M&SOM MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 13, No. 1, Winter 2011, pp. 89–107
ISSN 1523-4614 | EISSN 1526-5498 | 11 | 1301 | 0089

INFORMS
DOI 10.1287/msom.1100.0306

Strategic Capacity Rationing when Customers Learn

Qian Liu
Industrial Engineering and Logistics Management Department, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong, qianliu@ust.hk

Garrett van Ryzin
Graduate School of Business, Columbia University, New York, New York 10027, gjv1@columbia.edu

Consider a firm that sells products over repeated seasons, each of which includes a full-price period and a markdown period. The firm may deliberately understock products in the markdown period to induce high-value customers to purchase early at full price. Customers cannot perfectly anticipate availability. Instead, they use observed past capacities to form capacity expectations according to a heuristic smoothing rule. Based on their expectations of capacity, customers decide to buy either in the full-price period or in the markdown period. We embed this customer learning process in a dynamic program of the firm's capacity choices over time. One main result demonstrates the existence of a monotone optimal path of customers' expectations, which converges to either a rationing equilibrium or a low-price-only equilibrium. Further, there exists a critical value of capacity expectation such that the market converges to a rationing equilibrium if customers' initial expectations are less than that critical value; otherwise, a low-price-only equilibrium is the limiting outcome. These results show how firms can be stuck with unprofitable selling strategies from incumbent customer expectations. We also examine numerically how this critical value is affected by the firm's discount factor and customers' learning speed and risk aversion. Last, we show that the equilibrium under adaptive learning converges to that under rational expectations as the firm's discount factor approaches one.

Key words: consumer behavior; pricing and revenue management; dynamic programming
History: Received: September 4, 2007; accepted: May 1, 2010. Published online in *Articles in Advance* October 13, 2010.

1. Introduction

Faced with dynamic pricing, customers have an incentive to strategize the timing of their purchases and attempt to buy only when prices are low. One way for a firm to thwart this behavior is to deliberately create shortages to induce customers to buy early at higher prices. An earlier paper (Liu and van Ryzin 2008) addresses how a firm should optimally balance the benefit of capacity rationing against the opportunity cost of lost sales. It shows that a large high-value customer segment, large price differences over time, and risk aversion among customers all tend to make rationing an optimal strategy.

A key assumption made in Liu and van Ryzin (2008) is that customers have rational expectations[1] about the firm's capacity choice; that is, customers perfectly anticipate availability. The concept that customers rationally predict future product availability has been adopted in several recent papers on strategic customer behavior, for example, Aviv and Pazgal (2008), Cachon and Swinney (2009), Elmaghraby et al. (2008), Liu and van Ryzin (2008), Su (2007, 2008), and Jerath et al. (2007). Yet in reality, customers may only learn about a firm's strategy through repeat experiences. This is the situation this paper explores; we drop the rational expectations assumption and assume instead that customers adaptively learn over time. To give a concrete example of this phenomenon, we have heard department store managers lament the fact that they have, in effect, "trained customers to buy on sale" as a result of a longstanding practice of frequent promotional sales and end-of-season markdowns. Stores would like to regain the credibility of selling at full price but worry that without promotions or markdowns, customers may refuse to buy at all. In such a situation, should they attempt to change customers' expectations by restricting markdowns or the availability of goods? Or is it simply too difficult and costly to do so? What is the optimal response? These are the main questions we address in this paper.

The fundamental question of how customers form expectations of the future is addressed by adaptation-level theory; see Helson (1964), Sterman (1987), and Rubinstein (1998) for an overview. Experimental and empirical evidence support the hypothesis

[1] Note we use the term *rational expectations* to refer to the case where customers' expectations about the firm's capacity are rational; that is, based on all available information and correctly anticipating the firm's optimal stocking decision. Although it is related, we do not use the term in the sense it is used in macroeconomics.

that customers form expectations using an exponentially smoothed average of new and old information; see, for example, Akerman (1957), Nerlove (1958), Monroe (1973), Sterman (1989), Jacobson and Obermiller (1990), Stidham (1992), Greenleaf (1995), Kopalle et al. (1996), Rump and Stidham (1998), and Popescu and Wu (2007). Adaptive expectations imply that customers do not believe the permanence of recently revealed information; rather, they adjust to this new information gradually and merge it with past experience. To give an instance, Akerman (1957) and Nerlove (1958) look at how farmers react to price changes for their crops. When the current price increases, farmers do not believe its permanence until such an increased price remains in effect for a considerable period of time. Instead, farmers "discount" recent price changes and adapt their "memory" of prices accordingly. According to this exponential smoothing process, the effect of previous information decreases over time as weights decay exponentially; more recent information has a greater impact on expectations than the less recent information.

Other forms of dynamic expectation updates have also been employed to model customer demand. For example, Gans (2002) studies a problem of customer choice among a set of suppliers with uncertain service in quality level. He models customers' response to uncertain quality as a Bayesian multiarmed updating problem.

We use an exponential smoothing model to update customers' capacity expectations; that is, customers estimate the current capacity as a weighted average of the firm's capacity in the previous season and their prior estimate. In each season, based on their capacity expectations, customers decide whether to purchase at full price or to wait for a discount. The firm's decision is to make capacity choices for each season and profitably influence customers' expectations of capacity and thus their purchase behavior. This is modeled as a dynamic program, which links the firm's capacity decision and customers' capacity updates.

We establish several results for our model. First, we prove that the firm's capacity choices and customers' expectations converge to either a rationing equilibrium or a low-price-only equilibrium. We further show there exists a critical value of capacity expectation such that a rationing equilibrium is reached if customers' initial expectations are less than this critical value; otherwise, the firm eventually converges to serve the entire market at a low price without rationing. This helps explain how firms can be saddled with an unprofitable strategy because of their history. For example, in the department store scenario, mentioned above, customers may be used to finding products available whenever there is a sale, and it might take many seasons of lost sales to persuade customers that products are not as plentiful as before. The cost of lost sales incurred while customers learned this fact might dwarf the long-run benefit of finally getting them to change their behavior and buy at full price.

We also examine how the equilibrium obtained when customers adaptively learn about capacity expectations is related to the equilibrium when customers have rational expectations. Compared with the rational expectations case, the rationing equilibrium capacity under adaptive learning is always larger, as long as the discount factor is strictly less than one, whereas the long-run average profit at the rationing equilibrium under adaptive learning is lower than that under rational expectations. However, when the discount factor approaches one, the two equilibria are the same. This implies that the equilibrium under adaptive expectations learning is a good approximation of the rational expectations equilibrium when the discount factor is high enough.

Some recent literature also looks at dynamic expectation updating. Ovchinnikov and Milner (2009) study how a firm manages last-minute deals when customers learn and strategically respond to revenue management decisions. They consider different learning behaviors, including smoothing and self-regulating learning. However, they do not model each individual customer's waiting behavior as based on their payoff; rather, they assume that a random fraction of customers wait to buy. Gallego et al. (2008) investigate optimal pricing policies for a fixed stock of perishable goods when customers' learning behavior is taken into account. They analyze the influence of this learning behavior on a firm's equilibrium sales quantity based on a fluid model. Kleywegt (2006) considers a simple learning model, in which customers do not have full information about a firm's pricing policy or they learn about their valuations over time. Gaur and Park (2007) analyze a model of asymmetric consumer learning in which each customer updates his belief about product availability based on his own purchase experience and responds asymmetrically to positive and negative outcomes. However, the firm's fill rate is set at the beginning of the game and fixed throughout the time horizon; our model assumes the firm adjusts its capacity over time.

The remainder of the paper is organized as follows. In §2, we model the firm's capacity decisions when customers form expectations adaptively as a dynamic program. In §3, we show that the firm's optimal capacity converges to either a rationing equilibrium or a low-price-only equilibrium. Moreover, a critical value of capacity expectation determines which equilibrium is reached. In §4, we compare the adaptive learning equilibrium to the rational expectations outcome of Liu and van Ryzin (2008). In §5, we conduct numerical examples to study comparative statics for the critical value of capacity expectation. In §6, we study an alternative model based on fill-rate updates and discuss several extensions. Section 7 concludes the paper.

2. Adaptive Learning Model

2.1. Notation and Assumptions

Consider a firm that sells products over repeated sales seasons to a fixed population of customers. Each season (also called a *stage*) consists of two selling periods—a full-price (high-price) period and a markdown (low-price) period. At the beginning of each season, the firm makes its capacity decision. Resupply is not possible within a season. There are N customers in the market in each season. Customers are present when sales start and remain in the market until they either buy or the sales season ends. The market size N is deterministic and common knowledge to both the firm and the customers.[2] We assume N is large, so demand can be considered a continuous quantity, customers are nonatomic, and hence strategic interactions among customers can be safely ignored. Customers have unit demand and heterogeneous valuations for the firm's product that are uniformly distributed over $[0, U]$. The valuation distribution is common knowledge to both the firm and customers. We also assume customers have identical power utility functions $u(\cdot) = (\cdot)^\gamma$, where $0 < \gamma < 1$, implying they are risk averse.[3] The assumptions of uniform valuations and power utility function simplify the analysis. Nevertheless, we find numerically that the main results derived under these specialized assumptions hold for more general distributions of valuation and utility functions. Notation and parametric assumptions are summarized in Table 1.

Table 1 **List of Symbols**

Symbols	Definitions and parametric assumptions
p_H	Unit price in the full-price period
p_L	Unit price in the markdown period
c	Unit procurement cost; $c < p_L < p_H$
N	Population size
U	Upper bound of uniform valuations
δ	Firm's discount factor; $0 < \delta < 1$
θ	Customers' learning speed; $0 \le \theta \le 1$
γ	Customers' risk aversion; $0 < \gamma < 1$
t	Season (stage) index
$\underline{C}$	Lower bound of capacity; $\underline{C} = \frac{N}{U}(U - p_H)$
$\bar{C}$	Upper bound of capacity; $\bar{C} = \frac{N}{U}(U - p_L)$
C_t	The firm's capacity choice in season t
$\widehat{C}_t$	Customers' expectation of capacity in season t
C_s	Segmentation threshold capacity estimate; $C_s = \bar{C}\left(\frac{U - p_H}{U - p_L}\right)^\gamma$
$\widehat{C}_c$	Critical value of capacity expectation
C^0	Rationing equilibrium capacity under adaptive learning
C^*	Equilibrium capacity under adaptive learning; either $C^* = C^0$ or $C^* = \bar{C}$
C_R^0	Rationing equilibrium capacity under rational expectations
q_t	The firm's actual fill rate in the markdown period of season t
$\hat{q}_t$	Customers' expectation of fill rate in the markdown period of season t
$\hat{q}_c$	Critical value of fill rate expectation
q^0	Rationing equilibrium fill rate under adaptive learning
v_t	Threshold (cutoff) valuation in season t

2.2. The Model of Customer Behavior

During each season t, customers decide when to buy. They either buy early at full price and obtain one unit of product for sure, or they wait for a markdown, in which case availability is not guaranteed. Specifically, customers assess the fill rate (i.e., the probability of getting a unit), $\hat{q}_t$, and then weigh the payoff of purchasing immediately at the full price versus the expected payoff of waiting for a markdown. As shown in Liu and van Ryzin (2008), if customer utility functions are concave and strictly increasing, then for each fill rate there exists a unique cutoff value such that only customers with valuations greater than that value purchase at full price, and customers with valuations less than that value wait for a markdown. Customers with valuations equal to the cutoff value are indifferent. This indifference point (cutoff valuation), v_t, is defined (implicitly) by

$$(v_t - p_H)^\gamma = \hat{q}_t (v_t - p_L)^\gamma.[4] \tag{1}$$

Note that each customer's decision is based on his payoffs, and we do not consider waiting costs. Su (2007) models each customer's purchase decision by similar payoff functions, but waiting costs are included in the payoffs. In contrast, Ovchinnikov and Milner (2009) and Anderson and Wilson (2003) do not

[2] Deterministic aggregate demand is an assumption made in other literature on strategic customer behavior, such as Besanko and Winston (1990), Elmaghraby et al. (2008), and Su (2007).

[3] Risk aversion is required. As shown in Liu and van Ryzin (2008), when $\gamma = 1$ (i.e., risk-neutral customers), the rationing equilibrium becomes a limiting case; that is, only the full-price period incurs sales.

[4] We do not consider the case in which customers discount utility over time. As shown in Liu and van Ryzin (2008), when both discounting and rationing risk are taken into account, the results are qualitatively similar; discounting utility only adds to the incentive to buy early at a high price. But analytically characterizing optimal strategies becomes complicated. To keep the model simple and highlight the role of rationing, we don't consider the case where customers discount utility over time. We conjecture, however, that—as in the rational equilibrium case—our results qualitatively remain the same.

consider the buy-or-wait decision of individual customers; they model the aggregate waiting behavior and assume that a fraction of customers buy at full price and a fraction wait for a markdown. In particular, Ovchinnikov and Milner (2009) assume that a random number of customers, whose distribution is parameterized by a waiting parameter, will postpone purchases; Anderson and Wilson (2003) assume all the low-type customers will wait once the probability of obtaining discounted products is high.

In (1), the fill rate is the ratio of residual capacity to residual demand in the markdown period given by

$$\hat{q}_t = \frac{\widehat{C}_t - (N/U)(U - v_t)}{(N/U)(v_t - p_L)}, \tag{2}$$

where we use the notation "$\widehat{C}$" to distinguish customers' expectation of capacity from the firm's actual capacity decision.

