



University of Pennsylvania
ScholarlyCommons

Operations, Information and Decisions Papers


Wharton Faculty Research

7-2010

Intertemporal Pricing and Consumer Stockpiling

Xuanming Su
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/oid_papers

 Part of the [Educational Assessment, Evaluation, and Research Commons](#), and the [Sales and Merchandising Commons](#)

Recommended Citation

Su, X. (2010). Intertemporal Pricing and Consumer Stockpiling. *Operations Research*, 58 (4), 1133-1147.
<http://dx.doi.org/10.1287/opre.1090.0797>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/oid_papers/114
For more information, please contact repository@pobox.upenn.edu.

Intertemporal Pricing and Consumer Stockpiling

Abstract

We study a dynamic pricing problem for a class of products with stable consumption patterns (e.g., household items, staple foods). Consumers may stock up the product at current prices for future consumption, but they incur inventory holding costs. We model this situation as a dynamic game over an infinite time horizon: in each period, the seller sets a price, and each consumer chooses how many units to buy. We develop a solution methodology based on rational expectations. By endowing each player with beliefs, we decouple the dynamic game into individual dynamic programs for each player. We solve for the rational expectations equilibrium, where all players make optimal dynamic decisions given correct beliefs about others' behavior. In equilibrium, the seller may either charge a constant fixed price or offer periodic price promotions at predictable time intervals. We show that promotions are useful when frequent shoppers are willing to pay more for the product than are occasional shoppers. We also develop several model extensions to study the impact of consumer stockpiling on the seller's inventory, production, and rationing strategies.

Keywords

dynamic pricing, stockpiling, consumer inventory, promotions, rational expectations, price discrimination

Disciplines

Educational Assessment, Evaluation, and Research | Sales and Merchandising

Inter-temporal Pricing and Consumer Stockpiling

Xuanming Su

Haas School of Business, University of California, Berkeley, CA 94720

Abstract

We study a dynamic pricing problem for a class of products with stable consumption patterns (e.g., household items, staple foods). Consumers may stock up the product at current prices for future consumption, but they incur inventory holding costs. We model this situation as a dynamic game over an infinite time horizon: in each period, the seller sets a price, and each consumer chooses how many units to buy. We develop a solution methodology based on rational expectations. By endowing each player with beliefs, we decouple the dynamic game into individual dynamic programs for each player. We solve for the rational expectations equilibrium, where all players make optimal dynamic decisions given correct beliefs about others' behavior. In equilibrium, the seller may either charge a constant fixed price or offer periodic price promotions at predictable time intervals. We show that promotions are useful when frequent shoppers are willing to pay more than occasional shoppers for the product. We also develop several model extensions to study the impact of consumer stockpiling on the seller's inventory, production, and rationing strategies.

June 2009

1 Introduction

In this paper, we study a dynamic pricing problem in the context of consumer goods. Unlike traditional revenue management settings (e.g., airlines), consumer goods exhibit two important differences. First, they can be purchased in bulk and inventoried for future use. Second, demand is relatively stable, so consumers can predict their usage over time. Consequently, stockpiling behavior often arises as a strategic response to price promotions: when the price is low, consumers will plan ahead and stock up for future consumption. Such strategic behavior generates new challenges for revenue management of consumer products. The goal of this paper is to study the seller's optimal dynamic pricing strategies and consumers' optimal stockpiling strategies.

In our model, there is a monopolist selling to a market of consumers over an infinite time horizon. We formulate a dynamic game between the seller and consumers: during each time period, the seller sets a price for the product, and consumers choose how many units to buy. The product is consumed at a steady rate. Consumers may forward buy for future consumption, but they incur inventory holding costs. There is also a fixed cost associated with making any purchase. The population is heterogeneous, so different types of consumers will make different decisions. We assume that the seller wishes to maximize long-run average profits and consumers wish to maximize long-run average utility.

To study this dynamic game, we introduce a methodology based on rational expectations. Our fundamental premise is that players form expectations over others' behavior. These expectations allow us to decouple the dynamic game into dynamic programs for single agents. Given beliefs over consumer behavior, the seller's pricing problem becomes a Markov decision process. Similarly, given beliefs over the seller's pricing strategy and others' purchase behavior, each consumer's inventory problem can also be decoupled. We introduce the notion of a rational expectations (RE) equilibrium, in which consumers behave optimally given their beliefs, the seller prices optimally given consumer demand, and beliefs are consistent with actual outcomes. We demonstrate the existence of a RE equilibrium. Moreover, using duality theory, we show that an optimal dynamic strategy (or policy) can be mapped into a cycle of price/inventory states, with transitions through this cycle of states being driven by optimal pricing and stocking decisions. Based on this graph-theoretic analogy, we show that the RE equilibrium involves cyclical outcomes. In particular, there will be price cycles that repeat over time.

Our results provide several managerial insights on pricing strategies for consumer products that may be stockpiled. First, our model sheds some light on optimal pricing formats. When is it optimal to offer promotions and when is it preferable to charge “everyday-low-prices”? Interestingly, our model admits both types of pricing strategies in equilibrium. We show that periodic promotions should be used when frequent shoppers are willing to pay more for the product, compared to occasional shoppers. In such cases, promotions are a useful price discrimination device: occasional shoppers stockpile on promotion and frequent shoppers are willing to pay the regular price during their off-promotion visits. From the classic Economic Order Quantity (EOQ) model, note that frequent shoppers are likely the consumers with high holding costs, low fixed costs, and high consumption rates. Therefore, our results also suggest that promotions are effective for products where consumers with high holding costs (or low fixed costs or high consumption rates) also have higher valuations for the product. In addition, we derive comparative results on how the depth and frequency of promotions depend on market characteristics. Our results also illustrate how consumer stockpiling can generate the “post-promotion dip” phenomenon: immediately after a promotion, there is a significant slump in quantities sold.

We then study several model extensions and the key findings are summarized here. First, we consider how consumer stockpiling affects the seller’s inventory strategies. We find that the optimal inventory policy, in contrast to standard EOQ-type results, may involve different order quantities at variable time intervals. This observation suggests that stockpiling behavior on the consumer end can initiate the bullwhip effect in supply chains. Second, we study the seller’s production planning strategies. We find that the capacitated seller may sometimes need to produce in advance because consumer stockpiling during promotions can lead to demand spikes. This generates increased inventory costs and thus makes promotional pricing less attractive. Finally, we consider inventory rationing strategies (i.e., intentionally limiting quantities sold) and show that they can improve the profitability of price promotions. This is because the threat of stockouts during promotions may induce consumers to buy off-promotion at regular prices.

The remainder of the paper is structured as follows. We present a literature review in Section 2. We describe our model in Section 3, and present the problem formulation and technical analysis in Section 4. This is followed by a numerical study in Section 5. In Section 6, we analyze a special case that yields explicit solutions and practical insights. In Section 7, we study several model extensions. Finally, we offer concluding remarks in Section 8. All proofs are presented in the Appendix.

2 Literature Review

Revenue Management

This paper is closely related to the recent stream of literature on optimal dynamic pricing of finite inventories to strategic customers who form rational expectations of future prices. Representative papers, each adopting different modeling approaches, include Aviv and Pazgal (2008), Gallego, Phillips, Sahin (2008), Liu and van Ryzin (2008a), Su (2007), Gallien (2006), Elmaghraby, Gulcu, Keskinocak (2008), Xu and Hopp (2005), Ovchinnikov and Milner (2005), Levin, McGill, Nediak (2008), Zhou, Cho, Fan (2005), Yin, Aviv, Pazgal and Tang (2009). The book by Talluri and van Ryzin (2005) provides a comprehensive survey of the field of revenue management, and the survey by Shen and Su (2007) reviews emerging work in this area related to strategic purchasing behavior. In particular, Yin, Aviv, Pazgal and Tang (2009) is the most closely related to the current work. The authors also utilize the rational expectations approach in their equilibrium analysis. Another closely related paper is by Liu and van Ryzin (2008b), who show that rational expectations can emerge from an adaptive learning process. While the papers above consider inventory on the seller side, we focus on consumer inventory, which drives stockpiling decisions. We adopt a similar modeling approach and study strategic consumers who are active decision-makers. Through this common paradigm, this paper introduces consumer stockpiling behavior into the revenue management literature.

Consumer Response to Promotions

There is a large empirical marketing literature that studies consumer response to price promotions. Consistent with consumer stockpiling, during promotions, there should be increases in purchase quantity as well as duration till the next purchase. Empirical evidence for such “purchase acceleration” is provided by Shoemaker (1979) and Ward and Davis (1978), two of the earliest studies (see also Blattberg, Eppen and Lieberman, 1981; Neslin, Henderson and Quelch, 1985; Hendel and Nevo, 2006a). Another group of papers, Gupta (1988), Chiang (1991), Chintagunta (1993), and Bell, Chiang and Padmanabhan (1999), attempt to separate the primary demand effects (brand switching) from the secondary demand effects (purchase acceleration and stockpiling) of promotions. They do so by incorporating consumer brand choice decisions. All these studies find that the secondary demand effect is larger than the primary demand effect. There is another set of papers (see Gonul and Srin-

vasan, 1996; Erdem, Imai, and Keane, 2003; and Hendel and Nevo, 2006b) that structurally estimate a dynamic forward-looking model in which consumers form expectations of future prices. Broadly speaking, these papers demonstrate that consumer expectations have a significant impact on purchase decisions. These results motivate our rational expectations approach, whereby we posit that consumer stockpiling is influenced by their future price beliefs.

There are several theoretical models that characterize optimal consumer stockpiling strategies. In these models, the firms's prices are exogenously specified. Meyer and Assuncao (1990) study the consumer's purchase problem in response to random prices that are i.i.d. over time. Their analysis is based on Golabi's (1985) inventory control model, applied to consumer inventories. Krishna (1992) extend the Golabi model to the case of multiple brands. Assuncao and Meyer (1993) further extend this line of work by incorporating consumption decisions; they also consider prices that follow a first-order stochastic process. Krishna (1994) allows for an arbitrary distribution for the inter-arrival time between deals and characterizes the consumer's optimal purchasing policy. Ho, Tang, and Bell (1998) derive the optimal purchasing policy when consumers face fixed shopping costs. They show that under i.i.d. prices, the consumer's optimal purchase quantity is linear in the difference between current price and average price. The papers above assume that the seller's prices follow an exogenous stochastic process (e.g., i.i.d. or first-order Markov). Our paper adopts a different approach: we endogenously derive the seller's optimal pricing strategies in a rational expectations framework. In fact, we show that price cycles (significantly different from i.i.d. or first-order Markov prices) emerge in equilibrium.

Optimal Pricing and Consumer Stockpiling

There are only a few papers that study both the seller's pricing decisions and consumers' stocking decisions as we do. Blattberg, Eppen, and Lieberman (1981) introduces a model in which the seller offers periodic promotions to two types of consumers, with low or high holding costs. Low-cost consumers forward buy during promotions and thus alleviate the seller's burden of holding inventory, while high-cost consumers pay the full price. This suggests that offering periodic promotions is an effective way to transfer inventory carrying costs from the seller to the consumer. In their model, the seller is assumed to offer periodic promotions; the decision variables are the depth and frequency of promotions. In contrast, we endogenously derive the optimal temporal price structure and show that similar promotion patterns will arise in equilibrium.

In another paper, Jeuland and Narasimhan (1985) also assume that the seller offers periodic promotions. They also consider two groups of consumers with high or low holding cost, but they assume that high-cost consumers face a higher demand curve. Since high-cost consumers do not stockpile, offering promotions effective separately these two groups across time. This illustrates that periodic promotions serve as a price discrimination device. In our paper, we extend these insights by considering two additional dimensions of consumer heterogeneity. Apart from different holding costs, consumers may also have different consumption rates and different fixed shopping costs. We show that promotions are an effective price discrimination device when consumers with high holding cost, low fixed costs or high consumption rates also have higher valuations.

