MANAGEMENT SCIENCE
Vol. 54, No. 10, October 2008, pp. 1805–1821
ISSN 0025-1909 | EISSN 1526-5501 | 08 | 5410 | 1805

DOI 10.1287/mnsc.1080.0889

Inventory Management Under Market Size Dynamics

Tava Lennon Olsen
Olin Business School, Washington University in St. Louis, St. Louis, Missouri 63130,
olsen@wustl.edu

Rodney P. Parker
The University of Chicago Graduate School of Business, Chicago, Illinois 60637,
rodney.parker@chicagogsb.edu

We investigate the situation where a customer experiencing an inventory stockout at a retailer potentially leaves the firm's market. In classical inventory theory, a unit stockout penalty cost has been used as a surrogate to mimic the economic effect of such a departure; in this study, we explicitly represent this aspect of consumer behavior, incorporating the diminishing effect of the consumers leaving the market upon the stochastic demand distribution in a time-dynamic context. The initial model considers a single firm. We allow for consumer forgiveness where customers may flow back to the committed purchasing market from a nonpurchasing "latent" market. The per-period decisions include a marketing mix to attract latent and new consumers to the committed market and the setting of inventory levels. We establish conditions under which the firm optimally operates a base-stock inventory policy. The subsequent two models consider a duopoly where the potential market for a firm is now the committed market of the other firm; each firm decides its own inventory level. In the first model, the only decisions are the stocking decisions and in the second model, a firm may also advertise to attract dissatisfied customers from its competitor's market. In both cases, we establish conditions for a base-stock equilibrium policy. We demonstrate comparative statics in all models.

Key words: inventory; competition; Markov games; marketing and operations interface
History: Accepted by Paul H. Zipkin, operations and supply chain management; received November 9, 2006. This paper was with the authors 7 months for 2 revisions. Published online in *Articles in Advance* August 8, 2008.

1. Introduction

The treatment of consumers in classical inventory theory has typically been quite naïve. Whereas the aggregate consumer demand is often assumed to be uncertain, albeit with a known demand distribution, any further aspects of consumer behavior tend to be limited to assuming that unsatisfied customers will backlog, be lost, or become a mixture of these. However, a common consumer reaction to a stockout is to change retailers during a subsequent shopping excursion (e.g., Fitzsimons 2000). As stockout frequencies can be quite high in practice (see the consumer behavior discussion and citations in the online appendix (provided in the e-companion)[1] for research on typical values), the incorporation of the consumers' activities subsequent to experiencing a stockout is important.

Most commonly in the inventory literature, a unit stockout penalty cost is assessed to the firm for each customer whose demand is not satisfied from on-hand inventory immediately. This penalty cost has numerous interpretations (e.g., expedited delivery, premiums at alternative retailers, a more costly substitute, etc.) but commonly it is intended to represent the economic effects of a customer's lost goodwill. As Heyman and Sobel (1984) note, "[I]t is difficult to estimate such penalty costs, but usually, it is even harder to model explicitly the dependence of the demand process on the degree to which demands do not exceed stock levels." It is our objective to, indeed, explicitly model the diminution of demand caused by stockouts.

Schwartz (1966) appears to be the first research article to address the issue of future demands being affected by current poor inventory performance; it restricts attention to a deterministic demand rate. In Schwartz (1970), the model is extended to incorporate some uncertainty of the mean demand rate in continuous time, and some recognition (although not modeled) is given to the possibility of consumer forgiveness, a concept we formalize in our models. Liberopoulos and Tsikis (2006) extend this line of analysis to quantify the unit backorder cost in this economic order quantity context.

[1] An electronic companion to this paper is available as part of the online version that can be found at http://mansci.journal.informs.org/.

Philosophically, it is far more satisfying to explicitly capture the actual phenomenon of interest rather than rely on a proxy. However, one question we seek to address is, "How good is such a proxy?" In this paper, we exclude any unit stockout penalty cost and instead permit some customers to backlog, some to have lost demand in that period, and the remainder to leave the market altogether, thus creating a shrinkage in the demand distribution for the following period while maintaining the usual aspects of inventory models (stochastic demand, periodic review, unit holding costs, transition of physical inventory between periods). We focus on proving the optimality (or equilibrium existence) of base-stock policies under a model with market-size dependent demand. To be precise, we show that in each period there is an order-up-to level that is optimal (or an equilibrium) and independent of inventory (but not market size), and that starting inventory in each period will always lie below that level (so long as the initial inventory in period 1 is not too high). By characterizing sufficient conditions for optimality of such policies, we have characterized sufficient conditions for the existence of a proxy stockout cost in the analogous traditional inventory setting.

We initially consider a single enterprise concurrently making inventory decisions and marketing mix decisions. Two markets are specified in the model. The first is labeled the "committed" market, which consists of consumers who purchase regularly (demand is assumed to be affine in this market size). The second is a "latent" market, consisting of consumers who may have previously shopped with the firm and may do so again in the future, or customers who are aware of the firm but have yet to shop there. We permit a portion of the unsatisfied consumers (i.e., those experiencing a stockout) to be lost demand in that time period only, a portion to backlog into the following period, and a portion to leave the market entirely (i.e., flow from the committed to the latent market). The two marketing mix decisions that the firm makes are an incentive to persuade latent customers to become committed again at some cost, and an advertising decision to attract altogether new customers to join the committed market. Operating under this regime with some demand and parameter conditions, we discover the firm should operate under a base-stock inventory policy. In addition, we find we can determine a value for each committed and each latent customer.

Fergani's (1976) (unpublished) Ph.D. thesis is probably the most comprehensive attempt to capture the effects of inventory stockouts on future demand for a single firm. Fergani (1976) discusses three primary models. The first is a finite horizon Markov decision process model with a fraction of dissatisfied customers leaving the market. We independently analyzed this model but do not include our analysis here for the sake of brevity. Robinson (1990) considers an infinite horizon version of this model with a quite general demand function, establishing tractable upper and lower bounds on the optimal inventory policy. Fergani's (1976) second model incorporates an advertising mechanism to boost market size; the structure of advertising used is simple (with linear per unit costs) and is unrepresentative of current advertising literature. Fergani's (1976) third model assumes that market size is unknown at the beginning of every period but that a prior distribution is updated in a Bayesian fashion in every period. In comparison, we incorporate more general advertising functions, consumer forgiveness, and consumer incentives; these are all elements missing in Fergani's (1976) models.

Like the majority of the traditional inventory literature (see, e.g., Porteus 2002, Zipkin 2000), we do not allow the firm to set price. There are two equally compatible narratives for this setting. The first is simply that inventory decisions are made by a separate set of decision makers on a more frequent timeline than pricing decisions. The second is that the firm is a monopolist without setting a retail price (e.g., in a regulated environment) or that the consumers are not responsive to changes in price.[2]

After consideration of the single firm model, we then shift our attention to a competitive model, specifically a duopoly where dissatisfied customers leave the committed market of one firm and join the committed market of the other firm. That is, the "potential market" (in lieu of the latent plus external markets considered in the single-firm model) of one firm is the committed market of the other. In this Markov game (see, e.g., Başar and Olsder 1999, Heyman and Sobel 1984, Fudenberg and Tirole 1991, Parker and Kapuściński 2008), the state space at the beginning of each period consists of both firms' initial inventory levels and committed market sizes. Initially, we isolate the firms' decisions to stocking levels only, allowing for partial backlogging, lost sales, and customer defection as in the single-firm model. Under similar demand and parameter assumptions to the single-firm model, we show there is a base-stock equilibrium policy for each firm.

This duopoly model is a direct extension of the works of Hall and Porteus (2000) and Liu et al. (2007) to a nonperishable inventory setting. Those works study systems where duopolists compete by

[2] The potential shortcoming of this interpretation is that the firm should therefore choose an extremely high price; however, eventually nearly all customers will react against an extraordinarily high price. Assuming the firm adopts but does not set price (i.e., a price taker) reconciles this interpretation.

installing capacity in every period but where any service failures result in market diminution. Liu et al. (2007) extend Hall and Porteus' (2000) analysis with a more general demand function that we have also adopted in our models. The multiperiod nature of those models, with market reductions in a competitive framework, makes their work similar to ours. Their models describe a service system especially well, particularly where capacity may be changed at short notice. They also carry over to production settings where the "capacity" is now intended to represent *perishable* inventory in a newsvendor setting. A similar setting with perishable inventory is also considered in Henig and Gerchak (2003).

The work of Hall and Porteus (2000), Liu et al. (2007), and Henig and Gerchak (2003) all operate under the assumption that physical inventory or consumer backlogs are not carried between periods whereas we track physical inventory, backlogs, lost sales, and market defections. We are aware of no other dynamic game literature other than those described above that deals with the relationship between market sizes and stockouts, and we believe our paper to be the first in this setting to allow inventory and backlogs to carry over from period to period.

Finally, we consider a model where firms can actively try to attract dissatisfied customers from the other firm. We again show existence of a base-stock equilibrium policy. Therefore, our work is also related to inventory duopoly models that consider how customers are treated after experiencing a stockout. In particular, Parlar (1988), Lippman and McCardle (1997), and Netessine and Rudi (2003) have some fixed proportion of disappointed customers' demand transferring to the other retailer and the loyal but disappointed customers' demand considered lost; Avsar and Baykal-Gürsoy (2002) present the same treatment in the infinite horizon. Ahn and Olsen (2007), in a subscription model context, extend Lippman and McCardle's (1997) work to multiple periods. Netessine et al. (2006) present several (independent) treatments of customers' backlogging and transfer behavior in a dynamic environment. Olsen and Parker (2008) generalize and integrate these treatments in a single model with backlogging, lost sales, and transfers.

We allow the transfer of some portion of unsatisfied customers between markets but do not permit consumer searches within the same time period, an approach partly validated by Fitzsimons (2000) who finds that consumers having experienced a stockout are substantially more likely to visit an alternative retailer during a *subsequent* shopping outing, although we acknowledge that consumer search within the same period could certainly occur, too. It should be noted that we treat the consumer behavior of switching between markets as a black box and do not delve into the psychological elements underpinning these decisions (see Fitzsimons 2000 for an illustrative study of consumer choice, conflict, and behavioral responses to stockouts). Indeed, we assume that customer behavior is governed by Markov (memoryless) transition functions. Mahajan and van Ryzin (1999) summarize consumer choice models where the retailer can directly or indirectly control consumer substitution, the latter of which encompasses our approach.

We are aware of a few papers that address the issuing of incentives after customers experience a stockout. In particular, Netessine et al. (2006) examine a firm's incentive to persuade its own customers to remain loyal (i.e., backlog locally) rather than switch retailers after experiencing a stockout; DeCroix and Arreola-Risa (1998) consider a similar incentive for a monopolist. Anderson et al. (2006) find that the cost of such incentives to backlog do not tend to offset the increased revenues of these consumers; they recommend a targeted discounting strategy. Another paper dealing with the competitive aspects of customer defection is Gans (2002), where customers experience the quality of a service or product supplied by a firm and update a prior belief about that firm's quality in a Bayesian manner. Likewise, Gaur and Park (2007) incorporate consumer learning and retailer service levels in ascertaining competitive inventory policy.

We address the idea of offering incentives to customers after they experience inventory disappointment; however, we focus on firms attempting to draw customers from elsewhere to their markets. In the single-firm model, the firm persuades "latent" customers to become "committed" customers whereas in the duopoly model, each firm tempts dissatisfied customers from its competitor's market. Further, unlike our paper, none of the papers just noted have an underlying market from which demand is drawn. The consumer behavior literature contains interesting and relevant work that provides further empirical motivation for our work. In the interest of space, our survey of this literature may be found in the online appendix.

As outlined above, the overarching goal of our research is to investigate how realistic a model can be with respect to consumer behavior and still retain optimality (or equilibrium existence) of base-stock policies (and, hence, for the single-firm model, prove the existence of a proxy lost sales cost). Within this objective our contribution is fourfold. First, we explicitly model a range of consumer decisions in the face of stockouts. Second, we provide a much more detailed single-firm model than previously studied. In particular, we allow general advertising and explicitly capture consumer forgiveness through a latent market. This model is provided in §2. Third, we provide what we believe to be the first dynamic duopoly model that addresses the relationship between market size and stockouts while allowing inventory and backlog carryover from period to period. Finally, we extend our duopoly to allow firms to actively try to attract dissatisfied customers, an extension missing from the few works that do consider the relationship between market size and stockouts (none of which carry inventory between periods). Both duopoly models are given in §3. Concluding remarks appear in §4 and the appendix and online appendix contain all proofs.