By combining (1) and (2), the cutoff value and customers' expectation of capacity are related by

$$\widehat{C}_t = \frac{N}{U}\left(U - v_t + (v_t - p_L)\left(\frac{v_t - p_H}{v_t - p_L}\right)^{\gamma}\right), \quad p_H \le v_t \le U. \tag{3}$$

According to (3), we have

$$\underline{C} \le \widehat{C}_t \le C_s = \bar{C}\left(\frac{U - p_H}{U - p_L}\right)^{\gamma},$$

where we call C_s the *segmentation threshold capacity estimate*. If customers' expectation of capacity is less than C_s, the market is segmented and a fraction of customers buys at the full price; otherwise, all customers opt to buy in the markdown period.

We next show that the cutoff valuation, v_t, is uniquely determined by capacity expectation, $\widehat{C}_t$. This follows directly from Lemma 1 below, the proof of which is provided in the appendix.

Lemma 1. *$\widehat{C}_t(v_t)$ defined by* (3) *is strictly increasing and concave in $v_t \in [p_H, U]$.*

Lemma 1 implies that the inverse function of $\widehat{C}_t(v_t)$, denoted $g(\widehat{C}_t)$, exists; further, $g(\widehat{C}_t)$ is strictly increasing and convex in $\widehat{C}_t \in [\underline{C}, C_s]$. This is because the inverse function of a strictly increasing and concave function is strictly increasing and convex.[5] Therefore, the cutoff valuation v_t is uniquely characterized by $g(\widehat{C}_t)$ when $\underline{C} \le \widehat{C}_t \le C_s$. Once $\widehat{C}_t$ exceeds C_s, customer purchase behavior changes fundamentally; all customers wait to buy in the markdown period and segmentation of the market is no longer attainable. We define v_t equal to U when $C_s \le \widehat{C}_t \le \bar{C}$. Then, for any given capacity expectation $\widehat{C}_t \in [\underline{C}, \bar{C}]$, the cutoff valuation v_t is uniquely determined by

$$v(\widehat{C}_t) = \begin{cases} g(\widehat{C}_t) & \text{if } \underline{C} \le \widehat{C}_t \le C_s, \\ U & \text{if } C_s \le \widehat{C}_t \le \bar{C}. \end{cases}$$

To model the customers' learning process, we use an exponential smoothing function. We assume that this learning process takes place using capacity rather than fill rates and that all customers have the same expectation of capacity; that is, $\widehat{C}_t$ is a consensus and commonly held estimate among customers. Under these assumptions, the customers' expectation of capacity for season $t+1$, $\widehat{C}_{t+1}$, is a weighted average of the firm's capacity in season t, C_t, and customers' prior estimate, $\widehat{C}_t$

$$\widehat{C}_{t+1} = \theta C_t + (1-\theta)\widehat{C}_t, \tag{4}$$

where θ is the customers' learning speed. As θ increases, customers place more weight on new information and adjust their expectation more rapidly. An extreme case of $\theta = 1$ implies customers take the firm's capacity in the previous season as their estimate of the current capacity. The other extreme case of $\theta = 0$ implies no learning at all; customers simply stick to their initial beliefs.[6] We exclude this later case to avoid trivialities.

One may naturally question the plausibility of the assumption that customers observe and estimate capacity rather than fill rates. Customers might not be able to observe capacity directly, or even if they can, the firm might manipulate their perception of capacity by limiting the "on the rack" display. Yet the update of capacity in our model (4) is made ex post after season t is completed and hence would be influenced by a variety of information other than direct observation, including word of mouth, advertisements, websites, reports and surveys, etc. These sources of information may compensate for the inability to observe capacity directly. Another major motivation for using a capacity representation of the learning process is that it is analytically more tractable than the fill-rate update model. Moreover, we analyze a model of fill-rate updating in detail in §6.1, and many of the structural results are the same for both models.

Another desirable generalization of the learning process might be to allow customers to have heterogeneous estimates based on their own purchase histories

[5] The brief reason is the following: suppose $h(\cdot)$ is strictly increasing and concave, and $g(\cdot)$ is the inverse function of $h(\cdot)$. Then $g'(x) = 1/(h'(g(x))) > 0$, and $g''(x) = -h''((g(x))g'(x))/((h'(g(x)))^2) \ge 0$ because $g'(x) > 0$ and $h''(g(x)) \le 0$.

[6] We can also interpret $1-\theta$ as a memory parameter; thus, for $\theta = 1$, customers have no memory; for $\theta = 0$, customers have full memory and never adapt to observed information.

and stock-out experiences, rather than assuming that all customers share the same expectation of availability. Although such a model would be interesting to study, it would be much less tractable; we would have to track each customer's expectation of availability, and doing so would require a state vector with large dimension, leading to a significantly more complex model. Gaur and Park (2007) model an asynchronous learning process in which each customer updates his fill rate expectation based on purchase experience. However, they consider a homogeneous population of customers who have identical valuations; thus, customers' purchase decisions are determined by their perceived product availability only. Another difference is that they assume the firm's fill rate is fixed throughout the time horizon once it is selected at the beginning of sales. These assumptions simplify the analysis.

More importantly, assuming customers only form expectations based on their individual purchase experience is equally unrealistic. It would suggest, for example, that we could only know whether a car maker's automobiles are overstocked if we attempt to buy one—end-of-season liquidation advertisements blaring on our television sets notwithstanding. Or it could suggest that when we book a flight early, we do not learn whether it is sold out afterward—despite the fact that we observe how full the plane is once we travel. In this sense, the assumption of individual updating based only on purchase outcomes is the opposite extreme of our assumption of common knowledge among customers. Reality, of course, lies somewhere in between.

Distinct from assumptions about private or common knowledge of availability, it may also be desirable to allow customers to have different learning speeds. Again like updating based on private purchase experience, this type of heterogeneity would increase the complexity of the state space, though allowing for a small number of types (e.g., "slow" and "fast" learners) might not be overly difficult. Still, assuming a common learning speed captures the essential behavior we want to study and is simple and tractable.

2.3. The Model of the Firm

The capacity learning process (4) links the firm's capacity choices from one season to the next. Consequently, the firm's decision problem is naturally modeled as a dynamic program. Let $V(\widehat{C}_t)$ denote the maximum discounted profit, given that customers' capacity expectation is $\widehat{C}_t$. Future value is discounted by a discount factor of δ per season. $V(\widehat{C}_t)$ satisfies the following Bellman equation:

$$V(\widehat{C}_t) = \max_{C_t \in [\underline{C}, \bar{C}]} \{\Pi(\widehat{C}_t, C_t) + \delta V(\theta C_t + (1-\theta)\widehat{C}_t)\}, \quad (5)$$

where $\Pi(\widehat{C}_t, C_t)$ is the one-stage profit, given that the customers' capacity expectation is $\widehat{C}_t$ and the firm's capacity is C_t:

$$\Pi(\widehat{C}_t, C_t) = \frac{N}{U}(p_H - p_L)(U - v(\widehat{C}_t)) + (p_L - c)C_t. \quad (6)$$

The constraint $\underline{C} \leq C_t \leq \bar{C}$ ensures that there is no shortage during the full-price period and no overage during the markdown period.[7]

We note, however, that in some cases it may be profitable to deliberately understock products in the full-price period in order to change customer expectations as quickly as possible. For example, Greenleaf (1995) finds that offering "reverse" promotions—in which extremely high prices are charged in some periods—may be optimal in a recurring promotion model in order to raise future reference prices. Similarly, Byers and Huff (2005) find that it may be optimal to leave demand unfilled at higher prices on purpose to boost future demand. These possibilities are interesting, but our control constraints rule out such extreme actions.

To understand the basic tradeoff the firm faces, notice that customer decisions about when to buy depend only on their capacity expectation $\widehat{C}_t$, not on the firm's actual capacity choice C_t. Hence, if the firm were only interested in maximizing the current profit, it would choose the maximum capacity $\bar{C}$ and satisfy all demand during both the full-price and markdown periods. But the firm's capacity choice also influences customers' future expectations; a larger current capacity increases customers' expectation of capacity in future seasons, which encourages them to wait for markdowns. This will reduce future profits. So the key tradeoff is the short-term benefit of stocking amply to satisfy current demand versus the negative impact that such high stock levels have on customers' future expectations of capacity. The firm's goal is to seek a sequence of capacity choices over time that optimally balances these effects and maximizes its total discounted profit.

In (5), when the firm's optimal response to the state $\widehat{C}_t$ is to choose $C_t = \widehat{C}_t$, both the firm's optimal capacity decision and customers' expectation of capacity will be the same and will stay at that value thereafter. We call such a stable value an equilibrium of the

[7] Actually, a weaker constraint of $(N/U)(U - v_t(\widehat{C}_t)) \leq C_t \leq (N/U) \cdot (U - p_L)$ is sufficient to guarantee underage never occurs during the full-price period and overage never occurs during the markdown period. However, state-dependent constraints of this type significantly complicate the analysis of the model. For technical reasons, we assume the firm would stock at least the potential demand at full price, $\underline{C}$, and no more than the potential demand at markdown price, $\bar{C}$.

model (5). More precisely, let $C^*(\widehat{C})$ denote the optimal solution to (5), given state $\widehat{C}$. Then an equilibrium of (5) is a fixed point of the *state updating function*—

$$f(\widehat{C}) = \theta C^*(\widehat{C}) + (1-\theta)\widehat{C}.$$

That is, $\widehat{C}$ is an equilibrium if $f(\widehat{C}) = \widehat{C}$. At the equilibrium, $C^*(\widehat{C}) = \widehat{C}$. We henceforth do not differentiate the firm's capacity and customers' expectation of capacity at the equilibrium, both of which we call the *equilibrium capacity*.

3. Analysis of the Optimal Capacity Decision

Because customers exhibit fundamentally different purchase behavior once their expectation of capacity exceeds the segmentation threshold capacity estimate C_s, a direct analysis of the model (5) is quite complex. However, note that it can be easily solved as long as the state (i.e., capacity expectation) does not cross C_s. Therefore, we use a "separate-and-paste" approach to analyze the problem. We divide the entire state space into two subspaces, over which two isolated subproblems, called *the region 1* and *region 2 problems*, are defined. In this section, we first show that both the region 1 and region 2 problems are well behaved and can be completely analyzed using classical optimization techniques. In §3.2, we show that there always exists a monotone optimal state path for the original problem. Because of this property, the problem eventually reduces to either the region 1 or region 2 problem. We can then "paste" the results of subproblems together to solve the original problem. We find that the firm's optimal capacity choices converge to either a rationing equilibrium or a low-price-only equilibrium, depending on customers' initial expectations of capacity. In §3.3, we show that there exists a critical value of capacity expectation that determines which of these equilibria is reached. When customers' capacity expectation is less than that critical value, the rationing equilibrium is obtained; otherwise, the low-price-only equilibrium is the optimal long-run outcome.

3.1. Analysis of the Region 1 and Region 2 Problems

As discussed before, all customers wait to buy in the markdown period if their capacity expectation is greater than the segmentation threshold capacity estimate C_s. Then the firm's best response is to stock $\bar{C}$ to meet all potential demand at the markdown price. Because the switch from a segmented market to a nonsegmented market induces a jump of capacity, the value function $V(\widehat{C}_t)$, unfortunately, is not well behaved (i.e., $V(\widehat{C}_t)$ may be not differentiable at C_s). Nevertheless, the threshold function $v(\widehat{C}_t)$ has special structure: it is convex in the range of $[\underline{C}, C_s]$ and constant at $\bar{C}$ afterward. This motivates us to divide the entire state space into two subspaces—$[\underline{C}, C_s]$ and $[C_s, \bar{C}]$—over which two subproblems are defined. Specifically, the region 1 problem is defined as follows:

$$V_1(\widehat{C}_t) = \max_{C_t \in S_1(\widehat{C}_t)} \left\{ \frac{N}{U}(p_H - p_L)(U - v(\widehat{C}_t)) + (p_L - c)C_t + \delta V_1(\theta C_t + (1-\theta)\widehat{C}_t) \right\}, \quad (7)$$

where

$$S_1(\widehat{C}) = \left\{ C \mid \underline{C} \le C \le \bar{C} \text{ and } C \le -\frac{1-\theta}{\theta}\widehat{C} + \frac{C_s}{\theta} \right\}.$$

The region 2 problem is then

$$V_2(\widehat{C}_t) = \max_{C_t \in S_2(\widehat{C}_t)} \left\{ (p_L - c)C_t + \delta V_2(\theta C_t + (1-\theta)\widehat{C}_t) \right\}, \quad (8)$$

where

$$S_2(\widehat{C}) = \left\{ C \mid \underline{C} \le C \le \bar{C} \text{ and } C \ge -\frac{1-\theta}{\theta}\widehat{C} + \frac{C_s}{\theta} \right\}.$$

We require $C_t \in S_1(\widehat{C}_t)$ in the region 1 problem to ensure that any state falls within $[\underline{C}, C_s]$; and $C_t \in S_2(\widehat{C}_t)$ guarantees that the state space of the region 2 problem is $[C_s, \bar{C}]$. There are no overlapping states between these two subproblems except state C_s.

We first examine the region 2 problem, which is quite simple to analyze. Because all customers delay purchases in this region, the firm stocks $\bar{C}$ to satisfy all demand at the markdown price. The value function $V_2(\widehat{C}_t)$ is then state independent, and the equilibrium is trivially $\bar{C}$. However, the equilibrium $\bar{C}$ cannot be reached in finite time unless customers' capacity expectation is $\bar{C}$. This is because customers update expectations via $\widehat{C}_{t+1} = \theta C_t + (1-\theta)\widehat{C}_t$, and both C_t and $\widehat{C}_t$ are less than $\bar{C}$, though $\widehat{C}_t \to \bar{C}$ as $t \to \infty$.

