Optimal Dynamic Pricing with Patient Customers

Yan Liu
Research Center for Management Science and Data Analytics
School of Information Management and Engineering
Shanghai University of Finance and Economics
Shanghai 200433, China

William L. Cooper
Department of Industrial and Systems Engineering
University of Minnesota
Minneapolis, MN 55455

Appeared in *Operations Research*, 63(6), pp. 1307–1319, 2015

Abstract

We consider an infinite-horizon single-product pricing problem in which a fraction of customers is patient and the remaining fraction is impatient. A patient customer will wait up to some fixed number of time periods for the price of the product to fall below his or her valuation at which point the customer will make a purchase. If the price does not fall below a patient customer's valuation at any time during those periods, then that customer will leave without buying. In contrast, impatient customers will not wait, and either buy immediately or leave without buying. We prove that there is an optimal dynamic pricing policy comprised of repeating cycles of decreasing prices. We obtain bounds on the length of these cycles, and we exploit these results to produce an efficient dynamic programming approach for computing such an optimal policy. We also consider problems in which customers have variable levels of patience. For such problems, cycles of decreasing prices may no longer be optimal, but numerical experiments nevertheless suggest that such a decreasing cyclic policy (suitably chosen) often performs quite well.

Keywords: Pricing, Consumer Behavior, Dynamic Programming

1 Introduction

When faced with a posted price for a particular product, some customers may consider making a purchase but may judge that price to be too high. If such customers are willing to wait, then they may possibly purchase the product in the future if the price drops enough. In this paper, we focus on developing dynamic pricing policies that take into account such consumer behavior. Specifically, we consider a deterministic single-product pricing problem in which customers arrive dynamically and a portion of customers is patient. A patient customer is willing to wait up to some fixed amount of time and will purchase the product as soon as the price falls below his or her individual valuation.

It stands to reason that some customers may be willing to delay a purchase until the price is sufficiently low. As a practical matter, the specific behavior exhibited by our patient customers (waiting for the price to fall below their valuations) can be supported by a variety of web services and smartphone applications that track prices on, for example, Amazon.com. Such services and applications (sometimes called shopping bots) allow an individual customer to enter a particular price for a product of interest and to receive a notification whenever the price of that product falls below that particular price on Amazon.com. Upon receiving such a notification, the customer can then purchase the product in question.

Our main result is that for a seller with an infinite-horizon average-revenue objective, there is an optimal policy comprised of repeating cycles of decreasing prices. This result holds for general customer valuation distributions under the assumptions that (i) a fixed fraction of customers is patient and that such patient customers wait up to an arbitrary fixed number k of time periods, (ii) the remaining customers are impatient and willing only to make immediate purchases or to not buy, and (iii) that prices are selected from a finite set. Under additional mild assumptions, the result also holds when prices are selected from a continuous interval. We obtain bounds on the length of an optimal cycle, and we exploit a connection to a particular finite-horizon version of the problem to develop an efficient approach for computing an optimal policy.

The proof of our main result proceeds roughly as follows. We first reduce the infinite-horizon average-revenue problem to a related problem of maximizing the finite-horizon average revenue over a set of policies that do not have what are termed regeneration points. A regeneration point arises when prices are set so that there is a time point (the regeneration point) with the property that no customers who arrive before that time point buy at that time point or later. The maximization requires that we consider multiple different lengths for the finite horizon. The optimal policy for the infinite-horizon problem is then constructed by indefinitely repeating the optimal policy of the related finite-horizon problem.

The related problem is itself difficult, and our analysis of it constitutes the core of our technical contribution. We establish through a fairly lengthy argument that there is an optimal policy comprised of decreasing prices for the related finite-horizon problem. At a high level, we reach this conclusion by making a series of comparisons between policies in which prices are judiciously added, removed, or rearranged. A key concept that we introduce for our analysis is what we call a "strong markup," which occurs (roughly speaking) whenever there is a price increase immediately after a price that is a strict running minimum. We establish that any feasible policy with price increases is made no worse by re-ordering the prices from high to low between each successive pair of strong markups. Such a re-ordering yields a policy composed of a series of decreasing strings of prices, arranged so that price increases occur only immediately after running minima. We then argue that the best policy composed in such a way of decreasing strings is no better than a policy that is just (a slightly modified version of) its final decreasing string. We stress that the preceding is a loose explanation. Before moving forward, we note that often-used inductive methods for proving structural properties of optimal policies for pricing or inventory problems (see, for example, Section 5.4 of Porteus 2002) were not fruitful for our problem, so we instead needed to come up with a different approach.

We also consider problems in which customers have variable levels of patience. For such problems, there is again an optimal policy that is cyclic. However, in this case it may be necessary to consider cycles in which prices do not monotonically decrease. This turns out to render as generally intractable the problem of computing an optimal policy in the presence of variable patience levels. Consequently, it is important to identify good, computationally-tractable heuristics for such settings. Motivated by the optimality of repeating cycles of decreasing prices for problems with fixed patience levels, we propose using those policies as just such a heuristic. We establish that it is not difficult to compute the best policy with repeating cycles of decreasing prices for a problem with variable patience levels. In addition, we obtain a computable upper bound (expressed in terms of solutions to a suitable collection of problems with fixed patience levels) on the value of an optimal policy (among all policies), which allows us to bound the performance of the proposed heuristic relative to an optimal policy. Numerical experiments suggest that the heuristic is near optimal in many cases.

The remainder of the paper is organized as follows. Section 2 reviews related literature. Section 3 introduces the model and states our main result. Section 4 contains the reduction to the related finite-horizon problem. Section 5 outlines the key steps in the analysis of that finite-horizon problem. Section 6 contains refinements and extensions of our main result as well as a computational algorithm. Section 7 describes bounds for problems with variable patience levels. Section 8 presents results of numerical experiments. Section 9 contains concluding comments. Section A of the online appendix provides comparisons with some of the literature in a tabular format. Sections B–G of the online appendix contain complete details of the proofs.

2 Related Literature

The model of patient customer behavior that we employ is essentially that studied by Ahn et al. (2007), who consider both production and pricing decisions as well as inventory carrying costs. They consider only uniformly distributed customer valuations (linear demand curves in their terminology), but otherwise our model of customer behavior is conceptually the same as theirs. For pure pricing settings like those we consider without production decisions or inventory or capacity issues, they establish that it is optimal to alternate between a high price and a low price for a finite-horizon problem with uniformly distributed customer valuations and customers that wait either zero or one period. Such a pricing pattern can be viewed as a decreasing cyclic policy with two-period cycles. The assumption that customers wait either zero or one period is the same as assuming in our setup that a fraction of customers is impatient and the remainder is patient and willing to wait just one period. It follows from their result on finite-horizon problems that cycles of alternating high and low prices are also optimal for the infinite-horizon average-revenue problems that we consider (again under the added conditions that valuations are uniformly distributed and the "patience level" is $k = 1$). Ahn et al. also obtain a recursive expression for the best decreasing price sequence of any given fixed length for problems with linear demand curves. They motivate this development by commenting that for long finite horizons, policies comprised of repeated decreasing price cycles should work well. Our results show that policies of that form are in fact optimal for infinite-horizon average-revenue problems.

