

Asymmetric Consumer Learning and Inventory Competition

Vishal Gaur* and Young-Hoon Park†

February 2005

Abstract

We develop a model of consumer learning and choice behavior in response to uncertain service at the marketplace. Learning could be asymmetric, i.e., consumers may associate different weights with positive and negative experiences. Under this consumer model, we characterize the steady-state distribution of demand for retailers given that each retailer holds constant in-stock service level. We then consider a non-cooperative game at the steady-state between two retailers competing on the basis of their service levels. Our model yields a unique pure strategy Nash equilibrium. We show that asymmetry in consumer learning has a significant impact on the optimal service levels, market shares and profits of the retailers. When retailers have different costs, it also determines the extent of competitive advantage enjoyed by the lower cost retailer.

KEYWORDS: Asymmetric Consumer Learning; Customer Satisfaction; Inventory Competition; Retail Operations

*Stern School of Business, New York University, Suite 8-160, 44 West 4th St., New York, NY 10012; phone: (212) 998-0297; fax: (212) 995-4227; e-mail: vgaur@stern.nyu.edu.

†Johnson Graduate School of Management, Cornell University, 330 Sage Hall, Ithaca, NY 14853-6201; phone: (607) 255-3217; fax: (607) 254-4590; e-mail: yp34@cornell.edu.

1 Introduction

Optimal policies in traditional inventory theory are generally evaluated for exogenous demand functions with the assumption that consumer behavior and demand are unaffected by the operational decisions of the firm. In many cases, however, common assumptions of independence between consumer behavior and operational decisions are not desirable or reasonable. In practice, customers do react substantially and negatively to poor service (e.g., stockouts) which may lead them to switch retailers on subsequent trips (Fitzsimons 2000), and which may have a significant adverse effect on future demand (Anderson, Fitzsimons, and Simester 2003). On the other hand, satisfied customers are likely to continue to buy from the same firm (Anderson and Sullivan 1993). These findings raise an important research question: Can the inventory decision of a firm be improved by endogenously incorporating consumer response to uncertain service? Further, it is well-documented in the behavioral research literature that consumers are biased and react asymmetrically to satisfying and unsatisfying experiences (e.g., Kahneman and Tversky 1979; Tversky and Kahneman 1991). The impact of asymmetry in consumer response on inventory decisions has not been studied in the existing literature.

This paper presents a model to analyze competitive inventory decisions by explicitly recognizing the above findings. We make four key assumptions based on consumer behavior theory. First, a consumer is not well informed about product availability at the marketplace. Instead, the consumer learns about the firm's service level from her prior shopping experience and modifies her future shopping behavior in response to the service level provided by the firm. Thus, a consumer switches retailers as a result of the history of service she actually receives at various retailers, rather than as an immediate response to poor service at the current retailer. Second, consumers are biased and respond differently to positive and negative experiences. This bias reflects asymmetric

response to gains versus losses documented in behavioral literature. Third, we consider the notion of diminishing sensitivity over time, i.e., consumers weigh recent experiences more heavily than older experiences. Finally, consumers exhibit diminishing sensitivity to service level, i.e., they react less positively to satisfying shopping experiences when they perceive the service level to be very high. Under this consumer model, we analyze the effect of asymmetry in consumer learning on inventory competition between retailers at the marketplace.

We use a stylized model of a retailer in which each retailer maintains a constant fill-rate over an infinite time horizon. Retailers differ in their costs but are price-takers at the marketplace. Thus, fill-rate is the only dimension of competition in the market. Product availability is a fundamental issue in marketing and operations management. Many practitioners have documented the importance of fill-rate practices in retail firms, and researchers have considered fill-rate strategies in analytical models. For example, Germain and Cooper (1990) find that firms set target fill-rates as relatively long-term strategic variables; Elman (1989) documents in-stock levels for grocery retailers; and Bass (1989) conducts a similar study on general merchandise catalogers. Dana (2001) describes a nice experiment to estimate product availability at video rental stores and finds that the sampled movie titles were available 86% of the time at Blockbuster and 60% of the time at competing stores, consistent with Blockbuster's advertising campaign emphasizing high fill-rates (e.g., "Go Home Happy"). See also Bernstein and Federgruen (2003) for other examples.

We show that under the assumption of constant fill-rates, our choice model yields a steady-state distribution of the consumers' share of purchases at various retailers as a function of the service levels of all retailers and the degree of asymmetry in consumer learning. We then formulate a non-cooperative game at the steady-state between retailers competing on the basis of their service levels. The analysis of this game yields the following findings:

- A retailer always benefits from considering the dependence of its demand on its service level

and provides a higher service level than if its demand were exogenous. This result is consistent with the existing literature.

- There is a unique pure strategy Nash equilibrium in the inventory competition.
- Inventory competition results in a reduction in total industry profits and an increase in total inventory levels at the marketplace compared to the case where the retailers treat their demands as exogenous.
- Consumer learning bias is an important component of the optimal service levels, market shares and profits of the retailers. When consumers are biased towards negative experiences, the firms offer higher optimal service levels. Further, the total inventory in the industry increases and the total expected profit decreases. Individual firms' inventory and expected profit may increase or decrease depending on cost asymmetry.
- When retailers have different costs, consumer learning bias also determines the extent of competitive advantage enjoyed by the lower cost retailer. The firm with the higher cost suffers more erosion in market share and profit when consumers are biased towards negative experiences than when they are biased towards positive experiences.

Our results demonstrate that there is an interaction between consumer learning bias and the impact of differing costs for competing retailers. Low cost retailers have greater market power when consumers are biased towards negative experiences than when they are biased towards positive experiences. Further, we show that inventory decisions can be significantly improved by modeling the interaction between marketing (asymmetric consumer response to uncertain service) and operations (inventory costs). For optimal inventory planning decisions, therefore, it would be important to properly evaluate the extent of consumer response to uncertain service at the marketplace.

The remainder of the paper is organized as follows. Section 2 reviews related literature. Section 3 presents the consumer demand model and derives the steady-state market share of each retailer. In §4, we investigate a non-cooperative game between two retailers and further analyze the market equilibrium based on the parameters of interest. We discuss the limitations of our model and conclude with directions for future research in §5. All proofs are provided in Appendix B unless otherwise noted.

2 Literature Review

There has been considerable research in recent years on operational decisions under endogenous demand models in a variety of competitive settings. Demand endogeneity could arise from stock-outs causing overflow of customers to competing firms or from consumer choice driven by price, service or quality levels. Endogenous demand models have been studied at both aggregate market and individual consumer levels. Among individual-level models, significant advances have been made on consumer choice under perfect information as well as consumer learning under imperfect information. Distinguishing characteristics of models studied in this stream of research are as follows.

Gans (2002) models consumer learning and choice in response to random variation in the quality provided by competing suppliers. He develops an individual-level consumer demand model in which consumers use Bayesian updating to learn from their own experiences with the quality levels offered by suppliers. In each period, a consumer picks the supplier for which the consumer has the highest expectation of service level. Gans derives the steady-state characterization of this demand model when suppliers choose static quality policies and analyzes the competition between service providers competing on quality of service.

Hall and Porteus (2000) consider consumer response to operational decisions using finite-horizon models of both quality and inventory competition. Unlike Gans, they consider dynamic policies wherein a firm can change its service level in each period in response to changes in its market share. Moreover, in their consumer behavior model, a consumer switches suppliers as an immediate response to a service failure, rather than as a result of the history of service received at competing suppliers. The rate of switching is mediated by an external loyalty parameter.

In the inventory literature, Dana and Petruzzi (2001) consider a single-period model of inventory and price decisions when consumers have perfect information about the price and inventory level offered by a subject retailer but not about each other's valuation of the good. Consumers choose between visiting the retailer and an outside option with given valuation in order to maximize their expected utility. Deneckre and Peck (1995) consider a non-cooperative game between several retailers under a similar information structure. Dana (2001) considers a generalization of Deneckre and Peck, where consumers do not observe the firms' inventory decisions but realize that firms with higher prices are likely to offer higher service levels.

While the above models examine consumer responses to stocking decisions with and without consumer learning, researchers have also studied competitive inventory decisions using aggregate demand models. For instance, Bernstein and Federgruen (2003) analyze inventory and price competition in an infinite horizon setting using a demand model based on attraction models of market share. They show the existence of a Nash equilibrium of stationary strategies, i.e., service level and price are chosen by each retailer at the outset and kept constant throughout the time horizon. Tsay and Agrawal (2000) consider a single-period model of a two-stage distribution system with one manufacturer and two retailers. Retailers compete on price and service levels given that each retailer's demand is a linear function of the prices and service levels offered by both retailers. Li (1992) examines the optimal choice between make-to-stock and make-to-order policies when con-

sumers arrive according to a Poisson process and choose between firms based on price, quality and delivery time. Cachon and Netessine (2003) conduct a survey of applications of game theory to supply chain analysis, and specifically, newsvendor games.

Several researchers have considered demand substitution (overflow) due to stockouts, i.e., consumers switch from one retailer to another within the same period when their first-choice retailer is out of stock. Parlar (1988), Karjalainen (1992), Lippman and McCardle (1997), and Netessine and Rudi (2003) (see also references cited therein) analyze single-period models of demand substitution. The initial demand at each retailer is exogenous, but the demand in excess of inventory is allocated to competing retailers in deterministic proportions. These papers find the existence of Nash equilibria wherein retailers stock more inventory than if excess demand were not allocated to the competition. Netessine, Rudi and Wang (2003) model the allocation of excess demand in a multi-period setting, wherein unsatisfied demand at the end of a period may be lost, or may be back-ordered according to a menu of allocation rules for delivery in a subsequent period. Anupindi, Dada and Gupta (1998) present a methodology by which the parameters of the demand substitution model can be estimated from augmented sales data. Mahajan and van Ryzin (2001) consider an individual-level model of demand substitution in which a stochastic sequence of heterogeneous consumers choose dynamically between competing retailers using utility maximization criteria.