We then focus on the analysis of the region 1 problem. Using an argument similar to the one for the original problem (5), there exists a tradeoff between stocking more to increase current profits and stocking less to reduce customer expectations and induce more early purchases in future. In fact, we show that a rationing unique equilibrium capacity exists and characterize it in Proposition 1. The uniqueness of this equilibrium eliminates the optimality of manipulating customers' expectations whenever rationing is achievable; that is, it is not profitable to cycle between providing high and low availability.

To show the existence of equilibrium and further characterize it, we need the concavity of the value function $V_1(\widehat{C}_t)$, which is established in the following lemma.

Lemma 2. *The value function $V_1(\widehat{C}_t)$ defined by* (7) *is strictly concave and decreasing in $\widehat{C}_t$.*

The proof of Lemma 2 is in the appendix. By the concavity of the value function $V_1(\widehat{C}_t)$, it is easy to check that the inner maximization term of (7) is concave in C_t as well. The optimal capacity is then derived from the first-order conditions. Following the same approach as in Lemma 2, we can show that the value function $V(\widehat{C}_t)$ defined by (5) is decreasing in $\widehat{C}_t$. This monotonicity is quite intuitive; a lower expectation of capacity induces more customers to purchase early at full price, thus leading to a higher current profit. In addition, a lower current expectation of capacity leads to lower future expectation of capacity, so future profits benefit as well. However, the value function for the whole problem $V(\widehat{C}_t)$ is not concave; in fact, it is not even differentiable at some points. This is discussed further in §3.2.

Proposition 1 establishes the existence of a unique rationing equilibrium. Although it is hard to obtain the equilibrium in a closed form, we characterize it as a solution to an implicit function, allowing us to examine comparative statics of the equilibrium. This is established in Corollary 1 below. Both proofs are given in the appendix.

Proposition 1. *For the region* 1 *problem, there exists a unique equilibrium capacity that is equal to* $\min\{C^0, C_s\}$, *where C_s is the segmentation threshold capacity and C^0 is called the rationing equilibrium capacity determined by*

$$C^0 = \frac{N}{U}\left(U - v^0 + (v^0 - p_L)\left(\frac{v^0 - p_H}{v^0 - p_L}\right)^{\gamma}\right). \qquad (9)$$

In (9), *v^0 is the solution to*

$$\left(\frac{v - p_H}{v - p_L}\right)^{\gamma}\left(1 + \frac{\gamma(p_H - p_L)}{v - p_H}\right) - \frac{(p_L - c)(1/\delta - 1) + \theta(p_H - c)}{(p_L - c)(1/\delta - 1 + \theta)} = 0. \qquad (10)$$

Note that C^0 defined in (9) is strictly increasing in v^0 and C^0 is equal to C_s when $v^0 = U$. Therefore, the equilibrium for the region 1 problem is the rationing equilibrium capacity C^0 when $v^0 < U$; otherwise, it is C_s.

Corollary 1. *The rationing equilibrium capacity C^0 defined in* (9) *decreases in the firm's discount factor δ and customers' learning speed θ. And the per stage profit at C^0, $\Pi(C^0, C^0)$, increases in δ and θ.*

The intuition for the monotonicity in discount factor is straightforward; the larger the firm's discount factor, the more important are future profits. Hence, the long-run benefit of stocking less and creating future expectations of shortages dominates the short-term profit from stocking and selling more in the current season. As a result, a lower equilibrium capacity is expected as the firm's discount factor increases. The learning speed effect is less obvious but also intuitive. When customers adapt quickly to changes in capacity, the future benefit of creating expectations of shortages is realized quickly; that is, if we reduce capacity now, customers quickly adjust and we reap the benefit of their revised expectations (i.e., their desire to buy early at high prices) in the next few seasons. If, in contrast, they are slow to react, any current shortages we create take many seasons to work their way into revised expectations that will, ultimately, induce them to buy early at high prices. Hence, a high learning speed makes current shortages more beneficial and leads to a lower equilibrium capacity. The speed of learning also plays important roles in the papers by Ovchinnikov and Milner (2009) and Gallego et al. (2008). In both papers, they show that the firm's optimal policy also converges when customers learn slowly.

3.2. Convergence of the Optimal Capacity Policy

In this section, we "paste" the results for each region problem and establish monotone convergence to equilibria for the original problem. Proposition 2 states the relationship of the equilibria for the original problem to equilibria for each region problem; namely, the equilibria for the original problem must be equilibria for the region 1 or region 2 problems. The key reason that we are able to "paste" results of two subproblems is the monotonicity of optimal state paths, which is established in Proposition 3.

Proposition 2. *If there exists an equilibrium for the problem* (5), *then it is either C^0 or $\bar{C}$.*

The proof of Proposition 2 is given in the appendix, but the intuition is straightforward; an equilibrium capacity for the original problem (5), if it exits, must be an equilibrium for either the region 1 or region 2 problems as well, but C^0 or $\bar{C}$ is the only possible equilibrium for these subproblems.

Proposition 3 plays an important role in the separate-and-paste procedure. According to this proposition, there exists a monotone optimal state path. Therefore, the state of problem (5) will eventually remain in either region 1 or region 2. Hence, the equilibrium results for each region problem carry over to the original problem. The proof of Proposition 3 uses a sample-path analysis approach, which is provided in the appendix.

Proposition 3. *For the problem* (5), *there exists an optimal monotone state path; namely, given any initial state $\widehat{C}_1$, there exists an optimal state path, denoted $\{\widehat{C}_t\}_{t\geq 1}$, such that either $\widehat{C}_1 \geq \cdots \geq \widehat{C}_t \geq \cdots$ or $\widehat{C}_1 \leq \cdots \leq \widehat{C}_t \leq \cdots$.*

According to the monotone convergence theorem, any bounded monotone sequence converges. Because the state space in our model is compact, a direct consequence of Proposition 3 is that an optimal monotone state path must converge to some limiting state. We can further show that this limiting state is an equilibrium of the original problem, which must also be an equilibrium for the region 1 and region 2 problems. Theorem 1 below summarizes this result.

Theorem 1. *For the problem* (5), *given any state* $\widehat{C}_1 \in [\underline{C}, \bar{C}]$, *there exists a monotone optimal state path* $\{\widehat{C}_t\}_{t\geq 1}$, *which converges to either a rationing equilibrium capacity* C^0 *or a low-price-only equilibrium capacity* $\bar{C}$.

Proof. Given any state $\widehat{C}_1$, there exists a monotone optimal state path $\{\widehat{C}_t\}_{t\geq 1}$ by Proposition 3. Because $\{\widehat{C}_t\}_{t\geq 1}$ is bounded within $[\underline{C}, \bar{C}]$, the state sequence $\{\widehat{C}_t\}_{t\geq 1}$ converges to some state, denoted $\widehat{C}^*$.

Because of monotonicity of the state path $\{\widehat{C}_t\}_{t=1,2,\ldots}$, there exists a stage t' such that either $\underline{C} \leq \widehat{C}_t \leq C_s$, $\forall\, t \geq t'$ or $C_s < \widehat{C}_t \leq \bar{C}$, $\forall\, t \geq t'$.

Case 1. $\underline{C} \leq \widehat{C}_t \leq C_s$, $\forall\, t \geq t'$

In this case, the problem (5) eventually remains in region 1 after stage t'. Note that the region 1 problem and the original problem have the same dynamic program formulation, and any feasible policy for the region 1 problem is feasible for the original problem as well. Therefore, if the optimal policy for the original problem is feasible for the region 1 problem, it must be an optimal policy for the region 1 problem. Hence, $\{\widehat{C}_t\}_{t\geq t'}$ is an optimal state path for the region 1 problem and converges to state $\widehat{C}^*$, $\underline{C} \leq \widehat{C}^* \leq C_s$. That is, $\lim_{k\to\infty} \widehat{C}_{t'+k} = \widehat{C}^*$. Recall in the region 1 problem the state updating function, $f(\widehat{C}) = \theta C^*(\widehat{C}) + (1-\theta)\widehat{C}$, is continuous in $\widehat{C}$ (see the proof of Proposition 1). Then

$$f(\widehat{C}^*) = f\Big(\lim_{k\to\infty} \widehat{C}_{t'+k}\Big) = \lim_{k\to\infty} f(\widehat{C}_{t'+k}) = \lim_{k\to\infty} \widehat{C}_{t'+k+1} = \widehat{C}^*.$$

Because $\widehat{C}^*$ is an equilibrium of the region 1 problem, either $\widehat{C}^* = C^0$ or $\widehat{C}^* = C_s$ by Proposition 1. However, the value $V(\widehat{C}_t)$ strictly increases if the capacity decision is $\bar{C}$ instead of C_s for the original problem, which contradicts the optimality of the state path $\{\widehat{C}_t\}_{t\geq t'}$. Hence, it must be that $\widehat{C}^* = C^0$.

Case 2. $C_s < \widehat{C}_t \leq \bar{C}$, $\forall\, t \geq t'$

The problem (5) remains in region 2 after stage t'. Along the same line of argument as in Case 1, $\{\widehat{C}_t\}_{t\geq t'}$ is also an optimal path for the region 2 problem. Because the optimal policy at any state for the region 2 problem is $\bar{C}$, the state updating function $f(\widehat{C}) = \theta\bar{C} + (1-\theta)\widehat{C}$ is obviously continuous in $\widehat{C}$. It follows that the limiting state is an equilibrium of the region 2 problem. □

Theorem 1 shows that repeated interactions between the firm and its customers lead to an equilibrium, which is either a rationing equilibrium C^0 or a low-price-only equilibrium without rationing $\bar{C}$. Moreover, the fact that there is monotone convergence to these equilibria implies the optimal decisions do not oscillate among high and low capacity choices. In other words, the firm does not profit from manipulating customers' expectations by alternating between providing high and low availability during markdown periods. This differs from some previous results. For example, Ovchinnikov and Milner (2009) and Gallego et al. (2008) find that the firm alternates between offering high and low availability when customers learn quickly. However, the reason for such a difference is not well understood: Ovchinnikov and Milner do not model an individual's buy-or-wait decision but instead consider aggregate waiting behavior, and Gallego et al. (2008) demonstrate this result numerically only. It is therefore difficult to discern the underlying reason for this difference in the optimal policy. Policies of this type of convergence or cycling are also explored in dynamic pricing models with reference price effects—for example, Greenleaf (1995) and Kopalle et al. (1996). They find that the optimal prices either cycle or converge to a constant level, depending primarily on the customers' asymmetric responses to perceived gains and losses.

It is worth noting that not every optimal state path is monotone; there may exist an optimal nonmonotone state path. However, Proposition 4 shows that such a nonmonotone state path, if it exists, is of quite specialized form: it can change direction only at state C^0 and must converge to a low-price-only equilibrium. The proof is given in the appendix.

Proposition 4. *If there exists an optimal nonmonotone state path for the problem* (5), *it can change its direction only at state* C^0; *furthermore, such a path must converge to state* $\bar{C}$.

3.3. Incumbent Expectations and Equilibrium Outcomes

Both C^0 (rationing) and $\bar{C}$ (no rationing) can be equilibria for the adaptive learning model (5). Under what conditions does there exist a unique equilibrium? One can show that if a low-price-only equilibrium is more profitable than a rationing equilibrium in the long run, the low-price-only solution is the unique equilibrium for all initial expectations. (See Lemma 4 for a formal statement and proof.) However, if rationing is more profitable in the long run, both equilibria can be reached, depending on customers' initial expectations of capacity. If customers' expectations are larger than a critical value, the long-run equilibrium is no rationing, but if their initial expectations are smaller than the critical value, the outcome converges to a rationing equilibrium. This result is established in the following theorem.

Theorem 2. *There exists a critical value of capacity expectation, denoted $\widehat{C}_c$, such that the optimal capacity converges to $\bar{C}$ if the customers' capacity expectation $\widehat{C}$ is greater than $\widehat{C}_c$; otherwise, it converges to C^0.*

Proof. Suppose there exists an optimal state path from $\widehat{C}_1$, denoted $\{\widehat{C}_t\}_{t\geq 1}$, which converges to $\bar{C}$. We now show that any state path starting from a state greater than $\widehat{C}_1$ will converge to $\bar{C}$ as well. The proof is given on a case-by-case basis in terms of the location of the state $\widehat{C}_1$.

Case 1. $\widehat{C}_1 \leq C^0$

$\forall \widetilde{C}_1 > \widehat{C}_1$, suppose no optimal state path exists from $\widetilde{C}_1$ which converges to $\bar{C}$; then any optimal state path from $\widetilde{C}_1$, denoted $\{\widetilde{C}_t\}_{t\geq 1}$, must converge to C^0. Along the optimal state path, $\{\widehat{C}_t\}_{t\geq 1}$, there must exist state $\widehat{C}_m$ and $\widehat{C}_{m+1}$, $m \geq 1$, such that $\widehat{C}_m \leq C^0 < \widehat{C}_{m+1}$. By the optimality of $V(\widehat{C}_m)$,

$$\widehat{\Pi}(\widehat{C}_m, \widehat{C}_{m+1}) + \delta V(\widehat{C}_{m+1}) \geq \widehat{\Pi}(\widehat{C}_m, C^0) + \delta V(C^0),$$

which implies

$$V(C^0) - V(\widehat{C}_{m+1}) \leq \frac{p_L - c}{\theta\delta}(\widehat{C}_{m+1} - C^0). \tag{11}$$

Because $\{\widetilde{C}_t\}_{t\geq 1}$ converges to C^0 only,

$$\widehat{\Pi}(C^0, C^0) + \delta V(C^0) > \widehat{\Pi}(C^0, \widehat{C}_{m+1}) + \delta V(\widehat{C}_{m+1}),$$

which implies

$$V(C^0) - V(\widehat{C}_{m+1}) > \frac{p_L - c}{\theta\delta}(\widehat{C}_{m+1} - C^0). \tag{12}$$

But (12) contradicts (11). Therefore, there exists an optimal path from $\widetilde{C}_1$ that converges to $\bar{C}$.