The customer behavior we consider may be compared with that of what are often called strategic customers. Assumptions vary, but broadly speaking, strategic customers try to anticipate price changes and time purchases in an attempt to maximize their utility. In contrast, our patient customers do not look into the future and do not attempt to maximize utility. Rather, they may be viewed as making purchasing decisions that satisfice. One may imagine that our patient customers simply lack the motivation or the analytical sophistication of strategic customers. There is a considerable body of literature on satisficing dating back to early work of Simon (1955). Recent experimental studies that consider satisficing in consumer search and consumer choice include Caplin et al. (2011), Reutskaja et al. (2011), and Stüttgen et al. (2012). We note in passing that our use of the term "patient" differs from that found in some prior work. For example, Su (2007) uses the terms patient and impatient to distinguish how much it costs different customers to monitor prices; however, both patient and impatient customers in his paper attempt to maximize utility and hence would be called strategic in our terminology.

There is now a growing body of scholarship in the operations research and operations management literature that deals with intertemporal pricing and inventory control in the presence of strategic customers. There is also a considerable history of research on such problems in the economics literature. We will not attempt to summarize this work, but instead refer to reviews by Shen and Su (2007) and Aviv et al. (2009). Formulations with strategic customers often involve a game between a seller and customers. The presence of strategic customers may either reduce or increase the seller's revenue in comparison to problems without strategic customers; see, e.g., Li et al. (2014). There is also research that considers operations management problems in which some customers exhibit bounded rationality. See Huang et al. (2013) for an entry into this literature. Specifics vary across a wide array of models and problem contexts, but at a high level, boundedly rational customers do not make decisions that maximize their utility but instead use heuristics or "rules of thumb." The customers in our model can be viewed as acting in this manner.

Besbes and Lobel (2015) study a model similar to ours. They consider customers with varying levels of patience and establish that a cyclic pricing policy is optimal for an infinite-horizon average-revenue problem. In our work, optimal cycles are made up of decreasing prices whereas the prices in the optimal cycles of Besbes and Lobel need not be decreasing. They also obtain bounds on the cycle lengths. Their model of consumer behavior differs from ours. Specifically, they assume that customers are strategic and that each customer will purchase at the minimum price within his window of patience so long as that minimum is below his valuation. Such customers are forward-looking in the sense that they know the seller's future prices. Models with forward-looking customers are appropriate if customers have fully learned about the seller's pricing behavior, e.g., through previous interactions. In contrast, we assume that a customer buys as soon as the price falls below that customer's valuation. Our customers are not forward looking. A model such as ours is appropriate when customers do not know the seller's behavior or are not able to learn it.

Conlisk et al. (1984) study a monopolist that sells a durable good to customers who stay in the market (possibly indefinitely) until making a purchase. Each period, there arrives a fixed amount of such customers, each with one of two possible valuations for the good. Customers time their purchases to maximize their utility. The main results state that there is an optimal decreasing cyclic pricing policy in which the monopolist sets prices so that only high-valuation customers buy, except in the last period of a cycle when the price is set low enough that both high- and low-valuation customers make purchases. To draw some distinctions with our general model, we consider customers who are patient but not strategic and who wait at most a finite span of time to purchase. We also allow general valuation distributions, which means that customers in our model cannot be divided into only two categories of high and low valuation.

Table 4 in Section A of the online appendix summarizes central assumptions and results from our paper as well as from those papers cited above that establish optimality of cyclic pricing policies (Ahn et al. 2007, Besbes and Lobel 2015, Conlisk et al. 1984).

Some of the literature on intertemporal pricing focuses on finite-horizon problems in which all customers are present at the beginning of the horizon. For instance, Besanko and Winston (1990) consider such a problem in which customers time their purchases to maximize utility. Each customer has a valuation drawn from a uniform distribution. Besanko and Winston show that in a game-theoretic equilibrium, the prices set by the seller decline over time, resulting in what is termed price skimming, which allows the seller to exploit differences in customers' valuations to capture some of the consumers' surplus. The decreasing price cycles we obtain can be viewed similarly. The authors also make comparisons with a version of the problem with myopic customers (akin to our patient customers) who are assumed to buy as soon as the price falls below their valuations.

A few other recent papers have considered deterministic discrete-time pricing problems in which demand depends upon prices in multiple periods. Gümüş et al. (2013) consider a similar model to Ahn et al. (2007), but with multiple partially substitutable products. Popescu and Wu (2007) consider a setting in which demand depends upon the current price and a reference price, which is a function (an exponentially weighted average) of previous prices. Their setup yields a dynamic program with a single-dimensional state, which corresponds to the current reference price. They establish that prices converge monotonically to a steady-state price if customers are loss averse. If the starting reference price is higher than the steady state, then prices decrease in what can be viewed as a price skimming policy. For gain-seeking customers, they show that optimal prices cycle. Kopalle et al. (1996) obtain similar results. Nasiry and Popescu (2011) consider a pricing problem with loss-averse customers and a different model of reference price formation, called peak-end anchoring. They also find that optimal prices converge monotonically. For both models of reference price formation, optimal prices remain fixed in steady state when facing loss-averse customers. This stands in contrast to the cyclic price policies we identify. Caro and Martínez-de-Albéniz (2012) consider competing sellers that face customers whose previous consumption of products may affect their current demand for those products through a satiation effect. In their model, consumers' consumption levels converge to a steady state. They address both pricing and product design decisions, but they do not consider dynamic pricing, which is the focus of our study.

3 Model Setup and Central Result

We study a multi-period single-product pricing problem with deterministic demand and unlimited inventory. We assume that demand is a continuous quantity and that units are scaled so that in a single period a potential new demand of 1 arrives. This demand is comprised of infinitesimal customers, a fraction $\alpha \in (0,1]$ of whom are patient and a fraction $1-\alpha$ of whom are impatient.