It is to be noted that operations models analyzing the effect of service level on long-run demand and optimal inventory levels have a long history in the literature. Schwartz (1966; 1970) was probably the first to consider the concept that a stockout may not impose an immediate penalty on the retailer but may affect the distribution of its future demand. He computes optimal inventory order-up-to policies when mean demand and standard deviation of demand are stylized functions of the service level offered. Ernst and Cohen (1992) and Ernst and Powell (1995; 1998) extend this research to two-stage supply chains.

Our paper differs from the existing research in that we consider demand derived from an individual-level consumer learning and choice model. Our model has innovative features of consumer learning, asymmetric response to positive and negative experiences, and diminishing sensitivity over time and levels of satisfaction. While these findings on consumer behavior are well established in behavioral literature, their effects on operational decisions have not been evaluated in the existing research. Further, we employ the multinomial logit model to explain store choice. As a result, consumers in our model seek variety in their shopping trips and have a non-zero probability of shopping at the retailer for which they do not perceive the highest expected service level. Like Gans (2002) and Bernstein and Federgruen (2003), our paper is based on a steady-state characterization of the inventory competition, i.e., retailers choose their service levels at the start of the game and maintain the same service levels throughout. We obtain parametric results to demonstrate the impact of asymmetric learning on the optimal strategies and equilibrium outcomes not only when the retailers have identical costs but also when they have asymmetric costs.

3 Model Formulation

3.1 Notation and Assumptions

We consider a model with two retailers selling a single item to a fixed population of N consumers. At discrete time periods, $t = 1, \dots, \infty$, each consumer demands one unit of the item with probability ω , $0 < \omega \leq 1$, and chooses a retailer to visit. The selling price of the item, denoted r , is constant over time and identical across both retailers. Thus, there is no competition based on price. Instead, the retailers compete with each other based on service levels.

We define the service levels offered by the retailers similar to Dana and Petruzzi (2001) and Deneckere and Peck (1995). Let \tilde{X}_{st} denote the aggregate demand at retailer s at time t , λ_{st} denote

the mean of \tilde{X}_{st} , and q_{st} denote the inventory level of retailer s at time t . The *ex ante* service level provided by retailer s , denoted f_s , is defined as the ratio of expected sales to mean demand:

$$f_s \equiv \frac{\mathbb{E}[\min\{q_{st}, \tilde{X}_{st}\}]}{\lambda_{st}}, \quad \text{for all } t. \quad (1)$$

We use the fact that the *ex ante* expectation of the service level experienced by a consumer who decides to visit retailer s at time t is identical to f_s .¹ To see this, note that the unconditional probability that consumer i visits retailer s at time t is λ_{st}/N and the conditional probability that consumer i visits retailer s at time t , given that \tilde{X}_{st} equals k , is k/N . Therefore, from Bayes' Theorem, if consumer i decides to visit retailer s at time t , the expected service level observed by consumer i is given by:

$$\begin{aligned} & \mathbb{E} \left[\frac{\min\{q_{st}, \tilde{X}_{st}\}}{\tilde{X}_{st}} \middle| \text{consumer } i \text{ visits retailer } s \text{ at time } t \right] \\ &= \sum_{k=0}^N \frac{\min\{q_{st}, k\}}{k} \Pr[\tilde{X}_{st} = k | \text{consumer } i \text{ visits retailer } s \text{ at time } t] \\ &= \sum_{k=0}^N \frac{\min\{q_{st}, k\}}{k} \frac{\Pr[\tilde{X}_{st} = k](k/N)}{\lambda_{st}/N} \\ &= f_s. \end{aligned}$$

We use the terms service level and fill-rate interchangeably. We assume that f_s is time-invariant, i.e., each retailer makes a strategic decision about its service level and provides the same service level throughout the time horizon. We show that, under this assumption, our consumer behavior model yields steady-state distributions of aggregate demand at each retailer. Thus, we shall ignore transient behavior and formulate a single-period non-cooperative game between the retailers to determine equilibrium outcomes.

¹Service level can alternatively be defined as $\Pr[\tilde{X}_{st} \leq q_{st}]$, see for example Bernstein and Federgruen (2003). Since we consider an individual-level demand model, we use the definition in (1) because it gives the expected service level experienced by a consumer if she decides to visit a particular retailer.

The central aspect of this paper is that consumers do not know the value of f_s for any retailer in advance. Instead, each consumer forms a private belief about the service level offered by each retailer through her prior shopping experience at that retailer. She then chooses the retailer to shop at in each period and updates her beliefs after the visit. Thus, the demand faced by each retailer depends on its own and its competitor's service levels. Section 3.2 specifies the consumer choice model.

The remaining assumptions and notation are as follows. While the selling price is identical across retailers, the cost parameters may be different. For retailer s , let c_s denote the procurement cost per unit of the item, and s_s denote the salvage value per unit of the inventory left over at the end of each period. Here, $r > c_s > s_s$. We assume that inventory is not carried over from one period to the next. Hereafter, the word 'store' is used interchangeably with 'retailer'. We index the decisions of the subject retailer by s and its competitor by \bar{s} . Where convenient, we refer to the subject retailer as retailer 1 and its competitor as retailer 2. All results of this paper apply when there are more than two retailers at the marketplace.

3.2 Consumer Demand Model

We now describe how a consumer buys a focal product at the marketplace and utilizes her shopping experience in her future behavior with respect to store choice. Let p_{st}^i denote consumer i 's estimate of the service level at retailer s at time t . A visit to a retailer is called *satisfying* if the consumer does not experience a stockout and *unsatisfying* otherwise. In each period t when consumer i decides to purchase the item, she chooses the retailer to visit using the values of p_{st}^i , and then computes $p_{s,t+1}^i$ based on the outcome of her visit. The consumer choice and learning process are as follows.

Step 1: Store choice Our model of store choice is based on the multinomial logit model (e.g., McFadden 1974). We assume that consumer i 's indirect utility from retailer s at time t is given by the additive form,

$$u_{st}^i = w_{st}^i + \epsilon_{st}^i, \quad (2)$$

where w_{st}^i and ϵ_{st}^i are the deterministic and random components of u_{st}^i , respectively. The deterministic component is specified by $w_{st}^i = \alpha + \beta \cdot r + \log(p_{st}^i)$, where α is the consumer's nominal utility from purchasing the item, $\beta \cdot r$ incorporates the effect of price, and $\log(p_{st}^i)$ incorporates consumer i 's estimate of the service level offered by retailer s at time t . A logarithmic form is used for the service level because (1) when $p_{st}^i = 0$, we have $\log(p_{st}^i) = -\infty$, implying that if retailer s offers a zero service level, retailer s will never be preferred to its competitor; (2) an increase in p_{st}^i represents decreasing disutility from retailer s ; (3) when $p_{st}^i = 1$, we have $\log(p_{st}^i) = 0$ so that the consumer's utility is determined entirely by the nominal utility and the effect of the price. It is well-known that, under the assumptions of utility maximization and an independent and identically distributed type I extreme value distribution for ϵ_{st}^i , the probability that consumer i chooses retailer s at time t is given by:

$$\eta_{st}^i = \frac{p_{st}^i}{p_{st}^i + p_{st}^j}. \quad (3)$$

Thus, the probability that a consumer visits a given retailer is increasing in the consumer's estimate of the service level at that retailer. From (3), note that the consumer in our model seeks variety in her shopping trips and has a non-zero albeit smaller probability of visiting the store with the lower p_{st}^i . The multinomial logit model represents a very flexible choice model for consumer demand, such has been employed in a variety of research disciplines (see for example, Ben-Akiva and Lerman 1985; Guadagni and Little 1983; McFadden 1974; van Ryzin and Mahajan 1999).

Step 2: Consumer learning At each shopping occasion, the consumer updates p_{st}^i by the rule,

$$p_{s,t+1}^i = \begin{cases} (1 - \theta^u) \cdot p_{st}^i + \theta^u & \text{satisfying visit to store } s \text{ at time } t, \\ (1 - \theta^d) \cdot p_{st}^i & \text{unsatisfying visit to store } s \text{ at time } t, \\ p_{st}^i & \text{no visit to store } s \text{ at time } t, \end{cases} \quad (4)$$

where $\theta^u \in (0, 1)$ is a weight attached by the consumer to a satisfying store visit and $\theta^d \in (0, 1)$ is a weight attached to an unsatisfying store visit. If the visit to store s is satisfying for consumer i , her estimate of the service level at store s increases by $\theta^u \cdot (1 - p_{st}^i)$, otherwise it decreases by $\theta^d \cdot p_{st}^i$. The value of p_{st}^i remains unchanged if the consumer does not visit store s at time t . Since both θ^u and θ^d are between zero and one, we have $0 < p_{st}^i < 1$ for all t . We denote the ratio θ^u / θ^d by θ .

This model captures both positive and negative biases in consumer learning. We note that the consumer is biased towards positive experiences if $\theta^u > \theta^d$, and towards negative experiences if $\theta^u < \theta^d$. If θ^u is equal to θ^d , there is no bias and our model is identical to simple exponential smoothing.

