The Binary Customer Satisfaction Model in Inventory and Queueing Systems

By

Justin Sepehr Azadivar

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Industrial Engineering and Operations Research

in the

Graduate Division

of the

University of California, Berkeley

Committee in Charge:

Professor Zuo-Jun Max Shen, Chair
Professor J. George Shanthikumar
Professor Xuanming Su

Spring 2010

Abstract

The Binary Customer Satisfaction Model in Inventory and Queueing Systems

By

Justin Sepehr Azadivar

Doctor of Philosophy in Engineering - Industrial Engineering and Operations Research

University of California, Berkeley

Professor Zuo-Jun Max Shen, Chair

This dissertation introduces the Binary Customer Satisfaction Model for addressing logistics issues. In typical logistics problems, the arrival of customers through a demand process is considered external to the management decisions. In practice, it is typically the case that customers will respond to changes is service policy by changing their behavior. The Binary Customer Satisfaction Model provides a simple customer behavior model that directly interacts with the service policy and provides for analysis of managerial insights.

The Binary Customer Satisfaction Model assigns customers to one of two satisfaction states. Satisfied customers have one demand rate, while unsatisfied customers have a different demand rate. The Binary Customer Satisfaction Model accommodates situations where satisfied customers demand more, as well as those when satisfied customers demand less, a possibility undertreated in existing literature. Satisfaction changes for each customer when the customer demands service. Satisfaction occurs when the customer receives service, while unsatisfaction occurs when the customer attempts to receive service but is unable to.

The Binary Customer Satisfaction Model is generally applicable to a wide range of logistics problems. In this dissertation, we consider the application of the model to the newsvendor inventory problem as well as an M/M/s queueing model without a buffer. We also briefly consider extensions to these models and how the Binary Customer Satisfaction Model can inform management of these extended cases.

Key to these insights is a well-defined concept of myopic management policies. This dissertation defines myopic policies in such a way to allow explicit comparison between optimal and myopic policies, and quantitatively present the value of considering the effect of service policy on future customer behavior.

In both the inventory and queueing contexts, we find that the Binary Customer Satisfaction Model gives two major insights. The first confirms the intuitive result that, if satisfied customers are more likely to arrive, then it is worthwhile for a manager to provide a level of service that would appear too high to a myopic manager, as future increases in customer demand will offset the additional cost. Similarly, if satisfied customers are less likely to arrive, the manager should prepare a lower level of service.

The second main insight is that it is not enough to simply observe the demand to find an optimal policy when customer behavior depends on service policies. Even with full knowledge of the demand, and beginning with the optimal demand level, a myopic manager will choose a suboptimal service policy, which will in turn create a suboptimal policy, which will cause the myopic manager to move further away from optimality. This spiraling effect makes clear the importance of not blindly making policies based on empirical observation of demand.

Contents

Acknowledgments

I want to thank my advisors Max Shen and George Shanthikumar for their near-infinite patience in enabling my work.

CHAPTER 1

Overview

Typical logistics problems are framed as a manager in control of certain parameters, who must optimize those decisions in response to variables beyond the manager's control in order to optimize some objective. However, in multiperiod problems, it is often the case that the variables for future periods can be affected by the decisions in a given period. A simple optimization of the immediate objective in each period may lead to inferior results in future periods. More importantly, if this possibility is not considered, the manager will never recognize that she is making suboptimal decisions, as the objective is still optimized in each period.

In this chapter, we will introduce the Binary Customer Satisfaction Model in its general form in Section 1.1. We then discuss the method we use for quantitatively calculating the value of considering the impact of policies on future customer behavior in Section 1.2. We close the chapter with a discussion of some of the general results we observe, regardless of problem, in Section 1.3.

Chapter 2 conducts a literature review of some of the related work in the relationships between management policy and customer behavior. The section also highlights the specific contributions of this dissertation which are not present in most other literature.

The remainder of this dissertation will look at specific examples of problems involving the Binary Customer Satisfaction Model, and providing technical analysis. Chapter 3 applies the Binary Customer Satisfaction Model to the multiperiod newsvendor stochastic inventory control problem. Chapter 4 addresses using the Binary Customer Satisfaction Model in the M/M/s queueing context.

1.1. The Binary Customer Satisfaction Model

To illustrate this, we introduce the Binary Customer Satisfaction Model. The Binary Customer Satisfaction Model is a simple customer behavior model that gives rise to complex and interesting analytical results. The model can be summarized as follows:

1) There are a number of customers who each have the same behavior model.

2) Each customer is in either a “satisfied” state or an “unsatisfied” state.

3) The probability or rate at which a customer will demand service depends only on whether the customer is satisfied or unsatisfied.

4) The satisfaction state of a customer will not vary unless she demands service.

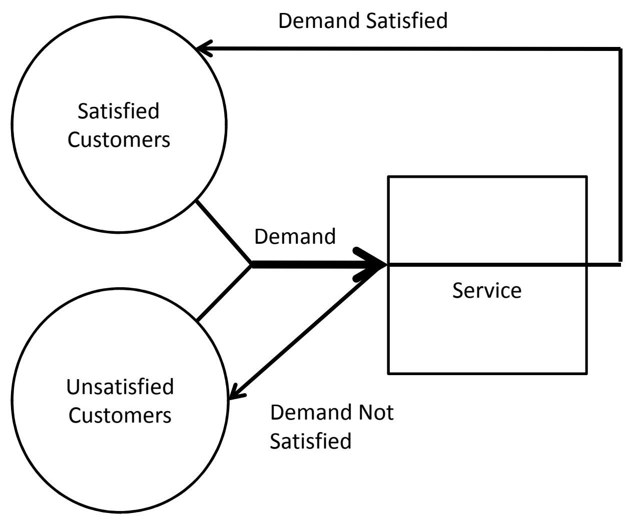

FIGURE 1.1. Binary Customer Satisfaction Overview

5) If a customer demands service, but is unable to receive service, the customer becomes unsatisfied.

6) If a customer demands service and receives it, the customer becomes satisfied.

See Figure 1.1 for a graphical description of this model.

The general structure of this model is not dependent on the specific problem being addressed. That is, a Binary Customer Satisfaction Model can be applied to the demand behavior of many different logistics problems. In this paper, we will apply it to simple inventory and queueing models.

This model has a number of advantages over more complex customer behavior models. The customer behavior is simplistic, but still sensible. Customer behavior in this model is not necessarily strategic, but real customers often do not behave in a strategic manner, either. More importantly, the behavior described by the Binary Customer Satisfaction Model is easy to communicate to a manager who has little Operations Research training. The customer behavior makes intuitive sense, and requires no special jargon or training to understand.

Finally, because the customer behavior is modeled on an individual basis, and simply, extending the model to include other customer behavior is easier than in some models which attempt to model large groups of customers in the aggregate. In many cases, we understand intuitively how customers may respond to certain situations, but lack a solid understanding of how that behavior gives rise to larger effects in the customer base as a whole. Modeling this behavior on an individual basis for analysis provides a means to predict the aggregate behavior it will cause.

These models can be grouped into "positive satisfaction" and "negative satisfaction" models. The distinction describes whether satisfaction makes customers more likely or less likely to demand respectively.

Positive satisfaction models are those in which customers who are satisfied are more likely to demand than less satisfied customers. Such models represent such factors as customer loyalty, where customers who receive satisfying experiences are more likely to demand from the manager in order to repeat positive experiences. Meanwhile, customers who fail to receive service are inclined to seek service elsewhere in the future. This model is best-suited to situations where customers are likely to need service on a routine basis unrelated to their previous service. For example, in the classic newsvendor problem, whether or not a customer reads a newspaper does not impact her need for news the next day, but it may impact from which source the customer seeks news.

When satisfied customers are less likely to demand service than unsatisfied customers, we have a negative satisfaction model. The negative satisfaction model is novel, as few papers consider the possibility that satisfied customers are less likely to demand service, but there are real scenarios which suggest the use of a negative satisfaction model. In many forms of service, providing the service obviates the need for additional service for some time. Failing to receive this service leaves the customer with a continuing need for service. As an example of this, consider a server who repairs machines that fail. A broken machine may incline the operator to seek service from the server. If that service is received, the operator does not need service for some time. On the other hand, if the operator fails to receive service, he is likely to seek service again, as her machine is still broken.

These two groups of models are described separately because, though they are modeled in an identical fashion, they represent different types of situations, and have different properties.

1.2. Myopic Comparisons

A primary contribution of this work is to define the means to quantify the value of considering the effect of current management decisions on future customer behavior. Within each model, we assume that the Binary Customer Satisfaction Model accurately represents customer behavior, and compare the effectiveness of two different policies. The first policy is the optimal policy, which optimizes the objective not only in the present, but also the future. We will often consider optimal policies to be the ones which generate the best objective in the steady-state.

We then compare this policy to what we refer to as "myopic policies." As the name suggests, the myopic policy is the policy that a myopic manager would make. Myopic managers do not consider the impact their management decisions will have on future customer behavior. Instead, they assume that the current customer behavior is the fixed customer behavior, as is done in traditional logistics problems.

We can then observe the values of the objectives for the myopic and optimal policies. Comparing these policies gives us a way to quantify the value of considering future customer behavior. If the value is sufficiently large, then it is a sign that the manager should consider investing resources into determining this effect so that it can be accounted for in the decision-making process.

The generation of a myopic comparison is somewhat complicated by the dynamics of these models. Since we are assuming the accuracy of the Binary Customer Satisfaction Model, even though the myopic manager is assuming that customer behavior is fixed, customer behavior may, in fact, change. The manager is myopic, not blind, and should not simply continue using the earlier policy. When a manager can change her policy in every period, this is not a concern, as the manager will simply choose the policy that optimizes the objective in each period. However, when the manager is not completely free to change the policy over time, a more nuanced definition of a myopic policy is required.

We introduce the idea of an "empirically myopic policy." An empirically myopic policy represents the ability of a manager to observe, empirically, the demand distribution over a significant period of time, and choose a policy accordingly. If the manager chooses a policy that affects the demand distribution over the long term in a way such that the policy is no longer optimal, an empirical study of the demand will convince the manager to change the policy. The manager will continue to adjust the policy until the demand her policy generates has a myopic solution equal to the policy already in use. We refer to such a policy as empirically myopic, and it forms the comparison with the optimal policy.

1.3. General Results

A number of general results are true of the Binary Customer Satisfaction Model, independent of the specific problem they are applied to. The precise details of how the result manifests varies from model to model, but the nature of the results remains the same.

1.3.1. Managers Should Over Serve (Positive Satisfaction) or Under Serve (Negative Satisfaction) Compared to Myopic Policies. The first general result is an intuitive one, which suggests that optimal managers will provide different levels of service than myopic managers to account for the impact of service level on future demand. Because a service level not only provides profit in a given period, but controls future demand, the impact of the decision on future demand will pull the optimal solution in one direction or the other.

In the case of positive satisfaction, the optimal manager should be prepared to provide more service than the myopic manager. Doing so incurs an immediate cost associated with retaining additional service capacity, but accommodating additional demand will lead to a larger number of satisfied customers. These customers, in turn, are more likely to demand service in the future, creating a larger demand distribution. This larger demand distribution will generate additional profit to offset the service capacity cost.

For negative satisfaction models, the optimal manager does not need to keep as much service capacity available as the myopic manager. Reducing service capacity reduces immediate profit, as fewer customers are served in the near-term, but this leaves more customers unsatisfied, and unsatisfied customers demand more in a negative satisfaction model. This tradeoff between higher demand but lower capacity is a way of minimizing the cost of variability. Since a customer who fails to receive service is more likely to come back sooner, keeping

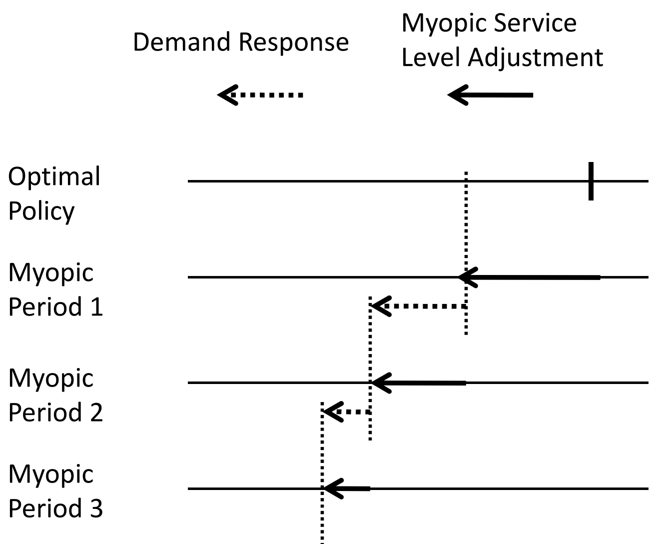

FIGURE 1.2. Iterative Empirically Myopic Policies (Positive Satisfaction)

capacity low is reasonable, as demand spikes do not need to be immediately accommodated, as those customers will continue to provide demand until they receive service.

1.3.2. Empirically Myopic Managers Drift Away from Optimality. The second, and more interesting, general result is that even with full understanding of the demand distribution, a myopic manager who uses empirically myopic policies will not only fail to reach optimality, but will in fact move away from optimality. This creates a spiraling effect, where attempts to accommodate changes in demand distribution actually exacerbate the problem.

