STOCKING RETAIL ASSORTMENTS UNDER DYNAMIC CONSUMER SUBSTITUTION

SIDDHARTH MAHAJAN

Fuqua School of Business, Duke University, Durham, North Carolina, 27708, mahajan@mail.duke.edu.

GARRETT VAN RYZIN

Graduate School of Business, Columbia University, New York, New York 10027, gjv1@columbia.edu.

(Received January 1998; revisions received February 1999; accepted January 2000.)

We analyze a single-period, stochastic inventory model (newsboy-like model) in which a sequence of heterogeneous customers dynamically substitute among product variants within a retail assortment when inventory is depleted. The customer choice decisions are based on a natural and classical utility maximization criterion. Faced with such substitution behavior, the retailer must choose initial inventory levels for the assortment to maximize expected profits.

Using a sample path analysis, we analyze structural properties of the expected profit function. We show that, under very general assumptions on the demand process, total sales of each product are concave in their own inventory levels and possess the so-called *decreasing differences* property, meaning that the marginal value of an additional unit of the given product is decreasing in the inventory levels of all other products. For a continuous relaxation of the problem, we then show, via counterexamples, that the expected profit function is in general not even quasiconcave. Thus, global optimization may be difficult. However, we propose and analyze a stochastic gradient algorithm for the problem, and prove that it converges to a stationary point of the expected profit function under mild conditions. Finally, we apply the algorithm to a set of numerical examples and compare the resulting inventory decisions to those of some simpler, naive heuristics. The examples show that substitution effects can have a significant impact on an assortment's gross profits. The examples also illustrate some systematic distortions in inventory decisions if substitution effects are ignored. In particular, under substitution one should stock relatively more of popular variants and relatively less of unpopular variants than a traditional newsboy analysis indicates.

1. INTRODUCTION

It is common knowledge that retail consumers are often willing to make substitutions if their initial choice of product is out of stock. That is, they may prefer to buy a different size, color or brand within a product category rather than go home empty handed. When such behavior is prevalent in a product category, it makes intuitive sense that a retailer's inventory decisions ought to account for the resulting substitution effects. Yet, most inventory models assume that demand processes for different variants are independent.

Past literature has considered manufacturer-controlled substitution, in which a supplier firm may choose to fill demand for one product with inventory of another product to avoid stocking out. For example a "fast" computer chip is used to satisfy the demand for products requiring a "slow" chip and so on. Single period, multiproduct versions of this problem have been studied quite extensively by Veinott, (1965), Pasternack and Drezner (1991), Bitran and Dasu (1992) and Bassok et al. (1997). But substitution in retail settings is fundamentally different. Substitution decisions are not *directed* by the retailer; rather, they are made by a large number of independently-minded (and self-interested) consumers. A retailer can only *indirectly* affect customer's decisions through his/her inventory policy.

Two recent papers, Smith and Agrawal (2000) and van Ryzin and Mahajan (1999), address the issue of consumer substitution using what we shall term a *static* model of substitution. While there are differences in the modeling and analysis in these papers, both assume a customer's choice is affected by the set of variants offered in the assortment, but not by the current inventory levels. In particular, it is assumed that:

1. the initial choice decision is independent of the current inventory levels of the variants, and
2. if a customer selects a variant that is out of stock, the customer does not undertake a second choice, and the sale is lost.

As a result, demand is independent of inventory levels, though it does depend on the initial choice of variants offered in the assortment. This static substitution model simplifies the resulting inventory and variety analysis, but it is a somewhat unsatisfying assumption, especially for categories such as cigarettes, soft drinks, grocery items, etc., in which consumers substitute readily when products are out of stock.

If one considers the effect of dynamic substitution on assortment planning, several interesting questions come to mind: Does the profit function possess any intuitively

Subject classifications: Inventory/production: consumer substitution. Marketing/Choice Models: choice and inventory decisions.
Area of review: Manufacturing Operations.

0030-364X/01/4903-0334 $05.00
1526-5463 electronic ISSN

appealing and/or simplifying structural properties? How can a retailer efficiently optimize profits in an assortment that exhibits significant substitutability? If assortments can be optimized efficiently, how much difference does it make in gross profits? What distortions are introduced in a retailer's inventory decisions by ignoring substitution effects? And what insights can such analysis provide about the relative costs and benefits of assortment variety within a category?

These are some of the questions we address. We consider a single-period inventory model of a merchandise category made up of multiple product variants, where each variant has a different unit selling price and unit procurement cost. A retailer chooses initial inventory levels for each variant before demand is realized. During the sales period, a sequence of heterogeneous consumers then chooses from among the in-stock variants (or they choose not to purchase), based on a utility maximization criterion. We use a sample path analysis, so minimal distributional assumptions are made on the number of customers and the utilities they assign to variants. In the first part of the paper, sample path properties of the profit function are studied and counterexamples to some hypothesized properties are presented. In the second part of the paper, a sample path gradient algorithm is developed and analyzed. Solutions from this optimization algorithm are then compared to several naive heuristics to investigate how assortment profits and inventory decisions are affected by dynamic substitution.

To our knowledge, Anupindi et al. (1997) and Noonan (1995) are the only other works to date that model the dynamic effect of stock-outs on consumer substitution behavior. In Anupindi et al. (1997), the authors propose estimates of demand for a category of two substitutable items in the presence of stock-outs, assuming that customer arrivals are described by a Poisson process, and that inventory is replenished at regular intervals. However, this work is concerned solely with estimation problems and does not consider the effect of dynamic substitution on assortment planning and inventory levels.

Noonan (1995) considers a problem very similar to ours, but the model and analysis are quite different. He assumes customers have a first choice and a second choice and that demand is generated in two stages. In the first stage, primary (first choice) demand is realized and satisfied as much as possible with available inventory. Then, in the second stage, unfilled primary demand is converted to secondary (second choice) demand for products based on *deterministic* proportions. The resulting total demand for each product is then analyzed in terms of multidimensional integrals over the space of initial demand realizations. Two computational approaches are proposed, although both appear to be limited to problems with small numbers of variants.

While the spirit of this work is quite similar to ours, our model has some distinct advantages. First, Noonan's (1995) choice model is somewhat stylized: it does not allow for more than one substitution attempt; the substitution demand is a deterministic fraction of the excess demand; and it does not model sales as a sequence of customer choices. Our utility maximization model avoids these approximations; the choice is made from all available variants, so there is no notion of a limit on the number of substitution attempts; substitution, like primary demand, is random; and the inventory and customer choice is modeled as a sequential process, so effects due to the arrival order of customers are accurately captured. Moreover, utility maximization is a theoretically satisfying model of choice that allows one to easily investigate consumer welfare and price effects. Finally, our model and analysis leads to a general and efficient algorithmic approach to the problem.

The remainder of the paper is organized as follows: In §2, we describe the model and the sample path view of the system. In §3, we prove our main structural result and discuss some counterexamples. In §4, we develop a relaxation of the problem to do computational work and outline a sample path gradient-based algorithm for computing inventory levels. In §5, we present other heuristic policies and compare the performance of the algorithm to these policies. In §6, we use the sample path gradient algorithm as a tool to study several assortment issues. Finally, conclusions are presented in §7.

2. MODEL FORMULATION

2.1. Notation

We begin with some notational conventions. The set of natural numbers (nonnegative integers) is denoted by Z and N denotes the set $\{1, \ldots, n\}$. 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. For example, x_t^j denotes the jth component of a vector x_t in a sequence $\{x_t : t \geqslant 1\}$. For a real vector y, $y^{[j]}$ denotes the jth largest component of y, that is $y^{[1]} \geqslant y^{[2]} \geqslant \cdots \geqslant y^{[n]}$. The notation $|A|$, denotes the cardinality of set A. $\mathbf{I}$ 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*. and w.p.1 is short for *with probability 1*.

2.2. Assumptions and Model Specification

The merchandise category consists of n substitutable variants, with selling prices p^j and a procurement cost of c^j, $j = 1, \ldots, n$. We consider a one-period (newsboy-like) inventory model in which the retailer's only decision is the vector of initial inventory levels $x = (x^1, \ldots, x^n)$ for each of the variants. Without loss of generality, there is no salvage value for any of the variants.

We take a sample path view 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 variants that are in-stock when he/she arrives. Let $x_t = (x_t^1, \dots, x_t^n)$ denote the vector of inventory levels observed by customer t, and note that $x_1 = x$, where x is the initial inventory decision mentioned above.

For any real vector y, let $S(y) = \{j \bigcup \{0\} : y^j > 0\}$ denote the set of variants with positive inventory (the set of *in-stock* variants) 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 $S(x_{t+1}) \subseteq S(x_t)$.

2.2.1. Choice Process. A customer's choice is based on a simple utility maximization mechanism as follows: each customer t assigns a utility U_t^j to purchasing variant $j \in N$ and to not purchasing, U_t^0. Let $U_t = (U_t^0, U_t^1, \dots, U_t^n)$, denote the vector of utilities assigned by customer t. These utilities should be interpreted as a measure of the *net benefit* to the consumer from purchasing each variant (or not purchasing). In particular, U_t^j is the utility net of the price paid for variant j and any other acquisition costs. As a canonical example, one can assume the no-purchase utility is zero and that consumers purchase only if the net utility of some variant is positive. In this case, U_j^t represents the consumer surplus. (See §5 for a specific example.)

Based on the inventory level x_t and utility vector U_t, customer t makes the choice, $d(x_t; U_t)$, that maximizes his/her utility:

$$d(x_t, U_t) = \arg\max_{j \in S(x_t)} \{U_t^j\}. \tag{1}$$

The resulting decision could be either to buy a variant or to not purchase at all, depending on whether $d(x_t, U_t) = j$ for some $j \in S(x_t)$, $j > 0$ or whether $d(x_t, U_t) = 0$, respectively.

Let $\omega = \{U_t : t = 1, \dots, T\}$ denote the the sample path. We assume ω is a sample from some probability space $(\Omega, \mathcal{F}, P)$. The only assumptions we make on this space are that the sequence is finite w.p.1, i.e., $P(T < +\infty) = 1$ and that each customer t makes a unique choice w.p.1; that is

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

This latter assumption is satisfied, for example, if the utilities have continuous distributions.

The retailer does not know the particular realization ω but does know the probability measure P, so we think of P as characterizing the retailer's knowledge of future demand. The retailer's objective is to choose initial inventory levels x that maximize total expected profit.

2.2.2. Some Special Cases. We next show how several common demand processes studied in the literature can be modeled as special cases of the above choice process:

Multinomial logit (MNL). The MNL model is the most common random utility model and is used widely in the economics and marketing literature. In the MNL, $U_t^0, U_t^1, \dots, U_t^n$ are mutually independent random variables of the form

$$U_t^j = u^j + \xi_t^j, \qquad j = 0, 1, \dots, n, \tag{2}$$

where u^j is a constant called the *nominal* (or expected) utility and ξ_t^j, $j = 0, 1, \dots, n$ are mutually independent noise terms that account for unobservable heterogeneity in taste. The noise components, ξ_t^j, are modeled as Gumbel (double exponential) random variables with distribution $P(\xi_t^j \leqslant z) = \exp(-e^{-(z/\mu+\gamma)})$ with mean zero and variance $\mu^2\pi/6$ (γ is Euler's constant, $\gamma \approx 0.5772$.) The Gumbel distribution is used primarily because it is closed under maximization (Gumbel [1958]). The nominal utility, in turn, can be further broken down as $u^j = a^j - p^j$, $j = 1, \dots, n$ and $u^0 = a^0$, where a^j is a quality index and p^j is the price for variant j.