Several other papers study the competitive pricing strategies of multiple firms facing consumers who stockpile. For simplicity, many of these papers assume that consumers may hold at most one unit of inventory. Salop and Stiglitz (1982) show that consumer stockpiling generates price dispersion in a mixed strategy equilibrium. During each period, some firms offer discounts to generate additional sales to consumers who stockpile, while others keep prices high. The equilibrium price distributions involve only two different prices and are i.i.d. across periods. Bell, Iyer, and Padmanabhan (2002) incorporate flexible consumption into the Salop and Stiglitz model and obtain similar equilibrium results. Hong, McAfee, and Nayyar (2002) extend the models of Varian (1980) and Narasimhan (1988) by allowing consumers to carry inventory. They show that equilibrium prices depend on consumer inventories and thus display negative serial correlation. In a vertical channel setup, Lal, Little and Villas-Boas (1996) analyze forward buying at the retailer level. There are two manufacturers, one retailer, and “loyals and switchers” at the consumer level. The authors find that allowing the retailer to forward buy decreases price competition between manufacturers. However, Anton and Das Varma (2005) obtain different conclusions when they study forward buying at the consumer level. Using a two-period duopoly model, they find that strategic stockpiling intensifies competition as firms attempt to capture future market share from rivals. Guo and Villas-Boas (2007) consider product differentiation using a two-period Hotelling setup. In their model, high-valuation consumers have a higher propensity to stockpile. Hence, stockpiling leads to a lower degree of differentiation and more intense price competition in future periods. In the papers described above, equilibrium prices are either i.i.d. across time or first-order Markov in a binary state space (since consumer inventory is either zero or one). We add to this literature by exploring a richer state space and hence a wider class of pricing strategies.

3 Model

There is a monopolist seller and a mass of consumers in the market. These players interact at discrete time periods over an infinite horizon. In each period, the seller sets a price for the product, and each consumer chooses the number of units to purchase. Units that are not consumed immediately may be inventoried for future use. However, there is an inventory holding cost. For the seller, production cost is normalized to zero without loss of generality. In this environment, the seller faces a infinite-horizon dynamic pricing problem, and individual consumers face a infinite-horizon inventory problem.

In our model, we allow for n different consumer types. Let $\theta \in \Theta \equiv \{1, \dots, n\}$ denote consumer type. Each consumer type is parameterized by four values, $R_\theta, v_\theta, K_\theta, h_\theta$, which we describe below. The *consumption rate* R_θ denotes the maximum number of units that may be consumed in each period. The *valuation* v_θ is the positive utility derived from each unit that is consumed. The *fixed cost* K_θ is incurred whenever the consumer makes a purchase. The *holding cost* h_θ is the cost of carrying one unit of the product over one unit of time. Let π_θ denote the proportion of type- θ consumers, so these proportions sum to one over all n types.

Consumer utility depends on three components: purchasing cost, valuation, and holding cost. Consider a type- θ consumer. First, we consider purchasing cost. Let $D_{\theta t}$ denote the number of units purchased at price p_t in period t by the type- θ consumer. Then, the total purchasing cost is $K_\theta \cdot 1_{\{D_{\theta t} > 0\}} + p_t D_{\theta t}$. We assume that all units purchased enter the consumer's inventory at the start of the period and thus can be consumed within the same period. Next, we consider consumption. Let $E_{\theta t} \in [0, R_\theta]$ denote the number of units consumed in period t . Then, this consumer derives positive consumption utility of $v_\theta \cdot E_{\theta t}$. We assume that consumption is spread evenly within the period. Finally, we consider holding cost. Let $I_{\theta t}$ denote the number of units held by the type- θ consumer at the end of period t . Then, the total holding cost for this period is $h_\theta \cdot (I_{\theta t} + E_{\theta t}/2)$. Note that the last term arises due to uniform consumption within the period, so on average, holding cost is incurred on half the units that were consumed. In summary, the type- θ consumer's utility accrued in period t is given by

$$U_{\theta t} = v_\theta \cdot E_{\theta t} - K_\theta \cdot 1_{\{D_{\theta t} > 0\}} - p_t \cdot D_{\theta t} - h_\theta \cdot (I_{\theta t} + E_{\theta t}/2). \quad (1)$$

In each time period t , the following chronology of events occurs.

1. First, the seller sets prices p_t .

2. Then, consumers make purchase decisions. Each type- θ consumer buys $D_{\theta t}$ units and the total demand is $D_t \equiv \sum_{\theta=1}^n \pi_{\theta} D_{\theta t}$.
3. Consumption occurs. Each type- θ consumer chooses the number of units $E_{\theta t}$ to consume. Note that the optimal consumption decisions $E_{\theta t}^* = \min\{R_{\theta}, I_{\theta, t-1} + D_{\theta t}\}$ are easy to obtain. Since consumers incur holding costs on leftover units, they prefer immediate consumption over delayed consumption. Hence, they will always consume whatever is available up to R_{θ} units.
4. Consumer inventory is updated by adding units purchased and subtracting units consumed, according to $I_{\theta t} = I_{\theta, t-1} + D_{\theta t} - E_{\theta t}$. By the same reasoning above, we have $I_{\theta t} = [I_{\theta, t-1} + D_{\theta t} - R_{\theta}]^+$.
5. The seller accrues profits in this period, as given by $\Pi_t = p_t D_t$.
6. Each consumer accrues utility $U_{\theta t}$ in this period according to (1) above.

Therefore, over the infinite horizon $t = 1, 2, \dots$, the seller makes pricing decisions and consumers make purchase and inventory decisions. We assume that the seller wishes to maximize long-run average profit $\Pi \equiv \lim_{T \rightarrow \infty} \frac{\sum_{t=1}^T \Pi_t}{T}$ and individual consumers seek to maximize long-run average utility $U_{\theta} \equiv \lim_{T \rightarrow \infty} \frac{\sum_{t=1}^T U_{\theta t}}{T}$. Although a similar analysis can be conducted using the discounted profit/utility criteria, we choose to use the long-run average profit/utility criteria in order to facilitate comparisons with the classical EOQ model, which is based on average cost.

4 General Formulation and Analysis

In this section, we formulate the seller's pricing problem, consumers' purchasing problem, and define our equilibrium concept for this dynamic game.

4.1 Seller's Problem

The seller's problem is to determine the optimal price in each period t , which depends on consumers' on-hand inventory. Let $\mathbf{I}_t = \{I_{\theta, t}\}_{\theta \in \Theta}$ denote the vector of on-hand inventories kept by each consumer type at the end of period t . Then, the optimal pricing strategy is a function $P^*(\mathbf{I})$ that yields the optimal price $p_t^* = P^*(\mathbf{I}_{t-1})$ in each period t .

We begin by deriving the seller's optimal pricing strategy, given his beliefs about consumer behavior. Suppose the seller believes that the type- θ consumer's demand function is $\hat{D}_\theta(\mathbf{I}, p)$, and thus the aggregate demand function is $\hat{D}(\mathbf{I}, p) \equiv \sum_{\theta=1}^n \pi_\theta \hat{D}_\theta(\mathbf{I}, p)$. In other words, during period t , when the vector of consumer inventories carried from the previous period is \mathbf{I}_{t-1} and when the price is p_t , the seller believes that type- θ demand will be $\hat{D}_{\theta t} = \hat{D}_\theta(\mathbf{I}_{t-1}, p_t)$ and that aggregate demand will be $\hat{D}_t = \hat{D}(\mathbf{I}_{t-1}, p_t) = \sum_{\theta} \pi_\theta \hat{D}_\theta(\mathbf{I}_{t-1}, p_t)$. These beliefs $\hat{D}_\theta(\mathbf{I}, p)$ are useful analytically as they allow us to isolate the seller's decisions and formulate the pricing problem as a dynamic program.

The seller's goal is to find the pricing policy $P(\mathbf{I})$ that maximizes the long-run average profit $\Pi \equiv \lim_{T \rightarrow \infty} \frac{\sum_{t=1}^T \Pi_t}{T}$. This is an average-reward-criterion dynamic programming problem, which can be formulated as follows: Find $\lambda, g(\mathbf{I})$ satisfying the Bellman equations

$$\lambda + g(\mathbf{I}) = \max_p \{p \hat{D}(\mathbf{I}, p) + g(\mathbf{I}'(\mathbf{I}, p))\} \quad \forall \mathbf{I}. \quad (2)$$

Here, $\mathbf{I}'(\cdot, \cdot)$ is the state transition function. In other words, when the previous inventory state is \mathbf{I}_{t-1} and the price p_t is chosen, the seller anticipates that the state will transition to $\mathbf{I}_t = \mathbf{I}'(\mathbf{I}_{t-1}, p_t)$ based on the beliefs about consumer demand $\hat{D}(\mathbf{I}, p)$ given above. Specifically, considering each type separately, we have $I'_\theta(\mathbf{I}, p) = [I_\theta + \hat{D}_\theta(\mathbf{I}, p) - R_\theta]^+$. For any pricing policy $P(\mathbf{I})$, we may interpret λ as the seller's long-run average revenue under this policy and $g(\mathbf{I})$ as a gain function, which captures short-term fluctuations from the long-run average when currently at state \mathbf{I} . In particular, at the solution $(\lambda^*, g^*(\mathbf{I}))$ to (2), we may interpret λ^* as the seller's optimal long-run average profit. Further, at each state \mathbf{I} , the maximizer p^* in the Bellman equation (2) is the optimal price at that state, so these maximizers collectively yield the desired pricing strategy $P^*(\mathbf{I})$.

In summary, given beliefs $\hat{D}(\mathbf{I}, p)$ about how demand will respond to price p in each inventory state \mathbf{I} , the seller's optimal pricing policy $P^*(\mathbf{I})$ can be obtained by solving the dynamic program (2).

4.2 Consumer's Problem

The consumer's problem in each period t is to determine how many units to purchase $D_{\theta,t}$, based on price p_t , on-hand inventory $I_{\theta,t-1}$, as well as inferences over others' inventories $\hat{I}_{-\theta,t-1}$ (which may affect future price). We assume that consumers form inferences over the inventories of other types, based on their beliefs of others' purchasing behavior (as described below). Hence we may express the consumer's optimal purchasing strategy $D_\theta^*(\mathbf{I}, p)$ as a function of inventories \mathbf{I} and the price p . In

the consumer's problem, we use $\mathbf{I} = (I_\theta, \hat{I}_{-\theta})$ to denote the vector of own inventory and inferences of others' inventories. The optimal purchasing strategy yields the optimal quantity $D_{\theta,t}^* = D_\theta^*(\mathbf{I}_{t-1}, p_t)$ in each period t .

As above, we begin by endowing the consumer with beliefs. Suppose all consumers believe that the seller sets price according to the function $\hat{P}(\mathbf{I})$. In addition, suppose that type- θ consumers are believed (by other types) to purchase according to the function $\hat{D}_\theta(\mathbf{I}, p)$. In other words, at the end of period t , all consumers anticipate the next period's price to be $\hat{p}_{t+1} = \hat{P}(\mathbf{I}_t)$ and that given this price, type- θ consumers will buy $\hat{D}_{\theta,t+1} = \hat{D}_\theta(\mathbf{I}_t, \hat{p}_{t+1})$ units in the next period. Therefore, all consumers share the same beliefs $\hat{P}(\mathbf{I})$ over the seller's actions as well as the same beliefs $\hat{D}_\theta(\mathbf{I}, p)$ and \hat{I}_θ over the actions and inventories of other types θ .

