On the Value of Commitment and Availability Guarantees When Selling to Strategic Consumers

Xuanming Su

Haas School of Business, University of California, Berkeley, Berkeley, California 94720,
xuanming@haas.berkeley.edu

Fuqiang Zhang

Olin Business School, Washington University in St. Louis, St. Louis, Missouri 63130,
fzhang22@wustl.edu

This paper studies the role of product availability in attracting consumer demand. We start with a newsvendor model, but additionally assume that stockouts are costly to consumers. The seller sets an observable price and an unobservable stocking quantity. Consumers anticipate the likelihood of stockouts and determine whether to visit the seller. We characterize the rational expectations equilibrium in this game. We propose two strategies that the seller can use to improve profits: (i) commitment (i.e., the seller, ex ante, commits to a particular quantity) and (ii) availability guarantees (i.e., the seller promises to compensate consumers, ex post, if the product is out of stock). Interestingly, the seller has an incentive to overcompensate consumers during stockouts, relative to the first-best benchmark under which social welfare is maximized. We find that first-best outcomes do not arise in equilibrium, but can be supported when the seller uses a combination of commitment and availability guarantees. Finally, we examine the robustness of these conclusions by extending our analysis to incorporate dynamic learning, multiple products, and consumer heterogeneity.

Key words: availability guarantee; commitment; product availability; search costs; stockouts; newsvendor; rational expectations
History: Received July 18, 2007; accepted December 15, 2008, by Ananth Iyer, operations and supply chain management. Published online in *Articles in Advance* March 6, 2009.

1. Introduction

Product availability is essential in attracting consumer demand. Consumers would not patronize a firm without some form of assurance that they can find what they are looking for. Despite technological and managerial advances, stockouts are a common phenomenon and product availability remains a key issue in marketing and operations (Balachander and Farquhar 1994, Anderson et al. 2006). Many industry studies have found that about 8% of retail items are out of stock at any one time, and the percentage can be much higher for advertised items (e.g., Andersen Consulting 1996, Gruen and Corsten 2008). A high likelihood of stockouts makes consumers less willing to visit the store. A lower demand, in turn, makes it more costly for the firm to maintain satisfactory levels of availability. This creates a chicken-or-egg issue, which may spiral toward market breakdown. In this paper, we study the role of product availability in shaping market interactions between the firm and its consumers.

The video rental industry provides an interesting example (see Dana 2000). According to a Time Warner survey, about 20% of customers were not able to rent their first-choice movie title on a typical visit. Some customers who find their preferred movie stocked out may leave the store empty handed, and the possibility of wasted trips may discourage future visits. To address this problem, Blockbuster initiated the "Go Home Happy" program, which guarantees availability of selected new releases in the form of a rain check for the same movie if it was unavailable (see Furman 1998, Dana and Spier 2001). Consumers rated the availability guarantees as a highly desirable attribute of a video rental store (see Video Software Dealers Association 1999).

The out-of-stock problem has also attracted much attention in the consumer goods industry. In a comprehensive study covering 32 product categories in 661 retail outlets across 29 countries, Gruen et al. (2002) show that overall out-of-stock rates are consistently in the range of 5%–10%, and often exceed 10% for items on promotion. Fitzsimons (2000) finds that consumers exposed to a stockout are substantially less likely to return to the same store on their next shopping trip. Although stockouts lead to decreased customer satisfaction, firms can adopt policies to mitigate their effects. For example, through consumer surveys, Verhoef and Sloot (2005) find that consumers appreci-

ate discount coupons, rain checks, and additional services, such as home delivery, for out-of-stock items.

The problem of product unavailability is not limited to traditional brick-and-mortar retail firms. In the case of mail-order catalog companies, customers even report stockouts as their chief complaint (see Fitzsimons 2000, *Consumer Reports* 1987). This is not surprising, because catalog companies, particularly specialty merchandisers, typically stock less than 80% of their advertised items (Balachander and Farquhar 1994). It is thus common practice in the catalog industry to offer financial compensation, such as price discounts or reduced shipping charges, to customers ordering out-of-stock items (Anderson et al. 2006). Similarly, some online retailers offer monetary incentives for customers to place a backorder on currently unavailable items (Breugelmans et al. 2006, Bhargava et al. 2006). To the extent that this strategy is effective in sustaining consumer demand, some e-tailers have even moved to a "stockless" mode of operation, relying entirely on using discounts to compensate consumers for stockout-related costs (Sun et al. 2008).

Consumers are concerned about product availability because stockouts are costly to them. Consumers often invest time and energy into their shopping trips. These ex ante costs are sunk and not recoverable. In many cases, these sunk costs can be quite substantial. When shopping for electronics, toys, or video rentals, customers typically travel to the store and locate their items on the shelf. Some consumers also spend time checking product reviews beforehand. For fashion apparel and footwear, customers often need to spend time trying on different sizes or styles before finalizing their selection (if available). For specialty items such as fragrance and skincare, customers need to evaluate the product at the store before making a purchase decision. Even for online purchases, the frictional costs for navigating through webpages and submitting orders could be substantial (see Hann and Terwiesch 2003). All these effort will be in vain if they culminate in an out-of-stock experience. Additionally, consumers incur ex post costs when a desired item is out of stock. Consumers may face increased decision-making cost (e.g., of having to choose a substitute), increased transaction cost (e.g., from going to another store to purchase the item), as well as other inconveniences (e.g., when delaying the purchase). This discussion brings us to the main premise of this paper: because stockouts are costly to consumers, product availability has a direct effect on demand.

Our model has two main elements: demand uncertainty and consumer costs. Demand uncertainty is important because it is the potential mismatch between supply and demand that generates stockouts. Consumer costs are also important because they are the reason why consumers are concerned about

stockouts in the first place. In our model, there is a newsvendor seller who makes pricing and inventory decisions. Consumers choose whether to visit the seller and incur sunk costs if they do. The seller's objective is to maximize expected profits, and consumers wish to maximize individual utility. Following our earlier work, Su and Zhang (2008), we adopt the rational expectations (RE) paradigm. Specifically, consumers form rational expectations of product availability and make their patronage decisions accordingly; similarly, the seller forms rational expectations of consumer behavior and acts accordingly. Using this setup, we characterize consumer behavior as well as optimal prices and inventory levels in the RE equilibrium. We also study the case where the seller can commit to some (presumably high) service level. In addition, we consider availability guarantees—that is, when the product is out of stock, the seller compensates consumers through monetary payments, store credit, or other forms of goodwill services. Throughout our analysis, we use the first-best (under which social efficiency is maximized) as a benchmark to evaluate market outcomes.

Let us now summarize our three main findings. First, we show that the ability to commit enhances the seller's profits. The intuition is as follows. Normally, consumers make patronage decisions based on their expectations of product availability. In contrast, when the seller commits to a particular service level, he is effectively taking the lead in this game and steering the equilibrium toward higher expectations and increased product availability. However, unless credibility is sustained by some external commitment device (e.g., the seller's reputation concerns, consumer laws, or third-party endorsements), consumers may not trust the seller's claims. This issue of credibility suggests that commitment may not always be feasible in practice.

Our second major finding is that availability guarantees are useful contractual mechanisms. By artificially inflating the *seller*'s costs of stockouts, such guarantees help the seller commit to higher service levels, which attracts customer patronage and increases willingness to pay. In fact, with availability guarantees, the seller can attain or even surpass profit levels under commitment. This finding suggests that availability guarantees can serve as a replacement for a commitment device. In our analysis, we consider various types of availability guarantees. When a product is out of stock, the seller may offer monetary compensation, or alternatives (such as store credit or discount coupons) through which the seller may exploit cost savings. That is, the cost of providing stockout compensation may be less than its value perceived by consumers. We find that such cost savings provide an additional avenue for the seller to improve profits. However, this

is a double-edged sword as it creates incentives for the seller to generate market inefficiencies. Specifically, the seller tends to overcompensate consumers during stockouts, because such an arrangement allows him to realize a higher-than-efficient price level.

Finally, we discuss the welfare implications of our results. As a benchmark for comparison, we study the first-best market outcomes under which social welfare is maximized. In general, we discover that efficiency is not attained in equilibrium. In most cases, the seller tends to set quantities too low and to offer excessively generous availability guarantees. Nevertheless, we find that the first-best can indeed be supported in equilibrium using a combination of commitment and availability guarantees. This suggests that when addressing the interplay between product availability and consumer behavior, both ex ante and ex post considerations are important. The seller should commit to high service level ex ante, and at the same time should insure consumers against ex post losses.