Markovian Second Choice. This is the model used by Smith and Agrawal (1996). Let the utilities be ordered so that $U_t^{[1]} > \dots > U_t^{[n+1]}$ (i.e., $U_t^{[j]}$ is the jth largest component of U_t). Let $q^j = P(U_t^{[1]} = U_t^j)$ denote the probability that a customer's first choice is j. The second choice depends of the first choice, namely

$$P(U_t^{[2]} = U_t^k \mid U_t^{[1]} = U_t^j) = p_j^k, \tag{3}$$

for all t, where $\sum_{k=0, k\neq j}^n p_j^k = 1$ and $j \neq 0$. There is no third choice under this model, so $U_t^0 > U_t^{[3]} > \dots > U_t^{[n]}$. A distribution of U_t satisfying these conditions corresponds to the Markovian model of Smith and Agrawal (1996).

Universal Backup. In this example, which is a special case studied by Smith and Agrawal (1996), every consumer has an identical second choice if their first choice is not available. For example, every consumer may be willing to settle for vanilla ice cream if their favorite flavor is out of stock.

Consider a category of n variants, and let Variant 1 be the universal backup or common second choice. This case is modeled by constructing a distribution where, w.p.1, either $U_t^0 > U_t^{[1]}$ (the customer prefers not to purchase), or $U_t^{[1]} \geqslant U_t^1 \geqslant U_t^{[2]}$ and $U_t^0 > U_t^{[3]} > \dots > U_t^{[n]}$, for all t. Then arriving customers may have different first choices; Variant 1 is the common second choice of everyone; and not purchasing is the common third choice.

Lancaster demand. This choice model is typical of the attribute space models used in the product variety work of Lancaster (1990). There are n products located along the interval $[0; 1]$ (the "attribute space"). Let the location of product j be denoted l^j. Consumer t has an "ideal point" L_t (a random variable) in the same attribute space $[0, 1]$. The utility of product j for customer t is then given by

$$U_t^j = a - b\|L_t - l^j\|, \tag{4}$$

where a represents the utility of a product that exactly matches the customer's ideal point and b specifies how fast the utility declines with deviations from the ideal point. One then makes a distributional assumption on the customer ideal points L_t (e.g., uniformly distributed on $[0, 1]$) to complete the specification of the model. Note that unlike in the MNL model, utilities in this model may be highly correlated.

INDEPENDENT DEMAND. Independent demand for variants can be modeled by constructing a distribution of U_t such that for all t, $U_t^{[1]} > U_t^0 > U_t^{[2]} > U_t^{[3]} > \cdots > U_t^{[n+1]}$. In this case, for each arriving customer only one variant is preferred to not purchasing. Since no substitution takes place under these conditions, the demand for each variant is independent of the inventory levels of other variants.

2.2.3. Profit Function. To express the total profit concisely, let $\eta^j(x, \omega)$ denote the number of sales of variant j made on the sample path ω given initial inventory levels x. Let $\eta(x, \omega) = (\eta^1(x, \omega), \ldots, \eta^n(x, \omega))$. Then the sample path profit, denoted $\pi(x, \omega)$ is given by

$$\pi(x, \omega) = p^T \eta(x, \omega) - c^T x. \tag{5}$$

where p and c are the vectors of prices and costs, respectively. We will also examine the individual profit functions for each variant,

$$\pi^j(x, \omega) = p^j \eta^j(x, \omega) - c^j x^j. \tag{6}$$

Note that $\pi(x, \omega) = \sum_{j=1}^n \pi^j(x, \omega)$.

The retailer's objective is to solve

$$\max_{x \geqslant 0} E[\pi(x, \omega)]. \tag{7}$$

Since the only complicated quantities in (5) are the functions $\eta^j(x, \omega)$, we focus on understanding their properties. To do so, we introduce a recursive formulation of the problem which is also used later in §4 to construct a computational algorithm.

2.3. Recursive Formulation

We first define the *system function*,

$$f(x_t, U_t) = x_t - e^{d(x_t, U_t)}, \tag{8}$$

where $d(x_t, U_t)$ is, as defined in §2.2, the decision made by customer t and e^j, $j = 1, \ldots, n$ are the n unit vectors and e^0 denotes the zero vector. Then $x_{t+1} = f(x_t, U_t)$, so that $f(\cdot)$ describes how the inventory evolves over time.

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

$$\eta_t^j(x_t, \omega) = x_t^j - f^j(x_t, U_t) + \eta_{t+1}^j(x_{t+1}, \omega), \tag{9}$$
$$x_{t+1} = f(x_t, U_t),$$

with initial conditions

$$\eta_{T+1}^j(x_{T+1}, \omega) = 0, \qquad j \in N,$$
$$x_1 = x.$$

Note $\eta_t^j(x_t, \omega)$ gives the number of sales of variant j on sample path ω from time t onward (the "sales-to-go"), and (9) decomposes this total sales-to-go as the sum of the sales of variant j resulting from customer t and the total sales-to-go of variant j for the remaining customers $t+1, t+2, \ldots T$. Of course, the total sales of variant j are simply the total sales-to-go for customers $1, \ldots, T$, so $\eta^j(x, \omega) = \eta_1^j(x_1, \omega)$. We can therefore use this recursion to investigate properties of the sales functions $\eta^j(x, \omega)$.

3. STRUCTURAL PROPERTIES

We next analyze several structural properties of the profit function (5). In particular, we show that the individual profit functions for each variant are concave in their own inventory and that the marginal value of a unit of inventory of any variant is decreasing in the level of inventory stocked of other variants. These facts imply that the optimal stock level for a given variant is decreasing in the stock levels of the other variants, which is intuitive since variants are substitutes. We then show via counterexamples, that the total profit function is not even component-wise concave in general; for a continuous version of the model, it is not quasiconcave. These facts suggest that global optimization may be difficult in general.

First, we require a preliminary result. One would expect that monotonicity of the inventory levels is preserved by the recursion (9). This is indeed true as shown in the following lemma.

LEMMA 1. *Let x and y be two starting inventory level vectors, such that $x \geqslant y$, and let x_t and y_t denote, respectively, the inventory at time t as determined by the recursion* (9) *given starting inventory levels x and y. Then for all sample paths ω,*

$$x_t \geqslant y_t, \qquad t = 1, \ldots, T.$$

PROOF. The proof is by induction. Fix a sample path ω. We are given that $x_1 \geqslant y_1$, so assume that $x_t \geqslant y_t$. We then need to show that $x_{t+1} \geqslant y_{t+1}$. There are two cases:

Case 1. $d(x_t, U_t) = j$ and $d(y_t, U_t) = 0$. Then $y_{t+1} = y_t$. If $j = 0$, then $x_{t+1} = x_t$ as well. Else, $x_{t+1}^l = x_t^l$ for all $l \in N$, $l \neq j$, so $x_{t+1}^l \geqslant y_{t+1}^l$ for all $l \in N$, $l \neq j$. Also, $y_t^j = 0 = y_{t+1}^j$ and $x_t^j > 0$, so $x_{t+1}^j = x_t^j - 1 \geqslant 0$.

Case 2. $d(x_t, U_t) = j$ and $d(y_t, U_t) = k$. If $j = k$, the result is trivial; if $j \neq k$, then it follows that $y_t^j = 0$ and that $x_{t+1}^j \geqslant y_{t+1}^j$, so $x_{t+1} \geqslant y_{t+1}$. $\square$

3.1. Parametric Monotonicity of the Sales Functions

We next turn to our main result, Theorem 1. Before doing so, we present some background material on parametric monotonicity. The following definitions and results can be found in Sundaram (1996) (see also Topkis 1978):

Let S and Θ be subsets of $\Re^n$ and $\Re^l$ respectively.

DEFINITION 1. A function $h: S \times \Theta \to \Re$ satisfies decreasing differences in (z, θ), if for all $z' \geqslant z$ and $\theta' \geqslant \theta$, $h(z'\theta) - h(z, \theta) \geqslant h(z', \theta') - h(z, \theta')$.

Suppose that S denotes a space of feasible actions and Θ denotes a parameter space. Consider the following optimization problem, for fixed $\theta \in \Theta$,

$$\max\{h(z, \theta) : z \in S\},$$

and define the *optimal action correspondence* (point-to-set mapping)

$$D^*(\theta) = \{z^* : h(z^*, \theta) \geqslant h(z, \theta)\ \forall z \in S\}.$$

We want to determine when the optimal action correspondence $D^*(\cdot)$ is monotone in θ. Let

$$z^*(\theta) = \max\{z : z \in D^*(\theta)\}. \tag{10}$$

The following lemma (see Sundaram 1996 and Topkis 1978 for proofs) shows that the property of decreasing differences can be used to establish parametric monotonicity.

LEMMA 2. *Suppose that the optimization problem*

$$\max\{h(z,\theta) : z \in S\}$$

has at least one solution for each $\theta \in \Theta$, and that h satisfies decreasing differences in θ. Then $z^*(\theta)$ is nonincreasing in the parameter θ.

To connect these parametric monotonicity results to our problem, let z be a scalar, $\theta \in \Re^n$ and define

$$h^j(z,\theta) = \eta^j(ze^j + \theta, \omega).$$

Thus, z changes the inventory level of variant j while θ changes the inventory levels of any variant. Our main result is:

THEOREM 1. *a. The function $h^j(z,\theta)$ is component-wise concave in z for all ω.*

b. The function $h^j(z,\theta)$ satisfies the property of decreasing differences in (z,θ) for all sample paths ω.

PROOF. The properties are quite related and follow from a few elementary observations. We first show component-wise concavity.

Component-wise concavity is equivalent to showing that the first differences,

$$\delta(k,\theta) = \eta^j(ke^j+\theta,\omega) - \eta^j((k-1)e^j+\theta,\omega),$$

are monotonically nonincreasing in $k = 1, 2, \ldots$. From the system equation (9), it is clear that $\delta(k,\theta)$ is either zero or one. That is, on the sample path ω we either sell the kth additional unit or we do not sell it. We next show that once $\delta(k^*,\theta) = 0$ for some k^*, then $\delta(k,\theta) = 0$ for all $k \geqslant k^*$ on the sample path ω. Since $\delta(k,\theta)$ is either 0 or 1, this in turn implies that $\delta(k,\theta)$ is monotonically nonincreasing in k. To prove this property, note from (9) that $\eta^j(k^*e^j+\theta,\omega) = \eta^j((k^*-1)e^j+\theta,\omega)$. That is, if $x_1 = k^*e^j+\theta$ then $x^j_{T+1} > 0$ (there is leftover stock of variant j). Hence $x^j_t > 0$ for all $t = 1, 2, \ldots, T$ (by Lemma 1). Therefore, each customer t always has variant j available, i.e., $j \in S(x_t)\ \forall t$. By Lemma 1, increasing k beyond k^* cannot decrease x^j_t, and therefore such increases have no affect on the choice set $S(x_t)$ at any time t, whence it follows that the choices of each customer t are unchanged, and therefore $\delta(k,\theta) = 0$ for all $k \geqslant k^*$.

