INVENTORY COMPETITION UNDER DYNAMIC CONSUMER CHOICE

SIDDHARTH MAHAJAN

Fuqua School of Business, Duke University, Durham, North Carolina 27706

GARRETT VAN RYZIN

Graduate School of Business, Columbia University, New York, New York 10027

(Received May 1999; revision received February 2000; accepted November 2000)

We analyze a model of inventory competition among n firms that provide competing, substitutable goods. Each firm chooses initial inventory levels for their good in a single period (newsboy-like) inventory model. Customers choose dynamically based on current availability, so the inventory levels at one firm affect the demand of all competing firms. This creates a strategic interaction among the firms' inventory decisions.

Our work extends earlier work on variations of this problem by Karjalainen (1992), Lippman and McCardle (1997) and Parlar (1988). Specifically, we model demand in a more realistic way as a stochastic sequence of heterogeneous consumers who choose dynamically from among the available goods (or choose not to purchase) based on a utility maximization criterion. We also use a sample path analysis, so minimal assumptions are imposed on this demand process. We characterize the Nash equilibrium of the resulting stocking game and prove it is unique in the symmetric case. We show there is a bias toward overstocking caused by competition; specifically, reducing the quantity stocked at any equilibrium of the game increases total system profits, and at any joint-optimal set of stocking levels, each firm has an individual incentive to increase its own stock. For the symmetric case, we show that as the number of competing firms increases, the overstocking becomes so severe that total system (and individual firm) profits approach zero. Finally, we propose a stochastic gradient algorithm for computing equilibria and provide several numerical examples.

1. INTRODUCTION

Consumers typically choose among competing goods based on criteria such as quality and price. However, the local availability of products or services is also an important factor in many consumer choice decisions. Rather than searching other locations or delaying their purchase, consumers may opt to substitute a competing good if their preferred good is unavailable. For example, Anupindi et al. (1998) report significant levels of substitution in an empirical test of vending machine products, and they cite a Food Marketing Institute (1993) finding that 82%–88% of grocery customers would substitute if their favorite brand-size was not available. Consumer substitution based on availability of either goods or capacity also occurs in consumer durable markets, in service markets, and in industrial raw materials markets. (See Dion et al. 1989, Jeuland 1979, and Mason and Wilkinson 1977.)

Dynamic consumer choice behavior creates strategic interactions among the stocking (or capacity) decisions of competing firms. This occurs because the inventory stocked (or capacity provided) by one firm affects the realized demand of its competitors. Examples of this sort of availability competition include retailers in close proximity who sell the same branded merchandise (e.g., consumer electronics and apparel); vendor-managed inventory (VMI), where competing manufacturers in a merchandise category independently manage the wholesale or retail stock of their own products (e.g., grocery items); competing airlines that fly the same route, where the availability of discount seats on one carrier affects the demand for discount seats on its competitors' flights.

Traditionally, inventory theory has not considered the effects of either consumer choice behavior or the resulting strategic effect on inventory decisions. (See Graves et al. 1993.) Recently, however, both phenomena have been addressed in the research literature. Work by Mahajan and van Ryzin (1998), Noonan (1995), Smith and Agrawal (2000), and van Ryzin and Mahajan (1999) has looked at profit maximizing decisions for a monopolist stocking a category of substitutable products. (See Mahajan and van Ryzin 1999 for a survey.) These works, however, do not examine the strategic effects resulting from consumer choice behavior. Availability as a strategic variable has been considered in the marketing literature by Balachander and Farquhar (1994), but this work does not consider inventory decisions directly.

Strategic stocking decisions have been considered in only a limited number of works, first by Parlar (1988) and then later by Karjalainen (1992) and Lippman and McCardle (1997). Our work builds on and extends this line of research. In this literature, as in our work, prices are assumed exogenous, and therefore price competition is

Subject classifications: Inventory/production: policies—marketing/pricing. Marketing: buyer behavior. Games/group decisions: noncooperative.
Area of review: LOGISTICS & SUPPLY CHAIN OPERATIONS.

0030-364X/01/4905-0646 $05.00
1526-5463 electronic ISSN

ignored. Primarily, this is done to simplify the problem and to focus on the effects of inventory competition. However, there are certainly cases (e.g., vending machines, airline tickets) in which inventory levels are a firm's primary decision at a disaggregate level (e.g., individual machine stock levels, capacity allocations for a specific flight departure) and prices are determined at a more aggregate level. (See Balachander and Farquhar 1994 for an analysis of availability and price competition in a simpler setting.)

Parlar (1988) considers a two-firm inventory game in which each firm manages the inventory of its good and the goods are substitutable. The goods receive random demand. A deterministic fraction of excess demand for each good substitutes to the alternative, if the alternative good has excess stock. A single-period inventory model is considered. The existence and uniqueness of a Nash equilibrium is shown. Also, it is shown that cooperation between firms improves the joint profit over that achieved under competition. Karjalainen (1992) considers a model similar to Parlar's for the n-firm case.

Lippman and McCardle (1997) analyze both a duopoly and a n-firm (oligopoly) model of inventory competition. They model aggregate demand for all firms in the market as a random variable that does not change with the number of firms. Each firm chooses the level of inventory to stock. Demand for each firm results from an initial allocation and a reallocation. These rules for splitting demand are known to all the firms. Rules for initial allocation are primarily of two kinds, deterministic and random. In the deterministic rule, a specified fraction of total demand is allocated to each firm. In the random allocation rule, how demand is to be split is determined based on the outcome of a random variable. Any excess demand at one firm is then reallocated among the other firms. As an example of a reallocation rule, the authors consider "herd behavior." One particular firm is chosen at random, and all excess demand is allocated to that firm. If excess demand is met, the process stops. If not, then another firm is chosen at random and the remaining excess is allocated. It is through these reallocation rules that the inventory decision of one firm affects the profits of the other firms.

The authors show existence of equilibrium under the assumption that the effective demand for each firm (i.e., demand after reallocation has occurred) is stochastically decreasing in the inventory levels of the other firms. They show uniqueness of the equilibrium in the symmetric case under continuity and monotonicity assumptions on the effective demand for each firm. For the two-firm case, if all excess demand is reallocated they show that competition leads to higher total inventory in the system. Finally, they show that profits tend to zero as the number of firms increases under herd behavior reallocation.

Our results closely parallel Lippman and McCardle's (1997), but we extend them in several important ways. First, our demand model is arguably more realistic and general than the somewhat stylized allocation and reallocation rules used by Lippman and McCardle. (An exception is the basic existence of a Nash equilibrium, which Lippman and McCardle show holds when each firm's demand is stochastically decreasing in the inventory levels of competing firms. Our demand model satisfies this condition. Other results in Lippman and McCardle require more restricted assumptions.) Such rules, while capturing certain substitution phenomena, are not the most natural and satisfying model of consumer choice behavior. In addition, Lippman and McCardle's results often depend on which particular allocation rules are used (e.g., herd behavior is required for the zero-profit result).

In contrast, demand for goods in our model is the result of the individual choice behavior of a sequence of heterogeneous consumers. These individual choices, which are made dynamically based on the on-hand stock the customer sees, determine the evolution of inventory levels over time and, consequently, the profits earned by each firm. Individual consumer choice itself is based on random utility maximization, which is a well-established model in the economics literature for how rational consumers choose from a discrete set of alternatives (Anderson et al. 1992). Moreover, because our analysis is based on sample path properties of the profit functions, the results apply to any such demand process that satisfies only some mild regularity conditions. For example, we show the uniqueness of the equilibrium in the symmetric case under a very general notion of symmetry. Thus, in a very general and realistic framework, we model both the dynamic choice behavior of consumers and its effect on inventory and firm profits.

