MANAGEMENT SCIENCE
Vol. 61, No. 1, January 2015, pp. 92–110
ISSN 0025-1909 (print) | ISSN 1526-5501 (online)
http://dx.doi.org/10.1287/mnsc.2014.2049
© 2015 INFORMS

Intertemporal Price Discrimination: Structure and Computation of Optimal Policies

Omar Besbes
Graduate School of Business, Columbia University, New York, New York 10027, ob2105@columbia.edu
Ilan Lobel
Stern School of Business, New York University, New York, New York 10012, ilobel@stern.nyu.edu

We study a firm's optimal pricing policy under commitment. The firm's objective is to maximize its long-term average revenue given a steady arrival of strategic customers. In particular, customers arrive over time, are strategic in timing their purchases, and are heterogeneous along two dimensions: their valuation for the firm's product and their willingness to wait before purchasing or leaving. The customers' patience and valuation may be correlated in an arbitrary fashion. For this general formulation, we prove that the firm may restrict attention to cyclic pricing policies, which have length, at most, twice the maximum willingness to wait of the customer population. To efficiently compute optimal policies, we develop a dynamic programming approach that uses a novel state space that is general, capable of handling arbitrary problem primitives, and that generalizes to finite horizon problems with nonstationary parameters. We analyze the class of monotone pricing policies and establish their suboptimality in general. Optimal policies are, in a typical scenario, characterized by nested sales, where the firm offers partial discounts throughout each cycle, offers a significant discount halfway through the cycle, and holds its largest discount at the end of the cycle. We further establish a form of equivalence between the problem of pricing for a stream of heterogeneous strategic customers and pricing for a pool of heterogeneous customers who may stockpile units of the product.

Keywords: pricing; optimization; intertemporal pricing; price commitment; price discrimination; dynamic pricing; strategic customers; stockpiling
History: Received August 7, 2012, accepted August 4, 2014, by Yossi Aviv, operations management.

1. Introduction

Dynamic pricing is widely used in practice by firms in a variety of industries, ranging from airlines and hotels to supermarkets and clothing outlets (Talluri and van Ryzin 2005). The drivers for dynamic pricing are multiple and range from the need to adjust prices to reflect the opportunity cost associated with scarce capacity to the stochastic nature of the demand environment and to the lack of information about the underlying demand. In such cases, prices are stochastic and difficult to predict for consumers. However, in various practical settings, dynamic pricing policies are highly predictable. For example, supermarkets frequently offer discounts in a predetermined fashion, even for categories with very stable demand (e.g., shampoo or coffee). Retail outlets often discount items during holiday weekends. Why do retail stores use such predictable discounting strategies? It does not appear to be designed to liquidate inventory, since stores typically increase order sizes to suppliers in anticipation of increased demand. A more likely explanation is that firms are engaging in a form of intertemporal price discrimination to capture surplus from a heterogeneous customer base. For example, customers who have high valuation and low patience buy when the need for a given product arises, whereas those who have a lower valuation but are more patient may wait until a promotional price is offered. A pricing policy meant to extract the maximum revenues will potentially adjust prices over time to capture low value customers while still taking advantage of the revenue opportunities associated with impatient customers.

In the present paper, we study the firm's problem of how to commit to a sequence of prices over time to maximize revenues given a customer population with heterogeneous patience levels and valuations. We consider a firm selling a single product to a stationary flow of heterogeneous customers that arrive over time. Each customer arrives with unit demand and is characterized by her valuation for the product and her willingness to wait before purchasing. The latter may be interpreted as the time the customer is willing to monitor the market. The customers' valuations and willingness to wait can take a very general form and in particular may be correlated. The firm's problem is to select prices to offer for all future periods, to maximize long-term average revenues. The

RIGHTSLINK

customers are assumed to be strategic; they anticipate the firm's prices and optimize their purchase timing over the time they are present in the system (based on their willingness to wait). If the lowest price the customer sees during her time in the system is below her valuation, she will purchase the product at that lowest price, otherwise she will leave the system without purchasing the product. Roughly speaking, this problem may be interpreted as a two-stage game in which the firm first commits to an infinite sequence of prices and the customers respond by selecting whether and when to purchase.

Main Contributions. At a high level, the present paper makes four main contributions.

(1) We establish that the pricing problem is amenable to analysis under fairly general assumptions. In particular, we show that an optimal policy is cyclic, and we derive a tight bound on the length of optimal cycles.

(2) We develop a general dynamic programming approach with a novel state-space structure that leverages the structure of the problem. Given the bound on the length of cycles of optimal policies, we establish that this procedure computes such policies efficiently.

(3) We demonstrate the algorithmic technique we develop can also be applied to finite-horizon problems with nonstationary demand, with customers incurring a uniform cost of waiting.

(4) We establish a clear connection between two settings that lead customers to time their purchases: varying patience levels for one-time purchases and varying storage capacities for repeat purchases (in which case consumers may stockpile). In particular, we show that the two problems are equivalent.

In more detail, we establish that for any joint distribution between patience levels and valuations, the firm may restrict attention to cyclic pricing policies. This initial result validates some of the cyclic policies being adopted in practice and enables one to narrow down the space of policies that need to be considered for optimization purposes. Using the structure of the firm's price optimization problem, we further show that one may restrict attention to cycles that have length at most twice the maximum willingness to wait of the customer population. This is a tight bound on the shortest cycle length of an optimal policy. This crisp result relies on a series of structural insights, which revolve around the concept of *effective price tables*, an object that summarizes the mapping from the prices actually paid to consumer segments and arrival times. In particular, the result relies on a reflection principle that establishes that a policy and its time reflection are revenue equivalent.

Although one may restrict attention to cyclic policies with bounded length, the number of potential price cycles is exponentially large, bringing to the foreground the question of how to compute optimal pricing policies. We develop a novel dynamic programming approach for the problem. Leveraging the above results on the cycle length bound and the underlying structure of the problem through the geometry of the effective price tables, we develop an algorithm to compute an optimal policy that is polynomial in the maximum willingness to wait and linear in the number of prices the firm can use. This approach is based on a dynamic program that uses a novel state and action space, which enables one to solve for an optimal policy recursively. This approach is general, and we show how it extends to the case of a finite horizon with nonstationary market size parameters and willingness-to-pay distributions, where customers may have a uniform cost of waiting.

We also study the structure of optimal policies. Given the attention that monotone cyclic polices have received in the literature (more on that in the review below), we study this subclass of policies in more detail. For this subclass of policies, we show that one may further restrict the length of cycles without loss of revenue. However, the class of monotone cyclic policies is, in general, suboptimal. We derive a class of problems in which they will always be suboptimal and show that they can yield arbitrarily poor performance compared to an unrestricted optimal policy. Optimal policies, in general, have a rich structure. In particular, we show that if the pricing policy cycle is relatively long (the period is close to twice the maximum willingness to wait), then the lowest price should be offered at the end of a cycle and the second lowest price should be offered halfway to the end of the cycle. We also show numerically that the optimal policy often takes the form of *nested sales*, where the firm offers small sales often, medium-sized sales less often, and its largest sale only once per selling season.

Finally, we show that one may apply all the results above to the problem of pricing for a pool of repeat consumers who may stockpile the item, a common problem encountered, for example, by grocery stores. In particular, we establish a fundamental connection between the problem of pricing for consumers with one-time purchases who time their purchase over a window and that of pricing for repeat consumers who may stockpile the product. We first analyze the stockpiling problem faced by consumers given a sequence of prices. We establish, through a proper accounting scheme from prices to units consumed, that in an optimal stockpiling policy, the effective price for a potential unit to be consumed in a given period is the minimum of the past prices over a window of length driven by the storage capacity. Based on this result, the firm's problem can be shown to be, roughly speaking, the mirror image of the problem with one-time purchases (with proper parameters). We then

RIGHTSLINK

establish that the two problems admit the same value and that an optimal policy for one problem is also optimal for the other one. In other words, the two problems are essentially equivalent.

Related Literature. How to optimally set prices over time given that consumers strategically time their purchases is a classical question in economics (see, e.g., Coase 1972, Stokey 1979, Conlisk et al. 1984, Besanko and Winston 1990, Sobel 1991) and one that has received significant attention in the revenue management and dynamic pricing community; see the recent reviews by Shen and Su (2007) and Aviv et al. (2011).

When consumers are strategic, various considerations come into play in their purchasing decisions, including the future prices, the evolution of valuations, and availability of the product. Aviv and Pazgal (2008) study dynamic pricing with capacity constraints in the presence of strategic customers. They illustrate the extent of revenue deterioration one may experience if one ignores the presence of strategic customers. Su (2007) finds that in a setting with limited inventory, both markdown and markup policies can be optimal depending on the problem instance. The paper by Ahn et al. (2007) studies joint pricing and manufacturing decisions when demand in a given period is a function of the price in multiple periods. Our model also possesses the latter feature, and our analysis, like their work, also exploits the regeneration of the system (or system reset) and the policy decomposition that follows. Our work and their work, however, deal with fairly different problems (we do not study manufacturing decisions) and, overall, the techniques we develop are different from theirs. Borgs et al. (2014), motivated by the question of how to sell online services, consider how to set prices to extract revenue while guaranteeing service availability to all customers willing to pay the price set by the firm. Their model also assumes consumers arrive over time and have windows of interest, but their focus is on handling time-varying service capacity constraints. Deb (2010) and Garrett (2011) capture the impact of customers' valuations evolving stochastically. In particular, Garrett (2011) shows that stochastic valuations can drive the optimal price path to be nonmonotone.

Our work takes a more fundamental starting point, attempting to isolate and capture the impact of a heterogeneous population on the optimal dynamic pricing policy, absent any other considerations. In this sense, our work builds upon the classical papers on intertemporal price discrimination. Stokey (1979, 1981) shows that a firm facing a heterogeneous population of customers can maximize revenue by either using a sequence of decreasing prices or by offering a single constant price, depending on the distribution of customers' valuations and the firm's ability to commit to a price path.

The paper that is most closely related to our work is by Conlisk et al. (1984), who show that if a new cohort of consumers arrives at every period, then the firm's optimal strategy is to use a cyclic pricing policy. In their model, a given consumer valuation can take one of two values —low or high—and customers may stay in the system forever. The paper shows the interesting phenomenon that seasonal monotone pricing arises naturally in stationary models as a result of the firm performing intertemporal price discrimination. The firm sells only to high-value consumers most periods but sells to low-value consumers as soon as enough of them accumulate in the system. The key differences in the present paper are as follows: (i) On the firm side, the firm commits to future prices, and maximizes finite horizon or long-term average revenues. (ii) On the demand side, there are infinitely many customers types (as opposed to two) that differ along valuations and the time they may spend in the system, and these do not discount future rewards. We show in the present paper that, as soon as one departs from the assumptions of Conlisk et al. (1984), there is no reason for monotone cyclic policies to be optimal in general, and the rich price dynamics observed in practice with multiple sales levels being offered at different times may be rationalized. Furthermore, we show that under the present model, the problem is equivalent to that of pricing for customers who stockpile, unifying settings where customers time their purchases.

In comparison to Board (2008), who also focuses on a setting with price commitment, the present model does not include discounting but can be seen as allowing for a richer description of customers since types are now three-dimensional: time of arrival to the system, patience level, and willingness to pay. In contrast, in Board (2008), customers were fully characterized by their arrival times and their willingness to pay. This additional dimension drives the cyclic pricing structure we derive (rich cycles are observed even under stationary demand) as opposed to Board (2008), where cyclic pricing follows from cyclic demand patterns.

In §6, we study the problem of how to do intertemporal pricing in the presence of consumers who stockpile the firm's goods. Two of the early papers on this topic are by Blattberg et al. (1981) and Jeuland and Narasimhan (1985), who show how firms can price discriminate between consumers with high and low holding costs by using dynamic pricing. In recent work, Su (2010) shows that, in a rational expectations model, the firm should use recurring promotions

RIGHTSLINK

when customers who shop frequently are willing to pay more than occasional shoppers.