Case 2. $\widehat{C}_1 > C^0$

Define $C(\widehat{C}_1) = \theta\bar{C} + (1-\theta)\widehat{C}_1$. One can easily check $C(\widehat{C}_1) > \widehat{C}_1$ if $\widehat{C}_1 > \underline{C}$. Given any state $\widetilde{C}_1 \in [\widehat{C}_1, C(\widehat{C}_1)]$, suppose there exists no optimal state path that converges to $\bar{C}$. Then such an optimal state path, denoted $\{\widetilde{C}_t\}_{t\geq 1}$, must converge to C^0.

Figure 1 illustrates all the possible scenarios of locations of states $\widehat{C}_1$, $\widehat{C}_2$, $\widetilde{C}_1$, and $\widetilde{C}_2$. We claim it is feasible to reach state $\widetilde{C}_2$ from $\widehat{C}_1$ in one step and to reach state $\widehat{C}_2$ from $\widetilde{C}_1$ in one step. This is trivially true for cases (a), (b), and (c) in Figure 1. In the case of $\widetilde{C}_1 > \widetilde{C}_2 > \widehat{C}_2 > \widehat{C}_1$ (Figure 1(d)), we can check that $(\widetilde{C}_2 - (1-\theta)\widehat{C}_1)/\theta \leq (\widetilde{C}_1 - (1-\theta)\widehat{C}_1)/\theta \leq \bar{C}$ and $(\widetilde{C}_2 - (1-\theta)\widehat{C}_1)/\theta > (\widehat{C}_1 - (1-\theta)\widehat{C}_1)/\theta = \widehat{C}_1 \geq \underline{C}$. Hence, it is feasible for reach $\widetilde{C}_2$ from $\widehat{C}_1$ in one step. The same argument shows that $\widetilde{C}_1$ can reach $\widehat{C}_2$ in one step as well.

By the optimality of $V(\widehat{C}_1)$,

$$\widehat{\Pi}(\widehat{C}_1, \widehat{C}_2) + \delta V(\widehat{C}_2) \geq \widehat{\Pi}(\widehat{C}_1, \widetilde{C}_2) + \delta V(\widetilde{C}_2),$$

which implies

$$V(\widetilde{C}_2) - V(\widehat{C}_2) \leq \frac{p_L - c}{\delta\theta}(\widehat{C}_2 - \widetilde{C}_2). \tag{13}$$

As hypothesized, there does not exist an optimal state at state $\widetilde{C}_1$ that converges to $\bar{C}$; therefore,

$$\widehat{\Pi}(\widetilde{C}_1, \widetilde{C}_2) + \delta V(\widetilde{C}_2) > \widehat{\Pi}(\widetilde{C}_1, \widehat{C}_2) + \delta V(\widehat{C}_2),$$

which implies

$$V(\widetilde{C}_2) - V(\widehat{C}_2) > \frac{p_L - c}{\delta\theta}(\widehat{C}_2 - \widetilde{C}_2). \tag{14}$$

But (14) contradicts (13); hence, there exists an optimal state path at state $\widetilde{C}_1$ that converges to $\bar{C}$.

Thus far, we have shown that if there exists an optimal state path at $\widehat{C}_1$ that converges to $\bar{C}$, then for any state between $\widehat{C}_1$ and $C(\widehat{C}_1)$, there exists an optimal state path that also converges to $\bar{C}$. Then set $\widehat{C}_1 \leftarrow C(\widehat{C}_1)$; along the same line of argument, there exists an optimal state path converging to $\bar{C}$ for any state within $[\widehat{C}_1, C(\widehat{C}_1)]$. Because $C(\widehat{C}_1)$ is strictly larger than $\widehat{C}_1$, eventually, $C(\widehat{C}_1)$ converges to $\bar{C}$. Hence, starting with any state greater than $\widehat{C}_1$, there exists an optimal state path converging to $\bar{C}$ as well if there exists an optimal state path at $\widehat{C}_1$ with the convergent state $\bar{C}$.

Applying the similar arguments, we can show that if there exists an optimal state path from $\widehat{C}_1$ that converges to C^0, then any state path from a state less than $\widehat{C}_1$ will converge to C^0, too. We thus conclude that there exists a critical value of state, $\widehat{C}_c$, such that the optimal state path converges to $\bar{C}$ if it starts from a state greater than $\widehat{C}_c$, and converges to C^0 otherwise. □

Note the proof of Case 1 in Theorem 2 holds for any state $\widetilde{C}_1$ (i.e., $\widetilde{C}_1 \geq \widehat{C}_1$ is not necessary). This implies that if there exists an optimal state path from a state less than C^0 that converges to $\bar{C}$, then for any state, there exists an optimal state path that converges to $\bar{C}$ as well. Therefore, the critical value of capacity expectation that determines which equilibrium is attained is either $\underline{C}$ or greater than C^0. The equilibrium is stable as long as perturbation is small enough not to cross the critical value of capacity expectation.

Theorem 2 is perhaps the most important result in our analysis. It provides an explanation for why retail firms may follow policies of providing high availability even though such policies may not be the most profitable in the long run. For instance, traditional department stores find themselves trapped in a pattern of selling most of their stock during holiday and end-of-season sale periods, whereas newer retailers like Zara are able to convince customers to buy at full price by stocking less and phasing out products

Figure 1 **Feasible Actions to Comprise an Alternative State Path**

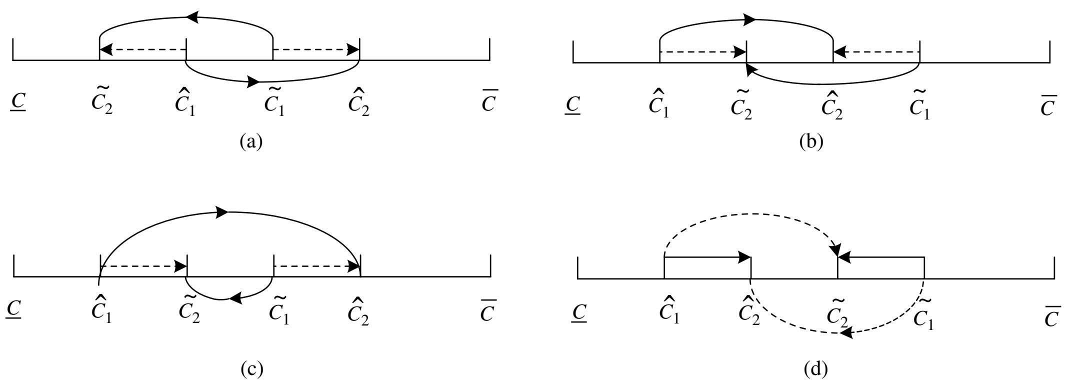

quickly. Similarly, General Motors was saddled with a pattern of dealer haggling, end-of-model-year discounting, and frequent promotions, yet its subsidiary upstart Saturn was able to enter the market with nonnegotiated prices and less-frequent discounting. Simply put, if customers expect high availability at discounted prices based on a long history of always finding products available on sale, and if these expectations change slowly, it can be quite costly to reach a new equilibrium in which customers anticipate difficulty finding products at less than full price. In practical terms, changing strategies might require limiting availability of discounts for a number of seasons, during which customers still refuse to buy at full price because they expect products to be available at a discount. It could take several years of experience to convince customers that the firm's availability policy has really changed, and the profit losses incurred during this process might simply be too severe to justify the change in strategy.

4. Comparison to Rational Expectation Equilibrium

Liu and van Ryzin (2008) study a similar capacity rationing problem but assume that customers can perfectly anticipate the firm's capacity (fill rate). That is, whatever capacity a firm chooses, customers immediately react and adjust their buying decisions accordingly. This corresponds to a case of rational expectations, and we call the resulting equilibrium the *rational expectation equilibrium*. In particular, Liu and van Ryzin (2008) show that the rational expectation equilibrium is either C_R^0 or $\bar{C}$, where C_R^0 is determined by

$$C_R^0 = \frac{N}{U}\left(U - v_R^0 + (v_R^0 - p_L)\left(\frac{v_R^0 - p_H}{v_R^0 - p_L}\right)^{\gamma}\right),$$

and v_R^0 is a solution to the equation

$$\left(\frac{v - p_H}{v - p_L}\right)^{\gamma}\left(1 + \frac{\gamma(p_H - p_L)}{v - p_H}\right) - \frac{p_H - c}{p_L - c} = 0. \qquad (15)$$

When the firms stocks $\bar{C}$, the outcome involves no rationing.

How is the adaptive learning equilibrium related to the rational expectation equilibrium? The answer is given in Theorem 3 below.

Theorem 3. *As the firm's discount factor $\delta \to 1$, the adaptive learning equilibrium converges to the rational expectation equilibrium.*

The proof of this theorem requires the following two lemmas, the proofs of which are provided in the appendix.

Lemma 3. *As the firm's discount factor δ approaches one, the rationing equilibrium C^0 under adaptive learning converges to the rationing equilibrium under rational expectations C_R^0.*

Lemma 4. *For the adaptive learning model* (5), *if the per stage profit at a rationing equilibrium, $\Pi(C^0, C^0)$, is less than that at a low-price-only equilibrium equal to $(p_L - c)\bar{C}$, then the equilibrium is uniquely attained at $\bar{C}$; otherwise, C^0 is the optimal long-run equilibrium when δ approaches one.*

Lemma 4 says when the discount factor is high enough, rationing is optimal if it brings more profits for the current period than no rationing. However, in general, the transition to a rationing strategy may not be optimal even though it leads to higher one-stage profits. As shown in Theorem 2, a critical value of capacity expectation determines when transitioning to rationing is optimal.

Liu and van Ryzin (2008) show that when customers have rational expectations, the optimal stocking quantity is C_R^0 if $\Pi(C_R^0, C_R^0) > (p_L - c)\bar{C}$; otherwise, the firm stocks $\bar{C}$ without rationing. Then based on the above two lemmas, Theorem 3 follows immediately. This result is intuitive too. As mentioned, a fundamental tradeoff under adaptive learning is between stocking more to increase short-term profits and stocking less to create expectations of shortages and thereby enhance long-run profits. Once the

Table 2 Equilibrium Capacities and Fill Rates Under Different Discount Factors

Discount factor	0.99	0.98	0.97	0.96
Equilibrium capacity	531.3	531.4	531.6	531.7
Equilibrium fill rate (%)	45.5	45.8	46.2	46.5

discount factor approaches one, the firm cares only about customers' long-run expectations and profits. Therefore, we would expect the adaptive learning equilibrium to be close to the rational expectation equilibrium.

We further find numerically that the adaptive learning equilibrium is a good approximation of the rational expectations equilibrium if the discount factor is high. Table 2 illustrates the equilibrium capacities under different discount factors when customers adaptively update their expectations of capacity.

In this example, the other parameters are set as $N = 1,000$, $p_H = 1$, $p_L = 0.8$, $c = 0.2$, $U = 2$, $\gamma = 0.5$, and $\theta = 0.5$. We can easily check that under rational expectations rationing is optimal, and the rationing equilibrium capacity C_R^0 is 531.1. Table 2 shows that the rationing equilibrium capacity under adaptive learning is indeed very close to the rational expectation outcome C_R^0.

5. Numerical Examples

In this section, we conduct extensive numerical examples to study comparative statics for the critical value of initial capacity expectations. In all the examples, we set $N = 1,000$, $p_H = 1$, $p_L = 0.8$, $c = 0.2$, and $U = 2$. The other parameters are $\gamma = 0.75$, $\theta = 0.7$ in Table 3(a); $\gamma = 0.75$, $\delta = 0.8$ in Table 3(b); $\delta = 0.8$, $\theta = 0.45$ in Table 3(c). In all these examples, the lower bound of capacity $\underline{C}$ is 500 and the upper bound of capacity $\bar{C}$ is 600. We numerically solved the problem (5) by a value iteration approach and graphed the optimal state updating functions. The critical value of capacity expectation is then determined by the point where the optimal state updating function jumps across the 45-degree line. The rationing equilibrium and associated one-stage profit are computed by (9) and (6), respectively. Table 3 summarizes the critical value of capacity expectation $\widehat{C}_c$, the rationing equilibrium capacity C^0, the fill rate equilibrium at C^0, and the long-run per stage profit at C^0, $\Pi(C^0, C^0)$, under different parameter settings of the firm's discount factor δ, the customers' learning speed θ, and their risk aversion γ.

We observe that the critical value of capacity expectation increases in the firm's discount factor (see Table 3(a)). For example, in the case of $\delta = 0.7$, only when customers' capacity expectation is less than 553.1 does rationing become the optimal long-run outcome. As δ increases to 0.9, the critical value of capacity expectation becomes 600; thus, rationing is always optimal regardless of customers' capacity expectation. Less discounting (i.e., a larger discount factor) means that future profits are more important to the firm. Therefore, there is increased incentive for a firm to endure the short-term losses of understocking to reap the long-term benefits of changing customers' capacity expectations. This leads to a larger critical value.