The remainder of this paper is organized as follows. Section 2 reviews the literature. In §3, we set up our basic model and characterize the seller's price, quantity, and profits in the RE equilibrium. In §4, we study the value of the ability to commit by comparing this to our basic model. In §5, we analyze availability guarantees: we first consider offering monetary compensation to consumers facing stockouts, and then proceed to consider using more general forms of compensation. In §6, we characterize the first-best and show that it can be attained using a combination of service-level commitment and availability guarantees. Section 7 presents several extensions of the basic model, and shows that our main results remain unchanged. Finally, we offer concluding remarks in §8. All proofs are presented in the online appendix (provided in the e-companion).[1]

2. Literature Review

Most closely related to our work is the stream of literature on using product availability as a strategic lever to attract or induce demand. Dana and Petruzzi (2001) consider a newsvendor setup, similar to ours, in which demand depends on the firm's price and inventory. In the model, consumers forgo an outside option when she visits the firm before knowing whether the product will be available. Alexandrov and Lariviere (2007) recently analyzed a related model in the context of restaurants, where walk-in customers may be denied a seat if the restaurant is full. In these settings, higher inventory levels and higher capacity attract consumers and increases their willingness

to pay. Here, a common assumption is that prices and inventory levels are observed by consumers. In a competitive setting, Deneckere and Peck (1995) study firms' pricing and inventory decisions when demand is stochastic. Similar to newsvendor models, the aggregate demand is random and each consumer chooses to visit one (and only one) firm based on each firm's price and inventory level, which are assumed to be publicly observable, similar to the papers above. The authors demonstrate that unless the number of firms is large, a pure-strategy equilibrium may fail to exist because firms have an incentive to undercut each other's prices, similar to models of Bertrand competition. Dana (2001) extends this analysis to the case of observable prices but unobservable inventory levels. He shows that in this case, prices tend to be positively related to inventory levels. Under Bertrand timing (where firms choose prices first), higher prices induce firms to hold higher inventories, whereas under Cournot timing (where firms choose inventories first), higher prices can be used to signal higher availability. In our present work, we consider both cases: unobservable inventory (in the basic model in §3) as well as observable inventory (which falls under the case with seller commitment in §4). Furthermore, we study the impact of availability guarantees on equilibrium outcomes.

This research follows recent interest in strategic customer behavior in the operations management community. By our terminology, customers are strategic in the sense that they are active decision makers and respond strategically to underlying economic mechanisms (such as availability guarantees). In our current setting, high inventory levels stimulate demand by providing an assurance of availability. However, there are also situations in which inventory has the opposite effect on customer demand. In a companion paper, Su and Zhang (2008), we consider customers who may strategically wait for markdowns at the end of the selling season, and we study the implications of such behavior on supply chain performance. In that setting, a high inventory level deflates regular-season demand because it increases the firm's likelihood of having leftover units, and thus increases the consumer's chances of obtaining the product on sale at the end of the selling season. Cachon and Swinney (2009) also study consumers' strategic waiting for sales. They demonstrate that the value of quick response strategies is much higher in the presence of strategic consumers. In another paper, Lai et al. (2009) examine the effectiveness of using price-matching policies to discourage strategic waiting for sales. Here, the seller agrees to reimburse the price difference to consumers if the items that they buy are subsequently marked down. Next, Jerath et al. (2007) demonstrate that when facing strategic customers, opaque selling (i.e., obscuring from customers the exact product specifications,

---

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

such as departure times in the context of airlines) can enhance seller profits. In the same spirit as our work, these papers are based on the rational expectations paradigm (see Muth 1961, Stokey 1981). Consumers make patronage or purchase decisions based on rational expectations of product availability, and the seller makes pricing and inventory decisions based on rational expectations of consumer behavior.

There is another stream of work that focuses on the dynamics of consumer forward-looking behavior. This began with the durable goods monopoly literature in economics. The seminal Coase conjecture (1972) states that when consumers rationally expect future price decreases, the monopoly is forced to price at marginal cost. Representative studies include Stokey (1979, 1981), Bulow (1982), and Besanko and Winston (1990). However, all these papers do not consider inventory constraints. In the operations literature, Aviv and Pazgal (2008) are the first to analyze the dynamic pricing problem with strategic customers and inventory constraints. Liu and van Ryzin (2008) offer a different perspective by capturing this scenario using a quantity-rationing model rather than a pricing model. Su (2007) considers a mixture of strategic and myopic consumers and demonstrates that this has a significant impact on the optimal pricing policy. Yin et al. (2007) study the effect of inventory information on consumer purchase behavior; this is also one of the main themes of our present work. For a comprehensive review of the literature on dynamic pricing with strategic customers, readers are referred to Shen and Su (2007).

3. Basic Model

We begin by describing our basic model. There is a newsvendor seller who announces the price p and stocks a certain quantity q at unit cost c. Leftover units have zero salvage value. The market demand X is random and follows distribution F (and density f). This market demand can be interpreted as a mass of infinitesimal consumers. We assume that F has an increasing failure rate, which is satisfied by most commonly used distribution functions. Each individual consumer has valuation v for the product and faces a search cost of h (e.g., transportation cost, shopping time, etc.) in order to visit the seller. If the product is out of stock, this sunk cost is not recoverable and the consumer incurs an additional inconvenience cost of h' (e.g., from forgoing or delaying the purchase). Therefore, consumers who visit the seller face two possible outcomes: they receive payoff $v - p - h$ if the product is available or payoff $-(h + h')$ if the product is sold out. Consumers who do not visit the firm receive zero payoff. Without loss of generality, we assume $h' = 0$; otherwise, we use $v + h'$ in place of v

and $h + h'$ in place of h above. For most of the analysis, we assume that consumers are homogeneous, so all consumers who visit the seller are equally likely to receive the product in the event of scarcity. For a nontrivial outcome to exist, we assume $0 < h < v - c$. Define E as the expectation operation and $\wedge$ as the minimum operation, i.e., $x \wedge y = \min(x, y)$. We use an upper bar to denote the complement function, e.g., $\bar{F} = 1 - F$.

In this model, the seller chooses price p and quantity q to maximize expected profit. The price is publicly known, but the quantity is not observed by individual consumers. On the other hand, consumers choose a reservation price r to maximize individual expected payoff. This reservation price provides a decision threshold: accordingly, consumers will visit the seller if and only if the seller's price p does not exceed this threshold r.

To study this strategic interaction, we shall use the notion of an RE equilibrium. This equilibrium concept has also been used on other operations models (see Stokey 1981, Cachon and Swinney 2009, Su and Zhang 2008). We describe the equilibrium concept in our context by separately considering the consumer's and the seller's decision problems.

Consider first the typical consumer's problem. Intuitively, the consumer's reservation price should depend on product availability: the more likely it is to obtain the product *after* visiting the seller, the higher the willingness to pay *before* agreeing to visit the seller in the first place. However, the consumer is unable to ascertain the probability of finding the product in stock because this probability depends on the seller's quantity q, which is not observable. Nevertheless, the consumer may possess some beliefs over this probability of availability. Suppose that the consumer expects to obtain the product with probability ξ_{prob} if she visits the seller. Under these beliefs, the consumer expects to earn surplus $(v - p)\xi_{\text{prob}} - h$ from visiting the seller (and zero otherwise). This implies that the consumer's reservation price r satisfies $(v - r)\xi_{\text{prob}} = h$.

Next, we consider the seller's problem. To maximize expected profit, the seller should set the price equal to the consumers' reservation price, because this is the highest price at which consumers are willing to visit the seller. However, the seller does not know consumers' reservation prices. Similarly as before, we endow the seller with beliefs ξ_r over consumers' reservation prices. Given these beliefs, the seller chooses price $p = \xi_r$ and stocks the critical fractile quantity $q = \bar{F}^{-1}(c/p)$ to maximize the newsvendor profit function $\Pi(q) = pE(X \wedge q) - cq$.

Given players' rational expectations, the strategic game between the seller and the consumers decomposes into two separate decision problems (a binary choice problem for consumers and a newsvendor

problem for the seller), both of which are easy to solve and have unique solutions. Moreover, each decision problem does not require information about the other side of the market. For example, given beliefs ξ_{prob}, consumers do not need to know the seller's cost c or demand distribution F. Similarly, given beliefs ξ_r, the seller does not need to know the consumer's valuation v or search cost h. Therefore, once all players have their beliefs in place, they will naturally settle into their optimal actions.