On the empirical side, Nair (2007) estimates a model of strategic purchase timing behavior in the context of video games, and recent work by Li et al. (2014) estimates the extent to which consumers are strategic in timing their purchases in the context of airline pricing. Pesendorfer (2002) and Hendel and Nevo (2006, 2013) study pricing of items in supermarkets, and estimate demand elasticity, accounting for demand accumulation, showing that such an effect is significant. The class of models we consider in §6 includes the demand model estimated in the latter paper in the absence of competition, and the class of models presented in §2 includes the perfect foresight model of Li et al. (2014).

Proofs. The proofs of Lemmas 1, 2 and 3; Theorems 1, 2, and 5; and Proposition 1 are presented in the appendix. The remaining proofs are available in the online supplemental appendix (available at http://dx.doi.org/10.1287/mnsc.2014.2049).

2. Model

We consider a monopolist facing a multi-period single-product pricing problem. Customers arrive with unit demand and are characterized by their valuation for the product $v \in [0, \bar{V}]$ as well as their willingness to wait $w \in \{0, 1, \ldots, S\}$ for some $\bar{V} \in \mathbb{R}^+$ and $S \in \mathbb{N}$. Customers are assumed infinitesimal and, in each period, the mass of the incoming customer population with patience w is given by γ_w. For each patience level w, the cumulative distribution of values is given by $F_w(\cdot)$. In other words, the demand stream is stationary. We do not impose any assumptions on the demand model $\{\gamma_w, F_w(\cdot)\}_{w=0,\ldots,S}$. In particular, the correlation between the customers' valuation for the product and willingness to wait is arbitrary. We consider a *deterministic* flow of customers, so that in every period $t = 1, 2, \ldots,$ the mass of customers arriving with patience w and valuation below or equal to v is exactly $\gamma_w F_w(v)$.

We let $\mathcal{D}$ denote the set of feasible prices available to the firm, which we assume to be an arbitrary nonempty closed subset of $[0, \bar{V}]$. The firm may select any pricing sequence $\mathbf{p} = \{p_t\}_{t\in\mathbb{N}}$ with elements in $\mathcal{D}$, and we assume that when it selects such a sequence, it commits to it. We let $\mathcal{P}$ denote the set of all such sequences.

Customers are assumed to be strategic and fully anticipate the firm's future prices, so a customer arriving in period t with a willingness to wait of w will compare the net utility stemming from purchasing in periods $\{t, t+1, \ldots, t+w\}$ and select the period that yields the highest net utility and purchase in that period only if the latter is nonnegative. Based on this process, the customer will consider purchasing the firm's product only at the period that has the lowest price in the window. For a given pricing policy $\mathbf{p} \in \mathcal{P}$, we say that a customer that arrives in period t with patience w faces an *effective price* of

$$e_{t,w}(\mathbf{p}) = \min_{t \le k \le t+w} p_k \tag{1}$$

and will purchase the product only if her valuation is above the effective price she encounters.

Given the consumer behavior outlined above and a pricing policy $\mathbf{p} \in \mathcal{P}$, the long-run average revenue collected by the firm is given by

$$R(\mathbf{p}) = \liminf_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \sum_{w=0}^{S} \gamma_w e_{t,w}(\mathbf{p}) \bar{F}_w(e_{t,w}(\mathbf{p})), \tag{2}$$

where $\bar{F}_w(v) = 1 - \lim_{v' \uparrow v} F_w(v')$, which represents the fraction of consumers with patience w that value the product at least v.[1] The firm solves a deterministic problem. The firm's objective is to select a sequence of prices to commit to in order to maximize its long-run average revenues; i.e., the firm solves

$$\sup_{\mathbf{p}\in\mathcal{P}} R(\mathbf{p}). \tag{3}$$

Discussion of the Assumptions. The present paper focuses on a setting in which consumers have a time window over which they consider purchasing, with an arbitrary link between the length of the windows and the willingness-to-pay distribution. It differs from the formulations of dynamic pricing problems in which the firm and/or the consumers use exponential discounting to trade off present versus future payoffs. It relates to models in which consumers have a homogeneous patience level (see, e.g., Ahn et al. 2007 and Yin et al. 2008) or the so-called deadline models (see, e.g., Mierendorff 2011, Pai and Vohra 2013). Versions of such models have also been estimated in recent empirical work (Li et al. 2014, Hendel and Nevo 2013). While the model we consider is of independent interest, the fact that it allows to establish a fundamental connection between the present problem and that of pricing to a pool of consumers who stockpile (§6) lends it further appeal. We also return in §7 to discuss how one may incorporate some cost of waiting, where customers pay a waiting cost that is linear on the number of periods they wait before purchasing, relating to the model in Su (2007).

Another important assumption underlying our results is the stationarity of the demand. This assumption is also made in other papers in the literature, including Conlisk et al. (1984), Su (2007), and Yin et al.

[1] The objective above is defined with a lim inf since the limit may be undefined for some pricing policies $\mathbf{p}$.

RIGHTSLINK

(2008). In §7, we show how one may consider nonstationary environments.

We highlight here that we focus on a setting in which the firm commits to future prices upfront; i.e., we assume that the firm has commitment power. Given that we will analyze different settings (long-run average objective, finite horizon objective, nonstationary demand, relationship to stockpiling problem), we anchor ideas around the simpler commitment case. One of the first papers on intertemporal price discrimination, Stokey (1979), as well as several recent papers by Board (2008), Deb (2010), Borgs et al. (2014), and Garrett (2011) study optimal dynamic pricing from the perspective of a monopolist who is able to commit to future prices. As discussed in the literature review, there is also a body of papers that analyze cases in which firms do not have commitment power.

3. Optimal Pricing Policies

In this section, we show that the pricing problem (3) admits an optimal solution. We show that one may restrict attention to cyclic policies, and that the maximum length of cycles to consider can be characterized only as function of the maximum willingness to wait.

3.1. Policy Decomposition and Optimality of Cyclic Policies

A first important concept that we introduce is that of *resetting* periods. Whenever prices are at their lowest, the entire system resets in the sense that all customers depart the system, either by making a purchase or by deciding not to purchase at all.[2] There is no value for a strategic customer to stay in the system past the date when the lowest price is being offered. Resetting, however, also occurs when prices are not at their lowest price overall. If the price offered today is lower than all the prices to be used in the next S periods, the system also resets since no customer is willing to wait more than S periods to make a purchase. This occurs whenever the current price p_t is equal to the effective price faced by a customer of maximum patience $e_{t,S}(\mathbf{p})$.

Definition 1. For any pricing policy $\mathbf{p}$, let $V(\mathbf{p}) \subseteq \mathbb{N}$ be the set of periods such that $p_t = e_{t,S}(\mathbf{p})$. We call the elements in $V(\mathbf{p})$ the *reset periods* of the system.

As a convention, we include 0 in the set $V(\mathbf{p})$ since the system is empty when the first customers arrive in period $t = 1$. We now introduce the subclass of cyclic pricing policies.

Definition 2. A pricing policy is *cyclic* if there exists some integer $L > 0$ such that $p_{t+L} = p_t$ for all $t \in \mathbb{N}$. The smallest $L > 0$ for which this holds is called the *cycle length* $L_{\mathbf{p}}$ of policy $\mathbf{p}$.

[2] Without loss of generality, one may assume that customers behave in this fashion.

With a slight abuse of notation, we represent a cyclic policy $\mathbf{p}$ by the finite sequence of prices $\mathbf{p} = (p_1, \ldots, p_{L_{\mathbf{p}}})$. Whenever we discuss a policy $(p_1, \ldots, p_T)$, we are referring to the policy for which this finite sequence of prices is repeated infinitely often.

Given the above, one may now introduce the notion of the components of an arbitrary pricing policy in $\mathcal{P}$. Purchasing patterns between a given pair of reset periods can be analyzed independently from prices offered before and after such reset periods since only the customers arriving in between those two reset periods are affected by these prices.

Definition 3. Let $V_i(\mathbf{p})$ be ith smallest element in the set $V(\mathbf{p})$. For any $i \in \mathbb{N}$, we say the finite sequence of prices $C_i(\mathbf{p}) = (p_{V_i(\mathbf{p})+1}, p_{V_i(\mathbf{p})+2}, \ldots, p_{V_{i+1}(\mathbf{p})})$ is the ith *component policy* of $\mathbf{p}$.

Lemma 1. *Suppose the set of prices $\mathcal{D}$ is finite. Then, for any policy $\mathbf{p}$, the number of time periods elapsed between any two reset periods is at most $S|\mathcal{D}|$, i.e.,*

$$\max_{i \in \mathbb{N}} \{V_{i+1}(\mathbf{p}) - V_i(\mathbf{p})\} \leq S|\mathcal{D}|.$$

In other words, when there is a finite number of prices, the component policies of an arbitrary policy $\mathbf{p}$ cannot be arbitrarily long. The result stems from the fact that whenever a period t is not a reset period, there must exist a period within $\{t+1, \ldots, t+S\}$ where the price offered is strictly less than p_t. Repeating this process recursively will necessarily lead to a reset period in finite time because prices may not decrease below $\min\{\mathcal{D}\}$. The next result highlights the connection between the performance of a policy and its components and will be a key building block in showing that the problem admits an optimal solution as well as restricting the set of policies that one needs to consider.

Lemma 2 (Policy Decomposition). *Suppose the set of prices $\mathcal{D}$ is finite. Then, the long-run average revenue $R(\mathbf{p})$ generated by a pricing policy $\mathbf{p}$ is at most the supremum of the long-run average revenues generated by each of its component policies, i.e.,*

$$R(\mathbf{p}) \leq \sup_{i \in \mathbb{N}} R(C_i(\mathbf{p})). \tag{4}$$

The idea behind the policy decomposition lemma is as follows: The average revenue from a pricing policy is nothing but a convex combination of all the average revenues obtained by the component policies. If the average revenue obtained in between a pair of reset periods is higher than in other periods, then one may replace these other prices by the ones from the component policy that yields higher average revenue. The key implication of this result is that one may restrict attention to the "best" component of a policy and obtain the following result.

RIGHTSLINK

Proposition 1. *Suppose the set of prices $\mathcal{D}$ is finite. Then, there exists a cyclic pricing policy with cycle length at most $S|\mathcal{D}|$ that achieves the supremum in (3).*

The proposition above has two immediate implications. The first one is showing that the supremum of the price optimization problem given in Equation (3) is attained. Therefore, the notion of an optimal pricing policy is well defined. The second implication is that there exists an optimal solution that is cyclic. The result relies on the fact that the maximum time elapsed between two reset periods is $S|\mathcal{D}|$, and hence the set of possible component policies is finite. This, in turn, implies that the supremum in Equation (4) can be replaced by a maximum over finitely many cyclic component policies.

Proposition 1 enables one to restrict attention to cyclic policies without loss of optimality.[3] However, the bound on the length of optimal cycles in Proposition 1 depends on the number of prices at the firm's disposal and may be a weak bound if the firm has many prices at its disposal. We next establish a *tight* bound that is independent of the number of prices.

3.2. Reflection Principle and Optimality of Short Cyclic Policies

Effective Price Tables. We next analyze in further detail the structure of the pricing problem. To do so, we introduce the notion of *effective price tables* for cyclic policies. One may represent the effective prices of a cyclic policy of length T through a matrix of size $T \times S$, in which each entry corresponds to the effective price faced by a customer with patience the row number and arriving in the period given by the column number.

Let us consider a numerical example with maximum willingness to wait $S = 3$, a cyclic policy with length $T = 8$, and decreasing prices $\mathbf{p} = (15, 12, 8, 7, 4, 3, 2, 1)$. The effective price table of the policy $\mathbf{p}$ is presented in Table 1. The performance of the policy may be easily computed given the table.

The Reflection Principle. We will now point to a general relationship between a cyclic policy $(p_1, \ldots, p_T)$ and its time reflection $(p_T, \ldots, p_1)$. Consider the time reflection of the above policy, which has cyclic increasing prices $\mathbf{p}^r = (1, 2, 3, 4, 7, 8, 12, 15)$. The corresponding effective price table is depicted in Table 2.