Second, we obtain some new structural insights. We provide comparative statics for the general, n-firm asymmetric case that show: (1) at any equilibrium set of stocking levels, total system profits can be improved if all firm reduce their inventory; and (2) at any joint optimum, each firm as an incentive to increase its own stocking levels. In particular, these results show that the tendency toward "competitive overstocking" shown by Lippman and McCardle (1997) for the duopoly case extends to our general oligopoly case. We also provide conditions under which industry profits tend to zero as the number of competing firms tends to infinity. Again, this extends the important result that the industry approaches the competitive (zero-profit) limit as the number of firms increases.

Lastly, our model and analysis approach leads to a computationally efficient (albeit not provably convergent) method to find equilibrium inventory levels based on modifications of the stochastic gradient algorithms developed in Mahajan and van Ryzin (1998). We illustrate the method with some numerical examples.

This paper is organized as follows. In §2 we describe the inventory game, the essential components of which are the customer choice process, a continuous model of inventory, and our sample path expressions for the profit functions of each firm. Preliminary results are contained in §3. In §4, we present our main results, which include demonstrating the existence of an equilibrium and proving its uniqueness for the case of symmetric goods. We also make comparisons between the Nash equilibria and a joint optimum inventory level vector, which is the inventory level decision assuming all goods are managed by a central decision maker. In §5, we analyze the symmetric case as the number of firms increases and prove that system profits converge to zero. Finally, in §6, we describe a numerical algorithm and illustrate its performance on several examples. Our conclusions are presented in §7.

2. A MODEL OF INVENTORY COMPETITION

2.1. Notation

We begin by introducing some notation. The set of natural numbers (nonnegative integers) is denoted by Z and the set N denotes $\{1, \ldots, n\}$. The notation $|A|$ denotes the cardinality of a finite set A. All vectors are in $\Re^n$ unless otherwise specified. The notation y^T stands for the transpose of a vector y. Where possible, components of vectors are denoted by superscripts while subscripts denote elements of a sequence. The set For example, x_t^j denotes the jth component of a vector x_t in a sequence $\{x_t\colon t \geqslant 1\}$. $\mathbf{I}$ denotes the identity matrix and $e^i \in Z^n$, $1 \leqslant i \leqslant n$ denotes the ith unit vector; that is, a column vector with a 1 in the ith position and a 0 elsewhere; we also extend this definition and let e^0 denote a column of all zeros. We use $\mathbf{1}$ to denote the column vector with a 1 as every element; a.s. means *almost surely*, and c.d.f. is short for *cumulative distribution function.*

2.2. Description of the Game

We consider a competitive version of a single-period (newsboy like) inventory model. There are n firms, and each firm stocks inventory of a single good. As described below, consumers choose among the n goods, so the inventory levels of one firm affect the profit earned by its competitors. Firm j has a selling price of p^j and a procurement cost of c^j, which are assumed exogenous.

Firm j's strategy is to decide the initial inventory level of good j, denoted x^j. We let $x = \{x^j\colon j \in N\}$ and $x^{-j} = \{x^1, \ldots, x^{j-1}, x^{j+1}, \ldots, x^n\}$ denote the decision of all firms other than j, so that $x = (x^j, x^{-j})$. Each firm has perfect information about the decisions of the other firms. Without loss of generality, there is no salvage value for any of the goods.

As in Mahajan and van Ryzin (1998), we use a sample path description of the inventory and sales process. Let T denote the number of customers on a sample path. Each customer $t = 1, \ldots, T$ chooses from the goods that are in stock when he/she arrives. We note that customers do not choose the time at which they arrive in our model; they simply arrive exogenously and then choose which product to buy, or they choose not to buy at all. One could imagine models in which a customer's timing is also part of his/her decision, but to keep the model tractable we ignore this possibility.

Let $x_t = (x_t^1, \ldots, x_t^n)$ denote the vector of inventory levels observed by customer t and note that $x_1 = x$, where x is the initial stocking decision mentioned above. For any real inventory vector y, let $S(y) = \{j \cup \{0\}\colon y_j > 0\}$ denote the set of goods with positive inventory (the set of *in-stock* goods) together with the *no-purchase* option, denoted by the element 0. Customer t can only make a choice $j \in S(x_t)$. Because inventory levels are nonincreasing over time, we have that $S(x_{t+1}) \subseteq S(x_t)$.

A customer's choice is based on utility maximization: Each customer t assigns a utility U_t^j to the goods $j \in N$ and to the no-purchase option, U_t^0. Let $U_t = (U_t^0, U_t^1, \ldots, U_t^n)$ denote the vector of utilities assigned by customer t. Based on the inventory level x_t and utility vector U_t, customer t makes the choice that maximizes his/her utility. Each customer t requires a certain quantity of good, Q_t, which could vary from customer to customer and could be nonintegral. The customer purchases from the inventory of the most preferred good first. If this good runs out, the customer purchases the remainder from the inventory of the second most preferred good until it runs out, and so on. This process continues until either the customer's requirement, Q_t, is met or the inventory of all goods valued higher than the no-purchase utility is exhausted. Note that if customers all demand unit quantities ($Q_t = 1$) and firms stock integral quantities of each good, then each customer either purchases one unit of some good or does not purchase, depending on the on-hand stock. The model above simply extends this natural integral case to allow for continuous inventory and demand.

While we do not directly model a customer's search cost in this framework, certain types of search costs can be captured through an appropriate choice of utilities. For example, suppose a customer has a preferred firm and is willing to search one additional firm only if the preferred firm has no inventory. Then one could construct a utility model in which U_t has only two positive elements, one with the utility of the primary firm and one with the utility (net of search cost) for the secondary firm. However, allowing a truly dynamic search process, in which customers sequentially select firms, inspect them for inventory (perhaps at a cost), and then decide to search more or stop, is beyond the scope of our choice model.

Let $\omega = \{(U_1, Q_1), (U_2, Q_2), \ldots, (U_T, Q_T)\}$, denote a sample path from some probability space $(\Omega, \mathcal{F}, P)$. The only assumptions made on this space are that the sequence is bounded w.p.1, i.e., $P(T \leqslant C) = 1$ for some finite C, and that each customer t makes a unique choice w.p.1.; that is

$$P(U_t^i \neq U_t^k) = 1 \quad \text{for all } i \neq k, \;\; i, k \in N \cup \{0\}.$$

This latter assumption is satisfied, for example, if the utilities have continuous distributions. Each firm does not know the particular realization ω but does know the probability measure P. So we think of P as the common knowledge each firm has of future demand.

Because consumer choice decisions depend on which goods are in stock, the profits of each firm depend in a nontrivial way on the stock levels of other firms. To express the profit of each firm concisely, let $\eta^j(x,\omega)$ denote the number of sales of good j made on the sample path ω given initial inventory levels x. Let $\eta(x,\omega) = (\eta^1(x,\omega),\dots,\eta^n(x,\omega))$. Then the sample path profit of firm j, denoted $\pi^j(x,\omega)$, is given by

$$\pi^j(x,\omega) = p^j\eta^j(x,\omega) - c^j x^j. \tag{1}$$

We denote the expected profit of firm j by $\pi^j(x) = E\pi^j(x,\omega)$. Therefore, the objective of each firm $j \in N$ is to solve

$$\max_{x^j \geqslant 0} \pi^j(x^j, x^{-j}), \tag{2}$$

where recall $x = (x^j, x^{-j})$.

