KRANNERT SCHOOL OF MANAGEMENT

Purdue University
West Lafayette, Indiana

Dynamic Inventory Control with Satisfaction-Dependent Demand

By

Justin Azadivar
Max Shen
George Shanthikumar

Paper No. 1249
Date: August, 2010

Institute for Research in the Behavioral, Economic, and Management Sciences

Dynamic Inventory Control with Satisfaction-Dependent Demand

Justin Azadivar, Max Shen, George Shanthikumar

1 Abstract

In this paper, we consider the discrete multiperiod newsvendor dynamic inventory control problem where customers follow a simple satisfaction-based demand process, where their probability of demand depends on whether their demand was satisfied the last time they demanded a product, and observe the differences between optimal policies and myopic policies which do not directly consider how inventory policies can affect future demand. We confirm the intuitive result that inventory managers should tend to order more than the myopic policy when satisfied customers are more likely to demand product, and less than the myopic policy when satisfied customers are less likely to demand. Moreover, we find that, when choosing a fixed order policy, even an empirically myopic solution with perfect demand distribution information will move away from the optimum towards a suboptimal solution.

2 Introduction

Logistics problems typically address designing systems and making decisions to optimize objectives based on externally determined behavior. In practice, this external behavior, in turn, will change to reflect the realities created by the logistics decisions. These changes must be anticipated by the decisionmaker.

In this paper, the primary topic of discussion is inventory control problems with demand dependent on previous service. In particular, a simple customer satisfaction model applied to the canonical newsvendor problem is used as a basis for analysis.

In the literal interpretation of the newsvendor problem, where a newsvendor purchases a certain number of newspapers to sell to news-hungry consumers, the newsvendor must balance the ability to profit by having newspapers available to meet demand and the cost associated with purchasing more newspapers than she can sell. This paper considers the further factor that a consumer seeking a newspaper who finds the newsvendor sold out will be less likely to make future attempts to purchase the newspaper, and so the newsvendor should purchase additional papers to increase future demand as well as meet current demand.

The Binary Customer Satisfaction Model presented here is a simplistic approach to modeling human behavior. Modeling customer behavior is difficult as the actual behavior of customers varies dramatically from customer to customer and from day to day. However, as a simple model, it is straightforward to communicate and augment as needed. Customers are either satisfied or unsatisfied, and their behavior in both cases is to a certain extent sensible, though not necessarily the result of a mathematical optimization problem. Since many types of customers do not behave in such a manner, this isn't always a failing of the model.

The key observation is that a newsvendor considering the possibility that satisfaction in the current period may affect future demand must sometimes order more than a newsvendor purchasing the optimal number to meet demand in the current period. More importantly, even with perfect information about the demand distribution, such as through empirical observation over a long period of time, optimizing the newsvendor problem for that distribution may move away from an optimal result and converge to a suboptimal result. This failure to consider how inventory decisions affect future customer behavior thus yields a "spiralling" effect similar to (if less dramatic than) the one described in the 2006 study by Cooper et. al. [2]

Section 3 of this paper discusses literature surrounding inventory problems with satisfaction-dependent demands. Section 5 introduces the Binary Customer Satisfaction Model and its application to the newsvendor problem (Section 5.1). It is applied to three variants on the problem, the general dynamic case (Section 5.2), the fixed-order quantity case (Section 5.3) with both positive satisfaction and negative satisfaction effects, and the case presented by approaching the infinite limit of the customer base (Section 5.4). Numerical examples and implementation issues for each variant are described in Section 6. Section 7 concludes the paper with consideration of potential avenues for extension.

3 Related Literature

In 1966, Schwartz [13] considered a canonical deterministic inventory control problem and introduced "perturbed demand" (PD) models as an alternative to "goodwill" costs associated with stockouts or backorders. In these models, a potential demand rate was given, but when a firm adopts a policy that fails to meet some fraction of demand, the actual demand is a function of that fraction. In [14] he addresses the creation of optimal policies on PD models. This was extended by Hill [8] to include holding and fixed costs in the basic inventory control model, while Caine and Plaut [1] address optima for fixed quantity, fixed interorder time, and fixed initial inventory cases, as well as a stochastic model. The specific formulations of the goodwill effect involve customers reducing their order rate in response to a disappointment.

Customer demand models can be largely divided into two classes, though they are not exclusive. The first class approaches customer demand in the aggregate, and defines the demand as a function of satisfaction. The aggregate models represent individual customer behavior without directly modeling it. Modeling the customers directly makes up the other class of models. Obviously, any individual customer behavior model will give rise to an aggregate demand model, while many aggregate demand models can be reduced to individual customer behavior models, so the distinction refers to the modeling approach more than the final model.

An example of an aggregate demand model is Ernst and Cohen [3], which extends supply chain coordination to the case where demand is not external, but rather influenced by service level at the retailer. The demand is modified by a multiplier that is linear and increasing in fill rate. The demand model used is

$$\tilde{D}(X) = (1 + \nu(\Omega' - \Omega))\tilde{D}_0$$

Where $\tilde{D}_0$ represents the base demand for the base fill rate (the proportion of demand that is satisfied) Ω, and Ω' is the actual fill rate. ν is a constant which quantifies the rate at which satisfaction increases demand. Ernst and Powell [4] consider the case where mean and variance respond differently to service level changes by assuming that new customers attracted by higher service levels have a different variance and/or mean than the old pool of customers. They then allow the demand multiplier to be raised to a fractional exponent γ to represent diminishing returns for increased service level:

$$\tilde{D}(X) = (1 + \nu(\Omega' - \Omega))^{\gamma}\tilde{D}_0$$

They later [5] apply this new model to consider incentives for the manufacturer in the supply chain problem to improve service level, finding that manufacturers can provide an incentive to retailers for increasing fill rate to coordinate the supply chain. These are aggregate demand models because the actual behavior of individual customers is not directly modeled. Instead, an ad hoc model is used that has the correct properties of how we intuitively think demand should respond to satisfaction. One of the drawbacks of these methods is that extensions to the model where the result is not intuitively obvious becomes difficult.

Robinson [12] considers a dynamic model where the demand follows a distribution whose mean and variance are modified by satisfied and unsatisfied customers. The simplest version of this model is a normal distribution whose mean, μ, changes over time.

$$\mu_{t+1} = a + b\mu_t + r_s s_t - r_d d_t$$

Here, a and b are constants representing demand processes outside the control of the firm, while s_t and d_t are the number of satisfied and dissatisfied customers, who change the aggregate demand mean linearly, proportional to constants r_s and r_s. The model includes periodic review to allow the firm, rather than simply choosing a policy whose equilibrium behavior is most profitable, to instead adapt the policy to the changing customer behavior. While the model recognizes individual behavior, the ability to change the mean is again somewhat ad hoc. Intuitively, one would expect there to be limits to the extent to which satisfaction or dissatisfaction can affect the mean.

Li [9] suggests a demand function that recognizes the utility of shorter lead times and concludes whether the firm should make-to-stock or make-to-order. He then creates an explicit competitive model to represent more directly what happens to "lost demand" in these models. Customers arrive according to a Poisson process and purchase randomly from the retailers that have stock, or, if none are available, from whichever firm can fill the order first. This provides a well-defined customer behavior model that is intuitive, but for many consumer applications, the absence of any loyalty and accessibility to full information provides room for development.