Note that the effective price tables of policies $\mathbf{p}$ and $\mathbf{p}^r$ demonstrate very different customer behavior. With the cyclic decreasing policy, all customers wait for a lower price except for customers who are completely impatient ($w = 0$) or already arrive at a period when the lowest price is being offered. In contrast, the only customers who wait in the case of the cyclic increasing policy, are the customers who can wait until the price falls to its lowest value ($p = 1$); everyone else either purchases at their arrival period or does not purchase at all. However, even though the pricing policies lead to very different consumer behavior, they yield identical revenue. This can be observed by counting the number of times each effective price appears for each value of w. For both pricing policies, the effective price of 15 appears only for $w = 0$ and appears only once; the effective price of 12 appears only for $w = 0$ and $w = 1$ and appears only once for each of these values, and so on. Cyclic decreasing policies and cyclic increasing policies appear at first brush to be very different policies and do indeed lead to different purchasing patterns; however they are in fact revenue equivalent. This is a general result that is not restricted to monotone cyclic policies that we formalize below.

Table 1 **Effective Price Table of a Cyclic Decreasing Policy p**

	$t=1$	$t=2$	$t=3$	$t=4$	$t=5$	$t=6$	$t=7$	$t=8$
$w=0$	15	12	8	7	4	3	2	1
$w=1$	12	8	7	4	3	2	1	1
$w=2$	8	7	4	3	2	1	1	1
$w=3$	7	4	3	2	1	1	1	1

Table 2 **Effective Price Table of a Cyclic Increasing Policy $\mathbf{p}^r$**

	$t=1$	$t=2$	$t=3$	$t=4$	$t=5$	$t=6$	$t=7$	$t=8$
$w=0$	1	2	3	4	7	8	12	15
$w=1$	1	2	3	4	7	8	12	1
$w=2$	1	2	3	4	7	8	1	1
$w=3$	1	2	3	4	7	1	1	1

Lemma 3 (Reflection). *A cyclic pricing policy $\mathbf{p} = (p_1, \ldots, p_T)$ and its time reflection $\mathbf{p}^r = (p_T, \ldots, p_1)$ yield the same revenue, i.e., $R(\mathbf{p}) = R(\mathbf{p}^r)$.*

The proof of the result resides in extending and formalizing the informal counting argument above. In particular, the following relationship holds for all t and w,

$$e_{t,w}(\mathbf{p}) = \min\{e_{t,w-1}(\mathbf{p}), e_{t+1,w-1}(\mathbf{p})\}. \tag{5}$$

That is, the effective price faced by a customer that arrives at time t with willingness to wait w is the lowest between the effective prices of a customer that arrives at the same period but is willing to wait only $w - 1$, and someone who arrives one period later and is willing to wait only $w - 1$. Starting from the fact that $\{e_{t,0}(\mathbf{p})\}_{t=1,\ldots,T}$ and $\{e_{t,0}(\mathbf{p}^r)\}_{t=1,\ldots,T}$ are reflections of each other, one can use Equation (5) recursively to show that for each value of w from 1 to S,

[3] Note that there always exist optimal noncyclic policies as well since any two policies $\mathbf{p}$ and $\mathbf{p}'$ that are identical except for finitely many periods will yield the same long-run average revenue.

RIGHTSLINK

$\{e_{t,w}(\mathbf{p})\}_{t=1,\ldots,T}$ and $\{e_{t,w}(\mathbf{p}^r)\}_{t=1,\ldots,T}$ contain exactly the same elements. We note that the result above relies in a critical fashion on the assumption of stationarity of demand.

We are now in a position to state one of our main results.

Theorem 1. *Suppose the set of prices $\mathscr{D}$ is finite or that $F_w(\cdot)$ is Lipschitz continuous for all $w = 0, \ldots, S$. Then, there exists an optimal cyclic pricing policy with cycle length at most 2S.*

A priori, it is not clear what drives the length of an optimal cycle. Theorem 1 states that the only cycles that need to be considered are short, in the sense that they need not exceed twice the maximum willingness to wait. When the set of prices is finite, the result can be obtained using the following argument. From the policy decomposition lemma, we immediately obtain that the lowest price should be used only once per pricing cycle in the shortest optimal cyclic policy. Less obviously, the same lemma also implies that the second-lowest price should be used only up to S periods before the lowest price is used. The key idea in the proof is to use the same logic on the reflected policy (which yields the same long-run revenues by Lemma 3). Doing so, one obtains that the second-lowest price should be used only up to S periods after the lowest point. If the optimal policy had length longer than $2S$, there would always be an option to further decompose the policy or its reflection without decreasing long-run average revenues. If the set of prices is closed but not finite and the cumulative distribution of values is Lipschitz continuous, we obtain the same result by finding a sequence of policies with cycle length up to $2S$ that rely only on finitely many prices such that the revenue of these policies converge to the revenue of the optimal policy.

The proof of the result above has an interesting structural implication for optimal policies. By ensuring that the second-lowest price does not appear more than S periods after or before the second-lowest price, it guarantees that the second-lowest price will be used roughly halfway through the selling season when the optimal policy is approximately $2S$ periods long. This idea is made precise in Proposition 4 in §5.

Furthermore, the bound above is sharp in the sense that for any S, there are instances with maximum willingness to wait S in which all cycles with length strictly lower than $2S$ are suboptimal. An instance of the problem is defined by the nonempty closed set of feasible prices $\mathscr{D} \subset [0, \bar{V}]$, the maximum patience level S in the set of nonnegative integers, the segment sizes $\{\gamma_w\}_{w=0,\ldots,S}$ in $\mathbb{R}_+^{S+1}$, the set of willingness-to-pay distributions $\{F_w(\cdot)\}_{w=0,\ldots,S+1}$ in the set of nondecreasing and right-continuous mappings from $[0, \bar{V}]$ into $[0, 1]$ with $F_w(0) = 0$ and $F_w(\bar{V}) = 1$. We let $\mathscr{I}$ denote the set of all possible instances.

Proposition 2. *For any S, there exists an instance of the problem in $\mathscr{I}$ with maximum willingness to wait S for which the shortest optimal cycle is $2S$ periods long.*

As we see next, Theorem 1 also enables one to construct an efficient algorithm to find an optimal policy.

4. Computing Optimal Cycles: A Geometric Approach

Theorem 1 established existence of an optimal cyclic pricing policy with length at most $2S$. While this implies that cycles under consideration can be fairly short, the number of such cycles is still very large (it is exponential in S when the price set is finite). In this section, we show, by leveraging the structure of the problem at hand, that the problem of finding an optimal pricing policy is tractable, and we construct an algorithm that efficiently finds an optimal cycle.

A Geometric View of the Effective Price Table. Our approach to the problem is centered around the geometry of the effective price table. Recall that for a cyclic policy $\mathbf{p}$ with length T, the effective prices table is a $T \times S$ table with the effective prices $e_{t,w}(\mathbf{p})$ for times $t = 1, \ldots, T$ and $w = 0, \ldots, S$. Since the policy is cyclic, we can assume without loss of generality that the smallest price is used in the last period, i.e., $p_T = \min_{1\le k\le T} p_k$. We use the notation $p_{(k)}$ to represent the kth lowest price used in policy $\mathbf{p}$ and $T_{(k)}$ to represent the first period in which $p_{(k)}$ is used. Therefore, our convention that p_T is the lowest price in the cycle is equivalent to $p_T = p_{(1)}$ or $T_{(1)} = T$. The first observation about the effective price table is that the set of elements where the effective price is $p_{(1)}$ forms a triangle, as can be seen in Figure 1. That is, for customers with $w = 0$, only the ones that arrive in period $t = T$ in the cycle are able to purchase at the lowest price. Among the ones with $w = 1$, customers who arrive in periods $t = T - 1$ or $t = T$ are able to buy at the lowest price, and so on. This triangle would have its left side truncated if $\mathbf{p}$ is a cyclic policy with length $T \le S$ (leading to a trapezoid).

The set of effective prices corresponding to the second-lowest price $p_{(2)}$ may also be described geometrically. It also takes the form of a triangle, but it is potentially truncated on both the left and right side, as illustrated by the striped object in Figure 2. It will be truncated on the right side if customers with high patience have access to the lowest price and hence belong to the set of customers who purchase at the end of the cycle. One may continue to represent the set of effective prices corresponding to the kth lowest price recursively.

A Dynamic Programming Recursion. The central idea in building our algorithm is as follows: no customer will ever skip over a low price to buy at a higher one. Recall that $T_{(2)}$ is the first period when the price is equal to $p_{(2)}$. Some customers that arrive between

RIGHTSLINK

Figure 1 (Color online) Set of Effective Prices Equal to $p_{(1)}$ in the Effective Price Table

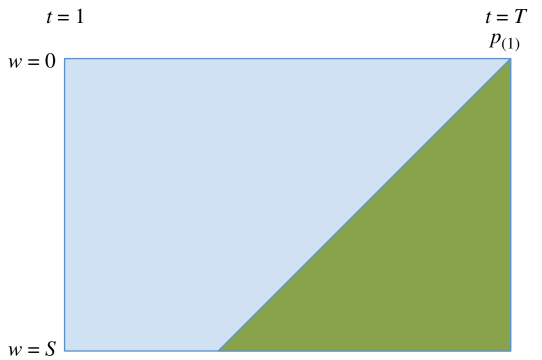

Figure 3 (Color online) Effective Price Table After Customers Who Purchase at p_T Are Removed

$t=1$ and $t=T_{(2)}$ might be patient enough to wait until the lowest price $p_{(1)}$. For these very patient customers, the prices offered in periods from $t=1$ up to $T_{(1)}-1$ are irrelevant. For everyone else arriving between $t=1$ and $t=T_{(2)}$, all prices offered after $t=T_{(2)}$ are irrelevant. The prices being offered after $T_{(2)}$ are either equal to or higher than $p_{(2)}$ or too far into the future. Conditionally on $p_{(2)}$ being the second lowest price and its position $T_{(2)}$, all prices between $t=1$ and $t=T_{(2)}-1$ may be computed independently from prices between $t=T_{(2)}+1$ and $t=T_{(1)}-1$.

The latter observation is the key step to formulate a dynamic programming recursion for the problem. The state space of this dynamic program is an unusual one and is best understood geometrically. If one ignores all the customers that are able to buy at the lowest price in the cycle (the triangle on the right in Figure 1), one is left with an effective price table as depicted in Figure 3.

Such a table takes either the shape of a triangle, if $T-1\le S$, or the shape of a trapezoid, as depicted, if $T-1>S$. The number of elements at row $w=0$ is $T-1$ since only the customers who arrive in period T are able to purchase at the lowest price p_T and therefore are part of the removed triangle.

We have argued that conditional on a period $T_{(2)}$ for the second- lowest price and its price $p_{(2)}$, all the prices before $T_{(2)}$ and after $T_{(2)}$ may be selected independently (subject to a lower bound on prices). This idea gives rise to the dynamic programming recursion. The key geometric insight that allows us to construct the state space can be gleaned from Figure 4: once one removes the customers that purchase at the second-lowest price $p_{(2)}$, we are left with two problems that are identical in structure to the original one but smaller in size. The state space therefore is composed of a pair (n, p), where n denotes the number of periods being considered and p represents the lowest price that can be used during those periods. We define a value function $W_n(p)$ to represent the maximum revenue that can be obtained over n periods assuming that the prices being offered over these periods is at least p and that, in period $n+1$, a price lower than p will be offered. Formally, the value function may be represented by

$$W_n(p)=\max_{p_1,\ldots,p_n\in\mathcal{D}}\sum_{w=0}^{\min\{n-1,S\}}\gamma_w\sum_{t=1}^{n-w}e_{t,w}(\mathbf{p})\bar{F}_w(e_{t,w}(\mathbf{p}))$$
$$\text{s.t.}\quad p_i\ge p\quad \text{for all } i\in\{1,\ldots,n\}.$$

Figure 2 (Color online) Set of Effective Prices Equal to $p_{(1)}$ and $p_{(2)}$ in the Effective Price Table