As noted earlier, the updating rule specified in (4) exhibits diminishing sensitivity over time and over levels of satisfaction. With respect to time, a recent experience is weighted more heavily than an older experience. With respect to levels of satisfaction, the marginal impact of a satisfying visit to a retailer is decreasing in the consumer's estimate of the service level at that retailer.

The parameters θ^u and θ^d can be estimated by comprehensive marketing research. In practice, both parameters may vary across product categories. For a necessity item such as milk, for instance, consumers may react more strongly to a stockout than to a satisfying visit since they might take it for granted that such items should always be in stock. For a search item such as fashion clothing, on the other hand, consumers may weigh a satisfying visit much more than a stockout of one of the variants because they expect to search for these items. Further, the parameters of consumer

learning may be specific to individual consumers. To achieve a parsimonious analytical framework, we consider a homogeneous population of consumers, i.e., the learning parameters are common across individuals.

We assume that if the consumer's visit to a retailer in period t is unsatisfying, then she decides not to purchase the item in period t and does not visit competing retailers in that period. In other words, we do not allow demand substitution across retailers within the same period. For research on inventory models with demand substitution, see Lippman and McCardle (1997), Mahajan and van Ryzin (2001), Netessine and Rudi (2003), Parlar (1988) and the references cited therein. For research using an alternative assumption that excess demand is backlogged, see Bernstein and Federgruen (2003). Since our paper is among the first attempts to develop a competitive inventory model with asymmetric consumer learning and consumer choice at the individual level, we seek to keep the model as parsimonious as possible to highlight the key aspects of inventory competition.

3.3 Steady-state Aggregate Demand

We first compute the steady-state distribution of demand for each retailer. To this end, we analyze the convergence of the sequence $\{p_{st}^i\}$ for any consumer i and retailer s as t goes to ∞ . Proposition 1 shows that if a retailer provides a constant non-zero service level, the retailer is visited infinitely often by each consumer.

Proposition 1 *Given $f_s > 0$, if $p_{s1}^i > 0$ for any consumer i , then consumer i visits retailer s infinitely often, and the expected time between successive visits by consumer i to retailer s is finite.*

We note that Proposition 1 holds for a retailer even when consumer i perceives a perfect service level at the competing retailer. As a result, the competitor of retailer s cannot force retailer s to exit the market by temporarily offering a very high service level. Therefore, each retailer can have

a non-zero expected market share as t goes to ∞ .

Proposition 2 shows that, for all consumers i , p_{st}^i converges in distribution to a random variable, p_s , as t goes to ∞ .

Proposition 2 *There exists a random variable p_s , $0 < p_s < 1$, such that p_{st}^i converges in distribution to p_s for each consumer i as $t \rightarrow \infty$. Further,*

$$E[p_s] = \frac{\theta^u \cdot f_s}{\theta^d + (\theta^u - \theta^d) \cdot f_s} = \frac{\theta \cdot f_s}{1 + (\theta - 1) \cdot f_s}. \quad (5)$$

According to (5), $\theta^u > \theta^d$ implies $E[p_s] > f_s$, and $\theta^u < \theta^d$ implies $E[p_s] < f_s$. Thus, the consumers' propensity to place different weights on positive and negative experiences results in their overestimating the service level provided by each retailer if $\theta^u > \theta^d$ (i.e., $\theta > 1$), and underestimating it otherwise. Thus, Proposition 2 shows that θ^u and θ^d shape near-term behavior in such a way as to create long-term differences between actual and perceived service levels. Therefore, to determine the inventory policies of retailers, it is critical to properly consider how consumers learn and update their private beliefs regarding the service levels at the marketplace.

4 Inventory Competition

In this section, we model the retailers' inventory decisions and analyze the competition at the steady-state as a non-cooperative game. We assume that the retailers' cost parameters as well as the market parameters, N , ω , θ^u and θ^d are common knowledge. Each retailer sets its service level at the outset. The service levels offered by the retailers jointly determine their long-run average market shares and profits. The long-run average profit of retailer s is given by

$$\lim_{T \rightarrow \infty} \frac{1}{T} \left[\sum_{t=1}^T \pi_{st} \right] = \lim_{T \rightarrow \infty} \frac{1}{T} \left[\sum_{t=1}^T r \min\{\tilde{X}_{st}, q_{st}\} - c_s q_{st} + s_s \max\{q_{st} - \tilde{X}_{st}, 0\} \right]$$

where \tilde{X}_{st} and q_{st} are as defined in §3.1, and π_{st} denotes the profit realized by retailer s at time t . Since Proposition 2 shows that p_{st}^i converges in distribution to p_s for all customers i , it can

easily be shown that there exists a random variable \tilde{X}_s such that the aggregate demand at retailer s , \tilde{X}_{st} , converges in distribution to \tilde{X}_s as t tends to infinity. Additionally, since q_{st} is measurable with respect to the information available up to time $(t-1)$, it is a Cauchy sequence and converges to a stationary value q_s in the limit as t tends to infinity. Then, the convergence of the profit function in distribution follows from weak convergence theory (Durrett 1996:§2.2). Therefore, the long-run average profit maximization problem of each retailer reduces to a single-period problem at the steady state. Letting π_s denote the profit of retailer s for the single-period steady-state problem, we have

$$\lim_{T \rightarrow \infty} \frac{1}{T} \left[\sum_{t=1}^T \pi_{st} \right] = \mathbb{E} \left[r \min\{\tilde{X}_s, q_s\} - c_s q_s + s_s \max\{q_s - X_s, 0\} \right] \equiv \mathbb{E}[\pi_s(q_1, q_2)].$$

According to the individual-level demand model described in §3, the probability that a given consumer visits retailer s in the steady-state is a function of p_s and $p_{\bar{s}}$, which are random variables. Thus, the steady-state aggregate demand for each retailer is the sum of the outcomes of Bernoulli trials across all the consumers at the marketplace. (Note that the values of p_{st}^i are not independent of each other even though they have the same limiting marginal distribution, p_s .) While this detailed model can be used in numerical studies, it is not directly amenable to obtaining analytical results regarding the retailers' inventory levels. Therefore, for analytical tractability, we approximate \tilde{X}_s by a continuous-valued random variable denoted X_s . Also let $X \equiv X_s + X_{\bar{s}}$ denote the total steady-state market demand, with probability density function $g(x)$ and cumulative distribution function $G(x)$. Note that $\mathbb{E}[X] = N\omega$ from the definition of the market in §3.1. We assume that $g(x)$ is positive on a compact subset of \mathfrak{R}^+ and 0 elsewhere in order that the optimal inventory policy is uniquely specified.

We further make the following three assumptions:

Assumption 1 *Using (3) and Proposition 2, the probability that any consumer chooses retailer s*

at the steady-state is approximated by the expression,

$$v_s = \frac{E[p_s]}{E[p_s] + E[p_{\bar{s}}]}. \quad (6)$$

Assumption 2 *The aggregate demand at each retailer s is scaled by that retailer's market share.*

In other words, $X_s = v_s X$ with expectation $E[X_s] = v_s E[X] = v_s N\omega$.

Assumption 3 $\theta \geq 0.5$.

Assumption 1 enables us to derive the steady-state aggregate demand at each retailer as a function of the service levels offered by both retailers. We interpret v_s as the market share of retailer s at the steady-state. Equation (6) is consistent with the attraction models of market share (e.g., Cooper and Nakanishi 1988; Bernstein and Federgruen 2003) and further incorporates our consumer learning model since $E[p_s]$ and $E[p_{\bar{s}}]$ are functions of $f_s, f_{\bar{s}}$ and θ as shown in Proposition 2.

Assumption 2 enables us to compute the inventory level corresponding to a given service level in closed form, and thus, obtain comparative statics results. Using Assumptions 1 and 2 in (1), the service level at retailer s corresponding to inventory level q_s at the steady-state can be written as:

$$f_s = \frac{E[\min\{q_s, X_s\}]}{v_s N\omega} = \frac{E[\min\{Q_s, X\}]}{N\omega},$$

where $Q_s = q_s/v_s$. Q_s can be interpreted as the level of inventory offered by retailer s to achieve service level f_s if retailer s were the only retailer in the market. Note that there is a one-to-one correspondence between f_s and Q_s . Finally, Assumption 3 limits extreme biases toward negative experiences in consumer behavior, i.e., scenarios where $\theta^d > 2\theta^u$.

To examine the implications of Assumptions 1 and 2, we conduct a comprehensive simulation study comparing the theoretical distribution of demand, obtained from Assumptions 1 and 2, with the empirical distribution of demand obtained from the individual-level model in §3. Appendix A

presents the results of this study. Further, to avoid potential limitations imposed by Assumption 2, we do not use this assumption in the numerical analysis to be described later in §4.2. Instead, we approximate the total market demand by a Poisson distribution with mean $N\omega$ and the demand for retailer s by a Poisson distribution with mean $N\omega v_s$ for $s = 1, 2$.