In equilibrium, beliefs must be consistent with outcomes. In other words, the seller's beliefs ξ_r must coincide with consumers' reservation price r, and consumers' beliefs over availability probability ξ_{prob} must agree with the actual in-stock probability corresponding to the seller's chosen quantity q. In fact, this probability is given by $A(q) \equiv E(X \wedge q)/E(X)$. The argument proceeds as follows, and is due to Deneckere and Peck (1995) and further explored by Dana (2001). Notice that the product is available with probability $(x \wedge q)/x$ when there are x consumers in the market. Because the market size is uncertain, conditional on his own presence in the market, an individual consumer's posterior demand density is $g(x) = xf(x)/E(X)$. Now, integrating over this posterior demand density, the availability probability becomes $\int ((x \wedge q)/x)g(x)\,dx = \int ((x \wedge q)f(x)/E(X))\,dx = A(q)$. We may also interpret the availability probability $A(q)$ as the service level of the seller.

Summarizing the preceding discussion, we have the following definition of an RE equilibrium.[2]

DEFINITION 1. An RE equilibrium consists of $(p, q, r, \xi_{\text{prob}}, \xi_r)$ satisfying

(i) $r = v - h/\xi_{\text{prob}}$;

(ii) $\bar{F}(q) = c/p$, $p = \xi_r$;

(iii) $\xi_{\text{prob}} = A(q)$, $\xi_r = r$.

The three conditions (i), (ii), and (iii) follow, respectively, from the consumer's optimal decision r given beliefs ξ_{prob}, the seller's optimal decisions p, q given beliefs ξ_r, and consistency conditions. Here, we consider the scenario where all players are endowed with accurate beliefs of market outcomes (i.e., rational expectations). The assumption that customers possess fully rational beliefs in equilibrium has been invoked

in other operations settings (see Liu and van Ryzin 2008 for a capacity-rationing model and Hassin and Haviv 2003 for queue-joining models). In §7.1, we extend our model to a dynamic setting and show that players can form rational expectations through adaptive learning (cf. Liu and van Ryzin 2007).

The next result characterizes the RE equilibrium and provides conditions under which it exists.

PROPOSITION 1. *There exists some $\bar{h} < v - c$ such that an RE equilibrium exists if and only if $h < \bar{h}$, where $\bar{h}$ increases in v and decreases in c. In any RE equilibrium, the seller's price p^* and quantity q^* satisfy* (i) $\bar{F}(q^*) = c/p^*$ *and* (ii) $(v - p^*)A(q^*) = h$. *Moreover, the seller's equilibrium profit is $\Pi^* = vE(X \wedge q^*) - cq^* - hE(X)$.*

This result implies that the market may break down when the search cost is above a certain threshold value $\bar{h}$. To avoid trivial outcomes, we impose the following assumption hereafter unless otherwise mentioned. This assumption ensures that the search cost h is below the threshold value, and guarantees the existence of an RE equilibrium.

ASSUMPTION 1. *Consumer search cost h satisfies $h < \bar{h}$.*

4. Value of Commitment

The RE equilibrium derived above assumes that consumers rely entirely on their rational expectations to draw inferences about product availability. In our model, suppose that the seller additionally makes explicit announcements regarding product availability. For example, the seller may announce to the market that he intends to stock quantity $\tilde{q}$ and maintain the availability probability $A(\tilde{q})$. How should consumers view these announcements $\tilde{q}$?

In our model above, consumers should simply ignore the seller's claims. This is because the seller has an incentive to convince consumers that service levels will be high, because higher availability implies higher willingness to pay. Therefore, the seller tends to exaggerate the announced quantity $\tilde{q}$, which is thus not credible. Consumers rely on their rational expectations, and the RE equilibrium above characterizes market outcomes.

Now, suppose there is some external commitment device that requires the seller to keep his word. That is, the seller's announced quantity $\tilde{q}$ must coincide with the true quantity q. The commitment device can be interpreted in several ways. For example, consumers may be able to directly monitor the seller's inventory and verify that the announcement $\tilde{q} = q$ is accurate. Alternatively, when stocking quantities and service levels are verifiable ex post, the seller may be averse to misrepresentation due to reputation concerns. Also, there may be legal institutions that enforce the seller's truth-telling constraint.

---

[2] In Definition 1, we have implicitly restricted attention to equilibria in which the market is functional; in other words, the seller stocks nonzero quantities and there are some consumers visiting the seller. Notice that there will always be another *nonparticipatory* equilibrium. Here, the trivial outcome is that the seller selects $q = 0$, and thus no consumer visits the seller. In fact, this outcome can always be supported by beliefs $\xi_{\text{prob}} = \xi_r = 0$. These are self-fulfilling prophecies that induce market failure solely because market participants do not believe it will work in the first place. When referring to an RE equilibrium in the rest of this paper, we are referring to a participatory equilibrium. Although there is always a nonparticipatory equilibrium, a participatory equilibrium may not exist.

We next analyze the market outcomes with seller commitment. In any equilibrium, consumers' beliefs ξ_{prob} must be consistent with the actual availability probability $A(q)$, which is determined by the seller's committed quantity q. Thus, they are willing to visit the seller as long as price does not exceed the reservation price $r = v - h/A(q)$. Corresponding to quantity q, it is optimal for the seller to charge price $p(q) = v - h/A(q)$. The seller's optimal quantity then follows from maximizing the following profit function $\Pi_C(q) = p(q)E(X \wedge q) - cq = (v - h/A(q))E(X \wedge q) - cq = vE(X \wedge q) - cq - hE(X)$. Here, the subscript C refers to the case with seller commitment. The following proposition characterizes market outcomes under seller commitment.

PROPOSITION 2. *Suppose that the seller can commit to quantity q. There exists $\bar{h}_C > \bar{h}$ such that an equilibrium exists if and only if $h < \bar{h}_C$. In equilibrium, the seller's price p_C^* and quantity q_C^* satisfy* (i) $\bar{F}(q_C^*) = c/v$ *and* (ii) $(v - p_C^*)A(q_C^*) = h$. *Moreover, the seller's equilibrium profit is $\Pi_C^* = vE(X \wedge q_C^*) - cq_C^* - hE(X)$.*

A direct comparison of Propositions 1 and 2 reveals that the seller's commitment power changes market outcomes dramatically. In fact, the next proposition demonstrates that commitment power is valuable to the seller. Under commitment, consumers are better able to assess their chances of securing the product if they do visit the seller (at a cost h). All else equal, this encourages consumer visits and increases their willingness to pay. This effect in turn increases profit margins and leads the seller to stock higher quantities, thus providing higher service levels in the first place. Therefore, both price and quantity are increased, and the overall effect on seller profits is positive.

PROPOSITION 3. *The equilibrium price, quantity, and seller profits are all higher with commitment. In other words, $p_C^* > p^*$, $q_C^* > q^*$, and $\Pi_C^* > \Pi^*$.*

The improved outcome is a direct consequence of commitment and cannot be sustained when the seller does not commit to quantity q. Consider the following hypothetical scenario without seller commitment. Suppose the seller sets price p_C^* and stocks quantity q_C^*, and consumers are willing to visit the seller. However, once consumers are willing to visit the seller at price p_C^*, the seller's profit function effectively becomes $\Pi(q) = p_C^*E(X \wedge q) - cq$ (i.e., with a fixed price), and the optimal stocking quantity is $\bar{F}^{-1}(c/p_C^*)$, which is less than $q_C^* = \bar{F}^{-1}(c/v)$. Therefore, when the seller does not commit to q_C^*, he has an incentive to decrease quantity below q_C^* once consumers are willing to pay p_C^*. This implies that consumers will not pay price p_C^* in the first place, and the outcome with price p_C^* and quantity q_C^* is not sustainable without seller commitment.

5. Availability Guarantees

Although we have seen that commitment power is valuable to the seller, a viable commitment device may not exist in practice. In such cases, is it possible for the seller to do better than the RE equilibrium outcome of §3? In particular, using the commitment case as a benchmark, is it possible to attain or even surpass the profit level Π_C^*? In this section, we show that availability guarantees can be used to stimulate demand and improve seller profits. We use the term "availability guarantee" to refer to mechanisms in which the seller promises to compensate consumers in the event of a stockout.

In practice, firms may offer availability guarantees that come in many different forms. For example, if a particular item is out of stock, retailers may offer to ship it to the consumer at no extra charge. This may be implemented through transshipment between retail locations, or from a central warehouse. As another example, consider movie rental companies. When a particular movie is stocked out, the company may offer a free rental on the same movie during a future visit. Finally, the seller may offer gift cards or discount coupons to placate consumers who do not find their desired items in stock. See Dong and Rudi (2004), Dana (2000), and Bhargava et al. (2006) for some anecdotal evidence of various forms of availability guarantees. Of course, the seller may also simply offer monetary compensation. In this section, we describe a general framework to analyze such availability guarantees.