Also note that because there exists a k^* such that $\delta(k,\theta) = 1$ for $k < k^*$ and $\delta(k,\theta) = 0$ for all $k \geqslant k^*$, a simple induction argument implies

$$\eta^j(ke^j+\theta,\omega) = \begin{cases} \theta^j + k & k < k^*, \\ \theta^j + k^* - 1 & k \geqslant k^*. \end{cases} \tag{11}$$

Note, $\theta^j + k^* - 1$ is the maximum number of sales of variant j given $\theta^i,\ i \neq j$ and ω. Moreover, it is not hard to show, using Lemma 1, that

$$\eta^j(ke^j+\theta,\omega) \quad \text{is nonincreasing in } \theta^i,\ i \neq j. \tag{12}$$

To show the decreasing differences property of Theorem 1, it suffices to show that the point k^* at which $\delta(k,\theta)$ first becomes zero does not increase when θ increases. In light of part 1, this is equivalent to showing that for each k, $\delta(k,\theta)$ does not increase in θ. But this follows from Lemma 1 and (11) by a contradiction argument.

Indeed, suppose for some $\hat\theta \geqslant \theta$ there exists a $\hat k \geqslant k^*$ such that $\delta(\hat k,\hat\theta) > \delta(k^*,\theta)$. Then by (11), we have

$$\eta^j(\hat k e^j + \theta) = \theta^j + k^* - 1,$$

$$\eta^j(\hat k e^j + \hat\theta) = \hat\theta^j + \hat k.$$

Without loss of generality, we can assume $\hat\theta^j = \theta^j$; else, redefine $\hat\theta := \hat\theta - (\hat\theta^j - \theta^j)e^j$ and $\hat k := \hat k + (\hat\theta^j - \theta^j)$. But $\hat k \geqslant k^*$ implies $\eta^j(\hat k e^j + \hat\theta,\omega) > \eta^j(\hat k e^j + \theta,\omega)$, where $\hat\theta^j = \theta^j$ and $\hat\theta^i \geqslant \theta^i, i \neq j$, which contradicts (12). □

These two properties lead immediately to the following corollary.

COROLLARY 1. *a. A critical stocking level (base-stock level) is optimal for maximizing the component-wise profits (6). b. The component-wise optimal inventory level for j is nonincreasing in $x_i,\ i \neq j$.*

Component-wise concavity of $h^j(x,\omega)$ implies that adding more stock of any particular variant produces decreasing marginal benefits. Therefore, if one is optimizing the profits of variant j only, a critical stocking level (base-stock level) is optimal. This result is merely a reflection of the fact that, given the inventories other than j are fixed, the profits earned by j as a function of x^j are simply the usual newsboy profits.

Part b of this corollary, which follows directly from Theorem 1 and Lemma 2, is more interesting. It says the optimal newsboy quantity for variant j decreases (or at least does not increase) if the inventory levels of other variants rise. This result is a quite natural property for the individual profit functions to possess and reflects the fact that variants are indeed substitutes. Unfortunately, as we show next, the total profit function does not satisfy these same properties.

3.2. Counterexamples of Properties of the Total Profit Function

The total profit function, (5), is less well behaved than the individual profit functions, (6), a fact we illustrate through several counterexamples. For these examples, it is sufficient to analyze the total sales function, defined by

$$T(x,\omega) = \sum_{j\in N} \eta^j(x\omega).$$

Additionally, let

$$H^j(z,\theta) = T(ze^j+\theta,\omega).$$

(This function plays a role analogous to $h(\cdot)$ in the previous section.) Note that if variants have identical prices and costs, then the total profit function is $\pi(x, \omega) = pT(x, \omega) - c\sum_{j=1}^{n} x^j$. Thus, if $T(x, \omega)$ is not component-wise concave or does not satisfy decreasing differences, then the profit function does not satisfy these properties in general. This is the content of the next theorem.

THEOREM 2. *There exist starting inventory level vectors x and sample paths ω for which* 1. $T(x, \omega)$ *is not component-wise concave in x, and* 2. $H^j(z, \theta)$ *does not satisfy the property of decreasing differences in (z, θ) for some $j \in N$.*

PROOF. 1. We consider a sample path with 5 customers and 2 variants, i.e., $T = 5$, and $n = 2$. The utility vectors, U_t, $t = 1, \ldots, 5$ assigned by each customer on the sample path to each variant and to the no-purchase option are:

	U_1	U_2	U_3	U_4	U_5
Variant 1	3	1	1	1	1
Variant 2	4	3	3	3	3
No-purchase	2	2	2	2	2

The total sales on the sample path for four different starting inventory level vectors x are:

x	(1,0)	(1,1)	(1,2)	(1,3)
$T(x, \omega)$	1	1	2	3

We see that the first increment of Variant 2 on the sample path does not increase total sales. This is because the first increment of Variant 2 switches customer 1 from Variant 1 to Variant 2, thus missing the only opportunity to sell Variant 1 on the sample path. However the second and successive increments of Variant 2, increase total sales by 1. This violates component-wise concavity of the total sales sample path function.

2. We consider a sample path with 4 customers and 3 variants, i.e., $T = 4$ and $n = 3$. The utility vectors, U_t, $t = 1, \ldots, 4$ assigned by each customer on the sample path to each variant and to the no-purchase option are:

	U_1	U_2	U_3	U_4
Variant 1	5	3	3	3
Variant 2	3	1	1	1
Variant 3	4	5	5	5
No-purchase	2	2	2	2

The total sales on the sample path for four different starting inventory level vectors x are:

x	(0,1,0)	(0,1,1)	(1,1,0)	(1,1,1)
$T(x, \omega)$	1	1	1	2

We see that when Variant 1 is at an inventory level of 0, an increment of Variant 3 does not result in an increase in total sales on the sample path. However when Variant 1 is at an inventory level of 1, an increment of Variant 3 increases total sales by 1 on the sample path. This violates the decreasing differences property of $H^3(\cdot)$. □

The fact that the total profit function is not even component-wise concave is rather startling and attests to the complexity of the profit function. Similarly, it is surprising that decreasing differences do not hold either. Of course, these are only sample path counterexamples. It is entirely possible that under more restricted distributional assumptions, one or both of these properties may hold for the expected profit function. However, such an analysis would most likely require different techniques.

3.3. Continuous Model and Quasiconcavity Counterexamples

So far the structural results we have shown have been derived assuming the inventory levels are integral. However, a continuous version of the problem is useful to address theoretical questions of quasiconcavity and to develop computational approaches using gradient-based algorithms. To this end, we develop a natural continuous relax-ation of the problem which has the integer problem as a special case. The fluid model may also be of independent interest in applications where demand and inventory are truly continuous (e.g., petrochemicals).

To define the continuous problem, inventory is viewed as a fluid and each customer t requires a continuous quantity of fluid, Q_t. As before, the choices are ordered based on the customer's utility vector U_t. The customer drains the inventory of the most preferred fluid first. If this fluid runs out, the customer drains the inventory of the second most preferred fluid and so on. This process continues until either the customer's entire requirement Q_t is met or the inventory of all fluids valued higher than the no purchase utility is exhausted. Note that if each customer requires exactly one unit of fluid and if the initial fluid levels are integral, then this model is equivalent to the discrete one considered above.

For the continuous problem, the sample path is $\omega = \{(U_1, Q_1), (U_2, Q_2), \ldots, (U_T, Q_T)\}$, and we modify the system function, $f(\cdot)$, as follows: Let the components of the vector U_t, be ordered as before so that $U_t^{[1]} > U_t^{[2]} > \cdots > U_t^{[n+1]}$. Let m denote the number of variants 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 variant 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

$$f^j(x, U_t, Q_t) = \begin{cases} (x_t^{[1]} - Q_t)^+ & b(j) = 1, \\ x_t^{[b(j)]} + \cdots + x_t^{[1]} - Q_t)^+ & \\ \quad -(x_t^{[b(j)-1]} + \cdots + x_t^{[1]} - Q_t)^+ & 1 < b(j) \leqslant m, \\ x_t^j & b(j) > m. \end{cases} \tag{13}$$

That is, if $b(j) > m$, then it implies that the no purchase option is preferred to variant j; i.e., $U_t^j < U_t^0$, so the fluid volume of variant j is not drained. If $b(j) \leqslant m$, then the fluid volume of variant j that is drained is given by the first expression in equation (13).

Define $\eta_t^j(x_t, \omega)$ analogously to be the total sales-to-go of fluid j from customers $t, t+1, \ldots, T$. Then, similar to equation (9), the evolution of the sample path can be described using the following recursive equation,

$$\eta_t^j(x_t, \omega) = x_t^j - f^j(x_t, U_t, Q_t) + \eta_{t+1}^j(x_{t+1}, \omega), \tag{14}$$

$$x_{t+1} = f(x_t, U_t, Q_t), \tag{15}$$

with initial conditions,

$$\eta_{T+1}^j(x_{T+1}, \omega) = 0, \qquad j \in N,$$
$$x_1 = x.$$

The sample path profit function $\pi(x, \omega)$ is defined by (5) as before, with the vector $\eta(x, \omega)$ corresponding to the sample path sales for the continuous model,

$$\pi(x, \omega) = p^T \eta(x, \omega) - c^T x.$$

With the continuous model so defined, we can investigate quasiconcavity properties of the profit function. Our main result here is a negative one. Namely, we demonstrate that, in general, the sample path profit function is not quasiconcave in the starting inventory level vector. The significance of this result is that without quasiconcavity, we cannot preclude the possibility that there may be local optima in the expected profit of the continuous problem (7).

Theorem 3. *There exist sample paths ω on which the sample path profit function for the continuous problem, $\pi(x, \omega)$ is not quasi-concave.*

Proof. We exhibit a sample path, shown in Table 1, on which the profit function is not quasiconcave. Three variants are stocked. Let $p^j = 2$, $j = 1, 2, 3$, and $c^j = 1$, $j = 1, 2, 3$. There are $T = 15$ customers who arrive on the sample path and each customer requires one unit of fluid, so $Q_t = 1$, $t = 1, \ldots, 15$. There are two initial inventory levels given by $y = (10, 0, 5)$, i.e., the retailer stocks 10 units of the first variant, no units of the second, and 5 units of the third variant, and $z = (0, 5, 10)$. Their convex combination given by $\alpha y + (1-\alpha) z$ at $\alpha = 0.8$ is the inventory level (8,1,6). Table 1 shows that the sample path sales function is not quasiconcave; that is, this sample path has strictly fewer sales using a convex combination of two inventory levels, in this case (8,1,6), than with either of the two inventory levels, (10,0,5) and (0,5,10). □

Table 1. Sample path for quasiconcavity counterexample.