Table 3 Critical Values and Equilibrium Outcomes in Capacity Learning Model

(a) Under different discount factors

δ	$\widehat{C}_c$	C^0	q^0 (%)	$\Pi(C^0, C^0)$
0.5	500	514.2	37.0	401.3
0.6	500	513.1	32.7	402.1
0.7	553.1	512.1	28.9	402.5
0.8	574.4	511.2	25.6	402.8
0.9	600	510.3	22.7	403.0

(b) Under different learning speeds

θ	$\widehat{C}_c$	C^0	q^0 (%)	$\Pi(C^0, C^0)$
0.1	526.3	515.9	44.2	399.4
0.3	542.5	512.8	31.4	402.2
0.5	586.3	511.7	27.5	402.7
0.7	600	511.2	25.6	402.8
0.9	600	510.8	24.4	402.9

(c) Under different risk aversions

γ	$\widehat{C}_c$	C^0	q^0 (%)	$\Pi(C^0, C^0)$
0.25	600	561.6	71.8	429.7
0.50	600	534.4	52.6	413.0
0.75	570.9	511.9	28.2	402.6

We find that the critical value of the customers' capacity expectation increases in their learning speed. The more quickly customers adapt their expectations to changes in the firm's actual capacity choices, the lower the transition cost of changing their expectations. Hence, changing expectations to the rationing equilibrium is optimal for a larger range of customers' capacity expectations.

We also observe that the critical value increases in the customers' risk aversion (a smaller γ implies more risk-averse customers). Intuitively, when customers become more risk averse, only a smaller amount of rationing risk is needed to induce segmentation, and thus the lost-sales cost of rationing is lower. With a lower rationing cost, there is more incentive to incur the (modest) transition cost to change customers' expectations, and again a larger critical value results.

6. Extensions

Our research can be extended in a number of directions. As mentioned earlier, it is arguably more natural to model customer learning using fill rate rather than capacity estimates. We next provide the analysis for this alternative fill-rate updating model.

6.1. The Fill-Rate Learning Model

As in the capacity updating model, we assume that customers update their expectations by an exponential smoothing process; that is, given customers' expectation of fill rate for the current season t, denoted $\hat{q}_t$, and the firm's actual fill rate in season t, q_t, a customer's estimate of the fill rate for the next season $t+1$, $\hat{q}_{t+1}$, is

$$\hat{q}_{t+1} = \theta q_t + (1-\theta)\hat{q}_t, \tag{16}$$

where θ is the customers' learning speed, as before. As we did for the capacity learning process, we can define $V(\hat{q}_t)$, the maximum discounted profit given customers' expectation of fill rate q_t, by the following Bellman equation:

$$V(\hat{q}_t) = \max_{0\le q_t\le 1}\left\{\frac{N}{U}(p_H-c)(U-v(\hat{q}_t))+\frac{N}{U}(p_L-c)(v(\hat{q}_t)-p_L)q_t + \delta V(\theta q_t+(1-\theta)\hat{q}_t)\right\}, \tag{17}$$

where $v(\hat{q}_t)$ is the threshold valuation for a fill-rate expectation of $\hat{q}_t$ and is determined by

$$v(\hat{q}_t) = \min\left\{\frac{p_H - p_L\hat{q}_t^{1/\gamma}}{1-\hat{q}_t^{1/\gamma}}, U\right\}. \tag{18}$$

Note that in the capacity updating model, it is easy to show that the value function for the region 1 problem (7) is concave because the one-stage profit function, $(N/U)(p_H - p_L)(U - v(\widehat{C}_t)) + (p_L - c)C_t$, is jointly concave in $(\widehat{C}_t, C_t)$. However, in the fill-rate learning model (17), the one-stage profit function, $(N/U)(p_H-c)(U-v(\hat{q}_t))+(N/U)(p_L-c)(v(\hat{q}_t)-p_L)q_t$, is not jointly concave in $(\hat{q}_t, q_t)$ in the region 1 problem. Despite this structural difference in the form of the profit function, we can still obtain some analytical results for convergence in the fill-rate learning model. To this end, we change the decision variable to be the fill-rate expectation in the next period; the Bellman equation (17) can be then expressed as

$$V(\hat{q}_t) = \max_{0\le\hat{q}_{t+1}\le 1}\left\{\frac{N}{U}(p_H-c)(U-v(\hat{q}_t))+\frac{N}{\theta U}(p_L-c)(v(\hat{q}_t)-p_L) \cdot(\hat{q}_{t+1}-(1-\theta)\hat{q}_t)+\delta V(\hat{q}_{t+1})\right\}. \tag{19}$$

The following lemma shows that customers' expectations of fill rates converge.

Lemma 5. *The state paths $\{\hat{q}_t\}_{t\ge 1}$ determined by* (19) *are monotone and thus converge.*

However, beyond existence, we were not able to further analyze properties of the equilibrium as we did in the capacity model. We thus used value iteration to study numerical solutions to the fill-rate learning model. Table 4 summarizes the critical values of fill-rate expectations and the rationing equilibrium

Table 4 Critical Values and Equilibrium Outcomes in Fill-Rate Updating Model

(a) Under different discount factors

δ	$\hat{q}_c$ (%)	q^0 (%)	C^0	$\Pi(q^0, q^0)$
0.5	0.0	78.5	522.1	360.5
0.6	0.0	68.0	520.4	382.5
0.7	68.7	57.4	518.5	392.9
0.8	83.9	43.4	515.7	399.6
0.9	100.0	30.4	512.5	402.4

(b) Under different learning speeds

θ	$\hat{q}_c$ (%)	q^0	C^0	$\Pi(C^0, C^0)$
0.1	0.0	87.2	523.3	314.0
0.3	0.0	65.1	519.9	386.0
0.5	80.0	52.2	517.5	396.0
0.7	83.8	43.8	515.8	399.5
0.9	86.6	38.1	514.5	401.1

(c) Under different risk aversion

γ	$\hat{q}_c$ (%)	q^0 (%)	C^0	$\Pi(q^0, q^0)$
0.25	100.0	76.3	564.1	449.8
0.50	97.2	65.0	539.4	445.4
0.75	77.8	54.7	518.0	443.8

outcomes under different values of the firm's discount factor, the customers' learning speed, and their risk aversion. These are the same data used in the capacity learning example of §5.

As seen in Table 4, when the critical value of fill-rate expectation, $\hat{q}_c$, is strictly greater than zero and less than one, both rationing and low-price-only outcomes can be equilibria. In particular, when the customers' expectation of fill rate is higher than this critical value, the firm's capacity choices converge to a low-price-only equilibrium; otherwise, capacity rationing is optimal.

Observe from Table 4 that the critical value of fill-rate expectation increases in the firm's discount factor, the customers' learning speed, and their risk aversion, just as in the capacity updating model. We also observe that the rationing equilibrium decreases in the firm's discount factor and the customers' learning speed, whereas it increases in the customers' risk aversion. These results indicate that the qualitative behavior of the equilibrium outcome does not depend critically on whether the learning state variable is fill rate or capacity.[8]

6.2. Other Valuation Distributions and Utility Functions

Although all the analytical results derived in the paper are based on uniform distribution of customer valuations and the power utility function, we tested the model numerically on other distributions and utility functions such as exponential and truncated normal distribution for valuation and exponential and log utility functions. The numerical results show that, as before, customers' expectations of capacity converge to either a rationing equilibrium or a low-price-only equilibrium; the rationing equilibrium capacity decreases in the firm's discount factor and the customers' learning speed, whereas it increases in customers' risk aversion; the critical value of capacity expectation increases in the firm's discount factor, the customers' learning speed, and risk aversion. Thus the major results appear quite robust to more general valuation distributions and utility functions.

6.3. Asynchronous Expectation Updates

Our current model assumes that all customers have identical expectations and synchronous updating of expectations based on common knowledge of overall availability. It would be desirable to analyze the opposite extreme in which customers have heterogeneous expectations and asynchronously update them based on individual purchase experience. Our conjecture, however, is that this asynchronous learning would not result in qualitatively different decisions and system behavior. Rather than having one expectation, customers would have a mix of expectations, and each would decide to buy early or wait based on their individual expectation. Their aggregate demand would thus be a weighted average of their individual demand, and the firm's decisions to provide more or less availability would shift each customers' expectations in the corresponding direction. As in the homogeneous case, we conjecture that the firm's decision would be to move either to a rationing or low-price-only equilibrium, likely depending on some threshold of a weighted measure of customers' expectations. We have performed some numerical experiments with two-segment models of this type that support these conjectures. We also observed that although expectations may start out heterogenous, eventually the firm's availability converges and customer expectations converge to the true availability, so ultimately the rationing and low-price-only equilibria are the same as in the homogeneous case. It would be nice to confirm these conjectures formally, but our preliminary analysis suggests the value function for this model is considerably more difficult to work with.

7. Conclusions

Rationing can serve as an optimal selling strategy when customers strategize over the timing of purchases. Threatened by rationing risk, customers have an incentive to purchase early at the full price. Yet they may not perfectly anticipate rationing risk. This raises the question of how a firm should respond in terms of its capacity decision. Our model shows that rationing can be sustained as an equilibrium only if changing customer expectations is not very costly. When customers adjust their expectations slowly, future profits are deeply discounted, or customers are not very risk averse, the firm could end up serving the entire market at the discount price, even though rationing is more profitable in the long run. We also find when the firm's discount factor approaches one, the equilibrium produced by assuming customers adaptively learn converges to the equilibrium under rational expectations.

The optimal policy in our model converges, although it may be cyclical in other works, for example, Ovchinnikov and Milner (2009). This difference appears because of the convexity/concavity of payoffs; if the value function is convex (as for example in the model of Ovchinnikov and Milner 2009), it pays to oscillate between high and low availability because

[8] Note, however, that the critical value of fill-rate expectation and the rationing equilibria are quite different in the two models, even though the same parameters are used in both cases. This is because the mapping between capacity and fill rate is not linear, leading to different capacity and fill-rate updates and different equilibrium values. So although the qualitative behavior is the same, the numerical predictions are quite different.

the average payoff from being in two extreme states dominates the payoff from being at the average of the two states. When the value function is concave (as in our case of the region 1 problem), the payoff from oscillating between states cannot dominate the payoff from staying in one state. Yet how the particular model primitives in each case lead to this convex versus concave payoff appears somewhat trickier to determine. It would be desirable to have better business intuition on whether oscillating or constant availability is optimal.

Acknowledgments
The authors thank Gérard Cachon, an associate editor, and three referees for many helpful suggestions, in particular for providing a simplified proof of Lemma 2 and the analysis of fill-rate update model. Qian Liu was supported by the Hong Kong Research Grant Council, Grant RGC616807, and Garrett van Ryzin was supported by NSF Grant 0752662.

Appendix

Proof of Lemma 1. Take the first and second derivatives of $\widehat{C}_t(v_t)$ in (3):

$$\frac{d\widehat{C}_t}{dv_t} = \frac{N}{U}\left(\left(\frac{v_t-p_H}{v_t-p_L}\right)^{\gamma}\left(1+\frac{\gamma(p_H-p_L)}{v_t-p_H}\right)-1\right),$$

$$\frac{d^2\widehat{C}_t}{d{v_t}^2} = -\frac{N}{U}\gamma(1-\gamma)(p_H-p_L)^2(v_t-p_L)^{-1} \cdot (v_t-p_H)^{-2}\left(\frac{v_t-p_H}{v_t-p_L}\right)^{\gamma} \le 0.$$

Then $d\widehat{C}_t/dv_t$ strictly decreases in $v_t > p_H$. Together with the fact that $d\widehat{C}_t/dv_t|_{v_t\to+\infty} \to 0$, we have $d\widehat{C}_t/dv_t > 0$ when $v_t > p_H$. Hence, $\widehat{C}_t(v_t)$ is strictly increasing and concave in $v_t > p_H$. Note that $\widehat{C}_t(v_t)$ is continuous at p_H; thus $\widehat{C}_t(v_t)$ is strictly increasing and concave in $v_t \ge p_H$. □

Proof of Corollary 1. We first claim that there is a unique solution to

$$\left(\frac{v-p_H}{v-p_L}\right)^{\gamma}\left(1+\frac{\gamma(p_H-p_L)}{v-p_H}\right) - \frac{(p_L-c)(1/\delta-1)+\theta(p_H-c)}{(p_L-c)(1/\delta-1+\theta)} = 0. \tag{20}$$

Denote $g(v) = ((v-p_H)/(v-p_L))^{\gamma}(1+\gamma(p_H-p_L)/(v-p_H))$. Then $g(v)$ decreases in v; and $g(v)|_{v\to v_H} \to +\infty$ and $g(v)|_{v\to+\infty} \to 1$. Note that $((p_L-c)(1/\delta-1)+\theta(p_H-c))/((p_L-c)(1/\delta-1+\theta)) > 1$. Therefore, there exists the unique solution, denoted v^0, to the Equation (20). Because $((p_L-c)(1/\delta-1)+\theta(p_H-c))/((p_L-c)(1/\delta-1+\theta))$ increases in δ and θ, the solution v^0 to (20) decreases in δ and θ. Because C^0 increases in v^0, C^0 decreases in δ and θ as well.

