{y}_j, p_i) , respectively.
- 3) Unique optimal inventory policies are such that $\bar{y}_i(p_i)$ is increasing and $\bar{y}_j(p_i)$ is decreasing in p_i .

Proof:

Result 11.1 follows directly from the assumptions. To prove 11.2, showing increasing differences in this case is equivalent to demonstrating supermodularity of the objective functions in (y_i, p_i) and (\tilde{y}_j, p_i) , respectively. This, again, is verified by taking the second-order cross-partial derivatives. The first derivatives are:

$$\begin{aligned} \frac{\partial G_i(y_i, y_j)}{\partial y_i} &= (m_i - u_i^B) \frac{\partial \eta_i}{\partial y_i} - (u_i^B + o_i) \Pr(\bar{D}_i < y_i) \left(1 - \frac{\partial \eta_i}{\partial y_i}\right) + u_i^B, \\ \frac{\partial G_j(y_i, y_j)}{\partial \tilde{y}_j} &= -(m_j - u_j) \frac{\partial \eta_j}{\partial \tilde{y}_j} + (u_j + o_j) \Pr(\bar{D}_j < -\tilde{y}_j) \left(1 - \frac{\partial \eta_j}{\partial \tilde{y}_j}\right) - u_j, \end{aligned}$$

and, furthermore:

$$\begin{aligned} \frac{\partial^2 G_i(y_i, y_j)}{\partial y_i \partial p_i} &= -\frac{\partial \eta_i}{\partial y_i} - \Pr(\bar{D}_i < y_i) \left(1 - \frac{\partial \eta_i}{\partial y_i}\right) + (u_i^B + o_i) f_{\bar{D}_i}(y_i) \left(1 - \frac{\partial \eta_i}{\partial y_i}\right) \frac{\partial \xi_i}{\partial p_i} + 1 \\ &= \left(1 - \frac{\partial \eta_i}{\partial y_i}\right) \left(\Pr(\bar{D}_i > y_i) + (u_i^B + o_i) f_{\bar{D}_i}(y_i) \frac{\partial \xi_i}{\partial p_i}\right) > 0. \end{aligned}$$

For the second cross-partial derivative we have

$$\frac{\partial G_j(y_i, y_j)}{\partial \tilde{y}_j \partial p_i} = -(u_j + o_j) f_{\bar{D}_j}(-\tilde{y}_j) \left(1 + \frac{\partial \eta_j}{\partial \tilde{y}_j}\right) \frac{\partial \xi_j}{\partial p_i} > 0,$$

and the proof is complete. Finally, 11.3 follows from Lemmas 1 and 2 and Theorem 4.2.2 in Topkis [32]. \square

5 Numerical experiments

In our numerical experiments, we investigate answers to the following questions:

1. What is the effect of backlogging on inventories and profits?
2. How much does accounting for different back-ordering behaviors in terms of differences affect profits and inventories across models?

3. What is the impact of problem parameters on profits and differences in profits across models?
4. What is the impact of offering customers an incentive to back-order?

To simplify the comparison, we work with single-period symmetric models (unless otherwise noted) since we have demonstrated that there is a stationary solution. Symmetry allows us to reduce the number of parameters. The following common parameters are used in this section: $r = 10$, $c = 5$, $\beta = 0.9$ and $D \sim N(100, \sigma)$ (truncated at zero with probability mass added to zero). We have also verified that insights point out in the same direction with nonsymmetric parameters as well.

A comparison of equilibrium inventories and profits

First, we want to compare inventories and profits in three models to see if consideration of customer back-ordering behavior is noticeable and if there are any consistent patterns among models in addition to those shown in Proposition 5. We have performed an extensive numerical study using the following set of problem parameters: $p \in \{1, 2, 3, 4, 5\}$, $h \in \{0.5, 1.5, 2.5, 3.5, 4.5, 5.5\}$, $\sigma \in \{20, 40, 60, 80, 100, 120\}$, $\rho \in \{-0.9, -0.6, -0.3, 0, 0.3, 0.6, 0.9\}$ and $\alpha \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$. In total, we analyzed $5 \times 6 \times 6 \times 7 \times 5 = 6300$ problem instances, capturing most reasonable parameter combinations. We report averages and standard deviations of inventories/profits for each model in Table 1. Evidently, different assumptions about customer back-ordering behavior result in inventories and profits that differ quite drastically from model to model, thus demonstrating that it is important to consider customer back-ordering behavior under competition.

	Model I	Model II	Model III
Mean (inventory)	128.43	113.94	98.72
Standard deviation (inventory)	28.33	31.80	37.99
Mean (profit)	355.52	406.35	417.54
Standard deviation (profit)	80.61	62.23	64.01

Table 1. Averages and standard deviations of profits/inventories.