Customer	Sequence of events	Inventory vector		
t		y	z	$0.8y + 0.2z$
0		(10,0,5)	(0,5,10)	(8,1,6)
1-3	First 3 arrivals choose Variant 3	(10,0,2)	(0,5,7)	(8,1,3)
4-5	Next 2 prefer Variants 2,1 in that order	(8,0,2)	(0,3,7)	(7,0,3)
6-10	Next 5 prefer Variants 1,3 in that order	(3,0,2)	(0,3,2)	(2,0,3)
11-12	Next 2 choose only Variant 2	(3,0,2)	(0,1,2)	(2,0,3)
13-15	Next 3 prefer Variants 1,2 in that order	(0,0,2)	(0,0,2)	(0,0,3)
	Total profit	11	11	9

4. OPTIMIZING ASSORTMENT INVENTORIES

A sample path gradient method is essentially like a steepest ascent method with a stochastic (noisy) gradient replacing a deterministic one. We describe the idea in more detail shortly. In developing the algorithm, we work with the continuous problem.

We first need a lemma to justify the interchange of the expectation and differentiation operations on a sample path ω. The required technical condition is that the purchase quantities, Q_t, be bounded, continuous random variables.

Lemma 3. *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 the gradient $\nabla E[\eta(x, \omega)]$, exists for all x and*

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

Proof. See the Appendix. □

Requiring continuously distributed purchase quantities is somewhat restrictive. However, we require this assumption only to prove convergence of the resulting algorithm. If this condition is violated, the algorithm gives stochastic subgradients instead, which can certainly be used effectively in practice, as shown in §5.

Note that in view of Theorem 3, it is difficult to find globally convergent algorithms for the continuous problem. We therefore restrict our attention to finding algorithms that converge to stationary points of the expected profit function (locally convergent algorithms). In particular, let the set D be defined as

$$D = \{x \in \Re^n : \nabla E[\pi(x, \omega)] = 0\}.$$

Our aim to find an inventory vector x such that $x \in D$. (We assume here that all stationary points are interior to the constraint set $\{x : x \geqslant 0\}$. One can easily modify the procedure to maintain nonnegativity by projecting onto this set at each iteration.)

Using Lemma 3, we see that finding $x \in D$, is equivalent to finding the roots of the equation,

$$E[p^T \nabla \eta(x, \omega) - c] = 0. \tag{16}$$

Before proceeding with the algorithm, we discuss how the sample path gradient $\nabla \eta(x, \omega)$ is calculated.

4.1. Calculating $\nabla \eta(x, \omega)$

Using the knowledge of the choice decision made by each customer, we calculate the inventory level vector observed by the subsequent customer on the sample path. This is done in a *forward pass*. After computing the inventory level vectors, $\{x_t : t = 1, \dots, T\}$, we compute the sample path gradient $\nabla \eta(x, \omega)$ in a *backward pass* on the sample path. Specifically, the steps are:

Forward Pass

$$x_1 = x$$
$$x_{t+1} = f(x_t, U_t, Q_t) \qquad \forall t = 1, \dots, T.$$

Backward Pass

$$\nabla \eta_T(x_T, \omega) = \mathbf{I} - \nabla f(x_T, U_T, Q_T), \tag{17}$$

$$\begin{aligned} \nabla \eta_t(x_t, \omega) = \mathbf{I} - \nabla f(x_t, U_t, Q_t) + \nabla f(x_t, U_t, Q_t) \\ \times [\nabla \eta_{t+1}(f(x_t, U_t, Q_t), \omega)], \end{aligned} \tag{18}$$

for $t = 1, \dots, T-1$, where the gradient $\nabla f(x_t, U_t, Q_t)$ is defined in the Appendix. Then $\nabla \eta(x, \omega) = \nabla \eta_1(x_1, \omega)$.

4.2. The Sample Path Gradient Algorithm

With the sample path gradient, $\nabla \eta(x, \omega)$ in hand, we can now describe the Sample Path Gradient algorithm which is used to compute inventory levels. We require an initial starting inventory level y and a sequence of step sizes, $\{a_k\}$, with the following properties.

$$\sum_{k=0}^{\infty} a_k = \infty \quad \text{and} \quad \sum_{k=0}^{\infty} a_k^2 < \infty.$$

(For example, $a_k = 1/k$.)

The algorithm proceeds as follows:

1. Initialize: $k = 0$ and $y_0 = y$.
2. At iteration k,
   a. Generate a new sample path ω_k.
   b. Calculate sample path gradient $\nabla \eta(y_k, \omega_k)$ and step size a_k.
   c. Update the starting inventory level for the next iteration, using the equation

$$y_{k+1} = y_k + a_k \nabla \pi(y_k, \omega_k).$$

3. $k := k + 1$ and GOTO Step 2.

Let $F_t(\cdot)$ denote the marginal distribution of Q_t, $t = 1, \dots, T$ and let $\{y_k\}$ denote, as above, the sequence of iterates generated by the Sample Path Gradient algorithm. Under relatively mild assumptions the Sample Path Gradient algorithm converges in the limit to an inventory level vector in the set D of stationary points of the expected profit function. This result is formalized in the next theorem.

THEOREM 4. *Suppose $F_t(\cdot)$ is Lipschitz for all $t = 1, \dots, T$. That is,*

$$\|F_t(x) - F_t(y)\| \leqslant K\|x - y\|,$$

for all $x, y \in \Re$. Then

a.

$$\lim_{k \to \infty} \nabla E[\pi(y_k, \omega_k)] = 0.$$

b. Every limit point of the sequence y_k is a stationary point of $E[\pi(y, \omega)]$.

PROOF. See the Appendix. □

4.3. Some Comments on Global Optimization

The above algorithm is only guaranteed to find stationary points and may not find a global optimum. One potential approach to extending the algorithm to find a global optimum is to imbed it in a so-called *path following* approach, as is done in Hanson and Martin (1996) for optimizing profit functions when the demand system is given by the Multinomial Logit random utility model to be discussed in the next section. (See Garcia and Zangwill (1981) for a general treatment of path following methods.) In utilizing this approach, we first perturb the problem with a parameter θ to make it concave. For example, if we add a sufficiently large θ to each utility U_t^j, $j \neq 0$ so that $U_t^j + \theta > U_t^0$ with probability one, then it is easy to show that the resulting profit function (5) is concave in x. The idea is then to decrease θ to zero, optimizing x at each point, to construct a "path" of inventory levels $x^*(\theta)$ from the global optimum of the perturbed but concave profit function to the global optimum of the original profit function.

However, there are potential pitfalls with this approach: A valid path may not exist; there could be a fork or bifurcation in the path and it is not clear, *a priori*, which "fork in the road" the algorithm will follow; the path may have discontinuous jumps, and so on. (See Hanson and Martin (1996) for a discussion of these potential problems.) Also, from a practical standpoint, path following approaches tend to be computationally intensive. Moreover, as shown below, in our numerical testing it appears that a straight-forward implementation of the Sample Path Gradient algorithm is quite robust in finding global optima.

5. NUMERICAL EXPERIMENTS AND COMPARISONS WITH HEURISTIC POLICIES

There are several reasons to engage in numerical analysis. The first, clearly, is to understand the performance of the proposed optimization algorithm. However, the sample path gradient algorithm described above is perhaps best viewed as a tool to help asses the impact of substitution behavior on both inventory decisions and profits. In this spirit, we compare the Sample Path Gradient algorithm to several other heuristic policies.

5.1. Heuristic Policies

We next propose two policies that are meant to mimic the types of heuristic decision rules used in practice. Our goals in examining these policies are: 1. to understand qualitatively any distortions that might be introduced in inventory decisions if one ignores (or approximates) substitution effects, and 2. to gauge the impact of substitution effects on profits and other performance measures.

In defining these policies, we assume we are given an estimate of the probability of choosing Variant j from a set S, which we denote $q_j(S)$. This probability depends on the choice model and is computed explicitly in the examples below. Further, in all our simulation testing we used the continuous model of §3.3 where the total number of customers, T, was Poisson, and the quantity demanded by each customer, Q_t, was exponentially distributed with mean one. The heuristics below are defined with these facts in mind, but can be easily generalized to other cases.

5.1.1. Independent Newsboy. The Independent Newsboy policy makes the simplifying assumption that demand for each variant is independent of the current on-hand inventory levels. It thus ignores dynamic substitution effects entirely. However, we allow demand to depend on the set S of variants that are stocked through its affect on the probabilities $q_j(S)$. Specifically, if the number of customers T is Poisson with mean λ, the number of customers who rank Variant j first is assumed to be a thinned Poisson process with mean $\lambda q_j(S)$ (which would be exact if there were no stockouts). Each customer demands an exponentially distributed quantity Q_t with mean one. Therefore, the total demand for Variant j under this assumption has mean $\lambda q_j(S)$ and variance $2\lambda q_j(S)$. The exact distribution is somewhat complex, but to simplify the heuristic we use a normal approximation. Indeed, in applying inventory models in practice it is quite common to use simple distributional assumptions of this sort.

Given S, the inventory decision for each variant is then treated as an independent newsboy problem. Using a normal approximation to the distribution of total demand, the optimal Independent Newsboy stock level, denoted x_I^j, is then

$$x_I^j = \lambda q^j(S) + z^j\sqrt{2\lambda q^j(S)}, \qquad j \in S, \tag{19}$$

where λ is the mean number of arriving customers,

$$z^j = \Phi^{-1}\left(1 - \frac{c^j}{p^j}\right) \tag{20}$$

and $\Phi(z)$ denotes the c.d.f. of a standard normal random variable.

The next step in implementing the Independent Newsboy heuristic is to determine the best subset S of variants to stock. In the general case this may require evaluating 2^n subsets (e.g., by simulation), which is impractical. However, for the MNL model a simplification based on our previous work in van Ryzin and Mahajan (1996) reduces the search to n sets, as described below.

5.1.2. Pooled Newsboy. In constrast to the Independent Newsboy—which assumes there is no dynamic substitution—this policy effectively assumes "complete substitution" among the products in the set S. That is, demand is pooled and an aggregate quantity is determined to maximize the resulting profits assuming that customers will freely substitute among all the variants in S. It also mimics a top-down approach in which a buyer orders an aggregate quantity for a category based on total demand and approximate gross margins and then allocates this total quantity in a rough-cut fashion to individual variants.

Specifically, let

$$q(S) = \sum_{j\in S} q^j(S),$$

denote the probability that a customer chooses at least one of the products in S. The total number of customers choosing the category S is then Poisson with mean $\lambda q(S)$. Each customer demands an exponentially distributed quantity Q_t with mean one as before. The total demand for the category S under this approximation therefore has mean $\lambda_q(S)$ and variance $2\lambda q(S)$. Under a normal approximation to this total demand distribution, the optimal aggregate inventory level for the category, denoted $x(S)$, is computed using

$$x(S) = \lambda q(S) + z\sqrt{2\lambda q(S)},$$

where z is the newsboy fractile. This fractile is determined using a weighted average price and cost as follows:

$$z = \Phi^{-1}\left(1 - \frac{\bar{c}}{\bar{p}}\right), \tag{21}$$

where

$$\bar{p} = \frac{\sum_{j\in S} q^j(S)p^j}{q(S)}, \qquad \bar{c} = \frac{\sum_{j\in S} q^j(S)c^j}{q(S)},$$

and $\Phi(z)$ denotes the c.d.f. of a standard normal random variable. The Pooled Newsboy inventory for variant j, denoted x_P^j, is then determined by allocating the aggregate inventory proportional to $q^j(S)$ as follows:

$$x_P^j = x(S)\frac{q^j(S)}{q(S)}, \qquad j \in S.$$

Again, one can then attempt to optimize over the subset S.