Next we show the per stage profit increases in θ and δ. The profit per stage at the v^0, denoted $\Pi(v^0)$, is given by

$$\Pi(v^0) = \frac{N}{U}(p_H-c)(U-v^0)+(p_L-c)(v^0-p_L)\left(\frac{v^0-p_H}{v^0-p_L}\right)^{\gamma}.$$

One can easily check that $\Pi(v^0)$ is strictly concave in v^0 and maximized at v_R^0, where v_R^0 is the solution to the equation

$$\left(\frac{v-p_H}{v-p_L}\right)^{\gamma}\left(1+\frac{\gamma(p_H-p_L)}{v-p_H}\right) - \frac{p_H-c}{p_L-c} = 0.$$

Because

$$\frac{(p_L-c)(1/\delta-1)+\theta(p_H-c)}{(p_L-c)(1/\delta-1+\theta)} < \frac{p_H-c}{p_L-c},$$

it must be $v^0 > v_R^0$ when $0<\delta<1$. Because $\Pi(v^0)$ is decreasing in v^0 when $v^0 > v_R^0$, together with the fact that $((p_L-c)(1/\delta-1)+\theta(1-c))/((p_L-c)(1/\delta-1+\theta))$ increases in δ and θ, we conclude $\Pi(v^0)$ increases in δ and θ as well. □

Proof of Lemma 2. For the value function of the region 1 problem

$$V_1(\widehat{C}_t) = \max_{C_t\in S_1(\widehat{C}_t)}\left\{\frac{N}{U}(p_H-p_L)(U-v(\widehat{C}_t))+(p_L-c)C_t + \delta V_{t+1}^1(\theta C_t+(1-\theta)\widehat{C}_t)\right\},$$

we can easily show that $(N/U)(p_H-p_L)(U-v(\widehat{C}_t)) + (p_L-c)C_t$ is strictly joint concave in $(C_t, \widehat{C}_t)$. The reason is the following: according to Lemma 1, $v(\widehat{C}_t)$ is strictly concave in $\widehat{C}_t$; thus, $N/U(p_H-p_L)(U-v(\widehat{C}_t))+(p_L-c)C_t$ is strictly joint concave in $(C_t, \widehat{C}_t)$. Then the strict concavity of $V_1(\widehat{C}_t)$ directly follows Theorem 4.8 in Stokey and Lucas (1989).

Because one-stage profit function $(N/U)(p_H-p_L)\cdot(U-v(\widehat{C}_t))+(p_L-c)C_t$ decreases in $\widehat{C}_t$, it follows immediately from Theorem 4.7 in Stokey and Lucas (1989) that $V_1(\widehat{C}_t)$ decreases in $\widehat{C}_t$.[9] □

Proof of Proposition 1. We first show there exists a unique equilibrium for the region 1 problem. By the Contraction Mapping Theorem, it is sufficient to show the following state updating function is a contraction mapping:

$$f(\widehat{C}) = \theta C^*(\widehat{C})+(1-\theta)\widehat{C},$$

where $C^*(\widehat{C})$ is the optimal solution to the value function (7) given state $\widehat{C}$. By Theorem 4.8 in Stokey and Lucas (1989), $f(\widehat{C})$ is a single-valued and continuous function.

Because $V_1(\widehat{C})$ is strictly concave in $\widehat{C}$, the inner maximization term of (7) is strictly concave in C, given $\widehat{C}$ as well. The first-order conditions then sufficiently characterize optimal policies. We denote the first derivative of the inner maximization term of (7) with respect to C by $g(C)$:

$$g(C) = p_L - c + \theta\delta\frac{dV_1}{d\widehat{C}}.$$

If $C^*(\widehat{C})$ is an interior point of the set $S_1(\widehat{C})$, then $C^*(\widehat{C})$ is the optimal solution; otherwise, the optimal solution is attained at boundary points.

[9] Theorem 4.7 in Stokey and Lucas (1989) says that $V_1(\widehat{C}_t)$ is increasing in $\widehat{C}_t$ if the one-stage profit function increases in the first argument $\widehat{C}_t$. Using the same approach, we can show that $V_1(\widehat{C}_t)$ is decreasing in $\widehat{C}_t$ if the one-stage profit function decreases in the first argument $\widehat{C}_t$.

Figure A.1 Illustration of Contraction Mapping of State Updating Function

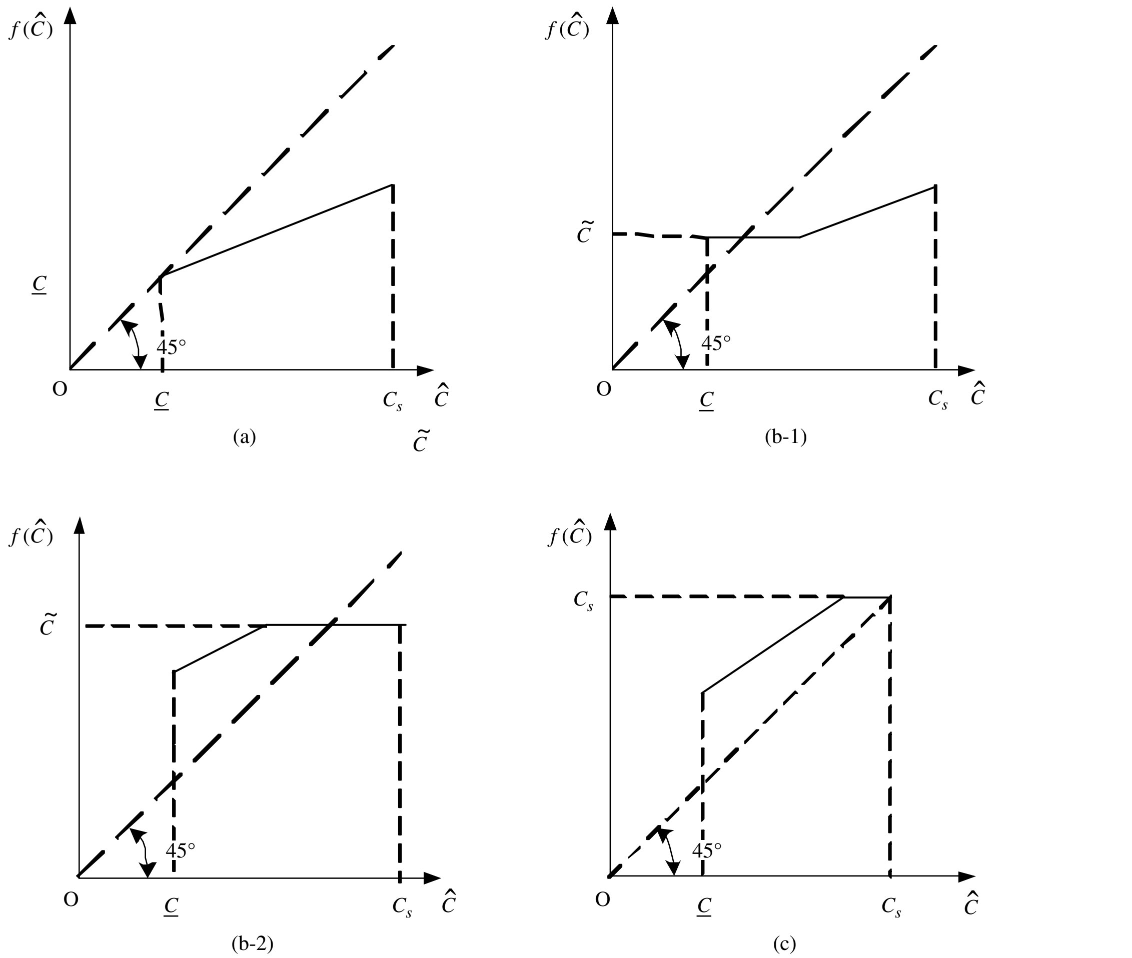

Case 1. $dV_1/d\hat{C}|_{\underline{C}} \le -(p_L - c)/\theta\delta$

Because $V_1(\hat{C})$ is decreasing in $\hat{C}$, $dV_1/d\hat{C}|_{\underline{C}} \le -(p_L - c)/\theta\delta$ implies $g(C) \le 0$ for any $\hat{C} \in [\underline{C}, C_s]$. It follows with $C^* = \underline{C}$ and $f(\hat{C}) = \theta\underline{C} + (1-\theta)\hat{C}$, which is illustrated in Figure A.1(a). Then, $\forall\, \hat{C}_1, \hat{C}_2 \in [\underline{C}, C_s]$, we have

$$|f(\hat{C}_1) - f(\hat{C}_2)| = (1-\theta)|\hat{C}_1 - \hat{C}_2|.$$

Therefore, $\hat{C} \mapsto f(\hat{C})$ is a contraction mapping if $\theta > 0$.[10]

Case 2. $dV_1/d\hat{C}|_{C_s} < -(p_L - c)/\theta\delta < dV_1/d\hat{C}|_{\underline{C}}$

Because of the strict concavity of $V_1(\cdot)$, there exists the unique $\tilde{C} \in (\underline{C}, C_s)$ such that $dV_1/d\hat{C}|_{\tilde{C}} = -(p_L - c)/\theta\delta$. Hence,

$$C^* = \begin{cases} -\dfrac{1-\theta}{\theta}\hat{C} + \dfrac{\tilde{C}}{\theta} & \text{if } \theta\bar{C} + (1-\theta)\hat{C} > \tilde{C} > \theta\underline{C} + (1-\theta)\hat{C}, \\ \underline{C} & \text{if } \tilde{C} \le \theta\underline{C} + (1-\theta)\hat{C}, \\ \bar{C} & \text{if } \tilde{C} \ge \theta\bar{C} + (1-\theta)\hat{C}. \end{cases}$$

[10] We exclude the case of $\theta = 0$, in which customer expectations stay constant at the initial 1; there is no learning at all.

and

$$f(\hat{C}) = \begin{cases} \tilde{C} & \text{if } \theta\bar{C} + (1-\theta)\hat{C} > \tilde{C} > \theta\underline{C} + (1-\theta)\hat{C}, \\ \theta\underline{C} + (1-\theta)\hat{C} & \text{if } \tilde{C} \le \theta\underline{C} + (1-\theta)\hat{C}, \\ \theta\bar{C} + (1-\theta)\hat{C} & \text{if } \tilde{C} \ge \theta\bar{C} + (1-\theta)\hat{C}. \end{cases}$$

Figure A.1(b-1) and (b-2) depict state-updating functions in this case. Then, $\forall\, \hat{C}_1, \hat{C}_2 \in [\underline{C}, C_s]$, we have

$$|f(\hat{C}_1) - f(\hat{C}_2)| \begin{cases} = 0 & \text{if } f(\hat{C}_1) = f(\hat{C}_2), \\ \le (1-\theta)|\hat{C}_1 - \hat{C}_2| & \text{otherwise.}^{11} \end{cases}$$

Again, $\hat{C} \mapsto f(\hat{C})$ is a contraction mapping if $\theta > 0$.

Case 3. $dV_1/d\hat{C}|_{C_s} \ge -(p_L - c)/\theta\delta$

[11] When $f(\hat{C}_1) \ne f(\hat{C}_2)$, it must be true that one of $\hat{C}_1$, $\hat{C}_2$ attains $\tilde{C}$. In other words, it cannot be the case that $f(\hat{C}_1)$ or $f(\hat{C}_2)$ is equal to $\theta\underline{C} + (1-\theta)\hat{C}$, and the other is $\theta\bar{C} + (1-\theta)\hat{C}$. This can be easily checked in Figure A.1(b-1) and (b-2).

Because $V_1(\widehat{C})$ decreases in $\widehat{C}$, $(dV_1)/d\widehat{C}|_{C_s} \geq -(p_L - c)/\theta\delta$ implies $g(C) \geq 0$ for any $\widehat{C} \in [\underline{C}, C_s]$. Then,

$$C^* = \begin{cases} \bar{C} & \text{if } \theta\bar{C} + (1-\theta)\widehat{C} \leq C_s, \\ -\dfrac{1-\theta}{\theta}\widehat{C} + \dfrac{C_s}{\theta} & \text{otherwise} \end{cases}$$

and

$$f(\widehat{C}) = \begin{cases} \theta\bar{C} + (1-\theta)\widehat{C} & \text{if } \theta\bar{C} + (1-\theta)\widehat{C} \leq C_s, \\ C_s & \text{otherwise.} \end{cases}$$

See Figure A.1(c) for the state updating function in this case. Then, $\forall\, \widehat{C}_1, \widehat{C}_2 \in [\underline{C}, C_s]$, it is easy to check

$$|f(\widehat{C}_1) - f(\widehat{C}_2)| \leq (1-\theta)|\widehat{C}_1 - \widehat{C}_2|.$$

Hence, $\widehat{C} \mapsto f(\widehat{C})$ is a contraction mapping if $\theta > 0$.