Next, we compare differences in inventories and profits among models. As we proved in Proposition 5, Model I results in higher inventory and lower profit than Model II does. Furthermore, we find that, in *all* experiments, Model II results in higher inventory and lower profit than Model III does. Relative differences in profits/inventories are reported in Tables 2 and 3. For example, the lower left cell in Table 2 (row “Model III” and column “Model I”) shows that the average value of $100\% \times (y^{III} - y^I) / y^{III}$ was -30.1% (superscripts denote model number).

	Model I	Model II	Model III
Model I	X	11.28%	23.14%
Model II	-12.71%	X	13.36%
Model III	-30.10%	-15.42%	X

Table 2. Differences in inventories.

	Model I	Model II	Model III
Model I	X	-14.30%	-17.45%
Model II	12.51%	X	-2.75%
Model III	14.85%	2.68%	X

Table 3. Differences in profits.

Observe that the difference between Model I and the other two models is quite large, which is explained by the lost-sales assumption in Model I. The difference between Models II and III is large in inventories but much smaller in terms of profits (there are, however, instances in which the difference exceeds 100%). A plausible explanation for this last observation is as follows: it has been shown in the literature (see, for example, Lippman and McCardle [18], Mahajan and van Ryzin [19] and Netessine and Rudi [24]) that competition in a similar problem setting typically leads to overstocking inventory. As we explained earlier, in Model III companies exist in a less competitive environment since customers switch between two companies *only if* the competitor has the product in stock, while in Model II the customers switch more often. Hence, due to competition, in Model II companies establish higher inventory levels, reducing the likelihood that customers will switch. In terms of equilibrium profits, Model I is naturally the least profitable since there is no back-ordering. Between the other two models, Model III generates higher profits since, as we just described, it is a less competitive environment and the companies tend to overstock less and as a result suffer less from the detrimental effects of competition. Hence, a reduction in the level of competition (by moving from Model II to Model III) results in higher profits for both firms. We discuss the implications of this comparison shortly.

The impact of problem parameters on profits

Next, we wish to understand the impact of problem parameters on profits to gain further insight into the impact of customer back-ordering behavior and determine if insights from single-period models with lost sales continue to hold in the presence of back-ordering. Using the same set of data as above, we study the impact of ρ , σ , p , h and α on absolute values of profits as well as on pair-wise differences in profits for three models (Table 4).

Parameter	G^I	G^{II}	G^{III}	$G^{II} - G^I$	$G^{III} - G^I$	$G^{III} - G^{II}$
ρ	↓	↓	↓	↑	↑	↑
σ	↓	↓	↓	↑	↑	↑
p	0	↓	↓	↓	↓	↓
h	↓	↓	↓	↑	↑	↑
α	↑	↑↓	↑	↓	↓	↑

Table 4. The impact of problem parameters on profits and absolute values of profit differences.

As one expects, profits in all models decrease in ρ , σ , p and h . Several previous papers (see Netessine and Zhang [25] and Anupindi and Bassok [3]) analyzing models with lost sales have also demonstrated that the profits of both players increase in α . This happens because the proportion of customers $(1 - \alpha)$ who do not switch is lost for both firms if they are not immediately satisfied. Hence, increasing α essentially increases the total demand that both firms face and therefore increases profits. We found that this result generally holds in Models I and III but not in Model II. In fact, in many instances the profits of both players in Model II decrease in α . To understand why this happens, we recall that Models II and III include full back-ordering, so increasing α does not change the total demand faced by the two firms. Rather, it benefits firms because the total

demand that the two firms face *in the current period* increases (so that back-order/holding costs decrease). In other words, increasing α reduces the number of customers who back-order with the retailer of the first choice instead of switching when a stock-out occurs. In Model II the customer who is not satisfied by his first-choice retailer is more likely to switch than in Model III. Hence, the negative impact of increasing α in Model II is stronger than in Model III. As a result, in Model III profits generally increase in α while in Model II the result is more ambiguous. The result for Model II does not correspond with the standard result found in the substitution literature that utilizes a single-period framework. To further illustrate this behavior, in Figure 1 we provide one specific numerical example (in which $h = 1$, $p = 3$, $\sigma = 50$ and $\rho = 0$). Note that profit in Model II is indeed decreasing in α for large values of α .




Figure 1: Equilibrium inventory (left) and profits (right) for the three models.