Each patient customer has a non-negative valuation drawn from a distribution $\mathcal{F}(\cdot)$. For each x, let $F(x) = \lim_{y\uparrow x} \mathcal{F}(y)$ be the left limit of $\mathcal{F}(\cdot)$ at x. Hence, $F(x)$ is the fraction of patient customers whose valuation is less than x, and $1-F(x)$ is the fraction of patient customers whose valuation is at least x. If $\mathcal{F}(\cdot)$ is continuous, then $F(\cdot) = \mathcal{F}(\cdot)$. Likewise, each impatient customer has a non-negative valuation drawn from a distribution $\mathcal{F}_0(\cdot)$, and we similarly define $F_0(x) = \lim_{y\uparrow x} \mathcal{F}_0(y)$. Let $G(x) = \alpha F(x) + (1-\alpha)F_0(x)$ be the overall fraction of customers with a valuation less than x.

Each period the seller offers one price. If the price offered is no greater than a particular customer's valuation, then that customer (whether impatient or patient) will make a purchase immediately and then leave. If the price is above that customer's valuation, then the customer's subsequent behavior depends upon whether the customer is patient or impatient. If the customer is impatient, then he will simply leave the market without purchasing. If the customer is patient, then he will wait for up to k more periods. In those k periods, the patient customer will make a purchase as soon as the price falls to or below his valuation. If the price remains above the customer's valuation for the full k periods, then the patient customer will leave without making a purchase. We assume that $k \geq 1$ is fixed.

Prices are selected from a finite set $\mathcal{P} = \{p(1), \ldots, p(m)\}$ with cardinality $m \geq 2$. We will be interested in both finite-horizon and infinite-horizon problems. Our analysis of the former will be useful for deriving our main results, which deal with the latter. For an infinite-horizon problem, a sequence $p = (p_1, p_2, \ldots) \in \mathcal{P}^\infty$ is called a pricing policy or simply a policy. For finite $L \in \mathbb{N} = \{1, 2, 3, \ldots\}$, we will also call $p = (p_1, \ldots, p_L) \in \mathcal{P}^L$ a policy for the finite-horizon L-period problem, and sometimes use the notation $L(p)$ to denote the length of p. For $p = (p_1, \ldots, p_{L_1}) \in \mathcal{P}^{L_1}$ and $q = (q_1, \ldots, q_{L_2}) \in \mathcal{P}^{L_2}$ with $L_1, L_2 \in \mathbb{N}$ we will use (p, q) to denote the policy $(p_1, \ldots, p_{L_1}, q_1, \ldots, q_{L_2}) \in \mathcal{P}^{L_1+L_2}$ that implements price p_t for $t \in \{1, \ldots, L_1\}$ and price q_{t-L_1} for $t \in \{L_1+1, \ldots, L_1+L_2\}$. Similarly, for $q \in \mathcal{P}^{L_2}$, we will use $(q, q, q, \ldots) \in \mathcal{P}^\infty$ to denote the infinite-horizon policy that charges q_t in periods $t, t+L_2, t+2L_2, t+3L_2, \ldots$ for $t = 1, \ldots, L_2$.

For $p \in \mathcal{P}^L$ with $L \in \mathbb{N} \cup \{\infty\}$, let

$$\rho_t(p) = p_t\Big[1 - G(p_t) + \alpha \sum_{i=1}^{k} [F(\min\{p_{t-i}, \ldots, p_{t-1}\}) - F(p_t)]^+\Big] \qquad \text{for } t = 1, 2, \ldots \tag{1}$$

where we use the convention that $p_t = 0$ for $t \leq 0$. Throughout, $x^+ = \max\{x, 0\}$. The quantity $\rho_t(p)$ represents the revenue accrued in period t under policy p. In (1), $1 - G(p_t)$ represents the number of customers that arrive in period t and immediately make a purchase. The expression $[F(\min\{p_{t-i}, \ldots, p_{t-1}\}) - F(p_t)]^+$ in (1) represents the fraction of patient customers that arrive in period $t - i$ and that have a valuation that is less than all the prices in periods $t - i, \ldots, t - 1$ but greater than or equal to the price in period t. Keeping in mind that α is the fraction of customers that is patient, we see that $\alpha[F(\min\{p_{t-i}, \ldots, p_{t-1}\}) - F(p_t)]^+$ represents the number of customers that initially arrive in period $t - i$ and subsequently make a purchase in period t.

We are now ready to present the objective function of the seller. For $p \in \mathcal{P}^\infty$ let

$$H(p) = \liminf_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} \rho_t(p)$$

denote the infinite-horizon long-run average revenue from implementing policy $p \in \mathcal{P}^\infty$. The limit inferior above will be a limit for the optimal policy we identify. The seller's goal is to select a pricing policy to maximize the long-run average revenue, that is, it wants to solve

$$\sup_p \{H(p) : p \in \mathcal{P}^\infty\}. \tag{2}$$

As in Besbes and Lobel (2015), we say that a policy $p \in \mathcal{P}^\infty$ is *cyclic* if there exists a positive integer L such that $p_{t+L} = p_t$ for all $t \in \mathbb{N}$. The smallest $L > 0$ for which this holds is the *cycle length* of policy p. A cyclic policy with cycle length L is of the form $p = (q, q, q, \ldots)$ where $q \in \mathcal{P}^L$. A cyclic policy p is said be a *decreasing cyclic policy* if $p = (q, q, q, \ldots)$ where $q = (q_1, \ldots, q_L)$ is such that $q_1 \geq q_2 \geq \cdots \geq q_L$. (We use decreasing to mean weakly decreasing.)

Our central result is the following.

Theorem 1 *There exists a decreasing cyclic policy with cycle length $L \in \{1, \ldots, m + k - 1\}$ that is an optimal solution to* (2).

In Section 6.2, under some additional mild assumptions, we obtain a lower bound of $k + 1$ on the cycle length. Based upon numerical results in Section 8, we also conjecture that there are broad conditions under which there is an optimal policy with decreasing cycles with length of either 1 (fixed price), $k + 1$, or $k + 2$.

In the next two sections, we build up to the proof of the preceding theorem. The proof involves reducing the problem to a suitable finite-horizon problem in Section 4 and then establishing that the finite-horizon problem has a decreasing optimal policy in Section 5.

4 Preliminary Analysis

The first step in the proof of the theorem above is to show that there is an optimal solution to the infinite-horizon problem (2) that cyclically repeats an optimal solution to the finite-horizon problem that is given in (5) below. To this end, we introduce some definitions. For $L \in \mathbb{N}$ and $p \in \mathcal{P}^L$ let $V_L(p)$ and $v_L(p)$ denote, respectively, the total revenue and average revenue accrued over horizon $1, \ldots, L$ under policy p; that is,

$$V_L(p) = \sum_{t=1}^{L} \rho_t(p) \qquad \text{and} \qquad v_L(p) = V_L(p)/L.$$

Let $V(p) = V_{L(p)}(p)$ and $v(p) = v_{L(p)}(p)$. Throughout, we will show the "length subscript" on $V(\cdot)$ or $v(\cdot)$ when we wish to emphasize the length of the finite-horizon policy in question.