5.2. Example 1

In these first two examples, we used the multinomial logit model (MNL) with utilities given by (2). A standard result of the MNL is that the choice probabilities are given by

$$q^j(S) = P(U_t^j = \max\{U_t^i : i \in S\}) = \frac{v^j}{\sum_{i\in S} v^i + v^0}, \tag{22}$$

where

$$v^j = \begin{cases} e^{u_j/\mu} & j \in S, \\ e^{u_0/\mu} & j = 0. \end{cases} \tag{23}$$

As mentioned, for the MNL model we optimized over the set S using a result from our previous work on static substitution in van Ryzin and Mahajan (1996). Assume variants are indexed so that $v_1 \geqslant v_2 \geqslant \cdots \geqslant v_1$ and let $A_k = \{1, 2, \ldots, k\}$ denote the set consisting of the k best variants. A structural result derived (van Ryzin and Mahajan 1996) shows that in the case where all variants have identical costs and prices, the optimal assortment in the static substitution case will always belong to one of the sets A_k, $k \in \{1, \ldots, n\}$. Motivated by this result, for the Independent Newsboy and Pooled Newsboy heuristic we simply tried out each one of these n sets in turn using simulation, and chose the one that yields the highest profit.

In our first example, we consider $n = 10$ possible variants. The nominal utilities are of the form

$$u^j = a^j - p^j, \qquad j = 1, \ldots, n,$$

$u^0 = a^0$ and the quality indices, a^j, are linearly decreasing:

$$a_j = 12.25 - 0.5(j-1) \qquad \forall j = 1, \ldots, 10, \tag{24}$$

and $a_0 = 4.0$. 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 and price are the same for all variants in Example 1, with $c^j = 3$ and $p^j = p$ for all j where p takes values in the range 3 to 9. This simplification facilitates comparison with other heuristic policies and also makes it easier to investigate how different performance measures vary with price. The number of customers in the sequence, T, is a Poisson random variable with mean 30 and the quantity Q_t demanded by each customer t is an exponential random variable with mean one.

We computed the inventory levels using the Sample Path Gradient algorithm and each of the heuristic policies above. We then simulated the performance of these resulting inventory levels. The number of sample paths simulated was determined as discussed in Banks et al (1996), so that total profits were within $\pm 1\%$ with 95% confidence. The number of sample paths ranged from 30,000 for $p = 3$ to about 200,000 for $p = 8$.

Because the Sample Path Gradient algorithm is not globally convergent, we tested its performance on a few examples for large numbers of starting points to gauge its performance. All converged to the same point. In all our testing, we have not yet found an example which exhibited convergence to different points. Moreover, as shown below, the Sample Path Gradient algorithm consistently outperforms all of the heuristic policies. For these reasons, we believe it is quite robust in finding a global maximum; however, global convergence cannot be guaranteed.

5.3. Performance Comparison for Example 1

The profit comparison for the three methods is given in Table 2. The results show that while the Sample Path Gradient algorithm outperforms the heuristic policies, these policies do quite well. The Pooled Newsboy heuristic in particular does very well, achieving 99.6% and 99.4% of the profits generated by the Sample Path Gradient algorithm, respectively, at price levels 5 and 8. The Independent Newsboy heuristic does not do as well, but the profits are still 99.0% and 97.6% of the maximum, respectively, at price levels 5 and 8.

We next qualitatively compare the inventory levels obtained under each heuristic policy relative to those obtained by the Sample Path Gradient algorithm.

Independent Newsboy. We first examine the case where the assortment consists of all 10 variants, that is $S = A_{10}$. While this is not the best subset, it illustrates some interesting properties of the Independent Newsboy policy. Results at price $p = 8.0$ are shown in the Figure 1. The total inventory under the Independent Newsboy policy is 20% higher than the sample path gradient inventory (30.3 vs. 25.2). The distribution of inventory is different as well. The Independent Newsboy inventory levels are lower for the first two variants and systematically higher for the last eight variants. That is, inventory is more evenly spread across variants. The Independent Newsboy profit is also lower at 83.8, a loss of 5.7% as compared to the Sample Path Gradient algorithm. (Note this profit is lower that reported in Table 2, because Table 2 if for the best subset, A_5.)

The intuition for this behavior is the following: substitution has two effects on the standard newsboy logic. First, because of substitution behavior, the demand for a given variant will be higher than the Independent Newsboy predicts, because it receives some additional substitute demand from other variants when they are out of stock. This effect increases the level of demand, which provides an incentive to *increase* inventory. On the other hand, the unit underage cost is lower than the independent newsboy predicts because an underage in one variant does not always result

Table 2. Gross profit of policies for Example 1.

Price	5.0	8.0
Sample Path Gradient	43.5	88.9
Independent Newsboy	43.0	86.8
Pooled Newsboy	43.3	88.4

Figure 1. Inventory of Independent Newsboy Policy for Example 1 ($p = 8$): All 10 Variants.

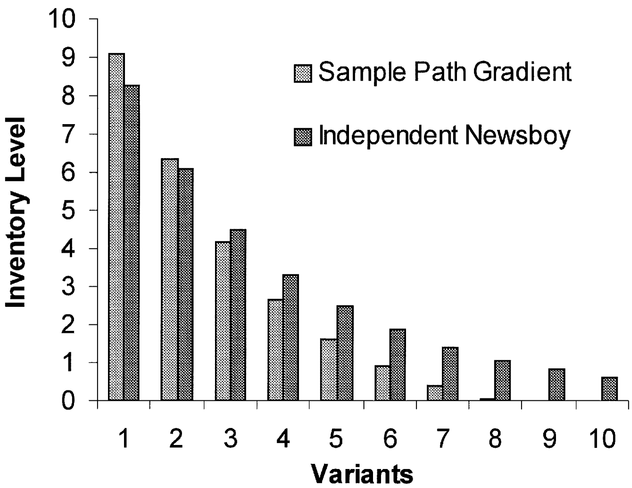

Figure 2. Inventory of Independent Newsboy policy for Example 1 using $A_5 (p = 8)$.

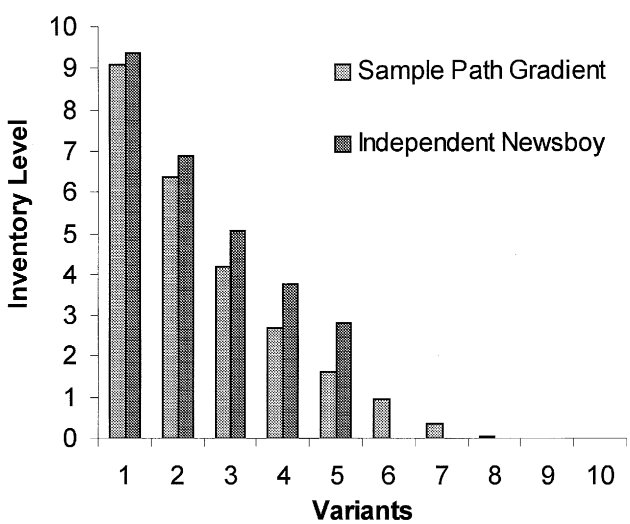

in a lost sale; customers may substitute rather than not purchase. This reduces the effective underage cost, which creates an incentive to *decrease* inventory. Which of these two opposing effects dominates depends on the item, with popular items having lots of additional substitution demand and fewer chances for "back-up" alternatives, while less popular items receive little substitution demand and have many attractive back-up alternatives if they are out of stock. The net result is that the independent newsboy is biased; it stocks too little of the popular items and too much of the less popular items.

If we next maximize over the sets $A_1, \ldots, A_{10}$ in turn, A_5 is the best set at price 8.0. Results are shown in Figure 2. Also the profit performance improves quite significantly over the A_{10} case, rising to 86.8 from 83.8. Choosing A_5 over A_{10} has two effects on the Independent Newsboy policy: 1. the total inventory is reduced from 30.3 to 27.9, which brings it closer to the total Sample Path inventory of

Figure 3. Inventory of Pooled Newsboy policy for Example 1.

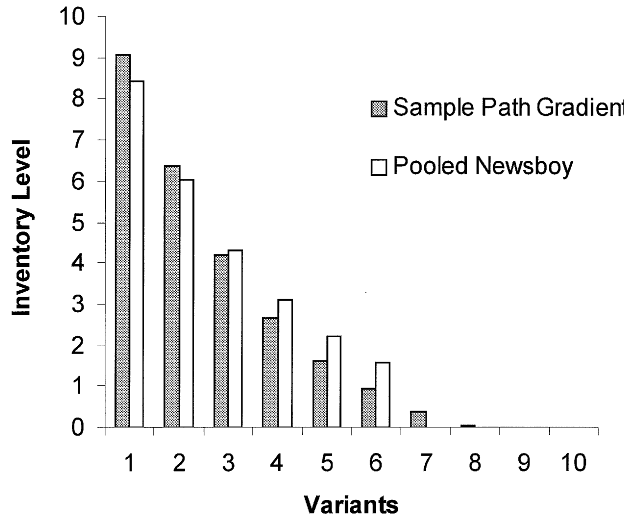

25.2; and 2. choosing A_5 forces the heuristic to carry more inventory of the popular variants and less (or none) of the unpopular variants as the Sample Path Gradient algorithm does. The combination seems to significantly improve the Independent Newsboy's performance.

Pooled Newsboy. Figure 3 shows the allocation under the Pooled Newsboy policy at $p = 8$. The total inventory is 25.7, which is remarkably close to the total inventory of 25.2 for the Sample Path Gradient. The optimal set of variants is A_6. However, inventory is more evenly spread across variants than is optimal. As seen in Table 2, the allocation scheme of the Pooled Newsboy achieves profit levels which are only slightly worse than the Sample Path Gradient algorithm. Indeed it is surprising that a simple allocation rule of this nature does so well. However, the way inventory is allocated does seem to matter. For example, we tested what happens when the same 25.7 units of inventory are allocated evenly across all variants, which is clearly not optimal; the profits decline by almost 12% in this case.

5.4. Example 2

We next consider a numerical example with two variants where we allow for different profit margins. In contrast to the equal-margin case of Example 1, the performance of the Sample Path Gradient algorithm is considerably better than both heuristics in this case and the performance of the pooled newsboy is the worst. Also, the stocking decisions of the various policies exhibit greater variability.

In this example, Variant 1 is less popular on average but has a high margin; Variant 2 is more popular on average but has a low margin. The actual numerical values are as follows: For Variant 1,

$$a^1 = 10, \qquad p^1 = 10, \qquad c^1 = 1, \tag{25}$$

Table 3. Profits of policies for Example 2.

Policy	Profit
Sample Path Gradient	79.5
Independent Newsboy	65.6
Pooled Newsboy	57.4

Figure 4. Inventory levels for Example 2.