2. The Single-Firm Model

In this section, we introduce and analyze a "single-firm" periodic review model. The firm begins every time period t knowing the current inventory level (x_t), the size of its committed market (θ_t), and the size of its latent market (β_t). The committed market consists of regular purchasers and the latent market is made up of former customers who left the committed market due to experiencing an inventory service failure. In reality, these markets must be estimated and will not be known exactly; here, for ease of exposition and analytic tractability, we assume they are indeed known. Internet retail providers are most likely to have a good estimate of these markets, although with frequent purchaser cards and modern data mining techniques, it appears likely that firms will become increasingly able to provide such estimates. It is also probable, in our estimation, that a Bayesian model similar to that considered in Fergani (1976) may be able to be layered on our model; we have not pursued such an extension and leave it as a potential subject of future research.

Let $D_t(\theta)$ be the uncertain demand in period t arising from a committed market of size θ. The firm is assumed to know the distribution of $D_t(\theta)$. When the period is clear, we will drop the subscript t for notational convenience. We make the following assumption:

Assumption 1. *Demand in period t is distributed as $D_t(\theta_t) = p_1\theta_t + (p_2\theta_t + p_3)\varepsilon_t$, where $p_1, p_2, p_3 \geq 0$ and ε_t is a mean zero random variable. The random variables $\{\varepsilon_t\}$ are independent and identically distributed (i.i.d.) and are drawn from a continuous distribution with a support that is a closed subset of $[-p_1/p_2, \infty)$, having cumulative distribution function (c.d.f.) $\Phi(\cdot)$ and density $\phi(\cdot)$.*

This demand form is analogous to that presented in Liu et al. (2007), where the reader is referred for further explanation and justification of this form. It contains additive and multiplicative demands as special cases.[3] Assumption 1 does not restrict the form of the distribution for demand; it does, however, imply that both the mean and standard deviation of demand are affine in market size and that the coefficient of variation of demand is nonincreasing in market size. This assumption will be seen later to add significant tractability, leading to a greater number of insights than would otherwise likely be possible.

The firm makes the following decisions simultaneously in each period: (1) an inventory stocking decision (y); (2) a marketing decision to persuade latent customers to return to the committed market (ρ); and (3) an advertising decision to increase the size of the committed market (ν). The flow decision ρ is the expected *proportion* of the latent market that is diverted to the committed market. The external advertising decision ν is the expected total flow of customers from outside both markets in response to advertising. We allow all these variables to be continuous. Thus, we are in effect assuming that market size (and, hence, demand) and market flows are large enough such that continuous flow-based approximations suffice. Figure 1(a) depicts the flows between the pools.

Suppose that control ρ is applied in period t. We assume the proportion of customers who switch from the latent pool to the committed pool is $R_t(\rho)$, where $\mathrm{E}[R_t(\rho)] = \rho$. The manipulation of ρ is presumed to be at a cost per latent customer of $C(\rho)$. Thus, if there are β customers in the pool of latent customers and a control of ρ is applied, then $R_t(\rho)\beta$ customers will choose to return to the committed pool and the total cost will be $C(\rho)\beta$. We have deliberately kept the form of this switching control general so that it may reflect coupons, targeted advertising, or some other marketing mechanism. The function $C(\rho)$ is assumed to be strictly convex, nonnegative, and increasing in ρ.

Likewise, we suggest that advertising externally can attract new customers to the committed pool from outside. A (similarly general) advertising cost of $K(\nu)$ will attract $U_t(\nu)$ customers to the committed pool, where $\mathrm{E}[U_t(\nu)] = \nu$. We assume there is a finite point $\bar{\nu}$ after which $K(\nu)$ is strictly convex, which will preclude infinite advertising in a period being optimal. Clearly, S-shaped advertising functions are a subcase of these assumptions. We do not model any interaction between $R_t(\rho)$ and $U_t(\nu)$, assuming that ν is indeed only targeted at external customers. We make the following assumption:

Assumption 2. *The sequences of random variables $\{R_t(\rho)\}$ and $\{U_t(\nu)\}$ are i.i.d., mutually independent, independent of all other random variables in the system, and have means $\rho = \mathrm{E}[R_t(\rho)]$ and $\nu = \mathrm{E}[U_t(\nu)]$, respectively.*

In each period t, let Γ_t be the random proportion of customers experiencing a stockout who choose not to backlog and Λ_t be the random proportion of

[3] We are grateful to a reviewer for suggesting we adopt this demand form. Our original form was multiplicative only ($p_1 = p_3 = 0$, $p_2 = 1$). In that case, a continuous distribution is not necessary in the finite horizon model.

Figure 1 **Schematic Figures of the Single-Firm and Duopoly Models with Notation**

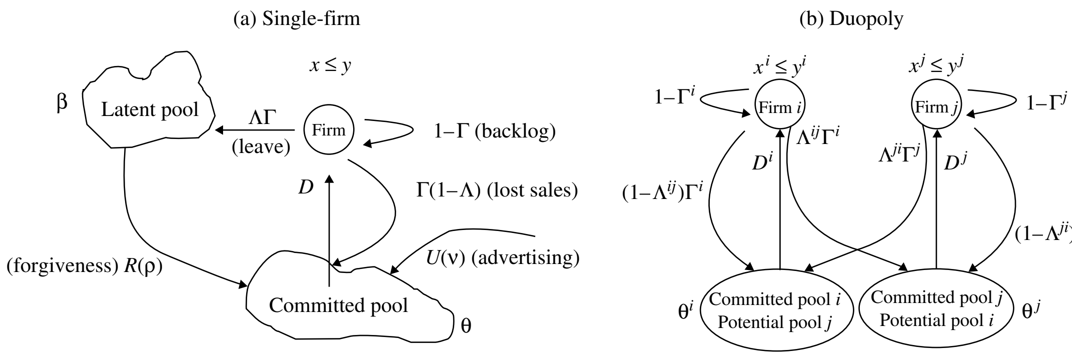

nonbacklogging unsatisfied customers who choose to leave the committed market. This formulation was chosen to reflect a greater desire for the firm's product by the customers who backlog. However, the parameters governing the division of the unsatisfied customers into backlogging, immediate lost sales, or market defection routes is arbitrary so long as market losses occur only if inventory is depleted. The routings can, in fact, be arranged in any manner (e.g., only backlogging customers defect or some random combination of backlogging and nonbacklogging customers defect). Further, some of these proportions may be zero so that total backlogging and total lost sales are both subcases of our model. We make the following assumption:

Assumption 3. *The sequences of random variables* $\{\Gamma_t\}$ *and* $\{\Lambda_t\}$ *are i.i.d., mutually independent, independent of all other random variables in the system, and have means* $\gamma = \mathrm{E}[\Gamma_t]$ *and* $\lambda = \mathrm{E}[\Lambda_t]$*, respectively.*

Assume controls ρ_t and ν_t are applied in period t. Then, the state transition functions are as follows:

$$x_{t+1} = (y_t - D_t(\theta_t))^+ - (1-\Gamma_t)(D_t(\theta_t) - y_t)^+$$
$$= y_t - D_t(\theta_t) + \Gamma_t(D_t(\theta_t) - y_t)^+, \quad (1)$$

$$\theta_{t+1} = \theta_t - \Lambda_t\Gamma_t(D_t(\theta_t) - y_t)^+ + R_t(\rho_t)\beta_t + U_t(\nu_t), \quad (2)$$

$$\beta_{t+1} = (1 - R_t(\rho_t))\beta_t + \Lambda_t\Gamma_t(D_t(\theta_t) - y_t)^+, \quad (3)$$

where we define $x^+ = x$ if $x \geq 0$ and $x^+ = 0$, otherwise. For future reference, we similarly define $x^- = -x$ if $x \leq 0$ and $x^- = 0$, otherwise. Equation (1) simply transfers any leftover physical inventory into the following period and likewise the backlogging proportion (at rate $1-\Gamma_t$) of the unsatisfied demand. Equation (2) states that the new committed market size consists of the old committed market size less outflow to the latent market due to stockouts plus inflow from the latent market due to forgiveness or incentives plus inflows due to external advertising. Equation (3) states the new latent market size is the old latent market size plus inflow from the committed market less outflow back to the committed market.

Assume $r > 0$ is the retail price and $h > 0$ is the per-unit holding cost in each period. We will assume a discount factor of α, $0 < \alpha < 1$. The objective is to maximize total discounted expected reward over either the finite or infinite horizon (this will be shown to be well-defined in the infinite horizon). In the finite horizon, assume there are T periods. Consider the firm's expected periodic profits in any period t, when controls (y_t, ρ_t, ν_t) are applied, $1 \leq t \leq T$:

$$r\mathrm{E}[\min(y_t, D_t(\theta_t)) + x_t^-] - h\mathrm{E}[(y_t - D_t(\theta_t))^+]$$
$$- C(\rho_t)\beta_t - K(\nu_t) \quad (4)$$
$$= -r\mathrm{E}[(D_t(\theta_t) - y_t)^+] - h\mathrm{E}[(y_t - D_t(\theta_t))^+]$$
$$+ r\mathrm{E}[D_t(\theta_t)] - C(\rho_t)\beta_t - K(\nu_t) + rx_t^-. \quad (5)$$

The revenue term in (4) consists of the sum of the backlog and the lesser of demand and available inventory. Clearly, the final term of (5) can be "rolled back" into period $t-1$ using (1) and discounting at rate α, thus producing a per-period reward of

$$-\tilde{r}\mathrm{E}[(D(\theta_t) - y_t)^+] - h\mathrm{E}[(y_t - D(\theta_t))^+] + r\mathrm{E}[D(\theta_t)]$$
$$- C(\rho_t)\beta_t - K(\nu_t),$$

where $\tilde{r} = r(1 - \alpha(1-\gamma))$.[4]

We will assume throughout that in the finite horizon model, the terminal value has been normalized by rx_{T+1}^-. In other words, if $\widetilde{V}_{T+1}(x, \theta, \beta)$ is the actual terminal value function, then we will use a terminal value of $V_{T+1}(x, \theta, \beta) = \widetilde{V}_{T+1}(x, \theta, \beta) - rx^-$. Thus, all assumptions on $V_{T+1}(x, \theta, \beta)$ should be translated into assumptions on the actual terminal value function $\widetilde{V}_{T+1}(x, \theta, \beta)$ by adding rx^- to $V_{T+1}(x, \theta, \beta)$.

If demand is nonrandom (i.e., $\varepsilon = 0$), an affine demand function (in committed market size) can be

[4] Note that no acquisition cost has been included in the model but could be easily incorporated along the lines suggested by Veinott (1966). For example, if the acquisition cost is $w(y_t - x_t) = wy_t - wx_t$, the first term is absorbed easily. The second term is accommodated using Equation (1), which will ultimately result in a modified revenue term, $\tilde{r} = r(1 - \alpha(1-\gamma)) - \alpha w\gamma$, which is positive when $r > w$.

seen to be necessary for concavity of the one-period reward function (and is likely also necessary for most forms of stochastic demand). As our focus is on the optimality of base-stock policies, we restrict attention to concave revenue functions. This is the primary reasoning behind the affine demand function in Assumption 1. However, this assumption also leads to a change of variable that significantly aids the model tractability. Note that linear demand assumptions were made (for similar reasons) in Fergani (1976), Hall and Porteus (2000), and Henig and Gerchak (2003), and the affine form used here was also used (for similar reasons) in Liu et al. (2007).

