THE REVIEW OF ECONOMIC STUDIES

Review of Economic Studies (2008) **75**, 391–413 0034-6527/08/00170391$02.00

Durable-Goods Monopoly with Varying Demand

SIMON BOARD

Department of Economics, University of Toronto and Department of Economics, University of California, Los Angeles

First version received May 2005*; final version accepted August* 2007 (*Eds.*)

This paper solves for the profit-maximizing strategy of a durable-goods monopolist when incoming demand varies over time. We first characterize the consumers' optimal purchasing decision by a cut-off rule. We then show that, under a monotonicity condition, the profit-maximizing cut-offs can be derived through a myopic algorithm, which has an intuitive marginal revenue interpretation. Consumers' ability to delay creates an asymmetry in the optimal price path, which exhibits fast increases and slow decreases. This asymmetry creates an upward bias in the level of prices, pushing them above the price charged by a firm facing the average level of demand. The optimal policy outperforms renting and can be implemented by a time-consistent best-price provision.

1. INTRODUCTION

In her classic 1979 paper, Nancy Stokey asked whether a firm should use time to discriminate between different types of consumers. By lowering its price over time, the firm can make extra sales to low-valuation agents; however, such a strategy will lead some high-valuation agents to postpone their purchases. When it can commit to its pricing policy, Stokey concluded that the firm should forgo the opportunity to discriminate, setting the price equal to the static monopoly price and holding it there forever. On the equilibrium path, sales only occur in the first period and no consumer ever delays.

Relaxing the assumption that demand is never refreshed, Conlisk, Gerstner and Sobel (1984) and Sobel (1991) allow a homogenous set of consumers to enter the market each period. They show that, once again, the profit-maximizing strategy is to set the price equal to the static monopoly price. There are sales in each period, but no consumer ever delays.

The purpose of the current paper is to analyse the firm's profit-maximization problem when incoming demand varies over time. We characterize the optimal sequence of prices and allocations, and examine how the firm uses time to discriminate between different generations. As incoming demand changes, prices fluctuate and some consumers choose to strategically delay their consumption. We show that this delay induces an asymmetry in the optimal price path and raises the average level of prices.

Variation in demand arises naturally in many markets. For example, each September, thousands of students return to universities and colleges across Canada. Canadian Tire, a large retailer, responds to this influx of new consumers by holding a back-to-school sale, reducing prices on furniture and stationery. While the price reduction helps increase profits from price-sensitive students, it also leads consumers who would have bought in July and August to delay their purchases. When choosing their September price, Canadian Tire must therefore trade off these two effects. Anticipating the back-to-school sale, Canadian Tire must also consider how to alter prices in July and August in order to mitigate consumer delay.

Allowing different generations to have different demand functions significantly complicates the firm's problem. Agent's types have two dimensions: their valuation and their birthdate, so the challenge is to identify which incentive compatibility constraints bind. Once we know which generations purchase in which periods (*e.g.* that consumers from August, but not July, will wait until September), then it is easy to derive the optimal prices from the firm's profit-maximization problem. Of course, which generations choose to delay are themselves endogenous and depend on the price path.

Rather than trying to directly solve for the optimal prices, as in Stokey (1979), we solve the problem by working in quantity space. First, we characterize the consumers' purchasing rule by a sequence of cut-offs: at any time t, a consumer purchases if his valuation lies above some time-t cut-off. We then use mechanism design to eliminate prices and describe the firm's profit as a function of these cut-offs.

Second, under a monotonicity condition on marginal revenue, we characterize the profit-maximizing cut-offs by a myopic algorithm, where the allocation at time t only depends upon the consumers who have entered the market up to time t. Consumers are forward looking, so prices depend upon the sequence of future demand; the optimal allocations, however, only depend on past demand. The myopic nature of this algorithm enables us to analyse any stochastic sequence of demand functions: monotone demand paths, demand cycles, i.i.d. demand draws, and so on. The optimal prices can then be derived from the consumers' utility-maximization problem.

The optimal myopic algorithm has an intuitive interpretation: The incoming demand curve in any period t can be associated with a marginal revenue curve, where marginal revenue is with respect to price, not quantity. In the first period, the firm sells the good to consumers with positive marginal revenue (net of costs). In each period thereafter, the firm adds the marginal revenue of the old consumers who have yet to buy the marginal revenue of the new agents, forming a cumulative marginal revenue function. The firm then sells the good to agents with positive cumulative marginal revenue.

Consumers' ability to delay induces an asymmetry in the optimal price path. When demand grows stronger over time, in that valuations tend to rise, the firm will want to increase its price. Agents then have no incentive to delay and the firm can discriminate between the different generations, charging the monopoly price against the incoming generation (the "myopic price"). When demand weakens over time, in that valuations tend to fall, the firm will want to decrease its price. Charging the myopic price, however, will lead to falling prices, causing some consumers to delay their purchases. In order to mitigate this delay, the firm slows the rate at which prices fall. The price path is flattened so much that prices always remain above the price chosen by a firm who faces the average level of demand (the "average demand price").

The asymmetry between increases and decreases in demand crucially affects the firm's optimal pricing policy in other environments. Suppose demand follows stationary cycles, as shown in Figure 1, where the lower panel describes which generations purchase in which periods. After the first cycle, prices follow a stationary shark-fin pattern, rising quickly and falling slowly. The pattern has two interesting features. First, the price is minimized in the last period of the slump, not in the period of lowest demand. Second, the asymmetry between increases and decreases in demand raises the price level, so that prices always exceed the average demand price. This means that the introduction of demand variation makes all consumers worse off and reduces social welfare.

The basic model makes two assumptions of note. First, there is no resale. There are many goods for which this is the right assumption: one-time experiences (*e.g.* a trip to Disneyland), intermediate products (*e.g.* aluminium), and goods with high transaction costs (*e.g.* fridges). However, such an assumption is not innocuous. With resale, price variations decrease in amplitude as

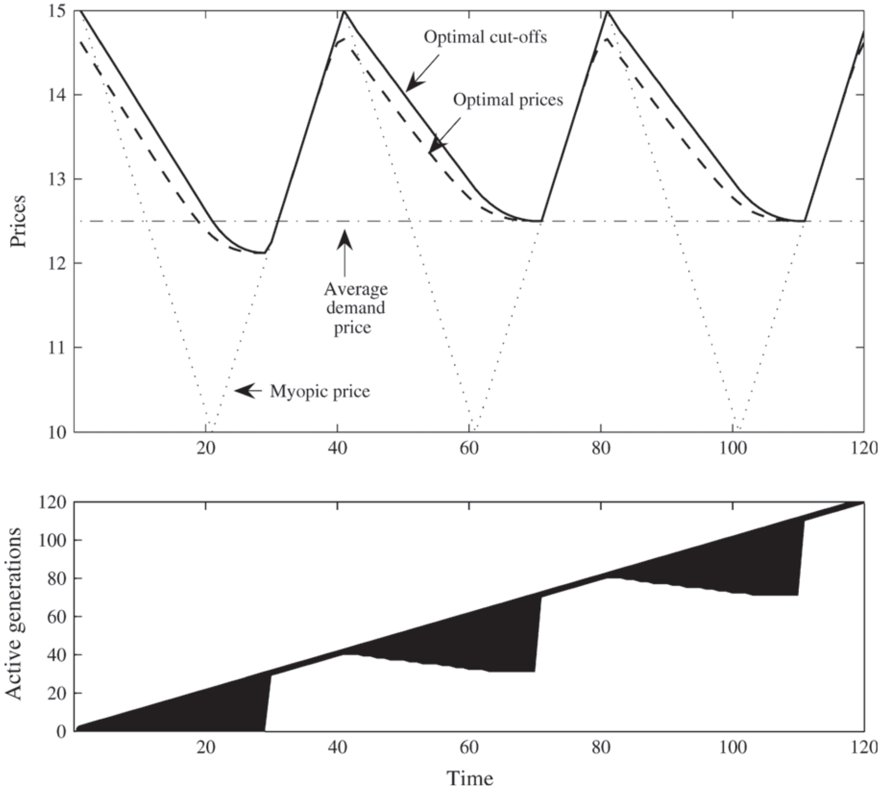

FIGURE 1

This figure assumes that valuations have measure 1 on $[0, b_t]$, $b_t \in [20, 30]$, and the discount rate is $\delta = 0{\cdot}75$. The top panel shows how prices and cut-offs vary when demand follows cycles. The bottom panel illustrates which generations are active in any given period. Examples 4 and 6 further analyse the linear demand example

the market gets older; without resale, the market fluctuations never abate. With resale, the price responds symmetrically to changes in demand; without resale, price movements are highly asymmetric. Introducing resale has no effect on allocations if demand falls over time, but otherwise lowers the firm's profits. As observed by Coase (1972) and Bulow (1982), renting is identical to commitment pricing with resale since both policies induce the same set of implementable allocations. Renting is therefore optimal only if demand is declining. This result also means that, in the presence of demand variation, the monopolist has an incentive to shut down any secondary market.

The second important assumption is that the firm can commit to a sequence of prices. With homogenous demand and no commitment, there are many equilibria ranging from those which are very bad for the firm to others which are close to the full-commitment outcome (Sobel, 1991). This paper should thus be viewed as establishing the best possible outcome for the firm. This seems particularly reasonable if the firm is concerned about its reputation across several durable-goods markets. In addition, there are contractual solutions to the commitment problem: we extend the result of Butz (1990) by showing that a best-price provision can implement the optimal scheme without pre-commitment.

The paper solves for the optimal price sequence of a durable-goods monopolist, but one can also look at more general mechanisms. For example, Segal (2003) shows that, without the entry of new generations, the optimal mechanism in a durable goods model can be implemented by a price mechanism if marginal costs are constant or if there are large numbers of consumers. Both these assumptions are satisfied in our model.

The precursors of this paper are the models of Stokey (1979), Conlisk *et al.* (1984) and Sobel (1991), as described in Section 3.2. Also closely related, Dudine, Hendel and Lizzeri (2006) analyse the problem of a storable-goods monopolist who faces varying demand. This differs from the current paper in two ways. First, Dudine *et al.* assume that the storage cost takes the form of a fixed fee, so all types of agents have the same incentive to intertemporally shift their consumption. This enables the authors to characterize the optimal prices both with and without commitment. Second, with a storable good, agents arbitrage a price spike by shifting consumption forward, causing prices to fall quickly and rise slowly.[1]

Our model is also related to the product life cycle literature. Rogers (1962) argues that the nature of new demand changes over the life cycle of a product: early adopters are "venturesome and daring", while laggards are "skeptical and traditional". The analysis of pricing of new goods has since focused on encouraging new demand via network effects, overlooking agents' intertemporal consumption decisions (*e.g.* Krishnan, Bass and Jain, 1999). This paper, in contrast, analyses the pricing of a new product where agents can strategically time their purchases.[2]

The remainder of the paper is organized as follows: Section 2 describes the model. Section 3 derives the firm's profit-maximization problem, which is solved in Section 4. Section 5 discusses applications including monotone demand paths and demand cycles. Section 6 extends the results to stochastic demand, while Section 7 analyses resale and renting. Section 8 concludes.