The solution concept we use is the Nash equilibrium (1951). Specifically, a vector of inventories, $\bar{x}$, is called a *Nash equilibrium in pure strategies* (or simply Nash equilibrium) if it satisfies

$$\pi^j(\bar{x}^j, \bar{x}^{-j}) = \max_{x^j \geqslant 0} \pi^j(x^j, \bar{x}^{-j}) \quad \forall\, j \in N. \tag{3}$$

In words, at a Nash equilibrium, no firm as an incentive to unilaterally change its stock level, provided all competing firms retain their current stock levels.

Finally, let

$$\pi(x) = \sum_{j=1}^{n} \pi^j(x)$$

denote the combined profit of all firms in the game. Let x^* denote a global maximum of the combined profit function, viz

$$\pi(x^*) = \max_{x \geqslant 0} \pi(x).$$

This monopoly version of the problem is analyzed in Mahajan and van Ryzin (1998). We will compare equilibrium and joint optimal inventory levels below.

Because the only complicated quantities in Equation (1) are the functions $\eta^j(x,\omega)$, we focus on understanding their properties. To do so, it is convenient to consider a recursive formulation of the problem.

2.3. A Recursive Formulation of Profit Functions

We first define a system function, $f(\cdot)$, which describes how the inventory evolves over time. Let the components of the vector U_t, be ordered so that $U_t^{[1]} > U_t^{[2]} > \cdots > U_t^{[n+1]}$. Let m denote the number of goods with utilities higher than the no purchase option. That is,

$$U_t^{[1]} > \cdots > U_t^{[m]} > U_t^0 = U_t^{[m+1]} > \cdots > U_t^{[n+1]}.$$

Let $b(j)$ denote the rank assigned to good j by customer t, with 1 being the highest rank. That is,

$$b(j) = k, \quad \text{if } U_t^j = U_t^{[k]}.$$

As before, let $x_t \in \Re^n$ denote the inventory vector and let

$$x_t^{[k]} = x_t^j, \quad \text{if } U_t^j = U_t^{[k]}.$$

Finally, let $f^j(\cdot)$ denote the jth component of the system function $f(\cdot)$. Then

$$f^j(x_t,U_t,Q_t) = \begin{cases} (x_t^{[1]} - Q_t)^+ & b(j)=1, 1\leqslant m \\ (x_t^{[b(j)]} + \cdots + x_t^{[1]} - Q_t)^+ & \\ \quad -(x_t^{[b(j)-1]} + \cdots + x_t^{[1]} - Q_t)^+ & 1 < b(j) \leqslant m \\ x_t^j & b(j) > m. \end{cases} \tag{4}$$

That is, if $b(j) > m$, then it implies that the no purchase option is preferred to good j, i.e. $U_t^j < U_t^0$, so no inventory of good j is consumed. If $b(j) \leqslant m$, then the inventory of good j that is consumed is given by the first two expressions in Equation (4).

Next, define a sequence of sales-to-go functions, $\eta_t^j(x_t,\omega)$ for $t = 1,\dots,T$ via the recursion

$$\eta_t^j(x_t,\omega) = x_t^j - f^j(x_t,U_t,Q_t) + \eta_{t+1}^j(x_{t+1},\omega), \tag{5}$$

$$x_{t+1} = f(x_t,U_t,Q_t), \tag{6}$$

with initial conditions

$$\eta_{T+1}^j(x_{T+1},\omega) = 0, \quad j \in N,$$

$$x_1 = x.$$

Note $\eta_t^j(x_t,\omega)$ gives the sales of good j on the sample path ω from customer t onward (the sales-to-go), and decomposes this total sales-to-go as the sum of sales of good j resulting from customer t and the sales-to-go of good j for the remaining customers $t+1, t+2, \dots, T$. The total sales of good j are simply the total sales-to-go for customers $1,\dots,T$, so $\eta^j(x,\omega) = \eta_1^j(x_1,\omega)$. One can therefore use this recursion to investigate properties of the sales functions $\eta^j(x,\omega)$.

With the inventory game and profit functions defined, we can begin to analyze the resulting equilibrium. We begin with some preliminary definitions and lemmas.

3. PRELIMINARY DEFINITIONS AND LEMMAS

Recall a function $h\colon S \to \Re^m$, $S \in \Re^n$ is said to be *Lipschitz with modulus* K_h if $\|h(y_2) - h(y_1)\| \leqslant K_h\|y_2 - y_1\|$ for all $y_1, y_2 \in S$. Let $F_t(\cdot)$ denote the marginal distribution of Q_t, $t = 1,\dots,T$, and let B be defined such that $F_t(B) = 1$. Lemma 1 is proved in Mahajan and van Ryzin (1998).

LEMMA 1. *For each firm j, the function $\eta^j(x,\omega)$ is Lipschitz with modulus $K_{\eta^j} = C_1 + 1$, where $C_1 = 2^{C(n+1)}$ and C is such that $P(T < C) = 1$.*

Lemma 2 justifies the interchange of the expectation and differentiation operations on a sample path ω (see Mahajan and van Ryzin (1998) for proof) based on results in Glasserman 1994:

LEMMA 2. *If the marginal c.d.f.s of the purchase quantities* Q_t, $P(Q_t \leqslant q)$, *are continuous and* $P(T \leqslant C) = 1$ *for some finite* C, *then for all* x: (i) *the gradient* $\nabla\eta(x, \omega)$ *exists* (*w.p.*1), (ii) *the gradient* $\nabla E[\eta(x, \omega)]$ *exists, and* (iii)

$$\nabla E[\eta(x, \omega)] = E[\nabla\eta(x, \omega)].$$

Next, we characterize the derivative of the sample path sales functions for each firm with respect to their own inventory level and the inventory levels of the other firms (see Mahajan and van Ryzin (1998) for proof):

LEMMA 3. *The partial derivatives for the sample path sales functions for each firm j satisfy*

$$\frac{\partial}{\partial x_t^j}\eta_t^j(x_t, \omega) \in \{0, 1\},$$

and

$$\frac{\partial}{\partial x_t^i}\eta_t^j(x_t, \omega) \in \{-1, 0\}, \quad i \neq j$$

for all $t = 1, \ldots, T$.

The second relation above implies that the sample path sales of good j satisfies the *decreasing difference* property because the cross partial derivative is nonpositive. (See Sunduram 1996.) This means the marginal value of inventory for firm j is decreasing in the inventory of any other firm $i \neq j$. (See Mahajan and van Ryzin 1998 for further analysis and discussion of this property.)

Our analysis also requires the following sample path random variable:

$$M^j(x^{-j}, \omega) = \sup\{\eta^j(x^j, x^{-j}, \omega) \colon x^j \geqslant 0\}.$$

In words, $M^j(x^{-j}, \omega)$ represents the largest quantity that firm j can sell on the sample path, ω, given that the inventory levels of the other goods is fixed at x^{-j}. Incrementing the inventory of good j beyond $M^j(x^{-j}, \omega)$ will only result in unsold inventory of j on the sample path ω.