Define the functions

$$y^{-1}(x,\theta)=\Phi((x-p_1\theta)/(p_2\theta+p_3))$$

and

$$y_f(z,\theta)=p_1\theta+\Phi^{-1}(z)(p_2\theta+p_3),$$

where the notation Φ^{-1} denotes the inverse of the cumulative distribution function Φ. We perform a change of variable, letting $z_t=y^{-1}(y_t,\theta_t)$ so that $y_t=y_f(z_t,\theta_t)$. Then,

$$y_t-D_t(\theta_t)=(p_2\theta_t+p_3)(\Phi^{-1}(z_t)-\varepsilon_t)$$

and z_t is the chosen critical fractile for satisfied demand. The transition functions may, therefore, be rewritten as follows:

$$x_{t+1}=(p_2\theta_t+p_3)(\Phi^{-1}(z_t)-\varepsilon_t+\Gamma_t(\varepsilon_t-\Phi^{-1}(z_t))^+), \quad (6)$$

$$\theta_{t+1}=\theta_t-(p_2\theta_t+p_3)\Lambda_t\Gamma_t(\varepsilon_t-\Phi^{-1}(z_t))^+ + R_t(\rho_t)\beta_t+U_t(\nu_t), \quad (7)$$

$$\beta_{t+1}=(1-R_t(\rho_t))\beta_t+(p_2\theta_t+p_3)\Lambda_t\Gamma_t(\varepsilon_t-\Phi^{-1}(z_t))^+. \quad (8)$$

Further, the expected periodic reward for any period t is

$$\begin{aligned}L(z_t,\rho_t,\nu_t,\theta_t,\beta_t)&\\ &\stackrel{\Delta}{=}(p_2\theta_t+p_3)(-\tilde{r}\mathrm{E}[(\varepsilon_t-\Phi^{-1}(z_t))^+]-h\mathrm{E}[(\Phi^{-1}(z_t)-\varepsilon_t)^+])\\ &\quad +rp_1\theta_t-C(\rho_t)\beta_t-K(\nu_t)\\ &=(p_2\theta_t+p_3)\tilde{L}(z_t)+rp_1\theta_t-C(\rho_t)\beta_t-K(\nu_t),\end{aligned}$$

where

$$\tilde{L}(z)\stackrel{\Delta}{=}-\tilde{r}\mathrm{E}[(\varepsilon_t-\Phi^{-1}(z))^+]-h\mathrm{E}[(\Phi^{-1}(z)-\varepsilon_t)^+].$$

Note that both the reward and transition functions are affine in θ_t and β_t.

We define

$$S(z)\stackrel{\Delta}{=}\mathrm{E}[\Lambda_t\Gamma_t(\varepsilon_t-\Phi^{-1}(z))^+] \quad (9)$$

as the expected proportion of lost customers when inventory is stocked to critical fractile z and

$$z^*_{my}\stackrel{\Delta}{=}\arg\max_z\tilde{L}(z), \quad (10)$$

where z^*_{my} is the optimal scaled myopic inventory quantity. If controls (z_t,ρ_t,ν_t) are applied in period t, then

$$\mathrm{E}[\theta_{t+1}]=\theta_t-(p_2\theta_t+p_3)S(z_t)+\rho_t\beta_t+\nu_t, \quad (11)$$

$$\mathrm{E}[\beta_{t+1}]=(1-\rho_t)\beta_t+(p_2\theta_t+p_3)S(z_t). \quad (12)$$

Note that $-S(z)$ and $\tilde{L}(z)$ are both concave in z with $-S(z)$ also being nondecreasing in z.

Our model contains limited memory; it is assumed that customers in either market are averaged across their tenures in the market. This is equivalent to assuming that consumers are memoryless about their previous experiences, good or bad, in the markets. This is partially justified by Anderson et al. (2006) who find no difference in response to a stockout between customers who have purchased previously and novice customers.

Another possible extension to our model is to apply a multiplicative stochastic shock to the latent market to reflect the chance that some latent customers will leave this market (either through moving away or through forgetting about the retailer) or to reflect other nonpurchasing customers becoming newly aware of the retailer (e.g., moving to the region from elsewhere). All the analysis is preserved (with some additional technical conditions on the average size of the shock) but little additional insight is gained with its inclusion. A similar shock cannot be applied to the committed market without destroying the analytical structure of the model.

2.1. Finite Horizon Results

Using the (normalized) terminal value function $V_{T+1}(x,\theta,\beta)$, recursively define the optimal profit-to-go or value function in period t, $1\le t\le T$, as

$$V_t(x,\theta,\beta)=\max_{\substack{z\ge y^{-1}(x,\theta)\\ 0\le\rho\le 1,\nu\ge 0}}(L(z,\rho,\nu,\theta,\beta)+\alpha\mathrm{E}[V_{t+1}(x_{t+1},\theta_{t+1},\beta_{t+1})]).$$

We seek to characterize the structure (with respect to (x,θ,β)) of the optimal decision variables z_t^*, ρ_t^*, and ν_t^*, which achieve the maximum in the above equation.

The affine nature of both the one-period revenue and transition functions, coupled with appropriate assumptions on the terminal value, will allow us to write $V_t(\cdot)$ as an affine function of (θ,β) and independent of x (so long as x is below the desired base-stock level). Above the desired base-stock level, $V_t(\cdot)$ will be bounded above by its value at the desired base-stock level. This will allow a simple characterization of the optimal decision variables as well as intuition into the components of the value function. Our inductive argument will rely on proving that starting inventory in the following period is below the desired base-stock level. We will also assume that initial inventory in period 1 is below the desired level; however, this (relatively mild) condition is for convenience only and the online appendix provides an extension to the proof of Theorem 1 where this assumption is relaxed. As it will be shown that the desired critical fractile is at least the myopic level, we state the assumption on period 1 inventory as follows:

Assumption 4. *Assume that initial inventory* $x_1 \leq y_f(z^*_{my}, \theta_1)$.

If the terminal value of committed customers is low, then the optimal decision will likely be to save on inventory costs and ramp down market size near the end of the horizon. As future market size is stochastic, this would likely imply that optimal policies depend (possibly in a nonsmooth fashion) on both the current market size and the number of periods to go. Similarly, if the terminal value of committed customers is quite high, then similar effects will likely occur with a growing market. Assumption 5 places the terminal value between "too high" and "too low." Of course, the effect of any assumption on terminal value becomes increasingly diminished as one moves further from the end of the horizon. We make the following assumption:

Assumption 5. *Assume, for any* x, $V_{T+1}(x, \theta, \beta) = a_{T+1}\theta + b_{T+1}\beta + c_{T+1}$, *where*

$$a_{T+1} = \frac{p_2 \tilde{L}(z^*_{my}) + rp_1}{(1-\alpha)^2}, \tag{13}$$

$$b_{T+1} = \alpha a_{T+1}, \quad \textit{and} \quad c_{T+1} = 0. \tag{14}$$

Salvage value functions are frequently used to (i) overcome undesirable and unrepresentative behavior at the end of a time horizon, (ii) endow a model with analytical tractability, or (iii) reflect economic reality. We use salvage value functions for reasons (i) and (ii) and note that such an assumption will not be needed in the infinite horizon model.

The final assumption in this section is a technical one that ensures that the future expected value of a committed customer is at least that of a latent customer. As it seems likely that the firm would prefer to keep customers in the committed market rather than lose them to the latent market, it is likely that the conditions needed to guarantee this condition are also reasonable:

Assumption 6.

$$1 - \rho^*_T - p_2 \lambda \gamma E[(\varepsilon_t - \Phi^{-1}(z^*_{my}))^+] \geq 0. \tag{15}$$

Assumptions analogous to Assumption 6 were made in both Hall and Porteus (2000) and Liu et al. (2007). Liu et al. (2007) contains further discussion and justification for their analogous assumption, referring to it as "very mild" (see their Condition 3 and the discussion surrounding it). We are now ready to present the main result of this section:

Theorem 1. *Recursively define*

$$m_t = \max_{0 \leq z \leq 1} (\tilde{L}(z) - \alpha S(z)(a_{t+1} - b_{t+1})), \tag{16}$$

$$a_t = p_2 m_t + \alpha a_{t+1} + rp_1, \tag{17}$$

$$b_t = \max_{0 \leq \rho \leq 1} (-C(\rho) + \alpha\rho(a_{t+1} - b_{t+1})) + \alpha b_{t+1}, \tag{18}$$

$$c_t = \max_{\nu \geq 0} (\alpha a_{t+1}\nu - K(\nu)) + p_3 m_t + \alpha c_{t+1}. \tag{19}$$

Then, under Assumptions 1–6,

$$z^*_t = \arg\max_{0 \leq z \leq 1} (\tilde{L}(z) - \alpha S(z)(a_{t+1} - b_{t+1}))$$

$$\rho^*_t = \arg\max_{0 \leq \rho \leq 1} (-C(\rho) + \alpha\rho(a_{t+1} - b_{t+1}))$$

$$\nu^*_t = \arg\max_{\nu \geq 0} (\alpha a_{t+1}\nu - K(\nu))$$

and for $x \leq y_f(z^*_t, \theta)$

$$V_t(x, \theta, \beta) = a_t\theta + b_t\beta + c_t.$$

For $x > y_f(z^*_t, \theta)$, $V_t(x, \theta, \beta)$ *is bounded above by* $V_t(y_f(z^*_t, \theta), \theta, \beta)$. *Further,* z^*_t, ρ^*_t, a_t, b_t, *and* $a_t - b_t$ *are nondecreasing in* t *with* $a_t - b_t \geq 0$ *and* $z^*_t \geq z^*_{my}$, *for all* t, $1 \leq t \leq T$.

The proof may be found in the appendix and follows by an inductive argument. Key to the main inductive step is to show that the desired base-stock level is both an affine function of the committed market size and guaranteed to be no less than the inventory at the start of the period. In some cases, inventory may be equal to the desired level; in these cases, no order will be placed (analogously to traditional inventory models with no fixed costs). The argument for why future inventory will be below the desired base-stock level proceeds as follows: If the committed market grows, then the affine nature of inventory in the market size will imply that the following period's inventory cannot be "too high." However, if the market shrinks, then there must have been dissatisfied customers; hence, inventory has been depleted and by definition again cannot be "too high."

This same basic argument is used (and formalized) in the proofs of Theorems 1–4 in this paper. In these theorems, we make the assumption (Assumption 4 for Theorem 1) that initial inventory in period 1 is below the desired level (although this assumption is relaxed for Theorem 1 in the online appendix). Hence, by the above argument, initial inventory will be below the desired level in all future periods. In this way, we avoid needing to prove that no order is placed if initial inventory is above the desired level, although we hypothesize that this is the case (and have formally proven it to be the case for Theorem 1).

Theorem 1 yields several observations. First, the overall value of the firm can be separated into elements of the value of the committed market, the value of the latent market, and any remaining value. An appealing interpretation is that the variables a_t and b_t are the *per customer* values in each of these markets. So, a_t is the discounted expected value of a current committed customer in period t, accounting for all possible expected movements over the remainder of the time horizon. This gives the firm some real intuition of how its inventory policies and customer responses to them affect the value of those customers to the firm in a tangible outcome: sales. Notice also that the base-stock level in period t, y_t^*, equals $p_1\theta_t + \Phi^{-1}(z_t^*)(p_2\theta_t + p_3)$ and, therefore, is affine in committed market size θ_t.

The difference, $a_t - b_t$, is the incremental benefit of having a committed rather than a latent customer. This increment is shown to be positive (using Assumption 6), which is natural because only a committed customer can purchase from the firm and the best a latent customer can do is to forgive and begin buying in the future. We also show that this increment is nondecreasing as the end of the horizon approaches, which we would argue is also natural as there are fewer and fewer opportunities for latent customers to become committed and committed customers have a sufficient salvage value.

The second observation is that the optimal inventory policy (operating under standing assumptions) is base-stock. The efficacy of this unadorned policy is well-known; it is a natural and appealing policy for implementation. The immediate conclusion is that the unit stockout cost used in classical theory as a surrogate for market shrinkage due to lost future demand and customer goodwill can indeed be a valid proxy. By explicitly modeling this market shrinkage rather than using the unit cost, we also arrive at the same structural optimal policy. This can be true under numerous modeling "accessories" (e.g., consumer forgiveness, advertising, coupons) or under minimal assumptions (as in the corollary below). The counterpoint to this statement is that the base-stock inventory policy may not be optimal under all circumstances. Thus, while the unit stockout cost can continue to be used in the future to approximate lost future demand, it should be used with some caution, noting whether the conditions appear justified.

The optimal level of advertising to the latent pool ρ_t^* depends on the future per-customer value difference $a_{t+1} - b_{t+1}$ but not on the size of either the committed or latent markets (although total advertising to the latent pool is proportional to the latent market size). Similarly, the optimal amount of external advertising ν_t^* depends on the future value of a committed customer a_{t+1} but not on the market size. This lack of dependence in market size is due to our affine model structure, which does not reflect economies of scale. Recall that an affine demand structure was necessary to prove concavity of the one-period profit function, so a model with economies of scale would need an entirely new method of analysis, which is beyond the scope of this paper.