We model availability guarantees as follows. Suppose the seller agrees to compensate customers who experience stockouts. Let u denote the consumer's utility from receiving the compensation, and let w denote the seller's cost of compensating one such consumer. We emphasize that these payoffs and costs are relevant only when stockouts occur. In our model, each utility-cost pair (u, w) represents a particular guarantee format. Observe that when the utility from stockout compensation u exceeds consumer search cost h, there may be a moral hazard problem in which consumers without any intention of buying still visit the seller with hopes of being compensated. Fortunately, firms will typically never offer guarantees that are inappropriately generous. Therefore, throughout the rest of our analysis, we assume that the seller knows that $u \leq h$ is satisfied. A sufficient (but not necessary) condition for this assumption is when the seller has an accurate estimate of h. In general, we distinguish between three different classes of availability guarantees: monetary compensation (where the seller pays cash to customers facing stockouts, so $w = u$), efficient guarantees (with $w \leq u$), and inefficient guarantees (with $w \geq u$). The choice between different formats for availability guarantees will be explored later.

Now, we first focus on a fixed type of guarantee (u, w) and investigate the RE equilibrium outcome under this guarantee.

Following the same reasoning as before, we see that, given beliefs ξ_{prob}, consumers' reservation prices r satisfy $(v - r)\xi_{\text{prob}} + u(1 - \xi_{\text{prob}}) = h$. On the other hand, considering the seller's cost and assuming that consumers are willing to buy, the seller's profit function is $\Pi(p, q) = pE(X \wedge q) - cq - w[E(X) - E(X \wedge q)]$, so the seller stocks the critical fractile quantity $q = \bar{F}^{-1}(c/(p + w))$. Together with consistency conditions, we can then characterize the resulting RE equilibrium corresponding to the (u, w) guarantee. Our next result gives the equilibrium conditions. The subscript G refers to availability guarantees.

PROPOSITION 4. *With availability guarantee (u, w), the RE equilibrium (p_G^*, q_G^*) is given by (i) $\bar{F}(q_G^*) = c/(p_G^* + w)$ and (ii) $(v - p_G^*)A(q_G^*) + uA(q_G^*) = h$. Moreover, the seller's equilibrium profit is $\Pi_G^* = vE(X \wedge q_G^*) - cq_G^* - hE(X) + (u - w)[E(X) - E(X \wedge q_G^*)]$. When $w = u$, there is a threshold value $\bar{h} < \bar{h}_G \leq \bar{h}_C$ such that an equilibrium exists if and only if $h < \bar{h}_G$.*

Note that availability guarantees may raise the threshold search cost for the existence of equilibrium, i.e., $\bar{h}_G > \bar{h}$ for the benchmark case $w = u$. Clearly, the threshold search cost will be higher with $u > w$ and lower with $u < w$. Furthermore, in the presence of search costs and demand uncertainty, availability guarantees serve as a useful contractual device to stimulate demand. We offer two intuitive explanations for why such guarantees are attractive measures to adopt. First, by offering a monetary compensation, the seller can charge a higher price and extract more profit. This is similar in spirit to some signaling games where firms "burn cash" or impose self-punishment; for example, Milgrom and Roberts (1986) use price and advertising expenditure to signal quality, and Dana (2001) uses price to signal availability. In our model, by offering availability guarantees, the seller is inflating his own costs in the event of stockouts. Doing so indirectly convinces consumers that the service level will be high, because stockouts are costly to the seller. Indeed, observe from (i) in Proposition 4 that the seller's underage cost is increased as a result of w. Moreover, note that the possibility of receiving u in the event of stockouts increases consumers' willingness to pay, raises the seller's profit margins, and creates incentives for the seller to improve service levels, thus reinforcing consumer faith in the seller. Second, because consumers recognize the possibility of stockouts before visiting the seller, they will never pay full valuation for the product. Indeed, to induce demand, the seller must pay consumers a premium to cover the possibility of facing a stockout. This could take two forms: an ex ante premium (in the form of a reduction in price) or an ex post premium (in the form of the stockout compensation). Notice that ex ante premia are always enjoyed by all consumers, whereas ex post premia apply only when stockouts occur. Thus, from the perspective of the seller, offering availability guarantees is a more cost-effective way to stimulate demand, compared to reducing the price.

5.1. Monetary Compensation
We next consider the special case of monetary compensation with $w = u$. That is, the seller offers a fixed monetary compensation $u > 0$ to consumers who experience a stockout. With this setup, we consider the following questions. First, how do monetary guarantees affect the equilibrium price and quantity? Second, how much compensation u should the seller offer? Specifically, should the seller offer maximal compensation $u = h$? Third, with an optimal compensation scheme, how much does the seller make, compared to our previous scenarios? The following two propositions provide answers to all these questions. The subscript M stands for monetary compensation.

PROPOSITION 5. *The seller's equilibrium price p_M^*, quantity q_M^*, and profit Π_M^* are increasing in u.*

PROPOSITION 6. *Suppose the seller offers maximal monetary compensation $u = h$. Then, the RE equilibrium satisfies $p_M^* = v - h > p_C^*$, $q_M^* = q_C^* = \bar{F}^{-1}(c/v)$, and $\Pi_M^* = \Pi_C^*$.*

These results demonstrate that the seller should offer maximal monetary compensation $u = h$. There are two noteworthy implications. First, this implies that the seller should fully reimburse the search cost to consumers who face stockouts. This places all consumers in an enviable risk-free situation: Regardless of whether each consumer finds the product in stock, the consumer will never incur any losses ex post. In fact, by offering to cover search costs, the seller may now price at consumers' full valuation (net of search costs), i.e., $p = v - h$; note that this is the highest possible price that the seller may charge. Therefore, with maximal compensation $u = h$, potential consumer losses are completely eliminated, but it is the seller who then extracts the entire consumer surplus. Second, the seller can attain the profit under commitment Π_C^* through availability guarantees. Without any monetary compensation, the seller has to contend with the RE equilibrium profit Π^*; in contrast, with maximal monetary compensation $u = h$, the seller is able to attain profit level Π_C^*. This suggests that monetary compensation can serve as a proxy for commitment power, both of which can be used by the seller to influence consumer behavior to enhance profit.

5.2. Efficient Guarantees
We shall now consider more general forms of availability guarantees beyond monetary compensation

Figure 1   **Different Formats of Availability Guarantees**

with $w = u$. Typically, firms have the opportunity to exploit cost savings by offering guarantees with $w \leq u$, such that consumers who face stockouts derive utility u from the compensation, but this costs the firm only w. This is possible, for instance, by offering free gifts or discount coupons because the firm's own profit margin drives a wedge between the two parameters u and w. However, because cost savings are not limitless, to maintain realism, we assume that to provide stocked-out consumers with utility u, there is a minimum cost $\underline{W}(u)$ that the firm must incur. We shall refer to this as the efficient frontier, and we assume that $\underline{W}(u)$ is increasing and convex in u. Figure 1 shows the feasible formats of availability guarantees (i.e., different (u, w) combinations) from which the firm can choose. The bold curve represents the efficient frontier, below which are infeasible (u, w) combinations; similarly, we do not allow $u > h$ to prevent consumer moral hazard. The dashed 45° line corresponds to guarantees based on monetary compensation, as analyzed previously. Feasible combinations below this line are efficient guarantee formats ($w \leq u$), whereas those above are inefficient formats ($w \geq u$).

As a first step in our analysis, we consider the following guarantee format: $u = h$ and $w = \underline{W}(h)$, i.e., point B in Figure 1. This maximal compensation fully insures consumers for their search cost h, and also exploits cost savings to some extent. In particular, these cost savings provide additional room for the seller to do better than the commitment benchmark of §4. We have already seen that this profit benchmark Π_C^* can be attained through the monetary compensation $u = h$. The next result confirms that exploiting cost savings allows the seller to do even better.

PROPOSITION 7. *Suppose the seller offers an availability guarantee with $u = h$ and $w = \underline{W}(h)$. Then, the RE equilibrium satisfies $p_G^* = v - h \geq p_C^*$, $q_G^* \leq q_C^*$, and $\Pi_G^* \geq \Pi_C^*$.*

This result shows that, similar to the case with maximal monetary compensation $u = h$, the equilibrium price also reaches the highest possible price of $v - h$, which exceeds the equilibrium price p_C^* in the commitment benchmark. However, different from the case with maximal monetary compensation $u = h$ is that the equilibrium quantity here is lower than q_C^* under commitment. This is because the cost saving discussed above decreases the seller's underage cost and induces him to stock less. Nevertheless, with lower stocking quantity, the seller still attains higher profits, compared to both the maximal monetary compensation $u = w = h$ case and the commitment benchmark case.

Finally, we close this section by analyzing the seller's optimal choice of guarantee formats (u, w). We introduce some notation. Let (u^{**}, w^{**}) denote the feasible (u, w) pair that maximizes the difference $u - w$. This can be interpreted as the socially efficient availability format, which generates the highest net benefit when stockouts occur. This socially efficient point is marked as point A in Figure 1.