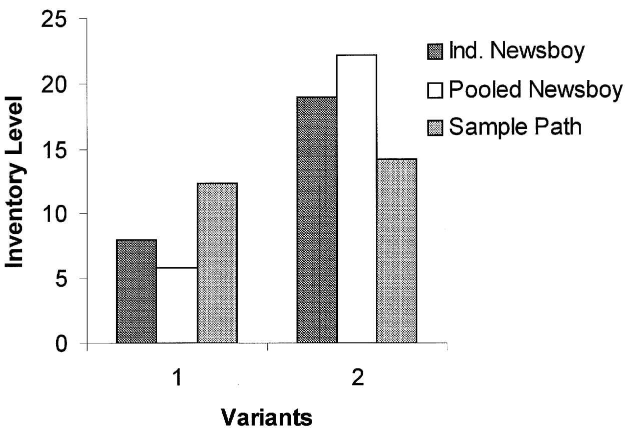

so the mean utility is $u^1 = a^1 - p^1 = 0$ and the margin is 9. For Variant 2,

$$a^2 = 4, \qquad p^2 = 2, \qquad c^2 = 1; \tag{26}$$

so the mean utility is $u^2 = a^2 - p^2 = 2$ and the margin is 1. As a result, while people will tend to prefer Variant 2 because it has a higher mean utility (net of price), Variant 1 is much more profitable.

Comparisons of the Sample Path Gradient method with the Independent and Pooled Newsboy model are shown in the Table 3. The Sample Path Gradient algorithm results in a 12% increase in profits over the Independent Newsboy heuristic and a 19% increase over the Pooled Newsboy heuristic.

The inventory levels for Example 2 are shown in Figure 4. Note that the Sample Path Gradient stocks more of the less-popular, high-margin variant (Variant 1) than the heuristic policies, while the heuristic policies both stock more of the popular, low-margin variant (Variant 2). What appears to be happening is that by stocking less of the low-margin variant, the Sample Path Gradient policy induces customers to "upgrade" to the high-margin variant even though it may not initially be their first choice (i.e., a sort of "bait-and-switch" effect).

5.5. Example 3

This example shows the performance of the three policies under a different choice model. In this example, we use a Lancaster-type (Lancaster 1990) demand model with attribute space [0 1]. There are four products with locations $l^1 = 0.8$, $l^2 = 0.2$, $l^3 = 0.3$ and $l^4 = 0.4$. Thus, Variants 2, 3

Figure 5. Attribute space and associated data for Example 3.

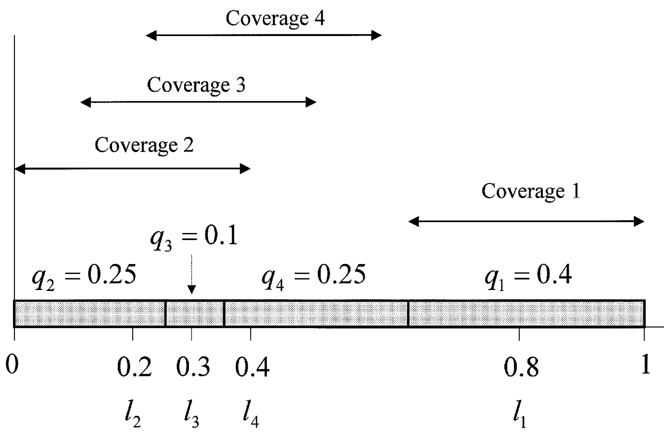

and 4 are closely clusted in attribute space, while Variant 1 is more distinct. For example, if the variants were shirts, Variants 2, 3 and 4 might be three shades of blue while Variant 1 is white. Customer t has an ideal point, L_t, that is uniformly distributed on [0, 1]. The utility of Variant j for customer t is then given by (4) where we used $a = 0.2$ and $b = 1$. We further assumed the no-purchase utility was identically zero for all customers.

The situation is depicted in Figure 5, where the four product locations are represented in the attribute space [0, 1]. The "coverage" intervals, which are of length 0.4, indicate the portions of attribute space where each variant has a nonzero utility. Customers with ideal point covered by an interval are willing to purchase the respective variant. Note that demand for Variant 1 is independent because its coverage interval does not overlap any other variant. Demand for Variants 2, 3 and 4, on the other hand, is dependent because the coverage intervals for these variants overlap.

Assuming all four variants are stocked ($S = \{1, 2, 3, 4\}$), this model gives choice probabilities of $q_1 = 0.4, q_2 = 0.25, q_3 = 0.1$ and $q_4 = 0.25$. This is shown in Figure 5, where the shaded bars at the bottom represent the regions where each of the variants is the closest and thus would be the first choice of customers lying in the intervals.

The heuristics and Sample Path Gradient were run with the above choice probabilities and Poisson arrivals with a mean of 30 and Q_t exponential with mean one as before. The price was $p = 100$ and the cost was $c = 1$ for all four variants. This extreme ratio of price to cost more clearly illustrates the inventory pooling effects and is not unreasonable for certain lost-sales inventory settings (e.g., c is the cost of holding a unit on the shelf for one day).

The inventory levels produced by each policy are shown in Figure 6; the average profits are shown in Table 4. Note that the Pooled Newsboy tends to understock all variants, because it assumes customers will freely substitute among all four variants. However, customers who prefer Variant 1 are not willing to substitute at all. As a result, Variant 1 is effectively independent, and thus both the Independent

Figure 6. Inventory levels for Example 3.

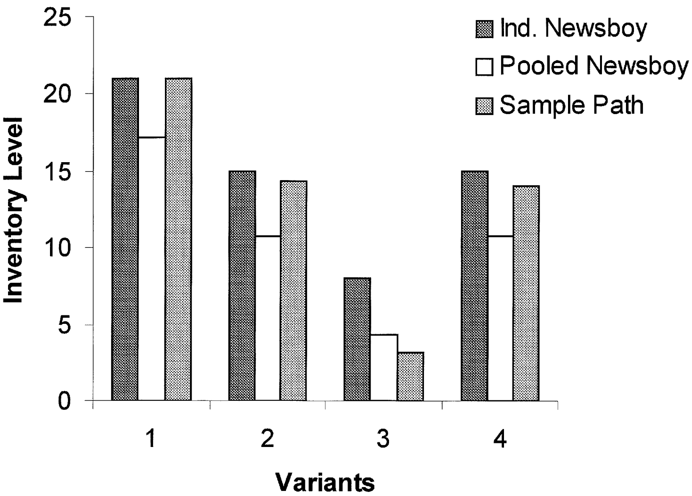

Table 4. Profits for policies for Example 3.

Policy	Profit
Sample Path Gradient	2944
Independent Newsboy	2940
Pooled Newsboy	2918

Newsboy and Sample Path Gradient policies give the same inventory level for Variant 1. On the other hand, there is significant substition among Variants 2, 3 and 4. The result is that the Pooled Newsboy overestimates the amount of pooling benefit due to substitution and performs relatively worse, with profits about 1% lower than the Sample Path Gradient.

Conversely, the Independent Newsboy stocks more of Variants 2, 3 and 4 than the Sample Path Gradient, because it does not account for the pooling among these variants. In particular, it is interesting to observe that the Sample Path Gradient concentrates its inventory on Variants 2 and 4, which together span the interval [0, 0.6] of the attribute space (see Figure 5)—the complement of the space spanned by Variant 1. However, because of the relatively low cost of inventory relative to lost sales (a ration of $p/c = 100$) the Independent Newsboy profits suffer a negligible loss for this overstocking.

6. PRICE AND SCALE EFFECTS

Under a static substitution model, we showed in van Ryzin and Mahajan (1999), that variety increases if the selling price of the items in a category rises. Similarly as the volume of customers increases, more variety is offered. We also showed that a particular definition of "fashion" based on the theory of majorization was a determinant of assortment profitability. The goal of the present analysis is to determine if these insights hold under a dynamic substitution model.

As in Example 1, we assume the price of all variants in the category is the same. Using the specific form of the utility function discussed earlier for Example 1, we see that the quality indices $\{a_j : j \in N\}$ and the price p completely characterize the vector $v = \{v^0, v^1 \dots, v^n\}$, where v^j is defined by (23).

To make comparisons across merchandise categories we need to understand which characteristics of the merchandise category affect the assortment decision. In van Ryzin and Mahajan (1999), it was shown that the "evenness" of the preference vector v played an important part. The theory of majorization was used to characterize this evenness or fragmentation of the preference vector. For a vector $y \in \Re^n$, let $[i]$ denote a permutation of the indices $\{1, 2, \dots, n\}$ satisfying $y^{[1]} \geqslant y^{[2]} \geqslant \cdots \geqslant y^{[n]}$. We then have the following definition of the partial order based on majorization:

DEFINITION 2. For $y, z \in< \Re^n$, y is said to be majorized by z, $y \prec z$ (z majorizes y), if $\sum_{i=1}^n y^{[i]} = \sum_{i=1}^n z^{[i]}$ and $\sum_{i=1}^k y^{[i]} \leqslant \sum_{i=1}^k z^{[i]}$, $k = 1, \dots, n-1$.

Intuitively, a nonnegative vector z that majorizes y tends to have more of its "mass" concentrated in a few components. In van Ryzin and Mahajan (1999), majorization was shown to provide the right measure for the degree of fragmentation in consumer preference, and the following definition was proposed:

DEFINITION 3. A merchandise category v is said to be more **fashionable** than w if $v \prec w$. If $v \prec w$, we refer to w as the **basic** category and v as the **fashion** category.

This definition says that for fashion categories, the vector of preferences v is more evenly spread across the variants. In our numerical test, we used the MNL model. The basic category w is given by (24) and the fashion category v is chosen so that $v \prec w$. Specifically,

$$a_j = 10.86 \qquad \forall j = 1, \dots, 10.$$

Results from our numerical work can be summarized as follows:

OBSERVATION 1. (i). *If $v \prec w$, and all other parameters are equal, then the basic category w is more profitable under optimal stocking decision.*

(ii). *Higher variety is offered when either price p or customer volume λ is increased.*

Part i of Observation 1, is consistent with our earlier results in van Ryzin and Mahajan (1999), using a static substitution model. Namely, fashion categories have higher overage and underage costs due to highly fragmented purchase choices which reduce their profitability. Figure 7 shows how the expected profit varies with price for the fashion and basic categories. We find that profit for the basic category is higher at all price levels.

Part ii implies that high selling prices create an incentive to stock higher levels of variety. As margins increase,

Figure 7. Profit as a function of price.

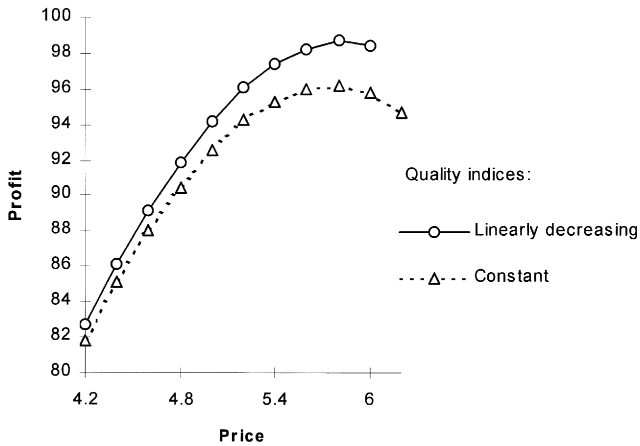

Table 5. Number of variants stocked vs. volume.