We recognize that Fergani (1976) offers a streamlined inventory model where future demand is affected by current stockouts and the demand may adopt a linear or affine form in the market size. His model does not include the latent market at all and, thus, we exclude those parameters in the following corollary. This implies there will only be an outflow of (dissatisfied) customers from the committed market and no resulting inflow (from the latent market or from external advertising[5]), thus this model is appropriate for a finite time horizon only. On the other hand, this streamlined model is burdened by few restrictive assumptions.

Corollary (Fergani 1976). *Setting $\rho = \nu = 0$, the system optimally operates under a base-stock inventory policy.*

2.2. Infinite Horizon Results

In the infinite horizon, if the market is either shrinking or growing, then there is transience toward zero or infinity, respectively. We therefore assume no external advertising is possible. Define the functions

$$z_f(\Delta) = \arg\max_{0 \le z \le 1}(\tilde{L}(z) - \alpha\Delta S(z)) \tag{20}$$

and

$$\rho_f(\Delta) = \arg\max_{0 \le \rho \le 1}(-C(\rho) + \alpha\Delta\rho). \tag{21}$$

In what follows, Δ will equal the value difference between a committed and a latent customer and $z_f(\Delta)$ and $\rho_f(\Delta)$ will be the optimal controls given this value difference. Note that $z_{my}^* = z_f(0)$. From the concavity of $-S(z)$ and $\tilde{L}(z)$, for $\Delta \ge 0$,

$$z_f(\Delta) = 1 - \frac{h}{\tilde{r} + \alpha\lambda\gamma\Delta + h}.$$

[5] In an extension, Fergani (1976) considers external advertising to "replenish" the market.

Thus, $z_f(\Delta)$ is increasing in Δ. Further, $\rho_f(0)=0$ and, for $\Delta>0$,

$$\rho_f(\Delta)=\min(C'^{-1}(\alpha\Delta),1).$$

Note that $\rho_f(\Delta)$ is nondecreasing in Δ. The solution to this equation is unique because $C(\cdot)$ is increasing and strictly convex.

Define

$$\Delta_{\max}=\frac{p_2\tilde{L}(z^*_{my})+rp_1}{1-\alpha},$$
$$z_{\max}=1-\frac{h}{\tilde{r}+\alpha\Delta_{\max}\gamma+h},$$
$$\rho_{\max}=\min(C'^{-1}(\alpha\Delta_{\max}),1).$$

These variables will be shown in the following lemma to indeed be upper bounds on their respective modifiers under the following assumption:

Assumption 7. $1-p_2S(z^*_{my})-\rho_{\max}\geq 0$.

This is a flow assumption, similar in nature to (15), to guarantee nonnegativity of the value difference Δ. Because $h>0$, this assumption implies that $\rho_{\max}<1$ and, hence, $\rho_{\max}=C'^{-1}(\alpha\Delta_{\max})$. The following fixed point lemma will aid in the infinite horizon proof.

Lemma 1. *Define the mapping*

$$T(\Delta)=p_2\tilde{L}(z_f(\Delta))-\alpha\Delta p_2S(z_f(\Delta))+rp_1+\alpha\Delta+C(\rho_f(\Delta))-\alpha\Delta\rho_f(\Delta).$$

Under Assumption 7, there is a unique fixed point Δ^* such that

$$\Delta^*=T(\Delta^*),\tag{22}$$

where $\Delta^*\in[0,\Delta_{\max}]$. Further, for any $\Delta\in[0,\Delta_{\max}]$,

$$z^*_{my}\leq z_f(\Delta)\leq z_{\max}\quad\text{and}\quad 0\leq\rho_f(\Delta)\leq\rho_{\max}.$$

The proof of Lemma 1 (found in the online appendix) follows from basic calculus and showing that $T(\cdot)$ is a contraction mapping. Let Π be the set of admissible policies. Define

$$V^*(x,\theta,\beta)=\sup_{\pi\in\Pi}\sum_{t=1}^{\infty}\alpha^{t-1}L(z_t,\rho_t,\theta_t,\beta_t),$$

where we redefine $L(z,\rho,\theta,\beta)\stackrel{\Delta}{=}L(z,\rho,\cdot,\theta,\beta)$. Then, $V^*(x,\theta,\beta)$ is the optimal discounted expected revenue function for the infinite horizon problem with initial state equal to (x,θ,β). We have the following result:

Theorem 2. *Assume Assumptions* 1–3 *and* 7. *Define*

$$a=(p_2(\tilde{L}(z_f(\Delta^*))-\alpha\Delta^*S(z_f(\Delta^*)))+rp_1)/(1-\alpha),\tag{23}$$
$$b=(-C(\rho_f(\Delta^*))+\alpha\Delta^*\rho_f(\Delta^*))/(1-\alpha),\tag{24}$$
$$c=p_3(\tilde{L}(z_f(\Delta^*))-\alpha\Delta^*S(z_f(\Delta^*)))/(1-\alpha),\tag{25}$$

where Δ^* is from Equation (22). *Then, $\Delta^*=a-b$ and a and b simultaneously solve*

$$a=p_2\max_{z\geq 0}(\tilde{L}(z)-\alpha S(z)(a-b))+rp_1+\alpha a,\tag{26}$$
$$b=\max_{0\leq\rho\leq 1}(-C(\rho)+\alpha\rho(a-b))+\alpha b.\tag{27}$$

Further, for $x\leq y_f(z_f(\Delta^*),\theta)$,

$$V^*(x,\theta,\beta)=a\theta+b\beta+c,$$

and $z_f(\Delta^*)$ and $\rho_f(\Delta^*)$ are an optimal stationary policy.

The proof (found in the appendix) follows by showing that $a\theta+b\beta+c$ satisfies the Bellman equation for V^*. We can offer some comparative statics for these optimality results:

Proposition 1. *The optimal "value increment" Δ^* increases in r (if $p_1\geq p_2$), $-h$, $-\gamma$, $-\lambda$, and α. The optimal stocking level $z_f(\Delta^*)$ increases with r, Δ^*, and $-h$. When $\lambda\gamma\Delta^*\geq r(1-\gamma)$, the optimal stocking level increases in α. The optimal incentive $\rho_f(\Delta^*)$ increases with increasing Δ^*, r, $-h$, $-\gamma$, $-\lambda$, and α.*

This result states that the value increment of a committed customer over a latent customer increases with the retail price and the discount factor. The former is obvious as committed customers will pay more when their demand is realized and satisfied. The latter arises because it lessens the effect of a defecting customer who has a chance of returning to the committed pool in the following period. The value increment will also increase (1) when the holding cost decreases because it lessens the cost of servicing a committed customer, and (2) when either γ or λ decreases because these govern the proportion of dissatisfied customers who leave the committed market. The optimal stocking level increases with an increase in the retail price or a decrease in the holding cost, because these changes indicate that a greater level of inventory service is economically warranted. The optimal stocking level increases with the optimal value increment because this represents preserving committed customers over losing them. The optimal incentive level increases by the same reasoning. In other words, a greater value increment, higher price, lower holding cost, and a smaller proportion of leaving customers are all reasons for the firm to spend more to convert a latent customer to a committed one.

The proof of these comparative statics follows from a standard application of the implicit function theorem and may be found in the online appendix. The observations following Theorem 1 for the finite horizon model carry over to the infinite horizon. Moreover, a specific value for the "equivalent" unit stockout cost in a traditional inventory model is found to be

$$\text{unit stockout cost}=\alpha\lambda\gamma(a-b)=\alpha\lambda\gamma\Delta^*.$$

The interpretation of this unit stockout cost is that it represents the value lost due to a stockout. It is discounted by α because the leaving customers will join the latent market in the following period. The parameters $\lambda\gamma$ represent the expected proportion of stocked-out customers who will leave, and $(a-b)$ is the expected lost (lifetime) value of the customers leaving the committed market for the latent market in the following period.

3. The Duopoly Model

In this section, we provide a competitive framework where two firms explicitly compete with each other for the retention of customers on the basis of their inventory performance. The "committed" market for one firm is now the "potential" market of the other firm. That is, when customers stock out at firm 1, they may join firm 2's market and vice versa. There is no external (outside the duopoly) advertising. Thus, the potential market has replaced both the latent and external markets of the previous section. This duopoly arrangement may be seen in Figure 1(b).

We first prove results for a duopoly where each firm makes a stocking decision only. We then extend the basic model to a model with incentive decisions as well as inventory stocking decisions. We define some commonalities shared by the two models before examining each separately. Though we only consider duopolies, the results would likely extend to oligolopies as well. However, one would need to carefully define and justify the flows of dissatisfied customers between firms. For the model with stocking decisions only, the separability that occurs would make this relatively straightforward; however, for the model where the flows depend on explicit decisions, much more care would be needed.

Much of the nomenclature is identical or analogous to the single-firm model, with the difference being a superscript identifying the firm. We will not redefine such notation if we believe its definition to be self-explanatory.

We have four state variables $(x_t^1, x_t^2, \theta_t^1, \theta_t^2)$, where x_t^i is firm i's inventory (or backlog) level at the beginning of period t and θ_t^i is the size of firm i's committed customer pool. We reserve indices i and j throughout to denote the two firms, where the use of both implies $j \neq i$. Let us first define the transition functions for each firm i:

$$x_{t+1}^i = y_t^i - D^i(\theta_t^i) + \Gamma_t^i (D_t^i(\theta_t^i) - y_t^i)^+, \tag{28}$$

$$\begin{aligned}\theta_{t+1}^i = \theta_t^i &- \Lambda_t^{ij}\Gamma_t^i(D_t^i(\theta_t^i) - y_t^i)^+ \\ &+ \Lambda_t^{ji}\Gamma_t^j(D_t^j(\theta_t^j) - y_t^j)^+,\end{aligned} \tag{29}$$

where Λ_t^{ij} is the proportion of unsatisfied firm i customers that defect to firm j in the following period.

Throughout this section, we assume the following.

Assumption 8. *For period t, $i=1,2$, $D_t^i(\theta_t^i) = p_1^i\theta_t^i + (p_2^i\theta_t^i + p_3^i)\varepsilon_t^i$, where $p_1^i, p_2^i, p_3^i \geq 0$. The sequences $\{\varepsilon_t^i\}$ and $\{\varepsilon_t^j\}$ are i.i.d., independent of each other, and follow the same distributional assumptions as in Assumption* 1.

Assumption 9. *The sequences of random variables $\{\Gamma_t^i\}$, $\{\Gamma_t^j\}$, $\{\Lambda_t^{ij}\}$, and $\{\Lambda_t^{ji}\}$ are i.i.d., mutually independent, independent of all other random variables in the system, and have means $\gamma^i = \mathrm{E}[\Gamma_t^i]$, $\gamma^j = \mathrm{E}[\Gamma_t^j]$, $\lambda^{ij} = \mathrm{E}[\Lambda_t^{ij}]$, and $\lambda^{ji} = \mathrm{E}[\Lambda_t^{ji}]$, respectively.*

Perform a change of variable so that

$$z_t^i = y_f^{-1}(y_t^i, \theta^i) = \Phi_i\left(\frac{y_t^i - p_1^i\theta^i}{p_2^i\theta^i + p_3^i}\right),$$

where Φ_i is the c.d.f. of ε^i (we use a subscript for i on Φ to aid notationally when its inverse is taken). Then, the transition functions can be rewritten as follows:

$$x_{t+1}^i = (p_2^i\theta_t^i + p_3^i)(\Phi_i^{-1}(z_t^i) - \varepsilon_t^i + \Gamma_t^i(\varepsilon_t^i - \Phi_i^{-1}(z_t^i))^+), \tag{30}$$

$$\begin{aligned}\theta_{t+1}^i = \theta_t^i &- (p_2^i\theta_t^i + p_3^i)\Lambda_t^{ij}\Gamma_t^i(\varepsilon_t^i - \Phi_i^{-1}(z_t^i))^+ \\ &+ (p_2^j\theta_t^j + p_3^j)\Lambda_t^{ji}\Gamma_t^j(\varepsilon_t^j - \Phi_j^{-1}(z_t^j))^+.\end{aligned} \tag{31}$$

For the model with no incentives, the periodic reward for period t for firm i is

$$\begin{aligned}&L^i(y_t^i, \theta_t^i) \\ &= -\tilde{r}^i\mathrm{E}[(D^i(\theta_t^i) - y_t^i)^+] - h^i\mathrm{E}[(y_t^i - D^i(\theta_t^i))^+] + r^i\mathrm{E}[D^i(\theta_t^i)] \\ &= (p_2^i\theta_t^i + p_3^i)\tilde{L}^i(z_t^i) + r^ip_1^i\theta_t^i,\end{aligned}$$

where

$$\tilde{L}^i(z_t^i) \stackrel{\Delta}{=} -\tilde{r}^i\mathrm{E}[(\varepsilon_t^i - \Phi_i^{-1}(z_t^i))^+] - h^i\mathrm{E}[(\Phi_i^{-1}(z_t^i) - \varepsilon_t^i)^+].$$

Note that both the reward and transition functions are affine in θ_t^i and θ_t^j. The model with incentives will have the additional advertising costs associated with attracting the other firm's customers in the periodic reward function; these will be written separately from $\tilde{L}^i(\cdot)$, which is defined as above in both models.