Hall and Porteus [7] considers a more explicit customer behavior model in the context of inventory capacity competition. Customers who fail to receive service from their retailer simply go to another retailer. The expected fraction of a firm's customer base that experiences service failures is defined by some decreaseing function $h(y)$, where y is the ratio of the firm's leased capacity and its customer base. The expected fraction of customers who leave the firm is given with a multiplier, $\gamma h(y)$, although there are no assumptions on the exact distribution of this fraction.

Gaur and Park [6] provides a more detailed model. Customers learn about product availability from their own experience in seeking the product (and lack information about others' experiences), and react in a biased manner to satisfaction and unsatisfaction. The customers each develop an estimate p of the service level of a given firm, and update this over time.

$$
\begin{aligned}
p_{t+1} = &(1-\theta^u)p_t + \theta^u && \text{satisfying visit at time } t \\
&(1-\theta^d)p_t && \text{unsatisfying visit at time } t \\
&p_t && \text{no visit at time } t
\end{aligned}
$$

θ^u and θ^d represent terms which represent how much weight the customers place on satisfying and unsatisfying visits. Here we have a fully specified model for individual behavior of customers, where the estimate of service level increases with satisfying visits and decreases with unsatisfying visits, and both processes give rise to diminishing returns.

Liberopoulos et. al. [10] consider a duopoly in which a single customer chooses between two firms randomly but influenced by a credibility factor, which then changes in response to whether or not the demand was satisfied. At any time period, the customer has a credibility level a, and chooses supplier 1 with probability $P(a)$ and supplier 2 with probability $(1 - P(a))$, where $P(a)$ is nondecreasing in a. In each period, if the customer chooses a supplier, and his order is satisfied, the credibility for that supplier will jump (a increases in supplier 1, and decreases if supplier 2). If the customer's order is not satisfied, the credibility drops. In this case, goodwill is modeled directly and customer behavior is based on goodwill. The firms make dynamic inventory decisions based on their credibilities.

Olsen and Parker [11] consider a division of customers into a "committed pool" of customers who demand from a firm, and "latent customers" who no longer demand from the firm, due to previous inventory failure. The demand is a function of the size of the committed pool, θ.

$$D(\theta) = p_1\theta + (p_2\theta + p_3)\epsilon$$

Where p_1, p_2, and p_3 are nonnegative parameters and ϵ is a mean-zero random error variable. Aside from the inventory decision, the manager also makes marketing and advertising decisions to increase the size of the committed pool. It can attempt to persuade customers to move back from the latent pool through marketing, or increase the size of the committed market through advertising to a larger external market. For example, the control ρ represents how much effort is spent in trying to recover latent customers, and $R(\rho)\beta$ customers are recovered. Similarly, a random fraction of those customers who experience a stockout move from the committed pool to the latent pool. Since these effects are multiplicative, the pools must be assumed to be large enough to allow a continuous approximation, and the collective behavior of the customers must be correlated to maintain coefficient of variance over different pool sizes. In the duopoly version of this model, the unsatisfied customers move to the other firm's customer pool, rather than to a latent pool.

4 Notation

N: The number of potential customers in the system.

S_k: The number of satisfied customers at the beginning of period k.

D_k: The total number of customers seeking service in period k. ("aggregate demand")

D_k^u: The number of unsatisfied customers seeking service in period k.

D_k^s**:** The number of satisfied customers seeking service in period k.

ν_k**:** The number of customers that receive service in period k.

p**:** The probability an unsatisfied customer will seek service in any given period.

αp**:** The probability a satisfied customer will seek service in any given period. $(0 < \alpha < \frac{1}{p})$

λ**:** The expected number of arrivals if all customers are unsatisfied (Np)

y_k**:** The inventory level chosen for time period k.

c**:** The per unit order cost

r**:** The per unit sales value

Φ_k**:** The profit in period k.

$\phi(Q)$**:** The expected stationary per period profit given an ordering policy Q.

5 Binary Customer Satisfaction Model

Here we introduce the Binary Customer Satisfaction Model. The basic modeling assumption is to assume that every potential customer is in one of two states: satisfied or unsatisfied. With each state comes a different customer behavior. Customer satisfaction is affected by receiving or failing to receive service.

This approach is similar to the supplier credibility approach in Liberopoulos et. al. [10]. There are only two credibility levels, but each customer has its own credibility level and there are multiple customers. "Goodwill," the independent benefit of satisfying customer demand, is being explicitly modeled here on an individual basis for each customer.

5.1 Formulation

Consider a simple newsvendor inventory control model. In each time period, each individual has an independent probability p of seeking the product if unsatisfied, and probability αp of seeking the product if satisfied $(0 < \alpha < \frac{1}{p})$. The value α represents the impact on customer behavior that being satisfied has. For values greater than 1, this represents an increased willingness to demand product if the customer's previous visit yielded satisfactory results. For values less than 1, this represents a decreased desire for the product, and may represent a customer no longer needing a product once she receives it. The customers are satisfied if they receive a product, and unsatisfied if they try to obtain a product and fail. If they do not seek product, their satisfaction status is unchanged.

We use S_k to represent the number of satisfied customers at the beginning of time period k. We will consider the size of the potential customer pool, N, to be fixed. Each period will be treated as a separate newsvendor problem, with demand level determined from a sum of binary distributions defined by S_k.

Each period k involves the following events in order:

1) The service provider observes the satisfaction level, S_k.

2) The service provider chooses an inventory level, y_k.

3) A demand D_k is generated,

$$D_k = D_k^s(S_k, \alpha p) + D_k^u(N - S_k, p) \quad (1)$$

where $D_k^i(x, y)$ are independent binary variables with x trials and a success rate y. (The arguments of D_k^i will be suppressed throughout this paper)

4) The service provider satisfies demand ν_k. We will assume the provider satisfies as much demand as possible[1], and hence

$$\nu_k = min\{D_k, y_k\}$$

5) The customers forming the demand return to the customer pool, and update the satisfaction level, S_{k+1}.

$$S_{k+1} = S_k - D_k^s + \nu_k \quad (2)$$

Observation 5.1. *In a sample path stochastic sense, the following are jointly true:*
D_k^s is linear increasing in S,
$S - D_k^s$ is linear increasing in S,
D_k^u is linear decreasing in S, and
D_k is linear increasing (if $\alpha > 1$), decreasing (if $\alpha < 1$) or constant (if $\alpha = 1$) in S.

We use Shanthikumar and Yao [15] for a definition of second-order stochastic properties for integer variables. Refer to the appendix for a detailed discussion of coupling and the sample path stochastic order.

Proof. We can prove this by defining the following probability space: In each period, a uniform [0,1] random variable is generated for each customer in the customer pool. These random variables are independent of each other across customers and time periods. If the random variable is less than the probability the customer will demand a product (p for unsatisfied customers and αp for satisfied customers), that customer will demand a product.