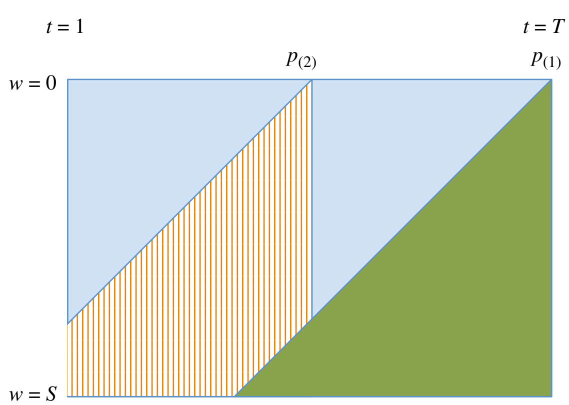

Figure 4 (Color online) Geometric View of the Dynamic Programming Recursion

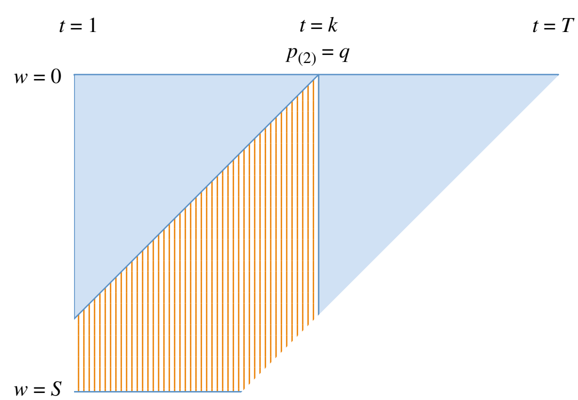

RIGHTSLINK

Using the observation that customers do not skip over low prices to buy at a higher one later, we obtain that the value function $W_n(p)$ satisfies the following Bellman equation:

$$W_n(p) = \max_{\substack{k\in\{1,\dots,n\}\\ p'\in\mathcal{D}: p'\geq p}} \left\{ \sum_{w=0}^{S} t_{n,k,w}\gamma_w p' \bar{F}_w(p') + W_{k-1}(p') + W_{n-k}(p') \right\}, \quad n \geq 1, p \in \mathcal{D}, \quad (6)$$

where $W_0(p) = 0$, and $t_{n,k,w}$ counts the number of periods between 1 and n in which the customers with patience w will purchase at the lowest price in the interval if such price is offered in period k. For a given n and k, the collection of $t_{n,k,w}$ for all w is represented by the shaded area in the middle of Figure 4. Mathematically,

$$t_{n,k,w} = (\min\{k, n-w\} - \max\{1, k-w\} + 1)^+,$$

where $x^+ = \max\{x, 0\}$. Once one has computed the value of $W_n(p)$ for a given n and all $p \in \mathcal{D}$, one may determine the optimal policy of length $n+1$ by adding the lowest price in the cycle at position $T = n+1$ (see Figure 1). The revenue obtained over the first T periods by the best policy of cycle length T is then

$$W_T^* = \max_{p\in\mathcal{D}} \left\{ \sum_{w=0}^{S} \min\{w+1, T\}\gamma_w p \bar{F}_w(p) + W_{T-1}(p) \right\}. \quad (7)$$

Computing an Optimal Policy. Since by Theorem 1 there exits an optimal policy that is cyclic with length at most $2S$, the optimal pricing policy can be determined by computing the average per-period revenue of an optimal policy for each T from 1 to $2S$, i.e.,

$$W^* = \max_{T\in\{1,\dots,2S\}} \frac{W_T^*}{T}, \quad (8)$$

leading to Theorem 2.

Theorem 2. *If the set of prices $\mathcal{D}$ is finite, then an optimal pricing policy can be computed in time $O(|\mathcal{D}|S^2)$.*

In other words, despite the fact that the number of cycles of length up to $2S$ is exponential in S, an optimal policy may be determined in polynomial time in S by exploiting the underlying structure the problem.

When the set of available prices $\mathcal{D}$ is not finite, one may still leverage the recursion above. If the willingness to pay distributions are Lipschitz, one may discretize the set $\mathcal{D}$ and use the regularity of the distributions to find a near-optimal policy efficiently. For any $\varepsilon > 0$, we say that a pricing policy $\mathbf{p}'$ is ε-optimal if $R(\mathbf{p}') \geq \sup_{\mathbf{p}\in\mathcal{P}} R(\mathbf{p}) - \varepsilon$.

Theorem 3. *Suppose the willingness to pay distributions $F_w(\cdot)$, $w = 0, \dots, S$, are Lipschitz with constant L and let $\Gamma = \sum_{w=0}^{S}\gamma_w$. For any closed set of prices $\mathcal{D} \subseteq [0, \bar{V}]$, an ε-optimal pricing policy can be computed in time $O(\Gamma\bar{V}LS^2/\varepsilon)$.*

Theorems 2 and 3 establish that the problem of finding optimal (for finite price sets) or ε-optimal pricing policies is tractable. In the next section, we use these results to compute optimal prices in some numerical instances, and further explore the structure of optimal cycles for intertemporal price discrimination.

5. Structure of Optimal Pricing Cycles

We first highlight that in the present context, the use of different prices over time is driven by the fact that customers with different patience levels have different willingness-to-pay distributions.

Proposition 3. *Suppose customers' valuations distributions are independent of their patience level, i.e., there exists some cumulative distribution function $G(\cdot)$ such that $F_w(\cdot) = G(\cdot)$ for all $w = 0, \dots, S$. Then, an optimal policy for the firm is to offer a constant price over time.*

5.1. Optimal Cycle Structure

In Figure 5, we report the optimal cycle obtained through the dynamic programming recursion for two instances. Panels (a) and (b) correspond to a case in which $S = 3$ and consumers have deterministic willingness to pay, with

$$(\gamma_w, v_w) = (1, 5-w) \quad \text{for } w = 0, 1, 2, 3. \quad (9)$$

In other words, in such an instance, all customer segments have equal size, the customers within a patience segment are all identical in terms of valuation, and more patient customers have lower valuations for the product. In addition, we assume the set of available prices is $\mathcal{D} = \{1, \dots, 5\}$. Panels (c) and (d) correspond to a case in which $S = 4$ and consumers have deterministic willingness to pay, with

$$v_w = 5 - w \quad \text{for } w = 0, 1, 2, 3, 4, \quad \text{and} \\ \gamma_0 = 4, \quad \gamma_w = 1 \quad \text{for } w = 1, 2, 3 \quad \text{and} \quad \gamma_4 = 3. \quad (10)$$

In addition, the price set is $\{1, \dots, 5\}$. Focusing first on the case $S = 3$ and in particular panel (b), we observe in this case that it is strictly suboptimal to use any cyclic policy with length strictly below $2S = 6$. This complements the result of Proposition 2 that highlighted that in general, it might be necessary to use cycles of length $2S$ to achieve optimality. As a matter of fact, panel (b) further illustrates that one may limit revenue collection by a significant amount by restricting attention to shorter cycles.

In Figure 6, we depict the purchase pattern along an optimal cycle for the case $S = 4$. Let us focus first

RIGHTSLINK

Figure 5 **(Color online) Optimal Cycle Structure**

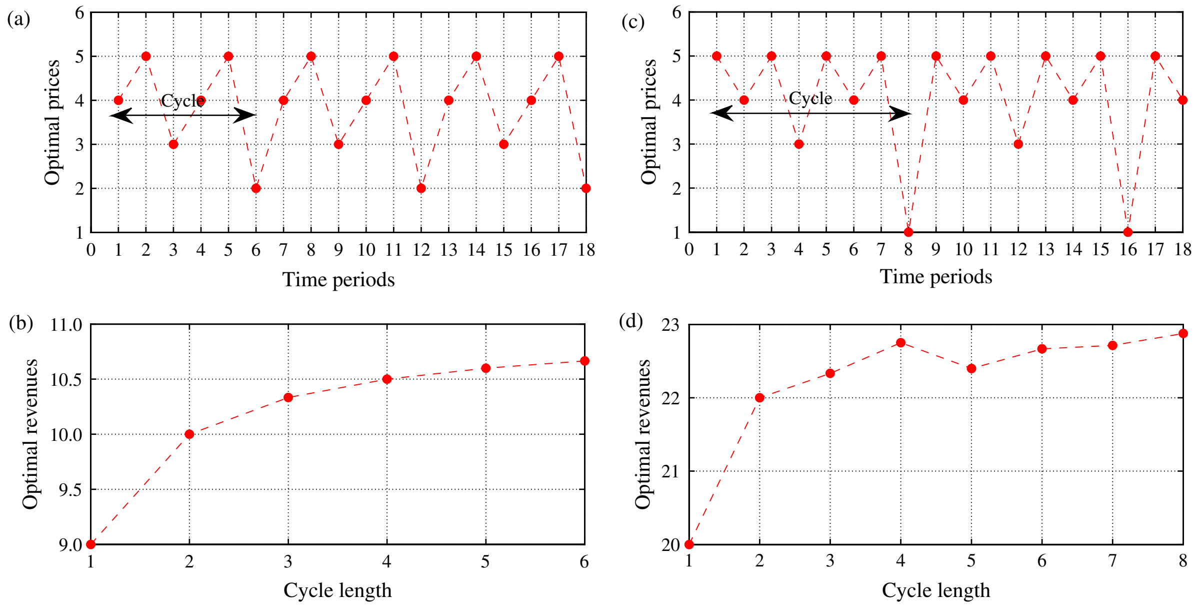

Notes. For the specifications with $S = 3$ given by (9), panel (a) displays prices of an optimal policy with cycle length 6 and panel (b) depicts the optimal revenues one may achieve as a function of the cycle length. For the specifications with $S = 4$ given by (10), panel (c) displays prices of an optimal policy with cycle length 8 and panel (b) depicts the optimal revenues one may achieve as a function of the cycle length.

on impatient customers with $w = 0$. All these customers see prices $p \leq v_0 = 5$, and hence all of them purchase the product upon their arrival (four customers per period). For customers with $w = 1$ (one per period), a customer arriving in the first period postpones his purchasing decision to second period because he would face a lower price there (and $p_2 \leq v_1 = 4$). In period 2, two customers with $w = 1$ purchase: one who arrived in period 1 and one who arrived in period 2. And so on. As we observe, in this optimal cycle, all customers with patience $w = 0$, and $w = 1$ end up purchasing; customers with patience $w = 2$ only purchase in periods 4 and 8, and six out of the eight customers arriving in a cycle purchase; only four customers with patience $w = 3$ (out of eight) and 15 customers with patience $w = 4$ (out of 25) purchase, and those do so at the lowest price in period 8.

Figure 6 **Sales Pattern Through an Optimal Cycle**

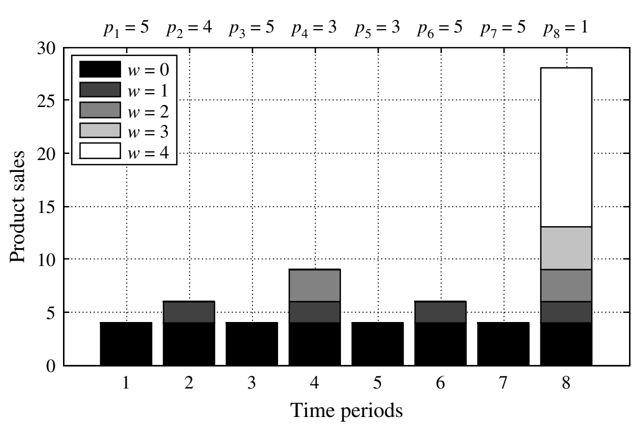

Notes. For the specifications with $S = 4$ given by (10), the figure depicts the number of items purchased during each period of a cycle, and the composition of customers purchasing as a function of patience level.

Nested Sales. Both panels (a) and (c) in Figure 5 depict optimal policies. We observe that the pricing structure within cycles is in general nonmonotone. In particular, optimal polices tend to alternate between sales and the full price (targeting the impatient high-value customers), and the sales are offered with multiple levels of discount depth. The lowest price is offered at the end of the cycle, and the second-lowest price in a cycle is offered exactly in the middle of the cycle (in period 3 for panel (a) and period 4 for panel (c)). The latter observation is more general in the following sense.