3.1. Duopoly with No Consumer Incentives

In this subsection, we assume each firm chooses its inventory levels, mindful of the potential of losing its own customers but with no conscious effort to attract customers from the other firm. This could be translated as the inward-looking "operations focused" model.

As in the single-firm model, we assume that there is a (normalized) salvage value associated with the end of horizon state vector $(x^i, x^j, \theta^i, \theta^j)$, as follows:

Assumption 10. *For any* x^i, x^j, $V^i_{T+1}(x^i, x^j, \theta^i, \theta^j) = a^i_{T+1}\theta^i + b^i_{T+1}\theta^j$, *where*

$$a^i_{T+1} = \frac{p^i_2 \tilde{L}^i(z^{i*}_{my}) + r^i p^i_1}{(1-\alpha)^2}, \tag{32}$$

$$b^i_{T+1} = \alpha a^i_{T+1}. \tag{33}$$

The intuition for this assumption is analogous to the single-firm model.

We define $V^i_t(x^i, x^j, \theta^i, \theta^j)$ to be the discounted expected value for firm i under a Markov equilibrium (if it exists) from period t onward, given a current state vector of $(x^i, x^j, \theta^i, \theta^j)$. Whereas this value will depend on the specific equilibrium chosen, we show that there is, in fact, a unique Markov equilibrium in each period and, hence, there is no ambiguity in the expression. Further, we assume that:

Assumption 11. *Assume*

$$1 - p^i_2 \lambda^{ij} \gamma^i \mathrm{E}[(\varepsilon^i - \Phi_i^{-1}(z^{i*}_{my}))^+] - p^j_2 \lambda^{ji} \gamma^j \mathrm{E}[(\varepsilon^j - \Phi_j^{-1}(z^{j*}_{my}))^+] \geq 0.$$

This condition, which ensures that we would prefer to keep customers rather than lose them to the competitor (see the proof of Theorem 3), has strong analogies to Assumption 6. It is also effectively the same as Condition 3 in Liu et al. (2007). Indeed, the assumptions of this section are effectively equivalent to those of Liu et al. (2007) *except* that we allow inventory (or backlogs) to be carried between periods (which is the significant contribution of the section), which in turn necessitates an assumption on the salvage value. As stated before, such an assumption becomes decreasingly important as one moves further from the end of the horizon and is not needed in the infinite horizon.

Theorem 3. *Recursively define*

$$z^{i*}_t = \underset{0 \leq z^i \leq 1}{\arg\max}\{\tilde{L}^i(z^i) - \alpha(a^i_{t+1} - b^i_{t+1}) S^i(z^i)\}, \tag{34}$$

$$a^i_t = p^i_2 \tilde{L}^i(z^{i*}_t) - \alpha p^i_2 (a^i_{t+1} - b^i_{t+1}) S^i(z^{i*}_t) + r^i p^i_1 + \alpha a^i_{t+1}, \tag{35}$$

$$b^i_t = \alpha p^j_2 (a^i_{t+1} - b^i_{t+1}) S^j(z^{j*}_t) + \alpha b^i_{t+1}, \tag{36}$$

$$c^i_t = p^i_3 \tilde{L}^i(z^{i*}_t) - \alpha p^i_3 (a^i_{t+1} - b^i_{t+1}) S^i(z^{i*}_t) + \alpha p^j_3 (a^i_{t+1} - b^i_{t+1}) S^j(z^{j*}_t) + \alpha c^i_{t+1}. \tag{37}$$

Then, under Assumptions 8–11 *for* $x^i_1 \leq y^i_f(z^{i*}_1, \theta^i)$ *and* $x^j_1 \leq y^j_f(z^{j*}_1, \theta^j)$, *the unique Markov perfect equilibrium policy is for the firms to order-up-to* $(y^i_f(z^{i*}_t, \theta^i), y^j_f(z^{j*}_t, \theta^j))$ *and this policy has value* $V^i_t(x^i, x^j, \theta^i, \theta^j) = a^i_t \theta^i + b^i_t \theta^j + c^i_t$. *Further,* z^{i*}_t, a^i_t, b^i_t, *and* $a^i_t - b^i_t$ *are nondecreasing in* t *with* $a^i_t - b^i_t \geq 0$.

Thus, so long as the inventory in the first period (only) is below the desired levels, there is an equilibrium in base-stock policies. As in the single-firm model, it is likely that one does not actually need to restrict first-period inventory, but the proof would become more involved because a bounding argument on $V^i_t(\cdot)$ is no longer sufficient.

The value function $V^i_t(\cdot)$ represents firm i's expected present value of the current and future rewards under the (unique) pure strategy Markov equilibrium given the current state. As is well-known, a Markov equilibrium is a subgame perfect Nash equilibrium in a finite horizon. In this particular model, due to Assumptions 8–11, we gain additive separability of each firm's value function into components dependent upon the market size state variables and independent of the beginning inventory state variables. We speculate the primary reason for the separability is that defecting customers do not search at the other retailer in the same period but join the competitor's market and may be served in the following period at the soonest. This assumption and resultant separability is also seen in Hall and Porteus (2000) and Liu et al. (2007).

Given the above separability, the Markov game effectively becomes two parallel Markov decision process models, where each firm can choose its inventory independent of the other firm's choices. As such, the solution to the infinite horizon model is well-defined and stationary. Further, one could use machinery similar to that of Theorem 2 to find the (unique) infinite horizon stationary values, where the stationary flow from the competitor $p^j_2 S^j(z^{j*})$ would replace the stationary flow decision $\rho_f(\Delta^*)$. We do not do so here in the interest of space. In this case, we are able to observe an "equivalent" unit stockout cost similar to the single-firm model. For firm i, it is

$$\text{unit stockout cost} = \alpha\lambda^{ij}\gamma^i(a^i - b^i) = \alpha\lambda^{ij}\gamma^i\Delta^{i*}.$$

3.2. Duopoly with Incentives for Dissatisfied Consumers

We now suppose that the firm may attract *dissatisfied* customers from the competition. That is, firm i may influence the mean (λ^{ji}) of the flow from firm j to firm i. We assume that there is a convex increasing advertising cost for firm i to attract an expected proportion λ^{ji} of firm j's dissatisfied customers. We intend this advertising effort to be directed toward all the customers of the competitor, but only the dissatisfied customers will be significantly affected by the message. Further, we assume that this cost may be written as $(p^j_2\theta^j_t + p^j_3)A_i(\lambda^{ji})$. As such, it is assumed to contain a term that is proportional to θ^j and a further term that is independent of θ^j, where the ratio between these terms is fixed. If $p^j_3 = 0$ (i.e., multiplicative demand), then this assumption simply implies that the advertising cost must be proportional to the competitor's market size.

For notational convenience, we will suppress explicit dependence of Λ_t^{ij} on the control λ^{ij}, but such dependence should be understood in the following. Further, the distributional assumptions of Assumption 9 continue to hold, where the Λ_t^{ij} are identically distributed conditional on having the same control λ^{ij} applied. The periodic reward will be as before with this additional incentive cost, as follows:

$$(p_2^i\theta_t^i + p_3^i)\tilde{L}^i(z_t^i) + r^i p_1^i \theta_t^i - \theta_t^j(p_2^j\theta_t^i + p_3^j)A^i(\lambda^{ji}).$$

Define

$$\tilde{S}^i(z) = \gamma^i \mathrm{E}[(\varepsilon^i - \Phi_i^{-1}(z))^+]. \tag{38}$$

Then,

$$\mathrm{E}[\theta_{t+1}^i] = \theta_t^i(1 - p_2^i\lambda_t^{ij}\tilde{S}^i(z^i)) + \theta_t^j p_2^j \lambda_t^{ji}\tilde{S}^j(z^j). \tag{39}$$

As in the single-firm model, Δ^i will represent the value difference between a committed (firm i) and potential (firm j) customer. We define the vector $\mathbf{\Delta} = (\Delta^1, \Delta^2)$. Further, define

$$z_f^i(\mathbf{\Delta}) = 1 - \frac{h^i}{\tilde{r}^i + \alpha\Delta^i\lambda_f^{ij}(\mathbf{\Delta})\gamma^i + h^i}, \tag{40}$$

$$\lambda_f^{ij}(\mathbf{\Delta}) = \min(A_j'^{-1}(\alpha\Delta^j\tilde{S}^i(z_f^i(\mathbf{\Delta}))), 1). \tag{41}$$

Analogous the single-firm model, $z_f^i(\mathbf{\Delta})$ and $\lambda_f^{ij}(\mathbf{\Delta})$ will represent equilibrium responses given customer value differences of $\mathbf{\Delta}$. Note that in contrast to the duopoly with inventory decisions only, here the decisions of the two firms truly represent an equilibrium decision. As such, we need to show that these responses are well-defined. This is done in the following lemma.

LEMMA 2. *For any positive pair $\mathbf{\Delta} = (\Delta^1, \Delta^2)$, there is a unique solution to Equations* (40) *and* (41).

The proof of Lemma 2 follows by showing that the response functions have opposite signs and can be found in the online appendix. Define the mapping

$$T^i(\mathbf{\Delta}) = p_2^i(\tilde{L}^i(z_f^i(\mathbf{\Delta})) - \alpha\Delta^i\lambda_f^{ij}(\mathbf{\Delta})\tilde{S}^i(z_f^i(\mathbf{\Delta}))) + r^i p_1^i + \alpha\Delta^i + p_2^j(A^i(\lambda_f^{ji}(\mathbf{\Delta})) - \alpha\Delta^i\lambda_f^{ji}(\mathbf{\Delta})\tilde{S}^j(z_f^j(\mathbf{\Delta}))). \tag{42}$$

Then, it will be shown that a fixed point solution such that $\Delta^1 = T^1(\mathbf{\Delta})$ and $\Delta^2 = T^2(\mathbf{\Delta})$ will be such that $\Delta^1 = a^1 - b^1$ and $\Delta^2 = a^2 - b^2$ in the infinite horizon equilibrium. The following lemma establishes preliminaries for the existence of such a fixed point. Its proof is primarily algebraic and may be found in the online appendix.

LEMMA 3. *Define*

$$\Delta_{\max}^i = \frac{p_2^i\tilde{L}^i(z_{my}^i) + r^i p_1^i}{1-\alpha},$$

$$\lambda_{\max}^{ij} = \min(A_j'^{-1}(\alpha\Delta_{\max}^j\tilde{S}^i(z_{my}^i)), 1),$$

$$z_{my}^i = 1 - \frac{h^i}{\tilde{r}^i + h^i}, \qquad z_{\max}^i = 1 - \frac{h^i}{\tilde{r}^i + \alpha\Delta_{\max}^i\gamma^i + h^i}$$

If $1 - p_2^i\lambda_{\max}^{ij}\tilde{S}^i(z_{my}^i) - p_2^j\lambda_{\max}^{ji}\tilde{S}^j(z_{my}^j) \geq 0$ then let

$$\Delta_{\min}^i = \frac{p_2^i\tilde{L}^i(z_{my}^i) + r^i p_1^i}{1 - \alpha(1 - p_2^i\lambda_{\max}^{ij}\tilde{S}^i(z_{my}^i) - p_2^j\lambda_{\max}^{ji}\tilde{S}^j(z_{my}^j))},$$

else let $\Delta_{\min}^i = 0$. Finally, let

$$\lambda_{\min}^{ij} = \min(A_j'^{-1}(\alpha\Delta_{\min}^j\tilde{S}^i(z_{\max}^i)), 1).$$

Then, if (Δ^i, Δ^j) is a fixed point of $T^i(\cdot)$, $T^j(\cdot)$, then $\Delta^i \in [\Delta_{\min}^i, \Delta_{\max}^i]$ and $\Delta^j \in [\Delta_{\min}^j, \Delta_{\max}^j]$. Further, for any $\Delta^i \in [\Delta_{\min}^i, \Delta_{\max}^i]$ and $\Delta^j \in [\Delta_{\min}^j, \Delta_{\max}^j]$,

$$z_{my}^i \leq z_f^i(\mathbf{\Delta}) \leq z_{\max}^i \quad \text{and} \quad \lambda_{\min}^{ij} \leq \lambda_f^{ij}(\mathbf{\Delta}) \leq \lambda_{\max}^{ij}.$$

To show there is a unique fixed point, we need to show that $|(\partial/\partial\Delta^j)T^i(\mathbf{\Delta})| < 1$. A relatively strong set of assumptions that implies this follows:

ASSUMPTION 12. *Assume*

$$p_2^j\frac{\Delta_{\max}^i}{\Delta_{\min}^j} \leq 2, \qquad \frac{\Delta_{\min}^i}{\Delta_{\max}^j} \geq \frac{1}{2}p_2^i,$$

$$p_2^i\Delta_{\max}^i < \min_{\lambda_{\min}^{ij} \leq \lambda \leq \lambda_{\max}^{ij}} A_j''(\lambda).$$

Further, for all z, assume

$$\frac{\phi_i(z)}{\bar{\Phi}_i(z)} \geq 1.$$

These assumptions say that the ranges of the Δ^i, Δ^j cannot be too far apart, that the second derivative of the cost function is sufficiently large relative to Δ^i, and that the hazard rate of ε^i, $\phi_i(z)/\bar{\Phi}_i(z)$, is at least one. These assumptions are driven by the algebra and unfortunately do not appear particularly intuitive. A weaker but somewhat more obscure assumption that can easily be shown to be implied by these assumptions is given as Assumption 13 in the appendix.