Given the beliefs above, the optimal purchasing policy $D_\theta^*(\mathbf{I}, p)$ that maximizes the type- θ consumer's long-run average utility $U_\theta \equiv \lim_{T \rightarrow \infty} \frac{\sum_{t=1}^T U_{\theta,t}}{T}$ can be obtained by solving an average-reward dynamic program. The formulation is as follows: Find $\mu_\theta, f_\theta(\mathbf{I}, p)$ satisfying the Bellman equations

$$\mu_\theta + f_\theta(\mathbf{I}, p) = \max_{D_\theta} \{U_\theta(I_\theta, p, D_\theta) + f_\theta(\mathbf{I}'(\mathbf{I}, p, D_\theta), P'(\mathbf{I}, p, D_\theta))\} \quad \forall \mathbf{I}, p. \quad (3)$$

Here, similar to (1),

$$U_\theta(I_\theta, p, D_\theta) = v_\theta \cdot E_\theta(I_\theta, D_\theta) - K_\theta \cdot 1_{\{D_\theta > 0\}} - p \cdot D_\theta - h_\theta \cdot (I'_\theta(I_\theta, p, D_\theta) + E_\theta(I_\theta, D_\theta)/2) \quad (4)$$

is the per-period utility function, where $E_\theta(I_\theta, D_\theta) = \min\{R_\theta, I_\theta + D_\theta\}$ is the quantity consumed. Next, $\mathbf{I}'(\cdot, \cdot, \cdot)$ and $P'(\cdot, \cdot, \cdot)$ are the state transition functions. In particular, for $\mathbf{I}'(\cdot, \cdot, \cdot)$, we have $I'_\theta(\mathbf{I}, p, D_\theta) = [I_\theta + D_\theta - R_\theta]^+$ for type θ and $I'_{-\theta}(\mathbf{I}, p, D_\theta) = [\hat{I}_{-\theta} + \hat{D}_{-\theta}(\mathbf{I}, p) - R_{-\theta}]^+$ for other types. Next, for the other transition function $P'(\cdot, \cdot, \cdot)$, we have $P'(\mathbf{I}, p, D_\theta) = \hat{P}(\mathbf{I}'(\mathbf{I}, p, D_\theta))$. At the solution $(\mu^*, f_\theta^*(\mathbf{I}, p))$ to (3), we may interpret μ^* as the consumer's optimal long-run average utility and $f_\theta^*(\mathbf{I}, p)$ as a gain function, which represents short-term fluctuations from the optimal long-run average when currently facing inventories \mathbf{I} and price p . At each state (\mathbf{I}, p) , the maximizer D_θ^* in the Bellman equation (3) is the optimal purchase quantity at that state, so these maximizers collectively yield the desired purchasing strategy $D_\theta^*(\mathbf{I}, p)$.

Here, we have formulated a dynamic program that yields the type- θ consumer's optimal purchasing policy $D_\theta^*(\mathbf{I}, p)$. This policy is optimal under common beliefs over the seller's pricing policy $\hat{P}(\mathbf{I})$ as well as common beliefs over other types' purchasing policies $\hat{D}_{-\theta}(\mathbf{I}, p)$.

4.3 Rational Expectations Equilibrium

We seek a rational expectations (RE) equilibrium in the game between the seller and the consumers. The basic requirements of our equilibrium concept is that each player (seller or consumer) forms beliefs over others' strategies and optimizes individual payoffs given these beliefs (as described in the previous subsections), and that these beliefs must be accurate. The definition of our RE equilibrium is provided next.

Definition 1 *A rational expectations (RE) equilibrium consists of a pricing policy $P^*(\mathbf{I})$ and purchasing policies $D_\theta^*(\mathbf{I}, p)$ satisfying the following three conditions.*

1. **(Seller optimality)** *Given beliefs over consumer demand $\hat{D}_\theta(\mathbf{I}, p)$ for each type θ , the seller's pricing policy $P^*(\mathbf{I})$ maximizes long-run average profits Π .*
2. **(Consumer optimality)** *Given beliefs over the seller's pricing policy $\hat{P}(\mathbf{I})$ and other types' purchasing policies $\hat{D}_{-\theta}(\mathbf{I}, p)$, each type- θ consumer's purchasing policy $D_\theta^*(\mathbf{I}, p)$ maximizes long-run average utility U_θ .*
3. **(Dynamic consistency)** *All beliefs are consistent with outcomes in all states and over all time. Specifically, $\hat{P}(\mathbf{I}) = P^*(\mathbf{I})$, $\hat{D}(\mathbf{I}, p) = \sum_{\theta=1}^n D_\theta^*(\mathbf{I}, p)$, $\hat{D}_\theta(\mathbf{I}, p) = D_\theta^*(\mathbf{I}, p)$, and $\hat{I}_{\theta t} = I_{\theta t}$.*

Before proceeding, we state two technical assumptions. First, we assume that the seller's set of allowable prices is bounded. This is not restrictive since the upper bound can be set as the highest per-unit consumer valuation without loss of optimality. Second, we assume that consumers may stock up for at most τ future periods. Again, this assumption is not restrictive since consumers incur holding costs; then, it is never optimal for the consumer to stock beyond some finite time point in the future. This assumption can also be interpreted as an expiration date for the product (e.g., perishable items) or a physical holding capacity (e.g., for bulky consumer products).

The following theorem guarantees the existence of a RE equilibrium in our model.

Theorem 1 *There exists an RE equilibrium $\{P^*(\mathbf{I}); D_\theta^*(\mathbf{I}, p)\}$ satisfying seller optimality, consumer optimality, and dynamic consistency.*

The proof appears in the appendix. The key idea is to consider a *static* version of the game between the seller and the consumers, in which the seller makes a *one-shot choice* of a pricing policy $P(\mathbf{I})$ and the

consumers also makes a one-shot choice of purchasing policies $D_\theta(\mathbf{I}, p)$, and the corresponding payoff functions are the long-run average payoffs Π and U_θ . Observe that a pure-strategy Nash equilibrium in this static game satisfies the conditions for our RE equilibrium defined above. Thus, the proof proceeds by establishing the existence of a pure-strategy equilibrium in this static game.

There is another useful structural property of the RE equilibrium. The next theorem indicates that in any RE equilibrium outcome, there will be price cycles.

Theorem 2

- (i) *Given any beliefs over consumer demand $\hat{D}_\theta(\mathbf{I}, p)$, the seller's optimal pricing policy $P^*(\mathbf{I})$ follows a cyclic pattern. That is, there are prices $\{\tilde{p}_i\}_{i=1}^m$ such that the seller chooses price $p_t = \tilde{p}_{t \bmod m}$ for every t .*
- (ii) *Given any beliefs over the seller's pricing policy $\hat{P}(\mathbf{I})$ and other types' purchasing policies $\hat{D}_{-\theta}(\mathbf{I}, p)$, each type- θ consumer's optimal purchasing policy $D_\theta^*(\mathbf{I}, p)$ generates inventory levels that follow a cyclic pattern. That is, there are inventory levels $\{\tilde{\mathbf{I}}_i\}_{i=1}^m$ such that the inventory level is $\mathbf{I}_t = \tilde{\mathbf{I}}_{t \bmod m}$ for every t .*
- (iii) *In particular, in any RE equilibrium, the seller's prices and the consumers' inventory levels follow a cyclic pattern.*

The proof is presented in the appendix and is based on a graphical argument. Here, we sketch the basic ideas. The seller's and consumers' problems described in Sections 4.1 and 4.2, which are dynamic programs, can be formulated as linear programs. We show that the dual solutions of these linear programs have a useful interpretation. They correspond to the cyclical trajectory of states that the system goes through (over time) as the seller sets prices optimally and consumers make purchases optimally in each period. Consequently, equilibrium outcomes must follow a periodic pattern. As we traverse the cycle of states, all outcomes (prices and inventory) fluctuate in a *predictable* manner that repeats itself over each cycle. The reader is referred to Appendix A for more details.

This result also provides a useful equivalence between dynamic pricing policies $P(\mathbf{I})$ and price cycles. With the latter interpretation, we may directly calculate the seller's long-run average revenue and the consumers' long-run average utility corresponding to any arbitrary cycle of prices and inventories. Although solving dynamic programming problems is computationally feasible, in some cases,

it may be analytically more convenient to look for equilibria over price cycles. This approach will be useful later.

5 Computational Study

The numerical study in this section proceeds as follows. As a starting point, we randomly generate a variety of hypothetical model scenarios. For each scenario, we computationally characterize the RE equilibrium using the theoretical results developed in the preceding section. With these results, we provide insights on how the seller’s optimal pricing strategies may be driven by model input parameters. Finally, we demonstrate that it is important to recognize consumer stockpiling by contrasting our computational results to a benchmark in which the seller neglects such behavior.

First, we describe our sample selection procedure. We simulate a total of 2000 model scenarios. In each model scenario, the seller sells to a market consisting of 4 equal-sized consumer segments. As described in Section 3, the market size is normalized to 1 and each consumer segment $i = 1, \dots, 4$ is fully characterized by the parameters v_i, R_i, h_i, K_i . The valuation v_i is drawn from a normal distribution with mean \$5.00 and standard deviation \$1.00. The per-period consumption rate R_i is either 1 unit (with probability 0.75) or 2 units (with probability 0.25). The per-period, per-unit holding cost h_i and the fixed cost K_i are both exponentially distributed with means \$1.50 and \$3.00 respectively. All random realizations above are independent across consumer segments and across scenarios. To simplify the state and action spaces, we assume that the seller charges integral unit prices and consumers purchase units in integral quantities.

Next, we describe how we compute the RE equilibrium for each model scenario. As explained in Section 4.3, the RE equilibrium can be viewed as a pure-strategy Nash equilibrium in the space of dynamic policies. That is, each player’s equilibrium strategy must be a best response to the dynamic policies adopted by all other players. We compute the equilibrium using the alternating-move Cournot adjustment process, which is an iterative algorithm used to compute Nash equilibria in games (see Fudenberg and Tirole, 1991). At each step in this process, one of the players (i.e., the seller or consumer segments 1, 2, 3, or 4, in turn) updates his current strategy to the optimal policy in response to current choices of other players (i.e, current beliefs). Each updating step requires us to solve either the dynamic program (2) for the seller or (3) for the consumer; we employ the value iteration algorithm here (see

Puterman, 1994). The computational procedure described above does not converge in 108 (i.e., 5.4%) of our simulated scenarios. Therefore, in our analysis below, we are left with a sample size of $N = 1892$ scenarios.

For each scenario, we are interested in the following quantities as they shed some light on the structure of the seller’s optimal pricing policy: (i) the seller’s average price in each period, (ii) the standard deviation of prices over each price cycle, (iii) the number of different prices offered, (iv) the price range, (v) the length of each price cycle, and (vi) the length of each cycle of inventory states. Note that when the seller charges a single fixed price, the standard deviation and price range are both zero, and each price “cycle” has length 1. However, the inventory-state cycle is generally longer than the price cycle because the former may have length exceeding 1 as a result of consumer stockpiling, even when the seller charges a fixed price. Finally, for each simulated scenario, we are also interested in the seller’s average profit per period. Table 1 provides summary statistics of the above quantities of interest, displaying the mean, standard deviation, and range of each quantity over our sample of observations.

	Mean	Std Dev	Min	Max
Average Price	2.4863	0.9361	1	5
Standard Deviation of Prices	0.1462	0.3237	0	2.3094
Number of Different Prices	1.2220	0.4487	1	4
Price Range	0.2600	0.5689	0	4
Length of Price Cycle	1.5872	1.5444	1	12
Length of Inventory States Cycle	3.2865	2.5763	1	12
Average Per Period Profit	1.9751	0.9950	0	7

Table 1: Summary statistics for our sample of observations.

What are some key market characteristics that shape the structure of the seller’s optimal pricing strategy? For example, under what market conditions should the seller charge higher average prices or more variable prices? When should the seller charge a fixed price? What factors influence the price range and the length of the price cycle? We shall address these questions using our computational results. In our sample, there is a large number of variables (i.e., 4 simulated parameters, v_i, R_i, h_i, K_i , for each of the 4 consumer segments). From this pool of variables, we identify potential predictors that might be of practical significance. For example, for each consumer attribute (e.g., the vector of fixed costs (K_1, K_2, K_3, K_4) for the 4 consumer segments), we consider the mean, standard deviation, and

range as possible predictors. Another set of potential predictors are correlations between each pair of consumer attributes. Having identified these potential predictors, we proceed to look for systematic relationships between them and the quantities of interest described in the preceding paragraph. Treating each quantity of interest (e.g., length of each price cycle) as a separate response variable, we run a simple linear regression with our predictors above. (For binary response variables, we fit a logistic regression model instead.) The results are summarized in Table 2. Each column represents a separate regression, and significant coefficients suggest a systematic relationship.