2. MODEL

Preliminaries. Time is discrete, $t \in \{1, \ldots, T\}$, where we allow $T = \infty$. Consumers and the firm have a common discount factor, which may vary over time, denoted by $\delta_t \in [0, \overline{\delta}] \subset [0, 1)$. Discounting up to time t is given by $\Delta_t = \prod_{s=1}^{t} \delta_s$, where $\Delta_t := 0$ for $t \geq T + 1$.

Preferences. A consumer with valuation $\theta \in [\underline{\theta}, \overline{\theta}]$ who purchases at time $t \in \{1, \ldots, T, \infty\}$ and price p_t obtains utility

$$(\theta - p_t)\Delta_t.$$

A consumer always has the option not to purchase, in which case they obtain zero utility and are said to buy at time $t = \infty$.

Demand. Each period, consumers of measure $f_t(\theta)$ enter the market. Let $F_t(\theta)$ be the absolutely continuous distribution function, where $F_t(\overline{\theta})$ is the total number of agents (which may vary over time), and denote the survival function by $\overline{F}_t(\theta) := F_t(\overline{\theta}) - F_t(\theta)$. The sequence of demand curves is deterministic and common knowledge.

1. Other authors introduce dynamics into durable goods models in different ways. Conlisk (1984) and Biehl (2001) have stochastic valuations, while cost variations have been analysed by Stokey (1979). When consumers and the firm have different discount rates, the optimal price may fall over time, as examined by Landsberger and Meilijson (1985). There is also a line of work where the good depreciates over time and consumers are allowed to scrap the product, such as Rust (1986), Waldman (1996), and Hendel and Lizzeri (1999).

2. In this paper, price fluctuations are simply driven by changes in elasticities, such as those during seasonal demand peaks (Nevo and Hatzitaskos, 2005), over the product life cycle (Gordon, 2007), or those due to shifting demographics (Lach, 2007). There may, however, be other reasons why firms hold sales. Price fluctuations may be internally generated (Conlisk *et al.*, 1984) and come from competition between stores (Varian, 1980) or from firms price discriminating against those unable to time their purchases. Sales may also be caused by aggregate demand shocks which undermine collusion (Rotemberg and Saloner, 1986) or intensify searching (Warner and Barsky, 1995).

Consumers' problem. Consider a consumer of type (θ, t) with valuation θ who enters in period t. Given a sequence of prices $\{p_t\}$, they have the problem of choosing a purchasing time $\tau(\theta, t) \in \{t, \ldots, T, \infty\}$ to maximize utility

$$u_t(\theta) = \sup_{\tau} (\theta - p_\tau)\Delta_\tau. \tag{2.1}$$

Let $\tau^*(\theta, t)$ be the earliest solution to this problem.[3]

Firm's problem. Assume that marginal cost is constant and normalize it to zero. Given a sequence of prices $\{p_t\}$, the firm's profit is

$$\Pi = \sum_{t=1}^{T} \int_{\underline{\theta}}^{\overline{\theta}} \Delta_{\tau^*(\theta,t)} p_{\tau^*(\theta,t)} dF_t, \tag{2.2}$$

where $\tau^*(\theta, t)$ maximizes the consumer's utility (2.1). The firm's problem is to choose $\{p_t\}$ to maximize (2.2).

3. SOLUTION TECHNIQUE

3.1. *Firm's problem*

Consumers choose their purchase times optimally, so we can apply the envelope theorem to the utility-maximization problem (2.1).[4] This yields utility

$$u_t(\theta) = \int_{\underline{\theta}}^{\theta} \Delta_{\tau^*(x,t)} dx + u_t(\underline{\theta}). \tag{3.1}$$

Since the seller will always choose prices $p_t \geq \underline{\theta}$ $(\forall t)$, it will be the case that $u_t(\underline{\theta}) = 0$ $(\forall t)$. Integrating by parts, *consumer surplus* from generation t is

$$\int_{\underline{\theta}}^{\overline{\theta}} u_t(\theta) dF_t = \int_{\underline{\theta}}^{\overline{\theta}} \Delta_{\tau^*(\theta,t)} \overline{F}_t(\theta) d\theta. \tag{3.2}$$

Welfare from generation t is defined by

$$W_t := \int_{\underline{\theta}}^{\overline{\theta}} \Delta_{\tau^*(\theta,t)} \theta dF_t \tag{3.3}$$

with total welfare $W := \sum_t W_t$. Since costs are zero, the welfare-maximizing pricing scheme sets all prices equal to zero. Profit equals welfare (3.3) minus consumer surplus (3.2),

3. We use the earliest purchasing time for ease of presentation. This choice does not affect profits: Lemma 1(b) implies that any optimal selection is decreasing in θ and therefore must coincide (a.e. θ).

4. The generalized envelope theorem of Milgrom and Segal (2002, Theorem 2) requires that the set of optimal purchasing times is non-empty. This follows from Lemma 1 in Section 3.3.

$$\Pi = \sum_{t=1}^{T} \int_{\underline{\theta}}^{\overline{\theta}} [\Delta_{\tau^*(\theta,t)}\theta - u_t(\theta)] dF_t$$

$$= \sum_{t=1}^{T} \int_{\underline{\theta}}^{\overline{\theta}} \Delta_{\tau^*(\theta,t)} m_t(\theta) d\theta, \tag{3.4}$$

where $m_t(\theta) := \theta f_t(\theta) - \overline{F}_t(\theta)$ is the *marginal revenue* with respect to price.[5]

Profit is thus the discounted sum of marginal revenues. Notice how the marginal revenue gained from agent (θ, t) is the same no matter when they choose to buy. That is, an agent's marginal revenue is determined when they are born and sticks to them forever. Intuitively, when the firm lowers the price in period t and awards the object to an extra few agents, they gain revenue from the extra sales, but lose money from the inframarginal agents who buy in that period. In addition, the firm must also lower all the prices in earlier periods in order to prevent agents delaying and so loses money from all the inframarginal agents who bought previously.

The *firm's problem* is thus to choose prices $\{p_t\}$ to maximize profit (3.4) subject to consumers choosing their purchasing time $\tau^*(\theta, t)$ to maximize utility (2.1).

Some definitions will be useful. Period H is said to have *higher demand* than period L if $m_H^{-1}(0) \geq m_L^{-1}(0)$, so the optimal static monopoly price is higher under $F_H(\theta)$ than $F_L(\theta)$.[6] There may, of course, be more people under the "low" demand, that is, $F_H(\overline{\theta}) \leq F_L(\overline{\theta})$, as is the case in the "back-to-school" example in the Introduction. Denote the marginal revenue from a set of generations $A \subset \{1, \ldots, T\}$ by $m_A(\theta) := \sum_{s \in A} m_s(\theta)$, and let $m_{\leq t}(\theta) := m_{\{1,\ldots,t\}}(\theta)$.

3.2. *Special cases*

We now consider three special cases that will serve as useful benchmarks. Example 1 comes from Stokey (1979), where there is a single demand curve of consumers and entry never occurs. Example 2 supposes that the monopolist can set a different price schedule for each generation. Example 3 is the homogenous entry model of Conlisk *et al.* (1984).

Example 1 (Single generation). *If $f_t(\theta) = 0$ for $t \geq 2$, then profit (3.4) reduces to*

$$\Pi = \int_{\underline{\theta}}^{\overline{\theta}} \Delta_{\tau^*(\theta,t)} m_1(\theta) d\theta. \tag{3.5}$$

To maximize (3.5), the firm would like to set purchasing times as follows:

$$\tau^*(\theta, 1) = 1, \qquad \text{if } m_1(\theta) \geq 0,$$
$$\tau^*(\theta, 1) = \infty, \qquad \text{if } m_1(\theta) < 0.$$

5. It is more common to use marginal revenue with respect to quantity, $\mathrm{MR}_t(\theta) := m_t(\theta)/f_t(\theta)$. However, since we will be adding demand across generations, summing demand curves horizontally, it is easier to work with marginal revenue with respect to price. In any case, both $m_t(\theta)$ and $\mathrm{MR}_t(\theta)$ have the same roots and are thus interchangeable in pricing formulae.

6. A sufficient condition for period H to have higher demand than period L is that $F_H(\theta)$ is larger than $F_L(\theta)$ in hazard order, $\overline{F}_H(\theta)/f_H(\theta) \geq \overline{F}_L(\theta)/f_L(\theta)$.

That is, the firm would like positive marginal revenue consumers to purchase immediately and the rest to never buy. If $m_1(\theta)$ is increasing in θ, this optimal policy can be implemented by setting

$$p_1^* = m_1^{-1}(0),$$

$$p_t^* \geq m_1^{-1}(0), \quad \text{if } t \geq 2.$$

Since the price is increasing, consumers buy in period 1 or never buy at all.

Example 2 (Discrimination between generations). *Next, suppose the firm could tell the different cohorts apart and set a price p_t^s for generation s in time t. The firm would then implement Stokey's solution for each cohort. That is,*

$$p_t^s = m_s^{-1}(0), \quad \text{if } t = s,$$

$$p_t^s \geq m_s^{-1}(0), \quad \text{if } t \geq s+1.$$

Hence, if demand grows over time, $m_t^{-1}(0) \geq m_{t-1}^{-1}(0)$, the seller can simply charge the myopic price, $p_t^* = m_t^{-1}(0)$, and need not discriminate.

Example 3 (Homogenous demand). *Finally, suppose demand is identical in each period, $f_t(\theta) = f_1(\theta)$ $(\forall t)$. The firm can implement the discriminatory optimum from Example 2 by setting $p_t^* = m_1^{-1}(0)$ $(\forall t)$.*

3.3. *The cut-off approach*

As suggested by Examples 1–3, rather than solving for prices directly, it is easier to solve for the optimal purchasing rule and back-out prices. This approach works since prices $\{p_t\}$ only enter into profits (3.4) via the purchasing rule $\tau^*(\theta, t)$—a standard feature of quasi-linear mechanism design problems. This is analogous to solving a standard monopoly model in quantities and using the demand curve to derive prices.

Lemma 1. *The earliest purchasing rule, $\tau^*(\theta, t)$, has the following properties:*

(a) Existence. $\tau^*(\theta, t)$ exists.
(b) θ-monotonicity. $\tau^*(\theta_H, t) \leq \tau^*(\theta_L, t)$, for $\theta_H \geq \theta_L$.
(c) Non-discrimination. If $\tau^*(\theta, t_L) \geq t_H$, then $\tau^*(\theta, t_L) = \tau^*(\theta, t_H)$, for $t_H \geq t_L$.
(d) Right-continuity. $\{\theta : \tau^*(\theta, t_L) \leq t_H\}$ is closed, for $t_H \geq t_L$.

Proof.