Proposition 4 (Nested Sales). *Suppose the shortest cyclic policy that solves Equation* (3) *has a cycle of length T that satisfies $S + 1 \leq T \leq 2S$. Then, there exists an optimal policy with cycle length T in which the lowest price appears last in the cycle and the second-lowest price belongs to $\{T - S, \ldots, S\}$.*

In other words, partial discounts will be found in the middle of a cycle when an optimal cycle is long in the sense that it is close to $2S$ periods long. This phenomenon seems to repeat itself between two discounts, as observed in periods 2 and 6 in panel (c). Conlisk et al. (1984) first noted that seasonal (cyclic) pricing variations would emerge in a setting with stationary demand when the firm was performing some

RIGHTSLINK

form of intertemporal price discrimination. However, in their model with no commitment power, two valuations, discounting, and consumers who may stay in the system forever, they find that such cycles take the form of cyclic monotone policies. In the present setting with heterogeneity over time windows, a continuum of valuations, commitment power, and the absence of discounting for consumers, we find that optimal policies often take the form of nested sales, where the firm offers small sales spread out through a selling season, a larger mid-season sale, and its largest sale at the end of a selling season.

Policy Sensitivity to Mix of Customer Classes. To better understand the impact of the mix of customer patience levels on the structure of the optimal policy, we re-solved the problem from the beginning of this section with $S=4$ (the one that has its solution depicted in panel (c) of Figure 5) but with a different γ_0 or a different γ_S.

We found that the minimum length of an optimal policy evolves nonmonotonically with both γ_0 and γ_S. With few impatient customers (γ_0 from 0 to 2), the optimal policy has a duration of five periods. With γ_0 equal to 3 or 4, the duration increases to 8, as depicted in panel (c) of Figure 5. As we increase the number of impatient customers after that, however, the duration of the optimal policy decreases to 4 (γ_0 equal to 5 or 6), 2 ($\gamma_0=7$), and eventually becomes stationary if $\gamma_0 \geq 8$. As the impatient segment of the population begins to dominate the market, the optimal policy naturally becomes increasingly static.

One should note that in general (beyond the current specific example), there is no reason why the optimal policy would use a single price when γ_0 is very high. When the price set is a continuum and the willingness-to-pay distributions are continuous, it might be worthwhile to deviate from time to time from the optimal price for targeting impatient customers to capture a fraction of the patient customers. However, one expects, of course, that as γ_0 increases, those price deviations would become ever smaller.

In the numerical experiments, the optimal policy was only four periods long if γ_4 is small (γ_4 is less than 2). With $\gamma_4=3$, the optimal policy grows to eight periods long, but this length is reduced to six if $\gamma_4=4$. With $\gamma_4 \geq 5$, the optimal policy becomes five periods long. In contrast to the large γ_0 case, the policy does not become increasingly stationary as we increase the fraction of patient customers. In fact, if γ_S is large, the optimal policy will typically be $S+1$ periods long, as we need to offer prices targeting patient customers only once every $S+1$ periods.

For all parameters tested, the optimal policy had a nested sales structure, except in the cases where the optimal policy was too short to display such structure (period equal to 1 or 2) and in the case with no impatient customers ($\gamma_0=0$).

The Value Induced by the Presence of Patient Customers. In the present setting, the fact that customers are patient induces the firm to consider more complex pricing strategies, but also allows the firm to conduct more targeted pricing. Consider a problem with a market with parameters $\{(\gamma_w, F_w(\cdot)): w=0,\ldots,S\}$. Suppose instead that the firm faced the same mix of customers, except that all of them would be impatient. In such a case, this is equivalent to the firm facing a market with parameters $\{\tilde{\gamma}_0, G_0(\cdot)\}$, where

$$\tilde{\gamma}_0=\sum_{w=0}^{S}\gamma_w, \quad G_0(\cdot)=\frac{1}{\tilde{\gamma}_0}\sum_{w=0}^{S}\gamma_w F_w(\cdot).$$

Now, an optimal policy for the latter setting is a fixed price that maximizes $p\bar{G}_0(p)$. Such a policy is feasible in the original problem (in which customers have different patience levels) and yields the same revenues. We deduce that the firm can necessarily garner at least as much revenues when customers are patient. In other words, the fact that customers are patient allows the firm to intertemporally price discriminate customers who have different willingness-to-pay distributions.

In the present setting, it is also easy to see that if the firm ignores the fact that customers are patient, then the losses are exactly equal to the performance gap between an optimal cycle of problem (3) and the best static price. This gap can be arbitrarily large (this follows from Proposition 7 below) and stems from misspecification of the customer model.

5.2. The Subclass of Monotone Cyclic Policies

We now study monotone cyclic policies. In particular, we bound their cycle length, characterize conditions under which they are strictly suboptimal, and analyze the revenue loss for a firm that restricts itself to such policies. The first result bounds the length of optimal policies within the set of all cyclic monotone policies.

Proposition 5. *For any cyclic policy that is monotone over a cycle, there exists a cyclic monotone policy with cycle length at most $S+1$ that yields at least as much revenue.*

In other words, when focusing on cyclic monotone policies, it is sufficient to focus on policies of length at most $S+1$. Longer monotone cycles would necessarily have multiple reset periods within each cycle, which is unnecessary by the policy decomposition lemma.

Exploring the structure of a cycle in the most general case is difficult given the combinatorial nature of the problem. However, one may further refine the analysis in special cases of interest. Next, we focus on the class of problems in which consumers have a deterministic patience-dependent willingness to pay.

RIGHTSLINK

Assumption 1. *Consumers with patience level w have a willingness to pay of v_w and $v_0 > v_1 > \cdots > v_S > 0$. Furthermore, the set of available prices $\mathcal{D}$ includes $\{v_0, v_1, \ldots, v_S\}$.*

Assumption 1 also imposes that consumers with lower patience have higher willingness to pay, which is natural in many settings. We define

$$\bar{R}(p) = \sum_{w=0}^{S} \gamma_w p(1 - F_w(p))$$

as the single-price revenue per period when the firm uses price p throughout.

Proposition 6. *Suppose Assumption 1 holds. Suppose that $\bar{R}(v_i)$ is nonmonotone in $i \in \{0, \ldots, S\}$ and let $j = \min\{i \in \{1, \ldots, S\}: \bar{R}(v_{i-1}) > \bar{R}(v_i)\}$. Then, an optimal cyclic policy either contains at most $j+1$ periods or is cyclic nonmonotone.*

Proposition 6 excludes the optimality of monotone policies that are longer than $j+1$. For example, if $\bar{R}(v_0) > \bar{R}(v_1)$, corresponding to a case in which the seller prefers to sell only to the impatient customers than to use a price that sells to both impatient and those with patience $w = 1$, then an optimal policy either contains at most two prices or is nonmonotone.

We next illustrate that monotone policies might in general leave significant revenues on the table; consider the following example with maximum patience level $S = 7$ and three segments of customers: impatient, moderately patient, and very patient. In particular, suppose $(\gamma_0, v_0) = (0.1, 10)$, $(\gamma_3, v_3) = (0.5, 2)$, $(\gamma_7, v_7) = (2, 0.5)$, and $\gamma_w = 0$ for $w = 1, 2, 3, 5, 6$. Note that this specification satisfies Assumption 1. Figure 7 depicts an optimal policy (top panel) as well as the best policy among monotone cyclic policies (bottom panel). The ratio of the performance of the latter compared to the optimal policy is of 87.53% in this instance. The example illustrates the need for nested sales for better price discrimination. In the example above, the natural choice of price to use for periods 1–3 is v_0; at period 4, the firm has to decide whether to target customers with intermediate patience, but if restricted to monotone policies, this switch in price implies that the firm will not be able to perfectly target impatient customers until the end of a cycle. For any candidate monotone policy, a similar trade-off will be present, i.e., the firm will have to decide whether to set a high price that will cause a significant portion of the customers not to purchase or to set a low one that will cause the firm to imperfectly target a significant segment of the customers until the end of the cycle. In contrast, a policy that is unconstrained does not face this trade-off. The optimal policy depicted on the top panel is able to target customers with moderate patience in the fourth period of a cycle while returning to target the impatient high-value customers in the next period.

As a matter of fact, we next show that cyclic monotone policies may yield arbitrary poor performance.

Proposition 7. *Let $\mathcal{M}$ denote the set of cyclic policies that are monotone over a cycle. Then,*

$$\inf_{\{\mathcal{D}, S, \{\gamma_w, F_w\}_{w=0,\ldots,S}\} \in \mathcal{I}} \frac{\sup_{\mathbf{p} \in \mathcal{M}} R(\mathbf{p})}{\sup_{\mathbf{p} \in \mathcal{P}} R(\mathbf{p})} = 0. \tag{11}$$

In other words, cyclic monotone policies may not guarantee a uniform finite fraction of revenues.

Figure 7 (Color online) Monotone vs. Optimal Policies

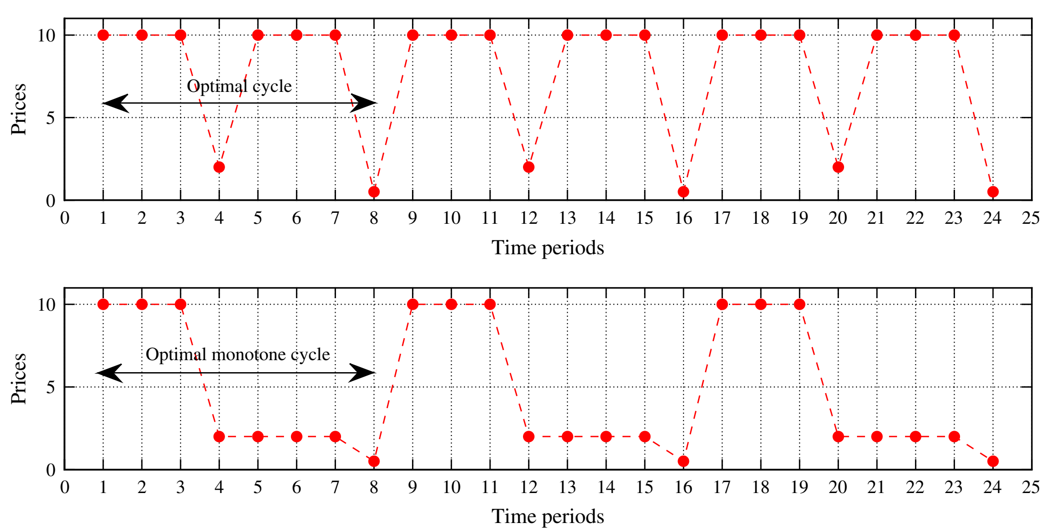

Note. The top panel displays prices associated with an optimal cycle and the bottom panel depicts the best monotone cyclic policy.

6. Intertemporal Pricing with Consumer Stockpiling

In the present section, we show that there is a close relationship between the problem analyzed in the previous sections (problem (3)) and that of pricing for a pool of heterogeneous consumers who may stockpile units of the product. In particular, we demonstrate that the framework developed, the effective price table geometry and the results that followed may be applied to another fundamental problem, that of pricing to a heterogeneous population of consumers who may stockpile the product. The latter problem has been studied in the economics and operations literatures and we refer the reader to Su (2010) and the references therein for further background.

Model of Consumer Stockpiling. We consider a monopolist facing a multi-period single-product pricing problem. The customer population is assumed to be present throughout, with unit demand per period. These are characterized by their valuation for the product $v \in [0, \bar{V}]$, which is constant from period to period, as well as their storage capacity $c \in \{0, 1, \ldots, C\}$, for some $\bar{V} \in \mathbb{R}^+$ and $C \in \mathbb{N}$. Customers are assumed infinitesimal and the mass of the customer population with storage capacity c is given by γ_c. For each storage level c, the cumulative distribution of values is given by $F_c(\cdot)$. We do not impose any assumptions on the demand model $\{\gamma_c, F_c(\cdot)\}_{c=0,\ldots,C}$. In particular, the correlation between the customers' valuation for the product and the storage capacity is arbitrary.

As earlier, we continue to let $\mathcal{D}$ denote the set of feasible prices available to the firm, which we assume to be an arbitrary nonempty closed subset of $[0, \bar{V}]$. The firm may select any pricing sequence $\mathbf{p} = \{p_t\}_{t\in\mathbb{N}}$ with elements in $\mathcal{D}$. We continue to denote by $\mathcal{P}$ the set of all such sequences.