Finally, we use the idea of exchangeable random variables to make precise the notion of symmetric goods. Let $k_1, \ldots, k_n$ denote a permutation of the indices $1, \ldots, n$. The random variables $X_1, \ldots, X_n$ are called *exchangeable* if each permutation, $(X_{k_1}, \ldots, X_{k_n})$, has the same n-dimensional probability distribution. We then have the following definition of symmetric goods.

DEFINITION 1. Goods $1 \leqslant j \leqslant n$, are called symmetric if the n random T-vectors $\{U_t^j \colon t = 1, \ldots, T\}$ are exchangeable.

If the random utility vectors are exchangeable, in a probabilistic sense it does not matter which good is associated with which utility value, and thus goods are symmetric. Utilities that are independent and identically distributed are exchangeable, though the above definition is more general.

4. EQUILIBRIUM ANALYSIS

With this background in place, we can establish our main results.

4.1. Existence of Equilibrium

We start by analyzing the existence of an equilibrium. Lippman and McCardle (1997) prove existence under the assumption that the demand for the good offered by any firm is stochastically decreasing in the inventory levels of the goods offered by the other firms. Our customer choice process results in the demand for each good being decreasing in the inventory levels of the other goods on a sample path basis. Therefore, it satisfies the stochastically decreasing property of Lippman and McCardle. In view of this fact, Theorem 1 is a special case of Lippman and McCardle's result. However, for convenience we include a self-contained proof for our case.

THEOREM 1. *There exists a pure strategy Nash equilibrium to the n firm inventory game.*

PROOF. Because at most C customers can arrive on the sample path w.p.1 and each customer demands at most B units of each good, the strategy space for the inventory level of firm j is given by $[0, BC]$. So the strategy space is a compact, convex subset of $\Re$.

We first show that $\pi^j(x^j, x^{-j})$ is continuous in (x^j, x^{-j}). From Equation (1), we need only to show that $E[\eta^j(x^j, x^{-j}, \omega)]$ is continuous in (x^j, x^{-j}). We have

$$\begin{aligned}
&\|E[\eta^j(x^j, x^{-j}, \omega) - \eta^j(y^j, y^{-j}, \omega)]\| \\
&\quad \leqslant E[\|\eta^j(x^j, x^{-j}, \omega) - \eta^j(y^j, y^{-j}, \omega)\|] \\
&\quad \leqslant (C_1 + 1)\|x - y\|,
\end{aligned} \tag{7}$$

where the last inequality follows from Lemma 1. From Equation (7), we see that $\pi^j(x^j, x^{-j})$ is continuous in (x^j, x^{-j}).

We next show that $\pi^j(x^j, x^{-j})$ is concave in x^j. This follows because $\eta^j(x^j, x^{-j}, \omega) = \min\{x^j, M^j(x^{-j}, \omega)\}$ and the minimum on the right is concave in x^j for fixed x^{-j} and ω. Therefore, $E[\eta^j(x^j, x^{-j}, \omega)]$ is concave in x^j as well. So $\pi^j(x^j, x^{-j})$ is concave in x^j. The theorem then follows from Fudenberg and Tirole (1991, Theorem 1.2). □

4.2. A Comparison of Equilibrium and Joint Optimal Stocking Levels

The next result provides comparitive statics on both the equilibrium and joint optimal stocking levels. First, let

$$A_x^{ij} = \left\{\omega \colon \frac{\partial\eta^j(x, \omega)}{\partial x^i} = -1\right\}.$$

The theorem requires the condition $P(A_x^{ij}) > 0$ to assert strict inequality in the results. The condition is quite general. Essentially, it states that for every inventory level vector x, there should be a set of sample paths with positive measure on which adding a unit of good i, decrements the sales of good j. This happens on sample paths if good i runs out of stock and there are customers arriving later who prefer good i to good j. Without this condition, Theorem 2 remains true with weak inequalities replacing the strict inequalities.

Theorem 2. *If $P(A_x^{ij}) > 0, \forall\, i, j, x \in (0, BC)^n$ then at any Nash equilibrium point $\bar{x}$,*

$$\frac{\partial \pi(\bar{x})}{\partial x^j} < 0 \quad \forall\, j \in N, \tag{8}$$

while at any joint optimum point x^*,

$$\frac{\partial \pi^j(x^*)}{\partial x^j} > 0 \quad \forall\, j \in N. \tag{9}$$

Proof. Because

$$\pi^j(x) = p^j E[\eta^j(x, \omega)] - c^j x^j,$$

using Lemma 2, which justifies the interchange of the derivative and the expectation operations, Lemma 3 and the condition $P(A_x^{ij}) > 0$, we see that

$$\frac{\partial \pi^j(x)}{\partial x^i} = p^j E\left[\frac{\partial \eta^j(x, \omega)}{\partial x^i}\right] < 0 \quad \forall\, i \neq j. \tag{10}$$

Any Nash equilibrium, $\bar{x}$, must satisfy the first-order conditions,

$$\frac{\partial \pi^j(\bar{x})}{\partial x^j} = 0, \quad \forall\, j \in N. \tag{11}$$

Using the fact that

$$\frac{\partial \pi(\bar{x})}{\partial x^i} = \sum_{j=1}^{n} \frac{\partial \pi^j(\bar{x})}{\partial x^i} \tag{12}$$

and Equations (10) and (11), we see that Equation (8) holds. Finally to show Equation (9), note that because any joint optimum point, x^*, satisfies the first-order condition

$$\frac{\partial \pi(x^*)}{\partial x^j} = 0, \quad \forall\, j. \tag{13}$$

This implies by Equation (10) that

$$\frac{\partial \pi^j(x^*)}{\partial x^j} = -\sum_{i=1,\, i\neq j}^{n} \frac{\partial \pi^i(x^*)}{\partial x^j} > 0. \qquad \square$$

To understand Theorem 2 intuitively, consider the symmetric case where all firms have the same selling prices, procurement costs and likelihood of being chosen by each customer. In this case, the combined system profit is split equally among the firms. (See §4.3 below.) In this case, at a Nash equilibrium point, $\bar{x}$, Equation (8) shows that each firm could improve its profit if all firms agreed to reduce their inventory levels. However, no firm has an incentive to do so unilaterally, because at the Nash equilibrium point every firm is making the best response to the other firm's inventory decisions. At a joint optimum point, x^*, each firm is making the highest profit because the combined profit is being maximized and then split equally. But from Equation (9), we see that each firm has an individual incentive to increase their inventory level at any such x^*. Thus, while x^* provides each firm with the highest profit, the less profitable Nash equilibrium $\bar{x}$ is the one that is actually achieved.

Finally, we note that Theorem 2 shows that competition leads to overstocking more generally, in the sense that for any equilibrium $\bar{x}$ there exists a vector $y < \bar{x}$ with higher total profits. This follows from Equation (8) because joint profits can be improved locally at any equilibrium point by slightly reducing any (or all) of the firms' stock levels (i.e., the vector $(-1, \ldots, -1)$ is an ascent direction). This is, of necessity, a local characterization, because in general there may be multiple equilibria and multiple joint optima. (For example, we cannot claim that any equilibrium is, component-wise, no smaller than every joint optimum.) Nevertheless, it provides broad theoretical support for the intuitive notion that a "competitive overstocking" effect is produced by inventory competition.

4.3. Uniqueness of Equilibrium

While Theorem 1 guarantees existence of an equilibrium for the inventory game, uniqueness is not guaranteed in general. We next show that when the goods are symmetric as in Definition **??** and all goods have identical selling prices and procurement costs (henceforth called the *symmetric* game), the equilibrium is indeed unique.