(a) The set of optimal purchasing rules is characterized by backwards induction. The minimal element can be attained by using the rule: stop when current utility is weakly greater than the continuation utility (Chow, Robbins and Siegmund, 1971, Theorem 4.2).
(b) $u_t(\theta)$ has strictly decreasing differences in (θ, τ), since $\delta_t < 1$, and is (weakly) supermodular in τ. Hence, every optimal selection is decreasing by Topkis (1998, Theorem 2.8.4).
(c) If $\tau^*(\theta, t_L) \geq t_H$, then the problems of agents (θ, t_L) and (θ, t_H) are identical. Hence, the earliest purchasing times coincide.
(d) Right continuity follows from the fact that $u_t(\theta)$ is continuous in θ. ‖

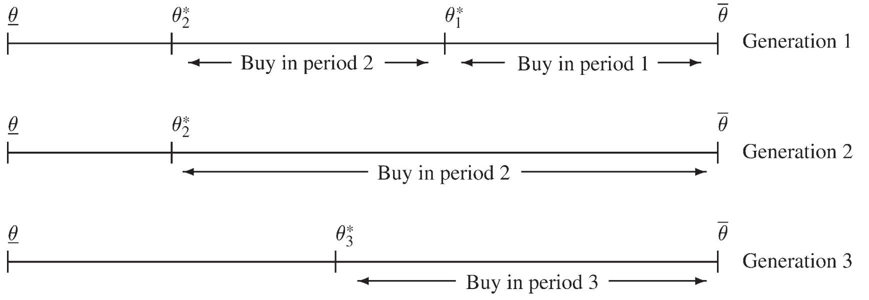

FIGURE 2

This figure shows a three-period example assuming $\theta_1^* > \theta_3^* > \theta_2^*$. Taking the cut-offs as given, it illustrates which agents purchase in which periods

Lemma 1 implies that the optimal purchasing time $\tau^*(\theta, t)$ can be characterized by a sequence of cut-offs. The *cut-off* θ_t^* is the lowest value agent from generation t who purchases in period t,

$$\theta_t^* := \min\{\theta : \tau^*(\theta, t) = t\}.$$

That is, consumers in the market in period t will buy if their valuation exceeds θ_t^*. For generation $t_L \leq t_H$, the *updated cut-off*, $\theta^*(t_H; t_L)$, is the lowest value agent from generation t_L who buys by time t_H,

$$\theta^*(t_H; t_L) := \min_{t_H \geq s \geq t_L} \theta_s^*. \tag{3.6}$$

If $t_H < t_L$, then set $\theta^*(t_H; t_L) = \infty$. Consumer θ from generation t_L will then buy in period t_H if

$$\theta \in [\theta^*(t_H; t_L), \theta^*(t_H - 1; t_L)).$$

A simple three-period example is shown in Figure 2.

We have shown that we can move from prices to cut-offs. The reverse is also true: prices can be backed out from any sequence of cut-offs. First suppose T is finite. In the last period, an agent with value θ_T^* is indifferent between buying and not, so the firm sets $p_T^* = \theta_T^*$. In earlier periods, an agent with value θ_t^* should be indifferent between buying in period t and waiting. Hence, prices are determined by the following algorithm:[7]

$$\begin{aligned} \Delta_t(\theta_t^* - p_t^*) &= \max_{\tau \geq t+1} (\theta_t^* - p_\tau^*)\Delta_\tau \\ &= \left(\theta_t^* - p_{\tau(\theta_t^*, t+1)}^*\right) \Delta_{\tau(\theta_t^*, t+1)}, \end{aligned} \tag{3.7}$$

where $\tau(\theta_t^*, t+1) = \min\{\tau \geq t+1 : \theta_t^* \geq \theta_\tau^*\}$. When T is infinite, there is no last period, but equation (3.7) remains valid. One can then calculate prices by truncating the problem, calculating prices for a finite T and letting $T \to \infty$, as shown by Chow *et al.* (1971, Theorems 4.1 and 4.3).

The *firm's problem* is thus to choose cut-offs $\{\theta_t^*\}$ to maximize profit (3.4).

4. OPTIMAL PRICING

This section presents the main results of the paper. Section 4.1 characterizes the optimal cut-offs using a simple myopic algorithm. Section 4.2 provides an alternative characterization of the

7. If the cut-offs lie in $(\underline{\theta}, \overline{\theta})$, these prices are unique. An equivalent way to obtain prices is to equate (2.1) and (3.1) under the optimal purchasing time $\tau^*(\theta, t)$.

optimal cut-offs that identifies which generations purchase in which periods. Finally, Section 4.3 shows that the optimal policy can be implemented by a time-consistent best-price provision.

4.1. *General solution*

Definition 1. Cumulative marginal revenue equals $M_t(\theta) := m_t(\theta) + \min\{M_{t-1}(\theta), 0\}$, where $M_1(\theta) := m_1(\theta)$.

We make the following weak monotonicity (WM) assumption.

Assumption (WM). *$M_t(\theta)$ is quasi-increasing* $(\forall t)$.[8]

The assumption that $m_t(\theta)$ is quasi-increasing (*i.e.* that revenue from generation t is quasi-concave in price) holds for most common distributions (*e.g.* normal, lognormal, exponential). If $m_t(\theta)$ is quasi-increasing and demand is increasing, in that $m_t^{-1}(0) \geq M_{t-1}^{-1}(0)$, then $M_t(\theta)$ is automatically quasi-increasing. If $m_t(\theta)$ is quasi-increasing and demand is decreasing, then, broadly speaking, $M_t(\theta)$ is quasi-increasing if the demand reduction is not too great. For example, with linear or semi-log demand, WM is satisfied if demand does not double between the lowest and highest periods.[9]

Theorem 1. *Under WM, the profit-maximizing cut-offs are given by $\theta_t^* = M_t^{-1}(0)$.*

Proof. Let Π_t be the profit from generation t, and $\Pi_{\geq t}$ be the profit from generations $\{t, \dots, T\}$. Denote the positive and negative components by $M_t^+(\theta) := \max\{M_t(\theta), 0\}$ and $M_t^-(\theta) := \min\{M_t(\theta), 0\}$. The proof will proceed by induction, starting with period $t = 1$.

Fix $\{\tau(\theta, t')\}$ for $t' \geq 2$, and consider the optimal choice of $\tau(\theta, 1)$. Notice that the non-discrimination condition in Lemma 1 implies that $\tau(\theta, 1) \in \{1, \tau(\theta, 2)\}$. Splitting up the profit equation (3.4), $\Pi = \Pi_1 + \Pi_{\geq 2}$,

$$\Pi = \int_{\underline{\theta}}^{\overline{\theta}} [\Delta_{\tau(\theta,1)} M_1^+(\theta) + \Delta_{\tau(\theta,1)} M_1^-(\theta)] d\theta + \Pi_{\geq 2}$$

$$\leq \int_{\underline{\theta}}^{\overline{\theta}} [\Delta_1 M_1^+(\theta) + \Delta_{\tau(\theta,2)} M_1^-(\theta)] d\theta + \Pi_{\geq 2}.$$

The second line solves for the optimal choice of $\tau(\theta, 1)$, setting $\tau^*(\theta, 1) = 1$ if $M_1(\theta) \geq 0$ and $\tau^*(\theta, 1) = \tau(\theta, 2)$ if $M_1(\theta) < 0$. This is independent of the value of $\tau(\theta, 2)$. If $M_1(\theta)$ is quasi-increasing, this purchasing rule can be implemented by setting $\theta_1^* = M_1^{-1}(0)$.

Continuing by induction, consider period t. Suppose $\theta_s^* = M_s^{-1}(0)$ for $s < t$. Fix $\{\tau(\theta, t')\}$ for $t' > t$, and consider the optimal choice of $\tau(\theta, t)$. The non-discrimination condition in Lemma 1 implies that $\tau(\theta, t) \in \{t, \tau(\theta, t+1)\}$. Splitting the profit equation $\Pi = \Pi_{\leq t-1} + \Pi_t + \Pi_{\geq t+1}$,

8. A function $M(\theta)$ is (strictly) quasi-increasing if $M(\theta_L) \geq 0 \Longrightarrow M(\theta_H) \geq (>) 0$ for $\theta_H > \theta_L$. Define the root of a quasi-increasing function by $M^{-1}(0) := \sup\{\theta : M_t(\theta) < 0\}$. If $M(\theta) < 0$ $(\forall \theta)$ then $M^{-1}(0) := \overline{\theta}$. If $M(\theta) \geq 0$ $(\forall \theta)$ then $M^{-1}(0) := \underline{\theta}$.

9. For linear demand, $f_t(\theta) = 1$ on $[0, 2b_t]$. Marginal revenue, $m_t(\theta) = 2(\theta - b_t)$, is increasing on $[0, 2b_t]$, so $M_t(\theta)$ is quasi-increasing if $\max_t b_t \leq 2 \min_t b_t$. For semi-log demand, $f_t(\theta) = \exp(-\theta/\lambda_t)/\lambda_t$. Marginal revenue, $m_t(\theta) = \exp(-\theta/\lambda_t)(\theta/\lambda_t - 1)$, is increasing on $[0, 2\lambda_t]$, so $M_t(\theta)$ is quasi-increasing if $\max_t \lambda_t \leq 2 \min_t \lambda_t$. These two examples were picked for their algebraic simplicity, but seem to be representative of many other distributions.

$$\Pi = \sum_{s=1}^{t-1} \int_{\underline{\theta}}^{\overline{\theta}} [\Delta_s M_s^+(\theta)] d\theta + \int_{\underline{\theta}}^{\overline{\theta}} [\Delta_{\tau(\theta,t)} M_t^+(\theta) + \Delta_{\tau(\theta,t)} M_t^-(\theta)] d\theta + \Pi_{\geq t+1}$$

$$\leq \sum_{s=1}^{t-1} \int_{\underline{\theta}}^{\overline{\theta}} [\Delta_s M_s^+(\theta)] d\theta + \int_{\underline{\theta}}^{\overline{\theta}} [\Delta_t M_t^+(\theta) + \Delta_{\tau(\theta,t+1)} M_t^-(\theta)] d\theta + \Pi_{\geq t+1}.$$

The optimal choice of purchasing time is $\tau^*(\theta,t) = t$ if $M_t(\theta) \geq 0$ and $\tau^*(\theta,t) = \tau(\theta,t+1)$ if $M_t(\theta) < 0$. If $M_t(\theta)$ is quasi-increasing, this purchasing rule can be implemented by setting $\theta_t^* = M_t^{-1}(0)$. ‖

In the first period, the firm can either sell to agent θ and gain $\Delta_1 m_1(\theta)$ or postpone selling to this agent and gain $\Delta_t m_1(\theta)$ supposing they eventually sell in period t. Since $\Delta_t < \Delta_1$, the firm should sell to the agent if and only if $m_1(\theta) \geq 0$. In period 2, and every subsequent period, the firm sums the marginal revenue of the new consumers and that of old agents who have yet to buy. This cumulative marginal revenue is given by $M_t(\theta) = m_t(\theta) + M_1^-(\theta)$, whereupon the firm sells to an agent with valuation θ if and only if their cumulative marginal revenue is positive, $M_t(\theta) \geq 0$.