Consider four sample paths on the same probability space. Path 1 has satisfaction at time k of S, paths 2 and 3 have satisfaction $S+1$, and path 4 has satisfaction $S+2$. We can define the values of the variables of interest with numbered superscripts, for example, D_k^{u1}, D_k^{u2}, D_k^{u3}, D_k^{u4}. Other variables are denoted similarly throughout these proofs.

We couple these four paths in the following way:

A) All customers but 2 behave identically over all sample paths.

B) The remaining 2 customers behave according to two uniform random variables as follows: For $i = 1, 2$, if customer i is satisfied, she will demand if U_i, a uniform[0,1] random variable, is less than αp. If unsatisfied, she will demand if that variable is less than p.

C) Neither customer is satisfied in path 1, the first customer is satisfied in path 2, the second customer is satisfied in path 3, and both are satisfied in path 4.

Thus:

[1] While there is nothing in the structure of the problem to prevent an inventory manager from choosing not to satisfy demand when she has inventory on hand, and the ability to affect future demand makes this choice non-trivial, the Appendix includes a proof that it is always optimal to satisfy demand when possible, so we will assume this behavior for the remainder of the paper

$$
\begin{aligned}
D_k^{s2} - D_k^{s1} = D_k^{s4} - D_k^{s3} \quad & = I(U_1 < \alpha p) \\
(S+1-D_k^{s2}) - (S-D_k^{s1}) = (S+2-D_k^{s4}) - (S+1-D_k^{s3}) & = 1 - I(U_1 < \alpha p) \\
D_k^{u2} - D_k^{u1} = D_k^{u4} - D_k^{u3} \quad & = -\, I(U_1 < p)
\end{aligned}
$$

$$
\begin{aligned}
D_k^2 - D_k^1 = D_k^4 - D_k^3 & = I(p \le U_1 < \alpha p) \quad \text{for } \alpha > 1 \\
& = -I(\alpha p \le U_1 < p) \; \text{for } \alpha < 1 \\
& = 0 \qquad\qquad\qquad\quad \text{for } \alpha = 1
\end{aligned}
$$

$I(x)$ is an indicator variable which equals 1 if its argument is true, and 0 otherwise. Since this indicator is nonnegative, the first three equations prove the first three results, while the last three present the cases to prove the last result. □

5.2 Dynamic Inventory Control

The solution to a single period newsvendor problem, as described above, is well-studied. To consider the impact that customer satisfaction has on inventory control decisions, we consider the multiperiod newsvendor problem with a Binary Customer Satisfaction Model. We will let r be the sales value, and c be the purchase cost of each unit. We will assume no scrap value, and $r > c > 0$.

The payoff for each period is[2]:

$$\Phi_k(y_k, D_k) = r\nu_k - cy_k$$

We define an ordering policy $Q_k(S)$ to be a function which takes as its argument the current satisfaction level S_k and returns an order quantity:

$$y_k = Q_k(S_k)$$

$Q_k(S)$ is defined on $S = 0, 1, 2, ..., N$.

The state is completely defined by S_k, which is only a function of the number of satisfied customers in the previous period, the demand which is a function of satisfaction, and the inventory decision (2). Given an ordering policy, this decision is also a function of satisfaction, so the state of each period depends only on the state of the previous period, defining a Markov chain. This makes the dynamic optimization of the ordering policy straightforward and typically computationally feasible.

If we ignore or assume away the existence of a time horizon, we can study the steady-state properties of ordering policies. As in the finite horizon case, the state is fully defined by the current satisfaction in the infinite horizon case, so we can assume the ordering policy is time-independent, and write it as $Q(S)$.

[2]For notational simplicity, we have assumed that values are not discounted over time. This assumption is not necessary for the results in this paper: Stochastic and comparative relationships still hold. Discounting would only affect the actual numerical results, and reduces the advantage of using optimal methods over myopic ones

It is not necessarily the case that a given ordering policy will yield a unique steady-state result[3]. If it does, we can refer to the steady-state distribution of satisfaction level as $\mathbf{S}(Q)$, a random variable whose distribution depends on the ordering policy Q. We can similarly define a demand distribution $\mathbf{D}(Q)$:

$$\mathbf{D}(Q) = \sum_{S=0}^{N} I(\mathbf{S}(Q) = S) D(S)$$

Where $D(S)$ is the random variable for demand given satisfaction level S, (equation (1)) and I is an indicator variable that returns 1 if its argument is true, and 0 otherwise.

We can then define the expected profit:

$$\phi(Q) = E[\Phi] = E[r \min\{Q, \mathbf{D}(Q)\} - cQ]$$

5.2.1 Optimal Policy

Finding the optimal policy is a straightforward dynamic optimization problem. If we begin at the time horizon K, and define a cost-to-go function G:

$$\begin{aligned} G_K(S) &= \max_{y \epsilon \{0..N\}} E[\Phi(y, D(S))] \\ G_k(S) &= \max_{y \epsilon \{0..N\}} E[\Phi(y, D(S)) + G_{k+1}(S - D^s(S) + \min\{D(S), y\})] \end{aligned} \tag{3}$$

Where $D^s(S)$ is the random variable for demand from satisfied customers given satisfaction level S. The arguments which give us these maxima make up the optimal ordering policy, $Q^*_k(S)$.

In the infinite horizon case, the steady-state optimum is the infinite limit as k approaches $-\infty$.

Observation 5.2. *$G_k(S)$ is nondecreasing in S when $\alpha > 1$.*

Proof. For $k = K$, this is true, as increased demand distribution yields a higher expected profit in the period for all values of y (except $y = 0$, when the profit will be equal regardless of demand distribution).

Assuming the observation to be true for $k + 1$, we note that both components of $G_k(S)$ increase in S for all values of y (again, excepting $y = 0$). This is true for Φ_k as it was for $k = K$. S_{k+1} is also larger in expectation, as $S - D^s$ is nondecreasing in S by Observation (5.1), and $min\{D(S), y\}$ increases in S in expectation as well. Thus, the cost-to-go function is also nondecreasing.

□

Observation 5.3. *$G_k(S)$ is nonincreasing in S when $\alpha < 1$.*

Proof. Again, for $k = K$, this is true, as increased satisfaction means decreased demand distribution by Observation (5.1), which in turn yields a lower expected profit except for $y = 0$.

Assuming the observation to be true for $k + 1$, Φ_k is nonincreasing as it was for $k = K$. The satisfaction of the next period is still nondecreasing in S, so both terms are nonincreasing in S. Thus, the cost-to-go function is nonincreasing.

□

[3] As an example of a policy which will not yield such a state, consider the policy of ordering 0 for satisfaction levels below $N/2$, and N otherwise. Depending on initial state, the steady state distribution for S will be either 0 or N

5.2.2 Myopic Comparison

We now compare the optimal solution described above to the myopic policy. The myopic policy is that which optimizes without regard for the impact on other time periods. In other words, to the myopic manager, every period is the single period newsvendor problem, and we can describe the myopic policy, $\hat{Q}(S)$:

$$\hat{Q}_k(S) = \operatorname*{arg\,max}_{y \epsilon \{0..N\}} E[r \min\{D(S), y\} - cy]$$

Note that even in the finite horizon case, the myopic policy is identical in all time periods.