The effect is most obvious in the case of positive satisfaction models. As noted in Section 1.3.1, the optimal manager will over serve compared to the myopic manager. If a myopic manager took over from an optimal manager, she would observe that the service level provided by the optimal manager was too high for the observed demand distribution, and would reduce the service level. As a result of the reduced service level, the number of satisfied customers would decrease, which will decrease demand in a positive satisfaction model. Observing this reduction of demand, the myopic manager would again reduce the service level, which will again reduce demand, and so on, until the reduction of demand is small enough that the service level does not need to be decreased to be myopically optimal, yielding an empirically myopic solution. See Figure 1.2, where in each review period, the

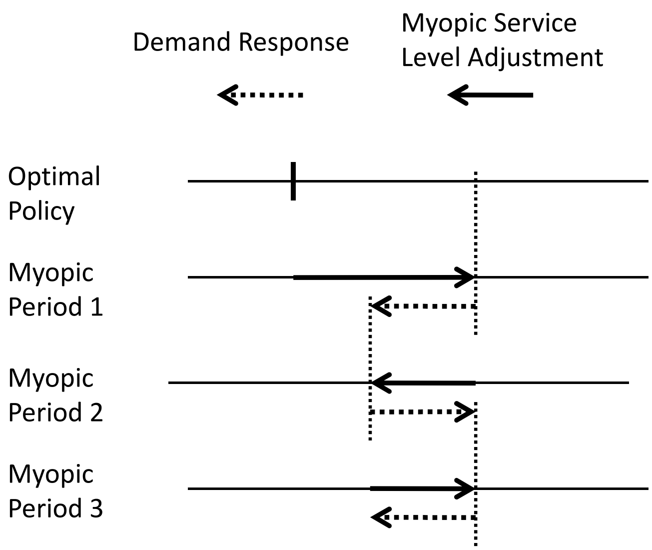

FIGURE 1.3. Iterative Empirically Myopic Policies (Negative Satisfaction)

myopic manager changes the service level in order to match the demand level of the previous period.

Not only does this myopic policy fail to provide the optimal profit, but without an understanding of the way customer behavior is affected by service decisions, the manager will never know that there is a problem. The policy provides the best profit in each period, given the empirically-observed demand distribution, which is enough to satisfy the myopic manager, as well as most observers.

In negative satisfaction models, the behavior is somewhat different. When a myopic manager takes over from an optimal manager, she will want to increase the service level to meet the demand. Doing so will increase the satisfaction level, which has the effect of decreasing the demand. The myopic manager would then want to decrease the service level to match this change, which will decrease satisfaction and increase demand. Because the change in demand is the opposite direction of the change in service level, the myopic manager will find herself alternating between high and low service levels, maybe (but not necessarily) settling on an empirically myopic policy. Figure 1.3 displays this concept graphically, as the manager must move back and forth to keep up with changing demand.

CHAPTER 2

Related Literature

In 1966, Schwartz [**14**] 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 [**15**] he addresses the creation of optimal policies on PD models. This was extended by Hill [**9**] 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. 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 more to the modeling approach than the final model.

An example of an aggregate demand model is Ernst and Cohen [**4**], 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 [**5**] 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 [**6**] 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 [**13**] 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 [**10**] 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 [**8**] considers a more explicit customer behavior model in the context of 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 decreasing 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 [**7**] 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. [**11**] 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 [**12**] 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.

Other models try to capture the behavior of customers who do not demand product, but make decisions based on the results of other customers who demand product. This "herding" behavior, if understood, can be manipulated to improve inventory management results. For example, Debo and Ryzin [**3**] look at the balance between controlling customer behavior by using stockouts to increase customer willingness to buy and maintaining the inventory to satisfy the increased demand. Since controlling customer behavior occurs using the same mechanism that capitalizing on customer behavior does, the inventory decision must be carefully tailored.

This paper adapts certain aspects of these modeling approaches and applies them to a discrete-customer, discrete-demand model, and find that sequential decision-making without full information can move systems away from optimality. We also discuss the possibility that satisfaction can actually reduce demand, which is not considered in most existing literature.

CHAPTER 3

Inventory Control

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 decision maker.

In this chapter, 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. Doing so will allow the newsvendor to ensure that demand remains high in future periods, improving profit.

Similarly, a service provider which holds a near-monopoly on ability to satisfy a customer's need can often have less capacity than may otherwise be indicated by optimizing a newsvendor problem, as the customer will have to return in the near future to demand service again. It is less necessary to be prepared for a period with an unexpectedly high demand.

The Binary Customer Satisfaction Model presented here is a straightforward 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 can be easily communicated and augmented 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 should sometimes order more or less than a newsvendor purchasing the optimal number to meet demand in the current period. Doing so allows the newsvendor to sacrifice profit in the current period in order to increase demand in future periods. 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 "spiraling" effect similar to (if less dramatic than) the one described in the 2006 study by Cooper et. al. [**2**]

Section 3.1 introduces the Binary Customer Satisfaction Model and its application to the newsvendor problem. It is applied to three variants on the problem, the general dynamic case (Section 3.2), the fixed-order quantity case (Section 3.3) with both positive satisfaction and negative satisfaction effects, and the case presented by approaching the infinite limit of the customer base (Section 3.4). Numerical examples for each variant that illustrate the impact of considering customer satisfaction are described in Section 3.5. Section 3.6 concludes the chapter with consideration of potential avenues for extension.

3.1. 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. [**11**]. There are only two credibility levels, but in our model, 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. We also allow for the possibility that satisfaction of demand can reduce future demand.

We begin by presenting the details of the model itself. We will then discuss optimization of the model under three different problem formulations. In all three cases, we compare the optimal behavior to the behavior of a decision maker who does not consider the effect of satisfaction on demand, and find that the value of demand manipulation can, in the long term, offset the cost of over or under ordering from the newsvendor optimum.

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. We refer to this as a "positive satisfaction effect." 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. This is referred to as a "negative satisfaction effect." 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 (i.e. there is no inventory carryover from period to period), 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,

$$D_k = D_k^s(S_k, \alpha p) + D_k^u(N - S_k, p) \tag{3.1}$$

where $D_k^i(x, z)$ are independent binomial 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 (for an explanation of this assumption, and why it is nontrivial, see Appendix A), 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}.

$$S_{k+1} = S_k - D_k^s + \nu_k \tag{3.2}$$

The following lemma describes first and second order stochastic ordering properties of this model that are used throughout the chapter.

LEMMA 3.1.1. *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 [**16**] for a definition of second-order stochastic properties for integer variables. Refer to appendix B for a detailed discussion of coupling and the sample path stochastic order as it is used in this paper.

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 use the following coupling for these four paths:

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

B) The remaining two 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.

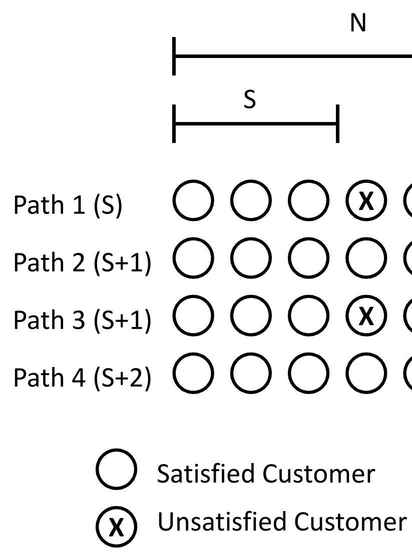

FIGURE 3.1. Example of Path Coupling for Lemma 3.1.1

C) Neither of these two customers 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.

This coupling is summarized in Figure 3.1. Each column represents one customer in four different sample paths. A single uniform [0,1] random variable determines the behavior of that customer on all four paths. If the customer is satisfied in a given path, the customer arrives if that random variable is less than αp. If unsatisfied, the customer arrives if the variable is less than p.

We then find the value of the difference of the variables in path 2 and 1, and the value of the difference of the variables in path 4 and 3:

$$D_k^{s2} - D_k^{s1} = D_k^{s4} - D_k^{s3} = I(U_1 < \alpha p) \tag{3.3}$$

$$(S + 1 - D_k^{s2}) - (S - D_k^{s1}) = (S + 2 - D_k^{s4}) - (S + 1 - D_k^{s3}) = 1 - I(U_1 < \alpha p) \tag{3.4}$$

$$D_k^{u2} - D_k^{u1} = D_k^{u4} - D_k^{u3} = -I(U_1 < p) \tag{3.5}$$

$I(x)$ is an indicator variable which equals 1 if its argument is true, and 0 otherwise.

For (3.3) and (3.4), we see that the differences of the variables are equal and nonnegative in each case. In (3.5), the differences are equal and nonpositive. As described in Appendix B, this proves linearity and the increasing or decreasing property.

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

In this case, we again get linearity because the differences are equal. Whether this difference is nonnegative, nonpositive, or zero depends on the value of α, which proves the final result of this lemma.

□

Having defined the model, we now examine how a manager should make inventory decisions in order to maximize profit. We also compare this behavior with a manager who makes myopic decisions based solely on the observable demand without considering the effect of the inventory decision on future demand, and compare the two policies, providing a way to calculate the value of considering the satisfaction effect on future demand. We will consider three optimization schemes for this model. The first is a dynamic inventory control problem (Section 3.2), in which a manager can change the order quantity each period. The second considers fixed-order policies (Section 3.3), where the manager must commit to a fixed order quantity over the long term. Finally, we examine the case when the number of customers approaches infinity (Section 3.4). In each case, we will see that using a myopic policy can result in a demand distribution that provides an inferior profit, even if the manager has full information about that distribution. Using an optimal policy, while providing less profit for a given period, allows the manager to keep the demand distribution higher and offset that loss.

3.2. Dynamic Inventory Control

First we consider a dynamic inventory model, where the manager can change the order quantity each period after observing the results of the previous period. We find that the optimal manager will order more (for a positive satisfaction effect) or less (for a negative satisfaction effect) than the myopic manager in order to create a more favorable demand for future periods.

The solution to a single period newsvendor problem with known demand is well-studied. To consider the impact that customer satisfaction has on inventory control decisions, we consider the multiperiod newsvendor problem with the Binary Customer Satisfaction Model described above. 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:

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

For notational simplicity, we have assumed that values are not discounted over time. This assumption is not necessary for the results in this paper, as stochastic and comparative relationships still hold even when values are discounted.

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 (see equation 3.2). 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 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 $Q(S)$.

It is not necessarily the case that a given ordering policy will yield a unique steady-state result. For example, a policy of ordering 0 for satisfaction levels below $N/2$, and N otherwise will, depending on the initial state, yield a steady state distribution for S of either 0 or N. If a unique steady-state result does exist, we can refer to the steady-state distribution of satisfaction level by $\mathbf{S}(Q)$, 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 (see equation 3.1) and I is an indicator variable that returns 1 if its argument is true, and 0 otherwise.

We can then define the expected profit:

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

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

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

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 $-\infty$.

LEMMA 3.2.1. *$G_k(S)$ 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 lemma 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 Lemma 3.1.1, and $min\{D(S), y\}$ increases in S in expectation as well. Thus, the cost-to-go function is also nondecreasing. ☐

LEMMA 3.2.2. *$G_k(S)$ is nonincreasing in S when $\alpha < 1$.*

PROOF. Again, for $k = K$, this is true, as increased satisfaction means decreased demand distribution by Lemma 3.1.1, 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.

☐

3.2.2. Myopic Comparison. Because we want to determine the value of considering satisfaction effects, we will define a myopic policy which ignores such effects and compare its effectiveness with the optimal policy described in Section 3.2.1. The myopic policy is the ordering policy which optimizes the profit in the current period 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, $\hat{Q}(S)$:

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

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:

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

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

We now prove the intuitive results that, when satisfaction can affect future demand, optimal managers will tend to order more than myopic managers when satisfied customers are more likely to arrive, and order less when they are less likely to arrive.

THEOREM 3.2.3. *When $\alpha > 1$, for any myopic policy $\hat{Q}_k(S)$, there exists an optimal policy $Q_k^*(S)$ 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 policies will be identical, 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 Lemma 3.2.1, 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 Lemma 3.2.1). Thus, if $\hat{Q}_k(S)$ is not optimal, then the optimal order quantity must be higher. □

Theorem 3.2.4. *When $\alpha < 1$, for any myopic policy $\hat{Q}_k(S)$, there exists an optimal policy $Q^*_k(S)$ such that $\hat{Q}_k(S) \geq Q^*_k(S)$.*

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

3.3. Fixed-Order Policies

We now consider a problem where the ordering policies are restricted to fixed-order policies. Contracts or logistics issues may limit the decision space of a manager to policies which order the same amount every period. It is in this context we find that even with perfect information about the demand distribution, a manager who does not consider the effect of satisfaction on demand may gradually move away from an optimal solution.

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, though generally concave for deterministic functions, does not necessarily have that property with stochastic variables.

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

$$\phi(Q) = E[r \min\{Q, \mathbf{D}(Q)\} - cQ] \tag{3.6}$$

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

As before, we want to create a myopic policy for comparison in order to provide a method to quantitatively define the value of considering the impact of customer satisfaction in choosing an order quantity. However, we find that creating a myopic comparison is not as straightforward as it was in the dynamic inventory control 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 chapter 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:

$$\hat{\phi}_{\hat{Q}}(Q) = E[r \min\{Q, \mathbf{D}(\hat{Q})\} - cQ] \tag{3.7}$$

$\hat{\phi}_{\hat{Q}}(Q)$ is the expected profit of order quantity Q given demand distribution $\mathbf{D}(\hat{Q})$. This differs from (3.6) 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 a myopic optimum. We formally define such policies as follows:

Definition 3.3.1. *A fixed-order policy $\hat{Q}$ is empirically myopic if it satisfies the following condition:*

$$\hat{Q} = F^{-1}_{\boldsymbol{D}(\hat{Q})}(1 - \frac{c}{r}) \tag{3.8}$$

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

Because the positive and negative satisfaction cases differ significantly in terms of analysis, we will consider each case separately. In both cases, we first describe the effect of the choice of order quantity on the satisfaction distribution. We then observe how the optimal policy varies from empirically myopic policies.