The impact of problem parameters on differences in profits

Sometimes firms are in a position to affect customer behavior. For example, by training service representatives better, the firm may be able to increase the number of customers who back-order instead of leaving in the case of a stock-out. For example, Anderson et al. [1] illustrate how various ways of handling stock-out situations can dramatically affect customers' back-ordering decisions and hence help a firm transition from Model I to Model II. Naturally, such measures are costly, so we want to understand when influencing customer back-ordering behavior is worth the expense. Furthermore, an information system that would allow customers to "observe" inventories at competing retailers would help firm to transition from Model II to Model III with corresponding profit increase. Again, since such systems are quite costly and challenging to implement in practice, we want to be able to describe situations in which their adoption is worthwhile.

From Table 4, we observe that higher ρ , σ , and h consistently lead to wider profit gaps. This result is quite intuitive: an increase in these parameters increases the cost of any mismatch between demand and supply, thereby accentuating the differences in profits among the three models. At the same time, increasing p leads to a decrease in profit gaps. This happens because an increase in p induces firms to stock more and hence the substitution effect becomes less pronounced since

firms are rarely out of stock. Finally, increasing α decreases the gap between Model I and the other two models but increases the gap between Models II and III (see also Figure 1). The first effect is due to the fact that in Model I increasing α has a stronger positive effect than in the other two models; customers in Model I have no alternative to back-ordering and hence increasing product substitutability significantly enhances profits. The increase in the gap between Models II and III is a direct result of the earlier finding that the profit in Model II is almost invariant to changes in α while the profit in Model III increases. To summarize, a transition from Model I to Model II (e.g., through better customer service) is most beneficial under conditions of high correlation, high demand uncertainty, high holding costs, low back-order penalty and low substitutability. A transition from Model II to Model III (e.g., through the adoption of the information system) is most beneficial under conditions of high correlation, high demand uncertainty, high holding costs, high substitutability and low back-order penalty.

The impact of the incentive to back-order

We now investigate the impact of the retailer's ability to offer customers a monetary incentive in order to reduce switching between retailers. We assume that in Models II and III, retailer i is able to reduce the proportion of switching customers α_{12} by offering a higher monetary incentive p_1 . We assume that there is a linear relationship between α_{12} and p_1 , $\alpha_{12} = 1 - 0.2 \times p_1$, so that $p_1 \in [0, 5]$ and $\alpha_{12} \in [0, 1]$. We also fix $\alpha_{21} = .5$ and $p_2 = 3$. For Model II, it is readily verified that condition (9) holds as long as $p_1 \geq 2.75$, i.e., this inequality defines values where we are guaranteed to have y_1 increasing and y_2 decreasing in p_1 . Figure 2 shows resulting equilibrium inventories and profits in Model II.




Figure 2: Equilibrium inventory (left) and profit (right) as a function of incentive in Model II.

We observe that as retailer i increases the incentive p_1 , the second retailer's inventory *and* profit go down. Furthermore, retailer i 's inventory appears to be convex and profit appears to be concave in the amount of the incentive, with the maximum profit achieved at the same time as minimum inventory is achieved. Notice that retailer j 's profit is much more sensitive to the incentive than retailer i 's profit. Furthermore, note that condition (9) is instrumental in defining the region in

which inventories are monotone increasing in the retailer’s own incentive and decreasing in the competitor’s incentive. Outside of this region we still observe monotonicity, but the direction is reversed. We now turn to Model III (see Figure 3).




Figure 3: Equilibrium inventory (left) and profit (right) as a function of incentive in Model III.

The behavior of inventory and profits in Model III is quite different. First, the equilibrium inventory levels are monotone in the incentive with the first retailer’s inventory increasing and the second retailer’s inventory decreasing. Such a behavior is, perhaps, more easily anticipated than in Model II. Secondly, in this particular instance, both players’ profits are decreasing, meaning that under the given assumptions it is optimal not to offer any incentive at all; by offering an incentive, retailer i lowers both players’ profitability. We offer the following explanation: since competition in Model III is so much lower than in Model II, the further introduction of the monetary incentive is simply not profitable, while in Model II the incentive is a valuable tool.

6 Concluding remarks

Traditionally, the operations literature considers rather simplistic customer behavior with respect to stock-out situations: the product is either back-ordered or the sale is lost. We have demonstrated how the presence of a competing retailer selling a substitute may complicate the environment, since customers face the choice of back-ordering with either one of two retailers. As a result, the effective demands faced by either retailer are rather complex functions of the competitor’s inventory levels. Three specific back-ordering models that account for substitution in the case of a stock-out are formulated in this paper. For these models we illustrate how customer switching behavior affects optimal inventory policies. We also consider a fourth model that allows us to incorporate effects such how inventories and service levels stimulate demand. Structural results include conditions for the existence and uniqueness of a Nash equilibrium and tractable analytical first-order optimality conditions for all models, as well as comparative statics results.