PROPOSITION 8. *Let the seller's profit-maximizing guarantee format be (u^*, w^*). Then, we have (i) $u^* \geq u^{**}$ and (ii) $w^* = \underline{W}(u^*)$.*

Interestingly, this proposition shows that although the profit-maximizing seller will fully exploit cost savings by choosing a guarantee format on the efficient frontier (part ii), he has an incentive to over compensate (part i). When left to his own devices, the seller's availability guarantees are too generous because the compensation paid to consumers during stockouts u^* exceeds the socially efficient level u^{**}. Therefore, the marginal value of the last unit of compensation provided by the seller is negative. To see this, observe that the slope of the efficient frontier at the seller's choice u^* is $\underline{W}'(u^*) \geq 1$. The consumer values this marginal unit of compensation at less than the resources required to provide it, and this leads to deadweight losses. This inefficiency can be attributed to the standard monopoly problem, wherein the monopolist sets too high a price; in our context, an excessively generous guarantee is offered to realize a high price.

5.3. Inefficient Guarantees

We provide a brief discussion about inefficient guarantees in this section. A direct consequence of our result in Proposition 8 is that inefficient guarantees are suboptimal for the seller. Although this finding is not surprising, we wish to caution readers about the possible limitations of this result. After all, it is conceivable that offering inefficient guarantees may allow the seller to better commit to providing high service levels, because the cost of compensating consumers is now even higher (relative to using efficient guarantees). We conjecture that such commitment strategies,

via offering inefficient guarantees, may be desirable in competitive environments. In a different context, Kim et al. (2001) discuss efficient versus inefficient rewards for customer loyalty.

6. First-Best

In this section, we take on a social welfare perspective. Assuming that there is a social planner who is able to influence the actions of consumers and firms (prior to demand realization), what should this social planner do? This is in contrast to the equilibrium analyses from previous sections, in which consumers and the seller were assumed to maximize individual payoffs. By comparing the outcomes of individual versus social optimization, our goal is to identify the sources of efficiency loss and to design mechanisms that achieve first-best.

We begin our analysis by defining the social welfare function. Recall that the seller sets the price p and quantity q, whereas the consumers select reservation prices r. In addition, the seller may choose the type of availability guarantee (u, w) to provide. Because prices merely constitute transfer payments and consumers need to visit the seller to realize nonzero social surplus, we may disregard p and r and focus on the social planner's choices of q and (u, w). Accordingly, the social welfare function may be expressed as $SW(q, u, w) = vE(X \wedge q) - cq - hE(X) + (u - w)[E(X) - E(X \wedge q)]$. Each of the four terms, respectively, correspond to (i) expected valuation enjoyed by the consumer population, (ii) production cost, (iii) search cost incurred by all consumers, and (iv) expected social surplus generated from compensating consumers facing stockouts. With this expression, we can find the first-best quantity q^{FB} and guarantee format (u^{FB}, w^{FB}) that maximize social welfare.

PROPOSITION 9. *The first-best quantity and guarantee format are given by* (i) $\bar{F}(q^{FB}) = c/(v - (u^{**} - w^{**}))$ *and* (ii) $(u^{FB}, w^{FB}) = (u^{**}, w^{**})$.

In general, the first-best compensation utility is below the maximal level, i.e., $u^{FB} = u^{**} \leq h$, and the first-best quantity satisfies $q^{FB} \leq \bar{F}^{-1}(c/v)$. To see the significance of these comparisons, let us consider two hypothetical scenarios, one where the seller bears the entire burden of potential stockouts, and another where consumers bear the entire burden. The former occurs when the seller offers maximal compensation $u = h$, which ensures that consumers never suffer negative ex post payoffs. The latter corresponds to the situation in which all consumers unconditionally visit the seller and are willing to pay their full valuation v, which then leads the seller to stock quantity $q = \bar{F}^{-1}(c/v)$. Observe that the first-best outcome falls between these two extremes. This suggests that social optimality requires the seller and the consumers

to jointly share the potential downside of supply-demand mismatches.

Unsurprisingly, we find that the first-best outcome is not attained in any of the scenarios considered above. In the absence of availability guarantees (i.e., when $u = 0$), the first-best is not achieved. This applies to our basic model, as well as our commitment benchmark. However, even when availability guarantees are considered, the RE equilibrium still does not sustain first-best outcomes. When we focus on monetary compensation, the seller's profit-maximizing choice of u is $u = h$. Although this attains the commitment benchmark outcomes, the first-best is not achieved because potential cost savings from alternative types of availability guarantees have not been made use of. When we consider efficient guarantees, recall our earlier finding that the seller's profit-maximizing compensation level exceeds the socially efficient level, i.e., $u^* \geq u^{FB}$. Again, the first-best is not attained in equilibrium.

Fortunately, as we illustrate below, the first-best can be attained in RE equilibrium by using a combination of commitment and availability guarantees. To see this, consider the situation in which the seller can commit to a stocking quantity q, and on top of this, the seller provides availability guarantee of type (u, w). In this case, given the seller's quantity q, consumers form their reservation prices r via the expression $(v - r)A(q) + u(1 - A(q)) = h$. Therefore, corresponding to quantity q, the seller charges the price $p(q)$, satisfying $(v - p(q))A(q) + u\bar{A}(q) = h$. It then follows that the seller's profit function is $\Pi_B(q, u, w) = p(q)E(X \wedge q) - cq - w[E(X) - E(X \wedge q)] = (v - u - (h - u)/A(q)) \cdot E(X \wedge q) - cq - w[E(X) - E(X \wedge q)] = vE(X \wedge q) - cq - hE(X) + (u - w)[E(X) - E(X \wedge q)] = SW(q, u, w)$, which coincides with the social welfare function. In this way, the seller's individual interests are aligned with social welfare, so he chooses $q_B^* = q^{FB}$ and $(u_B^*, w_B^*) = (u^{FB}, w^{FB})$, thus attaining social welfare. Here, the subscript B refers to "both commitment and availability guarantees." The following proposition summarizes these findings.

PROPOSITION 10. *Suppose that the seller can commit to quantity and implements a profit-maximizing availability guarantee format. Then, in RE equilibrium, the first-best is attained, with* $q_B^* = q^{FB}$ *and* $(u_B^*, w_B^*) = (u^{FB}, w^{FB})$. *Moreover, the equilibrium price is* $p_B^* = (v - u^{FB} - (h - u^{FB})/A(q^{FB})) \leq v - h$. *Finally, the seller extracts the entire social surplus, and consumers' expected surplus is zero.*

The main thrust of this result is to confirm that the seller can indeed implement the first-best in an RE equilibrium. Doing so requires a combination of two measures: (i) a commitment device for the seller to provide credible inventory information prior to

demand realization (i.e., the ex ante measure) and (ii) an incentive mechanism to induce demand by compensating consumers for potentially adverse outcomes afterward (i.e., the ex post measure). In isolation, each of the two strategies proposed above (i.e., establishing commitment power and providing availability guarantees) does not achieve first-best outcomes, but when appropriately combined, their inherent synergies can allow the seller to capture the entire first-best surplus.

7. Extensions

7.1. Learning

This section studies a multiperiod extension with consumer learning. In the static model of §3, the RE equilibrium outcomes are natural and intuitive, provided that players possess rational expectations. In other words, our basic model applies in the long run when equilibrium beliefs have formed and stabilized. Now, we also consider the transient setting. We outline a dynamic model of consumer learning and investigate whether market outcomes and beliefs indeed converge toward the static RE equilibrium analyzed earlier.

In this model extension, the seller interacts with a fixed market of consumers over discrete time periods $t = 1, 2, \ldots$. In each period t, consumers have some belief $\hat{\xi}_t$ over the availability probability. Similar to the basic model, these beliefs determine consumers' reservation price, so the seller can charge up to price $p_t = v - h/\hat{\xi}_t$. The seller also chooses a stocking quantity q_t, so the actual availability probability in period t will be $\xi_t = A(q_t)$. At the end of each period, the consumer population updates the belief for the next period to be $\hat{\xi}_{t+1} = \alpha \xi_t + (1-\alpha)\hat{\xi}_t$ for some $\alpha \in (0, 1)$. Such an exponential learning method has been widely used in the literature (see Liu and van Ryzin 2007 and the references therein). Although each period in this dynamic model can be viewed as a separate replica of our basic model, consumers' beliefs in one period do influence their beliefs in the next period.