3.3.1. 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 3.3.2. *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 $l \geq k$.*

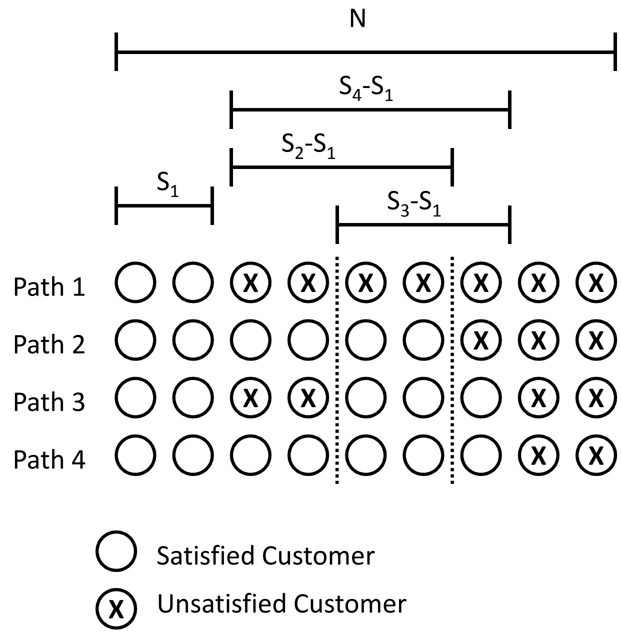

FIGURE 3.2. Example of Path Coupling for Theorem 3.3.2

PROOF. As with Lemma 3.1.1, we again use Shanthikumar and Yao [**16**] to define the stochastically nondecreasing and concave property. See Appendix B for more details.

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 $l + 1$.

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 $Q + 2$.

Because we have assumed that S_l is concave and nondecreasing, paths i ($i = 1, 2, 3, 4$) have satisfaction at time l of S_l^i such that:

$$\begin{aligned} S_l^1 &\leq S_l^2 \leq S_l^4 \\ S_l^1 &\leq S_l^3 \leq S_l^4 \\ S_l^2 + S_l^3 &\geq S_l^1 + S_l^4 \end{aligned} \tag{3.9}$$

The four paths are coupled as follows:

A) All customers but $S_l^4 - S_l^1$ of them 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.

Figure 3.2 summarizes this coupling visually. The dotted lines denote a period of overlap between satisfied customers in paths 2 and 3. This overlap is the amount by which $S_l^2 + S_l^3$ exceeds $S_l^1 + S_l^4$, and serves as a measure of concavity. This overlap does not exist when the satisfaction is linear, and represents the extent to which satisfaction can become less concave.

From Lemma 3.1.1, we recall that $S_l - D_l^s$ is linear and increasing in S_l. Taking the differences in the satisfaction update equations (3.2) for these paths gives us:

$$S_{l+1}^2 - S_{l+1}^1 = \sum_{i=1}^{S_l^2 - S_l^1} I(U_i > \alpha p) + \min\{D_l^2, Q+1\} - \min\{D_l^1, Q\} \tag{3.10}$$

$$S_{l+1}^3 - S_{l+1}^4 = -\sum_{i=1}^{S_l^4 - S_l^3} I(U_i > \alpha p) + \min\{D_l^3, Q+1\} - \min\{D_l^4, Q+2\} \tag{3.11}$$

Adding these together while assuming (3.9) gives us

$$\begin{aligned} &S_{l+1}^2 + S_{l+1}^3 - S_{l+1}^1 - S_{l+1}^4 = \\ &\sum_{i=S_l^4 - S_l^3 + 1}^{S_l^2 - S_l^1} I(U_i > \alpha p) + \min\{D_l^2, Q+1\} + \min\{D_l^3, Q+1\} - \min\{D_l^1, Q\} - \min\{D_l^4, Q+2\} \end{aligned} \tag{3.12}$$

If this equation is nonnegative almost surely under this coupling, then the S_{l+1} is concave.

Recall that demand is decreasing and linear in S, so the demand comparisons are opposite the satisfaction comparisons in (3.9), and will hold almost surely under this coupling. For example, path 1, having the lowest satisfaction, will always have the highest demand.

We now consider six different cases and find that (3.12) is nonnegative in each case.

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, and the last term is less negative, 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 (3.12), we find:

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

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 (3.3.2). To prove the nonincreasing aspect, simply observe that (3.10) is always nonnegative, and (3.11) must be nonpositive (as the differences in demands must be no greater than the difference in the number of satisfied customers not demanding, and paths 2 and 4 have higher inventory levels than paths 1 and 3, respectively).
□

We can extend this result to the steady-state distribution, as well.

COROLLARY 3.3.3. *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 then use the nondecreasing property from Theorem 3.3.2 to reach a fixed-order analogue of Theorem 3.2.4 showing a need to order no more, and possibly less, than the myopic manager to reach optimality.

THEOREM 3.3.4. *When $\alpha < 1$, 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 (3.6):

$$E[r \min\{Q^*, \mathbf{D}(Q^*)\} - cQ^*] \geq E[r \min\{\hat{Q}, \mathbf{D}(\hat{Q})\} - c\hat{Q}] \tag{3.13}$$

Since $\hat{Q}$ optimizes (3.7):

$$E[r \min\{\hat{Q}, \mathbf{D}(\hat{Q})\} - c\hat{Q}] \geq E[r \min\{Q^*, \mathbf{D}(\hat{Q})\} - cQ^*] \tag{3.14}$$

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

$$E[r \min\{Q^*, \mathbf{D}(\hat{Q})\} - cQ^*] \geq E[r \min\{Q^*, \mathbf{D}(Q^*)\} - cQ^*] \tag{3.15}$$

The right hand side of (3.13) equals the left hand side of (3.14), The right hand side of (3.14) equals the left hand side of (3.15), and the right hand side of (3.15) equals the left hand side of (3.13). Thus all three quantities are equal, and $\hat{Q}$ must also be optimal.

□

3.3.2. Positive Satisfaction. We now consider the case where $\alpha > 1$.

THEOREM 3.3.5. *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. The update equation (3.2) for S_{l+1} can be divided into two parts. The first part, $S_l - D_l^s$, is nondecreasing in S_l as noted in Lemma 3.1.1. The second part, $min\{Q, D_l\}$, is clearly increasing in Q. When $\alpha > 1$, D_l also increases in S_l, which is nondecreasing in Q by assumption, and the minimum of two nondecreasing functions is nondecreasing. Since both terms are nondecreasing in Q, S_{l+1} is nondecreasing in Q. □

Note that this theorem is weaker than Theorem 3.3.2, as it does not prove second-order properties. Because both the demand and the order quantity increase together, second-order properties that apply to deterministic functions do not necessarily translate into second-order properties for stochastic variables.

We can again extend this result to the steady-state situation.

COROLLARY 3.3.6. *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 3.3.4 also holds.

THEOREM 3.3.7. *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 to that of Theorem 3.3.4, 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 3.3.8. *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 (3.8), 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 average profit is always c lower than a policy of N. As a result, at some point, Q^{i+1} will not be larger than Q^i.

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.

3.4. Infinite Customer Base

We will now consider the situation where the number of customers approaches infinity, 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 approach infinity while keeping λ constant. We can redefine our satisfaction measure to be a proportion of satisfied customers, H (defined on [0,1]). Because binomial random variables allowed to approach infinity in this way are Poisson random variables, we get the following expression for demand:

$$D_k = D^s_k(H_k \alpha \lambda) + D^u_k((1 - H_k)\lambda)$$

Where $D^i_l(X)$ 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, respectively, for a given satisfaction level H.

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

THEOREM 3.4.1. *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

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

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. □

We now find that any order quantity will eventually induce a unique steady-state satisfaction level.

THEOREM 3.4.2. *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 Lemma 3.1.1. 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.

Therefore, 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 3.4.3. *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).

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

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

3.4.2. 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:

$$\hat{\phi}_{\hat{Q}}(Q) = E[r \min\{Q, \bar{H}(\hat{Q})\} - cQ] \tag{3.16}$$

As before, such a quantity need not exist or be unique. By the same argument in Theorem 3.3.4:

COROLLARY 3.4.4. *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 3.3.8 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.

3.5. Numerical Implementation and Examples

In this section, we illustrate the results found above through examples which show the difference between optimal and myopic policies, and how using myopic policies will generate an adverse demand distribution relative to optimal policies.

3.5.1. Dynamic Inventory Control. As noted in Section (3.2.1), optimizing the dynamic inventory control problem involves a dynamic programming problem. Because the state space is Markovian, the decision space does not expand over time and remains computationally feasible.

The myopic comparison involves finding the single period newsvendor solution for each satisfaction level.

Figure 3.3 describes the steady-state limit of an example dynamic inventory control problem with positive satisfaction ($N = 35$, $\alpha = 3$, $p = .07$ $r = 1.5$, $c = 1$). The optimal and myopic policies are compared in the top graph, with the optimal policy typically having a larger order quantity than the myopic one (as per Theorem 3.2.3). For a given satisfaction level, the optimal order quantity will have a smaller expected per period profit than the myopic policy, as shown in the second graph. However, choosing this higher order quantity will increase the distribution of the satisfaction level (the third graph), and the expectation of per period profit taken over the optimal distribution is 6.55% higher than in the myopic distribution.

Figure 3.4 shows the same steady-state limit, but for a problem with negative satisfaction ($N = 35$, $\alpha = 0.3$, $p = .21$ $r = 1.3$, $c = 1$). In this case, the optimal policy will usually have a smaller order quantity than the myopic one (Theorem 3.2.4), which gives a smaller satisfaction distribution. Since there is a negative satisfaction effect, a smaller satisfaction distribution indicates a larger demand distribution, and ordering less than the myopic policy improves the expected per period profit by 4.81%.

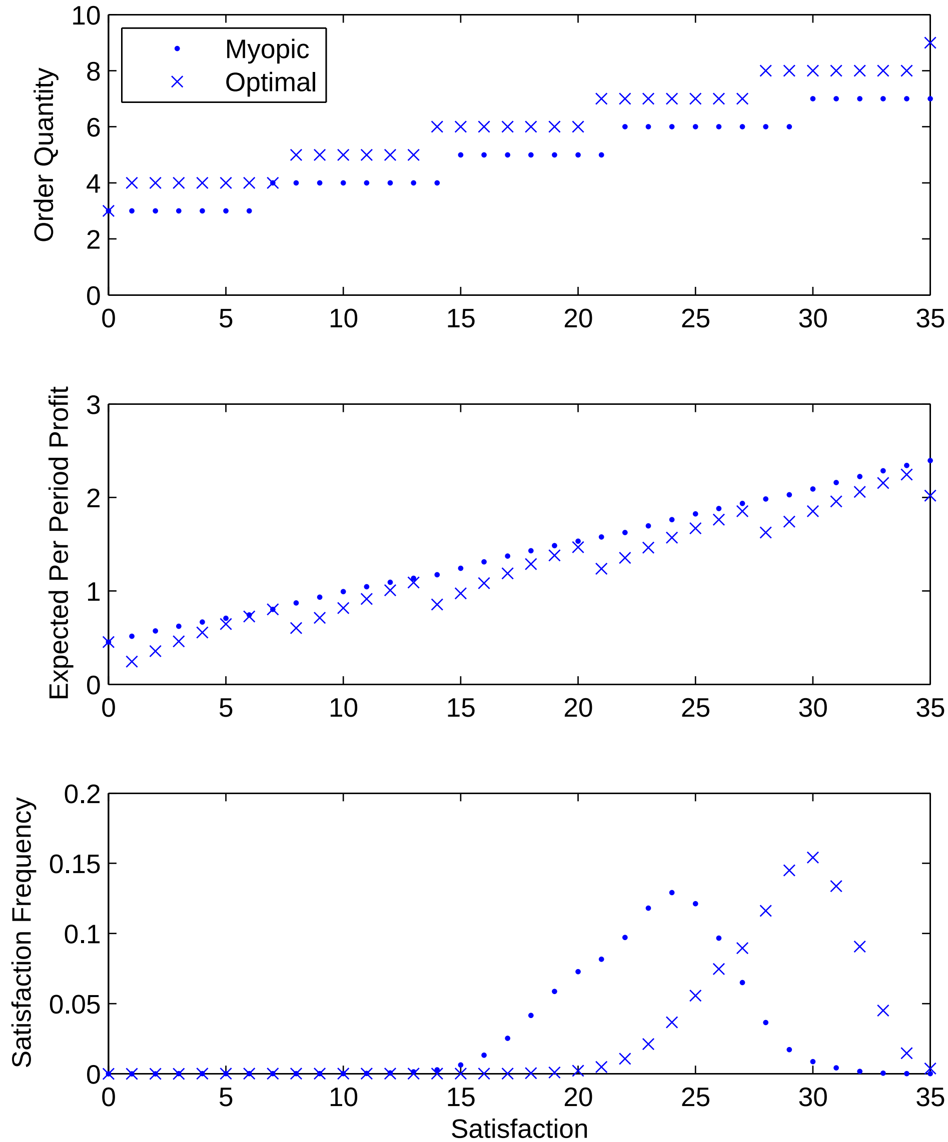

FIGURE 3.3. Dynamic Inventory Control (Positive Satisfaction)

3.5.2. Fixed-Order Policies. Optimizing (3.6) is a finite, discrete, one-dimensional optimization, and the solution can be computed through enumeration. Finding empirically myopic policies that satisfy (3.7) can be done by testing whether each order quantity satisfies myopic optimality for the distribution it creates.