This algorithm is completely myopic: the optimal cut-off point at period t only depends upon the consumers who have entered by time t. Consequently, the optimal cut-off at time t is independent of both the future demand and the discount rate.[10]

4.2. *Active generations*

In the Introduction we said that the key part of the problem is to work out which generations buy in which periods. At time t_H, we will say that a generation t_L is *active* if some members of generation t_L purchase in period t_H. Theorem 2 characterizes the set of active generations and uses this to provide an alternative derivation of the optimal cut-offs.

Definition 2. The upper active set is $\overline{A}(t_H) := \left\{t_L \leq t_H : \theta^*_{t_H} \leq \theta^*(t_H - 1; t_L)\right\}$.

Lemma 2. *The upper active set has the following properties:*

(a) $\overline{A}(t) = \{a, \ldots, t\}$ *for some* $a \leq t$.
(b) $\overline{A}(t) \supset \overline{A}(t-1)$ *or* $\overline{A}(t) = \{t\}$.

Proof. (a) $t \in \overline{A}(t)$ since $M_t^{-1}(0) \in [\underline{\theta}, \overline{\theta}] < \infty$. $\overline{A}(t)$ is connected since $\theta^*(t_H - 1; t_L)$ is increasing in t_L. (b) If $t' \in \overline{A}(t-1)$ then $\theta^*(t-1;t') = \theta^*(t-1;t-1)$. Therefore, if $\{t-1\} \in \overline{A}(t)$ then $t' \in \overline{A}(t)$. ‖

Lemma 2(a) says that the current generation is always active and that the set of active agents is connected. Lemma 2(b) says that once two generations are pooled, they are never separated. Define $\mathcal{A}(t) := \{\{a, \ldots, t\} : a \leq t\}$ as the collection of possible active sets at time t.

10. If WM fails, the optimal policy may no longer be myopic. To see this, suppose $M_t(\theta_L) > 0$ and $M_t(\theta_H) < 0$, for $\theta_H > \theta_L$. The firm could then serve both (θ_L, θ_H) in period t or could serve neither type. This calculation depends on whether these types would buy next period, and hence on future demand. Nevertheless, the result is quite robust: if WM holds in periods $\{1, \ldots, t\}$, then Theorem 1 will apply for the first t periods.

Theorem 2. *Suppose WM holds. Then, $M_t(\theta)$ is the lower envelope of $\{m_A(\theta): A \in \mathcal{A}(t)\}$ and the optimal cut-offs are given by*

$$\theta_t^* = \max_{A \in \mathcal{A}(t)} m_A^{-1}(0). \tag{4.1}$$

When $A = \overline{A}(t)$, this maximum is obtained. Moreover, if $M_t(\theta)$ is strictly quasi-increasing and continuous then $\overline{A}(t)$ is the maximal set in $\mathcal{A}(t)$ such that $m_A^{-1}(0) = \theta_t^*$.

Proof. Fix t and pick an arbitrary $A \in \mathcal{A}(t)$. That is, $A = \{a, \ldots, t\}$ for some $a \leq t$. By construction,

$$\sum_{s=a}^{t} m_s(\theta) = M_t(\theta) + \sum_{s=a}^{t-1} M_s^+(\theta) - M_{a-1}^-(\theta). \tag{4.2}$$

Hence, $M_t(\theta) \leq m_A(\theta)$. Cumulative marginal revenue can also be written as

$$M_t(\theta) = \sum_{s \leq t} m_s(\theta) \mathbf{1}_{\theta < \theta^*(t-1;s)} \tag{4.3}$$

so for any $\theta, \exists A \in \mathcal{A}(t)$ such that $M_t(\theta) = m_A(\theta)$. That is, $M_t(\theta) = \min\{m_A(\theta) : A \in \mathcal{A}(t)\}$.

Since $M_t(\theta) \leq m_A(\theta)$, if $M_t(\theta) \geq 0$ then $m_A(\theta) \geq 0, \forall A \in \mathcal{A}(t)$. That is, $M_t^{-1}(0) \geq \max_{A \in \mathcal{A}(t)} m_A^{-1}(0) \geq m_{\overline{A}(t)}^{-1}(0)$. To obtain the reverse inequality, equation (4.3) and the definition of $\overline{A}(t)$ imply that for small $\varepsilon > 0$, $M_t(\theta_t^* - \varepsilon) = m_{\overline{A}(t)}(\theta_t^* - \varepsilon) < 0$. Hence, $M_t^{-1}(0) \leq m_{\overline{A}(t)}^{-1}(0)$. Putting this together, $M_t^{-1}(0) = m_{\overline{A}(t)}^{-1}(0)$.

Fix t and define $A^*(t)$ to be the maximal set such that $m_{A^*(t)}^{-1}(0) = M_t^{-1}(0)$. Since $m_{\overline{A}(t)}^{-1}(0) = M_t^{-1}(0)$ it must be that $\overline{A}(t) \subset A^*(t)$. In order to obtain a contradiction, suppose $A^*(t) = \overline{A}(t) \cup B$ for some non-empty set B, where $b = \max\{t : t \in B\}$. Since $b \notin \overline{A}(t)$, $\theta_b^* < \theta_t^*$. $M_b(\theta)$ lies below $m_B(\theta)$ and is strictly quasi-increasing so there is a small $\varepsilon > 0$ such that $m_B(\theta_t^* - \varepsilon) \geq M_b(\theta_t^* - \varepsilon) > \varepsilon$. Moreover, $M_t(\theta)$ is continuous so ε can be chosen sufficiently small such that $m_{\overline{A}(t)}(\theta_t^* - \varepsilon) = M_t(\theta_t^* - \varepsilon) \in (-\varepsilon, 0)$. Hence, $m_{A^*(t)}(\theta_t^* - \varepsilon) = m_{\overline{A}(t)}(\theta_t^* - \varepsilon) + m_B(\theta_t^* - \varepsilon) > 0$, and $m_{\overline{A}(t)}^{-1}(0) > m_{A^*(t)}^{-1}(0)$, contradicting the supposition that $m_{A^*(t)}^{-1}(0) = M_t^{-1}(0)$. ||

Theorem 2 can be interpreted in two steps. First, the optimal cut-offs are determined by the marginal revenue of the active generations. This is as one would expect from the first-order condition of the firm's profit-maximizing problem, taking the set of active generations as exogenous. Second, the set of active generations $\overline{A}(t)$ maximizes the average monopoly price $m_A^{-1}(0)$. Intuitively, generation t_L is only active in period t_H if they come from a relatively high-demand generation, and therefore pulls up the average monopoly price.

The following example shows that the optimal policy can be deceptively simple.

Example 4 (Linear demand). *Suppose $f_t(\theta) = 1$ on $[0, 2b_t]$ where $b_t \in [10, 20]$.*[11] *Under Theorem 1, the monopolist sells quantity b_t each period (see Board, 2005, for a proof). This is the same as the quantity sold by a firm who could completely discriminate between different cohorts. This equivalence is analogous to the property that average quantity sold remains unaffected by third-degree price discrimination when demand is linear (Tirole, 1988, p. 139).*

11. The bounds on b_t ensure that $M_t(\theta)$ is quasi-increasing.

4.3. *Best-price provisions and time-consistent pricing*

Applying the revenue equivalence theorem, Proposition 1 shows that the optimal policy can be implemented by a best-price provision. Moreover, the best-price provision is time consistent so the firm need not commit to a sequence of prices at time 0, so long as they can promise to honour the best-price agreement.[12]

A *best-price provision* works as follows. In each period, the firm announces a price p_t^{BP}. If a consumer buys in period t and the price then falls, they are then given a rebate equal to the difference in the prices. In each subsequent period s, they are given a rebate equal to $\min\{p_t^{\mathrm{BP}},\dots,p_s^{\mathrm{BP}}\}-\min\{p_t^{\mathrm{BP}},\dots,p_{s-1}^{\mathrm{BP}}\}$. In discounted terms, the consumer purchasing in period t pays

$$\sum_{s=t}^{T}(\Delta_s-\Delta_{s+1})\min\{p_t^{\mathrm{BP}},\dots,p_s^{\mathrm{BP}}\}.$$

Proposition 1. *Suppose WM holds. Then, the firm's optimal policy under a best-price provision is to set $p_t^{\mathrm{BP}}=M_t^{-1}(0)$, inducing the same allocation and profits as Theorem 1. This policy is time consistent.*

Proof. Under a best-price provision, an agent of type (θ,t) will purchase as soon as the price falls below θ, so the set of implementable allocations is characterized by Lemma 1. Profits are given by equation (3.4) so the optimal cut-offs are given by Theorem 1 and implemented by setting $p_t^{\mathrm{BP}}=M_t^{-1}(0)$.

To prove time consistency, suppose that each period the firm chooses price p_t^{BP} to maximize its future profits. Since agents purchase as soon as the price falls below their valuation, the induced cut-off is given by $\theta_t^{\mathrm{BP}}=p_t^{\mathrm{BP}}$, and we can suppose the firm chooses the cut-offs θ_t^{BP} directly. Denoting the updated cut-off by $\theta^{\mathrm{BP}}(s;r):=\min_{s'\in\{r,\dots,s\}}\theta_{s'}^{\mathrm{BP}}$, receipts from generation r in period s equal

$$\begin{aligned}\Pi_{r,s}^{\mathrm{BP}}&=\theta^{\mathrm{BP}}(s;r)\overline{F}_r(\theta^{\mathrm{BP}}(s;r))-\theta^{\mathrm{BP}}(s-1;r)\overline{F}_r(\theta^{\mathrm{BP}}(s-1;r))\\&=\int_{\underline{\theta}}^{\overline{\theta}}\mathbf{1}_{\tau(\theta,r)\le s}m_r(\theta)d\theta-\int_{\underline{\theta}}^{\overline{\theta}}\mathbf{1}_{\tau(\theta,r)<s}m_r(\theta)d\theta\\&=\int_{\underline{\theta}}^{\overline{\theta}}\mathbf{1}_{\tau(\theta,r)=s}m_r(\theta)d\theta.\end{aligned}$$

Summing over different generations, profits from period t onwards are given by

$$\Pi_{\ge t}^{\mathrm{BP}}=\sum_{s=t}^{T}\sum_{r=1}^{s}\int_{\underline{\theta}}^{\overline{\theta}}\Delta_s\mathbf{1}_{\tau(\theta,r)=s}m_r(\theta)d\theta. \tag{4.4}$$

The firm chooses the cut-off, θ_t^{BP}, to maximize future profits (4.4). This choice also maximizes total profits, Π, since the difference

12. Butz (1990) reaches a similar conclusion in a model with declining demand and resale. When demand is allowed to increase, however, the presence of resale means that the best-price provision will generally *not* be time consistent. Intuitively, when demand in period 2 is higher than in period 1, the firm can raise production in the second period above the commitment solution without lowering p_2 below p_1, and therefore without triggering a rebate.

$$\Pi - \Pi_{\geq t}^{\mathrm{BP}} = \sum_{s=1}^{t-1} \sum_{r=1}^{s} \int_{\underline{\theta}}^{\overline{\theta}} \Delta_s \mathbf{1}_{\tau(\theta,r)=s} m_r(\theta) d\theta$$

is independent of θ_t^{BP}. ∥

5. APPLICATIONS

This section applies Theorems 1 and 2 to different paths of demand functions. We examine how cut-offs change over time, the amount of delay, and the time path of prices.