For $p \in \mathcal{P}^\infty$ we say that time $t \geq 2$ is an *I-regeneration point* of p if

$$\min\{p_t, p_{t+1}, \ldots, p_{t+k-1}\} \geq p_{t-1}.$$

For $L \in \mathbb{N}$ and $p \in \mathcal{P}^L$ we call time $t \in \{2, \ldots, L\}$ an *F-regeneration point* of p if

$$\min\{p_t, p_{t+1}, \ldots, p_{\min\{t+k-1,L\}}\} \geq p_{t-1}. \tag{3}$$

If t is a regeneration point, no customers who join the market prior to time t purchase in time t or later, because the prices from time t to time $t+k-1$ are greater than or equal to the price in time $t-1$. The modifiers "I-" and "F-" will help later to distinguish between infinite price sequences and finite price sequences. Observe that by definition, time $t = 1$ is not a regeneration point.

For any policy $p \in \mathcal{P}^\infty$, we call the sequence of prices between any two successive I-regeneration points a *component* of p. Let $r(i)$ denote the ith I-regeneration point. Then $p_{r(i)}$ is the price at the ith I-regeneration point. We refer to the sequence $C_i(p) := (p_{r(i)}, p_{r(i)+1}, \ldots, p_{r(i+1)-1})$ as the ith component of p for $i \geq 1$. By definition of $r(i)$ and $r(i+1)$, there are no I-regeneration points in time periods $r(i)+1, \ldots, r(i+1)-1$. For simplicity, we also define $r(0) = 1$ and refer to $(p_{r(0)}, \ldots, p_{r(1)-1})$ as $C_0(p)$. Both Ahn et al. (2007) and Besbes and Lobel (2015) use regeneration points and components in their work, and our reduction of (2) to (5) below parallels an argument of Besbes and Lobel.

The following lemma provides a link between the infinite-horizon and finite-horizon problems.

Lemma 1 *Fix $L \in \mathbb{N}$, and consider $q \in \mathcal{P}^L$ and $p = (q, q, q, \ldots) \in \mathcal{P}^\infty$.*

1. *$C_i(p) = q$ for $i = 0, 1, 2, \ldots$ if and only if q has no F-regeneration points.*

2. *If q has no F-regeneration points, then $H(p) = v(q)$.*

3. If q has no F-regeneration points, then p is cyclic with cycle length $L(q)$.

The lemma is an important ingredient in the proof of the following proposition.

Define $\kappa = k(m-1)+1$.

Proposition 1 *There exists an optimal solution p^* to (2) with the properties that (i) p^* is a cyclic pricing policy of the form $p^* = (q^*, q^*, q^*, q^*, \dots)$ with cycle length $L = L(q^*) \leq \kappa$ and (ii) the cycles and components of p^* coincide; i.e., $C_i(p^*) = q^*$ for $i = 0, 1, 2, \dots$.*

By Proposition 1, any optimal solution to problem (4) below is also an optimal solution to the original optimization problem (2).

$$\begin{aligned} &\max_p H(p) &&(4)\\ s.t.\quad &p = (q, q, q, \dots) \text{ and } C_i(p) = q, \text{ for } i = 0, 1, 2, \dots, \\ &q = (q_1, q_2, \dots, q_L) \in \mathcal{P}^L \text{ and } L \leq \kappa \end{aligned}$$

For $L \geq 2$, let

$$\begin{aligned} \Omega(L) &= \{q = (q_1, \dots, q_L) \in \mathcal{P}^L : q_{t-1} > \min\{q_t, \dots, q_{\min\{t+k-1, L\}}\} \quad \text{for } t = 2, \dots, L\} \\ &= \{q \in \mathcal{P}^L\text{: } q \text{ has no F-regeneration points}\} \end{aligned}$$

For $L = 1$, we define $\Omega(1) = \mathcal{P}$. Let $\Omega = \cup_{L=1}^{\kappa} \Omega(L)$.

We will also be interested in the following finite-horizon optimization problem.

$$\max_q \left\{ v(q) : q \in \Omega \right\} \tag{5}$$

Proposition 2 *A policy q solves* (5) *if and only if $p = (q, q, q, \dots)$ solves* (4).

Combining Propositions 1 and 2, we see that to prove Theorem 1, it will suffice to establish that the optimization problem (5) has a decreasing optimal solution of length $L \leq m + k - 1$. We take up this task below.

We close this section with a comment about the case of $k = 1$ in which patient customers are willing to wait one period. When $k = 1$ the definition (3) says that time $t \geq 2$ is an F-regeneration point of $q \in \mathcal{P}^L$ if $q_t \geq q_{t-1}$. Hence, the (finite) set Ω in (5) contains only strictly decreasing price sequences. In addition, $\kappa = m$ when $k = 1$. Therefore, the above propositions imply that a decreasing cyclic policy with cycle length $L \in \{1, \dots, m\}$ is optimal for (2) when $k = 1$. So at this point, Theorem 1 is proved for $k = 1$ without any more work. It remains only to prove the theorem when $k \geq 2$. Hence, we shall restrict our attention throughout the next section to cases with $k \geq 2$. Moving from $k = 1$ to $k \geq 2$ introduces a considerable additional level of difficulty. A key difference between settings with $k = 1$ and those with $k \geq 2$ is that for $k = 1$ any price increase produces a regeneration point, but for $k \geq 2$ it is possible to have a price increase that does not produce a regeneration point.

5 Proof of the Theorem

In this section we establish that there exists an optimal solution to (5) that is decreasing and has length at most $m + k - 1$, from which Theorem 1 follows. We begin by providing an overview.

In the first step of the argument, we introduce the notion of a "strong markup" and we use it to help decompose the feasible set Ω of (5) into two disjoint subsets, B and D, where B is comprised of the sequences in Ω with at least one strict increase in price, and D is comprised of the sequences in Ω that are decreasing. The second step of the proof involves showing that it is possible to improve the average revenue of any policy in the set B through a particular rearrangement of its prices. This rearrangement yields a policy that belongs to another set, which we call E. In the third step of the proof we establish that the best policy in E is not better than the best policy in D. As a consequence, we may restrict our attention to policies in D. The fourth step of the proof is to show that when maximizing over D it suffices to consider only those polices with length no greater than $m + k - 1$. We outline the four steps below, while leaving many of the technical details to the online appendix. We conclude the section by putting the pieces of the proof together.