Recall that $M^j(x^{-j}, \omega) = \sup\{\eta^j(x^j, x^{-j}, \omega) \colon x^j \geqslant 0\}$ denotes the largest quantity that firm j can sell on the sample path, ω, given x^{-j}. Note that in terms of $M^j(x^{-j}, \omega)$, the first-order conditions for firm j are

$$P\big(M^j(x^{-j}, \omega) > x^j\big) = \frac{c}{p}.$$

We shall make use of this fact shortly.

We next define two conditions that are required for the uniqueness result. These conditions specify how changing the inventory levels of goods affects the maximum sales of other goods. In particular, decreasing the inventory level of any one good potentially diverts demand to other goods or to the no-purchase option. We would like this diversion to occur in reasonable ways.

Incomplete Demand Diversion. For all $i \neq j$ and for all $0 < \alpha < x^i$

$$P(M^j(x^{-j} - \alpha e^i, \omega) < M^j(x^{-j}, \omega) + \alpha) > 0,$$

where e^i denotes the ith unit vector.

This condition is understood as follows. For any sample path we have that,

$$M^j(x^{-j} - e^i \alpha, \omega) \leqslant M^j(x^{-j}, \omega) + \alpha.$$

A decrease of good i by α units increases the maximum sales of good j, $M^j(\cdot,\omega)$, because customers who would previously have bought good i may now switch to good j. The increase in $M^j(\cdot,\omega)$, is by at most α units, but may be less if some customers choose the no-purchase option (or goods other than j) rather than good j. Incomplete demand diversion simply says that there is a positive probability that there is not a one-for-one exchange of sales between the two goods; that is, there is a positive probability that the increase in $M^j(\cdot,\omega)$ is less than α units.

NONTRIVIAL DEMAND DIVERSION. If $x, y \in \Re^n$, $y < x$, then

$$P(M^j(y^{-j},\omega) > M^j(x^{-j},\omega)) > 0.$$

In contrast to incomplete demand diversion, this condition requires that some diversion of demand takes place when inventory is reduced. That is, let x, y be inventory level vectors with $y < x$. If we compare the maximum sales on a sample path of good j, under these two starting inventory vectors x and y, one would expect that the maximum sales of good j would be higher for the lower inventory level vector y, because fewer competing goods are available. The nontrivial demand diversion condition simply says that there should be some probability of a strict increase in maximum sales if the inventory of competing goods is reduced.

Finally, we will need $M^j(x^{-j},\omega)$ to be absolutely continuous. Recall, a random variable X is *absolutely continuous* iff there is a nonnegative function $f = f_X$ defined on $\Re$ such that

$$F_X(x) = \int_{-\infty}^{x} f_X(t)\,dt,$$

where f_X is the usual density function of X. Absolute continuity rules out the existence of atoms in the cumulative distribution function, i.e., points where the distribution function makes a discontinuous upward jump.

Lemma 4 shows that under these two conditions, if firm i reduces its inventory by α, then firm j requires an increase in its inventory of strictly less than α to restore its inventory to the critical fractile (e.g., restore it to optimality).

LEMMA 4. *Suppose the incomplete demand diversion condition is satisfied, the distribution function of the random variable $M^j(\cdot,\omega)$ is absolutely continuous and $x \in (0, BC)^n$ satisfies*

$$P(M^j(x^{-j},\omega) > x^j) = r.$$

Then for any $i \neq j$ and any scalar α, $0 < x^i - \alpha < BC$,

$$P(M^j(x^{-j} - \alpha e^i,\omega) > x^j + \alpha) < r.$$

PROOF. Recall that for all sample paths $M^j(x^{-j} - \alpha e^i,\omega) \leqslant M^j(x^{-j},\omega) + \alpha$. Thus,

$$M^j(x^{-j} - \alpha e^i,\omega) > x^j + \alpha$$

implies the event

$$M^j(x^{-j},\omega) > x^j.$$

Incomplete demand diversion says that, with positive probability, the converse is not true. That is, $M^j(x^{-j},\omega) > x^j$ does not (w.p.1) imply $M^j(x^{-j} - \alpha e^i,\omega) > x^j + \alpha$. Therefore,

$$\begin{aligned} P(M^j(x^{-j} - \alpha e^i,\omega) > x^j + \alpha) &< P(M^j(x^{-j},\omega) > x^j) \\ &= r. \end{aligned}$$ □

Lemma 5 shows a monotonicity relationship between inventory levels and the probabilities of events described in Lemma 4.

LEMMA 5. *Let $x, y \in \Re^n$ with $y < x$. If nontrivial demand diversion is satisfied, then*

$$P(M^j(y^{-j},\omega) > y^j) > P(M^j(x^{-j},\omega) > x^j).$$

PROOF. From the definition of nontrivial demand diversion and the condition $y^j < x^j$, the event $M^j(x^{-j},\omega) > x^j$ implies the event $M^j(y^{-j},\omega) > y^j$. The lemma then follows. □

We use these two lemmas to establish the uniqueness result.

THEOREM 3. *Under the assumptions of Lemma 4 and Lemma 5, there exists a unique equilibrium to the symmetric n-firm inventory game and the equilibrium is symmetric.*

PROOF. Suppose $\bar{x}$ is a symmetric equilibrium point with $\bar{x}^j = x, \forall j = 1,\ldots,n$. Then $\bar{x}$ satisfies the first-order conditions

$$P(M^j(\bar{x}^{-j},\omega) > x) = \frac{c}{p} \quad \forall j \in N. \tag{14}$$

Uniqueness is established by showing that there are no symmetric equilibria other than $\bar{x}$ and that there are no asymmetric equilibria.

Suppose there was a symmetric equilibria x' other than $\bar{x}$ such that $x' < \bar{x}$ or $x' > \bar{x}$. From Lemma 5, we see that x' would not satisfy the equilibrium condition (14). So any symmetric equilibria $x' < \bar{x}$ or $x' > \bar{x}$ are ruled out.

We next rule out asymmetric equilibria by contradiction. Suppose x is an asymmetric equilibrium. Because x is asymmetric, there must exist i and j such that $x_i > x_j$. Let $\alpha = x_i - x_j > 0$. Because goods are symmetric, using Definition **??** it follows that any permutation of x must also be an asymmetric equilibrium. In particular, the vector $(x_1,\ldots,x_i - \alpha,\ldots,x_j + \alpha,\ldots,x_n)$ must be an equilibrium. But by Lemma 4,

$$\begin{aligned} &P(M^j(x_1,\ldots,x_i - \alpha,\ldots,x_{j-1},x_{j+1},\ldots,x_n,\omega) \\ &\qquad > x_j + \alpha) < \frac{c}{p}, \end{aligned}$$

which violates the first-order condition for good j and contradicts the fact that $(x_1,\ldots,x_i - \alpha,\ldots,x_j + \alpha,\ldots,x_n)$ is an equilibrium. □

5. INVENTORY COMPETITION WITH A LARGE NUMBER OF FIRMS

When the number of firms in the system becomes large, one would intuitively expect the increased competition to lead to lower profits. Lippman and McCardle (1997) showed this zero-profit result in the case of their "herd behavior" allocation rule. In this section, we prove the same result holds for our model under some additional restrictions. More precisely, we show that in the symmetric case when the number of firms gets arbitrarily large, competitive overstocking becomes so severe that the equilibrium inventory rises to the point where no firm earns a profit. To establish this result, we must make the somewhat restrictive assumption that for all customers purchasing any good is always preferred to the no-purchase option. As a result, the system loses a sale only if all firms are out of stock. This assumption is made primarily for analytical tractability. However, when the number of goods is large this assumption is perhaps a reasonable approximation (though a proof allowing for the general no-purchase option would be desirable).