The expected profit of retailer s at the steady-state corresponding to inventory levels q_1 and q_2 can now be written as:

$$\begin{aligned}
\mathbb{E}[\pi_s(q_1, q_2)] &\equiv \mathbb{E}[\pi_s(Q_1, Q_2)] &= r\mathbb{E}[\min\{q_s, X_s\}] - c_s q_s + s_s \mathbb{E}[\max\{q_s - X_s, 0\}] \\
& &= (r - s_s)v_s \mathbb{E}[\min\{Q_s, X\}] - (c_s - s_s)Q_s v_s \\
& &= v_s h_s(Q_s),
\end{aligned} \tag{7}$$

where $h_s(Q_s) = (r - s_s)\mathbb{E}[\min\{Q_s, X\}] - (c_s - s_s)Q_s$ is the familiar newsvendor profit function with exogenous demand X . Thus, the expected profit is a product of a scale variable and a scale-independent profit function. We note that similar profit functions have been used in other contexts, e.g., when price is endogenous (Agrawal and Seshadri 2000; Petruzzi and Dada 1999), when demand is an aggregate-level function of the service level (Dana and Petruzzi 2001; Bernstein and Federgruen 2003), and in a service competition when there are no economies of scale (Gans 2002). However, the form of the scale variable differs across all these models including ours.

We use the notation v'_s, v''_s, h'_s and h''_s to denote derivatives with respect to Q_s . The following properties of v_s are useful in the subsequent analysis.

Lemma 1 *For given $Q_{\bar{s}}$, v_s is strictly increasing and concave in Q_s .*

4.1 Myopic Versus Strategic Retailers

A retailer is called *strategic* if it recognizes that its demand distribution depends on its inventory level through consumer learning. The strategic retailer chooses Q_s to maximize (7) subject to (6)

treating v_s as endogenous. Let $F_s(Q_1, Q_2)$ denote the first derivative of π_s with respect to Q_s .

Thus, the first-order condition is:

$$F_s(Q_1, Q_2) \equiv \frac{\partial \mathbb{E}[\pi_s(Q_1, Q_2)]}{\partial Q_s} + \frac{\partial \mathbb{E}[\pi_s(Q_1, Q_2)]}{\partial v_s} \frac{dv_s}{dQ_s} = 0, \quad (8)$$

where

$$\frac{\partial \mathbb{E}[\pi_s(Q_1, Q_2)]}{\partial Q_s} = v_s h'_s(Q_s) = v_s [(r - s_s) \{1 - G(Q_s)\} - (c_s - s_s)],$$

and

$$\begin{aligned} \frac{\partial \mathbb{E}[\pi_s(Q_1, Q_2)]}{\partial v_s} \frac{dv_s}{dQ_s} &= h_s(Q_s) \frac{\mathbb{E}[p_{\bar{s}}]}{[\mathbb{E}[p_s] + \mathbb{E}[p_{\bar{s}}]]^2} \frac{\theta}{[f_s(\theta - 1) + 1]^2} \frac{1 - G(Q_s)}{N\omega} \\ &= h_s(Q_s) v_s (1 - v_s) \frac{\mathbb{E}[p_s]}{f_s^2 \theta} \frac{1 - G(Q_s)}{N\omega}. \end{aligned}$$

Let $Q_s^S(Q_{\bar{s}})$ denote the optimal value of Q_s for retailer s as a function of $Q_{\bar{s}}$, and $v_s^S(Q_{\bar{s}})$ and $q_s^S(Q_{\bar{s}})$ denote the corresponding market share and inventory level, respectively.

For comparison, a retailer is called *myopic* if it does not recognize the dependence of its demand distribution on its inventory level, but instead, takes its demand distribution as given and naively follows the traditional newsvendor policy. The first-order condition of the myopic retailer is:

$$\frac{\partial \mathbb{E}[\pi_s(Q_1, Q_2)]}{\partial Q_s} = v_s [(r - s_s) \{1 - G(Q_s)\} - (c_s - s_s)] = 0.$$

This has a unique solution independent of v_s ,

$$Q_s^M = G^{-1} \left(\frac{r - c_s}{r - s_s} \right).$$

with corresponding service level $f_s^M = \mathbb{E}[\min\{Q_s^M, X\}]/(N\omega)$, market share $v_s^M(Q_{\bar{s}})$, and inventory level $q_s^M(Q_{\bar{s}}) = v_s^M(Q_{\bar{s}})Q_s^M$. Here, the superscript M denotes the myopic retailer.

Proposition 3 contrasts myopic and strategic behavior. It shows that for a given service level offered by the competitor, a retailer is better off by taking the strategic decision than by taking the

myopic decision. This result is analogous to that found by Dana and Petruzzi (2001) in a single-period non-competitive newsvendor setting, and can be obtained as a special case of Proposition 2 in their paper.

Proposition 3 *For a given value of $Q_{\bar{s}}$, $Q_s^S > Q_s^M$, $v_s^S > v_s^M$, $q_s^S > q_s^M$ and $E[\pi_s(Q_s^S, Q_{\bar{s}})] > E[\pi_s(Q_s^M, Q_{\bar{s}})]$.*

Proof: Note that, for all Q_s such that $\partial E[\pi_s]/\partial Q_s \geq 0$, we have $dE[\pi_s]/dQ_s > \partial E[\pi_s]/\partial Q_s$. Therefore, $Q_s^S > Q_s^M$. Since $Q_{\bar{s}}$ is fixed, Lemma 1 gives $v_s^S > v_s^M$. Further, applying $q_s = v_s Q_s$, we get $q_s^S > q_s^M$. $E[\pi_s(Q_s^S, Q_{\bar{s}})] > E[\pi_s(Q_s^M, Q_{\bar{s}})]$ follows since Q_s^S is the unique value of Q_s that optimizes $E[\pi_s(Q_s, Q_{\bar{s}})]$. \square

4.2 Competitive Interaction

We now characterize the set of Nash equilibria and discuss their nature as a social outcome. We then describe the effects of bias in consumer learning and cost parameters of the retailers on the set of equilibria.

Consider the profit maximization problem of the strategic retailer. Let \bar{Q}_s be the value of Q_s at which $h_s(Q_s)$ intersects 0, so that $h_s(Q_s) > 0$ for all $Q_s < \bar{Q}_s$, and $h_s(Q_s) \leq 0$ otherwise. Clearly, $Q_s^M < \bar{Q}_s$ and $Q_s^S < \bar{Q}_s$. Further, Proposition 3 shows that $Q_s^S > Q_s^M$ for all $Q_{\bar{s}}$. Therefore, Q_s^S lies in the non-empty interval (Q_s^M, \bar{Q}_s) for all $Q_{\bar{s}}$. Lemma 2 shows that Q_s^S is uniquely defined for given $Q_{\bar{s}}$.

Lemma 2 *The best response function $\pi_s(Q_1, Q_2)$ of the strategic retailer is strictly concave in its inventory level Q_s for all $Q_s \in (Q_s^M, \bar{Q}_s)$.*

Lemma 2 is useful for proving the existence of a Nash equilibrium. Lemma 3 below is useful for proving the uniqueness of the Nash equilibrium.

Lemma 3 $dQ_s^S/dQ_{\bar{s}} > 0$.

We now obtain the following result.

Proposition 4 *There exists a unique pure strategy Nash equilibrium in the inventory game.*

We contrast the results in Propositions 3 and 4 to show the implications of inventory competition between the retailers. According to Proposition 3, strategic behavior dominates myopic behavior for each retailer. However, at the Nash equilibrium, when both retailers behave strategically, we find that this has a dramatic impact on the inventory levels and expected profits of the two retailers compared to the scenario when they behave myopically.

To examine this phenomenon in detail, consider the case when the retailers have equal costs. If both retailers are myopic, Q_1^M is equal to Q_2^M . Thus, each retailer receives the same market share and payoff (in units of expected profit). If both retailers are strategic, Q_1^S is equal to Q_2^S , which also yields the same market share and expected profit for each retailer. However, $E[\pi_s(Q_1^M, Q_2^M)] > E[\pi_s(Q_1^S, Q_2^S)]$ because $v_s^M = v_s^S$ and $h_s(Q_s^M) > h_s(Q_s^S)$. Thus, each retailer's expected profit is higher and inventory level is lower under myopic behavior than under strategic behavior from both retailers. Thus, inventory competition results in a reduction in the total industry profits and an increase in the total industry inventory level at the marketplace compared to the case if both retailers treated their demands as exogenous.

When the retailers have unequal costs, it is no longer true that both retailers are worse off due to competition. Table 1 illustrates a numerical example of the two retailers for different values of θ (0.5 and 2.0) and c_2 (0.2, 0.5, 0.8) with $c_1 = 0.2$.² The table quantifies the % change in the total

²The remaining parameters are as follows: $N = 5000$, $\omega = 0.2$, $r = 1$ and $s_1 = s_2 = 0$. We use these parameters for all numerical results reported in §4. In this analysis, we approximate total market demand by a Poisson distribution with mean $N\omega$. Thus, we relax Assumption 2. We also relax Assumption 3 by considering $\theta < 0.5$, and obtain numerical results consistent with the propositions in the paper.

industry inventory level and the % change in expected profits for each retailer by comparing the scenario when both retailers are myopic with the scenario when both retailers are strategic. We find that inventory competition increases the total industry inventory level regardless of the degree of cost asymmetry or the value of θ . For example, the % increases in inventory levels are 5.36%, 5.13% and 4.94%, respectively, when c_2 is 0.2, 0.5 and 0.8 and θ is 0.5.