	Average Profit	Average Price	Std Dev Price	Number of Different Prices	Pricing Format (0=fixed, 1=variable)	Price Range	Length of Price Cycle
Intercept	-2.6271 **	-0.5413 *	-0.0078	1.0914 **	-1.7383 *	0.1513	5.6238 **
Valuations, v							
Mean	0.8531 **	0.7600 **	-0.0065	0.0087	0.1805	-0.2222	-0.3209
Stdev	-0.2268	0.1510	0.2119	0.1758	0.7094	0.9444 *	-1.6052
Max	-0.1191	-0.1379	0.0774	0.0497	0.0988	0.4713 *	-0.5040
Min	0.0291	-0.0168	-0.0475	-0.0295	-0.1431	-0.1402	1.0430
Consumption rates, R							
Mean	1.9656 **	0.2048	-0.0284	-0.0758	-0.6051	0.1321	-1.5588
Stdev	0.1164	0.2268 *	-0.0977 *	-0.1603 *	-0.9701 *	0.1519	-0.7216
Holding costs, h							
Mean	-0.7654 **	-0.1407 *	0.0106	-0.0013	-0.0297	0.0843	-1.0817 **
Stdev	0.1341	0.0472	-0.0214	-0.0670	-0.2008	-0.1558	1.6476
Max	0.0643	-0.1959 *	-0.0247	-0.0508	-0.2014	-0.1581	-0.0388
Min	0.0683	0.0179	0.0095	0.0364	0.1210	0.0675	-0.4840
Fixed costs, K							
Mean	-0.3457 **	-0.0975 **	0.0517 **	0.0803 **	0.4054 **	0.0951	0.4004 *
Stdev	0.2655 **	0.2525 **	-0.1033 *	-0.0734	-0.5839	-0.2694	0.3423
Max	0.0146	-0.2946 **	-0.0951 **	-0.0814 *	-0.4925 **	-0.3116 **	0.4092
Min	-0.0620	-0.0785	0.0334	0.0090	0.1477	0.0984	-0.3275
Correlations							
v and R	0.2598 **	0.3222 **	0.0163	0.0091	0.0533	0.1103	-0.0290
v and h	-0.0519 **	-0.1851 **	0.0098	0.0150	0.0555	-0.0466	-0.3187
v and K	-0.0881 **	-0.3920 **	-0.0671 **	-0.0789 **	-0.4478 **	-0.1290 *	0.8657 **
R and h	-0.3649 **	-0.1386 **	0.0649 **	0.1117 **	0.7264 **	0.0085	0.5581 *
R and K	-0.2477 **	-0.2983 **	-0.0010	-0.0148	-0.1285	0.0692	-0.1708
h and K	0.0318	0.2107 **	-0.0277 *	-0.0587 **	-0.3140 **	0.0261	-0.5951 **

Table 2: Regression results. Each column displays the estimated coefficients of a separate regression model. The dependent variable is shown at the top of the column, and independent variables are displayed on the leftmost column. Statistical significance at the 5% and 1% levels are respectively marked by * and **.

We highlight three key observations.

1. The first two columns of Table 2 identify the main factors affecting the seller's profit and average prices. Our results show that profits and prices generally depend on the means of all four consumer attributes. Specifically, prices and profits are higher as average valuations increase

(since the coefficients 0.8531 and 0.76 are positive and significant), as average consumption rates increase, as average holding costs decrease, and as average fixed costs decrease. These effects match our intuition. Furthermore, note that prices and profits also depend heavily on each pairwise correlation between consumer attributes.

2. The next three columns study various measures of price variation: standard deviation of prices, number of unique prices offered, and whether the seller prefers variable over fixed pricing. Our results suggest that consumer fixed costs have a large role to play. In our sample, optimal prices appear to be more variable as average fixed costs increase (note that the coefficients 0.0517, 0.0803, 0.4054 are all highly significant and positive). In addition, our results also show that variable pricing becomes more attractive as consumer valuations and fixed costs become more negatively correlated (coefficient of -0.4478), as consumption rates and holding costs become more positively correlated (coefficient of 0.7264), and as holding costs and fixed costs become more negatively correlated (coefficient of -0.314).
3. The final two columns present regression results using a restricted sample consisting only of observations in which the seller uses variable pricing. We are interested in asking: when the seller varies prices over time, what factors affect the depth of discounts (reflected in the price range) and the time intervals between them (i.e., the price cycle length)? Our results suggest that these are more elusive questions. Compared to before, these regression models appear to have a smaller number of significant coefficients. The only predictor that is significant for both columns is the correlation between consumer valuations and fixed costs. As this correlation becomes more positive, our results suggest that the seller should offer smaller discounts (since $-0.129 < 0$) that are spaced between larger time intervals (since $0.8657 > 0$).

For the remainder of our numerical study, we estimate the benefits of using our model by comparing results against a benchmark. In this benchmark, the seller ignores consumer stockpiling and assumes that consumers make purchases every period as long as the price is sufficiently attractive. Thus, the seller's problem reduces to a standard monopoly pricing problem: the lower the price, the more consumers will buy. For each sample observation, given the randomly generated consumer attributes, we calculate the seller's preferred price and his corresponding profit when he ignores stockpiling. Since the seller fails to consider dynamic pricing, his profit will be lower in our benchmark

calculations. For each of our sample observations, we compute the percentage profit loss that the seller suffers when he neglects consumer stockpiling. This percentage difference can be interpreted as the benefit of using our model.

It is instructive to divide our sample observations into three categories and consider the benefits of our model for each of them. In the first category, the RE equilibrium involves a fixed price and no consumers stockpile. Observe that in these instances, the seller loses nothing by ignoring consumer stockpiling since it does not arise in equilibrium anyway. In the second category, the RE equilibrium involves a fixed price but some consumers do stockpile. For such cases, stockpiling allows consumers to exploit scale economies and reduce costs. When the seller fails to recognize this fact, he may underestimate willingness-to-pay and end up with a suboptimal price. (Certainly, there may also be scenarios in which the seller correctly chooses the optimal price even when he neglects consumer stockpiling.) Finally, in our third category of model scenarios, the seller’s prices vary dynamically and some consumers stockpile. In these cases, our model is used to derive the optimal price cycle, so the seller may suffer some losses when he ignores consumer stockpiling.

Table 3 presents statistics summarizing the seller’s percentage profit loss as a result of neglecting consumer stockpiling. The results are shown for each of the three category of observations and also for the entire sample. Note that each of the three categories, i.e., fixed-pricing-without-stockpiling, fixed-pricing-with-stockpiling, and variable-pricing-with-stockpiling, arises with frequencies 22%, 57%, and 21% respectively. For each category of observations (on separate rows), we show the average percentage loss as well as the standard deviation and the median. We also present these statistics for the subset of observations in which adopting our model yields a positive gain (i.e., the seller’s profit is strictly lower when he ignores consumer stockpiling).

		All observations in each category			Conditional on observations where our model performs strictly better than the benchmark		
		Frequency	Mean	Std Dev	Median	Mean	Std Dev
Fixed pricing without stockpiling	21.93	0	0	0	0	0	0
Fixed pricing with consumer stockpiling	57.19	14.51	21.30	0	35.67	18.97	33.33
Variable pricing with consumer stockpiling	20.88	15.14	16.09	10.00	20.48	15.52	14.96
Entire sample	100	11.46	18.71	0	29.61	19.17	25.00

Table 3: Seller’s percentage profit loss as a result of neglecting consumer stockpiling.

We draw three conclusions from these results.

1. In our sample, more than half of the observations involve a fixed price and consumer stockpiling. For these cases, the median profit loss from ignoring consumer stockpiling is zero, implying that very often, the seller can correctly obtain the optimal price without the aid of our model. Nevertheless, when our model does matter, the seller’s percentage loss averages 35.67% and exceeds one-third about half the time. These are rather significant proportions.
2. Next, consider the subset of observations in which the optimal pricing policy is cyclic and consumers stockpile. In these cases, when the seller neglects consumer stockpiling, the percentage losses have a mean of 15.14% and a standard deviation of 16.09%. Also, the seller’s loss exceeds 10% about half the time.
3. Finally, considering the entire sample of observations, our results show that the seller who ignores consumer stockpiling may lose about 11.46% of his profits (on average). Note that the standard deviation of 18.71% is rather high because there is a sizable mass point at 0 (since our model, admittedly, makes no difference for a significant proportion of the observations, e.g., when there is no stockpiling in equilibrium).

6 Practical Implications

In order to derive additional insights, we now consider a special case with two types of consumers. This special case of our model can be solved analytically. We characterize the RE equilibrium explicitly below. Then, we discuss practical implications of our results. Specifically, we: (i) discuss factors influencing the optimal choice of pricing formats, (ii) study how the frequency and depth of discounts depend on model parameters, and (iii) describe how equilibrium sales and inventory levels evolve over time.

We begin with some notation and terminology for the two-type special case. For each consumer type θ , it follows from the standard EOQ model that under constant prices, the consumer wishes to make purchases every T_θ periods, where $T_\theta = \sqrt{\frac{2K_\theta}{h_\theta R_\theta}}$. We can interpret T_θ as the *natural shopping frequency*, which minimizes holding and fixed costs, when prices are constant. This motivates the following terminology for our two-type model. We shall refer to our two types as *frequent shoppers*

and *occasional* shoppers, where the former shops more frequently under constant prices, i.e. $T_f = \sqrt{\frac{2K_f}{h_f R_f}} < \sqrt{\frac{2K_o}{h_o R_o}} = T_o$. Without loss of generality, we normalize time units so that $T_f \equiv 1$. In other words, the length of one time period in our discrete-time model can be interpreted as the natural shopping frequency of frequent shoppers; under constant prices, they visit the seller every period. Similarly, we assume that the natural shopping frequency of occasional shoppers is a positive integer, i.e. $T_o \in \mathbb{Z}^+$. In other words, while frequent shoppers prefer to shop every period, occasional shoppers prefer to shop every T_o periods, given constant prices.

Next, we characterize the willingness-to-pay of each consumer type. For each consumer type $\theta \in \{f, o\}$, recall from the EOQ model that the optimal (lowest) shopping cost is $\sqrt{2h_\theta K_\theta R_\theta}$ per unit time, which is equivalent to $\sqrt{\frac{2h_\theta K_\theta}{R_\theta}}$ per unit. (Note that this lowest cost is achieved only at the optimal shopping frequency.) Since per-unit valuation is v_θ , each type- θ consumer will never pay more than $v_\theta - \sqrt{\frac{2h_\theta K_\theta}{R_\theta}}$ for each unit of the product. We introduce the notation

$$\begin{aligned} WTP_f &\equiv v_f - \sqrt{\frac{2h_f K_f}{R_f}}, \\ WTP_o &\equiv v_o - \sqrt{\frac{2h_o K_o}{R_o}}, \end{aligned}$$

for the willingness-to-pay of frequent and occasional shoppers, respectively. Further, we denote $WTP_L \equiv \min\{WTP_f, WTP_o\}$, $WTP_H \equiv \max\{WTP_f, WTP_o\}$, and $\Delta \equiv WTP_H - WTP_L$.

Now, we are ready to characterize the RE equilibrium outcomes. Depending on model parameters, there are three possible types of equilibrium. We call them everyday-low-price (EDLP), everyday-high-price (EDHP), and promotional pricing (HILO). Under EDLP and EDHP, the seller sets a constant price at WTP_L and WTP_H respectively. Under HILO, the regular price is WTP_H and a sale price of WTP_L is offered periodically (i.e., the discount is Δ). The following proposition characterizes conditions under which each type of equilibrium arises.