Now the seller faces a dynamic optimization problem instead of a static newsvendor problem. Let $\Pi(q, \hat{\xi})$ denote the seller's profit from stocking q units in a period when consumer beliefs are $\hat{\xi}$. Because the seller charges price $p = v - h/\hat{\xi}$, the expected profit in this period is $\Pi(q, \hat{\xi}) = (v - h/\hat{\xi})E(X \wedge q) - cq$. The seller wishes to maximize total discounted profit over an infinite horizon $\sum_{t=1}^{\infty} \delta^{t-1}\Pi(q_t, \hat{\xi}_t)$, where the discount factor $\delta \in (0, 1)$. Therefore, we can write down the Bellman equations for the seller's dynamic program as follows: $V(\hat{\xi}) = \max_q[\Pi(q, \hat{\xi}) + \delta V(\alpha A(q) + (1-\alpha)\hat{\xi})]$. Then, the maximizer $q^*(\hat{\xi})$ in the equation above yields the seller's optimal quantity decision when consumer beliefs are $\hat{\xi}$. Furthermore, we may interpret the value function $V(\hat{\xi})$ as the seller's optimal total discounted profit, given current beliefs $\hat{\xi}$.

Under the seller's optimal policy, his quantity decisions q_t^* and consumers' beliefs $\hat{\xi}_t$ in each period t are governed by the following dynamic system (with initial beliefs $\hat{\xi}_1$): $q_t^* = q^*(\hat{\xi}_t)$ and $\hat{\xi}_{t+1} = \alpha A(q_t^*) + (1-\alpha)\hat{\xi}_t$ for $t = 1, 2, \ldots$. We are interested in the long-run behavior of this dynamic system. Do the seller's optimal quantities q_t^* and consumers' beliefs $\hat{\xi}_t$ converge to some steady state? The next proposition shows that the answer is yes and characterizes the long-run equilibrium quantity and belief.

PROPOSITION 11. *Let the initial belief be* $\hat{\xi}_1$. *There exists some threshold* $z \in [0, 1]$ *such that*
   (i) *if* $\hat{\xi}_1 < z$, *then* $q_t^* \to 0$ *and* $\hat{\xi}_t \to 0$;
   (ii) *if* $\hat{\xi}_1 > z$, *then* $q_t^* \to \tilde{q}$ *and* $\hat{\xi}_t \to \tilde{\xi}$, *where* $\tilde{\xi} = A(\tilde{q}) \in (0, 1)$ *and* $\tilde{q} > 0$ *is the larger solution of*

$$\left(v - \left(\frac{1-\delta}{1-(1-\alpha)\delta} \right) \frac{h}{A(q)} \right)\bar{F}(q) - c = 0. \quad (1)$$

In particular, under Assumption 1 (i.e., $h < \bar{h}$), *we have* $z < 1$.

This proposition shows that in our model of adaptive learning, there are two possible long-run equilibria. In case (i) of the proposition, quantities and beliefs converge toward zero, so the market eventually breaks down in the long run. In case (ii), however, there is some nonzero equilibrium that is eventually attained and then persists over time. Let us call this equilibrium (i.e., quantity $\tilde{q}$, consumer beliefs $\tilde{\xi} = A(\tilde{q})$, and price $\tilde{p} = v - h/\tilde{\xi}$) the *long-run equilibrium*. The proposition above also shows that this long-run equilibrium will emerge as long as consumer search cost h is not too large and consumers are sufficiently optimistic at the start (i.e., $\hat{\xi}_1$ is sufficiently large).

Interestingly, the long-run equilibrium coincides with a static RE equilibrium in our basic model. The next result draws this connection.

PROPOSITION 12. *Let* $q^*[h]$ *be the static RE equilibrium quantity when consumers face search cost* h. *Then, the long-run equilibrium quantity* $\tilde{q}$ *satisfies* $\tilde{q} = q^*[((1-\delta)/(1-(1-\alpha)\delta)) \cdot h]$.

This proposition shows that the long-run equilibrium in our dynamic model is precisely the RE equilibrium for our static model, with some appropriately adjusted consumer search cost. When the discount factor is δ, the long-run quantity $\tilde{q}$ matches the static RE quantity q^* with an adjusted search cost of $((1-\delta)/(1-(1-\alpha)\delta)) \cdot h \leq h$. This downward adjustment suggests that in the dynamic equilibrium, consumers are equivalently playing the static RE equilibrium, but with a smaller search cost. The intuition

is as follows. In a dynamic setting, the seller's quantity decision in any given period influences future beliefs and future profits, so there is a greater incentive to maintain high service levels. Such reputation effects decrease stockout experiences, so consumers can behave *as if* their search costs were lowered. From a modeling perspective, this result provides a tight link between the static and dynamic models. This link demonstrates that players can indeed reach a static RE equilibrium through a simple adaptive learning process.

Finally, notice that as the discount factor δ approaches 1 from below (recall that $\delta < 1$), the long-run equilibrium approaches the commitment benchmark. To see this, note that $(1-\delta)/(1-(1-\alpha)\delta) \to 0$ as $\delta \to 1$, so from Equation (1) of Proposition 11, the long-run quantity $\tilde{q} \to \bar{F}^{-1}(c/v)$, which is the seller's optimal quantity with commitment (cf. Proposition 2). This comparison suggests that in a dynamic setting with consumer learning, the seller's reputation concerns (when δ is close to, but less than, 1) can serve as a surrogate for commitment power.

7.2. A Multiproduct Model
This section demonstrates that the basic model can be extended to incorporate multiple products. Suppose the seller offers consumers a set of products indexed by $i = 1, \ldots, n$. We assume that the set of products is fixed. For instance, a retailer sells shirts from a fixed set of different sizes. There is a pool of consumers, and each consumer demands one unit of the product. However, a consumer does not know which product is suitable before she visits the seller, e.g., a consumer needs to try the shirt on to know the right size. Assume that there is at most one suitable product for each consumer. In other words, a lost sale occurs if a consumer's first choice is unavailable. This model has been used in van Ryzin and Mahajan (1999) and is a special case of the models studied by Smith and Agrawal (2000) and Mahajan and van Ryzin (2001). The consumers are divided into M types according to certain characteristics (e.g., height and weight). The fraction of type j $(1 \le j \le M)$ consumers is denoted by α_j $(\sum_j \alpha_j = 1)$ and let λ_{ij} denote the probability that product i is suitable for a type-j consumer $(\sum_i \lambda_{ij} = 1$ for all $j)$. The products share the same unit cost c. For ease of exposition, we assume that all consumers have an identical valuation v for a suitable product and zero otherwise.

First, consider the consumer's problem. If a type-j consumer chooses to visit the seller, she will get an expected surplus $\sum_i \lambda_{ij}(v - p_i)\xi_i - h$, where ξ_i is the availability probability and p_i is the price for product i. She receives zero surplus if she does not visit the seller.

Next, consider the seller's problem. Assume that all consumers choose to visit the seller. Then, the demand

for product i is $(\sum_j \alpha_j \lambda_{ij})X$. Define $\lambda_i = \sum_j \alpha_j \lambda_{ij}$ as the fraction of demand for product i. Because $\sum_j \alpha_j = 1$ and $\sum_i \lambda_{ij} = 1$, we know $\sum_i \lambda_i = 1$. Given the price p_i, the optimal stocking quantity Q_i is given by $\Pr(\lambda_i X \le Q_i) = (p_i - c)/p_i$. Given Q_i, the availability probability is then $\xi_i = E(\lambda_i X, Q_i)/E(\lambda_i X) = E(X, Q_i/\lambda_i)/E(X) = A(Q_i/\lambda_i)$, which only depends on p_i. For convenience, define $\hat{Q}_i = Q_i/\lambda_i$ as the *effective* quantity. Thus, once p_i is chosen, both the seller's optimal effective quantity $\hat{Q}_i$ (in equilibrium) and the availability probability $\xi_i = A(\hat{Q}_i)$ are determined.

Recall that in the basic model with a single product, the price p is determined in the RE equilibrium. In particular, the seller would charge a price such that the consumer's expected surplus is zero, i.e., $(v - p)\xi - h = 0$. Similarly, with multiple products, the seller will choose a price such that consumer surplus $\sum_i \lambda_{ij}(v - p)\xi_i - h = 0$ for all consumer types j. We assume that the seller charges a uniform price p across all product variants (e.g., the price is the same for shirts in different sizes or yogurt with different flavors). Although in theory the seller can price discriminate, in practice, such price differentials are often perceived as unfair and will attract consumer resentment (see Anderson and Simester 2008). We have the following equilibrium result.

PROPOSITION 13. *Suppose the seller charges a uniform price for all products. Then, the seller will choose the same $\hat{Q}$ for all products, and the RE equilibrium outcome $(p, \hat{Q})$ is characterized by* (i) $\bar{F}(\hat{Q}) = c/p$ *and* (ii) $(v - p)\xi(\hat{Q}(p)) - h = 0$.