Insert Table 1 about here

Table 1 also shows a counter-intuitive result that the higher cost retailer (retailer 2) gains market share and increases its expected profit when both retailers are strategic than when they are myopic. Correspondingly, the lower cost retailer (retailer 1) loses market share and decreases its expected profit due to inventory competition. The intuition for this result is as follows. When both retailers are myopic, the lower cost retailer provides a higher service level (than the higher cost retailer) and consequently acquires a higher market share. However, the higher cost retailer can, by stocking a little more inventory, improve its market share and thereby reduce the competitive advantage of the lower cost retailer. The lower cost retailer cannot effectively counter the higher cost retailer because market share is concave in service level resulting in diminishing returns. Table 1 further details that the % gain to the higher cost retailer from being strategic decreases as θ increases. Thus, the higher cost retailer has a greater incentive to behave strategically when consumers are biased towards negative experiences than when they are biased towards positive experiences.

4.2.1 Effect of Bias in Consumer Learning

The effect of θ on the service levels of the two retailers at equilibrium can be decomposed into the following two components by applying Implicit Function Theorem:

$$\frac{dQ_s}{d\theta} = \underbrace{- \frac{\partial F_s}{\partial \theta} \bigg/ \frac{dF_s}{dQ_s}}_{\text{Effect of Consumer Learning}} - \underbrace{\frac{dF_s}{dQ_s} \frac{dQ_{\bar{s}}}{d\theta} \bigg/ \frac{dF_s}{dQ_s}}_{\text{Effect of Inventory Competition}} .$$

The first term on the right-hand side can be interpreted as the direct effect of learning bias on the service level of retailer s . As θ increases, consumers tend to be more patient about negative experiences, while they react positively and significantly to satisfying shopping trips. Thus, θ could affect service levels in two ways. On the one hand, a higher value of θ implies that a retailer could provide a lower service level to retain market share. Thus, an increase in θ may result in a decrease in Q_s . On the other hand, a higher value of θ implies that a retailer could gain substantially more market share with a small increase in service level. Thus, an increase in θ may result in an increase in Q_s .

The second term on the right-hand side can be interpreted as the indirect effect of inventory competition. In this term, $\partial F_s / \partial Q_s < 0$ from Lemma 2. Further, it can easily be shown that $\partial F_s / \partial Q_{\bar{s}} > 0$. Thus, the indirect effect of θ on Q_s depends on the sign of $dQ_{\bar{s}} / d\theta$. If the service level of the competing retailer increases with θ , the effect on the service level of retailer s is positive, else negative.

We find that both the direct and the indirect effects of an increase in θ lead to a reduction in the service level of each retailer, as shown in Proposition 5.

Proposition 5 *The service levels offered by the two retailers at equilibrium are decreasing in θ .*

Now consider the effect of learning bias on the market shares, inventory levels and expected profits of the two retailers. When costs are equal, it is clear that market shares of the retailers are equal regardless of the value of θ . Thus, by applying Proposition 5, as θ increases, the inventory levels of the retailers decline, and their expected profits increase.

When costs are unequal, on the other hand, the effects of consumer learning on market shares, expected profits and inventory levels at the steady-state are not straightforward. We illustrate these effects through numerical analysis using values of θ between 0.5 and 5.0, and two cases with

different costs for the retailers: in the first case, we assign $c_1 = 0.2$ and $c_2 = 0.5$, and in the second case, $c_1 = 0.2$ and $c_2 = 0.8$. The remaining parameters are as specified earlier in §4.2.

We find that the expected profit of each retailer increases as θ increases. Further, Figure 1 shows the effect of θ on the market share of each retailer. We find that as θ increases, the market share of the lower cost retailer (retailer 1 in both cases) decreases while the market share of the higher cost retailer (retailer 2) increases. Thus, the higher cost retailer suffers more market erosion and the lower cost retailer enjoys greater market power when consumers are negatively biased than when they are positively biased. Viewing this result in conjunction with Table 1, we find that the higher cost retailer has a greater incentive to be strategic when θ is small. However, despite this, the lower cost retailer still enjoys greater power when θ is small. Figure 2 shows the effect of θ on the equilibrium inventory levels of both retailers. We find that the inventory levels of both retailers decline as θ increases.

In summary, both the analysis of the retailers with equal costs and the numerical examples of the retailers with unequal costs show that a negative bias in learning decreases the industry profits and increases the total inventory level. Further, the lower cost retailer has a greater competitive advantage and higher market share when consumers are negatively biased than when they are positively biased.

Insert Figures 1 and 2 about here

4.2.2 Effect of Cost Asymmetry

Since we assume that cost parameters could be different across retailers, we now examine the effect of cost asymmetry on the nature of the equilibrium. The dynamics are as follows. Holding c_2 and the market parameters (e.g., θ) constant, when c_1 increases, we would expect that the service level of retailer 1 should decline. Consider the effect of this decline on retailer 2. On the one hand,

retailer 2 could maintain its market share by reducing its service level as retailer 1 reduces its service level. On the other hand, retailer 2 could gain market share by maintaining or increasing its service level. Therefore, there are competing arguments for both an increase and a decrease in the service level of retailer 2 as a result of the increase in c_1 . Further any change in the service level of retailer 2 will have a feedback effect on the optimal service level of retailer 1 as well.

We obtain the following result regarding the changes in the service levels of the two retailers.

Proposition 6 *The service levels offered by the two retailers at equilibrium are decreasing in c_1 and c_2 .*

Proof: Similar to Proposition 5. Thus, omitted. □

We now illustrate the effects of an increase in c_1 on the market shares, inventory levels and expected profits of both retailers at equilibrium through numerical analysis. The unit cost of retailer 2, c_2 , is kept fixed at 0.2 while the unit cost of retailer 1, c_1 , is varied from 0.1 to 0.8. The remaining parameters are as in §4.2.1. Equilibria are computed for two values of θ , 0.5 and 2.0.

Figures 3 and 4, respectively, show the changes in the market shares and inventory levels of the two retailers at equilibrium as c_1 increases. The market share and inventory level of retailer 1 decline as expected. Interestingly, the inventory level of retailer 2 increases with c_1 . Combining this observation with Proposition 6, we find that, with an increase in c_1 , retailer 2 provides a lower service level but stocks more inventory as it gains market share. Another significant result is that the slopes of the market share and the inventory level with respect to c_1 vary with θ . In particular, Figure 3 shows that the impact of a change in c_1 on market shares is much less when $\theta = 2.0$ than when $\theta = 0.5$. Figure 4 shows that retailer 2 increases its inventory by a smaller amount when $\theta = 2.0$ than when $\theta = 0.5$. Thus, the degree of inventory competition intensifies significantly when consumers are negatively biased, and competitive pressure on inventory levels become more

important at the marketplace. This result is plausible mainly because consumers become more impatient about negative shopping experiences when θ is low.

Insert Figures 3 and 4 about here

5 Conclusions

We have proposed a model of consumer learning that reflects empirical findings established in behavioral research. Under this model, we obtain a closed-form solution for the share of purchases at the various retailers as a function of their service levels and the degree of asymmetry in consumer learning and then analyze inventory decisions for competing retailers. Our results show how asymmetric consumer learning affects the optimal service levels, market shares and expected profits of the retailers. More significantly, when retailers have unequal costs, asymmetric consumer learning affects the degree of competitive advantage enjoyed by the lower cost retailer.

Our paper extends the work of Gans (2002) on consumer learning by capturing the effects of bias in learning. It also extends the work of Dana and Petruzzi (2001) and Bernstein and Federgruen (2003) by demonstrating how consumer learning drives inventory competition. The proposed model can be extended to study several marketing-operations interface issues. For example, our consumer model could be applied to areas other than inventory competition, such as service competition or quality competition. Our model could also be extended to include the effects of price and switching costs on consumer behavior. Finally, our model illustrates that further analytical and empirical research which incorporates consumer behavior in operational models would be useful to find ways of managing the effects of consumer learning.

Our model has some limitations that can be addressed in future research. First, we approximate the steady-state market share of each retailer as shown in equation (6). Due to this assumption,

the variance of demand in our model is larger than the variance of true demand. Therefore, our model might yield more conservative service levels than required by consumer learning. Second, we assume that the coefficient of variation of the demand at each retailer does not change with market share. If, instead, the coefficient of variation of demand were to change as mean demand increases, it would affect the equilibrium service levels and expected profits of the retailers. Third, we do not consider time-varying service levels. Such inventory policies might create situations where one retailer is forced to exit the market, providing new insights.

References

- Agrawal, V., S. Seshadri. 2000. Effect of Risk Aversion on Pricing and Order Quantity Decisions. *Manufacturing & Service Operations Management* **2** 410-423.
- Anderson, E. T., G. J. Fitzsimons, D. I. Simester. 2003. Mitigating the Costs of Stockouts Working Paper, Graduate School of Business, University of Chicago.
- , M. W. Sullivan. 1993. The Antecedents and Consequences of Customer Satisfaction for Firms. *Marketing Science* **12** 125-143.
- Anupindi, R., M. Dada, S. Gupta. 1998. Estimation of Consumer Demand with Stock-out Based Substitution: An Application to Vending Machine Products. *Marketing Science* **17** 406-423.
- Bass, S. 1989. You Can't Sell What You Don't Have. *Catalog Business* (March 1) 24-25.
- Ben-Akiva, M., S. R. Lerman. 1985. *Discrete Choice Analysis*. MIT Press, Cambridge, MA.
- Bernstein, F., A. Federgruen. 2003. A General Equilibrium Model for Industries with Price and Service Competition. *Operations Research* (forthcoming).