We next derive this unique equilibrium. To do this, we formulate the Bellman equation (7) as a sequence problem:

$$\begin{aligned} V_1(\widehat{C}_1) = \max_{\{\widehat{C}_{t+1}\}_{t=1}^{\infty}} \quad & \sum_{t=1}^{\infty} \delta^{t-1} \widehat{\Pi}(\widehat{C}_t, \widehat{C}_{t+1}) \\ \text{s.t.} \quad & \widehat{C}_{t+1} \in \hat{s}_1(\widehat{C}_t), \quad t = 1, 2, \ldots \\ & \widehat{C}_1 \in [\underline{C}, C_s] \text{ given,} \end{aligned}$$

where

$$\begin{aligned} \hat{s}_1(\widehat{C}) &= \{C \mid \min\{\theta\bar{C} + (1-\theta)\widehat{C}, C_s\} \geq C \geq \theta\underline{C} + (1-\theta)\widehat{C}\}, \\ \widehat{\Pi}(\widehat{C}_t, \widehat{C}_{t+1}) &= \frac{N}{U}(p_H - p_L)(U - v(\widehat{C}_t)) - \frac{1}{\theta}(1-\theta)(p_L - c)\widehat{C}_t \\ &\quad + \frac{1}{\theta}(p_L - c)\widehat{C}_{t+1}. \end{aligned}$$

Then the optimal state sequence $\{\widehat{C}^*_{t+1}\}$, provided $\widehat{C}^*_{t+1}$ is in the interior of $[\underline{C}, C_s]$, satisfies the Euler equation:

$$\frac{\partial \widehat{\Pi}(\widehat{C}_t, \widehat{C}_{t+1})}{\partial \widehat{C}_{t+1}}\bigg|_{(\widehat{C}^*_t, \widehat{C}^*_{t+1})} + \delta \frac{\partial \widehat{\Pi}(\widehat{C}_t, \widehat{C}_{t+1})}{\partial \widehat{C}_t}\bigg|_{(\widehat{C}^*_{t+1}, \widehat{C}^*_{t+2})} = 0.$$

This leads to

$$v'(\widehat{C}^*_{t+1}) = \frac{(p_L - c)(1/\delta - 1 + \theta)}{(N/U)\theta(p_H - p_L)}. \tag{21}$$

However,

$$\widehat{C}(v) = \frac{N}{U}\left(U - v + (v - p_L)\left(\frac{v - p_H}{v - p_L}\right)^{\gamma}\right), \quad 1 \leq v \leq U.$$

By the inverse function theorem,

$$\begin{aligned} v'(\widehat{C}) &= \frac{1}{\widehat{C}'(v)} \\ &= \left\{ \frac{N}{U}\left[\left(\frac{v - p_H}{v - p_L}\right)^{\gamma}\left(1 + \frac{\gamma(p_H - p_L)}{v - p_H}\right) - 1\right]\right\}^{-1}. \end{aligned} \tag{22}$$

Combine (21) and (22), and some algebraic calculations yield Equation (10). Further, it is easy to check that the left-hand side of (10) is strictly decreasing in $v > p_H$ and has different signs at $v \to p_H$ and $v \to +\infty$. Therefore, there exists the unique solution $v^0 > p_H$ to (10).

When $v^0 < U$, C^0 defined by (9) is in the interior of the set $[\underline{C}, C_s]$. Moreover, the transversality condition

$$\lim_{t\to\infty} \delta^t \widehat{C}^*_t \cdot \frac{\partial \widehat{\Pi}(\widehat{C}_t, \widehat{C}_{t+1})}{\partial \widehat{C}_t}\bigg|_{(\widehat{C}^*_t, \widehat{C}^*_{t+1})} = 0$$

is satisfied. Hence, the Euler and transversality conditions are sufficient for the optimality of (v^0, C^0) (Stokey and Lucas 1989, Theorem 4.15).[12]

When $v^0 > U$, no interior solution satisfies the Euler equation. The equilibrium capacity is attained at boundary points of either $\underline{C}$ or C_s. We claim it can only be C_s. The reason is the following:

Apply the Envelope Theorem to the following Bellman equation:

$$\begin{aligned} V_1(\widehat{C}_t) = \max_{\widehat{C}_{t+1} \in \hat{s}_1(\widehat{C}_t)} &\left\{ \frac{N}{U}(p_H - p_L)(U - v(\widehat{C}_t)) - \frac{1}{\theta}(1-\theta)(p_L - c)\widehat{C}_t \right. \\ &\left. + \frac{1}{\theta}(p_L - c)\widehat{C}_{t+1} + \delta V_1(\widehat{C}_{t+1}) \right\}, \end{aligned}$$

then

$$V_1'(\widehat{C}_t) = -\frac{N}{U}(p_H - p_L)v'(\widehat{C}_t) - \frac{1}{\theta}(1-\theta)(p_L - c).$$

According to (22),

$$\lim_{\widehat{C}_t \to \underline{C}^+} v'(\widehat{C}_t) = \lim_{v \to 1^+} \frac{1}{\widehat{C}'(v)} = 0.$$

Because the value function $V_1(\widehat{C}_t)$ is right continuous at $\underline{C}$ by the Maximum Theorem,

$$\lim_{\widehat{C}_t \to \underline{C}^+} V_1'(\widehat{C}_t) = -\frac{1}{\theta}(1-\theta)(p_L - c) > -\frac{1}{\theta\delta}(p_L - c)\ .$$

This implies Case 1 never exists, thus the equilibrium is uniquely attained at C_s. □

Proof of Proposition 2. Suppose there exists an equilibrium $\widetilde{C}$ for the original problem (5), and $\widetilde{C} \neq C^0$, $\widetilde{C} \neq \bar{C}$. Then $\widetilde{C} < C_s$ because if $C_s \leq \widetilde{C} < \bar{C}$, the value gained at $\widetilde{C}$ is strictly less than the value at $\bar{C}$. The value at $\widetilde{C}$ is

$$V(\widetilde{C}) = \sum_{t=0}^{\infty} \delta^t \Pi(\widetilde{C}, \widetilde{C}).$$

By the definition of the equilibrium,

$$\Pi(\widetilde{C}, \widetilde{C}) + \delta V(\widetilde{C}) \geq \Pi(\widetilde{C}, C) + \delta V(\theta C + (1-\theta)\widetilde{C}), \quad \forall C \in [\underline{C}, \bar{C}].$$

We now claim $\forall\, \widehat{C} \in [\underline{C}, C_s]$; the optimal value of the region 1 problem at state $\widehat{C}$ is less than that of the original problem at the same state, namely, $V_1(\widehat{C}) \leq V(\widehat{C})$. This is because the optimal policy for the region 1 problem is feasible for the original problem, and the two problems have the same dynamic program formulation.

[12] Theorem 4.15 from Stokey and Lucas (1989) requires the one-stage profit function $\widehat{\Pi}(\widehat{C}_t, \widehat{C}_{t+1})$ is strictly increasing in the first argument $\widehat{C}_t$. But this result also holds when $\widehat{\Pi}(\widehat{C}_t, \widehat{C}_{t+1})$ is strictly decreasing in the first argument $\widehat{C}_t$.

On the other hand, as shown, $\underline{C} \le \widetilde{C} < C_s$. Therefore the policy $\{\widetilde{C}, \widetilde{C}, \ldots\}$ is feasible for the region 1 problem. By the principle of optimality, $V_1(\widetilde{C}) \ge V(\widetilde{C})$. It then follows that

$$\begin{aligned}\Pi(\widetilde{C}, C) + \delta V_1(\theta C + (1-\theta)\widetilde{C}) &\le \Pi(\widetilde{C}, C) + \delta V(\theta C + (1-\theta)\widetilde{C}) \\ &\le \Pi(\widetilde{C}, \widetilde{C}) + \delta V(\widetilde{C}) \\ &\le \Pi(\widetilde{C}, \widetilde{C}) + \delta V_1(\widetilde{C}).\end{aligned}$$

Therefore, $\widetilde{C}$ is an equilibrium of the region 1 problem as well. Together with the fact that $\widetilde{C} < C_s$, it must be $\widetilde{C} = C^0$. This contradicts the assumption that $\widetilde{C} \ne C^0$. □

Proof of Proposition 3. We first claim that if there exists an optimal state path $\{\widehat{C}_t\}_{t\ge 1}$ such that $\widehat{C}_t < \widehat{C}_{t+1} = \cdots = \widehat{C}_{t+m}$ and $\widehat{C}_{t+m} > \widehat{C}_{t+m+1}$, or $\widehat{C}_t > \widehat{C}_{t+1} = \cdots = \widehat{C}_{t+m}$ and $\widehat{C}_{t+m} < \widehat{C}_{t+m+1}$, where $m > 1$. Then replacing the subpath $\widehat{C}_{t+m-1} \to \widehat{C}_{t+m} \to \widehat{C}_{t+m+1}$ along $\{\widehat{C}_t\}_{t\ge 1}$ by $\widehat{C}_{t+m-1} \to \widehat{C}_{t+m+1}$ yields an optimal state path as well.

As defined, $V(\widehat{C}_1)$ is the total discounted value collected along the optimal state path $\{\widehat{C}_t\}_{t\ge 1}$. Let $V_1(\widehat{C}_1)$ be the value collected along the state path that substitutes $\widehat{C}_{t+m-1} \to \widehat{C}_{t+m} \to \widehat{C}_{t+m+1}$ on the optimal state path with $\widehat{C}_{t+m-1} \to \widehat{C}_{t+m+1}$ (i.e., take one step from $\widehat{C}_{t+m-1}$ to $\widehat{C}_{t+m+1}$). Denote $V_2(\widehat{C}_1)$ as the value along the state path that replaces $\widehat{C}_{t+m-1} \to \widehat{C}_{t+m} \to \widehat{C}_{t+m+1}$ on the optimal state path by $\widehat{C}_{t+m-1} \to \widehat{C}_{t+m} \to \widehat{C}_{t+m} \to \widehat{C}_{t+m+1}$ (i.e., stays at $\widehat{C}_{t+m}$ for one more stage). Then,

$$\begin{aligned}V(\widehat{C}_t) &= \Pi(\widehat{C}_t, C_t) + (\delta + \cdots + \delta^{m-1})\Pi(\widehat{C}_{t+1}, \widehat{C}_{t+1}) \\ &\quad + \delta^m \Pi(\widehat{C}_{t+m}, C_{t+m}) + \delta^{m+1} V(\widehat{C}_{t+m+1}), \\ V_1(\widehat{C}_t) &= \Pi(\widehat{C}_t, C_t) + (\delta + \cdots + \delta^{m-2})\Pi(\widehat{C}_{t+1}, \widehat{C}_{t+1}) \\ &\quad + \delta^{m-1} \Pi(\widehat{C}_{t+m}, C_{t+m}) + \delta^{m} V(\widehat{C}_{t+m+1}), \\ V_2(\widehat{C}_t) &= \Pi(\widehat{C}_t, C_t) + (\delta + \cdots + \delta^{m})\Pi(\widehat{C}_{t+1}, \widehat{C}_{t+1}) \\ &\quad + \delta^{m+1} \Pi(\widehat{C}_{t+m}, C_{t+m}) + \delta^{m+2} V(\widehat{C}_{t+m+1}),\end{aligned}$$

where

$$C_t = \frac{\widehat{C}_{t+1} - (1-\theta)\widehat{C}_t}{\theta} \quad \text{and} \quad C_{t+m} = \frac{\widehat{C}_{t+m+1} - (1-\theta)\widehat{C}_{t+m}}{\theta}.$$

By the optimality of $V(\widehat{C}_t)$, $V(\widehat{C}_t) \ge V_1(\widehat{C}_t)$, implying

$$V(\widehat{C}_{t+m+1}) \le \frac{\Pi(\widehat{C}_{t+m}, \widehat{C}_{t+m})}{\delta(1-\delta)} - \frac{\Pi(\widehat{C}_{t+m}, C_{t+m})}{\delta}.$$

Similarly, $V(\widehat{C}_t) \ge V_2(\widehat{C}_t)$, following with

$$V(\widehat{C}_{t+m+1}) \ge \frac{\Pi(\widehat{C}_{t+m}, \widehat{C}_{t+m})}{\delta(1-\delta)} - \frac{\Pi(\widehat{C}_{t+m}, C_{t+m})}{\delta}.$$

Therefore, it must be

$$V(\widehat{C}_{t+m+1}) = \frac{\Pi(\widehat{C}_{t+m}, \widehat{C}_{t+m})}{\delta(1-\delta)} - \frac{\Pi(\widehat{C}_{t+m}, C_{t+m})}{\delta}; \quad \text{and}$$

$$V(\widehat{C}_t) = V_1(\widehat{C}_t) = V_2(\widehat{C}_t).$$

Thus, if it is optimal to stay at state $\widehat{C}_{t+m}$ for m stages ($m > 1$), it is optimal as well to stay there for $m-1$ stages.

Now suppose that the optimal state path $\{\widehat{C}_t\}_{t\ge 1}$ is not monotone. It then suffices to restrict attention to the case in which there exists a subpath $\widehat{C}_t \to \widehat{C}_{t+1} \to \widehat{C}_{t+2}$ on the optimal state path such that either $\widehat{C}_t < \widehat{C}_{t+1}$ and $\widehat{C}_{t+1} > \widehat{C}_{t+2}$, or $\widehat{C}_t > \widehat{C}_{t+1}$ and $\widehat{C}_{t+1} < \widehat{C}_{t+2}$.