5.1 Step 1 of the Proof

To begin, we decompose the feasible set Ω of (5) as indicated above. We start with a few definitions.

For a policy $q \in \mathcal{P}^L$ we say that there is a *markup* at time $t \in \{2, \ldots, L\}$ if $q_t > q_{t-1}$.

Given $q \in \Omega(L)$, we let $S_1 = \{t \in \{2, \ldots, L\} : q_t > q_{t-1}\}$. If $S_1 \neq \emptyset$, let $e(1) = \min\{t \in S_1\}$. For $i \geq 2$, let

$$S_i = \{t \in \{e(i-1)+1, \ldots, L\} : q_t > q_{t-1}, q_{t-1} = \min\{q_1, \ldots, q_{t-1}\}, \text{ and } q_{t-1} < q_{e(i-1)-1}\}.$$

If $S_i \neq \emptyset$, then let $e(i) = \min\{t \in S_i\}$. We say that time $e(i)$ is the time of the ith *strong markup* of q. At time $e(1)$, the price strictly increases from a running minimum in the previous period $e(1)-1$. For $i > 1$, at time $e(i)$, the price strictly increases from a running minimum price in the previous period $e(i)-1$ that is *strictly* less than the price in period $e(i-1)-1$. From the definition, the time of the first strong markup $e(1)$ is also the time of the first markup. Note however that subsequent markups may not be strong markups. A policy has at least one strong markup if and only if it has at least one markup. The left panel of Figure 1 shows a policy with two strong markups (at times 3 and 10). In that figure, times 4 and 6 have markups, but not strong markups.

In a number of places throughout the remainder of the paper and in the online appendix, we will make comparisons between a policy that we will call q and another policy, say $\widehat{q}$. In those settings we will use $e(i)$ to mean the time of the ith strong markup of q. Hence, $\widehat{q}_{e(i)}$ refers to the price charged by policy $\widehat{q}$ at the time of the ith strong markup of policy q. If we wish to refer to the time of ith strong markup of $\widehat{q}$, we will use the notation $e(i|\widehat{q})$. This convention will apply to other quantities that will be introduced later as well (e.g., $t(i)$ and $M(i)$).

Define $n(q) = \max\{i : S_i \neq \emptyset\}$ to be the number of strong markups of policy q. If $S_1 = \emptyset$, then we define $n(q) = 0$. For $q \in \Omega(L)$, observe that $n(q) = 0$ if and only if $q_1 \geq q_2 \geq \cdots \geq q_L$. Let

$$\begin{aligned}
B^n(L) &= \{q \in \Omega(L) : n(q) = n\} \qquad n = 1, \ldots, L-1 \\
&= \{q \in \Omega(L)\text{: } q \text{ has exactly } n \text{ strong markups}\} \\
B(L) &= \cup_{n=1}^{L-1} B^n(L) \\
&= \{q \in \Omega(L)\text{: } q \text{ has at least one markup}\} \\
D(L) &= \{q \in \Omega(L) : n(q) = 0\} \\
&= \{q \in \Omega(L)\text{: } q \text{ is decreasing}\}.
\end{aligned}$$

The left panel of Figure 1 shows a policy in a $B^n(L)$. We note in passing that it is possible to have a decreasing policy $q = (q_1, \ldots, q_L)$ with F-regeneration points (if a price repeats for at least k consecutive periods or if the final two prices are the same); such a policy is not in $D(L)$.

From the preceding definitions we have $\Omega(L) = B(L) \cup D(L)$. Each policy in the feasible set Ω of problem (5) is either decreasing or in some $B^n(L)$. Defining $B = \cup_{L=1}^{\kappa} B(L) = \cup_{L=1}^{\kappa} \cup_{n=1}^{L-1} B^n(L)$ and $D = \cup_{L=1}^{\kappa} D(L)$, we have

$$\Omega = B \cup D, \tag{6}$$

where B is composed of those sequences in Ω with at least one strict price increase, and D is composed of those sequences in Ω that are decreasing. This completes the "decomposition" of Ω.

5.2 Step 2 of the Proof

Next, we introduce the collection of sets $\{E^n(L)\}$ and show that for any policy in a set $B^n(L)$, we can find another policy (obtained by re-arranging prices between each two successive strong markups into a decreasing sequence) in $E^n(L)$ that performs at least as well.

For $n \geq 1$, define $E^n(L)$ as the set of sequences $q \in \mathcal{P}^L$ with the property that $q = (s_0, s_1, \ldots, s_n)$ where $s_i = (q_{t(i)}, \ldots, q_{t(i+1)-1})$ for some times $\{t(i) : i = 0, \ldots, n+1\}$ such that $1 = t(0) < t(1) < \cdots < t(n+1) = L+1$ and

1. s_i is decreasing ($q_{t(i)} \geq q_{t(i)+1} \geq \cdots \geq q_{t(i+1)-1}$) for $i = 0, \ldots, n$
2. $q_{t(i)-1} < q_{t(i)}$ for $i = 1, \ldots, n$
3. $\{t \in \{t(i)+1, \ldots, t(i+1)-1\} : q_t < q_{t(i)-1}\} \neq \emptyset$ for $i = 1, \ldots, n$
4. $\min\{t \in \{t(i)+1, \ldots, t(i+1)-1\} : q_t < q_{t(i)-1}\} \in \{t(i)+1, \ldots, t(i)+k-1\}$ for $i = 1, \ldots, n$.

Examples of policies in an $E^n(L)$ are shown in the right panel of Figure 1 and in the left panel of Figure 2.

Condition 1 says that a sequence in $E^n(L)$ is comprised of $n+1$ decreasing subsequences. Conditions 1 and 2 ensure that there are markups only at times $\{t(i) : i = 1, \ldots, n\}$. Conditions 1–3 say that there is at least one price in subsequence s_i that is strictly lower than $q_{t(i)-1}$ and that $q_{t(i)-1}$ is the minimum price over times $1, \ldots, t(i)-1$. Condition 4 says that the first price in s_i that is strictly lower than $q_{t(i)-1}$ comes not too much later than time $t(i)-1$, and hence $t(i)$ is not an F-regeneration point. Policies in $E^n(L)$ have the property that all markups are strong markups. Our interest in $E^n(L)$ stems from the next proposition.