- Cachon, G. P., S. Netessine. 2003. Game Theory in Supply Chain Analysis. forthcoming in *Supply Chain Analysis in E-business era*. D. Simchi-Levi, S. D. Wu, M. Shen, Eds. Kluwer Publishers.
- Cooper, L. G., M. Nakanishi. 1988. *Market-Share Analysis: Evaluating Competitive Marketing Effectiveness*. Kluwer Academic Publishers, Boston, MA.
- Dana, J. D. 2001. Competition in Price and Availability When Availability is Unobservable. *RAND Journal of Economics*. **32** 497-513.
- , N. C. Petruzzi. 2001. The Newsvendor Model with Endogenous Demand. *Management Science* **47** 1488-1497.
- Deneckere, R., J. Peck. 1995. Competition Over Price and Service Rate When Demand is Stochastic: A Strategic Analysis. *RAND Journal of Economics* **26** 148-162.
- Debreu, D. 1952. A Social Equilibrium Existence Theorem. *Proceedings of the National Academy of Sciences* **38** 886-893.
- Durrett, R. 1996. *Probability: Theory and Examples*. 2nd Ed. Duxbury Press, Belmont, CA.
- Elman, D. 1989. Stock Service-Level Standard Urged. *Supermarket News* **39** February 6, 1.
- Ernst, R., M. A. Cohen. 1992. Coordination Alternatives in a Manufacturer/Dealer Inventory System Under Stochastic Demand. *Production and Operations Management*. **1** 254-268.
- , S. G. Powell. 1995. Optimal Inventory Policies Under Service-Sensitive Demand. *European Journal of Operational Research* **87** 316-327.
- , —. 1998. Manufacturing Incentives to Improve Retail Service Levels. *European Journal of Operational Research* **104** 437-450.

- Fitzsimons, G. J. 2000. Consumer Response to Stockouts. *Journal of Consumer Research* **27** 249-266.
- Gans, N. 2002. Customer Loyalty and Supplier Quality Competition. *Management Science* **48** 207-221.
- Germain, R., M. B. Cooper. 1990. How a Customer Mission Affects Company Performance. *Industrial Marketing Management* **19** 47-54.
- Guadagni, P. M., J. D. C. Little. 1983. A Logit Model of Brand Choice Calibrated on Scanner Data. *Marketing Science* **2** 203-238.
- Hall, J., E. Porteus. 2000. Customer Service Competition in Capacitated Systems. *Manufacturing & Service Operations Management* **2** 144-165.
- Kahneman, D., A. Tversky. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica* **47** 263-292.
- Karjalainen, R. 1992. The Newsboy Game. Working Paper, The Wharton School, University of Pennsylvania.
- Li, L. 1992. The Role of Inventory in Delivery-Time Competition. *Management Science*. **38** 182-197.
- Lippman, S. A., K. F. McCardle. 1997. The Competitive Newsboy. *Operations Research* **45** 54-65.
- Mahajan, S., G. van Ryzin. 2001. Inventory Competition Under Dynamic Consumer Choice. *Operations Research* **49** 646-657.

- McFadden, D. 1974. Conditional Logit Analysis of Qualitative Choice Behavior. in *Frontiers in Econometrics* P. Zarembka, ed. Academic Press, New York. 105-142.
- Netessine, S., N. Rudi. 2003. Centralized and Competitive Inventory Models with Demand Substitution. *Operations Research* **51** 329-335.
- , ———, Y. Wang. 2003. Dynamic Inventory Competition and Customer Retention. Working Paper, The Wharton School, University of Pennsylvania.
- Parlar, M. 1988. Game Theoretic Analysis of the Substitutable Product Inventory Problem with Random Demands. *Naval Research Logistics* **35** 397-409.
- Petruzzi, N. C., M. Dada. 1999. Pricing and the Newsvendor Problem: A Review with Extensions. *Operations Research* **47** 183-194.
- Schwartz, B. L. 1966. A New Approach to Stockout Penalties. *Management Science* **12** B538-B544.
- . 1970. Optimal Inventory Policies in Perturbed Demand Models. *Management Science* **16** B509-B518.
- Tsay, A., N. Agrawal. 2000. Channel Dynamics Under Price and Service Competition. *Manufacturing & Service Operations Management*. **2** 372-391.
- Tversky, A., D. Kahneman. 1991. Loss Aversion in Riskless Choice: A Reference-Dependent Model. *Quarterly Journal of Economics* **106** 1039-1061.
- van Ryzin, G., S. Mahajan. 1999. On the Relationship between Inventory Costs and Variety Benefits in Retail Assortments. *Management Science* **45** 1496-1509.

Appendix A: Simulation study to evaluate Assumptions 1 and 2

This appendix presents the results of a simulation study conducted to validate Assumptions 1 and 2 in §4. We simulated our consumer behavior model for a wide range of values of the parameters: $f_s, f_{\bar{s}} \in [0.5, 0.99]$ in increments of 0.01, $\theta^u, \theta^d \in [0.2, 0.8]$ in increments of 0.1, $N \in \{1000, 2000, 5000\}$, and $\omega \in [0.1, 0.5]$ in increments of 0.1. With regard to Assumption 1, the simulation results show that $\frac{E[p_s]}{E[p_s] + E[p_{\bar{s}}]}$ and $E\left[\frac{p_s}{p_s + p_{\bar{s}}}\right]$ are not statistically different. The statistical significance was ascertained based on paired- t test statistics (t -value = 0.33 and p -value = 0.7418). On average, we found that the difference is less than 0.0005, which shows that bias may be ignored. This shows that the expression for market share given in Assumption 1, $\frac{E[p_s]}{E[p_s] + E[p_{\bar{s}}]}$, is a valid approximation for $E\left[\frac{p_s}{p_s + p_{\bar{s}}}\right]$.

We then compared the theoretical distribution of demand, obtained from Assumption 2, with the empirical distribution of demand obtained from the simulation of the individual-level consumer behavior model defined in §3 across the values of model parameters defined above. The comparison of distributions is done using a Kolmogorov-Smirnov two-sample test of goodness of fit. The simulation results show that, for any retailer, the difference between the two distributions of demand is not statistically significant (p -value = 0.4360) when the retailer stocks more inventory than its mean demand. This corresponds to the retailer's fill-rate being generally higher than 0.7 for our parameter values. The test shows a poor fit for a retailer only when its service level drops below 0.7 and its competitor's service level increases above 0.98. This is an unlikely scenario. In cases of a poor fit, we find that our proposed approach gives a more conservative assessment of the competition: It results in under-estimating the inventory level (and over-estimating the profit) of the retailer with the lower fill-rate, and over-estimating the inventory level (and under-estimating the profit) of the retailer with the higher fill-rate. Thus, we find that Assumption 2 is valid across a wide range of parameter values and is practically relevant.

Appendix B: Proofs

Proof of Proposition 1: Given any time t , let P_{st}^i denote the probability that customer i will visit retailer s at time t or later and μ_{st}^i denote the expected number of time periods till the next visit to retailer s . We show that given $f_s > 0$, (i) if $p_{st}^i > 0$, then P_{st}^i is equal to 1 and μ_{st}^i is finite; (ii) p_{st}^i does not go to 0 as t increases. These two facts prove the required results.

Step 1: Consider a modified system wherein $p_{\bar{s}t}^i = 1$, i.e., the competitor of retailer s offers a 100% service level. Model this system as a Markov chain with two states representing retailers s and \bar{s} , respectively. In this chain, the probability that consumer i ever visits retailer s at time t or later is given by

$$\hat{P}_{st}^i = \sum_{\tau=t}^{\infty} \left[\left(\frac{1}{1+p_{st}^i} \right)^{\tau-t} \frac{p_{st}^i}{1+p_{st}^i} \right] = 1, \quad (9)$$

and the expected number of time periods till the next visit to retailer s is given by

$$\hat{\mu}_{st}^i = \sum_{\tau=t}^{\infty} \left[(\tau-t) \left(\frac{1}{1+p_{st}^i} \right)^{\tau-t} \frac{p_{st}^i}{1+p_{st}^i} \right] = \frac{1+p_{st}^i}{p_{st}^i}. \quad (10)$$

Here, we used the facts that $p_{st}^i > 0$, and that p_{st}^i is not updated until the next visit of consumer i to retailer s . From (9) and (10), it follows that state s is positive recurrent in the Markov chain, i.e., $\hat{P}_{st}^i = 1$ and $\hat{\mu}_{st}^i < \infty$.

Now consider the case when $p_{\bar{s}t}^i < 1$. The probability of visiting retailer s in this case is always greater than the probability of visiting retailer s in the above Markov chain, i.e.,

$$v_{st}^i = \frac{p_{st}^i}{p_{\bar{s}t}^i + p_{st}^i} > \frac{p_{st}^i}{1+p_{st}^i}. \quad (11)$$

Thus, it can be shown by induction over the number of time periods that $P_{st}^i \geq \hat{P}_{st}^i$ and $\mu_{st}^i \leq \hat{\mu}_{st}^i$. Combining with (9) and (10), we find that $p_{st}^i > 0$ implies that P_{st}^i is equal to 1 and μ_{st}^i is finite.