Let x denote the total inventory in the system in equilibrium, which is equally distributed among all n firms in the symmetric case. Let

$$\Delta = \frac{x}{n}$$

be the equilibrium inventory held by each firm. Let $Q = \sum_{t=1}^{T} Q_t$ be the aggregate quantity demanded on the sample path ω. Because by assumption the no-purchase option is excluded (each consumer prefers purchasing anything to not purchasing), the total sales of goods on the sample path equals $\min\{Q, x\}$. The total expected profit in the system is then given by

$$pE[\min\{Q, x\}] - cx.$$

Let x^0 be the inventory level which results in zero profits, so that

$$pE\big[\min\{Q, x^0\}\big] = cx^0.$$

We will show Theorem 4.

THEOREM 4. *Consider the symmetric game in which the no-purchase option has the lowest utility for all customers* (*w.p.*1). *Let x be the equilibrium total system inventory. Then as the number of firms becomes large, $x \to x^0$. That is, the total inventory converges to the level that achieves zero system* (*and individual firm*) *profits.*

PROOF. To prove the result, we condition on the value of T and the quantities $\{Q_t, t = 1, \ldots, T\}$, leaving the utility values $\{U_t, t = 1, \ldots, T\}$ uncertain. Effectively, each customer randomly permutes the indices of the firms and consumes product by proceeding down the permuted list until his/her quantity Q_t is satisfied or no product is left in the system.

Let A denote the set of firms that sell all their stock. Because customers prefer purchasing something to not purchasing, if $Q \geqslant x$ all firms will sell out and $A = N$. If $Q < x$, the situation is as depicted in Figure 1. Some firm's inventory is completely sold out (those in set A), others' are partially sold (denoted set B), and the remaining firms have no sales at all (denoted set C). A total quantity Q is sold by firms, either completely or partially. A key observation from Figure 1 is that there are at most T firms in the set B. This follows because each arriving customer t creates at most one new firm with some partially unsold inventory. Thus,

Figure 1. The sets A, B and C.

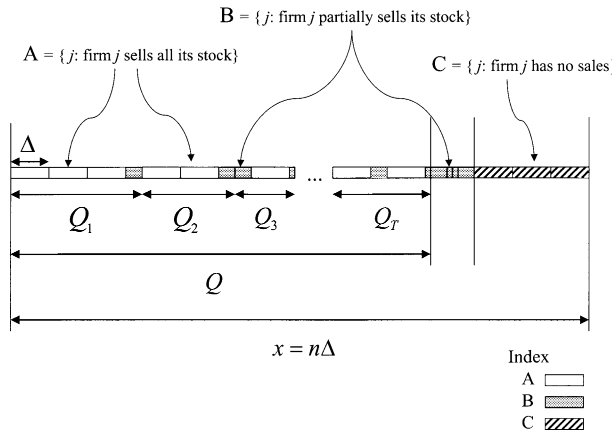

$$P(i \in B) \leqslant (T/n). \tag{15}$$

Moreover, because there are at most $\min\{Q/\Delta, n\}$ firms in the set A

$$P(i \in A) \leqslant \frac{\min\{Q/\Delta, n\}}{n} = \frac{\min\{Q, x\}}{x}. \tag{16}$$

Similarly, because there are at most $\max\{(x-Q)/\Delta, 0\}$ firms in the set C and at most T firms in the set B,

$$\begin{aligned} P(i \in A) &\geqslant \frac{n - \max\{(x-Q)/\Delta, 0\} - T}{n} \\ &= \frac{\min\{Q, x\}}{x} - \frac{T}{n}. \end{aligned} \tag{17}$$

Next, consider perturbing the inventory level of some firm i by a fraction α, $-1 < \alpha < 1$, from Δ to $\Delta(1+\alpha)$. Let $\pi_i(\alpha | Q, T)$ denote the profit to firm i under this perturbation given Q and T. If Δ is the equilibrium inventory (and hence x is the system equilibrium inventory), then any such perturbation should not increase firm i's profits. If firm i is in the set C, then firm i sells nothing and incurs a cost of $c\Delta(1+\alpha)$. If firm i is in set A, it earns revenues of $p\Delta(1+\alpha)$ and incurs cost of $c\Delta(1+\alpha)$. If it is in set B, the sales are indeterminate but are of order Δ and again the firm incurs cost of $c\Delta(1+\alpha)$. Thus, using Equations (15), (16), and (17) we have

$$\pi_i(\alpha|Q,T)$$
$$= p\Delta(1+\alpha)P(i\in A) - c\Delta(1+\alpha) + O(\Delta)P(i\in B)$$
$$= p\Delta(1+\alpha)\frac{\min\{Q,x\}}{x} - c\Delta(1+\alpha) + O\left(\frac{\Delta T}{n}\right)$$
$$= p(1+\alpha)\frac{\min\{Q,x\}}{n} - c\frac{x}{n}(1+\alpha) + O\left(\frac{xT}{n^2}\right).$$

Taking expectations with respect to Q and T yields

$$\pi_i(\alpha) = p(1+\alpha)\frac{E\min\{Q,x\}}{n} - c\frac{x}{n}(1+\alpha) + O\left(\frac{xE[T]}{n^2}\right).$$

As $n\to\infty$, the first two terms dominate. Because these terms are linear in α, the coefficient of α must approach zero, or else firm i would prefer the perturbation α. That is, we must have

$$pE\min\{Q,x\}\to cx,$$

as $n\to\infty$, which implies that $x\to x^0$. $\square$

Theorem 4 shows that the overstocking tendency, as described by Theorem 2, becomes quite extreme as the level of competition rises, causing firms to overstock to the point where they eliminate all their profits. Alternatively, the industry approaches the competitive extreme (zero profit) as n increases.

6. A COMPUTATIONAL ALGORITHM AND NUMERICAL RESULTS

To calculate inventory levels for the inventory game, we propose using a variation of the sample path gradient method developed in our earlier work (Mahajan and van Ryzin 1998) for the monopoly (joint optimal) case. We show therein that the joint optimal problem is not necessarily quasiconcave in general, so finding globally optimal points may be difficult. However, under the assumption that the quantities demanded, Q_t, have continuous distributions, interchanging expectation and differentiation on a sample path ω is justified, so sample path gradients can be used to find stationary points of the expected profit function. In the appendix, we describe in detail how the sample path gradient is calculated. Under relatively mild assumptions, a steepest decent algorithm using the sample path gradient converges to a stationary point of the expected profit function (see Mahajan and van Ryzin 1998).

6.1. Sample Path Gradient Algorithm for the Inventory Game

For finding equilibria, we combine this stochastic gradient algorithm with the idea of "simulated play." Specifically, each firm optimizes over its own inventory level using the sample path gradient algorithm described above, with the inventory levels of the other firms fixed. This is repeated in succession for all firms, and we cycle through the firms until the inventory level vectors converge to within a specified tolerance. In this way the sample path gradient method is used to generate a sequence of best responses until equilibrium is reached. In general, global convergence of this sort of "simulated play" approach is guaranteed only in special cases, for example when the best response function is a contraction mapping (see Bertsekas and Tsitsiklis 1996). In our problem, we have not been able to establish such conditions. Therefore we cannot guarantee convergence. Nevertheless, our experience with the method is that it is quite robust in finding equilbria.