Assumption (SM). *$m_t(\theta)$ is strictly increasing and continuous in θ, and $m_t^{-1}(0) \in (\underline{\theta}, \overline{\theta})$ ($\forall t$).*

This section uses strong monotonicity (SM), which implies WM. This SM assumption simplifies proofs and helps provide a cleaner characterization of demand cycles (Section 5.2). Broadly speaking, if $m_t(\theta)$ is quasi-increasing and demand variation is not too great, then SM will hold for all relevant valuations. For linear and semi-log demand, this means that demand cannot double between the lowest and the highest periods.[13]

Lemma 3. *Suppose SM holds. Then,*

(a) $m_1^{-1}(0) > m_2^{-1}(0)$ *implies that* $m_{\{1,2\}}^{-1}(0) \in (m_2^{-1}(0), m_1^{-1}(0))$;
(b) $m_1^{-1}(0) \geq m_2^{-1}(0)$ *implies that* $m_{\{1,2\}}^{-1}(0) \in [m_2^{-1}(0), m_1^{-1}(0)]$.

Proof. Omitted. ∥

5.1. *Monotone demand: fast rises and slow falls*

Consider the two-period model, illustrated in Figure 3. In the first period, $\theta_1^* = m_1^{-1}(0)$, independent of second-period demand. In the second period, the cut-off depends on the demand path. If demand is increasing, $m_1^{-1}(0) \leq m_2^{-1}(0)$, then $M_2^{-1}(0) = m_2^{-1}(0)$. Intuitively, the firm can sell to all generation 2 agents with positive marginal revenue without having to sell to any more generation 1 agents. Further lowering the cut-off is then undesirable since the firm will sell to negative marginal revenue agents from generation 2 and, if the cut-off is sufficiently low, negative marginal revenue agents from generation 1. Conversely, if demand is decreasing, $m_1^{-1}(0) \geq m_2^{-1}(0)$, then $M_2^{-1}(0) = m_{\leq 2}^{-1}(0)$. If the firm were to sell to all the generation 2 agents with positive marginal revenue, it must also sell to some generation 1 agents with negative marginal revenue who did not buy in the first period. Thus, the firm increases the cut-off until the total marginal revenue from both generations equals zero.

This two-period example also illustrates the asymmetry of the optimal price path. Suppose demand is either high or low, $m_H(\theta)$ or $m_L(\theta)$. When demand increases, prices increase rapidly: $p_1^* = m_L^{-1}(0)$ and $p_2^* = m_H^{-1}(0)$. Conversely, when demand decreases, prices decrease slowly:

13. See footnote 9. Versions of many of these results also extend to WM. This weaker assumption is sufficient for the "if" part of the monotone demand characterization (Proposition 2). Similarly, when demand follows cycles, cut-offs are stationary (Proposition 3) and, using Theorem 2, cut-offs and prices exceed the average demand price (Proposition 4).

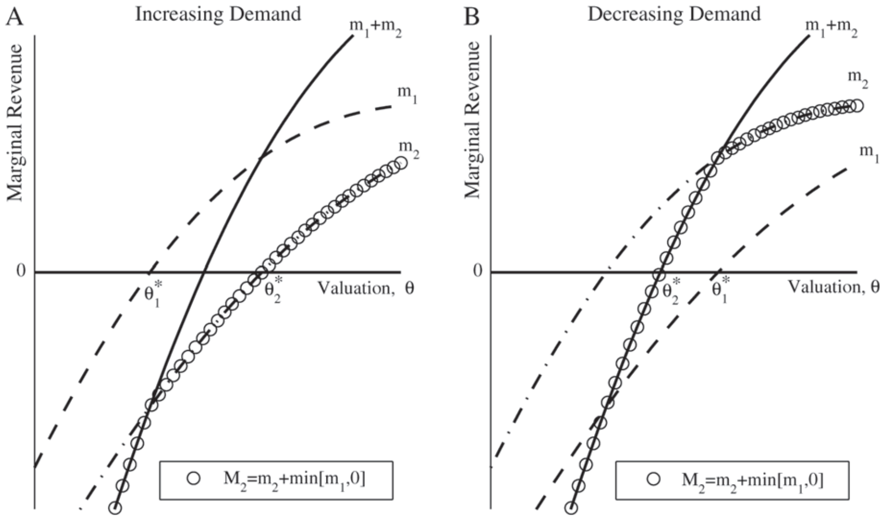

FIGURE 3

This figure shows a two-period example created using the Weibull distribution. It shows cumulative marginal revenue, equal to the lower envelope of $[m_{\{1,2\}}(\theta), m_2(\theta)]$, when demand (A) increases over time and (B) decreases over time

$p_1^* = (1-\delta)m_H^{-1}(0) + \delta m_{\{L,H\}}^{-1}(0) \leq m_H^{-1}(0)$ and $p_2^* = m_{\{L,H\}}^{-1}(0) \geq m_L^{-1}(0)$. These results can easily be extended to monotone demand paths lasting any number of periods.

Definition 3. Demand is increasing if $m_{t+1}^{-1}(0) \geq m_t^{-1}(0)$ $(\forall t)$. Demand is weakly decreasing if $m_{\leq t}^{-1}(0) \geq m_{t+1}^{-1}(0)$ $(\forall t)$.

Proposition 2. *Suppose SM holds. Optimal cut-offs are given by $\theta_t^* = m_t^{-1}(0)$ $(\forall t)$ if and only if demand is increasing. Moreover, these cut-offs can be implemented by prices*

$$p_t^* = m_t^{-1}(0).$$

Optimal cut-offs are given by $\theta_t^* = m_{\leq t}^{-1}(0)$ $(\forall t)$ if and only if demand is weakly decreasing. Moreover, these cut-offs can be implemented by prices

$$p_t^* = \sum_{s=t}^{T} \left(\frac{\Delta_s}{\Delta_t} - \frac{\Delta_{s+1}}{\Delta_t} \right) m_{\leq s}^{-1}(0). \tag{5.1}$$

Proof. We prove the result for the weakly decreasing demand case; the increasing demand case is similar. First, suppose demand is weakly decreasing. In period 1, $\theta_1^* = m_{\leq 1}^{-1}(0)$. Continuing by induction, suppose $\theta_s^* = m_{\leq s}^{-1}(0)$ for $s < t$ and consider period t. $M_t(\theta) = m_{\leq t}(\theta)$ on $[\underline{\theta}, \min\{\theta_1^*, \dots, \theta_{t-1}^*\}] = [\underline{\theta}, m_{\leq t-1}^{-1}(0)]$, using the induction hypothesis. Demand is weakly decreasing so Lemma 3 implies that $m_{\leq t}^{-1}(0) \leq m_{\leq t-1}^{-1}(0)$ and hence $M_t^{-1}(0) = m_{\leq t}^{-1}(0)$. Given these cut-offs, equation (3.7) implies that prices obey the AR(1) equation $(\theta_t^* - p_t^*) = (\theta_t^* - p_{t+1}^*)\delta_{t+1}$, yielding equation (5.1).

Next, suppose demand is not weakly decreasing. That is, $m_t^{-1}(0) > m_{\leq t-1}^{-1}(0)$ for some t. Theorem 2 means that $\theta_t^* \geq m_t^{-1}(0) > m_{\leq t}^{-1}(0)$, using Lemma 3. ‖

Two price paths, mentioned in the Introduction, will serve as useful benchmarks.

Definition 4. The myopic price is $p_t^M := m_t^{-1}(0)$. The average demand price is $p^A := \lim_{t\to T} m_{\leq t}^{-1}(0)$, assuming the limit exists.[14]

The myopic price is the price charged by a monopolist who only takes the current generation of consumers into account, ignoring the previous ones. The average demand price is charged by a monopolist who faces average demand $\frac{1}{T}\sum_{t=1}^{T} F_t(\theta)$ each period.

When demand is increasing, the price equals the myopic price in all periods. Conversely, when demand is decreasing, the price starts off below the myopic price and ends up above it. In fact, prices fall so slowly that they converge to the average demand price from above. That is,

$$\lim_{t\to T}[p_t^* - p^A] = \lim_{t\to T}\sum_{s=t}^{T}\left[\left(\frac{\Delta_s}{\Delta_t} - \frac{\Delta_{s+1}}{\Delta_t}\right)(m_{\leq s}^{-1}(0) - m_{\leq t}^{-1}(0))\right] = 0$$

since every convergent sequence is Cauchy.

5.2. *Demand cycles*

This subsection examines the optimal price path when demand follows deterministic cycles. The sequence of demand functions is described by K repetitions of $\{F_1(\theta), \ldots, F_T(\theta)\}$, where we denote period t of cycle k by $t_{[k]}$. An example of this was seen in Figure 1 in the Introduction, which also illustrates the lower active set, $\underline{A}(t)$, as defined below. One can see the pattern of sharp price increases and slow declines; these intuitively follow from Proposition 2. The picture also illustrates other regularities that occur after the first cycle:

1. Cut-offs and prices follow a stationary pattern.
2. Cut-offs and prices always lie above the average demand price.
3. The lowest cut-off and price occur in the last period of the slump.

Propositions 3–5 correspond to these results. First, it will useful to define the lower active set.

Definition 5. The lower active set is $\underline{A}(t_H) := \left\{t_L \leq t_H : \theta_{t_H}^* < \theta^*(t_H - 1; t_L)\right\}$.

The lower active set will often look very similar to the upper active set but is particularly useful in the analysis of demand cycles. Lemma 2 applies to the lower active set, as does a version of Theorem 2.

Theorem 2′. *Suppose SM holds. Then, the optimal cut-offs are given by*

$$\theta_t^* = \max_{A\in\mathcal{A}(t)} m_A^{-1}(0). \tag{5.2}$$

When $A = \underline{A}(t)$, this maximum is obtained. Moreover, $\underline{A}(t)$ is the minimal set in $\mathcal{A}(t)$ such that $m_A^{-1}(0) = \theta_t^*$.

14. Assume SM holds. The limit exists if demand is weakly decreasing, demand follows cycles, or T is finite.

Proof. Equation (5.2) is a restatement of (4.1). Using (4.3) and SM, $M_t(\theta_t^*) = m_{\underline{A}(t)}(\theta_t^*) = 0$. By SM, this uniquely defines $\theta_t^* = m_{\underline{A}(t)}^{-1}(0)$.