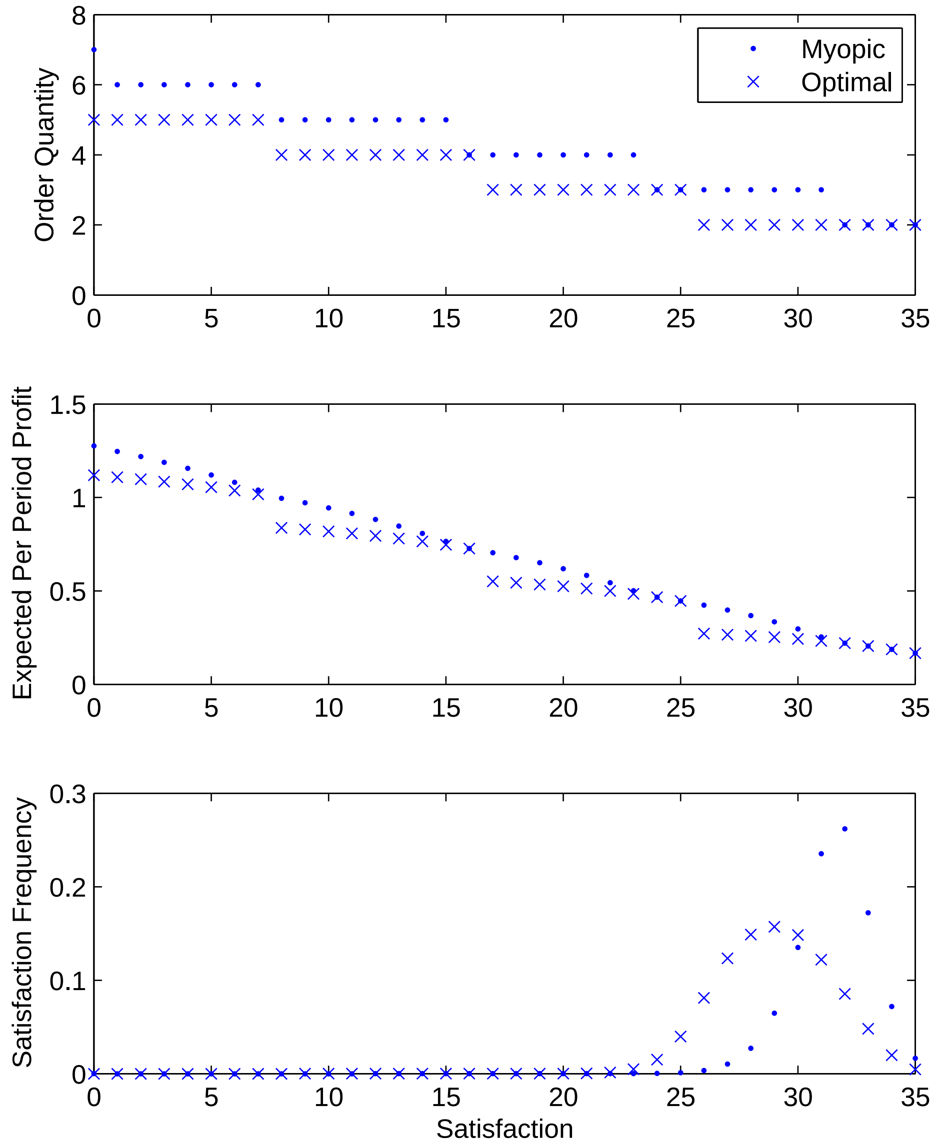

FIGURE 3.4. Dynamic Inventory Control (Negative Satisfaction)

Theorem (3.3.8) 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 parameters presented in figure 3.3. If we iteratively attempt to find the myopic optimum for a given demand distribution, we change the demand distribution and can fail to find the optimum, as shown in table 3.1. If we arbitrarily choose an order quantity, such as 10, we generate

TABLE 3.1. Iterative Myopic Fixed-Order Policies

Order Quantity	Long Term Profit	Myopic Optimum	Myopic Profit
10	2.60	8	3.10
8	2.61	7	2.73
7	2.49	7	2.49

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). This is a 4.82% improvement over the empirically myopic policy.

3.5.3. Infinite Customer Base. Theorem (3.4.2) suggest a search algorithm by which the satisfaction for a given order policy can be found when the number of customers approaches infinity. If we choose an arbitrary satisfaction level, the sign of the expected value of ΔS tells us whether the satisfaction level is higher or lower than the chosen level. A binary search can find the actual value to any level of precision.

Finding an empirically myopic fixed order quantity checks the values for agreement with the condition in (3.16).

In Figure 3.5 we see an example for a positive satisfaction effect ($\lambda = 3.5$, $\alpha = 3$, $r = 1.5$, $c = 1$). The top graph shows the steady-state satisfaction level as a function of order quantity. The vertical bars on that graph represent the myopic order quantities as a function of satisfaction level. Where these two functions intersect, the order quantity is the myopic policy for the demand function the order quantity generates, so such locations indicate empirically myopic policies. The bottom graph shows the expected per period profit for each order quantity. Note that the empirical myopic policy of 6 is lower than the optimal policy of 9, which improves the expected profit per period by 11.20%.

The arrows show the iterative myopic drift described in Corollary 3.4.4, starting from the optimal policy of 9 in this case. The myopic decision for the distribution is to order 7, which in turn generates a distribution that has a myopic policy of 6. Thus, if a manager simply observed the empirical demand distribution and optimized with respect to observations, she would drift from the optimum at 9 to a suboptimal policy at 6.

Figures 3.6 and 3.7 both show examples for a negative satisfaction effect. In Figure 3.6 ($\lambda = 10.5$, $\alpha = 0.3$ $r = 1.3$, $c = 1$), an empirically myopic policy exists at 3, while the optimal policy is at 2, with a 12.67% higher profit. However, unlike in the case of positive satisfaction, the myopic manager will not drift to the empirically myopic policy. Instead,

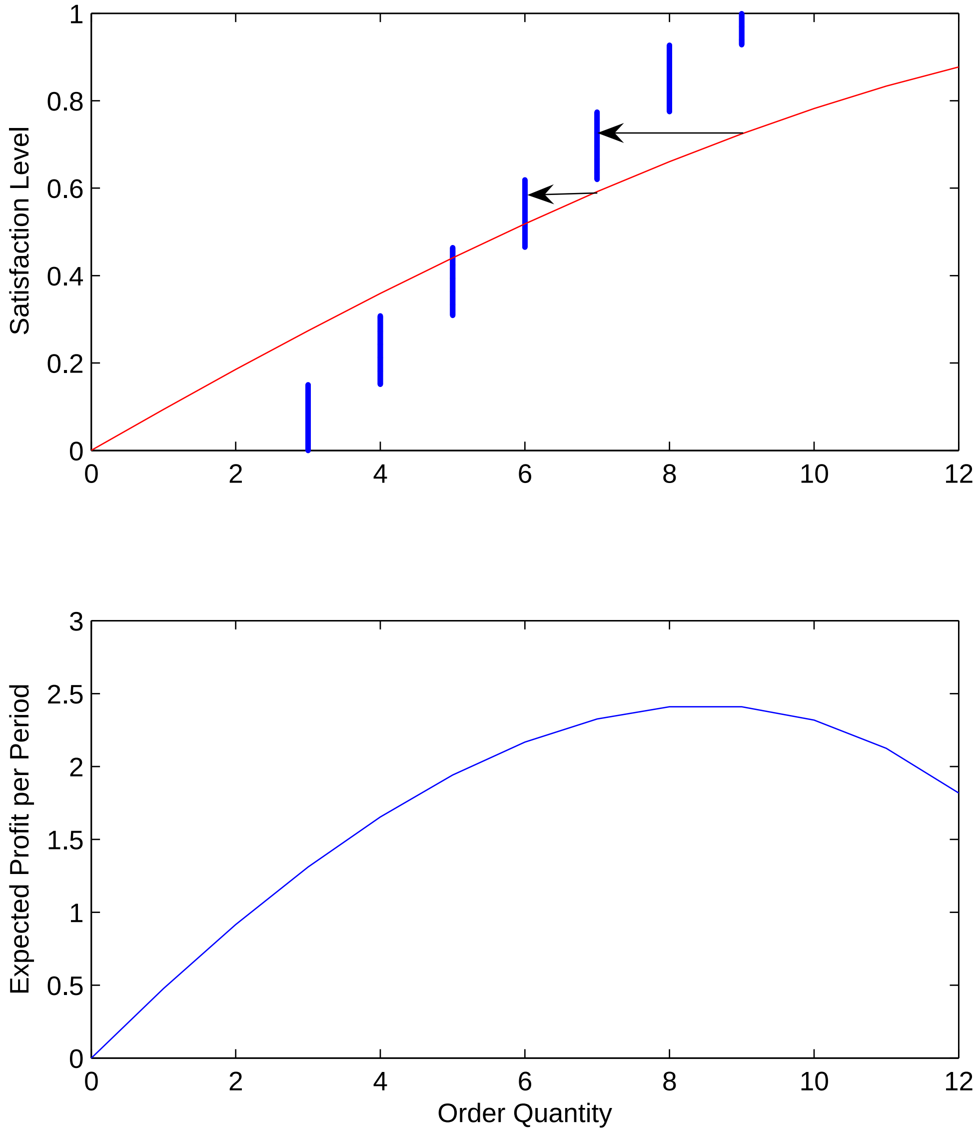

FIGURE 3.5. Infinite Customer Base (Positive Satisfaction)

she will oscillate between 2 and 4. Figure 3.7 gives an example where there is no empirically myopic policy at all, and a myopic manager will oscillate between 3 and 4.

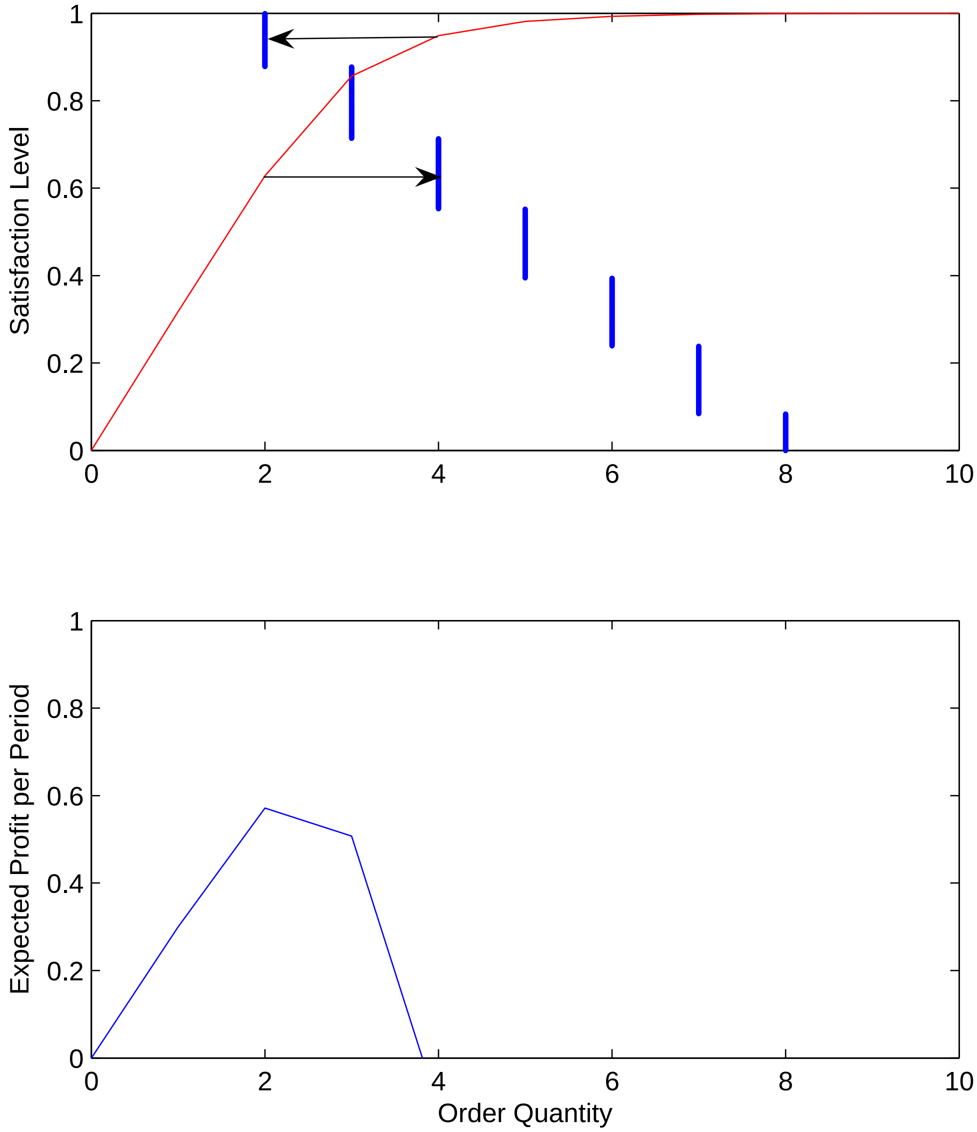

FIGURE 3.6. Infinite Customer Base (Negative Satisfaction)