Proposition 1 *(i) Suppose $WTP_f > WTP_o$. Then, only HILO or EDHP may arise in equilibrium.*

- *Under HILO, the low price WTP_o is offered once every T_o periods and the high price WTP_f is offered in other periods. Occasional shoppers buy only on sale, i.e., once every T_o periods. Frequent shoppers stock up for $\min\{\Delta/h_f, (T_o - 1)\}$ future periods when there is a sale and after this inventory runs out, they shop every period until the next sale.*
- *Under EDHP, the seller charges WTP_f every period and only frequent shoppers buy.*

(ii) Suppose $WTP_o > WTP_f$. Then, only EDLP or EDHP may arise in equilibrium.

- Under EDLP, the seller charges WTP_f every period. Frequent shoppers buy every period and occasional shoppers buy once every T_o periods.
- Under EDHP, the seller charges WTP_o every period. Only occasional shoppers buy and they shop once every T_o periods.

The proof is provided in the appendix. In the next few subsections, we shall proceed to discuss the practical implications of our equilibrium results in Proposition 1.

6.1 Optimal Pricing Format

We first discuss the seller’s choice of an optimal pricing format. In practice, periodic promotions (HILO) and constant prices (EDLP or EDHP) are both viable pricing formats. Within the boundary of our model, our results may offer some insight on when each pricing format is preferred. In fact, Proposition 1 clearly states that price promotions should be used only when frequent shoppers have a higher willingness-to-pay. Otherwise, it is optimal for the seller to offer fixed prices.

Intuitively, the “frequent shoppers spend more” condition implies that periodic promotions are an effective price discrimination device. By definition, occasional shoppers shop less frequently and thus they are likely to stock up for longer time horizons during promotions. On the other hand, frequent shoppers continue to visit the store even when there is no promotion. When frequent shoppers have a higher willingness-to-pay, the seller can then extract their entire surplus during these off-promotion visits, without losing sales to occasional shoppers (who have stocked up during promotions). Therefore, strategic stockpiling by occasional shoppers enables the seller to price discriminate and charge a higher off-promotion price to frequent shoppers.

When is this “frequent shoppers spend more” condition more likely to hold? Clearly, as frequent shoppers’ valuation v_f increases, this condition is more likely to hold. Recall that frequent shoppers are so defined because their optimal shopping interval $T_f = \sqrt{\frac{2K_f}{h_f R_f}}$ is short. In practice, consumers with high holding costs h , high consumption rates R , and low fixed costs K are likely to be frequent shoppers. If these consumers also have high valuations, the “frequent shoppers spend more” condition is more likely to hold, and price promotions thus become more attractive for the seller. We state this result as a corollary to Proposition 1.

Corollary 1 *Suppose the seller wishes to sell to both consumer segments. Then, all else equal, promotional pricing is the optimal pricing format when:*

- (i) consumers with higher holding costs have sufficiently high valuations,*
- (ii) consumers with higher consumption rates have sufficiently high valuations,*
- (iii) consumers with lower fixed costs have sufficiently high valuations.*

In practice, our model parameters h, K, R are closely related to demographic factors. For many household products (e.g., detergent), consumption rates R depend heavily on household size; hence, Corollary 1 suggests that promotions are preferred when larger households (with higher R) have higher valuations. In fact, even when all consumers have the same valuation (i.e., $v_f = v_o$), heterogeneity in consumption rates alone (i.e., $R_f > R_o$ but $K_f = K_o, h_f = h_o$) already implies that price promotions are optimal (since it follows that frequent shoppers have a higher willingness to pay). Next, note that holding costs h are influenced by income (when we consider financial holding cost) as well as house size (a larger house stores more products), especially for relatively heavy or bulky items (e.g., bottled water). In the case of canned or frozen foods, if high-holding-cost consumers (from lower income households or with smaller homes) have higher valuations, then Corollary 1 suggests that promotions are optimal. Finally, fixed shopping costs K are often incurred through traveling to the store and are thus related to consumers' location relative to the store. For a particular seller, Corollary 1 suggests that promotions are optimal when the consumers living close-by (with low K) have higher valuations.

Note that our specialized model in this subsection does have some limitations. We have assumed that the natural shopping frequencies T_o and $T_f = 1$ are integers, so the difference between consumer types must be rather significant for $T_o > T_f$ to hold. (For example, K_o must be at least 4 times as large as K_f , if all other parameters are the same between the two types.) Nevertheless, the numerical results of the previous section suggest that our conclusions do extend to more general cases. For example, Table 2 shows that the seller is more likely to prefer the HILO pricing format over fixed prices as the correlation between consumer valuations and fixed costs becomes more negative (i.e., significant coefficient of -0.4478); this is in agreement with our finding in Corollary 1(iii) above. Similarly, the other two coefficients corresponding to Corollary 1(i) and (ii) appear to have the correct sign (i.e., 0.0533 and 0.0555), although they are not statistically significant.

6.2 Frequency and Depth of Promotions

The previous subsection examines factors that influence the seller's optimal choice of pricing formats. Here, focussing on situations where the seller prefers promotional pricing, we study how the optimal frequency and depth of promotions vary with market characteristics.

Recall from Proposition 1 that under HILO pricing, the optimal discount (depth of promotion) is $\Delta = WTP_f - WTP_o$ and the optimal time interval between promotions (frequency of promotions) is T_o , as given below:

$$\begin{aligned}\Delta &= (v_f - v_o) + \sqrt{\frac{2h_o K_o}{R_o}} - \sqrt{\frac{2h_f K_f}{R_f}}, \\ T_o &= \sqrt{\frac{2K_o}{h_o R_o}}.\end{aligned}$$

It is natural to ask how these quantities of interest vary with consumer attributes along all three dimensions: holding cost, fixed cost, and consumption rate. In particular, as consumer attributes become more dispersed along each dimension (in the sense of a mean-preserving spread), how do the optimal promotional depth and frequency change? Our results below provide guidance for sellers operating in different markets, in which customers are primarily differentiated along different dimensions.

Proposition 2

- (i) *Suppose consumer holding costs go through a mean-preserving spread (i.e., h_f increases and h_o decreases). Then, the seller should offer smaller discounts and do so less frequently.*
- (ii) *Suppose consumer fixed costs go through a mean-preserving spread (i.e., K_f decreases and K_o increases). Then, the seller should offer larger discounts and do so less frequently.*
- (iii) *Suppose consumption rates go through a mean-preserving spread (i.e., R_f increases and R_o decreases). Then, the seller should offer larger discounts and do so less frequently.*

In summary, Proposition 2 indicates that a mean-preserving spread along any of the three dimensions will induce the seller to offer less frequent promotions. However, its impact on the magnitude of the discount is less clear cut. We find that mean-preserving spreads in consumption rates and in fixed costs both lead to deeper discounts. In contrast, a mean-preserving spread in holding costs leads to smaller discounts because an increase in h_f and a decrease in h_o will hamper the seller's ability to price discriminate against the frequent shoppers.

6.3 Inventory and Sales

So far, we have focused mainly on price outcomes, i.e., how prices should be adjusted over time. Now we turn attention to other measure of interest, such as quantities that are sold over time and consumer inventory levels over time. As above, we focus on the case where the RE equilibrium involves price promotions.

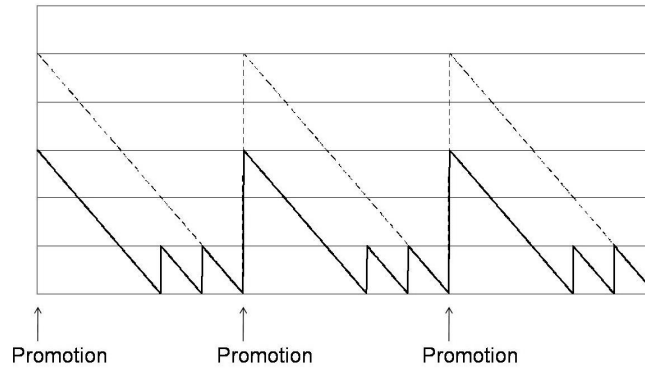


Figure 1: Inventory levels of frequent shoppers (solid line) and occasional shoppers (dashed line) over time.

Figure 1 shows an example of how consumer inventory levels fluctuate over time when the seller uses promotional pricing. In this example, promotions occur once every five periods. Occasional shoppers (represented by the dashed line) buy only on promotion and they stock up for consumption over all five periods. Frequent shoppers (represented by the solid line), on the other hand, stockpile for three periods when there is a promotion. Consequently, over the next two periods following the promotion, they do not visit the seller. In the fourth and fifth periods, frequent shoppers make purchases because they have run out of inventory. This cycle repeats itself every five periods.

Notice that periodic promotions induces inefficient stockpiling by the frequent shoppers. Their optimal shopping frequency (which minimizes shopping costs) is once every period, but instead, they choose to stock up during promotions and incur additional costs. Such stockpiling inefficiencies do not arise when the seller charges a fixed price. Under a fixed price, all consumers buy at their natural shopping frequency, which minimizes shopping costs.

Next, we look at how sales quantities fluctuate over time under promotional pricing. For the same example above, Figure 2 plots the quantities sold in each period over the price cycle. The figure

clearly demonstrates a phenomenon known as the post-promotional dip. Following the surge of sales at the start of the price cycle (in response to the attractive promotional price), there is an interval with no sales. During this time, consumers are comfortably stocked up and they have no incentive to buy. Toward the end of each price cycle, sales start to trickle in but only from frequent shoppers. At the next promotion, this cycle repeats itself. This type of sales pattern, characterized by a surge followed by zero and then mild sales within each price cycle, arises only under promotional pricing, strategic stockpiling, and consumer heterogeneity. All three elements are essential to generate the post-promotion dip.




Figure 2: Quantities sold over time.

7 Extensions

Now, we turn to several extensions of our model. In these extensions, we study how consumer stockpiling and the seller's pricing decisions interact with other supply-side strategies. In particular, we examine the seller's inventory, production, and rationing strategies.

7.1 Inventory management

In the basic analysis, we have assumed that the seller does not carry any inventory. For example, the seller may be a retailer who procures just enough units to meet consumer demand in every period. However, in practice, retailers often place orders in batches, as there may be economies of scale due to fixed ordering costs. In such cases, the seller needs to hold inventory between consecutive orders in order to satisfy demand before the next order arrives. How should the seller manage inventory

in the presence of consumer stockpiling? How do these inventory decisions interact with the pricing strategies studied above? We shall develop a model extension to address these questions.

In this extension, similar to the EOQ model, the seller incurs a fixed ordering cost of K_s and a per-unit inventory holding cost of h_s in every period. Let M denote the total number of consumers (i.e., market size). We build up on the two-type model introduced in Section 6. Recall that we have occasional and frequent shoppers with natural shopping intervals of T_o and T_f respectively, and the seller may choose between EDHP, EDLP, or HILO pricing strategies. Each of these pricing strategies generate different demand patterns. Our goal is to understand how the seller can minimize inventory-related costs while maintaining the required stock on hand to fulfill demand under each of the above classes of pricing strategies.

First, we consider the EDHP pricing strategies (i.e., the seller sells only to a single type of consumers). When the seller sells only to frequent shoppers, we have our basic EOQ model, which serves as a useful benchmark. The demand rate from frequent shoppers is $M\pi_f R_f$. Therefore, the seller's optimal cycle length is $T_s^* = \sqrt{\frac{2K_s}{h_s \cdot M\pi_f R_f}}$. In other words, the seller places an order once every T_s^* periods and sells to frequent shoppers at a constant rate over the entire order cycle. Next, the other alternative under EDHP is that the seller may sell only to occasional shoppers, who stock up once every T_o periods. In this case, when the seller orders precisely at these time intervals, he incurs no inventory cost because the entire stock is transferred to consumers at the start of each cycle.

Now, we consider the EDLP and HILO strategies (i.e., the seller sells to both consumer types). In these cases, the pricing strategies generate complex demand patterns that repeat every T_o periods.