The single period newsvendor problem optimum, and hence the myopic policy, is:

$$\hat{Q}_k(S) = F^{-1}_{D(S)}(1 - \frac{c}{r})$$

Where $F^{-1}_{D(S)}$ is the inverse distribution function for $D(S)$.

Theorem 5.4. *For $\alpha > 1$, for any value $\hat{Q}_k(S)$ that is myopically optimal, there exists an optimal solution $Q^*_k(S)$ such that $\hat{Q}_k(S) \leq Q^*_k(S)$.*

Proof. Because both the myopic and optimal managers optimize the same problem at time horizon K, such optima will be equal, and the theorem is true for $k = K$.

Recall that the optimal policy optimizes $E[\Phi_k + G_{k+1}(S_{k+1})]$, while the myopic policy optimizes $E[\Phi_k]$. But since S_{k+1} is nondecreasing in order quantity, by observation (5.2), the cost-to-go function is nondecreasing in order quantity. Therefore, if a manager orders less than $\hat{Q}_k(S)$, the expected profit will be lower, since lower values will occur for both the first term (because the myopic policy maximizes Φ_k) and the second term (by observation (5.2)). Thus, if $\hat{Q}_k(S)$ is not optimal, then the optimal order quantity must be higher. □

Theorem 5.5. *For $\alpha < 1$, for any value $\hat{Q}_k(S)$ that is myopically optimal, there exists an optimal solution $Q^*_k(S)$ such that $\hat{Q}_k(S) \geq Q^*_k(S)$.*

Because the argument is similar (though opposite), the proof has been omitted.

In other words, when satisfaction can affect future demand, optimal managers will tend to order more than the myopic manager when satisfied customers are more likely to arrive, and order less when they are less likely to arrive.

5.3 Fixed-Order Policies

We now consider a problem where the ordering policies are restricted to fixed-order policies. That is, in every period, the manager must order an amount Q which remains constant over time. The manager must choose Q to optimize long-term profit. We will look only at the infinite horizon problem, and consider the steady-state results of the inventory decision.

First, observe that a fixed-order policy Q will always generate a steady-state distribution on $[Q, N]$. We can refer to the steady-state random satisfaction generated by order quantity Q as $\mathbf{S}(Q)$, and the demand as $\mathbf{D}(Q)$. Unlike the dynamic policy, the neat symmetry between positive satisfaction effect ($\alpha > 1$) and negative satisfaction effect ($\alpha < 1$) no longer holds. While satisfaction increases in Q, which in turn causes demand to either increase ($\alpha > 1$) or decrease ($\alpha < 1$), ν, the number of customers receiving product in each period ($min\{D(Q), Q\}$), has dramatically different properties, as the minimum of two increasing concave functions in the same variable is not generally concave.

Optimization in this case simply involves a one-dimensional optimization of Q to maximize the expected steady-state profit function:

$$\phi(Q) = E[r \min\{Q, \mathbf{D}(Q)\} - cQ] \tag{4}$$

We refer to the optimum of this problem Q^*.

We also find that the myopic comparison is not as straightforward as it is in the dynamic case. While in the dynamic case, the manager can adapt the inventory policy to the new demand distribution generated in each period, the fixed-ordering manager cannot do so. It is also unreasonable to suppose that the manager will simply find the myopic policy in one period and use that order quantity.

The myopic approach we will use in this paper is to allow the manager to determine the distribution empirically, and then choose the inventory policy that optimizes the profit over that empirical demand distribution. A myopic manager would use some form of periodic review, in the sense that after a number of periods, if the new empirical distribution (based on the addition of new observations) would indicate a better profit with a new inventory level, the manager would change the order quantity to that new level. Such a manager is not ordering blindly, but is entirely backwards-looking, and does not consider the impact the inventory policy itself will have on the demand distribution.

We will define a class of policies as "empirically myopic." An empirically myopic ordering policy, $\hat{Q}$, generating a demand distribution $\mathbf{D}(\hat{Q})$, has the property that $\hat{Q}$ is the optimal order quantity Q for:

$$\hat{\phi}_{\hat{Q}}(Q) = E[r \min\{Q, \mathbf{D}(\hat{Q})\} - cQ] \tag{5}$$

$\hat{\phi}_{\hat{Q}}(Q)$ is the expected profit of order quantity Q given demand distribution $\mathbf{D}(\hat{Q})$. This differs from (4) in that the optimization is performed over a fixed demand distribution, rather than one that varies with the choice of $\hat{Q}$. If $\hat{Q}$ optimizes $\hat{\phi}_{\hat{Q}}(Q)$, then it is an empirically myopic quantity, as it will be the optimal ordering policy for the demand distribution it creates, without considering the possibility of creating other demand distributions. An empirically myopic manager who uses this policy will continue using it, as the empirical demand distribution will have the empirically myopic policy as an optimum. Such a policy will satisfy the following result:

$$\hat{Q} = F^{-1}_{\mathbf{D}(\hat{Q})}(1 - \frac{c}{r}) \tag{6}$$

Note that such a policy need not exist, and multiple such policies may be possible.

5.3.1 Negative Satisfaction

We begin by analyzing the situation where $\alpha \leq 1$. Satisfied customers are less likely to arrive than unsatisfied customers. An example of such a situation might be the support service for a program. The manager chooses a certain amount of support capacity (the "inventory") which is then used by customers. Customers who receive support no longer need support and are less likely to seek support in the future, while customers who do not receive support will still need support, and are more likely to seek support.

Theorem 5.6. *When $\alpha \leq 1$, for fixed order policies of order quantity Q beginning with a satisfaction S_k at time k, S_l is stochastically nondecreasing and concave (in a sample path sense) in Q for all $l \geq k$.*

Proof. For $l = k$, the theorem is trivially true, as S_k will not depend on Q. We now assume that the theorem is true for l and prove it for $l + 1$.

As with observation (5.1), we begin by constructing four sample paths. Path 1 uses order quantity Q, paths 2 and 3 use order quantity $Q + 1$, and path 4 uses order quantity $Q + 2$.

By assumption, paths i ($i = 1, 2, 3, 4$) have satisfaction at time l of S_l^i such that:

$$S_l^1 \leq S_l^2 \leq S_l^4 \tag{7}$$

$$S_l^1 \leq S_l^2 \leq S_l^4 \tag{8}$$

$$S_l^2 + S_l^3 \geq S_l^1 + S_l^4 \tag{9}$$

The four paths are coupled as follows:

A) All customers but $S_l^4 - S_l^1$ behave identically over all sample paths.

B) The remaining customers behave according to uniform random variables as follows: For $i = 1, ...S_l^4 - S_l^1$, if customer i is satisfied, she will demand if U_i, a uniform[0,1] random variable, is less than αp. If unsatisfied, she will demand if that variable is less than p.

C) None of these customers are satisfied in path 1, the first $S_l^2 - S_l^1$ customers are satisfied in path 2, the last $S_l^3 - S_l^1$ customers are satisfied in path 3, and all are satisfied in path 4.