LEMMA 4. *Under Assumption* 12 *or* 13, *the mappings $T^i(\mathbf{\Delta})$, $T^j(\mathbf{\Delta})$ have a unique fixed point.*

Let $\mathbf{\Delta}^*$ be the unique fixed point of the mappings $T^i(\mathbf{\Delta})$, $T^j(\mathbf{\Delta})$. Define

$$a^i = \frac{p_2^i(\tilde{L}^i(z_f^i(\mathbf{\Delta}^*)) - \alpha\Delta^{*i}\lambda_f^{ij}(\mathbf{\Delta}^*)\tilde{S}^i(z_f^i(\mathbf{\Delta}^*))) + r^i p_1^i}{1-\alpha}$$

$$b^i = \frac{p_2^j(\alpha\Delta^{*i}\lambda_f^{ji}(\mathbf{\Delta}^*)\tilde{S}^j(z_f^j(\mathbf{\Delta}^*)) - A_i(\lambda_f^{ji}(\mathbf{\Delta}^*)))}{1-\alpha}$$

$$c^i = \frac{1}{1-\alpha}\big(p_3^i(\tilde{L}^i(z_f^i(\mathbf{\Delta}^*)) - \alpha\Delta^{*i}\lambda_f^{ij}(\mathbf{\Delta}^*)\tilde{S}^i(z_f^i(\mathbf{\Delta}^*))) + p_3^j(\alpha\Delta^{*i}\lambda_f^{ji}(\mathbf{\Delta}^*)\tilde{S}^j(z_f^j(\mathbf{\Delta}^*)) - A_i(\lambda_f^{ji}(\mathbf{\Delta}^*)))\big).$$

so that $\Delta^{*i} = a^i - b^i$.

We are now ready to establish our main result, which shows that there is an equilibrium in stationary state-independent policies in the infinite horizon discounted game. An equilibrium in stationary policies is weaker than the Markov equilibrium of the previous section. It is one where all firms precommit to a fixed policy for the infinite horizon and then a one-shot game is played on the policy space. Such an equilibrium is not guaranteed to be subgame perfect (and is likely not). This is the same type of equilibrium used in most inventory games considered over the infinite horizon (e.g., Avsar and Baykal-Gürsoy 2002, Bernstein and Federgruen 2004, Cachon and Zipkin 1999) and it appears that a stronger result must await theory development in Markov games.

For the inventory game considered here, we must define what it means to follow a state-independent inventory policy. For ease of exposition, we assume that the firm may costlessly reduce down to the desired inventory level. However, we will also show that, operating under the equilibrium policy, the inventory level is never above the desired level and so this (somewhat unrealistic) option is never actually used. One could more realistically assume that the firm simply orders nothing and allows demand to draw down inventory if the firm finds itself above the stationary level, but (because this still never occurs) this complicates the proof with little extra value added.

Theorem 4. *Under Assumptions* 8, 9, *and* 12 *or* 13, $(z_f^i(\mathbf{\Delta}^*), z_f^j(\mathbf{\Delta}^*), \lambda_f^{ij}(\mathbf{\Delta}^*), \lambda_f^{ji}(\mathbf{\Delta}^*))$ *form an equilibrium in state-independent stationary policies in the infinite horizon discounted game. The expected discounted payoff function for firm i under this equilibrium for starting state* $(x^i, x^j, \theta^i, \theta^j)$ *with* $x^i \le y_f^i(z_f^i(\mathbf{\Delta}^*), \theta^i)$ *and* $x^j \le y_f^j(z_f^j(\mathbf{\Delta}^*), \theta^j)$ *equals* $a^i\theta^i + b^i\theta^j + c^i$.

We have the following comparative statics for the equilibrium that, similar to Proposition 1, are proven in the online appendix using the implicit function theorem on the mapping $T^i(\cdot)$. The sufficient condition used in this proposition will be shown to be quite weak in the associated proof.

Proposition 2. *When* $\tilde{S}^j(z_f^j(\mathbf{\Delta})) - \Delta^i\partial\tilde{S}^j(z_f^j(\mathbf{\Delta}))/\partial\Delta^i \ge 0$, (a) *firm i's* $(i \ne j)$ *equilibrium "value increment"* Δ^{i*} *increases in* r^i *(when* $p_1^i \ge p_2^i$*),* r^j, $-h^i$, $-h^j$, *and* $-\gamma^i$; (b) *firm i's* $(i \ne j)$ *equilibrium stocking level increases in* r^i, r^j, $-h^i$, *and* $-h^j$; *and* (c) *firm i's* $(i \ne j)$ *equilibrium incentive level increases in* r^i, r^j, $-h^i$, $-h^j$, *and* $-\gamma^i$.

This result establishes that a firm's valuation of one of its committed customers over the valuation of its potential customer, its equilibrium stocking level, and its equilibrium incentive level increase when either firm's retail price increases or either firm's holding cost decreases. The reasons for its own retail price and holding cost are straightforward. The reason for these changes in the other firm's retail price and holding cost is simply that the other firm will increase its value increment and stocking level and the original firm will, too, in response.

Our incentives are targeted toward dissatisfied customers from the competitor's market. Because both retailers can retain their own customers by performing well with their inventory decisions and limiting the number of stockouts as far as is economically sensible, the inventory decision is partly an incentive in itself. In research not reported here, we considered a model where one firm can directly attract any of the competitor's customers. Unfortunately, we found that the conditions needed to show base-stock equilibrium policies (the focus of this paper) are too restrictive to make the model of general interest. It may also be possible for the firm to work to retain its own dissatisfied customers (other than with available inventory). In that case, the interaction between the firm's actions and its competitor's actions would need to be carefully delineated. Future work should investigate such competitive models.

4. Conclusions and Extensions

Consumers in classical dynamic inventory models are assumed to be backlogged (most common), lost (next most common), partially backlogged and partially lost (relatively uncommon), or gone from the market thereby reducing future demand (rare). In the first three cases, the firm's economic burden from not satisfying customers is usually approximated using a simple unit stockout cost. Although there has been widespread agreement that one (significant) element of the unit stockout cost is to reflect the economic consequences of some of these dissatisfied customers leaving the firm's market (thus reducing future demand), there has been little research investigating how this phenomenon affects the optimal (or equilibrium) inventory policy. Notable exceptions to this statement include Fergani (1976), Hall and Porteus (2000), and Liu et al. (2007). In addition to explicitly modeling the effect of future stockouts on demand, we explicitly incorporate the three previously mentioned stockout alternatives (in contrast to Hall and Porteus 2000 and Liu et al. 2007), we include the possibility of consumer forgiveness (in contrast to Fergani 1976), and we consider the possibility of attracting new customers.

We first consider a single firm. This firm could be considered as one firm operating under perfect competition, a price-taking monopolist, or simply one where pricing decisions are made by separate decision makers on a longer time frame. Each firm decides stocking levels, the proportion of "latent" customers that can be convinced to become "committed," and the extent of external advertising to increase the committed pool of customers. We establish sufficient conditions under which the optimal inventory policy is base-stock for the finite and infinite time horizons. Although we do not consider the conditions strenuous, they do suggest that the unit stockout cost may not be a good proxy under all circumstances. When the conditions are supported, we find a closed-form solution for the unit stockout cost, representing the discounted lost value premium of those lost customers. In addition, we find "lifetime" values of committed and latent customers. The optimal base-stock level increases with the retail price, the proportion of nonbacklogging customers who leave, and the value premium the committed customers have over the latent customers; it decreases with the unit holding cost and proportion of stocked-out customers who wish to backlog.

The natural extension to the single-firm model is a duopoly where a customer leaving one firm's market joins the other firm's market and vice versa. In the initial duopoly, firms decide only upon inventory levels and conditions are found under which the firms will operate under a base-stock equilibrium policy. Due primarily to the fact that a leaving customer will join the other firm's market but not search within the same time period, the equilibrium separates in every period. In the subsequent duopoly model, an incentive decision is included with the inventory decision. The incentive decision is advertising targeted toward the other firm's dissatisfied customers. We establish conditions under which a base-stock inventory policy is an equilibrium in stationary policies.

As mentioned in §2, we assume the firms actually know the market sizes, whereas in reality they may only have estimates. A model with Bayesian updating such as that in Fergani (1976) could likely be used to accommodate this uncertainty and is an interesting topic for future research. It is also possible to describe each market as a vector where each element represents the number of consumers who have been in the market for a particular number of periods, and the sum of the elements is the total market size. Unfortunately, for our method of analysis to be sustained, overly restrictive assumptions are needed so that this extension was not pursued further and a more general model (with a different type of analysis) is left as the subject of future research. Finally, we assume that lead times are zero. As nonzero lead times with lost sales assumptions typically present a challenging problem, it is likely that incorporating lead time in our models will present similar challenges.

5. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://mansci.journal.informs.org/.

Acknowledgments

The authors appreciate the comments of Hyun-soo Ahn, Yehuda Bassok, Roman Kapuściński, Ben Polak, Evan Porteus, Martin Shubik, Matt Sobel, Eilon Solon, and Dick Wittink; seminar participants at University of California, Los Angeles, Cornell, Yale, IBM Research, Stanford, and the University of Chicago; and conference attendees at INFORMS, MSOM, and ORSIS. This research was supported in part by National Science Foundation Grant DMI-0245382. Financial support from the University of Chicago Graduate School of Business is gratefully acknowledged. The authors appreciate the support of Department Editor Paul Zipkin, the associate editor, and two anonymous referees, whose comments have considerably improved the paper.

Appendix

Proof of Theorem 1. The proof is inductive and we establish the basis in the online appendix. For period t, $1 \le t \le T-1$, assume:

- $V_{t+1}(x, \theta, \beta) = a_{t+1}\theta + b_{t+1}\beta + c_{t+1}$ for $x \le y_f(z^*_{t+1}, \theta)$ and is bounded above by $V_{t+1}(y_f(z^*_{t+1}, \theta), \theta, \beta)$ for $x > y_f(z^*_{t+1}, \theta)$;
- $z^*_{my} \le z^*_{t+1}$, $0 \le a_{t+1} - b_{t+1} \le a_{t+2} - b_{t+2}$, $a_{t+1} \le a_{t+2}$, and $b_{t+1} \le b_{t+2}$.