Step 2: Now consider only the subsequence of time periods $\{t_k\}$ when consumer i visits retailer

s. Let $Y_{st_k}^i$ be 1 if the consumer's visit at time t_k is satisfying, and 0 otherwise. We have

$$p_{s,t_k+1}^i = [(1 - \theta^u)p_{st_k}^i + \theta^u] Y_{st_k}^i + (1 - \theta^d)p_{st_k}^i (1 - Y_{st_k}^i). \quad (12)$$

Define a Markov chain over this subsequence of time periods with two states, $Y_{st_k}^i = 0$ and 1. The transition probabilities of this embedded Markov chain are determined by f_s . Since f_s is strictly positive, $Y_{st_k}^i = 1$ is a positive recurrent state. Therefore, (12) implies that $p_{st_k}^i > 0$ with probability 1 as k tends to ∞ . This further implies that $p_{st}^i > 0$ with probability 1 as t tends to ∞ because p_{st}^i is constant between successive visits to retailer s . On the other hand, if $f_s = 0$, then $p_{s,t_k+1}^i = (1 - \theta^d)p_{st_k}^i$ so that $p_{st_k}^i$ goes to 0 with probability 1 as t tends to ∞ . \square

Proof of Proposition 2: Let $X_t^{i\omega}$ be 1 if customer i visits the retailer at time t along the sample path ω , and 0 otherwise. Let $Y_t^{i\omega}$ be 1 if the visit is satisfying, and 0 otherwise. Let X_t^ω be the vector of store visits across all consumers at time t , and Y_t^ω be the vector of outcomes of store visits across all consumers at time t . A state of the world ω is a sequence of pairs $(X_t^\omega, Y_t^\omega), t = 1, \dots, \infty$.

Let $p_{st}^i(p_{s1}^i, \omega)$ be consumer i 's estimate of the fill-rate at retailer s at the start of period t , as a function of consumer i 's initial estimate of the fill-rate, p_{s1}^i and the sample path ω . Let F_{st}^i denote the distribution function of p_{st}^i . We wish to show that there exists a random variable p_s with distribution function F_s such that $F_{st}^i(x) \rightarrow F_s(x)$ for every x where F_s is continuous. To prove this, we show that $F_{st}^i(x)$ is a Cauchy sequence in $[0, 1]$, i.e., for all $\epsilon > 0$, there exists T such that $|F_{s,t+\tau}^i(x) - F_{st}^i(x)| < \epsilon$ for all $t \geq T$, for all τ .

We only need to consider the subsequence of time periods $\{t_k\}$ when consumer i visits retailer s . From Proposition 1, this subsequence is infinite. Thus, the subscript k is suppressed for convenience. The superscript i is also ignored to simplify the notation. The updating rule (4) gives the stochastic

recursion,

$$\begin{aligned} p_{s,t+1} &= [p_{st}(1 - \theta^u) + \theta^u] Y_{st} + p_{st}(1 - \theta^d)(1 - Y_{st}) \\ &= \theta^u Y_{st} + [(\theta^d - \theta^u)Y_{st} + (1 - \theta^d)] p_{st}. \end{aligned}$$

Expanding this for $p_{s,t+\tau}$, we get

$$p_{s,t+\tau} = p_{s,\tau+1}u(t, \tau) + v(t, \tau), \quad (13)$$

where $p_{s,\tau+1}$ is consumer i 's estimate of the service level at the start of period $\tau + 1$, and

$$\begin{aligned} u(t, \tau) &\equiv \prod_{k=1}^{t-1} [(\theta^d - \theta^u)Y_{s,t+\tau-k} + (1 - \theta^d)], \\ v(t, \tau) &\equiv \sum_{i=1}^{t-1} \prod_{k=1}^{i-1} [(\theta^d - \theta^u)Y_{s,t+\tau-k} + (1 - \theta^d)] \theta^u Y_{s,t+\tau-i}. \end{aligned}$$

Thus,

$$\begin{aligned} F_{s,t+\tau}(x) - F_{st}(x) &= \Pr [p_{s,\tau+1}u(t, \tau) + v(t, \tau) \leq x] - \Pr [p_{s1}u(t, 0) + v(t, 0) \leq x] \\ &\leq \Pr [v(t, \tau) \leq x] - \Pr [p_{s1}u(t, 0) + v(t, 0) \leq x], \end{aligned} \quad (14)$$

where the inequality follows since $p_{s,\tau+1}u(t, \tau) > 0$. Consider the second term on the right hand side of (14). Let $\delta_{max} = \max\{1 - \theta^u, 1 - \theta^d\}$. Since $[(\theta^d - \theta^u)Y_{st} + (1 - \theta^d)]$ is equal to $(1 - \theta^u)$ if $Y_{st} = 1$ and $(1 - \theta^d)$ otherwise, we have that $u(t, 0) \leq \delta_{max}^{t-1}$. Thus,

$$\begin{aligned} F_{s,t+\tau}(x) - F_{st}(x) &\leq \Pr [v(t, \tau) \leq x] - \Pr [p_{s1}\delta_{max}^{t-1} + v(t, 0) \leq x] \\ &\leq \Pr [x - \delta_{max}^{t-1} \leq v(t, 0) \leq x]. \end{aligned}$$

Here, $v(t, \tau)$ can be replaced by $v(t, 0)$ in the second inequality because the retailer maintains constant service level and Y_{st}^i are iid random variables. Since $\Pr [x - \delta_{max}^{t-1} \leq v(t, 0) \leq x] \rightarrow 0$ for all t sufficiently large, we obtain the required result. Thus, there exists a random variable p_s such that $p_{st}^i \xrightarrow{D} p_s$ as $t \rightarrow \infty$. From symmetry, the limiting distribution is identical for all consumers.

The expectation of p_s is now directly obtained from the updating rule (4) or from the expansion (13) since convergence in distribution implies convergence in expectation. \square

Proof of Lemma 1: We have

$$\frac{dv_s}{dQ_s} = \frac{E[p_{\bar{s}}]}{(E[p_s] + E[p_{\bar{s}}])^2} \frac{\theta}{[f_s(\theta - 1) + 1]^2} \frac{1 - G(Q_s)}{N\omega} > 0.$$

Thus, v_s is increasing in Q_s for given $Q_{\bar{s}}$. Further, with some algebraic manipulation, the second derivative of v_s with respect to Q_s can be written as

$$\begin{aligned} \frac{d^2v_s}{dQ_s^2} = & - \frac{E[p_{\bar{s}}]}{(E[p_s] + E[p_{\bar{s}}])^2} \frac{\theta}{[f_s(\theta - 1) + 1]^2} \\ & \times \left[\frac{g(Q_s)}{N\omega} + \left(\frac{1 - G(Q_s)}{N\omega} \right)^2 \frac{2\{\theta - (1 - \theta)E[p_{\bar{s}}]\}}{E[p_{\bar{s}}] + f_s\{\theta - (1 - \theta)E[p_{\bar{s}}]\}} \right]. \end{aligned}$$

All the terms in the above expression are positive, with the exception of $\theta - (1 - \theta)E[p_{\bar{s}}]$, which is negative if $E[p_{\bar{s}}] > \theta/(1 - \theta)$. However, $\theta \geq 0.5$ implies that $\theta/(1 - \theta) \geq 1$. Thus, $\theta - (1 - \theta)E[p_{\bar{s}}]$ is non-negative since $E[p_{\bar{s}}] \leq 1$ by definition. Therefore, v_s is concave in Q_s . \square

Proof of Lemma 2: From the profit function (7), we have

$$\frac{d^2E[\pi_s]}{dQ_s^2} = v_s''h(Q_s) + v_s h''(Q_s) + 2h'(Q_s)v_s'.$$

For $Q_s \in (Q_s^M, \bar{Q}_s)$, we have $h(Q_s) > 0$, $h'(Q_s) < 0$ and $h''(Q_s) < 0$. By Lemma 1, we further have $v_s > 0$, $v_s' > 0$ and $v_s'' < 0$. Therefore, $d^2E[\pi_s]/dQ_s^2 < 0$ for $Q_s \in (Q_s^M, \bar{Q}_s)$. \square

Proof of Lemma 3: Applying the Implicit Function Theorem to the first order condition (8), we get

$$\frac{dQ_s^S}{dQ_{\bar{s}}} = - \frac{dF_s/dQ_{\bar{s}}}{dF_s/dQ_s^S}.$$

It can easily be seen that $dF_s/dQ_{\bar{s}} > 0$. Further, $dF_s/dQ_s^S < 0$ from the concavity of $E[\pi_s]$ for all $Q_s \in [Q_s^M, \bar{Q}_s]$. Thus, $dQ_s^S/dQ_{\bar{s}} > 0$. \square

Proof of Proposition 4: Existence: The strategy spaces of the retailers are non-empty, compact, convex subsets of the real line and each retailer's response function is continuous and strictly concave in the inventory level. Therefore, from Debreu (1952), the result follows.

Uniqueness: We need to show that the reaction curves of the two retailers intersect at most once, so that there is at most one fixed point and the equilibrium is unique. Equivalently, we show that there is at most one point that satisfies the first order conditions of both retailers.

The first order conditions of the two retailers can be rewritten as

$$v_{\bar{s}} = - \frac{h'_s(Q_s)}{h_s(Q_s)\phi(Q_s)}, \quad \text{for } s = 1, 2,$$

where $\phi(Q_s) \equiv E[p_s] \cdot (1 - G(Q_s))/(f_s^2\theta N\omega)$. Suppose that the solution to these simultaneous equations is not unique, and there exist two distinct equilibria, (Q_1, Q_2) and (Q'_1, Q'_2) . Assume, without loss of generality, that $Q'_1 > Q_1$. This implies that $Q'_2 > Q_2$ since $dQ_2/dQ_1 > 0$ by Lemma 3.