From observation (5.1), we recall that $S_l - D_l^s$ is linear and increasing in S_l. Taking the differences in the satisfaction update equations (2) for these paths gives us:

$$S_{l+1}^2 - S_{l+1}^1 = \sum_{i=1}^{S_l^2 - S_l^1} I(U_i > \alpha p) + \min\{D_l^2, Q + 1\} - \min\{D_l^1, Q\} \tag{10}$$

$$S_{l+1}^3 - S_{l+1}^4 = -\sum_{i=1}^{S_l^4 - S_l^3} I(U_i > \alpha p) + \min\{D_l^3, Q + 1\} - \min\{D_l^4, Q + 2\} \tag{11}$$

Adding these together while assuming (7) gives us

$$S_{l+1}^2 + S_{l+1}^3 - S_{l+1}^1 - S_{l+1}^4 = \tag{12}$$

$$\sum_{i=S_l^4 - S_l^3 + 1}^{S_l^2 - S_l^1} I(U_i > \alpha p) + \min\{D_l^2, Q + 1\} + \min\{D_l^3, Q + 1\} - \min\{D_l^1, Q\} - \min\{D_l^4, Q + 2\}$$

If this equation is nonnegative almost surely under this coupling, then the theorem is true.

Recall that demand is nonincreasing and linear in S, so the demand comparisons are opposite the satisfaction comparisons, and will hold almost surely under this coupling. For example, path 1, having the lowest satisfaction, will always have the highest demand.

Case 1: $D_l^4 \geq Q + 2$.

Because path 4 has the highest satisfaction, all four paths demand at least D_1^4, which means all four inventory constraints will be binding, so ν_l is linear. The last four terms sum to 0, leaving only the sum term, which is nonnegative.

Case 2: $D_l^2 \geq Q+1$, $D_l^3 \geq Q+1$, $D_l^4 < Q+2$

In this case, only path 4 has a nonbinding inventory constraint, which gives a result larger than case 1.

Case 3: $D_l^2 \geq Q+1$, $D_l^3 < Q+1$

Both path 3 and path 4 will have nonbinding inventory constraints, while paths 1 and 2 do not. Since $D_l^3 \geq D_l^4$, this will also give a result no smaller than case 1.

Case 4: $D_l^2 < Q+1$, $D_l^3 \geq Q+1$

This case is essentially identical to case 3, except we compare $D_l^2 \geq D_l^4$.

Case 5: $D_l^1 \leq Q$

Here, none of the inventory levels will be binding, so $\nu_l = D_l$. If we add up these terms in (12), we find:

$$D_{l+1}^2 + D_{l+1}^3 - D_{l+1}^1 - D_{l+1}^4 = -\sum_{i=S_l^4-S_l^3+1}^{S_l^2-S_l^1} I(p \geq U_i > \alpha p) \geq -\sum_{i=S_l^4-S_l^3+1}^{S_l^2-S_l^1} I(U_i > \alpha p) \tag{13}$$

Thus, when added to the first sum, the equation remains nonnegative.

Case 6: $D_l^1 > Q$, $D_l^2 < Q+1$, $D_l^3 < Q+1$

The inventory constraint is binding on path 1, resulting in a value larger than in case 5.

This proves concavity for Theorem (5.6). To prove the nonincreasing aspect, simply observe that (10) is always nonnegative, and (11) must be nonpositive (as the differences in demands must be no greater than the difference in satisfaction, and path 4 has a higher inventory level). The same would be true if paths 2 and 3 were interchanged.

□

Corollary 5.7. *When $\alpha \leq 1$, for fixed order policies of order quantity Q, the steady state-distribution for S is stochastically nondecreasing and concave in Q in a sample path sense.*

We use the nondecreasing property from Theorem 5.6 to reach a fixed-order analogue of Theorem 5.5 while comparing the optimal and empirically myopic solutions.

Theorem 5.8. *For $\alpha < 1$, for any value $\hat{Q}$ that is empirically myopic, there exists an optimal solution Q^* such that $\hat{Q} \geq Q^*$.*

Proof. Suppose $Q^* \geq \hat{Q}$. Then, because Q^* optimizes (4):

$$E[r\min\{Q^*, \mathbf{D}(Q^*)\} - cQ^*] \geq E[r\min\{\hat{Q}, \mathbf{D}(\hat{Q})\} - c\hat{Q}] \tag{14}$$

Since $\hat{Q}$ optimizes (5):

$$E[r\min\{\hat{Q}, \mathbf{D}(\hat{Q})\} - c\hat{Q}] \geq E[r\min\{Q^*, \mathbf{D}(\hat{Q})\} - cQ^*] \tag{15}$$

But since demand is nonincreasing in Q, a minimum is nondecreasing in its arguments, and we assumed $Q^* \geq \hat{Q}$:

$$E[r\min\{Q^*, \mathbf{D}(\hat{Q})\} - cQ^*] \geq E[r\min\{Q^*, \mathbf{D}(Q^*)\} - cQ^*] \tag{16}$$

The right hand side of (14) equals the left hand side of (15), The right hand side of (15) equals the left hand side of (16), and the right hand side of (16) equals the left hand side of (14). Thus all three quantities are equal, and $\hat{Q}$ must also be optimal.

□

5.3.2 Positive Satisfaction

We now consider the case where $\alpha > 1$. Note that this theorem is weaker than Theorem 5.6, as it does not prove second-order properties.

Theorem 5.9. *When $\alpha > 1$, for fixed order policies of order quantity Q beginning with a satisfaction S_k at time k, S_l is nondecreasing in Q for all $l \geq k$.*

Proof. Again, for $l = k$, the theorem is trivially true.

We now assume the theorem to be true for $l - 1$. The update equation (2) can be divided into two parts. The first part, $S_{l-1} - D^s_{l-1}$, is nondecreasing in S_{l-1} as noted in observation (5.1). The second part, $min\{Q, D_{l-1}\}$, is clearly increasing in Q. When $\alpha > 1$, it also increases in S_{l-1}, as the minimum of two increasing functions will increase. □

Corollary 5.10. *When $\alpha > 1$, for fixed order policies of order quantity Q, the steady state-distribution for S is nondecreasing in Q.*

The positive satisfaction version of Theorem 5.8 also holds.

Theorem 5.11. *For $\alpha > 1$, for any value $\hat{Q}$ that is myopically optimal, there exists an optimal solution Q^* such that $\hat{Q} \leq Q^*$.*

The proof is again nearly identical, except we note that demand is nondecreasing in Q. However, this theorem gives rise to a key managerial insight into the shortcomings of relying solely on empirical demand distributions.

Theorem 5.12. *If $\alpha > 1$, an empirically myopic policy exists, and sequentially finding myopically optimal policies for the distributions created by the previous policy will eventually induce an empirically myopic policy, regardless of initial demand distribution.*

Proof. The initial demand distribution will have a myopically optimal inventory level, which we will denote as Q^0. This, in turn, generates the demand distribution $\mathbf{D}(Q^0)$. Thus, we can limit our consideration to demand distributions generated by inventory decisions. The sequence of inventory decisions updates as follows:

$$Q^{i+1} = F^{-1}_{\mathbf{D}(Q^i)}(1 - \frac{c}{r})$$

The theorem claims that this sequence converges to an empirically myopic policy. According to (6), if $Q^{i+1} = Q^i$, this is an empirically myopic policy.