Fix t and let $A^*(t)$ be the minimal set such that $m_{A^*(t)}^{-1}(0) = \theta_t^*$. In order to attain a contradiction, suppose $\underline{A} = A^* \cup B$ for some non-empty set B, where $b = \max\{t{:}t \in B\}$. Since $b \in \underline{A}(t)$, we have $\theta_b^* > \theta_t^*$ and $M_b^-(\theta_t^*) < 0$. Using equation (4.2), $0 = M_t(\theta_t^*) \leq m_{A^*(t)}(\theta_t^*) + M_b^-(\theta_t^*) < 0$, yielding a contradiction. ‖

Proposition 3. *Suppose demand follows deterministic cycles and SM holds. Then* $|\underline{A}(t)| \leq T$. *Hence, if* $k \geq 2$, *the cycles are stationary,* $\theta^*_{t_{[k]}} = \theta^*_{t_{[2]}}$.

Proof. Suppose $|\underline{A}(t)| > T$. Define the set A such that $\underline{A}(t) = A \cup B$ where B consists of the union of $\{1, \ldots, T\}$ sets and $|A| \leq T$. Then $m_{\underline{A}(t)}^{-1}(0) \leq \max\{m_A^{-1}(0), m_{\{1,\ldots,T\}}^{-1}(0)\}$ by Lemma 3, contradicting the fact that $\underline{A}(t)$ is the smallest set to achieve the maximum in Theorem 2′. ‖

Proposition 3 means that the cut-offs will be the same for each cycle $k \geq 2$. This substantially simplifies analysis: when $k \geq 2$, we can use modular arithmetic to write the collection of possible active sets at time t as $\mathcal{A}_T(t) = \{\{a, \ldots, t\} : a \in \{1, \ldots, T\}\}$. The cut-off and price are minimized at time

$$\underline{t} := \min\{\operatorname{argmin}_{t \in \{1,\ldots,T\}} \{\theta^*_{t_{[k]}} : k \geq 2\}\}.$$

A consumer who does not buy in period $\underline{t}$ will never buy. Hence, the market effectively resets at period $\underline{t}$, enabling us to throw out all previous generations. This means that the starting position of the cycle only matters for the first $\underline{t}$ periods, at which point the stationary cycles start.[15]

Prices are determined by equation (3.7). For cycles $k \in \{2, \ldots, K-1\}$, the prices are stationary with boundary condition $p^*_{\underline{t}} = \theta^*_{\underline{t}} = m^{-1}_{\{1,\ldots,T\}}(0)$, using Proposition 4(a). For the final cycle, where $k = K$ and $t > \underline{t}$, the relevant boundary condition is $p_T^* = \theta_T^*$. Since there is not as much scope for delay, prices in the last cycle may be higher than in previous cycles.

Proposition 4. *Suppose demand follows cycles with* $k \geq 2$ *and SM holds. Then,*

(a) *The lowest optimal cut-off equals the average demand price. Hence, agents buy later under optimal pricing.*
(b) *The lowest optimal price equals the average demand price.*
(c) *When compared to average demand pricing, the optimal price path induces higher profits, lower welfare for every generation, and lower utility for every type* (θ, t).

Proof.

(a) If $k \geq 2$ then $\{1, \ldots, T\} \in \mathcal{A}(t)$ and Theorem 2 implies that $\theta_t^* \geq m^{-1}_{\{1,\ldots,T\}}(0)$. Equation (4.2) implies

$$m_{\{1,\ldots,T\}}(\theta^*_{\underline{t}}) = M_{\underline{t}_{[k]}}(\theta^*_{\underline{t}}) + \sum_{s=\underline{t}_{[k-1]}+1}^{\underline{t}_{[k]}-1} M_s^+(\theta^*_{\underline{t}}) - M^-_{\underline{t}_{[k-1]}}(\theta^*_{\underline{t}}) = 0$$

using (1) the definition $\theta^*_{\underline{t}}$, (2) the fact that $\theta_s^* \geq \theta^*_{\underline{t}_{[k]}}$ for $s \in \{\underline{t}_{[k-1]} + 1, \ldots, \underline{t}_{[k]} - 1\}$, and (3) $\theta^*_{\underline{t}_{[k]}} \geq \theta^*_{\underline{t}_{[k-1]}}$. Thus, $\theta^*_{\underline{t}} = m^{-1}_{\{1,\ldots,T\}}(0) = p_t^A$.

15. Formally, for any $t_{[1]} > \underline{t}_{[1]}, \theta^*_{t_{[1]}} = \theta^*_{t_{[2]}}$. To see this observe that $\underline{A}(t_{[2]}) = \{a_{[2]}, \ldots, t_{[2]}\}$, where $a_{[2]} > \underline{t}_{[2]}$. Since $\{a_{[1]}, \ldots, t_{[1]}\} \in \mathcal{A}(t_{[1]})$, Theorem 2 implies that $\theta^*_{t_{[1]}} \geq \theta^*_{t_{[2]}}$. Conversely, $\mathcal{A}(t_{[1]}) \subset \mathcal{A}(t_{[2]})$, so $\theta^*_{t_{[2]}} \geq \theta^*_{t_{[1]}}$.

(b) Using part (a), $p_{\underline{t}}^* = \theta_{\underline{t}}^* = p_t^A$.
(c) Profit is higher under the optimal scheme by revealed preference. The utility of any customer (3.1) and welfare of any generation (3.3) are lower under the optimal scheme since the cut-offs are higher. ‖

Cut-offs and prices are always higher than if the monopolist faced average demand, while welfare and consumer surplus are always lower. That is, all consumers are made worse off by the ability of the monopolist to discriminate between generations. Intuitively, during high-demand periods, only high-demand generations are active and the price is high. However, during low-demand periods, both high and low generations are active and the price is middling. By delaying, the high-demand cohorts thus exert a negative externality on the low cohorts, raising prices.

With quasi-linear utility, the indirect utility function is convex in prices, suggesting that price variation benefits consumers because of the option value. However, Proposition 4 shows that when prices are endogenous, price variation may hurt all consumers. The result can also be contrasted with the standard view that the welfare effect of third-degree price discrimination is indeterminate (*e.g.* Tirole, 1988, p. 137). Inter- and intratemporal price discrimination clearly have quite different properties.

Define a cycle as *simple* if $m_t^{-1}(0) - m_{t-1}^{-1}(0)$ is nonzero and has at most two changes of sign, where $m_0(\theta) := m_T(\theta)$. That is, each cycle has one "boom" and one "slump".

Proposition 5. *Suppose demand follows cycles with $k \geq 2$ and SM holds. Then $\underline{t}$ obeys*

$$m_{\underline{t}}^{-1}(0) \leq m_{\{1,\ldots,T\}}^{-1}(0) \leq m_{\underline{t}+1}^{-1}(0). \tag{5.3}$$

When the cycle is simple, equation (5.3) uniquely defines $\underline{t}$.

Proof. First, Proposition 4(a) and Theorem 2 imply $m_{\{1,\ldots,T\}}^{-1}(0) = \theta_{\underline{t}}^* \geq m_{\underline{t}}^{-1}(0)$. Second, by the definition of $\underline{t}$, $\underline{A}(\underline{t}+1) = \{\underline{t}+1\}$. Hence, Theorem 2′ implies that $m_{\underline{t}+1}^{-1}(0) = \theta_{\underline{t}+1}^* \geq m_{\{1,\ldots,T\}}^{-1}(0)$. ‖

In a simple cycle, prices are minimized in the last period of the slump, just before demand returns to its long-run average. In literary terms: The darkest hour is just before dawn.

6. STOCHASTIC DEMAND

In this section, we extend the results of Section 4 to stochastic demand and discounting. We then apply these results to i.i.d. demand draws. The model is as follows.

Information. Demand and the discount rate are allowed to be uncertain, depending on the state of the world $\omega \in \Omega$. The information possessed by consumers and the firm is described by a filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, Q)$, where $\mathcal{F}$ are the measurable sets, $\mathcal{F}_t$ is the information partition at time t, which grows finer over time, and Q is the probability measure. The discount rate δ_t and incoming demand $f_t(\theta)$ are $\mathcal{F}_t$ adapted, so consumers and the firm know δ_t and $f_t(\theta)$ at time t.[16]

Consumers' problem. A consumer of type (θ, t) faces a sequence of $\mathcal{F}_t$-adapted prices $\{p_t\}$. They then have the problem of choosing a purchasing time $\tau(\theta, t) : \Omega \to \{t, \ldots, T, \infty\}$ to maximize expected utility,

16. Note, δ_t is $\mathcal{F}_t$ adapted if $\{\omega : \delta_t \leq x\} \in \mathcal{F}_t$ for all $x \in (0, 1)$.

$$u_t(\theta) = \sup_{\tau} E[(\theta - p_\tau)\Delta_\tau], \tag{6.1}$$

where "E" is the expectation over Ω. The purchasing time is a random variable taking values in $\{t, \dots, T, \infty\}$, where the decision to buy at time s depends on information available at time s, that is, $\{\omega : \tau(\theta, t) \le s\} \in \mathcal{F}_s$ $(\forall s)$. Again, let $\tau^*(\theta, t)$ be the earliest solution to this problem.

Firm's problem. The firm chooses $\mathcal{F}_t$-adapted prices $\{p_t\}$ to maximize profit,

$$\Pi = E\left[\sum_{t=1}^{T} \int_{\underline{\theta}}^{\overline{\theta}} \Delta_{\tau^*(\theta,t)} p_{\tau^*(\theta,t)} dF_t\right]$$

where $\tau^*(\theta, t)$ maximizes the consumer's utility (6.1).

6.1. *Optimal pricing*

We first analyse the consumer's problem, as in Section 3.3. The earliest purchasing rule exists and is unique up to its equivalence class (Klass, 1973, Theorem 6). As in Lemma 1, there exists a version of this optimal purchasing time that obeys "θ-monotonicity" and "non-discrimination", and can therefore be characterized by a sequence of $\mathcal{F}_t$-adapted cut-offs, $\{\theta_t^*\}$.

Assuming that $M_t(\theta)$ is quasi-increasing on the realized demand path, the optimal cut-offs are characterized by Theorems 1 and 2 (the proofs immediately extend). The myopic nature of the optimal policy means that uncertainty about future demand has no effect on the optimal allocation. This implies, for example, that if realized demand is monotone, then the optimal allocations are characterized by Proposition 2.