Customers are assumed to be able to fully anticipate the firm's future pricing and may time their purchases accordingly. In particular, consider a consumer with valuation v and storage capacity c. Let y_t denote the number of units purchased in period t. A consumer policy $\mathbf{y}$ consists of a purchasing sequence and it is said to be feasible if

$$I_0 = 0,$$
$$I_t = I_{t-1} + y_t - x_t, \quad t \geq 1,$$
$$x_t = \mathbf{1}\{\{y_t > 0\} \cup \{I_{t-1} > 0\}\}, \quad t \geq 1,$$
$$y_t \in \{0, 1, 2, \ldots, c+1\}, \quad I_t \geq 0, \quad I_t \leq c, \quad t \geq 1.$$

Here, x_t denotes the consumption in period t, and I_t denotes the inventory carried over from period t to $t+1$. The expression for x_t reflects the assumption that consumption always takes place if a unit is available, which is without loss of optimality for the consumers. We let $\mathcal{Y}_c$ denote the set of feasible policies for a consumer with storage capacity c.

An individual consumer with valuation v and storage capacity c maximizes her long-term average net utility, i.e., solves

$$\sup_{\mathbf{y}\in\mathcal{Y}_c} \liminf_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} (v x_t - p_t y_t). \tag{12}$$

In turn, the firm seeks to maximize the long-run average revenues it collects.

Optimal Consumer Stockpiling. We first analyze the consumer problem given a policy $\mathbf{p}$.

Proposition 8. *For any pricing sequence* $\mathbf{p}$, *problem* (12) *admits an optimal solution and the optimal long-run net utility is given by*

$$\liminf_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} (v - \tilde{e}_{t,c}(\mathbf{p}))^+,$$

where $\tilde{e}_{t,c}(\mathbf{p}) = \min\{p_{t-c}, p_{t-c+1}, \ldots, p_t\}$. *Furthermore, there is an optimal policy such that consumption in period* t *takes place if and only if* $v \geq \tilde{e}_{t,c}(\mathbf{p})$, *and the payment that was made for the unit consumed in period* t *is exactly* $\tilde{e}_{t,c}(\mathbf{p})$.

The proof relies on a detailed accounting of cost of a unit consumed in period t, the derivation of an upper bound on the performance of any policy, and the construction of a policy that achieves the bound.

Hence, we conclude that one may view the consumption problem of a consumer with storage capacity c in period t as one of facing an effective price of $\tilde{e}_{t,c}(\mathbf{p})$. In other words, while in the problem of one-time purchase with time windows studied in the earlier sections, the effective price faced by consumers was the minimum over a *future* time window, when consumers stockpile, the effective price associated with consumption in a given time period is the minimum price over a *past* time window, and the length of this time window is driven by the storage capacity. As Table 3 shows, the effective price table for stockpiling customer looks like the mirror image of the effective price table for who time their one time purchases. In particular, the effective prices now satisfy $\tilde{e}_{t,w}(\mathbf{p}) = \min\{\tilde{e}_{t,w-1}(\mathbf{p}), \tilde{e}_{t-1,w-1}(\mathbf{p})\}$ instead of Equation (5). As we will see, this connection enables us to adapt the framework developed in the previous section to this new setting.

Table 3 **Effective Price Table for a Stockpiling Customer for a Given Policy p**

	$t=1$	$t=2$	$t=3$	$t=4$	$t=5$	$t=6$	$t=7$	$t=8$
$w=0$	1	2	3	4	7	8	12	15
$w=1$	1	1	2	3	4	7	8	12
$w=2$	1	1	1	2	3	4	7	8
$w=3$	1	1	1	1	2	3	4	7

RIGHTSLINK

Optimal Pricing Policies. As we explicitly lay out in Proposition 8, it is possible to construct an optimal policy such that consumption in period t takes place if and only if $v \geq \tilde{e}_{t,c}(\mathbf{p})$, and the payment that was made for the unit consumed in period t is exactly $\tilde{e}_{t,c}(\mathbf{p})$. Any such policy is also optimal for any finite time horizon assuming payments are deferred to consumption times. We assume next that all consumers use such policies.

We also assume that payment is effectively made only when consumption occurs, which is without loss of generality given the long-run average revenue maximization objective. For a given pricing policy $\mathbf{p} \in \mathcal{P}$, given the consumers optimal policy for stockpiling outline above and the associated effective prices identified, the revenues collected by the firm over the first T periods may be written as

$$\sum_{t=1}^{T}\sum_{w=0}^{C}\gamma_c\int_0^{\bar{V}}\tilde{e}_{t,c}(\mathbf{p})\mathbf{1}\{v\geq\tilde{e}_{t,c}(\mathbf{p})\}\,dF(v)$$
$$=\sum_{t=1}^{T}\sum_{w=0}^{C}\gamma_c\tilde{e}_{t,c}(\mathbf{p})\bar{F}_c(\tilde{e}_{t,c}(\mathbf{p})).$$

Hence, the long-run revenue rate of the firm is given by

$$\tilde{R}(\mathbf{p})=\liminf_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}\sum_{c=0}^{C}\gamma_c\tilde{e}_{t,c}(\mathbf{p})\bar{F}_c(\tilde{e}_{t,c}(\mathbf{p})), \tag{13}$$

and the firm solves

$$\sup_{\mathbf{p}\in\mathcal{P}}\tilde{R}(\mathbf{p}). \tag{14}$$

Note that the expression for $\tilde{R}(\mathbf{p})$ is very similar to that for $R(\mathbf{p})$ in (2) with the notion of effective prices being different. One may develop a parallel concept to that of reset periods that was defined in §3. For any pricing policy, the set of periods such that $p_t = \tilde{e}_{t,C}(\mathbf{p})$ may now also be considered to be "reset periods." In any such period, all customers of type c, $c = 0, \ldots, C$, arrive with the same state of zero inventory (assuming they always break ties by purchasing at the period closest to the date of consumption). In the same manner as earlier, the system decouples from reset period to reset period. Using the policy decomposition idea (as in Lemma 2), one may establish again that when the set of prices is finite, the pricing problem (14) admits an optimal solution, and one may restrict attention to cycles of length at most $C|\mathcal{D}|$.

Theorem 4 (Equivalence). *The problem of pricing for consumers who stockpile with a population characterized by* $\{(\gamma_c, F_c)\colon c = 0, \ldots, C\}$ *(problem (14)) is equivalent to the problem of pricing to a stream of consumers who time their purchases over given time windows with characteristics* $\{(\gamma_c, F_c)\colon c = 0, \ldots, C\}$ *(problem (3)) in the following sense: both problems admit the same value function, and a cyclic policy that is optimal for one problem is also optimal for the other one.*

We prove the result below. Consider any cyclic pricing policy $\mathbf{p} \in \mathcal{P}$ of length T, with cycle $(p_1, p_2, \ldots, p_T)$. Let $\mathbf{p}^r$ denote its time reflection; it has cycle elements $(p_T, \ldots, p_2, p_1)$. Consider the effective price table $\{e_{t,c}(\mathbf{p}^r)\colon 1 \leq t \leq T, 0 \leq c \leq C\}$ corresponding to $\mathbf{p}^r$ in the one-time purchase problem with time windows and the effective price table $\{\tilde{e}_{t,c}(\mathbf{p})\colon 1 \leq t \leq T, 0 \leq c \leq C\}$ corresponding to $\mathbf{p}$ in the problem with consumer stockpiling. Note that for any t, c with $1 \leq t \leq T, 0 \leq c \leq C$, $e_{t,c}(\mathbf{p}^r) = \tilde{e}_{T-t+1,c}(\mathbf{p})$. In other words, each table is exactly the time reflection (or mirror image) of the other. This in particular implies that $R(\mathbf{p}^r) = \tilde{R}(\mathbf{p})$.

Using the reflection lemma (Lemma 3), one has that $R(\mathbf{p}^r) = R(\mathbf{p})$ and we deduce that, for any cyclic policy $\mathbf{p}$, $R(\mathbf{p}) = \tilde{R}(\mathbf{p})$. In other words, any cyclic policy that was optimal for problem (3) is also optimal for problem (14). □

Hence, the two problems are equivalent, and all the results regarding the bound on cycles, the structure of optimal policies, and the computation of optimal policies that were derived in §§3–5 apply directly to the problem of pricing for consumers who stockpile with objective (12). In addition to the direct results one obtains regarding the pricing policies that emerge for this new problem, the connection between the two fundamental problems established is also of independent interest.

7. The Generalized Finite-Horizon Case

In this section, we study a finite-horizon version of our problem and, within this context, consider several extensions of our model, including seasonality of customer demand and customers suffering a disutility from waiting to purchase a product.

Many products have obsolescence dates, such as the date when a clothing line goes out of season or when a technology product gets replaced by a newer generation model. In such situations, a consumer that arrives early in the selling season might have different preferences from one that arrives later on. Such a consumer might also be willing to wait for a lower price but suffer disutility from it. This section demonstrates that the dynamic programming approach developed in §4 may be generalized to such situations.

7.1. The Finite-Horizon Model

We consider a finite-horizon model with time $t = 1, \ldots, T$. At period t, a mass of customers $\gamma_{t,w}$ arrives with patience w, for each $w = 0, \ldots, S$. The fraction of these customers that value the product at most v

RIGHTSLINK

is given $F_{t,w}(v)$, for all $v \in [0, \bar{V}]$. Customers are assumed to have a uniform cost of waiting $c \geq 0$. That is, a customer with valuation v and patience w that arrives in period t and purchases in period $t' \in \{t, t+1, \dots, t+w\}$ earns utility $v - p_{t'} - c(t'-t)$. A potential customer who finds the current and future prices to be too high will choose to depart the system immediately without purchasing and will earn a net utility of zero. For any given price policy $\mathbf{p} = (p_1, p_2, \dots, p_T)$, the effective price faced by a consumer arriving in period t with patience w is

$$\hat{e}_{t,w}(\mathbf{p}) = \min_{t \leq k \leq t+w} p_k + c(k-t). \tag{15}$$

If there are multiple minimizers in the equation above, we assume the consumer breaks ties by buying according to the cheapest price. Let $k_{t,w}(\mathbf{p})$ be the index k that minimizes the equation above, i.e.,

$$k_{t,w}(\mathbf{p}) = \max \operatorname*{arg\,min}_{t \leq k \leq t+w} p_k + c(k-t),$$

where the max serves to break ties in favor of the cheapest price. Then, the customer delay is given by

$$d_{t,w}(\mathbf{p}) = k_{t,w}(\mathbf{p}) - t$$

and represents how many periods a customer who arrives at period t with patience w will choose to wait for a lower price. That is, among the customers who arrive at period t with patience w, the ones with valuation below $\hat{e}_{t,w}(\mathbf{p})$ will not purchase and the ones with valuation above or equal to $\hat{e}_{t,w}(\mathbf{p})$ will buy at period $t + d_{t,w}(\mathbf{p})$. Then, the firm's objective is given by

$$R_T(\mathbf{p}) = \sum_{t=1}^{T} \sum_{w=0}^{S} \gamma_{t,w} p_{t+d_{t,w}(\mathbf{p})} \bar{F}_{t,w}(\hat{e}_{t,w}(\mathbf{p})),$$

where $\bar{F}_{t,w}(v) = 1 - \lim_{v' \uparrow v} F_{t,w}(v')$. We assume that the set of prices $\mathcal{D}$ is a discretization of $[0, \bar{V}]$, with difference Δ between the prices such that c is a multiple of Δ. The firm's problem is how to optimize among policies $\mathcal{P} = \mathcal{D}^T$, i.e., $\sup_{\mathbf{p} \in \mathcal{P}} R_T(\mathbf{p})$.

Focusing on a finite-time horizon, the model we study in this section generalizes the model from §2 in two ways. First, it allows for a (uniform) cost of waiting for the customers. Second, it enables one to capture situations where the customer profiles evolve over the selling season, allowing for different mixes of patience and valuations for customers who arrive early versus customers who arrive near the end of the selling horizon. We continue to study posted price mechanisms in which the firm commits to all prices in advance.