This proposition indicates that the multiproduct model can be reduced to the single-product model studied in the previous sections. Therefore, the results obtained for the basic model also apply to the multiproduct setting.

7.3. Heterogeneous Consumer Valuations
This section extends the basic model in §3 to consider heterogeneous consumers. Now consumer valuation is continuously distributed on a support $[0, v]$. We index consumers by θ such that consumers' valuation v_θ increases in θ. Let $\Phi(v_\theta)$ be the proportion of consumers with a valuation v_θ or lower. (Note that the homogeneous model is a special case with a point distribution at v.) The total market demand is still X with a distribution function F, which is independent of the distribution Φ. Therefore, if the consumer segment that has a valuation v_θ or higher chooses to visit the seller, the seller's random demand will be $(1 - \Phi(v_\theta))X$. As before, there is a search cost h for all consumers. The seller sets an observable price p and an unobservable quantity q before the random demand X is realized. The consumers decide whether to visit the seller or not.

Suppose the consumers observe the seller's price and form a belief of the availability probability. Let θ be the marginal consumer who is indifferent between visiting and not visiting the seller. Define $s(v_\theta) = 1 - \Phi(v_\theta)$ as the market base for the seller (i.e., the fraction of the market that will choose to visit the seller). For consistency and fair comparison, we assume proportional rationing here as in the basic model: all consumers who choose to visit the seller have equal chances of getting the product. Under proportional rationing, in any RE equilibrium, the belief of availability probability ξ_{prob} must agree with the actual outcome corresponding to the seller's quantity q and the marginal customer's valuation v_θ: $\xi_{\text{prob}} = E(s(v_\theta)X \wedge q)/E(s(v_\theta)X) = E[X \wedge (q/s(v_\theta))]/E(X) = A(z)$. Here, we define $z \equiv q/s(v_\theta)$ and recall that the availability probability is $A(z) = E[X \wedge z]/E(X)$ as above.[3]

Suppose the seller offers an availability guarantee with utility-cost pair (u, w). For the same reasons in §5, we focus on $0 \le u \le h$. Below we derive and compare the RE equilibria under two scenarios, i.e., the scenarios with and without commitment. The following proposition characterizes the conditions for the existence of RE equilibrium without commitment. The subscript H refers to heterogeneous consumers.

PROPOSITION 14. *Consider the heterogeneous-consumer RE game without quantity commitment. Under availability guarantee (u, w):*

(a) There exists $\bar{h}_H \in (0, h]$ such that an RE equilibrium exists only if $h - u < \bar{h}_H$.

(b) For $h - u < \bar{h}_H$, let $z_1 < z_2$ be the two solutions to $v - c/\bar{F}(z) + (u - h)/A(z) = 0$ and z_3 be the unique solution to $(c/\bar{F}(z))E(X \wedge z) - cz - wE(X) = 0$. Then, an RE equilibrium exists if and only if $z_2 > z_3$. In any equilibrium, the seller's price p, quantity q, and the marginal valuation v_θ (and thus the market base $s(v_\theta)$) satisfy (i) $(v_\theta - p)A(q/s(v_\theta)) + u(1 - A(q/s(v_\theta))) = h$ and (ii) $\bar{F}(q/s(v_\theta)) = c/(p + w)$.

The above proposition states that an RE equilibrium does not exist in the heterogeneous-consumer

model when the search cost h is too high. Similar to the homogeneous model, a stockout compensation $u > 0$ might be needed to prevent market failure. The proposition also indicates that multiple RE equilibria may exist in the heterogeneous-consumer model. In fact, for any particular z ($z_3 < z < z_2$), there exists a corresponding RE equilibrium $(p, z, s(v_\theta))$ that satisfies (i) and (ii) above and yields a positive profit for the seller. This means that in the heterogeneous-consumer model, the seller first determines the optimal value for z (or the market base $s(v_\theta)$), and then plays the corresponding RE game with the consumers.

Next we consider the scenario with commitment. In this scenario, the seller can freely choose both price p and quantity q. The market base $s(v_\theta)$ is then given by the consumers' indifference condition, i.e., $(v_\theta - p)A(q/s(v_\theta)) + u\bar{A}(q/s(v_\theta)) = h$. This is in contrast with the no-commitment scenario above, where the seller can only choose the market base, and the rest of the decisions are governed by the equilibrium conditions. Note there is a one-to-one relationship between q and $z = q/s(v_\theta)$. Thus, solving the RE equilibrium is equivalent to identifying the optimal p and z that maximize the seller's profit.

With commitment power, the seller can replicate any RE equilibrium that may arise without commitment. This implies that the seller's profit in the commitment scenario is always higher than that in the no-commitment scenario. Furthermore, plugging $u = h$ into condition (i) in Proposition 14 gives $v_\theta - p = h$, which does not contain z. In this case, the z decision is decoupled from the p decision in the no-commitment scenario. As a result, the seller's optimal decisions (including the corresponding profit) in the no-commitment scenario will be the same as those in the commitment scenario. Therefore, we know that the seller's profit in the no-commitment scenario converges to the profit in the commitment scenario from below as u approaches h (see Figure 2 for an example).

What is the impact of availability guarantees on the seller's profit and social welfare? Do the results from the homogeneous-consumer model continue to hold in the heterogeneous-consumer model? The rest of this section aims to answer the above questions. Note

---

[3] Another rationing scheme is efficient (or parallel) rationing, which assumes that products are assigned to consumers in decreasing order of their valuations. Under efficient rationing, the marginal consumer (lowest-valuation consumer who chooses to visit the seller) has the lowest priority. Thus, the marginal consumer's availability probability is $G(z) \equiv \int_0^z g(x)\,dx$, where $g(\cdot)$ is the individual consumer's posterior demand density as defined in §3. It can be shown that the availability probability is lower under efficient rationing as compared to proportional rationing, i.e., $G(z) \le A(z)$. This implies that, ceteris paribus, efficient rationing will lead to a smaller market base for the seller. We conjecture that a similar analysis for the case of efficient rationing can be done with $G(z)$ in place of $A(z)$, but a full investigation is beyond the scope of this paper. See also Dana (1998), who studies both proportional and efficient rationing schemes in a capacity pricing model.

Figure 2   Seller Profit and Social Welfare Functions
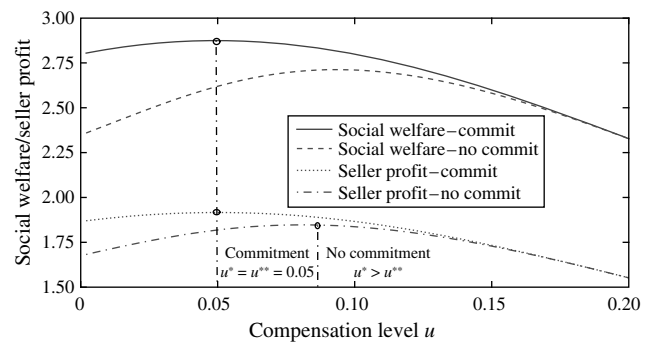
that with heterogeneous consumer valuations, it is infeasible for the seller to extract the entire social surplus by charging a single price. Nevertheless, we may still examine whether the seller's individual interests are aligned with social welfare. The following proposition demonstrates that the seller will choose the socially efficient guarantee with commitment.

PROPOSITION 15. *Consider the heterogeneous-consumer RE game with quantity commitment. Let the seller's profit-maximizing guarantee format be (u^*, w^*). Then, $u^* = u^{**}$ and $w^* = \underline{W}(u^*)$.*

Characterizing the seller's optimal guarantee in the no-commitment scenario is analytically challenging, so we use extensive numerical experiments to derive additional insight. In particular, we are interested in how the seller's profit and social welfare depend on the compensation level u in both scenarios. For sake of brevity, we present the results from a representative numerical example. In this example, Φ is uniform on $[0, 1]$ (i.e., $v = 1$, without losing generality) and other parameter values are as follows: $c = 0.4$, $h = 0.2$, the cost of availability guarantee is $\underline{W}(u) = 10u^2$, and X follows a gamma distribution with mean 90 and standard deviation 30. The socially efficient compensation level is $u^{**} = 0.05$. Figure 2 shows how the seller's profit and the corresponding social welfare vary in the compensation level, u. We can see from Figure 2 that both the seller's profit and the social welfare are higher with quantity commitment than their counterparts without commitment. In addition, there is $u^* = u^{**} = 0.05$ with commitment, i.e., the seller's profit-maximizing compensation level is socially efficient, which also delivers the highest level of social welfare. However, the seller will overcompensate (i.e., $u^* > u^{**}$) in the no-commitment scenario. From the numerical study, we observe that quantity commitment and availability guarantees can improve social welfare as well as the seller's profit in the model with consumer heterogeneity. Also, similar to our basic model, both quantity commitment and availability guarantees are required to realign the seller's individual interests with social welfare.