Given a set of cut-offs, prices can then be derived by backwards induction:

$$\Delta_t(\theta_t^* - p_t^*) = \max_{\tau \ge t+1} E[(\theta_t^* - p_\tau^*)\Delta_\tau \mid \mathcal{F}_t].$$

While cut-offs only depend on realized demand, prices depend on the cut-offs in all future states and thus on all possible demand realizations.[17]

6.2. *i.i.d. demand*

To illustrate how to apply the above results, suppose incoming cohorts are i.i.d. Each period, demand is drawn from $\{m_x(\theta)\}_x$ with probability measure $\mu(x)$ and support $x \in [0, 1]$, where higher indices imply higher demand: $m_{x_H}^{-1}(0) \ge m_{x_L}^{-1}(0)$, for $x_H \ge x_L$. Also suppose that SM holds $(\forall x)$.

In this setting, we can define the *average demand price* by $p^A := \left[\int m_x(\theta) d\mu(x)\right]^{-1}(0)$. This is the optimal strategy for a monopolist who must charge a constant price. Example 5 explicitly derives the optimal price schedule in a two-period model.

17. We have assumed that consumers and the firm possess the same information about future demand. However, the results are quite robust: due to the myopic nature of the optimal policy, the firm would still like to implement the cut-offs in Theorem 1 even if consumers have different information. If the firm is more informed, these cut-offs can be implemented by prices that are adapted to the consumers' information sets. If the firm is less informed, then there may be no prices adapted to the firm's information set that implement the optimal cut-offs. However, in this case, the firm should be able to extract consumers' information by playing off different agents against each other.

Example 5 ($T = 2$ with i.i.d. demand). *Denote the demand in the first and second period by* $m_1(\theta) = m_{x_1}(\theta)$ *and* $m_2(\theta) = m_{x_2}(\theta)$, *respectively, and suppose the discount rate is constant. Theorem 2 implies that* $\theta_1^* = m_{x_1}^{-1}(0)$ *and* $\theta_2^* = \max\left\{m_{\{x_1,x_2\}}^{-1}(0), m_{x_2}^{-1}(0)\right\}$. *The second-period price is* $p_2^* = \theta_2^*$, *while the initial price is given by*

$$p_1^* = \left(1 - \delta \int_0^{x_1} d\mu(x)\right) m_{x_1}^{-1}(0) + \delta \int_0^{x_1} m_{\{x_1,x\}}^{-1}(0) d\mu(x).$$

If demand of the first generation is low, then demand is likely to increase and p_1^* *is determined by the first-period demand,* $m_{x_1}^{-1}(0)$. *If demand of the first generation is high, then demand is likely to decrease and* p_1^* *is determined by the pooled demand,* $m_{\{x_1,x_2\}}^{-1}(0)$. *The extreme cases when* $x_1 \in \{0,1\}$ *are directly analogous to the monotone demand results in Proposition 2. The general principle is that* p_1^* *is affected by a state of the world if and only if the first generation is active in that state.*

When there are more periods, it becomes harder to solve the optimal stopping problem required to back-out prices. However, one can make comparisons as the number of periods grows large. This next result is the stochastic analogue to Proposition 4.

Proposition 6. *Suppose demand is i.i.d.,* $\{m_x(\theta)\}_x$ *are uniformly bounded, and SM holds* $(\forall x)$. *Then* $\liminf_{t\to\infty} \theta_t^* = p^A$ *and* $\liminf_{t\to\infty} p_t^* = p^A$ *almost surely.*

Proof. First, we show that $\liminf_{t\to\infty} \theta_t^* \geq p^A$ almost surely. Define $g_t(\theta) = \frac{1}{t}\sum_{s=1}^{t} m_s(\theta)$ and $g(\theta) = \int m_x(\theta) d\mu(x)$. The strong law of large numbers implies that $g_t(\theta) \overset{a.s.}{\to} g(\theta)$ pointwise $\forall\theta \in [\underline{\theta}, \overline{\theta}]$. The collection $\{m_x(\theta)\}_x$ is uniformly bounded and increasing so the Glivenko–Cantelli Theorem (Davidson, 1994, Theorem 21.5) states that $\sup_\theta |g_t(\theta) - g(\theta)| \overset{a.s.}{\to} 0$. Since $g(\theta)$ is strictly increasing, $m_{\{1,\ldots,t\}}^{-1}(0) = g_t^{-1}(0) \overset{a.s.}{\to} g^{-1}(0) = p^A$. Finally, $\{1,\ldots,t\} \in \mathcal{A}(t)$ so Theorem 2 implies that $\liminf_{t\to\infty} \theta_t^* \geq p^A$ a.s.

Next, we show that $\liminf_{t\to\infty} \theta_t^* \leq p^A$ almost surely. Suppose the cumulative marginal revenue evaluated at p^A exceeds zero at times $\{t_n\}$, so that $t_0 := 0$ and $t_n := \min\{m > t_{n-1} : M_t(p^A) \geq 0\}$. Similarly, let $s_0 := 0$ and $s_n := \min\{m > t_{n-1} : m_{\{1,\ldots,t\}}(p^A) \geq 0\}$. Since $m_{\{1,\ldots,t\}}(p^A)$ is a random walk with mean zero, $\Pr(s_1 < \infty) = 1$ (Durrett, 1995, Theorems 3.2.2 and 3.2.7). By definition of $M_t(\cdot)$, $t_1 = s_1$ and hence $\Pr(t_1 < \infty) = 1$. Thus, $\Pr(t_n < \infty) = \Pr(t_1 < \infty)^n = 1$ for all n, and $\Pr(\theta_t^* \leq p^A \text{ i.o.}) = \Pr(M_t(p^A) \geq 0 \text{ i.o.}) = 1$, as required. Putting this together, $\Pr(\liminf \theta_t^* = p^A) = 1$ and hence $\Pr(\liminf p_t^* = p^A) = 1$. ‖

In the long run, the price should be at least as high as the average demand price. Moreover, the price will often be strictly higher (*e.g.* when $m_t^{-1}(0) > p^A$). This means that prices will tend to increase when there is more demand variation or when the firm has more information about incoming demand.

7. RESALE AND RENTING

So far it has been assumed that there is no resale. This is reasonable if the good is a one-time experience or is associated with high transaction costs. This section analyses the opposite case: perfect resale.

Each period t, the consumer obtains discounted rental utility $(\Delta_t - \Delta_{t+1})\theta$, where $\Delta_{T+1} = 0$. As before, demand $f_t(\theta)$ enters the market each period, where the induced marginal revenue function satisfies SM. We also suppose $f_t(\theta)$ and Δ_t are deterministic, although the results immediately extend to the stochastic case. Consider two alternative policies:

1. The firm rents the good at price R_t each period.
2. The firm sells the good, committing to a price schedule $\{p_t^R\}$, where perfect resale is possible and the seller is allowed to buy goods back.

Coase (1972) and Bulow (1982) show that these policies induce the same profits and utilities.

Theorem 3. *Suppose SM holds and the monopolist either rents the good or they sell the good and allow resale. Then the profit-maximizing cut-offs are given by* $\theta_t^R = m_{\leq t}^{-1}(0)$.[18]

Proof. Both policies induce allocations of the form $\{\theta : \theta \geq \theta_t^R\}$ for some $\theta_t^R \in [\underline{\theta}, \overline{\theta}]$. When renting (policy 1), the time-t profit is given by

$$\int_{\underline{\theta}}^{\overline{\theta}} (\Delta_t - \Delta_{t+1})\mathbf{1}_{\theta \geq \theta_t^R} m_{\leq t}(\theta) d\theta.$$

The firm would therefore like to allocate the good to $\{\theta : m_{\leq t}(\theta) \geq 0\}$. Under SM, this is implemented by setting $\theta_t^R = m_{\leq t}^{-1}(0)$. When the firm commits to a sequence of sale prices and allows resale (policy 2), the approach in Section 3.1 can be used to derive the total profit,

$$\Pi^R = \sum_{t=1}^{T} \int_{\underline{\theta}}^{\overline{\theta}} (\Delta_t - \Delta_{t+1})\mathbf{1}_{\theta \geq \theta_t^R} m_{\leq t}(\theta) d\theta.$$

Again, the optimal policy is to set $\theta_t^R = m_{\leq t}^{-1}(0)$. ‖

The rental price under the optimal strategy is $R_t = (1 - \delta_{t+1}) m_{\leq t}^{-1}(0)$. The optimal price path with resale, p_t^R, can be derived from the AR(1) system

$$(\theta_t^R - p_t^R) = (\theta_t^R - p_{t+1}^R)\delta_{t+1}.$$

The resale price is thus the geometric sum of future rental values, as given by equation (5.1). Hence, the resale price converges to the average demand price, $p^A := \lim_{t \to T} m_{\leq t}^{-1}(0)$, if this limit exists. Intuitively, after enough time, the resale market grows very large and the firm loses the ability to discriminate.

7.1. *The effect of resale*

Figure 4 can be used to assess the effect of resale. With resale, troughs and peaks are treated symmetrically, and the cycles decrease in amplitude as the size of the resale market engulfs new production. In the limit, the resale price converges to the average demand price. Without resale, price cycles are highly asymmetric and are stationary. The reason for this difference is that with resale, a low-valuation consumer may buy if they anticipate the price to rise, so all previous generations remain active. Without resale, only the high-demand generations remain active.

Proposition 7. *Suppose SM holds.*

18. When SM fails, the firm should simply iron $m_{\leq t}(\theta)$ $(\forall t)$.

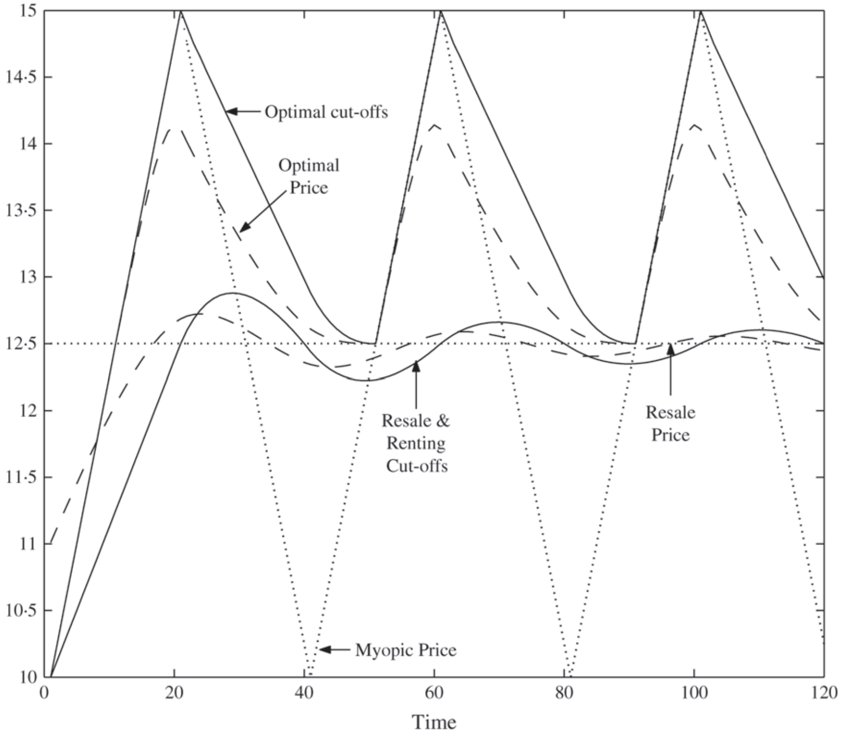

FIGURE 4

This figure shows how resale affects prices and allocations when demand follows cycles. It assumes that valuations have measure 1 on $[0, b_t]$, $b_t \in [20, 30]$, and the discount factor is $\delta = 0{\cdot}9$

(a) *Cut-offs and profits are lower with resale than without resale.*
(b) *Cut-offs and profits are unaffected by resale if and only if demand is weakly decreasing.*

Proof.