Figure 8 (Color online) Effective Prices Table of Generalized Finite-Horizon Model

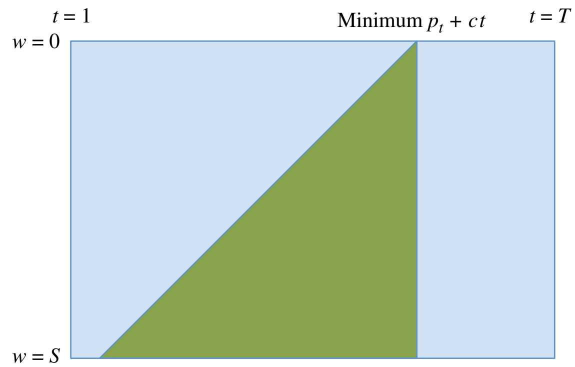

7.2. Computing Optimal Policies

In this subsection, we show how to compute optimal pricing policies for the generalized finite-horizon model. In particular, we show that a technique similar to the one developed in §4 applies to this problem.

Customers still buy according to effective prices in the generalized finite-horizon model, with the only difference being that the effective prices are modified by the cost of waiting, as given by Equation (15). Consider the effective price table and note that effective prices satisfy

$$\hat{e}_{t,w}(\mathbf{p}) = \min\{\hat{e}_{t,w-1}(\mathbf{p}), \hat{e}_{t+1,w-1}(\mathbf{p}) + c\}. \tag{16}$$

We can apply the algorithmic techniques from §4 by taking the delay costs into account. Consider the sequence of prices $\{p_k + ck: 1 \leq k \leq T\}$, and let t be the index that minimizes the value of $p_k + ck$. Then, given the recursion in Equation (16), the set of customers who find period t the most attractive to purchase will form a (potentially truncated) triangle, as in the case with no cost of waiting (see Figure 8).

Besides the modification of the effective prices to take the cost of delay into account, the other major difference in the algorithm for the generalized finite horizon problem is that we can no longer assume without loss of generality that a low price will be used at the end of a cycle (the optimal policy might not be cyclic in a finite-horizon problem). To accommodate this, we need to double the size of our state space and consider the value of rectangular regions in addition to the triangles and trapezoids introduced in §4 (as shown in Figure 8, the price that divides the set of customers may not be at the end of the selling season). The value of a region of triangular (or trapezoidal) shape starting from period m, ending in period n, and restricted to prices above $p + ci$ at position i, is given by

$$Z_{m,n}(p) = \max_{p_m, \dots, p_n \in \mathcal{D}} \sum_{w=0}^{\min\{n-m, S\}} \sum_{t=m}^{n-w} \gamma_{t,w} p_{t+d_{t,w}(\mathbf{p})} \bar{F}_{t,w}(\hat{e}_{t,w}(\mathbf{p}))$$
$$\text{s.t.} \quad p_i \geq p + ci \quad \text{for all } i \in \{m, \dots, n\},$$

RIGHTSLINK

and the value of a rectangular region starting from period m, ending in period n, and restricted to prices above p, is given by

$$\tilde{Z}_{m,n}(p) = \max_{p_m,\ldots,p_n \in \mathcal{D}} \sum_{w=0}^{\min\{n-m,S\}} \sum_{t=m}^{n} \gamma_{t,w} p_{t+d_{t,w}(\mathbf{p})} \bar{F}_{t,w}(\hat{e}_{t,w}(\mathbf{p}))$$
$$\text{s.t.} \quad p_i \geq p + ci \quad \text{for all } i \in \{m,\ldots,n\}.$$

When we construct the state space now, the price constraint is on the delay-modified effective price. The price p of the state space no longer represents a real price but instead represents what price at a fictional time 0 would impose a given constraint on the delay-modified effective prices. Therefore, we need to expand the set of price constraints in the state space to $\{-cT, -cT+\Delta, \ldots, 0, \Delta, \ldots, \bar{V}\}$. For example, if the lowest effective price in the entire horizon occurs at time t with price q, the appropriate constraint on the price at a generic time t' with price q' is $q' \geq q + c(t-t')$ regardless of whether $t > t'$ and regardless of whether $q > q'$.

We can solve this problem by solving two dynamic programs. We first compute the value of all triangular (or trapezoidal) regions according to

$$Z_{m,n}(p) = \max_{\substack{k\in\{m,\ldots,n\}\\ p'\in\mathcal{D}:\, p'\geq p+kc}} \left\{ \sum_{w=0}^{S} \sum_{t=\max\{m,k-w\}}^{\min\{k,n-w\}} \gamma_{t,w} p' \bar{F}_{t,w}(p' + c(k-t)) + Z_{m,k-1}(p'-kc) + Z_{k+1,n}(p'-kc) \right\},$$

since, as before, the value of a triangular region can be decomposed into smaller triangular regions. The value of a rectangular region with lowest effective price at periods t can be obtained by combining a triangular (or trapezoidal) region associated with periods before t, the customers served at period t, and a rectangular region associated with periods after t, as seen in Figure 8. We can thus compute the values of rectangular regions by applying

$$\tilde{Z}_{m,n}(p) = \max_{k\in\{m,\ldots,n\} p'\in\mathcal{D}:\, p'\geq p+kc} \left\{ \sum_{w=0}^{S} \sum_{t=\max\{m,k-w\}}^{k} \gamma_{t,w} p' \bar{F}_{t,w}(p' + c(k-t)) + Z_{m,k-1}(p'-kc) + \tilde{Z}_{k+1,n}(p'-kc) \right\}.$$

The revenue of the firm over the entire horizon is given by $\tilde{Z}_{1,T}(-cT)$, where the $-cT$ represents the fact that lowest delay-modified effective price of the entire horizon is unconstrained. We thus obtain Theorem 5.

Theorem 5. *An optimal pricing policy can be computed in time* $O(S^2T^3(cT + \bar{V}/\Delta))$.

The algorithm for the generalized finite horizon runs in polynomial time, but it is significantly more computationally demanding than the original dynamic program. The source of this additional computational burden is the larger state space required in the new dynamic program. The state space is larger for three reasons: we need to analyze two types of geometric figures, triangles and rectangles; we need to associate a state with each start and end period, rather than just a duration, to account for nonstationarity; we need to enlarge the set of prices in the state space to correctly restrict the delay-modified effective prices in a given state.

8. Conclusions

We studied the problem of how to choose a sequence of prices, under price commitment, given a customer population that arrives over time that is heterogeneous with regard to both valuation and patience. We established that the problem of finding optimal pricing policies is a tractable one with very few assumptions on the distribution of customers' willingness to wait and valuation, and proposed a novel geometrical approach for solving the problem. From a structural perspective, there are optimal pricing policies that are cyclical with a "short" period, in the sense that the cycle length is at most twice the maximum willingness to wait of the customer population. In addition, the class of cyclic monotone policies is generally a suboptimal one and there is an opportunity cost associated with restricting attention to cyclic monotone policies that can be significant. Optimal policies often take the form of nested sales, where the firm oscillates between targeting high-value impatient customers and targeting more patient customer classes. We have further established a form of equivalence between the above problem and that of pricing to a pool of heterogeneous consumers who may stockpile units of the products over time. This equivalence enables one to obtain the same set of structural and algorithmic results for this problem. The framework and results we present in this paper lay the ground for a potential new approach to a class of intertemporal pricing problems. Avenues for future research include the expansion of the set of problems that may be tackled through the present approach. For example, the question of how one may coordinate pricing and inventory decisions in the presence of strategic customers is a natural extension. Now rationing becomes a concern for the customers, and inventory could potentially serve as a commitment device in a model without commitment. Another interesting direction of future research pertains to the potential use of general models of strategic consumers with heterogeneous preferences for estimation purposes.

RIGHTSLINK

Supplemental Material

Supplemental material to this paper is available at http://dx.doi.org/10.1287/mnsc.2014.2049.

Acknowledgments

The authors are grateful to the department editor, an associate editor, and two reviewers for their comments that helped improve the initial version of this paper. Ilan Lobel thanks the National Science Foundation for its financial support [Grant CCF-1216004].

Appendix. Selected Proofs

Proof of Lemma 1. For any period t, if t is not a reset period, then $p_t > e_{t,S}(\mathbf{p})$, implying there exists some period $t^1 \in \{t+1, \ldots, t+S\}$ such that $p_t > p_{t^1}$. If t^1 is also not a reset period, by the same logic, there exists some $t^2 \in \{t^1+1, \ldots, t^1+S\}$ such that $p_{t^1} > p_{t^2}$. Because the set of available prices $\mathcal{D}$ is finite, one may repeat this argument only $|\mathcal{D}|$ times. Thus, at least one period in $\{t, \ldots, t+S|\mathcal{D}|\}$ is a reset period. □

Proof of Lemma 2. For any pricing policy $\mathbf{p}$ and time T, let $j(T,\mathbf{p})$ denote the integer that satisfies $V_{j(T,\mathbf{p})} < T \le V_{j(T,\mathbf{p})+1}$. That is, $j(T,\mathbf{p})$ defines the component policy being offered at time T. The value of policy $\mathbf{p}$ is given by

$$\begin{aligned} R(\mathbf{p}) &= \liminf_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T}\sum_{w=0}^{S}\gamma(w)e_{t,w}(\mathbf{p})\bar{F}(e_{t,w}(\mathbf{p})) \\ &= \liminf_{T\to\infty} \frac{1}{T}\Bigg(\sum_{i=0}^{j(T,\mathbf{p})} R(C_i(\mathbf{p}))L_{C_i(\mathbf{p})} \\ &\qquad + \sum_{t=V_{j(T,\mathbf{p})}+1}^{T}\sum_{w=0}^{S}\gamma(w)e_{t,w}(\mathbf{p})\bar{F}(e_{t,w}(\mathbf{p}))\Bigg), \end{aligned}$$

where the revenue of the policy up to time T is composed of the revenue of the component policies, for components from 0 to $j(T,\mathbf{p})$, plus a leftover term starting from period $V_{j(T,\mathbf{p})+1}$ up to period T for any leftover periods not completely covered in a component policy by time T. The average revenue obtained from the leftover component

$$\lim_{T\to\infty}\frac{1}{T}\sum_{t=V_{j(T,\mathbf{p})}+1}^{T}\sum_{w=0}^{S}\gamma(w)e_{t,w}(\mathbf{p})\bar{F}(e_{t,w}(\mathbf{p})) = 0,$$

since the per-period revenue is bounded and the leftover term includes at most $S|\mathcal{D}|$ periods by Lemma 1. Therefore, the value obtained from policy $\mathbf{p}$ is nothing but an average of the values obtained by the component policies weighted by their lengths, i.e.,

$$R(\mathbf{p}) = \liminf_{T\to\infty}\frac{1}{T}\sum_{i=0}^{j(T,\mathbf{p})} R(C_i(\mathbf{p}))L_{C_i(\mathbf{p})},$$

and the revenue from the policy $\mathbf{p}$ is bounded by the supremum among the revenues of all the component policies of $\mathbf{p}$. □

Proof of Proposition 1. By Lemma 1, every component pricing policy is a cyclic policy with length at most $S|\mathcal{D}|$. Thus, the set of possible component policies is finite and its cardinality is bounded by $\sum_{i=1}^{S|\mathcal{D}|}|\mathcal{D}|^i$, implying that the supremum in Equation (4) is attained. Therefore, by the policy decomposition lemma, there exists a cyclic policy with length at most $S|\mathcal{D}|$ that maximizes the seller's revenue. □

Proof of Lemma 3. Let $\mathbf{p}$ be any cyclic policy with length T. It is convenient to define the natural extension of $\mathbf{p}$ to the set of nonpositive indices. Since $\mathbf{p}^r$ is a reflection of $\mathbf{p}$, $p_k^r = p_{T+1-k}$ for any k. For any time t and willingness-to-wait w, the effective prices of the policy $\mathbf{p}$ and its reflection $\mathbf{p}^r$ satisfy

$$\begin{aligned} e_{T+1-t-w,w}(\mathbf{p}^r) &= \min_{T+1-t-w\le k\le T+1-t} p_{T+1-k} \\ &= \min_{t\le k'\le t+w} p_{k'} = e_{t,w}(\mathbf{p}), \end{aligned} \tag{17}$$

where the second equality is obtained by a change of variables $k' = T+1-k$. The revenue obtained from customers with willingness-to-wait w under the reflected policy $\mathbf{p}^r$ is

$$\begin{aligned} R_w(\mathbf{p}^r) &= \frac{\gamma(w)}{T}\sum_{t=1}^{T} e_{t,w}(\mathbf{p}^r)\bar{F}_w(e_{t,w}(\mathbf{p}^r)) \\ &= \frac{\gamma(w)}{T}\sum_{t=1}^{T} e_{T+1-t-w,w}(\mathbf{p}^r)\bar{F}_w(e_{T+1-t-w,w}(\mathbf{p}^r)) \\ &= \frac{\gamma(w)}{T}\sum_{t=1}^{T} e_{t,w}(\mathbf{p})\bar{F}_w(e_{t,w}(\mathbf{p})) = R_w(\mathbf{p}), \end{aligned}$$

where the second equality follows from the cyclic nature of the policy and the third is derived from Equation (17). By summing over all w from 0 to S, we obtain the desired result, i.e., $R(\mathbf{p}^r) = \sum_{w=0}^{S} R_w(\mathbf{p}^r) = \sum_{w=0}^{S} R_w(\mathbf{p}) = R(\mathbf{p})$. □