In what follows, $\{a_k\}$ is a sequence that satisfies $\sum_{k=1}^{\infty} a_k = \infty$ and $\sum_{k=1}^{\infty} a_k^2 < \infty$ (e.g. the sequence $a_k = 1/k$). Again, the exact sample path gradient formulae are given in the appendix. The algorithm is:

Step 1. Initialize outer loop: $y_1 = y$, $j = 1$
Step 2. Initialize inner loop: $k = 1$.
Step 3. For Firm j at iteration k

- (i) Generate a new sample path ω_k
- (ii) Calculate sample path gradient for Firm j: $\dfrac{\partial\pi(y_k^j, y^{-j}, \omega_k)}{\partial y^j}$ and step size a_k
- (iii) Update the starting inventory level for the next iteration, using the equation

$$y_{k+1}^j = y_k^j + a_k\frac{\partial\pi(y_k^j, y^{-j}, \omega_k)}{\partial y^j}.$$

Step 4. Check convergence of y^j. If not within tolerance, then $k := k+1$ and GOTO Step 3.
Step 5. Check convergence of y. If not within tolerance, then $j := j+1$ (modulo) n and GOTO Step 2.

The outer loop above repeatedly cycles through the n firms. The inner loop solves for the best response of firm j given the current values y^{-j} of the other firms. The procedure terminates when the best response of each firm results in essentially no change in their decision.

6.2. Numerical Experiments

To illustrate the numerical algorithms and confirm the theoretical properties of the game, we applied the above algorithm to some numerical examples.

In the examples, we used the standard multinomial logit model (MNL), which is the most common random utility model in the economics and marketing literature (See Anderson et al. 1992). In the MNL, the utility takes the form

$$U_t^j = u_j + \xi_t^j \tag{18}$$

for $j\in S(x_t)$, where u_j is called the *nominal* (or expected) utility and ξ_t^j is an noise term that allows for unobservable heterogeneity in taste. The noise component, ξ_t^j, is modeled as a Gumbel (1958) (double exponential) random variable with distribution $P(\xi_t^j \leqslant z) = \exp(-e^{-(\frac{z}{\mu}+\gamma)})$ with mean zero and variance $\frac{\mu^2\pi^2}{6}$. (γ is Euler's constant,

$\gamma \approx 0.5772$.) The nominal utility can be further broken down as $u_j = y + a_j - p_j$ and $u_0 = y + a_0$, where y stands for consumer income, a_j is a quality index and p_j is the price for good j. We analyzed three symmetric cases, with $n = 2, 10, 20$ firms and $a_j = 7.06$, $j = 1, \ldots, 10$, and $a_0 = 4.0$.

We also analyzed an asymmetric case with $n = 2$ firms to see what effect the degree of asymmetry has on equilibrium stocking levels. These quality indices are shown in Table 1 with $a_0 = 4.0$ as before. These values resulted in the same probability of no-purchase in each case but in varying probabilities of purchase between the two firms (higher a_j corresponding to higher purchase probability).

In all cases, the error terms ξ_t^j are i.i.d., Gumbel distributed with parameter $\mu = 1.5$, so the variance of ξ_t^j is 1.18. The procurement cost for a unit is set at $c_j = 1$ for all j, and we assume uniform prices as well, with $p_j = p = 2$, for all firms j. Finally, the number of customers in the sequence, T, was a Poisson random variable with mean 30. Each customer on the sample path desires a quantity of goods, Q_t, which was exponentially distributed with a mean of 1.

We calculated inventory levels using both the equilibrium and joint-optimal sample path gradient algorithms. We then simulated the performance of these inventory levels, to determine the profit resulting from each solution. The number of sample paths simulated was determined as discussed in Banks et al. (1996) so that a 95% confidence interval was within ±1% of the simulated profit.

6.3. Numerical Results

For the symmetric case, we see from Table 2 that as the number of firms increases from 2 to 20, the equilibrium profits in the inventory game relative to the joint optimal profits drops from 99.07% to 73.14%. This is a significant change in performance and confirms the result in Theorem 4 that intensifying competition makes all firms worse off.

Table 1. Quality indices for the asymmetric example.

a_1	a_2
2.00	8.07
4.00	8.00
6.51	7.46
7.06	7.06

Theorem 2 and the discussion in §4.2 suggest that the system under competition stocks more inventory relative to the joint optimum. This behavior is confirmed by the numerical results in Table 2. Note that the competitive inventory relative to the joint optimal solution increases from 108.5% with 2 firms to 144.6% with 20 firms.

We next consider the $n = 2$, asymmetric game. Recall, the level of asymmetry is determined by the attribute indices a_j, where $a_1 = a_2$ is the case of symmetric goods and $a_1 \ll a_2$ is a case of highly asymmetric goods. From Table 3, we see that as the goods become more asymmetric, the competitive profits are somewhat closer to the joint optimal profits, though the gap is relatively small in all cases. The degree of competitive overstocking declines as the problem becomes more asymmetric, falling from 108.5% to 99.4%. This is to be expected because in the limiting case where one good has a attribute index a_j that is arbitrarily higher than the other firm, the problem reduces to a single-firm (monopoly) problem.

In summary, these examples confirm the basic theoretical behavior of the game and illustrate that equilibria can be computed using relatively simple modifications of previous algorithms.

7. CONCLUSIONS

Strategic inventory behavior, in which firms stock goods to take advantage of a competitor's stock-outs or to prevent diversion of demand to competitors, provides an alternative approach for explaining the inventory and capacity decisions observed in a variety of industries. Our work extends this idea, first introduced by Parlar (1988) and later analyzed by Karjalainen (1992) and Lippman and McCardle (1997) to the general, dynamic consumer choice model introduced in our earlier analysis of monopoly assortment decisions Mahajan and van Ryzin (1998).

Our main findings are the following. Equilibrium inventory levels exist under mild regularity conditions, and they are unique in the case of n symmetric firms. Qualitatively, equilibrium inventory levels are in general higher than joint optimal levels, in the sense that the system profit is locally decreasing in the inventory level of each firm at any equilibrium. On the other hand, at any joint optimal solution, each firm has an individual incentive to increase its inventory. This establishes, in a quite general setting, the intuitive result that competition leads to excess stocking. Moreover, in the symmetric case where all customers prefer purchasing something to not purchasing, we found that as the number of firms, n, increases, the excess stocking reaches the

Table 2. Effect of number of firms, n, on profit and inventory.

	Profit			Inventory		
n	Joint Optimal	Competitive	% Opt.	Joint Optimal	Competitive	% Opt.
2	18.4	18.3	99.1	23.7	25.7	108.5
10	22.1	18.5	83.9	27.9	37.3	133.4
20	22.7	16.6	73.1	28.5	41.2	144.6

Table 3. Effect of degree of asymmetry on profit and inventory.

	Profit			Inventory		
a_1, a_2	Joint Optimal	Competitive	% Opt.	Joint Optimal	Competitive	% Opt.
2.00, 8.07	18.8	18.9	100.3*	23.9	23.8	99.4**
4.00, 8.00	18.6	18.5	99.3	23.9	24.6	103.1
6.51, 7.46	18.5	18.2	98.5	23.7	25.7	108.1
7.06, 7.06	18.4	18.3	99.0	23.7	25.7	108.5

* 100.3% due to simulation error
** 99.4% due to tolerance error in stochastic gradient algorithm

point where no firm earns a profit; that is, the industry becomes competitive as $n \to \infty$.