Suppose that $Q^{i+1} > Q^i$. Since $\mathbf{D}(Q)$ is nondecreasing in Q, $\mathbf{D}(Q^{i+1}) \geq \mathbf{D}(Q^i)$, and thus $Q^{i+2} \geq Q^{i+1}$. If the inventory levels are equal, this is an empirically myopic policy. If not, we repeat this process to find Q^{i+3} and so on. But note that if $Q^{i+1} = N$, then $Q^{i+2} = Q^{i+1}$, as demand is bounded by N, and $N+1$ can never be a myopically optimal solution. As a result, we will eventually reach a myopic policy.

By the same argument, if $Q^{i+1} < Q^i$, $Q^{i+2} \leq Q^{i+1}$, and if $Q^{i+1} = 0$, then $Q^{i+2} = Q^{i+1}$.

□

Note that this problem with using empirical demand distributions is independent of any difficulties in acquiring reliable data. Even with error-free knowledge of the demand distribution and/or beginning with the demand distribution that will be part of the optimal solution, failing to account for the impact of inventory policy on demand can lead to a suboptimal policy.

5.4 Infinite Customer Base

We will now consider an infinite extension of this problem, while still considering the infinite horizon. We will allow the number of customers to approach infinity while maintaining an arrival rate bounded in expectation. We define a base arrival rate $\lambda = Np$. We then let N go to infinity while keeping λ constant. We can redefine our satisfaction measure to be a proportion of satisfied customers, H (defined on [0,1]). We then get the following expression for demand:

$$D_k = D_k^s(H_l \alpha \lambda) + D_l^u((1 - H_l)\lambda)$$

Where $D_l^i(X)$ is an independent poisson variable with parameter X.

We can define a one-period change in satisfaction ΔS as follows:

$$\Delta S(y, H) = \min\{y, D(H)\} - D^s(H)$$

Where D and D^s are the demand and satisfied demand for a given satisfaction level H.

The results for the finite customer case above still apply in this limit.

Theorem 5.13. *In a situation with an infinite customer base, the optimal policy at steady state is a fixed order policy.*

Proof. Note that the expected number of arrivals for a given time period is bounded by αpN (if $\alpha > 1$) or pN (if $\alpha < 1$). This bound does not change as N approaches infinity (as $pN = \lambda$). Further, note that whether a customer is satisfied or not only changes when the customer arrives, so this bound applies to the absolute value of ΔS.

Consider a time block of size K (that is, a set of K time units), and let $N = K^2$. If we let K approach infinity, clearly N will as well.

In doing so, by the law of large numbers, the expected change in satisfaction fraction H over the time block will be bounded by a number approaching

$$\begin{array}{ll} \dfrac{\alpha\lambda K}{N} & \alpha > 1 \\ \dfrac{\lambda K}{N} & \alpha < 1 \end{array}$$

which approaches 0. Thus, in the limit, if an order quantity $Q(H)$ is optimal for a satisfaction H_k at time k, it will be optimal for H_{k+K} at time $k + K$, as the difference in satisfaction between the two periods will approach 0.

□

Theorem 5.14. *For any fixed order quantity Q, there exists a unique satisfaction level $\bar{H}(Q)$ such that $E[\Delta S(Q, \bar{H}(Q))] = 0$ and:*

$E[\Delta S(Q, H)] < 0$ for $H > \bar{H}(Q)$

$E[\Delta S(Q, H)] > 0$ for $H < \bar{H}(Q)$

Proof. We've seen that $D(H) - D^s(H) = D^u(H)$ is linear and decreasing in H and $D^s(H)$ is linear and increasing, from observation (5.1). Thus, we can write:

$$\Delta S(Q, H) = \min\{Q, D(H)\} - D^s(H) = \min\{Q - D^s(H), D^u(H)\}$$

Both arguments of the minimum clearly decrease in H for a given Q, thus $\Delta S(Q, H)$ is nonincreasing.

Thus, it is sufficient to prove the existence of $\bar{H}(Q)$ for this lemma to hold. This can be proven from the intermediate value theorem, as $E[\Delta S(Q, H)]$ must be nonnegative at $H = 0$, and negative at $H = 1$, and the function is continuous in H.

□

The consequence of this theorem is that over an infinite amount of time, we see that for any order quantity, the satisfaction will drift to a certain value.

Corollary 5.15. *In a situation with an infinite customer base, the steady-state satisfaction of fixed order policies of quantity Q, $\bar{H}(Q)$, satisfies the following condition:*
$E[\min\{Q, D(\bar{H}(Q))\}] = E[D^s(\bar{H}(Q))] = \alpha\lambda\bar{H}(Q)$

That is, the average number of satisfied customers demanding service is the same as the average number of customers receiving service (who then become satisfied).

5.4.1 Optimal Policy

Finding the optimal policy is again a one-dimensional optimization problem. Q^* is the value of Q that optimizes:

$$\phi(Q) = E[r \min\{Q, D(\bar{H}(Q))\} - cQ] \tag{17}$$

5.4.2 Myopic Comparison

The results for the finite fixed-order quantity case apply here as well. Again, we need to define the empirically myopic policy under the assumption that the inventory manager chooses an inventory level based on the demand distribution she observes. Here, any myopic policy $\hat{Q}$ will optimize Q in the following objective:

$$\hat{\phi}_{\hat{Q}}(Q) = E[r \min\{Q, \bar{H}(\hat{Q})\} - cQ] \tag{18}$$

As before, such a quantity need not exist or be unique. By the same argument in Theorem (5.8):

Corollary 5.16. *In the infinite customer base case, if $\hat{Q}$ exists, $\hat{Q} \leq max\{Q^*\}$ if $\alpha > 1$, and $\hat{Q} \geq min\{Q^*\}$ if $\alpha \leq 1$*

Theorem (5.12) also holds, but since the proof relied on the existence of a finite N, the proof for the infinite case is slightly different. Instead, we can use a "softer" upper limit, and simply observe that the empirically myopic order quantity cannot increase infinitely, as the demand is never greater than a poisson variable with mean $\alpha\lambda$, which will have a finite newsvendor optimum.

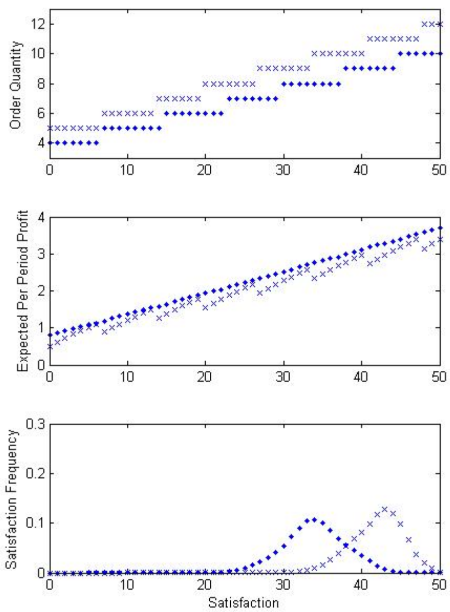

(a) Finite customer base, steady state, positive satisfaction effect. ($N = 50$, $\alpha = 3$, $p = .07$ $r = 1.5$, $c = 1$) Myopic policies are labeled with dots, while optimal policies are labeled with x's. The optimal policy's expected steady-state profit is 7.34% higher than that of the myopic policy.