3.6. Conclusion

This chapter 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

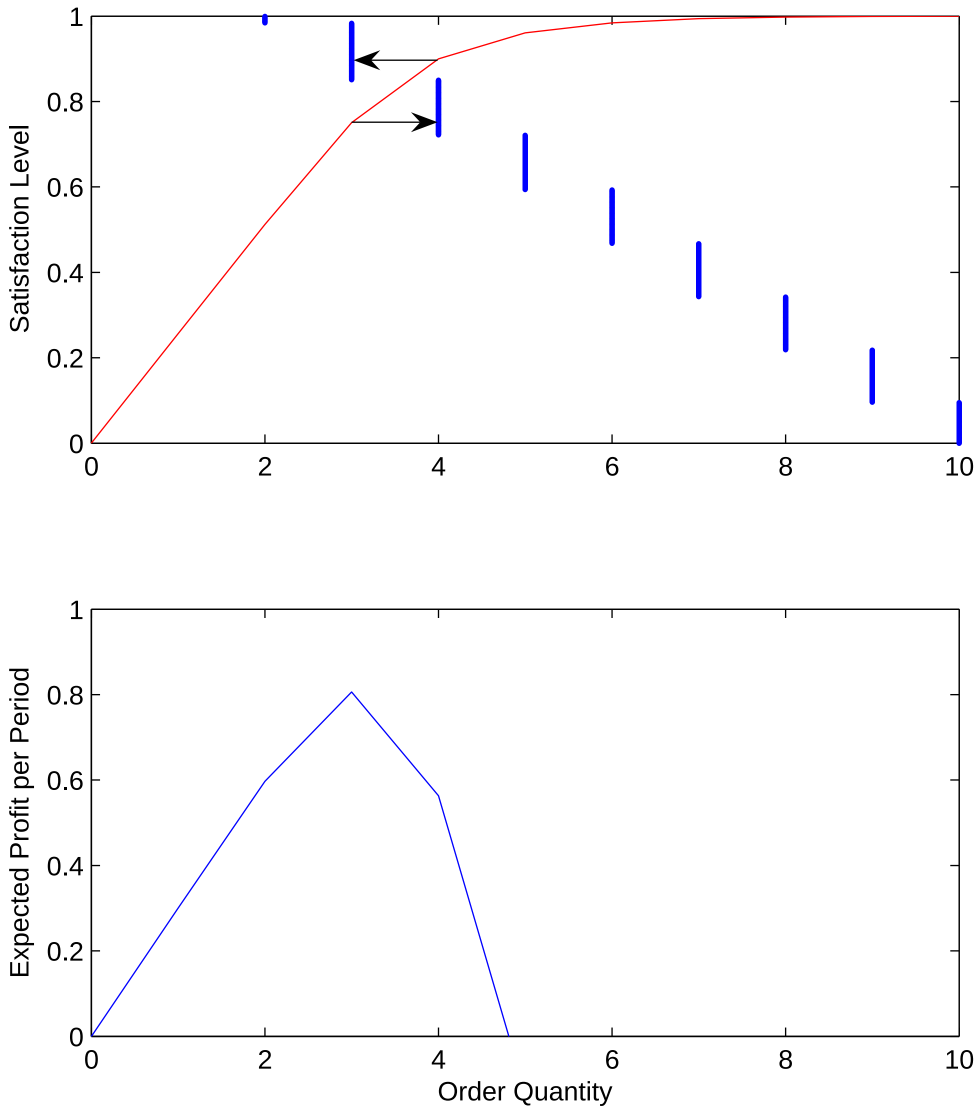

FIGURE 3.7. Infinite Customer Base (Negative Satisfaction, No Empirically Myopic Policy)

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. Managers who behave optimally will sacrifice near-term profit to create superior demand distributions over the long-term. We have also provided a way to quantitatively calculate the value of considering future demand when making inventory 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. A better understanding of how much it is worth to consider the effect of inventory decisions on future demand can inform decisions about how much should be spent studying these effects.

CHAPTER 4

Queueing

In this chapter, we will investigate the importance of considering the impact of service level on future demand in the queueing context. The basic model is a different version of the fixed-order quantity model in the inventory control context (Section 3.3). Instead of the ability to provide service renewing in each time period, we allow for the continuous-time use of server resources, where each arrival leaves the server unable to provide service for some time.

While many of the results from the inventory context have analogues in the queueing context, because customers in service do not demand product, long service times will tend to create adjustments. Additional assumptions are needed to get many of these results.

In Section 4.1, we introduce the application of the Binary Customer Satisfaction Model to the queueing context. Section 4.2 defines what optimization means in this context, while Section 4.3 applies that optimization to the case where the number of customers approaches infinity. Section 4.4 provides some numerical examples to illustrate the Binary Customer Satisfaction Model in the queueing context. Section 4.5 summarizes the broad conclusions we draw in this chapter.

4.1. Binary Customer Satisfaction for Queueing

We now apply the Binary Customer Satisfaction Model to the queueing context. The basic model is identical to the newsvendor adaptation, in that customers are in either a satisfied or unsatisfied state, and their demand behavior depends on their state. Instead of having a different probability of arrival, customers now have a different arrival rate, depending on whether they are satisfied or not.

The base model is the simple M/M/s queue, without a buffer. There is a constant number of servers, Q, and each server's service time is an exponential random variable with rate μ. Customers arrive according to exponential random variables whose rate is determined by whether they are satisfied or unsatisfied. Satisfied customers have arrival rate $\alpha\lambda$, while unsatisfied customers have arrival rate λ. The value of α will be greater than 1 if satisfied customers are more likely to arrive ("positive satisfaction effect") and less than 1 if satisfied customers are less likely to arrive ("negative satisfaction effect"). When customers do arrive, if they find an available server, they immediately enter service and will be satisfied after they receive service. If they do not find an available server, they balk and become unsatisfied (there is no buffer in this model). The total number of customers in the system is fixed (N).

At any given time, the state can be described by the number of satisfied customers (excluding those currently in service), S, and the number of customers currently being

TABLE 4.1. Rate of State Transition Events

Event	Rate	State Change ($R_t < Q$)	State Change ($R_t = Q$)
Satisfied Customer Arrival	$S_t\alpha\lambda$	R_t++, S_t- -	S_t- -
Unsatisfied Customer Arrival	$(N - S_t - R_t)\lambda$	R_t++	none
Service Completion	$R_t\mu$	R_t- -, S_t++	R_t- -, S_t++

served, R. We define the state of the system at time t by (S_t, R_t). The rate at which events occur is summarized in the Table 4.1.

The notation X++ in the table indicates that variable X increases by 1, while the notation X- - indicates that variable X decreases by 1. Note that the change in state depends on whether there are any available servers. We will occasionally refer to the number of unsatisfied customers, $N - S - R$, as U.

LEMMA 4.1.1. *In the steady-state,*

$$\alpha\lambda E[S] = \mu E[R]$$

PROOF. In steady state, the expected change in the number of satisfied customers must equal 0 over the long term. $\alpha\lambda E[S]$ is the average combined rate at which satisfied customers demand service, after which the demanding customer will no longer be satisfied (either being in service or being unsatisfied). $\mu E[R]$ is the average combined rate at which customers complete service after which they become satisfied. Since there is no way for unsatisfied customers to become satisfied directly, these two rates must be equal in order for the system to be at steady state. □

Because a portion of the customer base will be in service at any given time, the satisfaction effect does not translate directly into demand effect as it did for the inventory case. That is, a positive satisfaction effect ($\alpha > 1$) does not necessarily translate into an increased demand with increasing satisfaction. The relationship is still linear, but whether it is increasing or decreasing now depends on the ratio of λ to μ in addition to the value of α.

THEOREM 4.1.2. *The average arrival rate for a given expected satisfaction level is*

$$\lambda(N + ((1 - \frac{\lambda}{\mu})\alpha - 1)E[S]) \tag{4.1}$$

PROOF. At any given time, the arrival rate is

$$\lambda U + \alpha\lambda S$$

Since $N - S - R = U$, we can rewrite this in terms of S and R.

$$\lambda N + (\alpha - 1)\lambda S - \lambda R$$

If we consider the expectation of this equation, we can apply Lemma 4.1.1 to get the result of the theorem. □

It's important to note that this is not the conditional expected arrival rate for any particular state or group of states. This will complicate analysis of this model and its myopic comparison.

THEOREM 4.1.3. *When* $\alpha < 1$*, in the steady-state,* $E[S]$ *and* $E[R]$ *are concave and nondecreasing in* Q*.*

PROOF. We will consider sample paths indexed by q. Path q is a representation of what occurs when there are q servers. We will show that, for any path q, samples paths $q+1$ and $q+2$ can be constructed such that $E[R]$ is concave and nondecreasing in path index q. By Lemma 4.1.1, this will prove the same result for $E[S]$.

For each customer, we define a series of service completion times. At these times, if that customer is in service, it completes service. If the customer is not in service, nothing occurs at this time. Similarly, we define a series of arrival times. The customer arrives at these times if the customer is not in service and its satisfaction level allows an arrival. When $\alpha > 1$, being satisfied always allows an arrival at these times, while being unsatisfied allows an arrival when a uniform random number on [0,1] is less than $\frac{1}{\alpha}$. Similarly, when $\alpha < 1$, being unsatisfied always allows an arrival, while being satisfied allows an arrival when the uniform random number is less than α. We refer to these arrival times either as "general" or "conditional." The potential arrival and service times for each customer are the same across sample paths.

In order to make the comparison between paths q and $q+1$, we will consider the addition of a server one service completion at a time. In other words, we will examine what happens to the system when an additional server is added and performs service only once. This will give us a modified version of path q, which we will refer to as $\tilde{q}^1$. We then add another service by the additional server to path $\tilde{q}^1$, and refer to it as $\tilde{q}^2$, and so on. For notational consistency, we will refer to path q interchangeably with path $\tilde{q}^0$.

We compare paths $\tilde{q}^i$ and $\tilde{q}^{i+1}$ based on the difference in behavior in the customer who is served by the new server. This customer will take a sojourn for some length of time in path $\tilde{q}^{i+1}$ before returning to the state she held in path $\tilde{q}^i$. During this sojourn, we will keep track of two quantities: The number of additional service completions, and the number of eliminated balks.

We will show that the expected number of additional service completions is nonnegative, and increases in the number of balks. We will also show that the number of balks does not increase, and can decrease, for each additional server. Together, these facts give us that the number of additional service completions increases for each additional server, but the amount by which it increase decreases with each additional server. Since the number of service completions is proportional to R (and hence to S), this will give us the concave and nondecreasing result for the theorem.

We will describe the state of the customer based on how it differs between paths $\tilde{q}^i$ and $\tilde{q}^{i+1}$. We will consider State I to be the situation where the customer behavior is identical in both paths.

We allow server $q+1$ to serve an additional customer, which will serve the first available arrival during a period when all servers were busy on path $\tilde{q}^i$. This customer enters service instead of becoming unsatisfied. This removes 1 balk and adds 1 service completion. How this affects the rest of the system depends on the path for that customer. (Recall that for this case, conditional arrivals only occur if the customer is unsatisfied) This will put the customer in State II.

State II: The customer is in service in path $\tilde{q}^{i+1}$, but was unsatisfied in path $\tilde{q}^i$.

Event IIA: The customer has a service completion time. In path $\tilde{q}^i$, the customer was not in service, but in path $\tilde{q}^{i+1}$, a service completion occurs. This moves the customer to State III.

Event IIB: A general or conditional arrival time occurs. The customer entered service in path $\tilde{q}^i$, but is already in service in path $\tilde{q}^{i+1}$. We allow server $q+1$ to turn over the customer to the server that originally served the customer in path $\tilde{q}^i$, giving back the service completion. (Because service times are memoryless, this will not affect the system) The customer reverts to State I.

Event IIC: A general or conditional arrival time occurs. The customer attempted its next arrival when all servers were full and balked in path $\tilde{q}^i$. The customer remains in State II, but another balk is removed.

State III: The customer is satisfied in path $\tilde{q}^{i+1}$, but was unsatisfied in path $\tilde{q}^i$. Service completion times won't matter on either path.

Event IIIA: A general arrival time occurs. The customer arrives at the same time in both paths $\tilde{q}^i$ and $\tilde{q}^{i+1}$, with identical results. The customer reverts to State I

Event IIIB: A conditional arrival time occurs. The customer attempted its next arrival when all servers were full and balked in path $\tilde{q}^i$. The customer remains in state III, but another balk is removed.

Event IIIC: A conditional arrival time occurs. The customer entered service on path $\tilde{q}^i$, but does not arrive on path $\tilde{q}^{i+1}$. This costs a service completion, makes a server available on path $\tilde{q}^{i+1}$ that was not available on path $\tilde{q}^i$, and moves the customer to State IV.

State IV: The customer is satisfied in path $\tilde{q}^{i+1}$, but was in service in path $\tilde{q}^i$. Conditional arrival times won't matter on either path. However, note that there is an additional server available on path $\tilde{q}^{i+1}$ that was not available on path $\tilde{q}^i$, so we have to consider the behavior of other customers as well.

Event IVA: A service completion occurs for the customer in State IV. The customer became satisfied on path $\tilde{q}^i$, and it already is satisfied on path $\tilde{q}^{i+1}$. The customer reverts to State I. This also frees up a server on path $\tilde{q}^i$, to bring it in line with the behavior on path $\tilde{q}^{i+1}$.

Event IVB: A general arrival occurs. The customer enters service on path $\tilde{q}^{i+1}$, and already is in service on path $\tilde{q}^i$. Path $\tilde{q}^{i+1}$ gains a service completion back, and the customer reverts to State I. The extra server on path $\tilde{q}^{i+1}$ is now busy, as it is on path $\tilde{q}^i$.

Event IVC: A different customer in State I who balked in path $\tilde{q}^i$ enters service with the freed server on path $\tilde{q}^{i+1}$. We can bring the two paths in line by swapping customers, so that the customer who was in service on path $\tilde{q}^i$ is also in service on path $\tilde{q}^{i+1}$, while the different customer who balked in path $\tilde{q}^i$ (and hence was unsatisfied afterwards) becomes satisfied in path $\tilde{q}^{i+1}$. This puts that customer in State III, gives path $\tilde{q}^{i+1}$ an additional service, and removes a balk.