Some open questions and extensions warrant further investigation. On a technical side, it would be nice to see if the zero-profit limit is achieved for the case where a no-purchase option is included. We suspect that the result still holds, but the proof seems more complex. It would be interesting to endogenize price and consider combined price and inventory competition. Another avenue to investigate would be cooperative game theory approaches, which could yield potentially different and interesting insights. Finally, we have applied our inventory competition model to a n-manufacturer, 1-retailer problem to study vendor managed inventory (VMI) and other coordination mechanisms in a supply chain with competition among manufacturers (Mahajan and van Ryzin 1999).

APPENDIX: SAMPLE PATH GRADIENT CALCULATION

Consider the system function defined by Equation (4), where recall $b(j)$ denotes the rank assigned to good j by customer t. Note that $\nabla f(x_t, U_t, Q_t)$ is an $n \times n$ matrix given as

$$\nabla f(x_t,U_t,Q_t) = \left[\nabla f^1(x_t,U_t,Q_t), \nabla f^2(x_t,U_t,Q_t), \ldots, \nabla f^n(x_t,U_t,Q_t)\right],$$

where

$$\nabla f^j(x_t, U_t, Q_t) = \begin{pmatrix} \dfrac{\partial}{\partial x_t^1} f^j(x_t, U_t, Q_t) \\ \vdots \\ \dfrac{\partial}{\partial x_t^n} f^j(x_t, U_t, Q_t) \end{pmatrix}, \tag{A1}$$

and if $b(j) \leqslant m$, then

$$\frac{\partial}{\partial x_t^l} f^j(x_t, U_t, Q_t) = \begin{cases} \mathbf{1}_{\{x_t^{[b(j)-1]}+\cdots+x_t^{[1]}<Q_t\}} \\ \quad \times \mathbf{1}_{\{x_t^{[b(j)]}+\cdots+x_t^{[1]}>Q_t\}} & \text{if } b(j) > b(l) \\ \mathbf{1}_{\{x_t^{[b(j)]}+\cdots+x_t^{[1]}>Q_t\}} & \text{if } b(j) = b(l) \\ 0 & \text{if } b(j) < b(l), \end{cases} \tag{A2}$$

while if $b(j) > m$, then

$$\nabla f^j(x_t, U_t, Q_t) = e^j.$$

Next, using the system function we calculate the inventory level vector observed by each customer on the sample path in a *forward pass* simulation of the choice process. After computing these inventory level vectors, $\{x_t\colon t = 1, \ldots, T\}$, we compute the sample path gradient $\nabla\eta(x, \omega)$ in a *backward pass* on this sample path. Specifically, the steps are as follows.

FORWARD PASS

$$x_1 = x$$

$$x_{t+1} = f(x_t, U_t, Q_t) \quad \forall t = 1, \ldots, T.$$

BACKWARD PASS

$$\nabla\eta_T(x_T, \omega) = \mathbf{I} - \nabla f(x_T, U_T, Q_T) \tag{A3}$$

$$\nabla\eta_t(x_t, \omega) = \mathbf{I} - \nabla f(x_t, U_t, Q_t) + \nabla f(x_t, U_t, Q_t) \times \left[\nabla\eta_{t+1}(f(x_t, U_t, Q_t), \omega)\right], \tag{A4}$$

for $t = 1, \ldots, T-1$, where the gradient $\nabla f(x_t, U_t, Q_t)$ is defined above. Then $\nabla\pi(x, \omega) = \nabla\eta_1(x_1, \omega) - c$.

REFERENCES

Anderson, S. P., A. de Palma, J. F. Thisse. 1992. *Discrete Choice Theory of Product Differentiation*. The MIT Press, Cambridge, MA.

Anupindi, R., M. Dada, S. Gupta. 1998. Estimation of consumer demand with stock-out based substitution: an application to vending machine products. *Marketing Sci.* **17** 406–423.

Balachander, S., P. H. Farquhar. 1994. Gaining more by stocking less: a competitive analysis of product availability. *Marketing Sci.* **13** 3–22.

Banks, J., J. S. Carson, B. L. Nelson. 1996. *Discrete Event System Simulation*. Prentice Hall, Englewood Cliffs, NJ.

Bertsekas, D. P., J. N. Tsitsiklis. 1996. *Parallel and Distributed Computation*. Athena Scientific, Belmont, MA.

Dion, P. A., P. C. Dorin, L. M. Hasey, J. Lundin. 1989. Assessing the costs of industrial market stockouts: an empirical investigation. Working Paper, Bryant College, Smithfield, RI.

Food Marketing Institute. 1993. Variety or duplication: a process to know where you stand. Research Department, Food Marketing Institute, Washington, DC.

Fudenberg, D., J. Tirole. 1991. *Game Theory*, The MIT Press, Cambridge, MA.

Glasserman, P. 1994. Perturbation analysis of production networks. In *Stochastic Modeling and Analysis of Manufacturing Systems*. D. Yao (ed.), Springer-Verlag, New York.

Graves, S. C., A. H. G. Rinnooy Kan, P. H. Zipkin. 1993. *Logistics of Production and Inventory*, Vol. 4 in *Handbooks in Operations Research and Management Science*. Elzevier, Amsterdam.

Gumbel, E. J. 1958. *Statistics of Extremes*. Columbia University Press, New York.

Jeuland, A. P. 1979. The interaction effect of preference and availability on brand switching and market share. *Management Sci.* **25** 953–965.

Karjalainen, R. 1992. The newsboy game. Working Paper, Wharton School, University of Pennsylvania.

Lippman, S. A., K. F. McCardle. 1997. The competitive newsboy. *Oper. Res.* **45** 54–65.

Mahajan, S., G. van Ryzin. 2001. Stocking retail assortments under dynamic consumer substitution, *Oper. Res.* **49**, 334–351.

———, ———. 1999. Retail inventories and consumer choice. In *Quantitative Models for Supply Chain Management*. Tayur et al. (eds.), Kluwer, Dordrecht, The Netherlands.

———, ———. 1999. Supply chain coordination under horizontal competition. Submitted to *Management Sci.*

Mason, J. B., J. B. Wilkinson. 1977. Supermarket product availability and the consumer response. In *Consumer Satisfaction, Dissatisfaction and Complaining Behavior*. R. L. Day (ed.), Indiana University, Bloomington, 153–158.

Nash, J. 1951. Non-cooperative games. *Ann. Math.* **54** 286–295.

Noonan, P. S. 1995. When consumers choose: a multi-product, multi-location newsboy model with substitution. Goizueta Business School Working Paper, Emory University, Atlanta, GA.

Parlar, M. 1988. Game theoretic analysis of the substitutable product inventory problem with random demands. *Naval Res. Logist.* **35** 397–409.

Smith, S. A., N. Agrawal. 2000. Management of multi-item retail inventory systems with demand substitution. *Oper. Res.* **48** 50–64.

R. K. Sundaram. 1996. *A First Course in Optimization Theory*. Cambridge University Press.

van Ryzin, G., S. Mahajan. 1999. On the relationship between variety benefits and inventory costs in retail assortments. *Management Sci.* **45** 1496–1509.