Mean # arrivals	# Variants stocked	Depth
10	3	0.899
30	5	0.933
100	9	0.829
1000	10	0.808

a wide variety is offered to minimize the likelihood of customers not purchasing. It also states that there are scale economies to offering variety. Results are shown in Table 5, where we include, as a measure of depth, the total number of units stocked per unit of mean demand.

Thus, as the volume of business increases, more variants are stocked. This is because, as the store traffic grows the relative costs of inventory overage and underage decrease, and more variety can be profitably offered. Again, these results are consis-tent with those obtained in van Ryzin and Mahajan (1999) for the static substitution model. We find that the depth of the assortment remains almost constant, showing only a slight decrease as the breadth is increased.

7. CONCLUSION

We have proposed a model to understand inventory decisions in retail assortments when consumers choose dynamically based on the on-hand stock. We believe the model is appealing for several reasons. First, consumer choice is based on maximization of stochastic (heterogeneous) utilities, which is a widely accepted mechanism in 44 economics for how rational consumers select from a mutually exclusive set of alternatives. Second, the stochastic processes assumed are essentially completely general. Finally, the model also leads to an efficient computational approach using sample path gradients.

Though we have performed only limited testing of heuristics, our results provide some interesting findings and suggest avenues for future research. The results showed that under substitution, one should stock relatively more of popular variants and relatively less of unpopular variants than a traditional newsboy analysis indicates. Intuitively, this is due to excess substitution demand combined with a reduced underage cost from having substitute variants as backups. This raises an interesting question of how to approximate these effects heuristically. At the same time, the simple Pooled Newsboy heuristic performs remarkably well in the equal-margin case of Example 1. Perhaps treating an entire category as if it were a single variant and then performing a simple allocation of the aggregate inventory is a reasonable way to manage such assortments in practice. On the other hand, Example 2 showed that this simple heuristic can produce bad decisions if margins are unbalanced. Also, the numerical results derived using the sample path gradient algorithm support the theoretical insights from the static substitution model in van Ryzin and Mahajan (1999).

The ability to incorporate price and choice effects in such a general framework suggests several other extensions. For example, one can easily study jointly setting price and assortment decisions to maximize profits as is suggested by Figure 7. A further extension is to use our model to study competitive effects in which each variant is managed by a separate firm seeking to maximize its own profits. This is the subject of a forthcoming paper Mahajan and van Ryzin (1999).

APPENDIX

We first prove two lemmas which are used subsequently.

LEMMA 4. *Let x and y be two starting inventory level vectors. Let C be such that $P(T \leqslant C) = 1$. Then we have that*

$$\|x_t - y_t\| \leqslant C_1 \|x - y\|,$$

for all $t = 1, \ldots, T+1$ for the continuous model, where $C_1 = 2^{C(n+1)}$.

PROOF. We first show that

$$\|x_2 - y_2\| \leqslant 2^{n+1} \|x_1 - y_1\| k. \tag{27}$$

Here $x = x_1$ and $y = y_1$. Since the square root is a subadditive function, we have that

$$\|x_2 - y_2\| \leqslant \sum_{i=1}^{n} \|x_2^i - y_2^i\|. \tag{28}$$

We start by showing equation (27) for the case $m = n+1$, i.e., the first customer values the no-purchase option the least.

Let j be such that $b(j) = 1$. We use equation (13), to write that

$$\|x_2^j - y_2^j\| \leqslant \|(x_1^{[1]} - Q_1)^+ - (y_1^{[1]} - Q_1)^+\| \tag{29}$$

$$\leqslant \|(x_1^{[1]} - y_1^{[1]})\|, \tag{30}$$

where we have used that for any two vectors $z, w \in \Re$,

$$\|z^+ - w^+\| \leqslant \|z - w\|. \tag{31}$$

We use induction to show that if $b(j)=k$, then

$$\|x_2^j - y_2^j\| \leqslant (2^k-1)\|x_1-y_1\|. \tag{32}$$

From equation (29), the result holds for $k=1$. We assume that the result holds for $r=1,\ldots,k-1$. We show that it holds for $r=k$. We use equations (13) and (31). We have that

$$\begin{aligned}
\|x_2^j - y_2^j\| &= \|[(x_1^{[k]}+\cdots+x_1^{[1]}-Q_1)^+ \\
&\quad -(x_1^{[k-1]}+\cdots+x_1^{[1]}-Q_1)^+] \\
&\quad -[(y_1^{[k]}+\cdots+y_1^{[1]}-Q_1)^+ \\
&\quad -(y_1^{[k-1]}+\cdots+y_1^{[1]}-Q_1)^+]\| \\
&\leqslant \|(x_1^{[k]}+\cdots+x_1^{[1]}-Q_1)^+ \\
&\quad -(y_1^{[k]}+\cdots+y_1^{[1]}-Q_1)^+\| \\
&\quad +\|[(x_1^{[k-1]}+\cdots+x_1^{[1]}-Q_1)^+ \\
&\quad -(y_1^{[k-1]}+\cdots+y_1^{[1]}-Q_1)^+]\| \\
&\leqslant \|x_1^{[k]}-y_1^{[k]}\|+2(2^{k-1}-1)\|x_1-y_1\| \\
&\leqslant (2^k-1)\|x_1-y_1\|,
\end{aligned}$$

where the last inequality follows by the induction hypothesis and equation (31). Using equations (28) and (32), we have that

$$\|x_2-y_2\| \leqslant (2^{n+1}-n-2)\|x_1-y_1\| \leqslant 2^{n+1}\|x_1-y_1\|.$$

If $m<n+1$, then for all $\{j : b(j)\leqslant m\}$, equation (32) holds as before. For $\{j : b(j)>m\}$, we have that

$$\|x_2^j-y_2^j\| = \|x_1^j-y_1^j\| \leqslant \|x_1-y_1\|. \tag{33}$$

In particular equation (32) is satisfied. So equation (27) holds.

Using equation (27), we have that

$$\|x_3-y_3\| \leqslant 2^{n+1}\|x_2-y_2\| \leqslant 2^{2(n+1)}\|x_1-y_1\|.$$

Therefore,

$$\|x_t-y_t\| \leqslant 2^{C(n+1)}\|x_1-y_1\|,$$

for all $t=1,\ldots,T+1$. □

LEMMA 5. *If h_1 and h_2 are Lipschitz with modulus K_{h1} and K_{h2} and a and b are scalars, then ah_1+bh_2 is Lipschitz with modulus $aK_{h1}+bK_{h2}$.*

PROOF. We use Lipschitz condition and the triangle inequality. □

PROOF OF LEMMA 3. We use the following terminology from Glasserman (1994). Let $\mathcal{M}$ denote the set of all mappings from $\Omega\times\Re^n$ to $\Re^n$. For $Y\in\mathcal{M}$, $Y(x)$ is its value at x and the argument ω is omitted. Let

$$\mathcal{D}=\{Y\in\mathcal{M} : Y \text{ is differentiable at } x \text{ a.s. } \forall x\in\Re^n\},$$

and Lip^1 be defined as the set

$$\mathrm{Lip}^1=\{Y\in\mathcal{M} : Y \text{ is a.s. Lipschitz, } E[K_Y]<\infty\}.$$

Then the following result is as stated in Glasserman (1994).

If $Y\in \mathrm{Lip}^1\bigcap\mathcal{D}$, and each $Y(x)$ is integrable then the derivatives $\{\nabla E[Y(x)], x\in\Re^n\}$ exist and

$$E[\nabla Y(x)]=\nabla E[Y(x)]$$

for all $x\in\Re^n$.

We have that $P(\eta(x,\omega)\leqslant BC)=1$ where B was defined to be such that $F_t(B)=1$ for all t and $P(T\leqslant C)=1$. Therefore $\eta(x,\omega)$ is integrable for all x. Also, since Q_t has a continuous distribution, we have from equation (13) that $f^j(x_t,U_t,Q_t)$ is a.s. differentiable for each $j\in N$ and each $t=1,\ldots,T$. So from equation (14) we have that $\eta_t^j(x_t,U_t,Q_t)$ is a.s. differentiable for each $j\in N$ and each $t=1,\ldots,T$. To use the result, we need to show that $\eta(x,\omega)$ is Lipschitz almost everywhere on Ω', i.e., there exists K_η such that,

$$\|\eta(x,\omega)-\eta(y,\omega)\| \leqslant K_\eta\|x-y\|.$$

By definition, $\eta(x,\omega)=x-x_{T+1}$, so

$$\begin{aligned}
\|\eta(x,\omega)-\eta(y,\omega)\| &= \|(x-y)-(x_{T+1}-y_{T+1})\| \\
&\leqslant \|x-y\|+\|x_{T+1}-y_{T+1}\| \\
&\leqslant (C_1+1)\|x-y\|. \quad \square
\end{aligned}$$

DEFINITION OF ∇f. To derive $\nabla f(x_t,U_t,Q_t)$, we use equation (13). We see that $\nabla f(x_t,U_t,Q_t)$ is an $n\times n$ matrix given as

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

where

$$\nabla f^j(x_t,U_t,Q_t)=\begin{pmatrix}\frac{\partial}{\partial x_t^l}f^j(x_t,U_t,Q_t)\\ \vdots\\ \frac{\partial}{\partial x_t^n}f^j(x_t,U_t,Q_t)\end{pmatrix}, \tag{34}$$

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

$$\frac{\partial}{\partial x_t^l}f^j(x_t,U_t,Q_t)=\begin{cases}\mathbf{1}_{\{x_t^{[b(j)-1]}+\cdots+x_t^{[1]}<Q_t\}} & \text{if } b(j)>b(l),\\ \times\mathbf{1}_{\{x_t^{[b(j)]}+\cdots+x_t^{[1]}>Q_t\}} & \\ \mathbf{1}_{\{x_t^{[b(j)]}+\cdots+x_t^{[1]}>Q_t\}} & \text{if } b(j)=b(l),\\ 0 & \text{if } b(j)<b(l),\end{cases} \tag{35}$$

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

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

We next prove a lemma, which shows that the derivative of the sample path sales function for the continuous problem can only take one of three values 0, 1 or −1.