8. Conclusion

In this paper, we study the impact of product availability on market outcomes. High service levels attract consumer demand but must correspondingly be supported by the firm's inventory investment. We develop a modeling framework based on the rational expectations paradigm, and characterize the seller's equilibrium price and inventory levels when facing rational consumers who may change their shopping behavior in response to the anticipated likelihood of stockouts. We study two strategies that can be used to attract consumers. The first strategy is to use an external *commitment* device, i.e., the seller credibly

promises to keep service levels high; the second is to offer *availability guarantees*, i.e., the seller promises to compensate consumers in the event of a stockout. Our main conclusion is the seller may enhance profits by using either strategy and may additionally realize first-best outcomes via a combination of both strategies.

This research can be extended in several directions. First, we have considered only a single product, whereas in practice, most firms offer multiple substitutable products, which offer opportunities for substitution when stockouts occur. For example, consumers at a video rental store may switch to their second-choice movie when their first-choice movie is not available. Similarly, fashion apparel that come in different designs and colors offer additional options when stockouts occur. This suggests that with ex post consumer substitution, product variety can serve as a strategic lever for firms to alleviate concerns about product availability. A complete investigation is left for future research. Second, consumer heterogeneity is present in most markets. Although we have briefly studied the case where consumers have different valuations, it would also be interesting to consider the case where consumers have different search costs. The interaction (correlation) between these two dimensions of consumer heterogeneity may also have strategic implications on firms' operational policies. Finally, our analysis has assumed that consumers are "ultra-rational" in that they form rational expectations and make optimal decisions with perfect accuracy. How can our model be extended to capture estimation and decision errors? We believe that incorporating consumer bounded rationality (cf. Su 2008) is a fruitful area for future research.

9. Electronic Companion

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

References

Alexandrov, A., M. A. Lariviere. 2007. Are reservations recommended? Working paper, Northwestern University, Evanston, IL.

Andersen Consulting. 1996. Where to look for incremental sales gains: The retail problem of out-of-stock merchandise. The Coca-Cola Retailing Research Council, Atlanta.

Anderson, E. T., D. Simester. 2008. Does demand fall when customers perceive that prices are unfair? The case of premium pricing for large sizes. *Marketing Sci.* **27**(3) 492–500.

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

Aviv, Y., A. Pazgal. 2008. Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing Service Oper. Management* **10**(3) 339–359.

Balachander, S., P. H. Farquhar. 1994. Gaining more by stocking less: A competitive analysis of product availability. *Marketing Sci.* **13**(1) 3–22.

Besanko, D., W. L. Winston. 1990. Optimal price skimming by a monopolist facing rational consumers. *Management Sci.* **36**(5) 555–567.

Bhargava, H. K., D. Sun, S. H. Xu. 2006. Stockout compensation: Joint inventory and price optimization in electronic retailing. *J. Comput.* **18**(2) 255–266.

Breugelmans, E., K. Campo, E. Gijsbrechts. 2006. Opportunities for active stock-out management in online stores: The impact of the stock-out policy on online stock-out reactions. *J. Retailing* **82**(3) 215–228.

Bulow, J. I. 1982. Durable-goods monopolists. *J. Political Econom.* **90**(2) 314–332.

Cachon, G. P., R. Swinney. 2009. Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Sci.* **55**(3) 497–511.

Coase, R. H. 1972. Durability and monopoly. *J. Law Econom.* **15**(1) 143–149.

Consumer Reports. 1987. Mail-order companies. *Consumer Reports* **52** 607–614.

Dana, J. D. 1998. Advance-purchase discounts and price discrimination in competitive markets. *J. Political Econom.* **106**(2) 395–422.

Dana, J. D. 2000. Blockbuster video. Kellogg School of Management, Northwestern University, Evanston, IL.

Dana, J. D. 2001. Competition in price and availability when availability is unobservable. *RAND J. Econom.* **32**(3) 497–513.

Dana, J. D., N. C. Petruzzi. 2001. The newsvendor model with endogenous demand. *Management Sci.* **47**(11) 1488–1497.

Dana, J. D., K. E. Spier. 2001. Revenue sharing and vertical control in the video rental industry. *J. Indust. Econom.* **49**(3) 223–245.

Deneckere, R., J. Peck. 1995. Competition over price and service rate when demand is stochastic: A strategic analysis. *RAND J. Econom.* **26**(1) 148–162.

Dong, L., N. Rudi. 2004. Who benefits from transshipment? Exogenous vs. endogenous wholesale prices. *Management Sci.* **50**(5) 645–657.

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

Furman, P. 1998. At Blockbuster video, a fast fix moves flicks. *Daily News* (July 27).

Gruen, T. W., D. Corsten. 2008. A comprehensive guide to retail out-of-stock reduction in the fast-moving consumer goods industry. Grocery Manufacturers of America, Washington, D.C.

Gruen, T. W., D. Corsten, S. Bharadwaj. 2002. Retail out of stocks: A worldwide examination of causes, rates, and consumer responses. Grocery Manufacturers of America, Washington, D.C.

Hann, I.-H., C. Terwiesch. 2003. Measuring the frictional costs of online transactions: The case of a reverse buying site. *Management Sci.* **49**(11) 1563–1579.

Hassin, R., M. Haviv. 2003. *To Queue or Not to Queue: Equilibrium Behavior in Queueing Systems*. Kluwer Academic Publishers, Norwell, MA.

Jerath, K., S. Netessine, S. K. Veeraraghavan. 2007. Revenue management with strategic customers: Last-minute selling and opaque selling. Working paper, University of Pennsylvania, Philadelphia.

Kim, B.-D., M. Shi, K. Srinivasan. 2001. Reward programs and tacit collusion. *Marketing Sci.* **20**(2) 99–120.

Lai, G., L. G. Debo, K. Sycara. 2009. Buy now and match later: Impact of posterior price matching on profit with strategic consumers. *Manufacturing Service Oper. Management*, ePub ahead of print February 10, http://msom.journal.informs.org/cgi/content/abstract/msom.1080.0248v1.

Liu, Q., G. van Ryzin. 2008. Strategic capacity rationing to induce early purchases. *Management Sci.* **54**(6) 1115–1131.

Liu, Q., G. van Ryzin. 2007. Strategic capacity rationing when customers learn. Working paper, Columbia University, New York.

Mahajan, S., G. van Ryzin. 2001. Stocking retail assortments under dynamic consumer substitution. *Oper. Res.* **49**(3) 334–351.

Milgrom, P., J. Roberts. 1986. Price and advertising signals of product quality. *J. Political Econom.* **94**(4) 796–821.

Muth, J. F. 1961. Rational expectations and the theory of price movements. *Econometrica* **29** 315–335.

Shen, Z.-J., X. Su. 2007. Customer behavior modeling in revenue management and auctions: A review and new research opportunities. *Production Oper. Management* **16**(6) 713–728.

Smith, S. A., N. Agrawal. 2000. Management of multi-item retail inventory systems with demand substitution. *Oper. Res.* **48**(1) 50–64.

Stokey, N. L. 1979. Intertemporal price discrimination. *Quart. J. Econom.* **93**(3) 355–371.

Stokey, N. L. 1981. Rational expectations and durable goods pricing. *Bell J. Econom.* **12**(1) 112–128.

Su, X. 2007. Intertemporal pricing with strategic customer behavior. *Management Sci.* **53**(5) 726–741.

Su, X. 2008. Bounded rationality in newsvendor models. *Manufacturing Service Oper. Management* **10**(4) 566–589.

Su, X., F. Zhang. 2008. Strategic customer behavior, commitment, and supply chain performance. *Management Sci.* **54**(10) 1759–1773.

Sun, D., J. K. Ryan, H. Shin. 2008. Why do we observe stockless operations on the Internet? Stockless operations under competition. *Production Oper. Management* **17**(2) 139–149.

van Ryzin, G., S. Mahajan. 1999. On the relationship between inventory costs and variety benefits in retail assortments. *Management Sci.* **48**(1) 50–64.

Verhoef, P. C., L. M. Sloot. 2005. Out-of-stock: Reactions, antecedents, management solutions, and a future perspective. M. Krafft, M. K. Mantrala, eds. *Retailing in the 21st Century: Current and Future Trends*. Springer, Berlin, 239–254.

Video Software Dealers Association. 1999. Annual report on the home video market. VSDA, Encino, CA.

Yin, R., Y. Aviv, A. Pazgal, C. S. Tang. 2007. The implications of customer purchasing behavior and in-store display formats. Working paper, University of California, Los Angeles.