Let $V(\widehat{C}_1)$ again be the total discounted value collected along the optimal state path $\{\widehat{C}_t\}_{t\ge 1}$; let $V_1(\widehat{C}_1)$ be the value collected along the state path that substitutes $\widehat{C}_t \to \widehat{C}_{t+1} \to \widehat{C}_{t+2}$ on the optimal state path with $\widehat{C}_t \to \widehat{C}_{t+2}$ (i.e., take one step from $\widehat{C}_t$ to $\widehat{C}_{t+2}$; the feasibility of this is argued below); and let $V_2(\widehat{C}_1)$ be the value collected along the state path that replaces $\widehat{C}_t \to \widehat{C}_{t+1} \to \widehat{C}_{t+2}$ on the optimal state path by $\widehat{C}_t \to \widehat{C}_{t+1} \to \widehat{C}_{t+1} \to \widehat{C}_{t+2}$ (i.e., stays at $\widehat{C}_{t+1}$ for one more stage). Then

$$\begin{aligned}V(\widehat{C}_t) &= \Pi(\widehat{C}_t, C_t) + \delta\Pi(\widehat{C}_{t+1}, C_{t+1}) + \delta^2 V(\widehat{C}_{t+2}), \\ V_1(\widehat{C}_t) &= \Pi(\widehat{C}_t, C'_t) + \delta V(\widehat{C}_{t+2}), \\ V_2(\widehat{C}_t) &= \Pi(\widehat{C}_t, C_t) + \delta\Pi(\widehat{C}_{t+1}, \widehat{C}_{t+1}) \\ &\quad + \delta^2\Pi(\widehat{C}_{t+1}, C_{t+1}) + \delta^3 V(\widehat{C}_{t+2}),\end{aligned}$$

where

$$C_t = \frac{\widehat{C}_{t+1} - (1-\theta)\widehat{C}_t}{\theta}, \quad C_{t+1} = \frac{\widehat{C}_{t+2} - (1-\theta)\widehat{C}_{t+1}}{\theta},$$

$$C'_t = \frac{\widehat{C}_{t+2} - (1-\theta)\widehat{C}_t}{\theta}.$$

One can easily show that C'_t is a feasible action. By the optimality of $V(\widehat{C}_t)$, that is, $V(\widehat{C}_t) \ge \max\{V_1(\widehat{C}_t), V_2(\widehat{C}_t)\}$. Along the same line of argument as before, it must be that $V(\widehat{C}_{t+2}) = (1/(1-\delta))\Pi(\widehat{C}_{t+1}, \widehat{C}_{t+1}) + (1/(\theta\delta))(p_L - c) \cdot (\widehat{C}_{t+1} - \widehat{C}_{t+2})$, and $V(\widehat{C}_t) = V_1(\widehat{C}_t) = V_2(\widehat{C}_t)$. Therefore, we can replace the nonmonotone subpath $\widehat{C}_t \to \widehat{C}_{t+1} \to \widehat{C}_{t+2}$ by $\widehat{C}_t \to \widehat{C}_{t+2}$ and achieve at least as large as the optimal value $V(\widehat{C}_t)$. □

Proof of Proposition 4. If there exists an optimal nonmonotone state path $\{\widehat{C}_t\}_{t\ge 1}$, then along it there exists a subpath $\widehat{C}_t \to \widehat{C}_{t+1} \to \cdots \to \widehat{C}_{t+m} \to \widehat{C}_{t+m+1}$ ($m \ge 1$) such that either $\widehat{C}_t < \widehat{C}_{t+1} = \cdots = \widehat{C}_{t+m}$ and $\widehat{C}_{t+m} > \widehat{C}_{t+m+1}$, or $\widehat{C}_t > \widehat{C}_{t+1} = \cdots = \widehat{C}_{t+m}$ and $\widehat{C}_{t+m} < \widehat{C}_{t+m+1}$. According to the proof of Proposition 3, it is also optimal to stay at $\widehat{C}_{t+1}$ for $m-1$ stages if staying at $\widehat{C}_{t+1}$ for m stages is optimal ($m > 1$). The optimal value at state $\widehat{C}_t$ can then be expressed as follows:

$$V(\widehat{C}_t) = \Pi(\widehat{C}_t, C_t) + \delta\Pi(\widehat{C}_{t+1}, C_{t+1}) + \delta^2 V(\widehat{C}_{t+m+1}), \tag{23}$$

where

$$C_t = \frac{\widehat{C}_{t+1} - (1-\theta)\widehat{C}_t}{\theta}, \; C_{t+1} = \frac{\widehat{C}_{t+m+1} - (1-\theta)\widehat{C}_{t+1}}{\theta}.$$

As shown in Proposition 3, $V(\widehat{C}_{t+m+1}) = \Pi(\widehat{C}_{t+1}, \widehat{C}_{t+1}) / (\delta(1-\delta)) - \Pi(\widehat{C}_{t+1}, C_{t+1})/\delta$. If we substitute it into (23), we have $V(\widehat{C}_t) = \Pi(\widehat{C}_t, C_t) + (\delta/(1-\delta))\Pi(\widehat{C}_{t+1}, \widehat{C}_{t+1})$. However, $V(\widehat{C}_t) = \Pi(\widehat{C}_t, C_t) + \delta V(\widehat{C}_{t+1})$; therefore,

$$V(\widehat{C}_{t+1}) = \frac{\Pi(\widehat{C}_{t+1}, \widehat{C}_{t+1})}{1-\delta}.$$

This implies $\widehat{C}_{t+1}$ is an equilibrium of the original problem. Note that $\widehat{C}_{t+1} < \bar{C}$, it must be $\widehat{C}_{t+1} = C^0$.

Next we show if there exists a nonmonotone optimal state path along which

$$\widehat{C}_t \to \underbrace{C^0 \to \cdots \to C^0}_{m} \to \widehat{C}_{t+m+1} (m \geq 1),$$

it must be $\widehat{C}_t > C^0$ and $\widehat{C}_{t+m+1} > C^0$. The reason is the following: suppose $\widehat{C}_t < C^0$ and $\widehat{C}_{t+m+1} < C^0$; then $\{\widehat{C}_i\}_{i \geq t+m+1}$ must be a decreasing sequence and converge to a state strictly less than C^0. This contradicts Theorem 1. Therefore, $\widehat{C}_t > C^0$ and $\widehat{C}_{t+m+1} > C^0$. Further, because $\{\widehat{C}_i\}_{i \geq t+m+1}$ is an increasing sequence, it will converge to $\bar{C}$. □

Proof of Lemma 3. The cutoff value v^0 under the adaptive learning model is the solution to

$$\left(\frac{v - p_H}{v - p_L}\right)^{\gamma} \left(1 + \frac{\gamma(p_H - p_L)}{v - p_H}\right) - \frac{(p_L - c)(1/\delta - 1) + \theta(p_H - c)}{(p_L - c)(1/\delta - 1 + \theta)} = 0. \tag{24}$$

This equation becomes (15) as $\delta \to 1$. Therefore, $v^0 \to v_R^0$ and $C^0 \to C_R^0$ as $\delta \to 1$. □

Proof of Lemma 4. Let $\widetilde{V}(C^0)$ and $\widehat{V}(C^0)$ be the values of 5 at C^0 when the capacity decisions take C^0 and $\bar{C}$ for any stage, respectively; then

$$\widetilde{V}(C^0) = \Pi(C^0, C^0) + \delta\Pi(C^0, C^0) + \cdots, \tag{25}$$

$$\widehat{V}(C^0) = \Pi(C^0, \bar{C}) + \delta\Pi(\theta\bar{C} + (1-\theta)C^0, \bar{C}) + \cdots. \tag{26}$$

Note the first term in (25) is strictly less than the first one in (26); that is, $\Pi(C^0, C^0) < \Pi(C^0, \bar{C})$. Because $\Pi(C^0, C^0) \leq (p_L - c)\bar{C} \leq \Pi(\widehat{C}, \bar{C})$, $\forall\, \widehat{C} \in [\underline{C}, \bar{C}]$, from the second term on, (25) is at most as large as the counterpart in (26) on a term-by-term basis. Therefore, it must be $\widetilde{V}(C^0) < \widehat{V}(C^0)$. It then follows that C^0 cannot be an equilibrium.

We next show that C^0 is the optimal equilibrium when δ approaches 1 and $\Pi(C^0, C^0) < (p_L - c)\bar{C}$. Suppose there exists state $\widehat{C}_1 \in [\underline{C}, \bar{C}]$ such that an optimal state path from $\widehat{C}_1$, denoted $\{\widehat{C}_t\}_{t \geq 1}$, converges to $\bar{C}$. By the monotonicity of this optimal state path, there exists $m \geq 1$ such that $\widehat{C}_t \geq C_s$, $\forall\, t \geq m$ (recall that C_s is the state which divides the problem into the region 1 and region 2 problems); moreover, the optimal policy $C_t^* = \bar{C}$, $\forall\, t \geq m$. That is,

$$V(\widehat{C}_1) = \Pi(\widehat{C}_1, C_1^*) + \delta\Pi(\widehat{C}_2, C_2^*) + \cdots + \delta^{m-2}\Pi(\widehat{C}_{m-1}, C_{m-1}^*) + \sum_{t=m}^{\infty} \delta^{t-1}(p_L - c)\bar{C}.$$

Construct an alternative state path from $\widehat{C}_1$, denoted $\{\widetilde{C}_t\}_{t \geq 1}$ ($\widetilde{C}_1 = \widehat{C}_1$), which visits C^0 at stage $t = n$ and stays there afterward; then the value along this path $\{\widetilde{C}_t\}_{t \geq 1}$ denoted by $\widetilde{V}(\widetilde{C}_1)$ is

$$\widetilde{V}(\widetilde{C}_1) = \Pi(\widetilde{C}_1, C_1) + \delta\Pi(\widetilde{C}_2, C_2) + \cdots + \delta^{n-2}\Pi(\widetilde{C}_{n-1}, C_{n-1}) + \sum_{t=n}^{\infty} \delta^{t-1}\Pi(C^0, C^0).$$

Since $V(\widehat{C}_1) \geq \widetilde{V}(\widetilde{C}_1)$, some algebraic arrangement yields

$$\begin{aligned} &\frac{1}{1-\delta}\left(\delta^{m-1}(p_L - c)\bar{C} - \delta^{n-1}\Pi(C^0, C^0)\right) \\ &\geq \Pi(\widetilde{C}_1, C_1) + \delta\Pi(\widetilde{C}_2, C_2) + \cdots + \delta^{n-2}\Pi(\widetilde{C}_{n-1}, C_{n-1}) \\ &\quad - \Pi(\widehat{C}_1, C_1^*) - \delta\Pi(\widehat{C}_2, C_2^*) - \cdots - \delta^{m-2}\Pi(\widehat{C}_{m-1}, C_{m-1}^*). \end{aligned} \tag{27}$$

As $\delta \to 1^-$, the left-hand side of (27) goes to $-\infty$ because $\Pi(C^0, C^0) > (p_L - c)\bar{C}$. However, the right-hand side of (27) is a finite number. Hence, there does not exist such a state $\widehat{C}_1$ from which an optimal state path converges to $\bar{C}$, and therefore all the optimal state paths from any state converge to C^0 as $\delta \to 1^-$. □

Proof of Lemma 5. In problem (19), because $v(\hat{q}_t)$ is increasing in $\hat{q}_t$, the function under maximization has positive cross-derivative and hence is supermodular in $(\hat{q}_t, \hat{q}_{t+1})$. Hence, the optimal decision, $\hat{q}_{t+1}$, is monotone increasing in the state $\hat{q}_t$. Because the optimal decision is monotone increasing in the state in each period, and the optimal decision in the current period becomes the state for the next period, we conclude the optimal decisions, $\{\hat{q}_t\}_{t \geq 1}$, are monotone and converge. □

References

Akerman, G. 1957. The cobweb theorem: A reconsideration. *Quart. J. Econom.* **71**(1) 151–160.

Anderson, C. K., J. G. Wilson. 2003. Wait or buy? The strategic consumer: Pricing and profit implications. *J. Oper. Res. Soc.* **54**(3) 299–306.

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

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

Byers, R. E., D. J. Huff. 2005. No soup for you! Intentionally unsatisfied demand and future production. Presentation, Manufacturing Service Oper. Management 2005 Conference, June 26–28, Evanston, IL.

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

Elmaghraby, W., A. Gülcü, P. Keskinocak. 2008. Designing optimal preannounced markdowns in the presence of rational customers with multi-unit demands. *Manufacturing Service Oper. Management* **10**(1) 126–148.

Gallego, G., R. Phillips, O. Sahin. 2008. Strategic management of distressed inventory. *Production Oper. Management* **17**(4) 402–415.

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

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

Greenleaf, E. A. 1995. The impact of reference price effects on the profitability of price promotions. *Marketing Sci.* **14**(1) 82–104.

Helson, H. 1964. *Adaptation-Level Theory.* Harper & Row, New York.

Jacobson, R., C. Obermiller. 1990. The formation of expected future prices: A reference price for forward-looking consumers. *J. Consumer Res.* **16**(4) 420–432.

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

Kleywegt, A. J. 2006. Dynamic pricing with buyer learning. Extended Abstract, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta.

Kopalle, P. K., A. Rao, J. L. Assunção. 1996. Asymmetric reference price effects and dynamic pricing policies. *Marketing Sci.* **15**(1) 60–85.

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

Monroe, K. B. 1973. Buyer's subjective perceptions of price. *J. Marketing Res.* **10**(1) 70–80.

Nerlove, M. 1958. Adaptive expectations and cobweb phenomena. *Quart. J. Econom.* **72**(2) 227–240.

Ovchinnikov, A., J. M. Milner. 2009. Revenue management with end-of-period discounts in the presence of customer learning. Working paper, University of Virginia, Charlottesville.

Popescu, I., Y. Wu. 2007. Dynamic pricing strategies with reference effects. *Oper. Res.* **55**(3) 413–429.

Rubinstein, A. 1998. *Modeling Bounded Rationality*. MIT Press, Cambridge, MA.

Rump, C., S. Stidham. 1998. Stability and chaos in input pricing for a service facility with adaptive customer response to congestion. *Management Sci.* **44**(2) 246–261.

Sterman, J. D. 1987. Expectation formulation in behavioral simulation models. *Behav. Sci.* **32**(3) 190–211.

Sterman, J. D. 1989. Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Sci.* **35**(3) 321–339.

Stidham, S. 1992. Pricing and capacity decisions for a service facility: Stability and multiple local optima. *Management Sci.* **38**(8) 1121–1139.

Stokey, N., R. E. Lucas. 1989. *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, MA.

Su, X. 2007. Inter-temporal 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.