For $z_t \le z^*_{t+1}$, $x_{t+1} \le y_f(z^*_{t+1}, \theta_{t+1})$ by the following reasoning. Observe

$$x_{t+1} = (p_2\theta_t + p_3)(\Phi^{-1}(z_t) - \varepsilon_t + \Gamma_t(\varepsilon_t - \Phi^{-1}(z_t))^+).$$

If $(\varepsilon_t - \Phi^{-1}(z_t))^+ = 0$, then there were no dissatisfied customers so that $\theta_{t+1} \ge \theta_t$ and therefore $x_{t+1} \le (p_2\theta_t + p_3)\cdot \Phi^{-1}(z_t) \le (p_2\theta_{t+1} + p_3)\Phi^{-1}(z_t) \le (p_2\theta_{t+1} + p_3)\Phi^{-1}(z^*_{t+1}) \le y_f(z^*_{t+1}, \theta_{t+1})$. For the case where $(\varepsilon_t - \Phi^{-1}(z_t))^+ > 0$, $x_{t+1} < 0 \le y_f(z^*_{t+1}, \theta_{t+1})$. Therefore, for $z_t \le z^*_{t+1}$, $\mathrm{E}V_{t+1}(x_{t+1}, \theta_{t+1}, \beta_{t+1}) = a_{t+1}\mathrm{E}\theta_{t+1} + b_{t+1}\mathrm{E}\beta_{t+1} + c_{t+1}$.

For $z_t > z^*_{t+1}$,

$$\begin{aligned}
&\mathrm{E}V_{t+1}(x_{t+1}, \theta_{t+1}, \beta_{t+1})\\
&\quad\le a_{t+1}\mathrm{E}\theta_{t+1} + b_{t+1}\mathrm{E}\beta_{t+1} + c_{t+1}\\
&\quad= a_{t+1}(\theta_t - (p_2\theta_t + p_3)S(z_t) + \rho_t\beta_t + \nu_t)\\
&\qquad+ b_{t+1}((1-\rho_t)\beta_t + (p_2\theta_t + p_3)S(z_t)) + c_{t+1}\\
&\quad= a_{t+1}\theta_t - (p_2\theta_t + p_3)S(z_t)(a_{t+1} - b_{t+1})\\
&\qquad+ \beta_t(\rho_t(a_{t+1} - b_{t+1}) + b_{t+1}) + a_{t+1}\nu_t + c_{t+1}.
\end{aligned} \tag{43}$$

Define $f_t(z) \overset{\Delta}{=} \tilde{L}(z) - \alpha S(z)(a_{t+1} - b_{t+1})$. We will show that $z_t^* = \arg\max_z f_t(z)$. Note that in this case, by the concavity of $-S(z)$, $z_t^* \geq z_{my}^*$. Now,

$$\begin{aligned}(p_2\theta_t + p_3)\tilde{L}(z_t) &+ \alpha \mathrm{E} V_{t+1}(x_{t+1}, \theta_{t+1}, \beta_{t+1}) \\ &\leq (p_2\theta_t + p_3) f_t(z_t) + \alpha\beta_t(\rho_t(a_{t+1} - b_{t+1}) + b_{t+1}) \\ &\quad + \alpha a_{t+1}(\theta_t + \nu_t) + \alpha c_{t+1}.\end{aligned}$$

By the concavity of $-S(z)$ and $\tilde{L}(z)$ and the nondecreasing nature of $a_t - b_t$ (by the induction assumption), $\arg\max_z f_t(z) \leq \arg\max_z f_{t+1}(z) = z_{t+1}^*$. Therefore, by the concavity of $f_t(\cdot)$, $f_t(z_t) \leq f_t(z_{t+1}^*)$ for $z_t > z_{t+1}^*$. Consequently, we can exclude consideration of $z_t > z_{t+1}^*$. Therefore,

$$\begin{aligned}V_t(x, \theta, \beta) &= \max_{\substack{z_{t+1}^* \geq z \geq y^{-1}(x,\theta) \\ 0 \leq \rho \leq 1,\, \nu \geq 0}} [(p_2\theta + p_3)\tilde{L}(z) + rp_1\theta - C(\rho)\beta - K(\nu) \\ &\qquad + \alpha a_{t+1}\mathrm{E}[\theta_{t+1}] + \alpha b_{t+1}\mathrm{E}[\beta_{t+1}] + \alpha c_{t+1}],\end{aligned}$$

where $y^{-1}(x, \theta) = \Phi((x - p_1\theta)/(p_2\theta + p_3))$. Applying the same logic as in (43),

$$\begin{aligned}V_t(x, \theta, \beta) &= (p_2\theta + p_3) \max_{z \geq y^{-1}(x,\theta)} f_t(z) + \theta(rp_1 + \alpha a_{t+1}) \\ &\quad + \beta \max_{0\leq\rho\leq 1}(-C(\rho) + \alpha(\rho(a_{t+1} - b_{t+1}) + b_{t+1})) \\ &\quad + \alpha c_{t+1} + \max_{\nu \geq 0}(\alpha a_{t+1}\nu - K(\nu)) = a_t\theta + b_t\beta + c_t\end{aligned}$$

for $x \leq y_f(z_t^*, \theta)$, where a_t, b_t, and c_t are as defined in Equations (17)–(19). Because $z_{t+1}^* = \arg\max_z(\tilde{L}(z) - \alpha S(z) \cdot (a_{t+2} - b_{t+2}))$, $z_{t+1}^* \geq z_t^*$ through the induction assumption ($a_{t+2} - b_{t+2} \geq a_{t+1} - b_{t+1}$) and again by the concavity of $-S(z)$. Now,

$$\begin{aligned}a_t &= p_2 \max_z f_t(z) + \alpha a_{t+1} + rp_1 \\ &= p_2(\tilde{L}(z_t^*) - \alpha S(z_t^*)(a_{t+1} - b_{t+1})) + \alpha a_{t+1} + rp_1,\end{aligned}$$

which is increasing in t because $(1 - S(z)) \geq 0$ for all z and $(a_{t+1} - b_{t+1})$ is also increasing in t. Also,

$$\begin{aligned}b_t &= \max_\rho(-C(\rho) + \alpha\rho(a_{t+1} - b_{t+1})) + \alpha b_{t+1} \\ &= -C(\rho_t^*) + \alpha\rho_t^*(a_{t+1} - b_{t+1}) + \alpha b_{t+1},\end{aligned}$$

which is also increasing along similar reasoning to a_t. Now,

$$\rho_t^* = \arg\max_\rho(-C(\rho) + \alpha\rho(a_{t+1} - b_{t+1}))$$

so $\rho_t^* \leq \rho_{t+1}^*$ from the induction assumption. Further,

$$\begin{aligned}&a_t - b_t \\ &= p_2\tilde{L}(z_t^*) + rp_1 + C(\rho_t^*) + \alpha(a_{t+1} - b_{t+1})(1 - \rho_t^* - p_2S(z_t^*)) \\ &\leq p_2\tilde{L}(z_t^*) + rp_1 + C(\rho_{t+1}^*) + \alpha(a_{t+1} - b_{t+1})(1 - \rho_{t+1}^* - p_2S(z_t^*)) \\ &\leq p_2\tilde{L}(z_t^*) + rp_1 + C(\rho_{t+1}^*) + \alpha(a_{t+2} - b_{t+2})(1 - \rho_{t+1}^* - p_2S(z_t^*)) \\ &\leq p_2\tilde{L}(z_{t+1}^*) + rp_1 + C(\rho_{t+1}^*) + \alpha(a_{t+2} - b_{t+2})(1 - \rho_{t+1}^* - p_2S(z_{t+1}^*)) \\ &= a_{t+1} - b_{t+1},\end{aligned}$$

where the first inequality is due to the optimality of ρ_t^*, the second inequality is because $a_{t+1} - b_{t+1} \leq a_{t+2} - b_{t+2}$, and the third is by definition of z_{t+1}^*. Finally, $a_t - b_t \geq 0$ because $a_{t+1} - b_{t+1} \geq 0$ and $1 - \rho_t^* - p_2S(z_t^*) \geq 1 - \rho_T^* - p_2S(z_{my}^*) \geq 0$. □

Proof of Theorem 2. By definition, a and b simultaneously solve

$$a = p_2 \max_z(\tilde{L}(z) - \alpha S(z)(a - b)) + \alpha a + rp_1, \tag{44}$$

$$b = \max_{0\leq\rho\leq 1}(-C(\rho) + \alpha\rho(a - b)) + \alpha b. \tag{45}$$

Let ε, Γ, and Λ be some random realization of demand, nonbacklogging proportions, and leaving proportions, respectively. Define $X(z) = \Phi^{-1}(z) - \varepsilon + \Gamma(\varepsilon - \Phi^{-1}(z))^+$ and $Y(z) = \Lambda\Gamma(\varepsilon - \Phi^{-1}(z))^+$. From Theorem 4.2.3 in Hernández-Lerma and Lasserre (1996), the optimal stationary solution must satisfy

$$\begin{aligned}V^*(x, \theta, \beta) = \max_{\substack{z \geq y^{-1}(x,\theta) \\ 0\leq\rho\leq 1}} \big((p_2\theta + p_3)\tilde{L}(z) &+ rp_1\theta - C(\rho)\beta \\ &+ \alpha\mathrm{E}[V^*(y_f(X(z), \theta), \theta(1 - Y(z)) \\ &+ R(\rho)\beta, (1 - R(\rho))\beta + Y(z)\theta)]\big).\end{aligned}$$

We first consider the relaxed problem where there is no lower bound on z. Substituting $\widetilde{V}^*(x, \theta, \beta) = a\theta + b\beta + c$ into the right-hand side of the relaxed version of the above equation yields

$$\begin{aligned}&\max_{z,\, 0\leq\rho\leq 1} \big((p_2\theta + p_3)\tilde{L}(z) + rp_1\theta - C(\rho)\beta \\ &\qquad + \alpha(a(\theta(1 - S(z)) + \rho\beta) + b((1 - \rho)\beta + S(z)\theta) + c)\big) \\ &= (p_2\theta + p_3)\max_z(\tilde{L}(z) - \alpha(a - b)S(z)) + \theta(rp_1 + \alpha a) \\ &\quad + \beta\max_{0\leq\rho\leq 1}(-C(\rho) + \alpha(a - b)\rho + \alpha b) + \alpha c \\ &= \theta a + \beta b + c.\end{aligned}$$

Therefore, the optimal stationary policy for the relaxed problem is $(z_f(\Delta^*), \rho_f(\Delta^*))$. If this solution is also feasible for the original problem, then it must also be optimal for the original problem. Let $z^* = z_f(\Delta^*)$. This policy is feasible if future inventory is less than or equal to the desired future order-up-to point. Now, future inventory equals $(p_2\theta + p_3) \cdot (\Phi^{-1}(z^*) - \varepsilon + \Gamma(\varepsilon - \Phi^{-1}(z^*))^+)$ and the future desired order-up-to point equals

$$\begin{aligned}y_f(z^*, \theta(1 - Y(z^*)) + R(\rho)\beta) &= p_1(\theta(1 - Y(z^*)) + R(\rho)\beta) + \Phi^{-1}(z^*) \\ &\quad \cdot (p_2(\theta(1 - Y(z^*)) + R(\rho)\beta) + p_3).\end{aligned}$$

If $\varepsilon \leq \Phi^{-1}(z^*)$ so that there are no unsatisfied customers, then $Y(z) = 0$, and using the fact that demand must be nonnegative, the future desired order-up-to point is bounded below by

$$p_1\theta + (p_2\theta + p_3)\Phi^{-1}(z^*) \geq (p_2\theta + p_3)(\Phi^{-1}(z^*) - \varepsilon),$$

where the right-hand side equals future inventory (in this case). If $\varepsilon > \Phi^{-1}(z^*)$, then future inventory is negative but, again using nonnegativity of demand, the future desired order-up-to point is nonnegative. Thus, in both cases, future inventory is at most the future desired order-up-to point and z^* is indeed feasible. Hence, $(z_f(\Delta^*), \rho_f(\Delta^*))$ is the optimal stationary policy. □

Alternate (Weaker) Assumption to Assumption 12

A weaker assumption to Assumption 12 (which can easily be shown to be implied by Assumption 12) is as follows:

ASSUMPTION 13. *For all* $\Delta^i \in [\Delta^i_{\min}, \Delta^i_{\max}]$ *and* $\Delta^j \in [\Delta^j_{\min}, \Delta^j_{\max}]$,

$$\left|\frac{-\alpha^2 p_2^i \Delta^i \tilde{S}^i(z_f^i(\mathbf{\Delta}))^2 \phi_i(z_f^i(\mathbf{\Delta}))(g^i(\Delta^i, \lambda_f^{ij}(\mathbf{\Delta})))^2}{n^i(\mathbf{\Delta})} + \frac{\alpha^2 p_2^j h^j \Delta^i (\gamma^j)^2 (\lambda_f^{ji}(\mathbf{\Delta}))^2 \Pr(\varepsilon^j > z_f^j(\mathbf{\Delta})) A_i''(\lambda_f^{ji}(\mathbf{\Delta}))}{n^j(\mathbf{\Delta})}\right| < 1,$$

where $n^i(\mathbf{\Delta}) = A_j''(\lambda_f^{ij}(\mathbf{\Delta}))\phi_i(z_f^i(\mathbf{\Delta}))(g^i(\Delta^i, \lambda_f^{ij}(\mathbf{\Delta})))^2 + h^i\alpha^2\Delta^i \cdot \Delta^j(\gamma^i)^2 \Pr(\varepsilon^i > z_f^i(\mathbf{\Delta}))$ *and* $g^i(\Delta^i, \lambda^{ij}) = \tilde{r}^i + \alpha\Delta^i\lambda^{ij}\gamma^i + h^i$.