- Under EDLP, the seller sells $M\pi_o R_o T_o$ units to occasional shoppers (who stock up) once every T_o periods, plus an additional $M\pi_f R_f$ units to frequent shoppers every period. In other words, the demand pattern is $M\pi_f R_f$ units every period, except at the start of the cycle when the demand is $M\pi_o R_o T_o + M\pi_f R_f$ units.
- Under HILO, occasional shoppers stock up for the entire cycle length of T_o periods but frequent shoppers stock up for a shorter time. Let S denote the length of time at the end of each cycle when frequent shoppers visit the seller. Then, the demand at the start of each cycle is $M\pi_o R_o T_o + M\pi_f R_f (T_o - S)$ units, and the demand at each of the final S periods is $M\pi_f R_f$ units. In between, the seller faces zero demand since consumers are all comfortably stocked up.

How should the seller manage inventory so that he has sufficient stock to fulfill the demand patterns described above? For each cycle of T_o periods, we find the inventory policy that minimizes cost while sustaining sales to consumers. Our results are summarized below.

Proposition 3 Let $T_s = \sqrt{\frac{2K_s}{h_s \cdot M\pi_f R_f}}$.

- (i) Under EDLP, the seller should divide each cycle of T_o periods into N segments whose lengths $\tau_1, \tau_2, \dots, \tau_N$ differ by at most one. At the start of each i -th segment, the seller places an order to fulfill all demand that arrives within that segment. Therefore, the seller orders $M\pi_o R_o T_o + M\pi_f R_f \tau_1$ units for the first segment and $M\pi_f R_f \tau_i$ units for each remaining segment, where $i = 2, \dots, N$. The possible candidates for N are the two integers closest to T_o/T_s .
- (ii) Under HILO, the seller places an order at the start of each cycle. In addition, the seller should divide the final S periods into N segments whose lengths $\tau_1, \tau_2, \dots, \tau_N$ differ by at most one, and place an order at the start of each segment. Therefore, the seller orders $M\pi_o R_o T_o + M\pi_f R_f (T_o - \sum_{i=1}^N \tau_i)$ units at the start of the cycle and $M\pi_f R_f \tau_i$ units for each remaining segment, where $i = 1, \dots, N$. The possible candidates for N are zero and the two integers closest to S/T_s .

This result highlights some major effects of consumer stockpiling on the seller's inventory strategies. First and foremost, when the seller sells to heterogeneous consumers with a propensity to stockpile, we show that the optimal inventory policy may involve different order quantities at possibly variable time intervals. In contrast, the basic EOQ model (which neglects consumer stockpiling) recommends placing equal-sized orders at regular intervals. We find that sellers may need to use a combination of *jumbo* orders (to fulfill demand from stockpiling consumers) as well as additional smaller *regular* orders. In addition, under HILO strategies, there may be a longer time lag after placing a jumbo order before placing the next (regular) order. However, such a lag is not necessary under EDLP pricing as all orders are still placed at regular intervals.

Another finding is that the optimal policy parameters depend heavily on the demand from frequent shoppers. Note that $T_s = \sqrt{\frac{2K_s}{h_s \cdot M\pi_f R_f}}$, which is based on frequent shopper demand, plays a key role in the results of Proposition 3. Our model suggests that market studies focusing on aggregate demand may not be sufficient. For inventory management purposes, it is also important to estimate consumer stockpiling tendencies in order to distill the demand rates of frequent shoppers.

Finally, our results suggest that consumer stockpiling can generate the well-known bullwhip effect, even in EDLP environments. As shown in Proposition 3, every-day-low-pricing does not eliminate consumer stockpiling, and the seller will find it optimal to place a jumbo order once in a while. These jumbo orders introduce variability that may propagate upwards along the supply chain. While price stability can help to control demand fluctuations, our model cautions that stockpiling on the consumer end can continue to inject variability into the system.

7.2 Production planning

Now, we consider the seller's production planning strategies. In the basic analysis, we have implicitly assumed that the seller can instantaneously produce enough units to meet demand in each period. However, in practice, the seller may face capacity constraints. For example, during a promotion, the seller may not be able to produce enough units in the same period to meet all the demand. In such cases, the seller may need to build up and carry inventory into the promotional period. As a result, the seller incurs inventory holding costs, which may influence his choice of a pricing strategy.

Here, we continue to employ the two-type model analyzed in Section 6. Recall that when frequent shoppers are willing to pay more, the seller chooses between HILO and EDHP, but in the reverse case, the seller chooses between EDLP and EDHP. Here, we shall study how limited production capacity and inventory holding costs will influence the seller's choice between the two possible pricing formats in each case. We shall assume that the production capacity is large enough, so that it is feasible for the seller to fulfill total demand over the entire time horizon, although advance production may be necessary.

Let $\Pi(EDLP)$, $\Pi(EDHP)$, $\Pi(HILO)$ denote the seller's long-run average profit under each pricing format. We have the following result.

Proposition 4

- (i) Suppose $WTP_f > WTP_o$ (i.e., the seller chooses between HILO and EDHP). With limited production capacity, the value of $\Pi(EDHP) - \Pi(HILO)$ is increased.
- (ii) Suppose $WTP_o > WTP_f$ (i.e., the seller chooses between EDLP and EDHP). With limited production capacity, the value of $\Pi(EDHP) - \Pi(EDLP)$ is increased.

This result indicates that the seller is more inclined to choose EDHP when there is limited production capacity. Intuitively, as a result of stockpiling behavior, selling to all consumers (via either EDLP or HILO) involves periodic demand surges. To fulfill these periodic surges of high demand (e.g., in the case of HILO, when there is a promotion), the seller may need to build up inventory, which incurs holding costs. Figure 3 shows an example of the capacitated seller’s inventory process under HILO (dashed line) and under EDHP (solid line); holding costs are clearly higher in the former case. As the capacity constraint becomes tighter, the seller has to begin production earlier and thus incurs higher holding costs. It may even be worthwhile for the seller to reduce holding costs by offering smaller and/or more frequent discounts so that the periodic demand surges become smaller. As long as there is limited production capacity, our results suggest that the seller will have an incentive to discourage consumer stockpiling.

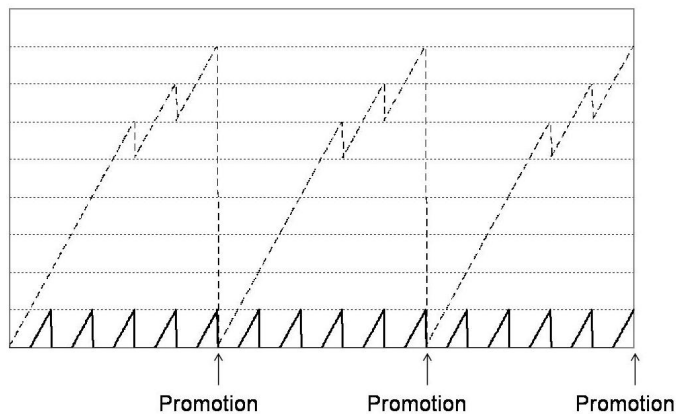


Figure 3: Seller’s inventory level with production capacity constraint, under HILO (dashed line) and EDHP (solid line).

7.3 Rationing strategies

In this extension, we consider the possibility that the seller may use inventory rationing to discourage strategic consumer stockpiling. In some situations, it may be profitable for the seller to withhold inventory and fulfill only a portion of available demand. We focus on rationing during price promotions (HILO), since it is clearly suboptimal to do so when the seller charges a fixed price. By limiting availability during a promotion, the seller encourages consumers to buy off-promotion and thus may

improve profits.

Suppose the seller has a finite amount of inventory to sell during a promotional period. Within this period, units are initially offered at the regular price and any remaining units are then put on sale (cf. Liu and van Ryzin, 2008a, 2008b). Consumers who wait for the sale price may find the product sold out. Note that consumers know when promotions will occur, but they do not know whether they will find the product available on sale. Since there is a fixed shopping cost, some consumers may opt to buy at the regular price and forego the *chance* of obtaining the product on promotion. Let ρ denote the probability that the product is available at the promotional price. In equilibrium, the seller can realize this probability by making the corresponding quantity available; that is, if the equilibrium demand during the promotional period is $D_{regular}$ and $D_{promotion}$ at the regular and promotional price, the seller stocks $D_{regular} + \rho D_{promotion}$ units, so that after the regular-priced units are sold, the availability probability is indeed ρ . Here, we assume proportional rationing; i.e., all consumers have equal access to the product.

The following proposition shows that promotional pricing becomes more attractive when we allow inventory rationing. Recall from Proposition 1(i) that when frequent shoppers have a higher willingness-to-pay, the seller chooses between HILO and EDHP pricing formats. Now, if rationing is allowed, the seller always prefers HILO.

Proposition 5 *Suppose $WTP_f > WTP_o$. If inventory rationing is allowed, the seller always uses HILO. The regular price is WTP_f and the promotional price of WTP_o is offered once every T_o periods but only with an availability probability of ρ . The two possible candidates for this probability are $\rho = 1$ (i.e., no rationing) and $\rho = \tilde{\rho} < 1$, where $\tilde{\rho} \equiv 1 / \left(1 + \frac{\Delta}{h_f}\right)^2$. The latter is preferred if the proportion of occasional shoppers π_o is sufficiently small.*

The intuition for the inventory rationing strategy proceeds as follows. The candidate availability probability $\tilde{\rho}$ obtained above is the cutoff probability below which frequent shoppers will prefer to buy at the regular price rather than wait for the discount and take the chance of not getting the product. Under this probability $\tilde{\rho}$, frequent shoppers always buy at the regular (high) price. Consequently, this inventory rationing strategy strictly dominates EDHP because, apart from regular sales to frequent shoppers, there are also promotional sales to occasional shoppers. Therefore, the seller always uses HILO and chooses between rationing (i.e., $\rho = \tilde{\rho}$) and no-rationing (i.e., $\rho = 1$). Here, the tradeoff

is that under rationing, some sales are lost when occasional shoppers are rationed, but doing so is essential to induce frequent shoppers to buy at the regular price. Therefore, inventory rationing is viable when the proportion of occasional shoppers is small.

This discussion shows that inventory rationing increases the profitability of price promotions. While periodic promotions allow the seller to practice inter-temporal price discrimination (i.e., across time periods), inventory rationing additionally allows for price discrimination within a time period. In practice, some retailers offer “while supplies last” promotions. At these promotions, although prices are low, demand is not completely satisfied. Our results suggest that these promotions cater to occasional shoppers with a relative low willingness-to-pay.

8 Conclusion

In this paper, we study the monopolist seller’s inter-temporal pricing problem when consumers strategically stockpile for future consumption. We develop a solution methodology based on rational expectations, and use it to provide a graphical interpretation of dynamic pricing policies in terms of cycles in a directed graph. With this result, we show that the equilibrium may involve either fixed prices or periodic promotions. Our findings shed light on the optimal choice of pricing formats. Specifically, we show that periodic promotions are preferred when frequent shoppers are willing to pay more than occasional shoppers for the product. We also use several model extensions to study the interaction between consumer stockpiling and the seller’s inventory, production, and rationing strategies.

There are several limitations in this study that present additional research opportunities. First, we have assumed an exogenous consumption rate for each consumer. For some product categories, consumption is flexible and individuals may consume more when the price is low (see Ailawadi and Neslin, 1998, and Bell, Iyer, and Padmanabhan, 2002). It would be useful to understand the effect of flexible consumption on equilibrium pricing strategies studied here. Second, we have also assumed that the seller charges a fixed price for every unit. It is conceivable that nonlinear pricing may be used to influence stockpiling behavior. Then, should quantity discounts be offered, and if so, when? Third, this analysis is also relevant in a supply chain setting. Instead of end-consumers, downstream buyers in a supply chain may also stockpile when they anticipate price fluctuations. This scenario is more complex because downstream buyers do not “consume” the products at a steady rate.

Instead, downstream buyers (e.g., retailers) sell to end-consumers who may also stockpile the product themselves. It is interesting to study optimal stockpiling strategies, in response to price fluctuations, at different locations along the supply chain. This brings us to our final suggested research direction: the current analysis can be extended to a competitive setting. An oligopoly counter-part to the present analysis serves to verify the robustness of our results. The models by Salop and Stiglitz (1982) and Pesendorfer (2002) may provide useful starting points. We conjecture that with multiple sellers, the equilibrium involves mixed strategies and randomized sales.