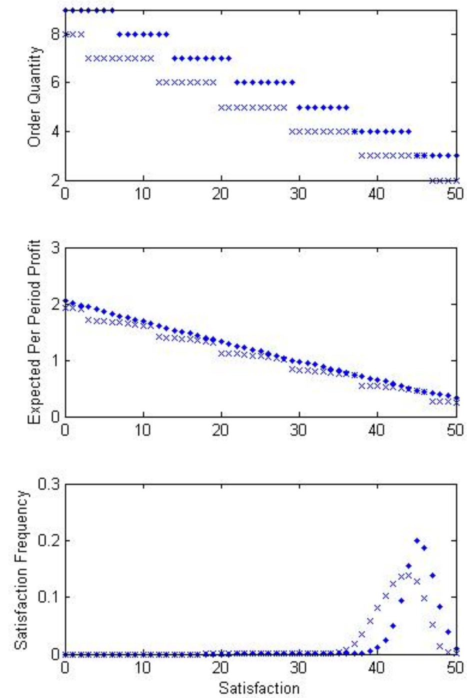

(b) Finite customer base, steady state, negative satisfaction effect. ($N = 50$, $\alpha = 0.3$, $p = .21$ $r = 1.3$, $c = 1$) Myopic policies are labeled with dots, while optimal policies are labeled with x's. The optimal policy's expected steady-state profit is 4.11% higher than that of the myopic policy.

6 Numerical Implementation and Examples

6.1 Dynamic Inventory Control

As noted in Section (5.2.1), optimizing the Dynamic Inventory Control problem involves a straightforward dynamic programming problem. Because the state space is Markovian, the decision space does not expand over time and remains computationally feasible.

The myopic comparison is even simpler, as it involves finding the single period newsvendor solution for each satisfaction level.

Figures 1a and 1b describe the steady-state limit, showing how the myopically suboptimal decision changes the demand distribution to improve the profit over time.

Observe that, in the case of positive α (Figure 1a), for a given satisfaction level, the optimal policy typically orders more than the myopic policy, yielding a lower per period profit. In exchange, the optimal policy generates a higher satisfaction (and thus demand) distribution. The per period profit multiplied by the satisfaction distribution, which is the overall average expected profit, becomes higher. Similarly, for $\alpha < 1$ (Figure 1b), ordering less than the myopic policy yields the same result by creating a lower satisfaction distribution, which is a higher demand distribution.

6.2 Fixed-Order Policies

Optimizing (4) is a finite, discrete, one-dimensional optimization, and the solution can be computed easily. Finding empirically myopic policies that satisfy (5) can be done by testing whether each order quantity satisfies myopic optimality for the distribution it creates.

Theorem (5.12) can be demonstrated visually more effectively in the infinite customer situation below, but it also applies to the finite case. As an example, consider the case presented in figure 1a. If we arbitrarily choose an order quantity, such as 10, we generate a satisfaction distribution, and in turn, a demand distribution. We also generate a profit of 2.60. With this distribution, an order quantity of 8 would generate a profit of 3.10, and is the myopic policy. However, using a fixed order quantity of 8, with the consequent satisfaction and demand distributions, the actual profit is the reduced value of 2.61, because the demand distribution decreases. And even this distribution has a myopic policy of 7, which, if implemented, performs even worse than the old policy of 8 with a profit of 2.49. This is an empirically myopic policy, as 7 is also the myopic policy for the distribution it generates. Meanwhile, the actual optimal policy is at 9, with a profit of 2.65 (though again, if a myopic policy is applied to this distribution, the empirically myopic policy will again become the suboptimal policy of 7).

6.3 Infinite Customer Base

Using the condition in corollary (5.15) and/or the search algorithm suggested in lemma (5.14), the satisfaction for a given order quantity can be estimated analytically. Simulation is also an option.

Finding a myopic fixed order quantity checks the values for agreement with the condition in (18).

In Figures 1 and 2 we see two examples of this process. The top graphs represent the satisfaction levels as a function of order quantity, while the bottom graphs represent the profit on the same order quantity axis. The vertical bars are the myopic order quantities for given satisfaction levels. In both cases, the maximum profit is clearly shifted from the empirically myopic policies, which can be identified as the locations where the satisfaction levels and myopic optima intersect.

Note that in Figure 1, if the satisfaction generated by a non-myopic policy is used to choose an empirical policy, the policy will be closer to a myopic policy, and the process can repeat, in a demonstration of theorem (5.12). If we begin with the optimal policy of ordering 9, the generated satisfaction level of 0.72 will indicate a myopic policy of 7. This, in turn, generates a satisfaction level of 0.59, which indicates a myopic policy of 6. This policy is also the myopic policy for its generated satisfaction level of 0.52, which makes it an empirically myopic policy.

The same does not necessarily hold for the negative satisfaction example in Figure 2. While an initial distribution with a myopic order quantity of 3 will establish an empirically myopic policy, other initial distributions may never settle on a single value. For instance, if the initial distribution is that generated by the optimal order quantity of 2, the myopic order quantity is 4. This, in turn, generates a satisfaction level of 0.95, and the myopic order quantity is again 2, and the cycle will repeat without yielding the empirically myopic policy of 3.

Figure 3 provides an example where no empirically myopic policy exists. Attempts to decide policy empirically will always have to be readjusted, no matter what initial empirical distribution is used.

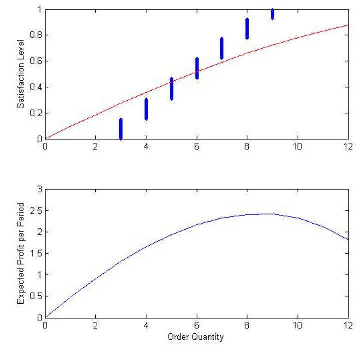

Figure 1: Infinite customer base, steady state, positive satisfaction effect. ($\lambda = 3.5$, $\alpha = 3$, $r = 1.5$, $c = 1$) The function plotted in the top graph represent the steady-state satisfaction values for given order quantities. The vertical bars represent the myopical solutions for given satisfaction levels. The intersections between these two represent empirically myopic order quantities. The function plotted in the bottom graph represents steady-state expected profit as a function of order quantity. The optimal order quantity is at 9, as compared to the highest myopic order quantity at 6, which gives the optimal order quantity a 11.20% higher profit than the best empirically myopic policy.

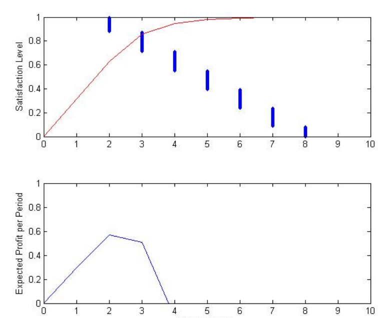

Figure 2: Infinite customer base, steady state, negative satisfaction effect. ($\lambda = 10.5$, $\alpha = 0.3$ $r = 1.3$, $c = 1$) The optimal order quantity is at 2, as compared to the empirically myopic order quantity at 3, which gives the optimal order quantity a 12.67% higher profit than the best myopic policy.