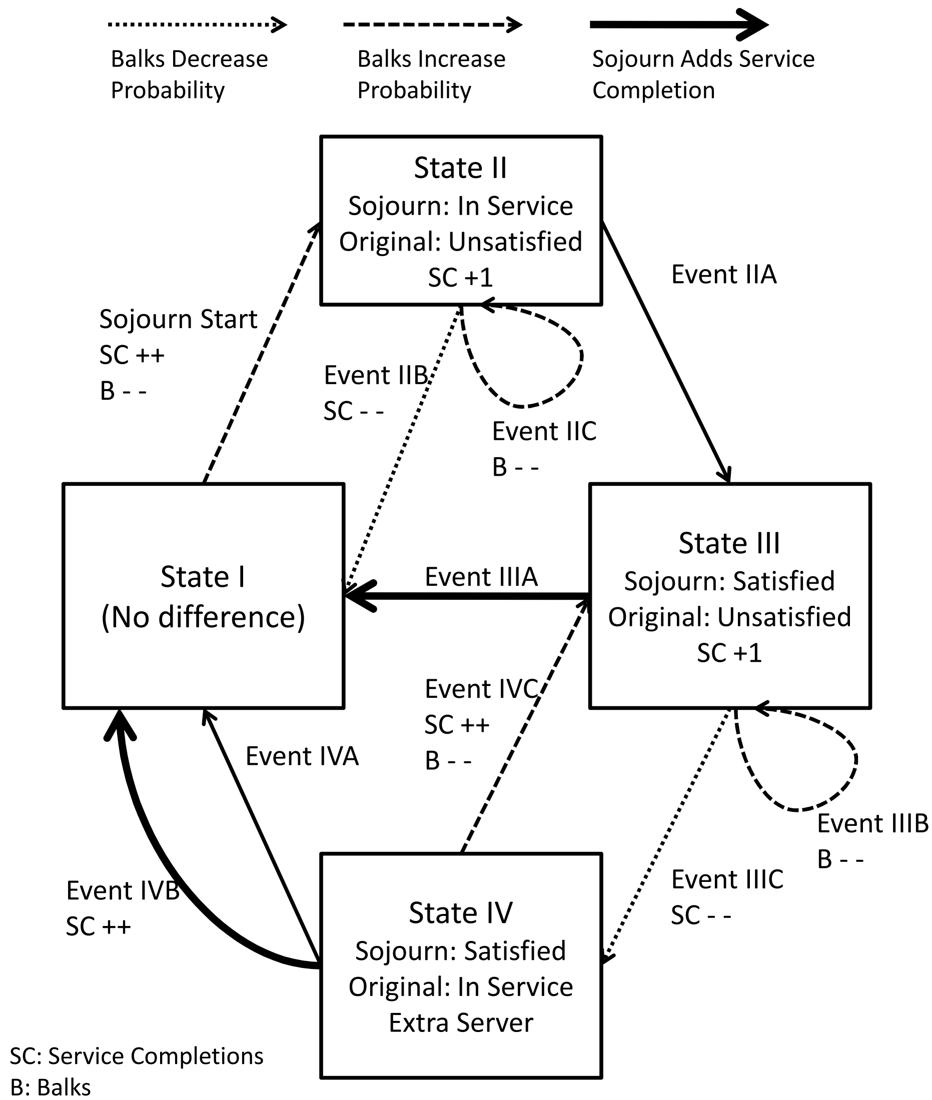

FIGURE 4.1. Customer Sojourn for Additional Service (Negative Satisfaction)

Regardless of what path is used to return to state I, we observe that at least 1, and potentially many more balks are removed in path $\tilde{q}^{i+1}$ as compared to path $\tilde{q}^{i}$. Either 1 or 0 service completions are added.

These transitions are summarized in Figure 4.1. The change in number of service completions depends on the number of times events IIIA and IVB occur.

This set of outcomes repeats for each time server $q+1$ serves an additional customer. The number of times the server serves additional customers depends on the timing of balks on path q, which means that on path $q+2$, there will be fewer opportunities than for server $q+1$ to serve customers, as there are fewer balks. Thus, the transition from I to II occurs less frequently for higher values of q.

When in state II, when there are fewer balks, event IIB (which reverts to State I and leaves no net service increase) occurs more frequently, and event IIA (which moves to State III) occurs less frequently.

When in state III, when there are fewer balks, event IIIC (which moves to State IV) occurs more frequently, and event IIIA (which reverts to State I while preserving the net service increase) occurs less frequently.

When in state IV, when there are fewer balks, event IVA (which reverts to State I with no net service increase) occurs more frequently, and event IVB (which reverts to state I with a net service increase) occurs less frequently. Event IVC also occurs less frequently.

In order to gain an additional service completion, the customer must go from State I to State II to State III. The service completion can be gained from either State III or State IV with the same rate $\alpha\lambda$ in each, but the customer can return to state I without a net service increase in State IV as well. This makes State III more favorable for gaining a service completion, and we see that with fewer balks, Event IIIC occurs more frequently and Event IVC occurs less frequently, putting more customers in State IV.

Since fewer balks corresponds to higher values of q, we see that the number of sojourns decreases, as does the probability that a service completion is gained within each sojourn, with higher values of q. This will give us the concave nondecreasing result we seek. ☐

THEOREM 4.1.4. *When $\alpha > 1$, in the steady-state, $E[S]$ and $E[R]$ are nondecreasing in Q.*

PROOF. We now consider the case when $\alpha > 1$. The second-order condition cannot be proven the same way, as each customer satisfied by server $q+1$ can create additional arrivals. However, we can still observe the first-order result.

As before, we consider State I to represent when the behavior of a customer is the same in paths $\tilde{q}^{i}$ and $\tilde{q}^{i+1}$. Again, we begin with a transition to State II when a customer who balks in path $\tilde{q}^{i}$ is served by server $q+1$. The number of service completions has increased by one, while the number of balks is reduced by one. Recall that now, conditional arrivals only occur if the customer is unsatisfied.

State II: The customer is in service in path $\tilde{q}^{i+1}$, but was unsatisfied in path $\tilde{q}^{i}$.

Event IIA: The customer has a service completion time. In path $\tilde{q}^{i}$, the customer was not in service, but in path $\tilde{q}^{i+1}$, a service completion occurs. This moves the customer to State III.

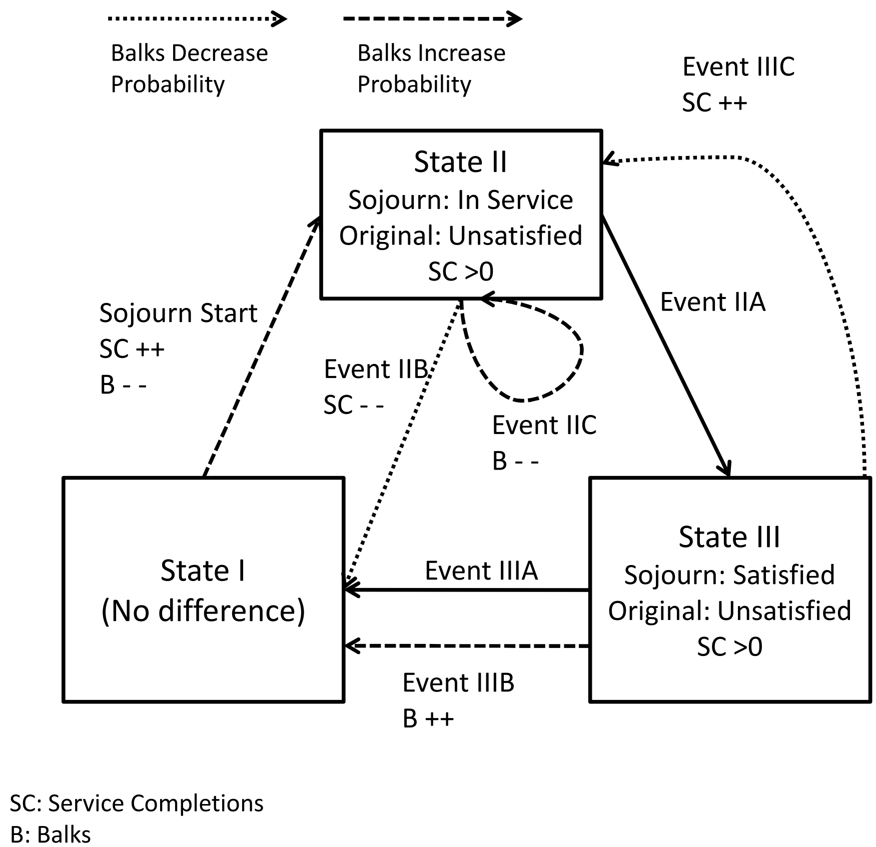

FIGURE 4.2. Customer Sojourn for Additional Service (Positive Satisfaction)

Event IIB: A general or conditional arrival time occurs. The customer entered service in path $\tilde{q}^i$, but is already in service in path $\tilde{q}^{i+1}$. We allow server $q+1$ to turn over the customer to the server that originally served the customer in path q, giving back the service completion, but freeing up the additional server and leaving the remainder of the system unchanged. The customer reverts to State I.

Event IIC: A general or conditional arrival time occurs. The customer attempted its next arrival when all servers were full and balked in path $\tilde{q}^i$. The customer remains in State II, but another balk is removed.

State III: The customer is satisfied in path $\tilde{q}^{i+1}$, but was unsatisfied in path $\tilde{q}^i$. Service completion times won't matter on either path.

Event IIIA: A general arrival time occurs. The customer arrives at the same time in both paths $\tilde{q}^i$ and $\tilde{q}^{i+1}$, with identical results. The customer reverts to State I

Event IIIB: A conditional arrival time occurs. The customer attempted its next arrival when all servers were full and balked in path $\tilde{q}^{i+1}$. The customer in path $\tilde{q}^{i+1}$ is now unsatisfied, moving the customer back to State I, and adding a balk to path $\tilde{q}^{i+1}$.

Event IIIC: A conditional arrival time occurs. The customer entered service once again on path $\tilde{q}^{i+1}$, but does not arrive on path $\tilde{q}^i$. This gains another service completion for path $\tilde{q}^{i+1}$, and moves the customer back to State II. This also removes an available server for path $\tilde{q}^{i+1}$, which has implications for future sojourns.

Since State II has a customer in service rather than unsatisfied, the expected value of R is greater for $\tilde{q}^{i+1}$. Similarly, the expected value of S is greater as State III has a customer satisfied rather than unsatisfied.

However, since the number of additional service completions depends on the availability of other servers, which may be increased from q to $q+1$ (because server $q+1$ is not always busy), we cannot derive second-order conditions from this approach.

Figure 4.2 illustrates these transitions. Note that, because of the possibility of repeated transitions between states II and III, we cannot associate a specific value for the number of additional service completions with each state. This makes determining second-order properties difficult. □

4.2. Optimization

With the model defined, we now consider optimization. The manager has control over the number of servers to use, but cannot change that number easily. Therefore, we will consider the steady-state results, and choose the number of servers to optimize various objectives.

Let us assume a cost c per server per unit time. There are a number of different objectives that could be used, but for now, let us assume a revenue p per service completion. This gives an expected profit per unit time of:

$$\phi(Q) = p\mu E[R] - cQ \tag{4.2}$$

From Lemma 4.1.1, we know the expected value of R is proportional to the expected value of S (and hence to $S+R$), so optimization for any of these objectives will be similar. The state is a continuous-time, discrete-state Markov process, and can be optimized numerically. In cases where $\alpha < 1$, the task is made easier by the following theorem, which allows a binary search algorithm.

THEOREM 4.2.1. *When $\alpha < 1$, $\phi(Q)$ is concave in Q.*

PROOF. This theorem follows from Theorem 4.1.3, as $\phi(Q)$ is the sum of the concave nondecreasing $p\mu E[R]$ and the linear $-cQ$. □

4.2.1. Myopic Policies. In order to quantify the value of considering future demand, we again wish to compare the optimal policy to a myopic policy. We will use Theorem 4.1.2, which gives the average arrival rate for a process. To represent myopic decisionmaking, we assume that a myopic manager will attempt to choose an order quantity to maximize the profit from the queueing system if the arrival rate was a Poisson process with rate fixed at the value given in (4.1). We can refer to this value as $\hat{\Lambda}(Q)$.

$$\hat{\Lambda}(Q) = \lambda(N + ((1 - \frac{\lambda}{\mu})\alpha - 1)E[S(Q)])$$

Where $E[S(Q)]$ is the expected value of satisfaction when the order quantity is Q. We will again define empirically myopic policies as policies where the policy is give the maximum profit under the arrival rate given by the policy.

DEFINITION 4.2.2. *A policy $\hat{Q}$ is empirically myopic if it optimizes*

$$\hat{\phi}(Q) = p\mu E_{\hat{\Lambda}(\hat{Q})}[R(Q)] - cQ \tag{4.3}$$

Where $E_{\hat{\Lambda}(\hat{Q})}[R(Q)]$ is the expected number of customers in service given ordering policy Q and fixed arrival rate $\hat{\Lambda}(\hat{Q})$.

Note that this definition of myopic policies falls a bit short of the definition in the inventory case (Section 3.3), as it does impose some blindness on the manager. This myopic manager ignores the fact that the presence of customers in service will reduce the arrival rate, though recognizing this does not require an understanding of the customer behavior model. The distribution of satisfied and unsatisfied customers among the potential demand pool is not equal over all possible values of R, so there is no straightforward way to perfectly capture the myopic manager's behavior. This failing does not apply in the infinite customer case below, as the number of customers in service has an effect on arrival rate that approaches 0.

4.3. Infinite Customer Base

We now allow the number of customers to approach infinity while keeping the arrival rate bounded. We find that many of the fixed-order quantity results from the inventory case only apply in the infinite customer base model in the queueing context. This is primarily because the reduction in demand due to customers in service approaches 0 only when the demand rates of each individual customer approaches 0.