The assumption arises from the need for $|(\partial/\partial\Delta^j)T^i(\mathbf{\Delta})| < 1$ and the expression on the left will be shown to be $|(\partial/\partial\Delta^j)T^i(\mathbf{\Delta})|$ (see Lemma 5 in the online appendix).

PROOF OF THEOREM 4. Let $\mathbf{\Delta}^*$ be the unique fixed point of the mappings $T^i(\mathbf{\Delta})$, $T^j(\mathbf{\Delta})$ (which exists by the given assumptions and Lemma 4). Further, for any x, define the function

$$V^i(x, \theta^i, \theta^j) = a^i\theta^i + b^i\theta^j + c^i.$$

To show that $(z_f^i(\mathbf{\Delta}^*), z_f^j(\mathbf{\Delta}^*), \lambda_f^{ij}(\mathbf{\Delta}^*), \lambda_f^{ji}(\mathbf{\Delta}^*))$ form an equilibrium in stationary policies in the infinite horizon discounted game, we must show that

$$V^i(x^i, x^j, \theta^i, \theta^j) = \underset{\substack{z\ge 0\\ 0\le\lambda^{ji}\le 1}}{\text{eqm}} \Big[(p_2^i\theta^i + p_3^i)\tilde{L}^i(z) - (p_2^j\theta^j + p_3^j)A^i(\lambda^{ji}) + r^i p_1^i\theta^i + \alpha \mathrm{E}[V^i(x_{t+1}^i, x_{t+1}^j, \theta_{t+1}^i, \theta_{t+1}^j)]\Big],$$

where z is unrestricted due to our assumption that inventory may be drawn down costlessly. Further, we must show that for $x^i \le y_f^i(z_f^i(\mathbf{\Delta}^*), \theta^i)$ and $x^j \le y_f^j(z_f^j(\mathbf{\Delta}^*), \theta^j)$, $x_{t+1}^i \le y_f^i \cdot (z_f^i(\mathbf{\Delta}^*), \theta_{t+1}^i)$ and $x_{t+1}^j \le y_f^j(z_f^j(\mathbf{\Delta}^*), \theta_{t+1}^j)$.

We begin with the final point, which implies that a stationary state-independent order-up-to policy is feasible for the system where inventory may not be removed costlessly. Pick $0 \le z \le 1$. For $z_t^i \le z$, $x_{t+1}^i \le y_f^i(z, \theta_{t+1}^i)$ (and, hence, z is also feasible in the following period) due to the following reasoning. We have that

$$x_{t+1}^i = (p_2^i\theta_t^i + p_3^i)(\Phi_i^{-1}(z_t^i) - \varepsilon_t^i + \Gamma_t^i(\varepsilon_t^i - \Phi_i^{-1}(z_t^i))^+)$$

and

$$y_f^i(z, \theta_{t+1}^i) = p_1^i\theta_{t+1}^i + \Phi_i^{-1}(z)(p_2^i\theta_{t+1}^i + p_3^i).$$

If $\varepsilon_t^i \le \Phi_i^{-1}(z)$ so there are no unsatisfied customers, then $\theta_{t+1}^i \ge \theta_t^i$ and, using the fact that demand must be nonnegative, the future desired order-up-to point $y_f^i(z, \theta_{t+1}^i)$ is bounded below by

$$p_1^i\theta_t^i + \Phi_i^{-1}(z_t^i)(p_2^i\theta_t^i + p_3^i) \ge (p_2^i\theta_t^i + p_3^i)(\Phi_i^{-1}(z_t^i) - \varepsilon_t^i),$$

where the right-hand side equals future inventory x_{t+1}^i (in this case). If $\varepsilon_t^i > \Phi_i^{-1}(z)$, then future inventory x_{t+1}^i is negative but, again using nonnegativity of demand, the future desired order-up-to point $y_f^i(z_f^i(\mathbf{\Delta}^*), \theta_{t+1}^i)$ is nonnegative. Thus, in both cases, $x_{t+1}^i \le y_f^i(z, \theta_{t+1}^i)$.

Now,

$$\begin{aligned} &a^i\mathrm{E}\theta_{t+1}^i + b^i\mathrm{E}\theta_{t+1}^j \\ &\quad = a^i(\theta^i - (p_2^i\theta^i + p_3^i)\lambda^{ij}\tilde{S}^i(z^i) + (p_2^j\theta^j + p_3^j)\lambda^{ji}\tilde{S}^j(z^j)) \\ &\qquad + b^i(\theta^j - (p_2^j\theta^j + p_3^j)\lambda^{ji}\tilde{S}^j(z^j) + (p_2^i\theta^i + p_3^i)\lambda^{ij}\tilde{S}^i(z^i)). \end{aligned} \quad (46)$$

Fixing the opponents strategy at (z^j, λ^{ij}),

$$\begin{aligned} &\max_{\substack{z\ge y_i^{-1}(x^i,\theta^i)\\ 0\le\lambda^{ji}\le 1}} \Big[(p_2^i\theta^i + p_3^i)\tilde{L}^i(z) - (p_2^j\theta^j + p_3^j)A^i(\lambda^{ji}) + r^i p_1^i\theta^i + \alpha\mathrm{E}[V^i(x_{t+1}^i, x_{t+1}^j, \theta_{t+1}^i, \theta_{t+1}^j)]\Big] \\ &= (p_2^i\theta^i + p_3^i)\max_{z\ge y_i^{-1}(x^i,\theta^i)}[\tilde{L}^i(z) - \alpha\lambda^{ij}\tilde{S}^i(z^i)\Delta^{i*}] + \theta^i(r^i p_1^i + \alpha a^i) \\ &\quad + (p_2^j\theta^j + p_3^j)\max_{0\le\lambda^{ji}\le 1}[-A^i(\lambda^{ji}) + \alpha\lambda^{ji}\tilde{S}^j(z^j)\Delta^{i*}] + \alpha\theta^j b^i + \alpha c^i \\ &= (p_2^i\theta^i + p_3^i)\max_{z\ge y_i^{-1}(x^i,\theta^i)} M^i(z, \lambda^{ij}) + \theta^i(r^i p_1^i + \alpha a^i) + (p_2^j\theta^j + p_3^j) \\ &\quad \cdot \max_{0\le\lambda^{ji}\le 1} B^i(\lambda^{ji}, z^j) + \theta^j\alpha b^i + \alpha c^i, \end{aligned}$$

where $M^i(z, \lambda^{ij}) \triangleq \tilde{L}^i(z) - \alpha\lambda^{ij}\tilde{S}^i(z)\Delta^{i*}$ and $B^i(\lambda^{ji}, z^j) \triangleq -A^i(\lambda^{ji}) + \alpha\lambda^{ji}\tilde{S}^j(z^j)\Delta^{i*}$. However,

$$z_f^i(\mathbf{\Delta}^*) = \arg\max_{z\ge w} M^i(z, \lambda_f^{ij}(\mathbf{\Delta}^*)) \text{ for } w \le y_i^{-1}(x^i, \theta^i) \quad \text{and}$$

$$\lambda_f^{ji}(\mathbf{\Delta}^*) = \arg\max_{0\le\lambda^{ji}\le 1} B^i(\lambda^{ji}, z_f^j(\mathbf{\Delta}^*)).$$

Therefore, $(z_f^i(\mathbf{\Delta}^*), \lambda_f^{ji}(\mathbf{\Delta}^*))$ is an optimal response to $(z_f^j(\mathbf{\Delta}^*), \lambda_f^{ij}(\mathbf{\Delta}^*))$. □

References

Ahn, H.-S., T. L. Olsen. 2007. Inventory competition with subscriptions. Working paper, University of Michigan, Ann Arbor.

Anderson, E. T., G. J. Fitzsimons, D. Simester. 2006. Measuring and mitigating the cost of stockouts. *Management Sci.* **52**(11) 1751–1763.

Avsar, Z., M. Baykal-Gürsoy. 2002. Inventory control under substitutable demand: A stochastic game application. *Naval Res. Logist.* **49** 359–375.

Başar, T., G. J. Olsder. 1999. *Dynamic Noncooperative Game Theory*, 2nd ed. Society for Industrial and Applied Mathematics, Philadelphia.

Bernstein, F., A. Federgruen. 2004. Dynamic inventory and pricing models for competing retailers. *Naval Res. Logist.* **51** 258–274.

Cachon, G. P., P. H. Zipkin. 1999. Competitive and cooperative inventory policies in a two-stage supply chain. *Management Sci.* **45**(7) 936–953.

DeCroix, G. A., A. Arreola-Risa. 1998. On offering economic incentives to backorder. *IIE Trans.* **30**(8) 715–721.

Fergani, Y. 1976. A market-oriented stochastic inventory model. Unpublished Ph.D. dissertation, Stanford University, Stanford, CA.

Fitzsimons, G. J. 2000. Consumer response to stockouts. *J. Consumer Res.* **27** 249–266.

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

Gans, N. 2002. Customer loyalty and supplier quality competition. *Management Sci.* **48**(2) 207–221.

Gaur, V., Y.-H. Park. 2007. Asymmetric consumer learning and inventory competition. *Management Sci.* **53**(2) 227–240.

Hall, J., E. Porteus. 2000. Customer service competition in capacitated systems. *Manufacturing Service Oper. Management* **2**(2) 144–165.

Henig, M. I., Y. Gerchak. 2003. Product availability competition and the effects of shortages in a dynamic duopoly. Working paper, Tel-Aviv University, Tel-Aviv, Israel.

Hernández-Lerma, O., J. B. Lasserre. 1996. *Discrete-Time Markov Control Processes: Basic Optimality Criteria*. Springer-Verlag, New York.

Heyman, D., M. Sobel. 1984. *Stochastic Models in Operations Research*, Vol. 2. McGraw-Hill, New York.

Liberopoulos, G., I. Tsikis. 2006. Backorder penalty cost coefficient "b:" What could it be? Working paper, University of Thessaly, Volos, Greece.

Lippman, S. A., K. F. McCardle. 1997. The competitive newsboy. *Oper. Res.* **45**(1) 54–65.

Liu, L., W. Shang, S. Wu. 2007. Dynamic competitive newsvendors with service-sensitive demands. *Manufacturing Service Oper. Management* **9**(1) 84–93.

Mahajan, S., G. J. van Ryzin. 1999. Retail inventories and consumer choice. S. Tayur, R. Ganeshan, M. Magazine, eds. *Quantitative Models for Supply Chain Management*. Kluwer Academic Press, Boston, 491–551.

Netessine, S., N. Rudi. 2003. Centralized and competitive inventory models with demand substitution. *Oper. Res.* **51**(2) 329–335.

Netessine, S., N. Rudi, Y. Wang. 2006. Inventory competition and incentives to back-order. *IIE Trans.* **38**(11) 883–902.

Olsen, T. L., R. P. Parker. 2008. Dynamic inventory competition with stockout-based substitution. Working paper, The University of Chicago Graduate School of Business, Chicago.

Parker, R. P., R. Kapuściński. 2008. Managing a non-cooperative supply chain with limited capacity. Working paper, The University of Chicago Graduate School of Business, Chicago.

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

Porteus, E. L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Stanford, CA.

Robinson, L. W. 1990. Appropriate inventory policies when service affects future demands. Working paper, Johnson Graduate School of Management, Cornell University, Ithaca, NY.

Schwartz, B. L. 1966. A new approach to stockout penalties. *Management Sci.* **12**(12) B538–B544.

Schwartz, B. L. 1970. Optimal inventory policies in perturbed demand models. *Management Sci.* **16**(8) B509–B518.

Veinott, A. F. 1966. The status of mathematical inventory theory. *Management Sci.* **12**(11) 745–777.

Zipkin, P. H. 2000. *Foundations of Inventory Management*. McGraw-Hill Higher Education, New York.