Proposition 3 *For any policy $q \in B^n(L)$ with $L \leq \kappa$, there exists a policy $\bar{q} \in E^n(L)$ such that $V_L(q) \leq V_L(\bar{q})$ and $v_L(q) \leq v_L(\bar{q})$.*

The policy $\bar{q}$ above is constructed from q by rearranging the prices $(q_{e(i)}, q_{e(i)+1}, \ldots, q_{e(i+1)-1})$ between each two consecutive strong markups $e(i)$ and $e(i+1)$ of q into a decreasing sequence and also rearranging the prices $(q_{e(n)}, q_{e(n)+1}, \ldots, q_L)$ after the last strong markup into a decreasing sequence. So $\bar{q}$ will consist of $n+1$ decreasing sequences, and will be such that its markups and strong markups coincide.

To see how this works, fix $q \in B^n(L)$ and consider the string of prices $(q_{e(1)}, q_{e(1)+1}, \ldots, q_{e(2)-1})$ between $e(1)$ and $e(2)$. At least one of these prices is strictly less than $q_{e(1)-1}$ by the definition of $e(2)$. Let the time of the first such price in the string be called $e(1)+w+1$ (this naming convention is used to match the developments in the proof of the proposition). It must be that $w \leq k-2$, because otherwise $e(1)$ would be an F-regeneration point. Also by the definition of $e(2)$, we must have $q_{e(1)+w+1} \geq q_{e(1)+w+2} \geq \cdots \geq q_{e(2)-1}$. Hence, to rearrange $(q_{e(1)}, q_{e(1)+1}, \ldots, q_{e(2)-1})$ into a decreasing sequence, we need only rearrange the prices $(q_{e(1)}, \ldots, q_{e(1)+w})$ into a decreasing sequence and leave the other prices unchanged. See Figure 1. With this rearrangement, revenues obtained in periods $1, \ldots, e(1) - 1$ remain unchanged. Likewise, the revenues collected in periods $e(1) + w + 2, \ldots, L$ remain unchanged as well. (To see this, note that the price $q_{e(1)+w+1}$ in period $e(1)+w+1$ is unchanged by the rearrangement and that $q_{e(1)+w+1}$ is the minimum price over periods $1, \ldots, e(1)+w+1$ both with and without the rearrangement. Therefore, by the definition of revenue function $\rho_t(\cdot)$, the revenue collected in periods after $e(1) + w + 1$ does not change.) However, the total revenue accrued in periods $e(1), \ldots, e(1) + w + 1$ increases with the rearrangement. This increase can be attributed to customers who initially arrive in periods $e(1), \ldots, e(1) + w + 1$ (but not to customers who initially arrive earlier). Complete details are in the online appendix.

We close this section with the observation that upon defining $E = \cup_{L=1}^{\kappa} \cup_{n=1}^{L-1} E^n(L)$, it follows from Proposition 3 that maximizing over B yields an average revenue that is no greater than that obtained by maximizing over E.

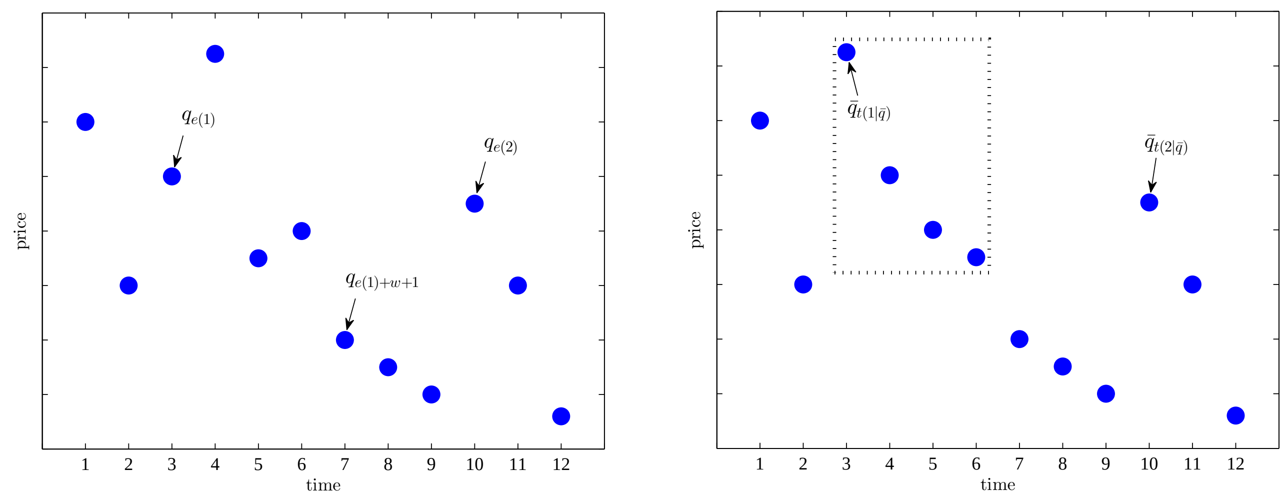

Figure 1: Suppose $k \geq 5$. The left panel shows a policy in $B^n(L)$ with $n = 2$ and $L = 12$. (If $k \leq 4$ then time $t = 3$ would be an F-regeneration point, and hence the policy would not be in $B^n(L)$.) The right panel shows the rearrangement of prices described above for periods $(e(1), \ldots, e(1) + w) = (3, \ldots, 6)$. Here, the rearrangement yields a policy $\bar{q}$ in $E^n(L)$. In general, we would need to make similar rearrangements between each pair $e(i)$ and $e(i + 1)$ and after $e(n)$.

5.3 Step 3 of the Proof

Here we establish that the best policy in E is no better than the best policy in D.

Proposition 4 *Consider* $d \in \arg\max_{p \in D} v(p)$ *and* $q \in \arg\max_{p \in E} v(p)$*. Then* $v(q) \leq v(d)$*.*

The proof of the preceding proposition is quite long. The main idea is as follows. We know that $q \in \arg\max_{p \in E} v(p)$ is an element of some $E^n(L)$. Consider the price $q_{t(n)}$ that begins the final piece s_n of q at time $t(n) = t$. If we modify q by adding "enough" copies (say L') of $q_{t(n)}$ starting at time $t+1$ and we push the prices $q_{t(n)+1}, \ldots, q_L$ later by L' periods, then we obtain a new policy of length $L + L'$ that has an F-regeneration point at time t and that also decreases from time t onward. See Figure 2. It turns out that the average revenue from the decreasing policy that uses only the prices from t onward is at least the average revenue of q. Were this not the case, then q could be improved by removing the final string s_n.