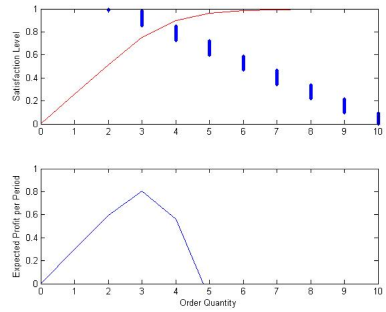

Figure 3: Infinite customer base, steady state, negative satisfaction effect. ($\lambda = 13$, $\alpha = 0.3$ $r = 1.3$, $c = 1$) An example of a problem without an empirically myopic policy.

7 Conclusion

This paper has used a specific customer behavior model to examine and quantify the importance of considering future demand when making inventory policies. The Binary Customer Satisfaction Model, while simple, uses discrete decisions and is applied to both continuous and discrete state spaces. Customer behavior in this model is both intuitive and easy to communicate to managers who may not be overly familiar with the jargon of operations research. Its simplicity also provides it the flexibility to function in many different problem environments, and the intuitive behavior of its customers can be easily extended to add additional intuitive properties, such as multiple levels of satisfaction, or communication between customers, without predetermining the aggregate properties of the demand base.

In the context of the multiperiod newsvendor problem, we've shown that when even the simplest customer behavior models respond to customer experience with the inventory system, full knowledge of demand is insufficient to generate optimal policies. Empirical policies, no matter how efficiently data is gathered, will not be sufficient to find optimal policies if demand is treated as external to the system. The response of customers to the inventory policy itself must be studied and factored into decisions.

This model can be extended to add additional behavioral properties to the customers, such as different classes of customer or communication between customers, or applied to other problems, such as those including competition or supply chains, or even different classes of problems, such as queueing or facility location. We have also assumed full information, an assumption which can be relaxed to study how the need for learning can affect these problems.

References

[1] G. J. Caine, R. H. Plaut. Optimal Inventory Policy when Stockouts Alter Demand. *Nav. Res. Logist. Quart* 23 (1976) 1-13.

[2] Cooper, W. L., T. Homem-de-Mello, A. J. Kleywegt. Models of the spiral-down effect in revenue management. *Oper. Res.* 54(5) (2006) 968-987.

[3] R. Ernst, M. Cohen. Coordination Alternatives in a Manufacturer/Dealer Inventory System Under Stochastic Demand. *Production Oper. Management* 1 (1992) 254-268.

[4] R. Ernst, S. Powell. Optimal Inventory Policies Under Service-Sensitive Demand. *Eur. J. Oper. Res.* 87 (1995) 316-327.

[5] R. Ernst, S. Powell. Manufacturer Incentives to Improve Retail Service Levels. *Eur. J. Oper. Res.* 104 (1998) 437-450.

[6] V. Gaur, Y-H. Park. Asymmetric Consumer Learning and Inventory Competition. *Management Science*, 53(2) (2007) pp. 227-240.

[7] J. Hall, E. Porteus. Customer Service Competition in Capacitated Systems. *Manufacturing Service Oper. Management* 2 (2000) pp. 144-165.

[8] Thomas W. Hill Jr. A Simple Perturbed Demand Inventory Model with Ordering Cost. *Management Sci* 23(1) (1976) pp. 38-42.

[9] L. Li. The Role of Inventory in Delivery-Time Competition. *Management Sci.* 38(2) (1992) 182-197.

[10] G. Liberopoulos, I. Tsikis, G. Kozanidis. Inventory Control in a Duopoly: A Dynamic Non-Cooperative Game-Theoretic Approach. Working Paper, Department of Mechanical and Industrial Engineering, University of Thessaly. 2005.

[11] T. L. Olsen, R. P. Parker, Inventory Management Under Market Size Dynamics. *Management Sci.* 54(10) (2008) 1805-1821.

[12] L. W. Robinson. Appropriate Inventory Policies When Service Affects Future Demand. Working Paper #88-08, Johnson Graduate School of Management, Cornell University. 1991.

[13] B. L. Schwartz. A New Approach to Stockout Penalties. *Management Science* 12(12) (1966) pp. B538-B534.

[14] B. L. Schwartz. Optimal Inventory Policies in Perturbed Demand Models. *Management Sci* 16(8) (1970) pp. B509-B518.

[15] Shanthikumar, J.G. and Yao, D.D. Stochastic Convexity and Its Applications in Parametric Optimization of Queueing Systems. *Proceedings of IEEE 27th Conference on Decision and Control.* Austin, TX, 1988, 657-662.

A Optimality of Full Satisfaction

Throughout this paper, we have assumed that the inventory manager is satisfying all demand possible when it is realized. While in a single period newsvendor (and hence under the direction of the myopic manager), this choice is obvious, its applicability to a case where demand satisfaction in the current period can affect future demand distributions is not trivial.

Lemma A.1. *Given an inventory level y_k and demand D_k, it is always optimal to allocate $\nu_k = \min\{y_k, D_k\}$.*

For the case when $\alpha \geq 1$, this result is obvious. Since such an allocation maximizes current period profit and maximizes future satisfaction and hence demand, profit is maximized.

For the case when $\alpha < 1$, the result becomes less obvious, as while current period profit is maximized, future satisfaction is maximized which minimizes future demand. However, forgoing current profit for future demand is never profitable over the long term, as each sale forgone increases the number of unsatisfied customers by 1, which will never increase future demand by more than 1 for a single period, and the lost profit will not necessarily be recouped, and cannot be surpassed.

B Sample Path Stochastic Orders and Coupling

Shanthikumar and Yao [15] provide a definition for stochastic orders in a sample path sense for random variables with integer parameters using a coupling approach. As used in this paper, we define a variable X parameterized by integer parameter θ to be stochastically ordered in the following way.

Definition B.1. *$X(\theta)$ is increasing and concave in a sample path stochastic sense if, for any parameters θ^1, θ^2, θ^3, and θ^4 such that:*

$$\begin{gathered}\theta^1 \leq \theta^2 \leq \theta^4 \\ \theta^1 \leq \theta^3 \leq \theta^4\end{gathered}$$

There exists a probability space with variables $\hat{X}(\theta^i) =^d X(\theta^i)$ such that the following holds true almost surely:

$$X(\theta^2) - X(\theta^1) \geq X(\theta^4) - X(\theta^3) \geq 0 \tag{19}$$

In other words, there exists a sample path under which variables with the same distribution can be coupled such that the increasing convex property holds almost surely. Other orders are defined similarly, simply by changing the directions of the inequalities in (19). For example:

$$\begin{aligned} X(\theta^2) - X(\theta^1) = X(\theta^4) - X(\theta^3) \leq 0 & \qquad \text{linear and decreasing} \\ X(\theta^2) - X(\theta^1) \leq X(\theta^4) - X(\theta^3) & \qquad \text{convex} \end{aligned}$$