Previous research on horizontal competition has demonstrated that competing retailers tend to

overstock products to prevent customers from switching to a competitor. In this paper we demonstrate that in single-period models with lost sales the overstocking effect might be overstated relative to a multi-period setting with back-ordering. The option to back-order the product reduces the overstocking effect, because customers are not necessarily lost. We show that the differences in profits and inventories under various back-ordering scenarios are significant, so accounting for back-ordering behavior has important ramifications. Our results also indicate that increasing the proportion of customers who are willing to switch between retailers may not increase either retailer's profits, as is the case in single-period models without back-ordering. In particular, when customers are able to back-order with the retailer of their second choice (Model II), increasing the number of switching customers at the same time reduces the number of customers who are willing to back-order with the first-choice firm, and this negative effect dominates. In practice managers should consider the interaction of these two effects.

We find that under the lost-sales assumption (Model I), firms always stock more and earn less than when customers back-order with the retailer of their second choice (Model II). In practice, firms can affect a customer's propensity to back-order through better customer service and therefore transition from Model I to Model II. We demonstrate that this transition is most effective under conditions of high demand uncertainty, high correlation, high holding costs but low back-order penalty and low substitution rates. Furthermore, we find that when customers back-order with the retailer of their second choice (Model II) both firms stock more and earn less than when customers back-order with the retailer of their first choice (Model III). This finding leads to the counterintuitive insight that revealing their inventories to the competitor's customers may be beneficial for competing firms. In practice, an information system that makes competitors' inventories visible enables a transition from Model II to Model III. We demonstrate that investment into such a system is most prudent under conditions of high demand uncertainty, high correlation, high holding cost, high substitution rates and low back-order penalty. From a practical point of view, there may be difficulties in implementing the information system needed for inventory transparency because the first-choice retailer has an incentive to report to the customer (e.g., by manipulating the information system) that the other retailer is out of stock. These difficulties may require administration of this information system by a third party, e.g., an automobile manufacturer could require competing dealers to adopt such a system and then supervise its use.

Our analysis demonstrates that a firm's stocking decisions and profitability greatly depend on its ability to retain customers, i.e., to induce them to back-order the product rather than switch to a competitor. If, however, a retailer decides to offer a monetary incentive aimed at retaining more customers, a corresponding correction should be made to its inventory policy. Namely, we provide conditions in which a retailer's own inventory should be adjusted up while the competitor's inventory should be adjusted down. Our numerical experiments demonstrate notable differences in profitability among the models. Inventory visibility in Model III results in a higher profitability for the retailers; in this model the customers' ability to observe the competitor's inventory reduces com-

petition since customers switch between retailers *only* when the competitor has inventory. Hence, this model generates the highest profits while having the lowest equilibrium inventories. We also find that one should exercise caution when offering a monetary incentive to retain customers: when inventories are visible to customers (Model III), such an incentive might be detrimental to both competitors.

Our analysis gives rise to further questions regarding customer back-ordering behavior under competition. First, it is important to be able to calculate the optimal amount of the incentive to be offered. As mentioned earlier, such a problem merits a separate study and is not a primary focus of this paper. Second, our model takes demand for the product from first-choice customers as a given. A more realistic approach might be to introduce a customer choice model which describes how customers decide on the first-choice retailer over the second-choice retailer. Also, we constrained the effects that inventory has on demand into a single period, i.e., customers do not have “memory” that allows them to base their choice of a retailer on previous experience. In this sense, Gans [11] provides a more advanced model. Finally, we do not address the important question of coordinating the competing retailers as car manufacturers often attempt to do through centralized information systems. Again, such a question is an important one and merits a separate study.

References

- [1] Anderson, E., G.J. Fitzsimons and D. Simester. 2003. Mitigating the costs of stockouts. Working paper, University of Chicago.
- [2] Andraski, J.C. and J. Haedicke. 2003. CPFR: time for the breakthrough? *Supply Chain Management Review*, May/June, 54-60.
- [3] Anupindi, R. and Y. Bassok. 1999. Centralization of stocks: retailers vs. manufacturers. *Management Science*, Vol. 45, 178-191.
- [4] Avsar, Z.M. and M. Baykal-Gursoy. 2002. Inventory control under substitutable demand: a stochastic game application. *Naval Research Logistics*, Vol. 49, 359-375.
- [5] Balakrishnan, A., M.S. Pangburn and E. Stavrulaki. 2004. “Stack them high, let ’em fly”: lot-sizing policies when inventories stimulate demand. *Management Science*, Vol. 50, 630-644.
- [6] Bernstein, F. and A. Federgruen. 2004. Dynamic inventory and pricing models for competing retailers. *Naval Research Logistics*, Vol. 51, 258-274.
- [7] Cachon, G. P. and P.H. Zipkin. 1999. Competitive and cooperative inventory policies in a two-stage supply chain. *Management Science*, Vol. 45, No. 7, 936-953.
- [8] Chen, F. 1999. Decentralized supply chains subject to information delays. *Management Science*, Vol. 45, 1076-1090.
- [9] Cheung, K.L. 1998. A continuous review inventory model with a time discount. *IIE Transactions*, Vol. 30, 747-757.