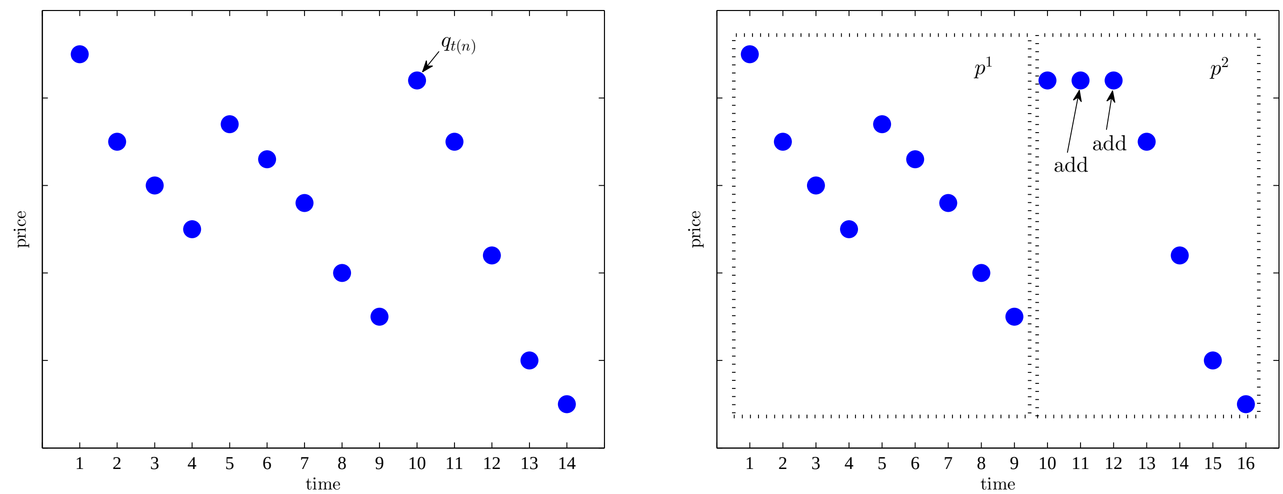

Figure 2: Suppose $k = 5$. The proof of Proposition 4 considers a policy q that maximizes $v(\cdot)$ over E. The left panel shows a case in which this policy q is a element of $E^n(L)$ with $n = 2$ and $L = 14$. Inserting $L' = 2$ additional copies of $q_{t(n)}$ as shown in the right panel yields policy of length $L + L' = 16$ that has an F-regeneration point at time $t = 10$. The average revenue of the policy in the right panel is a convex combination of the average revenues of policies p^1 and p^2 shown there.

5.4 Step 4 of the Proof

We now turn to the problem of maximizing $v(p)$ over $p \in D$. The next proposition states that for this problem, it suffices to consider only those sequences with length less than $k + m$. Of course, the longest possible *strictly* decreasing sequence of prices is m, but recall that "decreasing" should be interpreted as non-increasing, that is, the price can strictly decrease or remain the same as time moves forward (and so the result is not vacuous).

Proposition 5 *There exists* $d \in \arg\max\{v(p) : p \in D\}$ *such that* $L(d) \leq m + k - 1$.

The key to the proposition is the fact that no prices need repeat after period k in an optimal policy. Then, since there are m distinct prices in the price set $\mathcal{P}$, there is an optimal decreasing policy with length less than $m+k$. (In this discussion, "optimal" means optimal for the problem $\max\{v(p) : p \in D\}$.) To see why there is no benefit from repeating a price after period k, consider an $L \geq m+k$ and a policy q that maximizes $v(p)$ over $D(L)$. It must be that $q_j = q_{j+1}$ for some $j \geq k$. By removing q_{j+1} and shifting all later prices one period earlier, we arrive at a new policy q' which is in $D(L-1)$. Some careful thought reveals that the total revenue of q' is exactly $q_{j+1}[1 - G(q_{j+1})]$ less than that of q. This amount is no greater than $\max\{x[1-G(x)] : x \in \mathcal{P}\}$, which is itself no greater than the average revenue of the sequence q because the average revenue of the best policy in $D(L)$ is at least the optimal revenue of a one-period problem (or else the optimal solution to the one-period problem is itself better than q). Therefore, by removing such price repetitions from q, the average revenue gets no worse or else there is a policy in $D(1)$ that is itself better than q. Complete details are in the online appendix.

5.5 Putting the Pieces Together

We are now ready to present the proof of Theorem 1.

Proof of Theorem 1. By (6), the feasible set Ω in problem (5) can be expressed as $\Omega = B \cup D$. By Propositions 3, 4, and 5 we have

$$\max\{v(p) : p \in B\} \leq \max\{v(p) : p \in E\} \leq \max\{v(p) : p \in D\} = \max\{v(p) : p \in \cup_{L=1}^{m+k-1} D(L)\}\,.$$

Therefore, for $d \in \arg\max\{v(p) : p \in \cup_{L=1}^{m+k-1} D(L)\}$ we have $d \in \arg\max\{v(p) : p \in \Omega\}$. Consequently, $(d, d, d, \dots) \in \mathcal{P}^\infty$ is an optimal solution to (2) by Propositions 1 and 2. $\square$

6 Computation, Refinements, and Extensions

In this section we develop a dynamic programming method for computing an optimal policy. The structure identified in Theorem 1 plays a crucial role in the approach. We also present some additional results, including an extension of Theorem 1 to problems with continuous price sets.

6.1 A Computational Approach

As seen in the proof of Theorem 1, to find an optimal policy, it suffices to maximize $v(p)$ over $p \in \cup_{L=1}^{m+k-1} D(L)$. This can be done by maximizing $v(p)$ over $p \in D(L)$ for each individual

$L = 1, \ldots, m+k-1$ and then selecting the best of those $m+k-1$ maximizers. The search space can be reduced to just $L = k+1, \ldots, m+k-1$ when the assumption in Proposition 6 below holds.

Next we present a dynamic programming approach with a one-dimensional state variable. The dynamic program computes a decreasing sequence $\delta^L = (\delta_1^L, \ldots, \delta_L^L)$ of length L that maximizes $v(p)$ over all decreasing sequences p of length L. The monotonicity of prices established in our main theorem is what allows us to use a one-dimensional state variable. Without such monotonicity, we would need a k-dimensional state vector, rendering the problem computationally intractable for moderate or large k. To this end, fix L and let $\pi_t(x)$ denote the optimal revenue from time t onward in a L-period problem given that we have used a decreasing policy in periods $i = 1, \ldots, t-1$ and the price in period $t-1$ was $x \in \mathcal{P}$. The algorithm is as follows.

Step 1: Let $\pi_{L+1}(x) = 0$ for each $x \in \mathcal{P}$.

Step 2: For $t = L, \ldots, 2$ recursively compute

$$\pi_t(x) = \max_{z: z\in\mathcal{P}, z\leq x} \left\{ z\Big[1 - G(z) + \alpha \min\{k, t-1\}[F(x) - F(z)]\Big] + \pi_{t+1}(z) \right\} \qquad \text{for each } x \in \mathcal{P}$$

and let $z_t(x)$ be a maximizer of the right side above.