Appendix A: Graphical interpretation of equilibrium strategies

In this appendix, we provide an alternative linear programming formulation of the underlying Markov decision processes in our model. These alternative formulations yield an intuitive graphical interpretation of the equilibrium strategies.

Consider the seller's problem in Section 4.1. There is an alternative formulation of the dynamic programming problem in terms of a linear program. We call this the primal problem.

$$\min_{\lambda, g(\mathbf{I})} \quad \lambda \tag{5}$$

$$\text{such that } \lambda + g(\mathbf{I}) \geq p\hat{D}(\mathbf{I}, p) + g(\mathbf{I}'(\mathbf{I}, p)) \quad \forall \mathbf{I}, p. \tag{6}$$

It is well-known that the solution to the primal problem also satisfies the Bellman equations (2). At this point, it is convenient to write down the dual of the above linear program. We call this the dual problem.

$$\max_{x(\mathbf{I}, p)} \quad \sum_{\mathbf{I}} \sum_p p\hat{D}(\mathbf{I}, p)x(\mathbf{I}, p) \tag{7}$$

$$\text{such that } \sum_p x(\mathbf{I}, p) - \sum_{\{(\mathbf{I}_0, p_0): \mathbf{I}'(\mathbf{I}_0, p_0)=\mathbf{I}\}} x(\mathbf{I}_0, p_0) = 0 \quad \forall \mathbf{I}, \tag{8}$$

$$\sum_{\mathbf{I}} \sum_p x(\mathbf{I}, p) = 1, \tag{9}$$

$$x(\mathbf{I}, p) \geq 0 \quad \forall \mathbf{I}, p. \tag{10}$$

The dual problem has a useful and intuitive representation. Let us visualize a directed graph \mathcal{G} where each node represents a particular inventory state \mathbf{I} and each edge represents a transition out of state \mathbf{I} as a result of applying price p in that state. Thus, each dual variable $x(\mathbf{I}, p)$ is associated with a particular directed edge in this graph. The dual problem is to find a distribution of positive

weights $x(\mathbf{I}, p)$, summing to one, over all edges, so that the objective (7) is maximized. Notice that the constraint (8) is a flow balance condition: the sum of all weights flowing out each state \mathbf{I} , which is the first summation, must equal the sum of all weights flowing in, which is the second summation. The following result states that the optimal dual solution can be viewed as a cycle of states in the graph \mathcal{G} .

Lemma 1 *There exists an optimal solution to the dual problem $x^*(\mathbf{I}, p)$ such that the edges corresponding to all (\mathbf{I}, p) for which $x^*(\mathbf{I}, p) > 0$ form a cycle. Further, all such strictly positive dual variables are equal.*

Intuitively, the cycle corresponding to the optimal dual solution $x^*(\mathbf{I}, p)$ indicates the trajectory of states that the system goes through over time. In other words, given that the ordered set of strictly positive dual variables is $\{x^*(\mathbf{I}_1, p_1), x^*(\mathbf{I}_2, p_2), \dots, x^*(\mathbf{I}_m, p_m)\}$, we know that the system will cycle over the states $\mathbf{I}_1 \rightarrow \mathbf{I}_2 \rightarrow \dots \rightarrow \mathbf{I}_m \rightarrow \mathbf{I}_1 \rightarrow \dots$ by the seller applying price p_i in state \mathbf{I}_i . Therefore, the seller's prices also follow a cyclic pattern $p_1, p_2, \dots, p_m, p_1, p_2, \dots$, and so on. Similarly, the consumer's problem described in Section 4.2 can also be formulated as a linear program, and the dual problem has an analogous interpretation.

We stress that the dual formulation is important because it is the strictly positive variables at the dual solution that highlight the cycle of states. In the primal problem, such visualization is not possible because the primal solution yields the value functions for each state. In contrast, it is the binary nature of the dual solutions (either zero or strictly positive) that picks out the price and inventory cycles.

Appendix B: Proofs

Proof of Theorem 1 Consider the following static game between the seller and the consumers. The seller chooses a pricing strategy $P(\mathbf{I})$ and consumers choose a purchasing strategy $D_\theta(\mathbf{I}, p)$. We omit the state space for brevity and write P and D_θ . The corresponding payoff functions (which are functions of the above pricing and purchasing strategies) are the long-run average payoffs $\Pi(P; D_\theta)$ for the seller and $U_\theta(D_\theta; P, D_{-\theta})$ for the type- θ consumer. In this static game, a pure-strategy Nash equilibrium satisfies the conditions of our RE equilibrium in Definition 1. Therefore, it suffices to demonstrate the existence of a pure-strategy Nash equilibrium in our static game.

We proceed by studying the payoff functions $\Pi(P; D_\theta)$ for the seller and $U_\theta(D_\theta; P, D_{-\theta})$ for the type- θ consumer. Consider first the seller. Given D_θ , the seller's best response P^* that maximizes $\Pi(P; D_\theta)$ can be solved via dynamic programming as described in Section 4.1, so choosing the pricing strategy P is equivalent to choosing the value functions $g(\mathbf{I})$ and objective value λ . Thus, we may also write the payoff function as $\Pi(\lambda, g(\mathbf{I}); D_\theta)$. As we show in Appendix A, the seller's dynamic program is equivalent to the linear program (5)-(6). It is well known (see Rockafellar, 1970) that the optimal objective value in the above program is equal to that in its Lagrangian relaxation below:

$$\min_{\lambda, g(\mathbf{I})} \lambda + \sum_{\mathbf{I}} \sum_p L(\mathbf{I}, p) [\lambda + g(\mathbf{I}) - pD(\mathbf{I}, p) - g(\mathbf{I}'(\mathbf{I}, p))], \quad (11)$$

where $L(\mathbf{I}, p)$ are some Lagrange multipliers. Therefore, we may express the seller's payoff function as

$$\Pi(\lambda, g(\mathbf{I}); D_\theta) = - \left\{ \lambda + \sum_{\mathbf{I}} \sum_p L(\mathbf{I}, p) [\lambda + g(\mathbf{I}) - pD(\mathbf{I}, p) - g(\mathbf{I}'(\mathbf{I}, p))] \right\}, \quad (12)$$

which is linear in the seller's strategy λ and $g(\mathbf{I})$. Using the same argument, we can show that the consumer's payoff function is also linear in the consumer's own strategy. Since the payoff functions are linear and hence quasi-concave in each corresponding player's own strategy, and the strategy spaces are compact and convex, it then follows (see Fudenberg and Tirole, 1994) that our static game has a pure-strategy Nash equilibrium. ■

Proof of Theorem 2 This result follows from Lemma 1 and the derivations in Appendix A. ■

Proof of Lemma 1 By the flow balance constraint (8), the edges (\mathbf{I}, p) with strictly positive weights must form cycles. If there is more than one such cycle, we can redistribute the weights onto the single cycle with maximal value, where the value of a cycle is defined by the sum $\sum p \hat{D}(\mathbf{I}, p) x(\mathbf{I}, p)$ over edges in that cycle. This implies that there must be an optimal dual solution with weights over a single cycle. Finally, by flow balance, the weights over each edge along the cycle must be equal. ■

Proof of Proposition 1 (i) Suppose $WTP_f > WTP_o$. If the seller wishes to sell only to frequent shoppers, it is clear that he will use EDHP as described. It remains to consider the case where the seller wishes to sell to both types. In order to sell to occasional shoppers, there must be time periods in which the price is WTP_o or lower.

We first show that the low price must be exactly WTP_o . If it were lower, occasional shoppers will still prefer to shop at their natural frequency (once every T_o periods), knowing that the seller will prefer to offer a sale again whenever they run out of inventory (i.e. after T_o periods); thus, the seller is worse off by lowering the sale price below WTP_o .

Next, we show that the low prices WTP_o are indeed offered once every T_o periods. If the time between sales were longer, there would be time periods in which only frequent shoppers buy. If such periods earned a higher average profit, the seller would prefer EDHP described above. Since the seller does not prefer EDHP, such time periods (selling only to frequent shoppers) must earn a lower average profit, and they can be eliminated by offering sales every T_o periods. Similarly, the interval can not be shorter because occasional shoppers will not be willing to pay price WTP_o at this more costly shopping frequency.

Finally, we show that the price during all other periods must be WTP_f . This holds because only frequent shoppers are potential buyers and the seller will extract the entire surplus.

We have shown that the seller's pricing strategy is the HILO strategy described in the proposition. In this case, occasional shoppers buy only during a sale and they stock up for the entire T_o periods. On the other hand, frequent shoppers purchase for z periods' consumption during each sale, and we determine z below. After z periods, they buy every period (only for current consumption) until the next sale. To compute z , note that the frequent shopper's cost in each price cycle, if he stocks up for z periods during the sale, is

$$\left(K_f + z \cdot R_f \cdot WTP_o + \frac{h_f R_f z^2}{2} \right) + (T_o - z) \cdot R_f \cdot v_f. \quad (13)$$

The first term captures the cost to cover consumption for the z periods after the sale, and the remaining terms are the cost to cover consumption for the remaining $(T_o - z)$ periods until the next sale (this latter cost must equal the entire consumption valuation since the seller captures all surplus). To minimize cost, we equivalently minimize the following expression over z :

$$\left(K_f + z \cdot R_f \cdot (WTP_o - WTP_f) + \frac{h_f R_f z^2}{2} \right) - z \sqrt{2h_f K_f R_f}. \quad (14)$$

The minimizer is

$$z^* = \sqrt{\frac{2K_f}{h_f R_f}} + \frac{\Delta}{h_f} = 1 + \frac{\Delta}{h_f}. \quad (15)$$

Thus, during a sale, the frequent shopper stocks up for Δ/h_f future periods (may be rounded up or down), up to $(T_o - 1)$ future periods, as stated in the proposition.

(ii) Now, we consider the other case with $WTP_o > WTP_f$. If the seller wishes to sell to only one consumer type, he will use EDHP as described and only occasional shoppers will buy at their desired frequency. Next, if the seller wishes to sell to both types, there must be time periods in which the price is at WTP_f or lower. In any such period with the low price, suppose occasional shoppers stock up for τ_o periods and frequent shoppers stock up for $\tau_f < \tau_o$ periods. Then, during the $\tau_o - \tau_f$ periods in which only occasional shoppers still have inventory, the seller will optimally sell to frequent shoppers at price WTP_f (who buy every period). Thus, all consumption over these τ_o periods have been purchased at prices not exceeding WTP_f . The seller can simply use EDLP as described and obtain at least as much profit. Under the constant price WTP_f , frequent shoppers buy every period and occasional shoppers buy once every T_o periods. ■

Proof of Corollary 1 From Proposition 1, when the seller wishes to sell to both types, the two possible pricing formats are HILO and EDLP. If $WTP_f > WTP_o$, the former is more profitable; otherwise, the latter is more profitable. When all other dimensions of heterogeneity are equal across types, consumers with higher holding costs are the frequent shoppers. If they have sufficiently high valuations (i.e. v_f sufficiently large), the condition $WTP_f - WTP_o > 0$ holds and HILO is thus more profitable. This shows (i), and a similar argument applies to (ii) and (iii). ■

Proof of Proposition 2 First, note that the optimal promotional depth is $\Delta = (v_f - v_o) + \sqrt{\frac{2h_o K_o}{R_o}} - \sqrt{\frac{2h_f K_f}{R_f}}$, which decreases as h_f increases and h_o decreases, but increases when K_f decreases and K_o increases, as well as when R_f increases and R_o decreases. Next, note that the optimal time between promotions is $T_o = \sqrt{\frac{2K_o}{h_o R_o}}$, which increases as h_o, R_o decreases and when K_o increases. A larger T_o implies less frequent promotions. ■

Proof of Proposition 3 (i) The seller places an order at the start of each cycle to meet demand from occasional shoppers who stock up. Suppose there are N orders within each cycle. These orders must be evenly spaced since the holding cost incurred between consecutive orders is convex in the time between consecutive orders. Therefore, their lengths τ_1, \dots, τ_N may differ by at most one. Finally, the candidate values for N must be the integers closest to T_o/T_s since these solutions generate inventory cycle lengths that are closest to the EOQ solution $T_s = \sqrt{\frac{2K_s}{h_s \cdot M \pi_f R_f}}$, which minimizes cost.