(a) Since $\{1, \dots, t\} \in \mathcal{A}(t)$, $\theta_t^R \le \theta_t^*$ by Theorem 2. Using Theorem 1, profits without resale are

$$\Pi = \int_{\underline{\theta}}^{\overline{\theta}} \sum_{s=1}^{T} \Delta_s M_s^+(\theta) d\theta = \int_{\underline{\theta}}^{\overline{\theta}} \sum_{t=1}^{T} (\Delta_t - \Delta_{t+1}) \sum_{s=1}^{t} M_s^+(\theta) d\theta$$

noting that $\Delta_s = \sum_{t=s}^{T} (\Delta_t - \Delta_{t+1})$ and changing the order of summation. Profits with resale are

$$\Pi^R = \int_{\underline{\theta}}^{\overline{\theta}} \sum_{t=1}^{T} (\Delta_t - \Delta_{t+1}) m_{\le t}^+(\theta) d\theta$$

$$= \int_{\underline{\theta}}^{\overline{\theta}} \sum_{t=1}^{T} (\Delta_t - \Delta_{t+1}) \left[\max \left\{ 0, M_t(\theta) + \sum_{s=1}^{t-1} M_s^+(\theta) \right\} \right] d\theta$$

$$\leq \int_{\underline{\theta}}^{\overline{\theta}} \sum_{t=1}^{T} (\Delta_t - \Delta_{t+1}) \left[\max\{0, M_t(\theta)\} + \sum_{s=1}^{t-1} \max\{0, M_s^+(\theta)\} \right] d\theta$$

$$= \int_{\underline{\theta}}^{\overline{\theta}} \sum_{t=1}^{T} (\Delta_t - \Delta_{t+1}) \left[\sum_{s=1}^{t} M_s^+(\theta) \right] d\theta = \Pi,$$

where the second line uses equation (4.2) and the third uses Jensen's inequality.

(b) When demand is weakly decreasing, allocations and profits are identical, by Proposition 2. When demand is not weakly decreasing, there exists a t such that $m_t^{-1}(0) > m_{\leq t-1}^{-1}(0)$. Theorem 2 means that $M_t^{-1}(0) = \theta_t^* \geq m_t^{-1}(0) > m_{\leq t}^{-1}(0)$, using Lemma 3. Thus, $M_t(\theta) < 0$ and $m_{\leq t}(\theta) \geq 0$ for $\theta \in [m_{\leq t}^{-1}(0), M_t^{-1}(0))$ and the inequality in (a) is strict. ||

With resale, Coase (1972) and Bulow (1982) argue that renting achieves the same profits as selling, as shown in Theorem 3. Without resale, Proposition 7 implies that selling strictly outperforms renting, unless demand is weakly decreasing. Consequently, when there is demand variation, the firm has an incentive to shut down any secondary market.

From a welfare perspective, the effect of resale is generally ambiguous. However, Example 6 shows that, with linear demand, resale improves welfare.

Example 6 (Linear demand). *Suppose* $f_t(\theta) = 1$ *on* $[0, 2b_t]$, *where* $b_t \in [10, 20]$. *With resale, the optimal quantity sold in period* t *is* b_t, *which is the same as without resale (Example 4). Resale increases allocative efficiency and hence welfare (see Board, 2005, for a proof).*

With linear demand curves, resale increases welfare, reduces profits, and therefore increases consumer surplus. Yet it is not the case that all consumers are better off. To see this, consider a two-period model with increasing demand. The first (low-demand) cohort prefers no resale since they then avoid being pooled with the second (high-demand) cohort. For the same reason, the second generation prefers resale. However, if demand follows cycles, as in Section 5.2, all consumers will prefer resale, if they are born sufficiently late. With resale, prices converge to the average demand price; without resale, prices always exceed the average demand price (Proposition 4).

8. CONCLUSION

This paper has studied a durable-goods monopolist facing fluctuations in incoming demand. Such fluctuations occur in many markets. When launching a new product, demand is uncertain, with the tastes of laggards unlike those of early adopters. With seasonal goods, the preferences of transitory customers are different from those of repeat buyers. And over the longer term, demand changes as a function of demographics and inter-generational tastes.

In this paper, we derived a technique to solve the firm's two-dimensional mechanism design problem, characterizing the optimal allocations by a simple myopic algorithm. We then used this solution to examine how different demand paths are propagated by the firm's pricing strategy. When new demand grows stronger, the price rises quickly, unhindered by the presence of old consumers. In contrast, when new demand becomes weaker, the price falls slowly as consumers delay their purchases. This asymmetry between rises and falls leads to an increase in the price level, reducing welfare below that under average demand pricing.

The model has made a number of simplifying assumptions, abstracting from other aspects of the firm's decision problem. First, we assumed that demand is exogenously given. In practice, the arrival of new customers can be affected by marketing and the timing of new product launches. Second, while we allowed the firm to learn about the process-generating demand, we assumed that it always knows current demand. This excluded the possibility that the firm may learn about demand through its pricing policy. Third, we assumed that all consumers are equally good at arbitraging price fluctuations. In general, one would like to design pricing policies that perform well when some customers are less able to time their purchases. Hopefully, future research will address some of these issues.

Acknowledgements. I am indebted to my advisor, Jeremy Bulow, for his guidance and enthusiasm, and to Richard Saouma for initiating my interest in sales. Excellent suggestions by two referees and the editor, Juuso Välimäki, also substantially improved the paper. I have received helpful comments from David Ahn, David Baron, Willie Fuchs, Ilan Kremer, John McMillan, Robert McMillan, Barry Nalebuff, Ilya Segal, Andy Skrzypacz, and Jeff Zwiebel. I also wish to thank the seminar audiences at the following universities: Cambridge, Essex, London School of Economics, Northwestern, New York, Oxford, Penn, Rochester, California San Diego, Stanford, Texas at Austin, Toronto, and Yale.

REFERENCES

BIEHL, A. R. (2001), “Durable-Goods Monopoly with Stochastic Values”, *RAND Journal of Economics*, **32** (3), 565–577.

BOARD, S. (2005), “Durable-Goods Monopoly with Varying Demand” (Working paper, University of Toronto; http://www.chass.utoronto.ca/∼sboard/research.html).

BULOW, J. (1982) “Durable-Goods Monopolists”, *Journal of Political Economy*, **90** (2), 314–332.

BUTZ, D. A. (1990), “Durable-Good Monopoly and Best-Price Provisions”, *American Economic Review*, **80** (5), 1062–1076.

CHOW, Y. S., ROBBINS, H. and SIEGMUND, D. (1971), *Great Expectations: The Theory of Optimal Stopping* (Boston, MA: Houghton Mifflin).

COASE, R. H. (1972), “Durability and Monopoly”, *Journal of Law & Economics*, **15** (1), 143–149.

CONLISK, J. (1984), “A Peculiar Example of Temporal Price Discrimination”, *Economics Letters*, **15** (1–2), 121–126.

CONLISK, J., GERSTNER, E. and SOBEL, J. (1984), “Cyclic Pricing by a Durable Goods Monopolist”, *Quarterly Journal of Economics*, **99** (3), 489–505.

DAVIDSON, J. (1994), *Stochastic Limit Theory* (Oxford: Oxford University Press).

DUDINE, P., HENDEL, I. and LIZZERI, A. (2006), “Storable Good Monopoly: The Role of Commitment”, *American Economic Review*, **96** (5), 1706–1719.

DURRETT, R. (1995), *Probability: Theory and Examples* (Belmont, CA: Duxbury Press).

GORDON, B. (2007), “A Dynamic Model of Consumer Replacement Cycles in the PC Precessor Industry” (Working paper, Columbia University; http://www.columbia.edu/∼brg2114/research.html).

HENDEL, I. and LIZZERI, A. (1999), “Adverse Selection in Durable Goods Markets”, *American Economic Review*, **89** (5), 1097–1115.

KLASS, M. (1973), “Properties of Optimal Extended-Valued Stopping Rules for S_n/n”, *Annals of Probability*, **1** (5), 719–757.

KRISHNAN, T. V., BASS, F. M. and JAIN, D. C. (1999), “Optimal Pricing Strategy for New Products”, *Management Science*, **45** (12), 1650–1663.

LACH, S. (2007), “Immigration and Prices”, *Journal of Political Economy*, **115** (4), 548–587.

LANDSBERGER, M. and MEILIJSON, I. (1985), “Intertemporal Price Discrimination and Sales Strategy under Incomplete Information”, *RAND Journal of Economics*, **16** (3), 424–430.

MILGROM, P. and SEGAL, I. (2002), “Envelope Theorems for Arbitrary Choice Sets”, *Econometrica*, **70** (2), 583–601.

NEVO, A. and HATZITASKOS, K. (2005), “Why does the Average Price of Tuna Fall During Lent?” (Working paper, NBER, No. 11572).

ROGERS, E. M. (1962), *Diffusion of Innovations* (New York: The Free Press).

ROTEMBERG, J. J. and SALONER, G. (1986), “A Supergame-Theoretic Model of Price Wars during Booms”, *American Economic Review*, **76** (3), 390–407.

RUST, J. (1986), “When is it Optimal to Kill off the Market for Used Durable Goods?”, *Econometrica*, **54** (1), 65–86.

SEGAL, I. (2003), “Awarding Monopoly Franchises”, *American Economic Review*, **93** (3), 509–529.

SOBEL, J. (1991), “Durable Goods Monopoly with Entry of New Consumers”, *Econometrica*, **59** (5), 1455–1485.

STOKEY, N. L. (1979), “Intertemporal Price Discrimination”, *Quarterly Journal of Economics*, **93** (3), 355–371.

TIROLE, J. (1988), *The Theory of Industrial Organization* (Cambridge, MA: MIT Press).

TOPKIS, D. M. (1998), *Supermodularity and Complimentarity* (Princeton, NJ: Princeton University Press).

VARIAN, H. R. (1980), “A Model of Sales”, *American Economic Review*, **70** (4), 651–659.

WALDMAN, M. (1996), “Durable Goods Pricing When Quality Matters”, *Journal of Business*, **69** (4), 489–510.

WARNER, E. J. and BARSKY, R. B. (1995), “The Timing and Magnitude of Retail Store Markdowns: Evidence from Weekends and Holidays”, *Quarterly Journal of Economics*, **110** (2), 321–352.