- [10] DeCroix, G.A. and A. Arreola-Risa. 1998. On offering economic incentives to backorder. *IIE Transactions*, Vol. 30, 715-721.
- [11] Gans, N. 2002. Customer loyalty and supplier quality competition. *Management Science*, Vol. 48, 207-221.
- [12] Heyman, D.P. and M.J. Sobel. 1984. *Stochastic models in Operations Research*, Vol.II: Stochastic Optimization. McGraw-Hill.
- [13] Karjalainen, R. 1992. The newsboy game. Working paper, University of Pennsylvania.
- [14] Kirman, A.P. and M.J. Sobel. 1974. Dynamic oligopoly with inventories. *Econometrica*, Vol. 42, 279-287.
- [15] Lee, H., and S. Whang. 1999. Decentralized multi-echelon supply chains: incentives and information. *Management Science*, Vol. 45, 633-642.
- [16] Li, Q. and A. Ha. 2003. Accurate response, reactive capacity and inventory competition. Working paper, The Hong Kong University of Science and Technology.
- [17] Lilien, G.L., P. Kotler and K.S. Moorthy. 1992. *Marketing models*. Prentice Hall, New Jersey.
- [18] Lippman, S.A. and K.F. McCardle. 1997. The competitive newsboy. *Operations Research*, Vol. 45, 54-65.
- [19] Mahajan, S. and G. van Ryzin. 1999. Retail inventories and consumer choice. In *Quantitative models for supply chain management*, S. Tayur, R. Ganeshan and M. Magazine (eds.), Kluwer Academic Publishers, Massachusetts.
- [20] Mahajan, S. and G. van Ryzin. 2001. Inventory competition under dynamic consumer choice. *Operations Research*, Vol. 49, 646-657.
- [21] Moinzadeh, K. and C. Ingene. 1993. An inventory model of immediate and delayed delivery. *Management Science*, Vol. 39, No. 5, 536-548.
- [22] Moulin, H. 1986. *Game theory for the social sciences*. New York University Press, New York.
- [23] Nagarajan, M. and S. Rajagopalan. 2003. Optimal policies for substitutable goods. Working paper, Marshall School of Business, U.S.C.
- [24] Netessine, S. and N. Rudi. 2003. Centralized and competitive inventory models with demand substitution. *Operations Research*, Vol. 51, 329-335.
- [25] Netessine, S. and F. Zhang. 2005. Positive vs. negative externalities in inventory management: implications for supply chain design, *Manufacturing and Service Operations Management*, Vol. 7, 58-73.
- [26] Parlar, M. 1988. Game theoretic analysis of the substitutable product inventory problem with random demands. *Naval Research Logistics*, Vol. 35, 397-409.
- [27] Porteus, E. 2000. Responsibility tokens in supply chain management. *Manufacturing and Service Operations Management*, Vol. 2, 203-219.
- [28] Rockafellar, T.R. 1972. *Convex analysis*. Princeton University Press, New Jersey.

- [29] Rudi, N. 2000. Some models of risk pooling. PhD dissertation, University of Pennsylvania.
- [30] Shapley, L. 1953. Stochastic games. *Proceedings of the National Academy of Sciences*, Vol. 39, 1095-1100.
- [31] Topkis, D.M. 1998. Minimizing a submodular function on a lattice. *Operations Research*, Vol. 26, 305-321.
- [32] Topkis, D.M. 1998. *Supermodularity and complementarity*. Princeton University Press, New Jersey.
- [33] Wang, Y. and Y. Gerchak. 2001. Supply chain coordination when demand is shelf-space dependent. *Manufacturing and Service Operations Management*, Vol. 3, No. 1, 82-87.
- [34] Wang, Q. and M. Parlar. 1994. A three-person game theory model arising in stochastic inventory control theory. *European Journal of Operational Research*, Vol. 76, 83-97.
- [35] Wolfe, H.B. 1968. A model for control of style merchandise. *Industrial Management Review*, Vol. 9, No. 2, 69-82.