(ii) The seller places an order at the start of each cycle to meet demand from occasional shoppers who stock up. If there are no further orders within each cycle, then we have the case with $N = 0$. If there are further orders, this should occur only during the final S periods because there is no demand before then. These additional orders help fulfill demand from frequent shoppers during the final S periods, and the optimal structure of these orders follows from the same argument as in part (i) above. ■

Proof of Proposition 4 (i) It suffices to show that the holding costs incurred due to limited production capacity is higher under HILO compared to EDHP. Consider the units sold to frequent shoppers under HILO. The holding costs incurred as a result of producing only these units already exceed the holding costs under EDHP, since the latter faces the same demand quantity but at a constant rate. Thus, the total holding costs under HILO must be even higher.

(ii) Now, we need to show that holding costs incurred due to limited production capacity is higher under EDLP compared to EDHP. The same argument above applies. ■

Proof of Proposition 5 Consider the HILO pricing strategy with rationing where the availability probability is ρ . If the frequent shopper does not wait for the sale, his average utility is zero since the seller extracts all surplus. If he waits for the sale, he incurs fixed cost K_f and then there is a probability ρ of stocking up for z periods and the remaining probability $1 - \rho$ of earning zero. Thus, the expected utility from waiting is

$$\begin{aligned}
\tilde{U}_f &= -K_f + \rho \left(zR_f \cdot v_f - zR_f \cdot WTP_o - \frac{h_f R_f z^2}{2} \right) \\
&= -K_f + \rho \left(zR_f \cdot (v_f - WTP_o) - \frac{h_f R_f z^2}{2} \right) \\
&= -K_f + \rho \left(zR_f \cdot \left(\Delta + \sqrt{\frac{2h_f K_f}{R_f}} \right) - \frac{h_f R_f z^2}{2} \right) \\
&= -K_f + \rho \left(zh_f R_f \cdot \left(\frac{\Delta}{h_f} + \sqrt{\frac{2K_f}{h_f R_f}} \right) - \frac{h_f R_f z^2}{2} \right) \\
&= -K_f + \rho \cdot \frac{h_f R_f \tilde{z}^2}{2},
\end{aligned}$$

recalling from Proposition 1 that $T_f = \sqrt{\frac{2K_f}{h_f R_f}} = 1$ and the frequent shopper's optimal choice of z is $\tilde{z} = 1 + \Delta/h_f$. Therefore, the shopper will not wait for the discount if and only if $\tilde{U}_f \leq 0$, or equivalently, $\rho \leq \frac{2K_f}{h_f R_f} / \tilde{z}^2 = 1/\tilde{z}^2 = \tilde{\rho}$, where $\tilde{\rho} = 1 / \left(1 + \frac{\Delta}{h_f}\right)^2$ as given in the proposition.

To maximize profit, the seller will not use any $\rho < \tilde{\rho}$ (since this only reduces sales to occasional shoppers) and will also not use any $\rho \in (\tilde{\rho}, 1)$ (since frequent shoppers will wait anyway, it is more profitable with $\rho = 1$). Notice also that EDHP is dominated by HILO with $\rho = \tilde{\rho}$ described above. Therefore, to maximize profit, the seller always uses HILO and chooses between $\rho = 1$ (no rationing) and $\rho = \tilde{\rho}$. Finally, observe that using $\rho = 1$ yields some consumer surplus to frequent shoppers (when they stock up during the promotion) while using $\rho = \tilde{\rho}$ allows the seller to capture all surplus from frequent shoppers. On the other hand, using $\rho = 1$ sells to all occasional shoppers while using $\rho = \tilde{\rho}$ sells only to some of them. Therefore, using $\rho = \tilde{\rho}$ is preferred when the proportion of frequent shoppers is large, or equivalently, when the proportion of occasional shoppers is small. ■

References

- [1] Ailawadi, K.L., S.A. Neslin. 1998. The Effect of Promotion on Consumption: Buying More and Consuming It Faster. *J. Mktg. Res.* 35(3): 390-398.
- [2] Anton, J.J., G. Das Varma. 2005. Storability, Market Structure and Demand-Shift Incentives. *RAND J. Econ.* 36(3): 520-543.
- [3] Assuncao, J.L., R.J. Meyer. 1993. The Rational Effect of Price Promotions on Sales and Consumption. *Mgmt. Sci.* 39(5): 517-535.
- [4] Aviv, Y., A. Pazgal. 2008. Optimal pricing of seasonal products in the presence of forward-looking consumers. *Mfg. & Service Oper. Mgmt.* 10(3): 339-359.
- [5] Bell, D.R., J. Chiang, V. Padmanabhan. 1999. The Decomposition of Promotional Response: An Empirical Generalization. *Mktg. Sci.* 18(4): 504-526.
- [6] Bell, D.R., G. Iyer, V. Padmanabhan. 2002. Price Competition Under Stockpiling and Flexible Consumption. *J. Mktg. Res.* 39(3): 292-303.
- [7] Blattberg, R.C., G.D. Eppen, J. Lieberman. 1981. A Theoretical and Empirical Evaluation of Price Deals for Consumer Nondurables. *J. Mktg.* 45(1): 116-129.
- [8] Chiang, J. 1991. A Simultaneous Approach to Whether, What and How Much to Buy Questions. *Mktg. Sci.* 10(4): 297-315.

- [9] Chintagunta, P. 1993. Investigating Purchase Incidence, Brand Choice and Purchase Quantity Decisions of Households. *Mktg. Sci.* 12(2): 184-209
- [10] Elmaghraby, W., A. Gulcu, P. Keskinocak. 2008. Designing optimal preannounced markdowns in the presence of rational customers with multiunit demands. *Mfg. & Service Oper. Mgmt.* 10(1): 126-148.
- [11] Erdem, T., S. Imai, M.P. Keane. 2003. Brand and Quantity Choice Dynamics Under Price Uncertainty. *Quant. Mktg. Econ.* 1(1): 5-64.
- [12] Fudenberg, D., J. Tirole. 1991. *Game Theory*. MIT Press, Cambridge, MA.
- [13] Gallego, G., R. Phillips, O. Sahin. 2008. Strategic management of distressed inventory. *Prod. Oper. Mgmt.* 17(4): 402-415.
- [14] Gallien, J. 2006. Dynamic mechanism design for online commerce. *Oper. Res.* 54(2): 291-310.
- [15] Golabi, K. 1985. Optimal Inventory Policies when Ordering Prices are Random. *Oper. Res.* 33(3): 575-588.
- [16] Gonul, F., K. Srinivasan. 1996. Estimating the Impact of Consumer Expectations of Coupons on Purchase Behavior: A Dynamic Structural Model. *Mktg. Sci.* 15(3): 262-279.
- [17] Guo, L., J.M. Villas-Boas. 2007. Consumer Stockpiling and Price Competition in Differentiated Markets. *J. Econ. Mgmt. Strat.* 16(4): 827-858.
- [18] Gupta, S. 1988. Impact of Sales Promotions on When, What, and How Much to Buy. *J. Mktg. Res.* 25: 342-355.
- [19] Hendel, I., A. Nevo. 2006a. Sales and consumer inventory. *RAND J. Econ.* 37(3): 543-561.
- [20] Hendel, I., A. Nevo. 2006b. Measuring the Implications of Sales and Consumer Inventory Behavior. *Econometrica.* 74(6): 1637-1673.
- [21] Ho, T.H., C.S. Tang, D.R. Bell. 1998. Rational Shopping Behavior and the Option Value of Variable Pricing. *Mgmt. Sci.* 44(12): S145-S160.

- [22] Hong, P., R.P. McAfee, A. Nayyar. 2002. Equilibrium Price Dispersion with Consumer Inventories. *J. Econ. Theory*. 105: 503-517.
- [23] Jeuland, A.P., C. Narasimhan. 1985. Dealing - Temporary Price Cuts - By Seller as a Buyer Discrimination Mechanism. *J. Busi.* 58(3): 295-308.
- [24] Krishna, A. 1992. The Normative Impact of Consumer Price Expectations for Multiple Brands on Consumer Purchase Behavior. *Mktg. Sci.* 11(3): 266-286.
- [25] Krishna, A. 1994. The Impact of Dealing Patterns on Purchase Behavior. *Mktg. Sci.* 13(4): 351-373.
- [26] Lal, R., J.D.C. Little, J.M. Villas-Boas. 1996. A Theory of Forward Buying, Merchandising, and Trade Deals. *Mktg. Sci.* 15(1): 21-37.
- [27] Levin, Y., J. McGill, M. Nediak. 2008. Optimal Dynamic Pricing of Perishable Items by a Monopolist Facing Strategic Consumers. *Prod. Oper. Mgmt.* Forthcoming.
- [28] Liu, Q., G. van Ryzin. 2008a. Strategic capacity rationing to induce early purchases. *Mgmt. Sci.* 54(6): 1115-1131.
- [29] Liu, Q., G. van Ryzin. 2008b. Strategic capacity rationing when customers learn. Working paper.
- [30] Meyer, R.J., J. Assuncao. 1990. The Optimality of Consumer Stockpiling Strategies. *Mktg. Sci.* 9(1): 18-41.
- [31] Narasimhan, C. 1988. Competitive Promotional Strategies. *J. Business.* 61(4): 427-449.
- [32] Neslin, S., C. Henderson, J. Quelch. 1985. Consumer Promotions and The Acceleration of Product Purchases. *Mktg. Sci.* 4(2): 147-165.
- [33] Ovchinnikov, A., J.M. Milner. 2005. Strategic response to wait-or-buy: revenue management through last-minute deals in the presence of customer learning. Working paper.
- [34] Pesendorfer, M. 2002. Retail Sales: A Study of Pricing Behavior in Supermarkets. *J. Busi.* 75: 33-66.

- [35] Puterman, M.L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY.
- [36] Salop, S., J.E. Stiglitz. 1982. The Theory of Sales: A Simple Model of Equilibrium Price Dispersion with Identical Agents. *Amer. Econ.Rev.* 72(5): 1121-1130.
- [37] Shen, Z.J., X. Su. 2007. Customer Behavior Modeling in Revenue Management and Auctions: A Review and New Research Opportunities. *Prod. & Oper. Mgmt.* 16(6): 713-728.
- [38] Shoemaker, R.W. 1979. An Analysis of Consumer Reactions to Product Promotions. *Educators' Conference Proceedings*. Beckwith, N., M. Houston, R. Mittelstaedt, et al. (eds.) Chicago: Amer. Mktg. Assoc. 244-248.
- [39] Su, X. 2007. Intertemporal Pricing with Strategic Customer Behavior. *Mgmt. Sci.* 53(5): 726-741.
- [40] Talluri, K.T., G.J. van Ryzin. 2005. *The Theory and Practice of Revenue Management*. Springer. New York, NY.
- [41] Varian, H.R. 1980. A model of sales. *Amer. Econ. Rev.* 70(4): 651-659.
- [42] Ward, R.W., J.E. Davis. 1978. Coupon Redemption. *J. Advertising Res.* 18(4): 51-58.
- [43] Xu, X., W.J. Hopp. 2004. Customer heterogeneity and strategic behavior in revenue management: a martingale approach. Working paper.
- [44] Yin, R., Y. Aviv, A. Pazgal, C.S. Tang. 2009. Optimal Markdown Pricing: Implications of Inventory Display Formats in the Presence of Strategic Customers. *Mgmt. Sci.* Forthcoming.
- [45] Zhou, Y.P., M. Fan, M. Cho. 2005. On the threshold purchasing behavior of customers facing dynamically priced perishable products. Working paper.