Proof of Theorem 1. At first, assume the set of prices $\mathcal{D}$ is finite. Then, by Proposition 1, there exists an optimal cyclic policy. Let $\mathbf{p}$ be an optimal policy such that its length $T = L_{\mathbf{p}}$ is minimal among all optimal policies. Assume, without loss of generality, that $p_T = \min_{1\le k\le T} p_k$. Let $p' = \min_{1\le k\le T-1} p_k$ denote the second-lowest price in the policy. There can be no reset periods in $\{1, \ldots, T-1\}$, otherwise a shorter optimal policy would exist by the policy decomposition lemma. Therefore, the lowest price in the policy is used only once per cycle, i.e., $p_T < p'$. Let t' and t'' represent respectively the first and last periods in $k \in \{1, \ldots, T-1\}$ such that $p_k = p'$. Then, $t' \ge T-S$, since otherwise t' would be a reset period. Consider now the time- reflected policy $\mathbf{p}^r$, which yields the same revenue as $\mathbf{p}$ by the reflection lemma. In the reflected policy, the first time price p' is used is at $T+1-t''$, and the next time the lowest price in the policy $(\min_{1\le k\le T-1} p_k)$ is used is at time $T+1$ (the first period of the next cycle). If $T+1-t''$ was a rest period, one could construct a policy that garners weakly higher revenues with shorter cycles, which would contradict the fact that the policy we started with was optimal with minimal length among optimal policies. We deduce that $T+1-t''$ cannot be a reset period of the reflected policy, i.e., $t'' \le S$. Combining the bounds on t' and t'', we obtain that $S \ge t'' \ge t' \ge T-S$ and, therefore, $T \le 2S$. This completes the proof of the theorem for the case in which $|\mathcal{D}| < \infty$.

Now let the set of prices $\mathcal{D}$ be an arbitrary closed subset of $[0, \bar{V}]$. Let $\{\mathbf{p}^i\}_{i\in\mathbb{N}}$ be a sequence of feasible price

RIGHTSLINK

sequences in $\mathcal{P}$ such that $R(\mathbf{p}^k) \geq R^* - 1/k$, where $R^* = \sup_{\mathbf{p}\in\mathcal{P}} R(\mathbf{p})$. Let $V^k = \lceil \bar{V}k \rceil / k$ and let $\hat{p}^k$ be a sequence of prices such that each p_t^k is equal to p_k rounded down to the closest element in the set $\{0, 1/k, 2/k, \ldots, V^k - 2/k, V^k - 1/k\}$. Since the functions $F_w(\cdot)$ are Lipschitz continuous for all w and the effective prices $e_{t,w}(\mathbf{p})$ are also Lipschitz continuous functions of the relevant prices $p_t, \ldots, p_{t+w}$, the revenue function $R(\cdot)$ is Lipschitz continuous in the infinity norm, so there exists a constant L such that $|R(\mathbf{p}^k) - R(\hat{\mathbf{p}}^k)| \leq L \sup_{t\in\mathbb{N}} |p_t - \hat{p}_t|$. Therefore, $R(\hat{\mathbf{p}}^k) \geq R^* - L/k$. By the first part of the proof, we can obtain a policy $\tilde{\mathbf{p}}^k$ such that $R(\tilde{\mathbf{p}}^k) \geq R(\hat{\mathbf{p}}^k)$ where $\tilde{\mathbf{p}}^k$ is a component policy of $\hat{\mathbf{p}}^k$ and $\tilde{\mathbf{p}}^k$ has cycle length at most $2S$. By construction, any price in $\tilde{\mathbf{p}}^k$ is at most $1/k$ away from the set $\mathcal{D}$. Using the Lipschitz continuity of $R(\cdot)$ again, we can construct a cyclic pricing policy $\bar{p}^k \in \mathcal{P}$ with period at most $2S$ such that $R(\bar{\mathbf{p}}^k) \geq R^* - 2L/k$. By considering the limit as k goes to infinity, we observe that the supremum revenue among all cyclic policies in $\mathcal{P}$ with length at most $2S$ is also R^*. Since the set $\mathcal{D}$ is closed and bounded, the set of cyclic policies with length at most T with $1 \leq T \leq 2S$ is compact. The problem of maximizing the average revenue over all policies with cycle exactly $T = 1, \ldots, 2S$ is a maximization problem of a continuous function over a compact set. Therefore, the supremum over the set of cyclic policies with length at most $2S$ is attained. □

Proof of Theorem 2. The text preceding the statement of the theorem explains why the Bellman equation in Equation (6) applies to the value function $W_n(p)$, and we now construct an algorithm that uses it to find optimal pricing policies. The first step of the algorithm is to compute the value of $\sum_{w=i}^{j} \gamma(w) p \bar{F}_w(p)$ for all i, j, and p, which takes $O(|\mathcal{D}|S^2)$ steps if done recursively. Now, notice that for any fixed n and k, the sequence $t_{n,k,w}$ always takes the form $(1, 2, 3, \ldots, n-1, n, n, \ldots, n, n, n-1, \ldots, m+1, m)$, with a first part where numbers increase by 1 in each step, a second part where numbers stay flat, and a third part where they decrease by 1 in each step. Therefore, we can recursively compute the values of $\sum_{w=0}^{S} t_{n,k,w} \gamma(w) p \bar{F}_w(p)$ for all n, k, and p in $O(|\mathcal{D}|S^2)$ by utilizing the previously computed values of $\sum_{w=i}^{j} \gamma(w) p \bar{F}_w(p)$ as building blocks. Next, define the auxiliary value function $\tilde{W}_n(p)$ as

$$\tilde{W}_n(p) = \max_{k\in\{1,\ldots,n\}} \left\{ \sum_{w=0}^{S} t_{n,k,w} \gamma(w) p \bar{F}_w(p) + W_{k-1}(p) + W_{n-k}(p) \right\}$$

so that $W_n(p) = \max_{p'\in\mathcal{D}: p' \geq p} \tilde{W}_n(p')$. Let $W_0(p) = \tilde{W}_0(p) = 0$ for any $p \in \mathcal{D}$. Suppose the values of $W_n(p)$ are known for all values of n up to $T-1$ and all values of $p \in \mathcal{D}$. Then, the values of $\tilde{W}_T(p)$ can be computed by the equation above for $n = T$ and all $p \in \mathcal{D}$ in time $O(|\mathcal{D}|T)$. With the values of $\tilde{W}_n(p)$ on hand for all p, we can compute the values of $W_n(p)$ for all $p \in \mathcal{D}$ in time $O(|\mathcal{D}|)$. By Theorem 1, it is sufficient to consider policies of cycle length up to $2S$. Repeating this process for all T from 1 to $2S-1$ takes time $O(|\mathcal{D}|S^2)$. With the value functions computed, the optimal policy of length T can be computed for each T by Equation (7) and the actual overall optimal policy can determined by Equation (8). Since these last two operations take $O(|\mathcal{D}|S)$ time, the overall computational complexity of this algorithm is $O(|\mathcal{D}|S^2)$. □

Proof of Theorem 5. This algorithm works very similarly to the infinite time horizon algorithm from Theorem 2, except that this generalized model does not have some of the symmetry features of the original problem that allow for speedier computation. The value of

$$\sum_{w=0}^{S} \sum_{t=\max\{m,\,k-w\}}^{\min\{k,\,n-w\}} \gamma_{t,w} p' \bar{F}_{t,w}(p' + c(k-t))$$

can be computed for any m, n, k and p in $O(S^2)$ as the second summation has at most S terms. Therefore, it can be computed for all possible parameters in time $O(|\mathcal{D}|S^2T^3)$, where $|\mathcal{D}| = cT + \bar{V}/\Delta$. We can then use dynamic programming to compute the value of $Z_{m,n}(p)$ for all possible parameters in less time than $O(S^2T^3(cT + \bar{V}/\Delta))$. We can then repeat the same process with the values of

$$\sum_{w=0}^{S} \sum_{t=\max\{m,\,k-w\}}^{k} \gamma_{t,w} p' \bar{F}_{t,w}(p' + c(k-t))$$

to obtain the value of $\tilde{Z}_{m,n}(p)$ for all problem parameters, thus finding the optimal pricing policy. □

References

Ahn H, Gümüs M, Kaminsky P (2007) Pricing and manufacturing decisions when demand is a function of prices in multiple periods. *Oper. Res.* 55(6):1039–1057.

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

Aviv Y, Levin Y, Nediak M (2011) Counteracting strategic consumer behavior in dynamic pricing systems. Netessine S, Tang CS, eds. *Consumer-Driven Demand and Operations Management Models: A Systematic Study of IT-Enabled Sales Mechanisms*, Chap. 12 (Springer, New York), 323–352.

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

Blattberg R, Eppen G, Lieberman J (1981) A theoretical and empirical evaluation of price deals for consumer nondurables. *J. Marketing* 45(1):116–129.

Board S (2008) Durable-goods monopoly with varying demand. *Rev. Econom. Stud.* 75(2):391–413.

Borgs C, Candogan O, Chayes J, Lobel I, Nazerzadeh H (2014) Optimal multiperiod pricing with service guarantees. *Management Sci.* 60(7):1792–1811.

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

Conlisk J, Gerstner E, Sobel J (1984) Cyclic pricing by a durable goods monopolist. *Quart. J. Econom.* 99:489–505.

Deb R (2010) Intertemporal price discrimination with stochastic values. Working paper, University of Toronto, Toronto, ON.

Garrett D (2011) Durable goods sales with dynamic arrivals and changing values. Working paper, Toulouse School of Economics, Toulouse, France.

Hendel I, Nevo A (2006) Measuring the implications of sales and consumer inventory behavior. *Econometrica* 74(6):1637–1673.

Hendel I, Nevo A (2013) Intertemporal price discrimination in storable goods markets. *Amer. Econom. Rev.* 103(7):2722–2751.

Jeuland A, Narasimhan C (1985) Dealing—Temporary price cuts—By seller as a buyer discrimination mechanism. *J. Bus.* 58(3): 295–308.

Li J, Granados N, Netessine S (2014) Are consumers strategic? Evidence from the air-travel industry. *Management Sci.* 60(9): 2114–2137.

RIGHTSLINK

Mierendorff K (2011) Optimal dynamic mechanism design with deadlines. Working paper, University of Zurich, Zurich.

Nair H (2007) Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video-games. *Quant. Marketing Econom.* 5(3):239–292.

Pai MM, Vohra R (2013) Optimal dynamic auctions and simple index rules. *Math. Oper. Res.* 38(4):682–697.

Pesendorfer M (2002) Retail sales: A study of pricing behavior in supermarkets. *J. Bus.* 75(1):33–66.

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

Sobel J (1991) Durable goods monopoly with entry of new consumers. *Econometrica* 59(5):1455–1485.

Stokey NL (1979) Intertemporal price discrimination. *Quart. J. Econom.* 93(3):355–371.

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

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

Su X (2010) Intertemporal pricing and consumer stockpiling. *Oper. Res.* 58(4, Part 2):1133–1147.

Talluri KT, van Ryzin GJ (2005) *Theory and Practice of Revenue Management* (Springer-Verlag, New York).

Yin R, Aviv Y, Pazgal A, Tang CS (2008) Optimal markdown pricing: Implications of inventory display formats in the presence of strategic customers. *Management Sci.* 55(8): 1391–1408.

RIGHTSLINK