Let $\Lambda = \lambda N$ be the base customer arrival rate when all customers are unsatisfied. We allow N to approach infinity while holding Λ constant. Doing so allows us to ignore the impact of customers in service on the arrival rate. We let $H(Q) = \frac{E[S(Q)]}{N}$ be the steady-state proportion of customers satisfied when the ordering policy is Q. The arrival rate, $\bar{\Lambda}$, is then given by:

$$\bar{\Lambda}(Q) = \Lambda(1 + (\alpha - 1)H(Q)) \tag{4.4}$$

THEOREM 4.3.1. *For an infinite customer base, when $\alpha < 1$, for any value $\hat{Q}$ that is empirically myopic, there exists an optimal solution Q^* such that $\hat{Q} \geq Q^*$.*

PROOF. Suppose $\hat{Q} \leq Q^*$. Then, since Q^* is optimal,

$$p\mu E_{\bar{\Lambda}(Q^*)}[R(Q^*)] - cQ^* \geq p\mu E_{\bar{\Lambda}(\hat{Q})}[R(\hat{Q})] - c\hat{Q}$$

Since Q optimizes equation 4.3,

$$p\mu E_{\bar{\Lambda}(\hat{Q})}[R(\hat{Q})] - c\hat{Q} \geq p\mu E_{\bar{\Lambda}(\hat{Q})}[R(Q^*)] - cQ^*$$

Finally, from equation 4.4, we see that the arrival rate is decreasing in $H(Q)$, and from Theorem 4.1.3, we know that H is increasing in Q, and a decreased arrival rate will reduce the average number of customers in service.

$$p\mu E_{\bar{\Lambda}(\hat{Q})}[R(Q^*)] - cQ^* \geq p\mu E_{\bar{\Lambda}(Q^*)}[R(Q^*)] - cQ^*$$

These three inequalities form a cycle of inequalities, so all three profits must be equal, and $\hat{Q}$ must also be optimal.

□

THEOREM 4.3.2. *For an infinite customer base, when $\alpha > 1$, for any value $\hat{Q}$ that is empirically myopic, there exists an optimal solution Q^* such that $\hat{Q} \leq Q^*$.*

The proof is almost identical to the previous one, so we have omitted it here.

Since the fraction of customers in service at any given time approaches 0 as N approaches infinity, the change in satisfaction over a finite time approaches 0. Nevertheless, we want a way to characterize the direction of change, which we do by defining a value $\Delta S(Q, H)$, which is the expected rate at which satisfaction changes as a function of the current satisfaction level and service level.

$$\Delta S(Q, H) = \mu E_{\Lambda(1+(\alpha-1)H}[R(Q)] - \Lambda\alpha H$$

The first term is the rate at which satisfied customers leave service, given the satisfaction level and service level, while the second term is the rate at which satisfied customers seek service at the given satisfaction level. At steady-state, this value will be 0.

THEOREM 4.3.3. *For a given service policy Q, there exists a unique satisfaction level $\bar{H}(Q)$ such that $\Delta S(Q, \bar{H}(Q)) = 0$ and:*
$\Delta S(Q, H) < 0$ *for* $H > \bar{H}(Q)$
$\Delta S(Q, H) > 0$ *for* $H < \bar{H}(Q)$

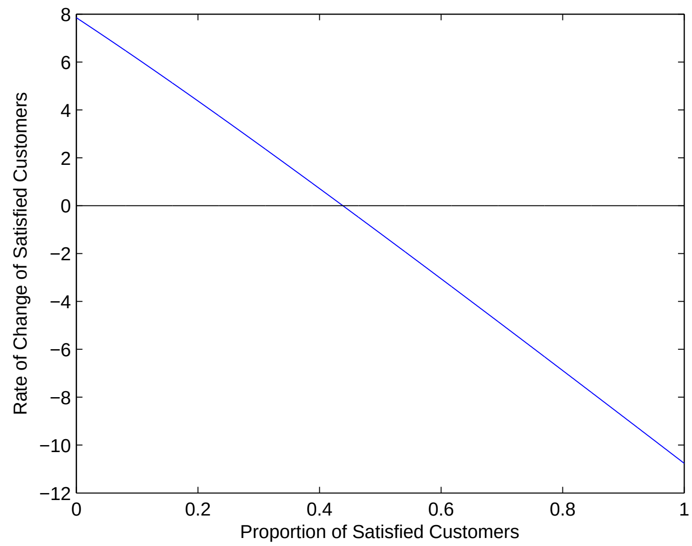

FIGURE 4.3. $\Delta S(Q,H) < 0$ as a function of H

PROOF. We observe that for $H = 0$, $\Delta S(Q,H)$ only has the service completion term, which is positive. For $H = 1$, the satisfied customer arrival term must be at least as large as the service completion term, as the number of service completions cannot be higher than the number of arrivals, and for $H = 1$, all arrivals are satisfied arrivals. Thus, the term must be nonpositive. Since $\Delta S(Q,H)$ is continuous in H, it must equal 0 at some value of H.

We also observe that $\Delta S(Q,H)$ is monotonically decreasing in H for a fixed value of Q. The satisfied arrival term is clearly monotonically decreasing in H. If $\alpha < 1$, the service completion term is also monotonically decreasing in H, as the arrival rate decreases in H. If $\alpha > 1$, the service completion term is increasing in H, but cannot increase by more than the number of additional arrivals, which will always be smaller than the number of additional satisfied arrivals, as the number of unsatisfied customer arrivals decreases in H. Thus, $\Delta S(Q,H)$ is still monotonically decreasing in H.

Since $\Delta S(Q,H)$ must be 0 somewhere on the interval (0,1], and is monotonically decreasing, $\bar{H}(Q)$ is unique and has the properties in the theorem. □

Figure 4.3 provides an example of how $\Delta S(Q,H) < 0$ varies with H ($\Lambda = 10$, $\mu = 1$, $\alpha = 2$, $Q = 10$). The point where the rate of change intersects with 0 is the value of $\bar{H}(Q)$. For this example, $\bar{H}(Q) = 0.44$. For any satisfaction level to the left of $\bar{H}$, the number of satisfied customers is increasing over time, while for satisfaction levels to the right of $\bar{H}$, the number of satisfied customers is decreasing.

The importance of this theorem is that, as the number of customers approaches infinity, for any given service policy Q, the proportion of satisfied customers will gradually move to $\bar{H}(Q)$. This gives us a function for the steady-state value of H in terms of Q. We can apply theorems 4.1.3 and 4.1.4 to this function. This provides us with the spiraling result for the queueing context, as we had for the inventory context.

THEOREM 4.3.4. *In the infinite customer case, if $\alpha > 1$, an empirically myopic policy exists, and sequentially finding myopically optimal policies for the arrival rates created by the previous policy will eventually induce an empirically myopic policy, regardless of initial satisfaction level.*

Suppose we begin at an initial demand satisfaction level H^0. There exists a myopic policy for the arrival rate generated by satisfaction level H^0, which we will refer to as $\hat{Q}^1$, which in turn generates a satisfaction level $\bar{H}(\hat{Q}^1)$. Subsequent values of $\hat{Q}^i$ optimize Q against the arrival rate generated by the previous value.

$$\hat{Q}^{i+1} = \arg\max\{p\mu E_{\bar{\Lambda}(\hat{Q}^i)}[R(Q)] - cQ\}$$

Observe that $\bar{\Lambda}(\hat{Q}^i)$ is increasing in $\hat{Q}^i$, since $\alpha > 1$. Thus, if $\hat{Q}^{i+1} \geq \hat{Q}^i$, then $\hat{Q}^{i+2}$ will be maximizing with a larger arrival rate, and so $\hat{Q}^{i+2} \geq \hat{Q}^{i+1}$ as well, and so on. But since the arrival rate is bounded by $\alpha\Lambda$, there exists a maximum value of $\hat{Q}$. Thus, at some point, not only must $\hat{Q}^{i+1} \geq \hat{Q}^i$, but it the two quantities must also be equal. This is the definition of an empirically myopic policy.

Similarly, if $\hat{Q}^{i+1} \leq \hat{Q}^i$, then $\hat{Q}^{i+2} \leq \hat{Q}^{i+1}$. Since Q cannot be less than 0, again, at some point, $\hat{Q}^{i+1} = \hat{Q}^i$.

As we will see below, an empirically myopic policy need not be optimal. As a result, even with completely accurate and reliable information about the arrival rate at any given time, a manager who only tries to choose the best value of Q for the observed arrival rate may move away from optimality and settle on a sub-optimal solution.

4.4. Numerical Results

4.4.1. Finite Customer Base. Using the state transitions in Table 4.1, we can construct a continuous time, discrete space Markov chain in two dimensions and solve for the steady-state distribution analytically. The ability to compare to myopic policies is limited in the finite customer case, but we can still observe the impact of service level on satisfaction. We provide three examples, examining the impact of service level on average demand and the expected number of satisfied customers. Recall that the expected number of satisfied customers is proportional to expected number of customers in service and the expected number of service completions.

The first example is displayed in Figure 4.4 ($N = 30$, $\alpha = 3$, $\lambda = 0.2$, $\mu = 1$). As described in Theorem 4.1.4, the number of satisfied customers is increasing in Q. The average arrival rate is also increasing in Q. Compare this with Figure 4.5 ($N = 30$, $\alpha = 1.5$, $\lambda = 0.5$, $\mu = 1$). This is still an example with a positive satisfaction effect, and the number of satisfied customers is increasing in Q. However, looking at Theorem 4.1.2, since the arrival rate is so high compared to the service rate and the satisfaction effect, many of the additional satisfied customers spend much more time in service, to the point where the average demand

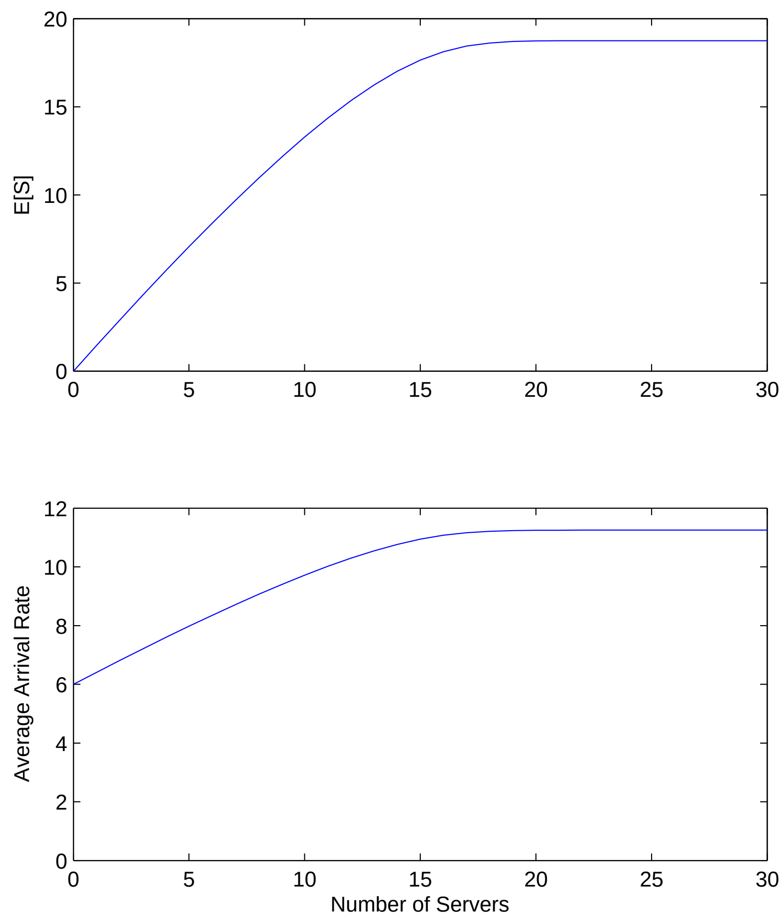

FIGURE 4.4. Positive Satisfaction Effect, Positive Demand Effect

is actually decreasing in Q, even when the number of satisfied customers are increasing, and those customers have a higher arrival rate.

The last example appears in Figure 4.6 ($N = 30$, $\alpha = 0.5$, $\lambda = 0.4$, $\mu = 1$). This is a negative satisfaction effect, and we can observe the number of satisfied customers is concave and nondecreasing, as per Theorem 4.1.3. The average arrival rate is convex and decreasing, and inspection of Theorem 4.1.2 shows that it is not possible to have a positive demand effect with increased satisfaction.

4.4.2. Infinite Customer Base. In the infinite customer case, we are able to make comparisons to the myopic policies as was done in the inventory context. For a given value of Q, $\bar{H}(Q)$ can be found by finding the value of H which satisfies Theorem 4.3.3. For any

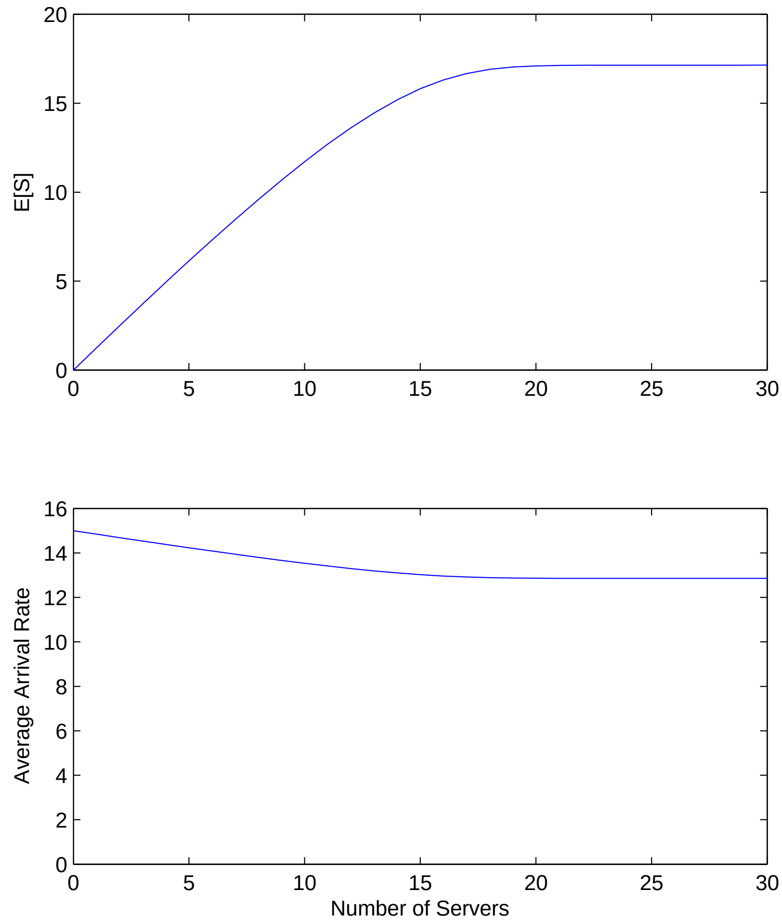

FIGURE 4.5. Positive Satisfaction Effect, Negative Demand Effect

given value of H, the myopic policy can be found as a one-dimensional optimization on Q for the fixed arrival rate.