Note that $h'_s(Q_s)$ is negative and decreasing in Q_s , and $h_s(Q_s)$ and $\phi(Q_s)$ are both positive and decreasing in Q_s . Thus, $-h'_s(Q_s)/[h_s(Q_s)\phi(Q_s)]$ is positive and increasing in Q_s . Therefore, we have

$$- \frac{h'_s(Q'_s)}{h_s(Q'_s)\phi(Q'_s)} > - \frac{h'_s(Q_s)}{h_s(Q_s)\phi(Q_s)}, \quad \text{for } s = 1, 2.$$

Adding the inequalities for $s = 1$ and 2 gives

$$v_1(Q'_1, Q'_2) + v_2(Q'_1, Q'_2) > v_1(Q_1, Q_2) + v_2(Q_1, Q_2).$$

But this is an impossibility since $v_1(Q_1, Q_2) + v_2(Q_1, Q_2) = 1$ for all (Q_1, Q_2) . Therefore, it must be that $(Q_1, Q_2) = (Q'_1, Q'_2)$ and there is at most one Nash equilibrium. \square

Proof of Proposition 5: Q_1 and Q_2 are implicit functions of θ defined by the first order conditions of the two retailers given in (8). The derivative of (8) with respect to θ gives

$$\frac{\partial F_s}{\partial \theta} + \frac{dF_s}{dQ_s} \frac{dQ_s}{d\theta} + \frac{dF_s}{dQ_{\bar{s}}} \frac{dQ_{\bar{s}}}{d\theta} = 0, \quad \text{for } s = 1, 2.$$

By solving these simultaneous equations, the derivative of Q_s with respect to θ is obtained as

$$\frac{dQ_s}{d\theta} = - \frac{\frac{dF_{\bar{s}}}{dQ_{\bar{s}}} \frac{\partial F_s}{\partial \theta} - \frac{dF_s}{dQ_{\bar{s}}} \frac{\partial F_{\bar{s}}}{\partial \theta}}{\frac{dF_1}{dQ_1} \frac{dF_2}{dQ_2} - \frac{dF_1}{dQ_2} \frac{dF_2}{dQ_1}}. \quad (15)$$

Recall that

$$\frac{dF_s}{dQ_s} < 0, \quad \text{and} \quad \frac{dF_s}{dQ_{\bar{s}}} > 0. \quad (16)$$

Additional inequalities are established by the following Lemmas:

Lemma 4

$$\left. \frac{\partial F_s}{\partial \theta} \right|_{F_s=0} < 0.$$

Proof: Differentiating condition (8) with respect to θ and simplifying, we get

$$\begin{aligned} \left. \frac{\partial F_s}{\partial \theta} \right|_{F_s=0} &= h_s(Q_s) \frac{E[p_s]v_{\bar{s}}}{f_s \theta^2 (E[p_s] + E[p_{\bar{s}}])[f_s(\theta - 1) + 1]} \frac{1 - G(Q_s)}{N\omega} \\ &\quad \times \left[E[p_s] \left(1 - \theta - \frac{1}{f_s} \right) + E[p_{\bar{s}}] \left(-1 - \theta + \frac{1}{f_{\bar{s}}} \right) \right]. \end{aligned} \quad (17)$$

Here,

$$\begin{aligned} &E[p_s] \left(1 - \theta - \frac{1}{f_s} \right) + E[p_{\bar{s}}] \left(-1 - \theta + \frac{1}{f_{\bar{s}}} \right) \\ &= \frac{\theta f_s}{f_s(\theta - 1) + 1} \frac{f_s - 1 - f_s \theta}{f_s} + \frac{\theta f_{\bar{s}}}{f_{\bar{s}}(\theta - 1) + 1} \frac{1 - f_{\bar{s}} - f_{\bar{s}} \theta}{f_{\bar{s}}} \\ &= -\theta + \frac{\theta(1 - f_{\bar{s}} - f_{\bar{s}} \theta)}{f_{\bar{s}}(\theta - 1) + 1} \\ &= \frac{-2f_{\bar{s}} \theta^2}{f_{\bar{s}}(\theta - 1) + 1} \\ &< 0. \end{aligned}$$

Since all other terms in (17) are non-negative, the result follows. \square

Lemma 5

$$\frac{dF_1}{dQ_1} \frac{dF_2}{dQ_2} - \frac{dF_1}{dQ_2} \frac{dF_2}{dQ_1} > 0. \quad (18)$$

Proof: From (8), note that

$$\begin{aligned} \frac{dF_s}{dQ_s} &= h_s''(Q_s)v_s + 2h_s'(Q_s)v_s' + h_s(Q_s)v_s'', \\ \frac{dF_s}{dQ_{\bar{s}}} &= h_s'(Q_s) \frac{dv_s}{dQ_{\bar{s}}} + h_s(Q_s) \frac{d^2v_s}{dQ_s dQ_{\bar{s}}}. \end{aligned}$$

Simplifying (18) using the fact that $dv_{\bar{s}}/dQ_s = -v_s'$, we get

$$\begin{aligned} \frac{dF_1}{dQ_1} \frac{dF_2}{dQ_2} - \frac{dF_1}{dQ_2} \frac{dF_2}{dQ_1} &= \left(h_1''(Q_1)v_1 \frac{dF_2}{dQ_2} + 2h_1'(Q_1)h_2''(Q_2)v_1'v_2 + h_1(Q_1)h_2''(Q_2)v_1''v_2 \right) \\ &+ (4h_1'(Q_1)h_2'(Q_2)v_1'v_2' - h_1'(Q_1)h_2'(Q_2)v_1'v_2') \\ &+ \left(2h_1(Q_1)h_2'(Q_2)v_1''v_2' + h_1(Q_1)h_2'(Q_2)v_1' \frac{d^2v_1}{dQ_1 dQ_2} \right) \\ &+ \left(2h_1'(Q_1)h_2(Q_2)v_1'v_2'' + h_1'(Q_1)h_2(Q_2)v_2' \frac{d^2v_2}{dQ_1 dQ_2} \right) \\ &+ \left(h_1(Q_1)h_2(Q_2)v_1''v_2'' - h_1(Q_1)h_2(Q_2) \frac{d^2v_1}{dQ_1 dQ_2} \frac{d^2v_2}{dQ_1 dQ_2} \right). \end{aligned}$$

Denote the terms in the five sets of brackets as A, B, C, D, E , respectively. A is positive because $h_s(Q_s)$ and v_s are positive and concave, $h_s'(Q_s) < 0$, $v_s' > 0$ and $dF_s/dQ_s > 0$. B is positive since $h_s'(Q_s) < 0$ and $v_s' > 0$.

The following additional facts are useful to analyze C, D and E .

$$\begin{aligned} v_s' &= v_s v_{\bar{s}} \phi_s \\ v_s'' &= (v_{\bar{s}} - v_s)v_s' \phi_s + v_s v_{\bar{s}} \phi_s' \\ \frac{d^2v_s}{dQ_s dQ_{\bar{s}}} &= (v_s - v_{\bar{s}})v_s' \phi_s = (v_s - v_{\bar{s}})v_s' \phi_{\bar{s}}, \end{aligned}$$

where $\phi_s = E[p_s] \cdot (1 - G(Q_s))/(f_s^2 \theta N \omega)$ is positive and decreasing in Q_s . Thus, C gives

$$\left(2v_1''v_2' + v_1' \frac{d^2v_1}{dQ_1 dQ_2} \right) h_1(Q_1)h_2'(Q_2) = \left(v_1''v_2' + v_1 v_2 v_2' \frac{d\phi_1}{dQ_1} \right) h_1(Q_1)h_2'(Q_2) > 0.$$

D is analogous to C . Thus, it can be shown that $D > 0$. E gives

$$\begin{aligned}
 v_1'' v_2'' - \frac{d^2 v_1}{dQ_1 dQ_2} \frac{d^2 v_2}{dQ_1 dQ_2} &= v_1'' v_2'' - \phi_1 \phi_2 v_1' v_2' (v_1 - v_2)(v_2 - v_1) \\
 &= v_1'' v_2'' + \phi_1 \phi_2 v_1' v_2' (v_1 - v_2)^2 \\
 &> 0.
 \end{aligned}$$

This proves the required inequality. □

Applying (16) and Lemmas 4 and 5 to (15), it follows that $dQ_s/d\theta < 0$ for $s = 1, 2$. □

	% Change in Total Industry Inventory Level $\left[\frac{q_1^S + q_2^S}{q_1^M + q_2^M} - 1 \right]$			% Change in Expected Profit of Each Retailer $\left[\frac{\pi_1^S}{\pi_1^M} - 1, \frac{\pi_2^S}{\pi_2^M} - 1 \right]$		
	$c_2 = 0.2$	$c_2 = 0.5$	$c_2 = 0.8$	$c_2 = 0.2$	$c_2 = 0.5$	$c_2 = 0.8$
$\theta = 0.5$	5.36	5.13	4.94	[-1.17,-1.17]	[-6.27,1.60]	[-12.04,6.26]
$\theta = 2.0$	2.39	2.12	1.85	[-0.15,-0.15]	[-0.46,0.04]	[-1.08,0.54]

Table 1: Comparison of Strategic Behavior with Myopic Behavior of Retailers ($c_1 = 0.2$)




Figure 1: Effect of Learning Bias on Market Share




Figure 2: Effect of Learning Bias on Inventory Level




Figure 3: Effect of Cost Asymmetry on Market Share ($c_2 = 0.2$)




Figure 4: Effect of Cost Asymmetry on Inventory Level ($c_2 = 0.2$)