LEMMA 6. *The partial derivatives for the sample path sales functions for the continuous problem satisfy*

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

and

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

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

PROOF. We start with $\nabla\eta_T^j(x_T,\omega)$. Using equation (17), we see that

$$\nabla\eta_T^j(x_T,\omega)=e^j-\nabla f^j(x_T,\omega_T).$$

Using equations (34) and (35) we see that the result holds for $\nabla\eta_T^j(x_T,\omega)$. We assume that the result holds for $k=t+1,\dots,T$ and show that it holds for $k=t$. Using equation (18), we have that

$$\begin{aligned}\nabla\eta_t^j(x_t,\omega)&=e^j-\nabla f^j(x_t,U_t,Q_t)\\&\qquad+\nabla f(x_t,U_t,Q_t)[\nabla\eta_{t+1}^j(f(x_t,U_t,Q_t),\omega)]\\&=e^j+\nabla f(x_t,U_t,Q_t)\\&\qquad\times(\nabla\eta_{t+1}^j(f(x_t,U_t,Q_t),\omega)-e^j).\end{aligned}\tag{36}$$

Let $(\partial/\partial x^i)f^j(x_t,U_t,Q_t)$ denote the ith partial derivative of $f^j(\cdot)$. Then we show the following:

$$\sum_{j=1}^{n}\frac{\partial}{\partial x^i}f^j(x_t,U_t,Q_t)=1,\tag{37}$$

for all $i=1,\dots,n$. We fix i. Let $S=\{r\in N: b(i)<b(r)\leqslant m\}$. If $b(i)>m$, then

$$\frac{\partial}{\partial x^i}f^j(x_t,U_t,Q_t)=\begin{cases}0 & j\neq i,\\ 1 & j=1.\end{cases}$$

So equation (37) holds in this case. If $b(i)\leqslant m$, then

$$\begin{aligned}\sum_{j=1}^{n}&\frac{\partial}{\partial x^i}f^j(x_t,U_t,Q_t)\\&=\sum_{r\in S}[\mathbf{1}_{\{x_t^{[b(r)-1]}+\cdots+x_t^{[1]}<Q_t\}}\mathbf{1}_{\{x_t^{[b(r)]}+\cdots+x_t^{[1]}>Q_t\}}]\\&\quad+\mathbf{1}_{\{x_t^{[b(i)]}+\cdots+x_t^{[1]}>Q_t\}},\end{aligned}$$

which equals 1. So equation (37) holds.

Using the induction assumption on $\eta_{t+1}^j(x_{t+1},\omega)$ we note that the jth element of the column vector $\nabla\eta_{t+1}^j(f(x_t,U_t,Q_t),\omega)-e^j$ can only equal -1 or 0. Also all remaining elements of the column vector other than the jth element can equal either 0 or 1.

From the above observation, and equations (37) and (36) we see that the result holds. □

PROOF OF THEOREM 4. We use a convergence theorem from Bertsekas and Tsitsiklis (1996), Chapter 4.

THEOREM 5 (BERTSEKAS AND TSITSIKLIS 1996). *Consider the algorithm*

$$y_{k+1}=y_k+a_k\nabla\pi(y_k,\omega k),$$

where the step sizes a_k *are nonnegative and satisfy*

$$\sum_{k=0}^{\infty}a_k=\infty\quad\text{and}\quad\sum_{k=0}^{\infty}a_k^2<\infty,$$

and $\nabla\pi(y_k,\omega_k)$ *is a random term. Let* $\mathcal{F}_k$ *denote the history of the algorithm until iteration* k. *Then*

$$\mathcal{F}_k=\{y_0,\dots,y_k,\nabla\pi(y_0,\omega_0),\dots,\nabla\pi(y_k,\omega_k),a_0,\dots,a_k\}.$$

Suppose there exists a function $h:\Re^n\to\Re$ *with the following properties*:

a. $h(x)>0$ *for all* $x\in\Re^n$.

b. *The function* h *is continuously differentiable and there exists some constant* L *such that*

$$\|\nabla h(x)-\nabla h(y)\|\leqslant L\|x-y\|,$$

for all $x,y\in\Re^n$.

c. *There exists a positive constant* m *such that*

$$m\|\nabla h(y_k)\|^2\leqslant[\nabla h(y_k)]^T E[\nabla\pi(y_k,\omega_k)|\mathcal{F}k],$$

for all $k=0,1,2,\dots$.

d. *There exist positive constants* K_1 *and* K_2 *such that*

$$E[\|\nabla\pi(y_k,\omega_k)\|^2|\mathcal{F}_k]\leqslant K_1+K_2\|\nabla h(y_k)\|^2.$$

Then the following holds with probability 1,

i. *The sequence* $h(y_k)$ *converges.*

ii. $\lim_{k\to\infty}\nabla h(y_k)=0$.

iii. *Every limit point of* y_k *is a stationary point of* h.

By using $h(x)=E[\pi(x,\omega)]$, we apply this theorem result to prove Theorem 4. Condition a can be satisfied by adding a sufficiently large constant to the profit function.

We next verify condition b From equation (16) and Lemma 5 we see that it is sufficient to show that there exists some constant L' such that

$$\|\nabla\eta^j(x,\omega)-\nabla\eta^j(y,\omega)\|\leqslant L'\|x-y\|.\tag{38}$$

From Lemma 6, we have that

$$\begin{aligned}\|\nabla\eta^j(x,\omega)-\nabla\eta^j(y,\omega)\|&\leqslant 2\sqrt{n}\\&\quad\times P(\nabla\eta^j(x,\omega)\neq\nabla\eta^j(y,\omega)).\end{aligned}\tag{39}$$

Using equations (18) we have that

$$P(\nabla\eta^j(x,\omega)\neq\nabla\eta^j(y,\omega))\leqslant\sum_{t=1}^{T}[p(\nabla f(x_t,U_t,Q_t)$$

$$\neq \nabla f(y_t, U_t, Q_t)), \tag{40}$$

and

$$P(\nabla f(x_t, U_t, Q_t) \neq \nabla f(y_t, U_t, Q_t)) \\ \leqslant \sum_{j=1}^{n}\sum_{i=1}^{n} P\left(\frac{\partial}{\partial x^i} f^j(x_t, U_t, Q_t) \neq \frac{\partial}{\partial x^i} f^j(y_t, U_t, Q_t)\right) \tag{41}$$

To illustrate the computation involved in the RHS of equation (41), we take a particular case when the choice made by customer t is such that $b(j) < b(i)$. Then from equation (35), we have that

$$P\left(\frac{\partial}{\partial x^i} f^j(x_t, U_t, Q_t) \neq \frac{\partial}{\partial x^i} f^j(y_t, U_t, Q_t)\right) \\ \leqslant P(x_t^{[b(j)-1]} + \cdots + x_t^{[1]} < Q_t) \\ \times P(y_t^{[b(j)-1]} + \cdots + y_t^{[1]} > Q_t) \\ + P(x_t^{[b(j)]} + \cdots + x_t^{[1]} > Q_t) \\ \times P(y_t^{[b(j)]} + \cdots + y_t^{[1]} < Q_t). \tag{42}$$

From Lemma 4, we have that

$$\|x_t^i - y_t^i\| \leqslant \|x_t - y_t\| \leqslant C_1\|x-y\|,$$

for all $i = 1, \ldots, n$. Therefore,

$$\left\|\sum_{i=1}^{b(j)}(x_t^{[i]} - y_t^{[i]})\right\| \leqslant \sum_{i=1}^{b(j)} \left\|x_t^{[i]} - y_t^{[i]}\right\| \leqslant C_1 b(j)\|x-y\|. \tag{43}$$

Using equations (42) and (43), and the Lipschitz property of the c.d.f. $F(\cdot)$ of Q_t, we see that

$$P\left(\frac{\partial}{\partial x^i} f^j(x_t, U_t, Q_t) \neq \frac{\partial}{\partial x^i} f^j(y_t, U_t, Q_t)\right) \\ \leqslant KC_1[b(j) + b(j-1)]\|x-y\|. \tag{44}$$

Carrying out a similar analysis for other components of the $\nabla f(x_t, U_t, Q_t)$ matrix and using equations (40), (41) and (44), we have that

$$P\big(\nabla\eta^j(x,\omega) \neq \nabla\eta^j(y,\omega)\big) \\ \leqslant KC_1\left[\frac{nC(2n^2+1)}{3} + (n-1)^2 + n^2\right]\|x-y\|. \tag{45}$$

From equations (45) and (39) it follows that the constant L' in equation (38) is

$$L' = 2\sqrt{n}KC_1\left[\frac{nC(2n^2+1)}{3} + (n-1)^2 + n^2\right].$$

We have now shown the Lipschitz property of the derivative of $h(x) = E[\pi(x,\omega)]$.

For the proof of part c, we note that

$$E\big[\nabla\pi(y_k, \omega_k)|\mathfrak{F}_k\big] = E\big[\nabla\pi(y_k, \omega_k)\big], \tag{46}$$

since the gradient of the sample path sales function for an inventory level y_k and sample path ω_k is not affected by how the inventory level vector y_k, is reached from the inventory level y_0.

Using Lemma 3 and fixing the value of $m = 1$, it follows that condition c is satisfied at equality for all $k = 0, 1, 2, \ldots$.

To prove condition d, we use Lemma 6. We see that

$$\|\nabla\pi(y_k, \omega_k)\| = \|p^T\nabla\eta^T(y_k, \omega_k) - c\| \\ < \left\|-\left(\sum_{i=1}^{n} p_i\right)\mathbf{1} - c\right\|, \tag{47}$$

for all $k = 0, 1, 2, \ldots$. This is possible since the partial derivative $(\partial/\partial y_k^i)\eta^j(y_k, \omega_k)$ cannot be lower than -1, while the partial derivative $(\partial/\partial y_k^i)\eta^j(y_k, \omega_k)$ cannot be lower than 0, but even that has been chosen to be -1. Then with

$$K_1 = \left(\left\|-\left(\sum_{i=1}^{n} p_i\right)\mathbf{1} - c\right\|\right)^2 \tag{48}$$

and K_2 arbitrarily small and positive, we see that condition d holds. ☐

REFERENCES

Agrawal, N, S. A. Smith. 1996. Estimating negative binomial demand for retail inventory management with unobservable lost sales. *Naval Res. Logist.* **43** 839–861.

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. 1997. A Dynamic model of consumer demand with stock-out based substitution. Working Paper, Kellogg School of Management, Northwestern University.

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

Bassok, Y., R. Anupindi, R. Akella. 1997. Single period multiproduct inventory models with substitution. *Oper. Res.* Forthcoming.

Bertsekas, D. P., J. N. Tsitsiklis. 1996. *Neuro-Dynamic Programming*, Athena Scientific, Belmont, MA.

Bitran, G., S. Dasu. 1992. Ordering policies in an environment of stochastic yields and substitutable demands. *Oper. Res.* **40**(5), 177–185.

Garcia, C. B., W. I. Zangwill. 1981. *Pathways to Solutions, Fixed Points and Equilibria*, Prentice-Hall, NJ.

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

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

Hanson, W., Kipp Martin. 1996. Optimizing multinomial logit profit functions. *Management Sci.* **42** 992–1003.

Lancaster, K. 1990. The economics of product variety: A survey. *Marketing Sci.* **9** 189–210.

Mahajan, S., G. J. van Ryzin. 1999. A multi-firm stocking game under dynamic consumer substitution. *Oper. Res.* Forthcoming.

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

Pasternack, B., Z. Drezner. 1991. Optimal inventory policies for substitutible commodities with stochastic demand. *Naval Res. Logist.* **38** 221–240.

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

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

Topkis, D. M. 1978. Minimizing a submodular function on a lattice. *Oper. Res.* **26** 305–321.

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

Veinott Jr., A. 1965. Optimal policy for a multi-product, dynamic, nonstationary inventory problem. *Management Sci.* **12**(3) 206–222.