The first example is a positive satisfaction example, given in Figure 4.7 ($\alpha = 2$, $\Lambda = 10$, $\mu = 1$, $p = 1.3$, $c = 1$). The top graph displays $\bar{H}(Q)$. The vertical bars represent the myopic policies for given values of H. The intersection of these two functions represents an empirically myopic policy. In this particular case, the optimal value for Q is 13, with a profit of 1.53. However, a myopic manager observing the arrival rate caused by having 13 servers will determine that 11 servers is the correct number. Then, observing the new arrival rate because of this decision, she will remove another server to reach the empirically myopic policy of 10, which has a profit of 1.40. Thus, the optimal policy has a profit that is 9.35% higher than the empirically myopic policy.

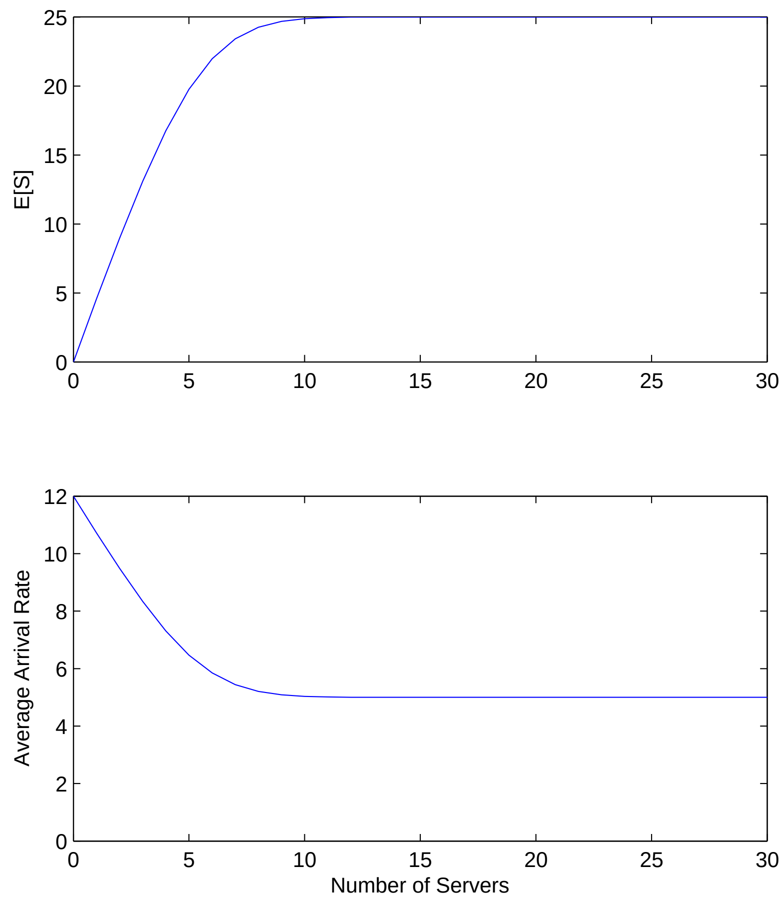

FIGURE 4.6. Negative Satisfaction Effect, Negative Demand Effect

The next example has a negative satisfaction effect, and is given in Figure 4.8 ($\alpha = 0.6$, $\Lambda = 20$, $\mu = 1$, $p = 1.3$, $c = 1$). The optimal value is 9, but the myopic manager will want to add a server to accommodate the arrival rate. This gives the optimal policy a profit that is 3.52% higher than the empirically myopic policy.

Another negative satisfaction example, given in Figure 4.9, does not have an empirically myopic policy ($\alpha = 0.5$, $\Lambda = 20$, $\mu = 1$, $p = 1.3$, $c = 1$). The optimal policy is at 7. A myopic manager seeing the arrival rate of this policy will add two servers for a total of 9. The arrival rate generated by this policy, however, will suggest the myopic manager should, in fact, have 8 servers. But the arrival rate generated by 8 servers again indicates that 9 servers is the appropriate policy to the myopic manager. Thus, the myopic manager will be continuously changing the number of servers to keep up with the changes in arrival rate.

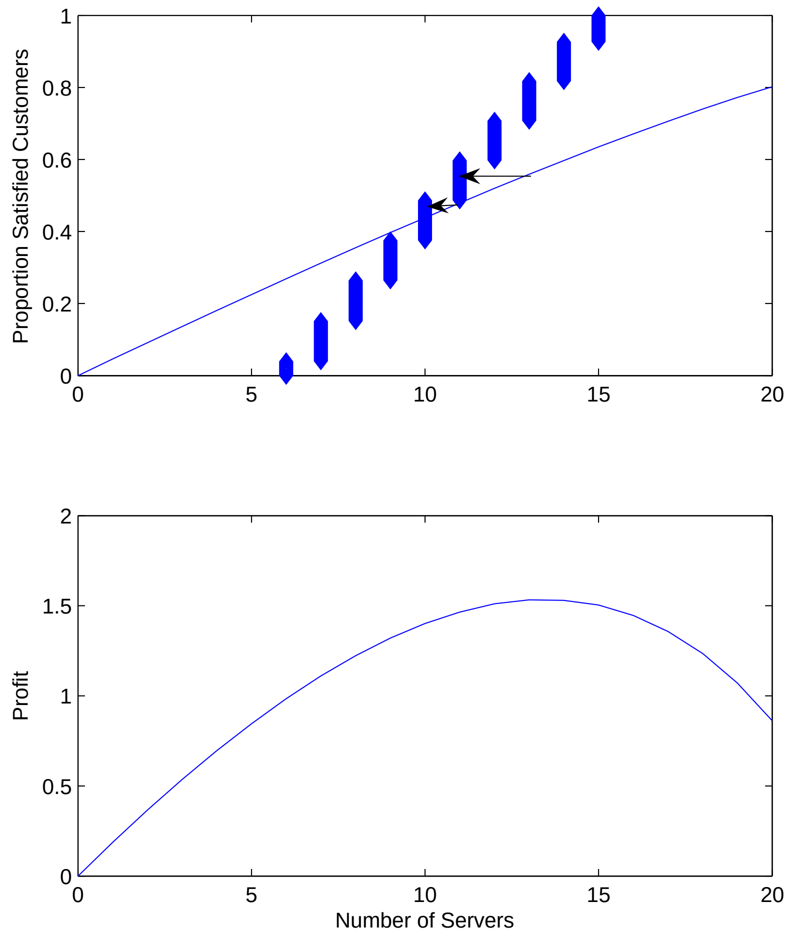

FIGURE 4.7. Infinite Customer Base, Positive Satisfaction

4.5. Conclusions

We have extended the Binary Customer Satisfaction Model to the queueing context. In the queueing context, we find that many of the same results hold, but because customers in service affect the demand rate, more assumptions are necessary to reach these results. In particular, we had to assume the number of customers approaches infinity to reach most results.

However, the basic intuition of the Binary Customer Satisfaction Model still holds. We find that a full understanding of the demand process at a given time is insufficient to reach optimal solutions when service affects future customer behavior. We have provided a means of quantifying this suboptimality, to help justify whether or not finding out more about customer behavior is worth additional cost.

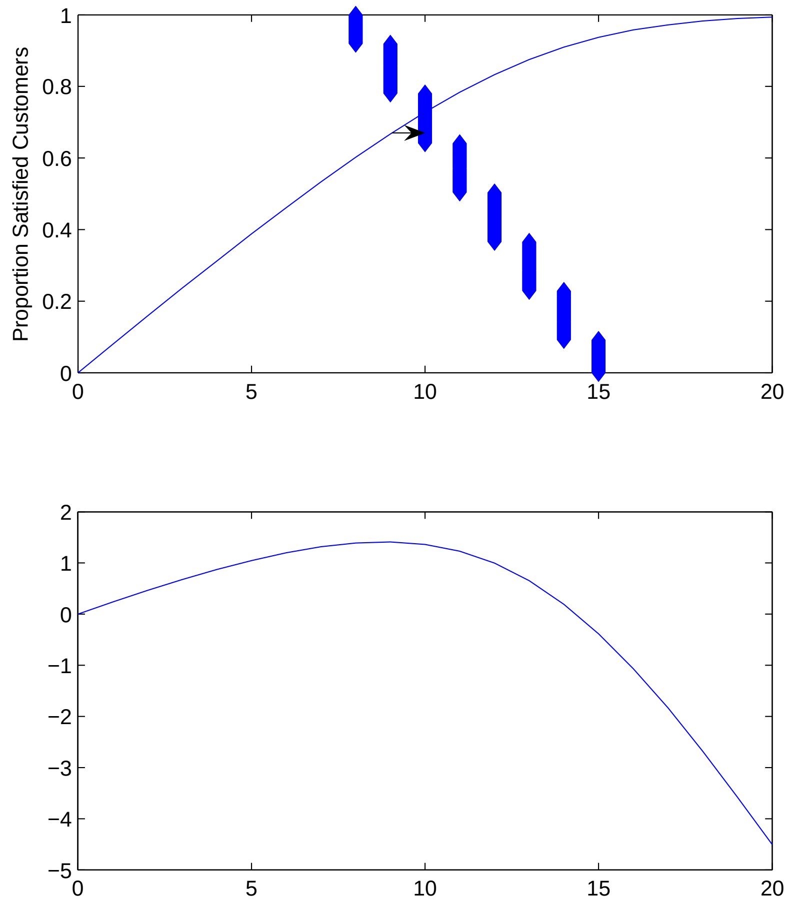

FIGURE 4.8. Infinite Customer Base, Negative Satisfaction

As with the inventory model, the queueing model can be extended into other forms of queues. A buffer can be added, for instance, or customers can arrive from different groups, each with their own behavior models.

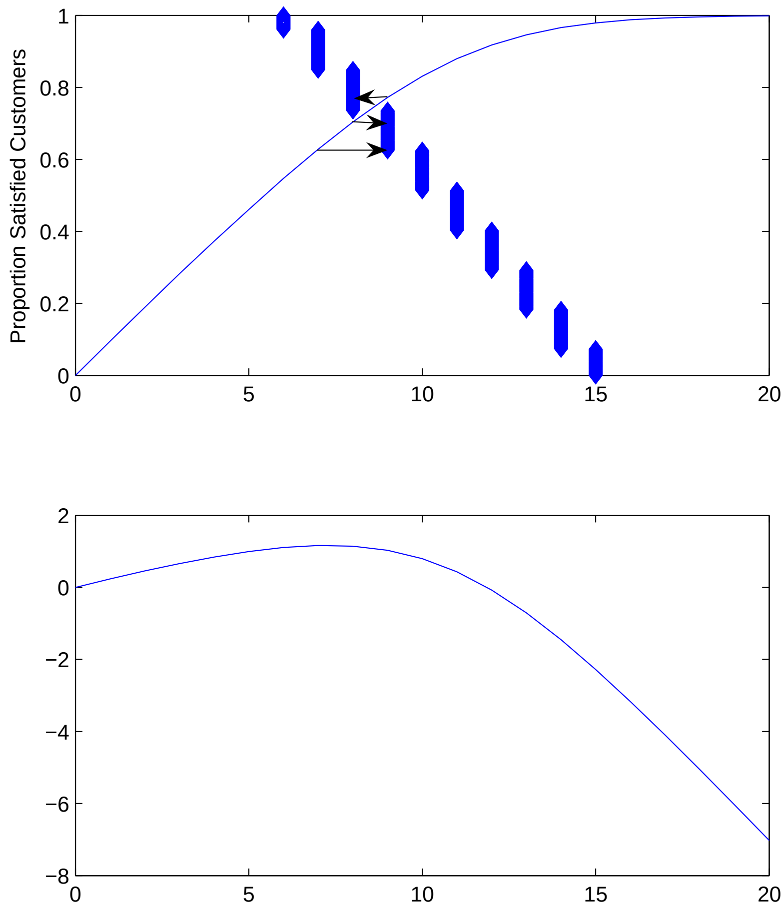

FIGURE 4.9. Infinite Customer Base, Negative Satisfaction, No Empirically Myopic Policy

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] L. Debo, G. J. Ryzin. Creating Sales with Stockouts. Working Paper No. 09-25, The University of Chicago Booth School of Business. 2009.

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

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

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

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

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

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

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

[11] 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.

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

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

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

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

[16] 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.

APPENDIX A

Optimality of Providing Service

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.0.1. *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. As a result, the lost profit will not necessarily be recouped, and cannot be surpassed.

APPENDIX B

Sample Path Stochastic Orders and Coupling

Shanthikumar and Yao [**16**] 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.0.2. $X(\theta)$ *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:*

$$\theta^1 \leq \theta^2 \leq \theta^4$$
$$\theta^1 \leq \theta^3 \leq \theta^4$$

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

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

In other words, there exists a sample path under which variables with the same distribution can be coupled such that the increasing concave property holds almost surely. Other orders are defined similarly, simply by changing the directions of the inequalities in (B.1). For example:

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