Step 3: Compute

$$\pi_1 = \max_{z: z\in\mathcal{P}} \Big\{ z[1 - G(z)] + \pi_2(z) \Big\}$$

and let z_1 denote an optimal solution, i.e., $z_1 \in \arg\max \Big\{ z[1 - G(z)] + \pi_2(z) \Big\}$.

Step 4: The policy $\delta^L = (\delta_1^L, \delta_2^L, \ldots, \delta_L^L)$ given by $\delta_1^L = z_1$ and $\delta_t^L = z_t(\delta_{t-1}^L)$ for $t = 2, \ldots, L$ maximizes $V(p)$ — and hence also $v(p)$ — over the set of decreasing policies p of length L. In addition, $V(\delta^L) = \pi_1$ and $v(\delta^L) = \pi_1 / L$.

We note in closing that the decreasing policy δ^L produced by this dynamic programming algorithm may have F-regeneration points if the final two prices in δ^L are identical or if a price is repeated k or more times consecutively in δ^L. If either of these two cases occur, then there is an $L' < L$ and a policy $q \in D(L')$ with $v(q) \geq v(\delta^L)$. Hence, it poses no problem if $\delta^L \notin D(L)$.

6.2 A Lower Bound on Optimal Cycle Lengths

The next result gives conditions under which we can further restrict the space in which we search for an optimal policy. Let $q^j \in \arg\max\{V_j(q) : q \in D(j)\} = \arg\max\{v_j(q) : q \in D(j)\}$ for $j = 1, \ldots, k+1$. For clarification, we note that the notation q^j is used differently in the proof of Proposition 4.

Proposition 6 *If $v_2(q^2) \geq v_1(q^1)$, then $v_1(q^1) \leq v_2(q^2) \leq \cdots \leq v_{k+1}(q^{k+1})$ and therefore, there exists an optimal solution $q \in D(L)$ to (5) with $L \geq k+1$. If $v_2(q^2) > v_1(q^1)$, then $v_1(q^1) < v_2(q^2) < \cdots < v_{k+1}(q^{k+1})$.*

The condition $v_2(q^2) \geq v_1(q^1)$ means that the best policy $q^2 = (q_1^2, q_2^2)$ in $D(2)$ is at least as good as the best policy $q^1 = (q_1^1)$ in $D(1)$. If we consider a problem with two time periods, then the condition implies that under policy q^2, more revenue is accrued from customers who initially arrive in period 1 than from customers who initially arrive in period 2. [Were this not the case, then the policy in $D(1)$ that sets price q_2^2 in period 1 would be better than — that is, yield higher average revenue than — q^2, which would contradict $v_2(q^2) \geq v_1(q^1)$.] In the proof of the proposition, we establish by induction that $v_{j+1}(q^{j+1}) \geq v_j(q^j)$ for $j = 1, \ldots, k$. The condition $v_2(q^2) \geq v_1(q^1)$ provides the base case in that argument. A key piece of the inductive step is to show that for a j-period problem, under policy q^j, the revenue accrued from customers who originally arrive in period 1 exceeds the average revenue per period over periods $2, \ldots, j$ accrued from just those customers who initially arrive in periods $2, \ldots, j$. It is simple to compute $v_2(q^2)$ and $v_1(q^1)$, and therefore it is easy to check whether it holds that $v_2(q^2) \geq v_1(q^1)$.

6.3 Continuous Price Set

Our next result establishes that the decreasing cyclic structure in Theorem 1 holds even when the price set is infinite and prices are selected from $\mathcal{P} = [0, \bar{P}]$ where $\bar{P} < \infty$. In this case (2) becomes $\sup_p \{H(p) : p \in [0, \bar{P}]^\infty\}$.

Proposition 7 *Suppose $\mathcal{P} = [0, \bar{P}]$ and $F(\cdot), G(\cdot)$ are Lipschitz continuous on $\mathcal{P}$ with continuous and strictly positive derivatives on $(0, \bar{P})$. Then there is a decreasing cyclic policy that is an optimal solution to (2).*

The uniform distribution function is Lipschitz continuous, and therefore Proposition 7 applies when customers' valuations are uniformly distributed.

The proof of the proposition involves passing to a limit through a sequence of problems, each with a finite price set that is a discretization of $[0, \bar{P}]$. The idea is to let the discretization of the price set become progressively finer, yielding progressively larger price sets that more closely approximate $[0, \bar{P}]$. For each problem in the sequence, we can apply Theorem 1 to see that there is an optimal decreasing cyclic policy for the discretized problem in question. One difficulty that arises is that the cardinality of the price set increases as the discretization becomes finer, and hence it is not immediately apparent what will keep the cycle length finite when we pass to the limit.

This issue is resolved in Section G of the online appendix, where Lemma G.2 allows us to establish a uniform upper bound on the cycle lengths of the optimal policies of the discretized problems.

7 Variable Patience Levels: Bounds and Heuristics

In this section we consider a problem involving customers with different patience levels. The model is the same as that described in Section 3, except we now assume that customers' patience levels range from 0 to $K < \infty$, and each level k accounts for a fraction of α_k of customers where $\sum_{k=0}^{K} \alpha_k = 1$. A customer with a patience level k will wait up to k periods to make a purchase. So, a customer with patience level $k = 0$ is impatient in the sense described earlier. For $p \in \mathcal{P}^\infty$ let

$$\widetilde{\rho}_t(p) = p_t\Big[1 - G(p_t) + \sum_{k=1}^{K}\sum_{i=k}^{K} \alpha_i [F(\min\{p_{t-k}, \ldots, p_{t-1}\}) - F(p_t)]^+\Big] \quad \text{for } t = 1, 2, \ldots \tag{7}$$

where here $G(p_t) = \alpha_0 F_0(p_t) + (1 - \alpha_0)F(p_t)$. For $p \in \mathcal{P}^\infty$ let $\widetilde{H}(p) = \liminf_{T\to\infty} T^{-1} \sum_{t=1}^{T} \widetilde{\rho}_t(p)$. The seller now wants to solve

$$\widetilde{H} = \sup_p \{\widetilde{H}(p) : p \in \mathcal{P}^\infty\}. \tag{8}$$

Proposition 8 *There exists a cyclic pricing policy with cycle length at most $K(m-1)+1$ that is an optimal solution to* (8).

The proof of this result is identical to that of Proposition 1, and is omitted. It is seemingly a difficult task to identify additional structure of an optimal policy. It is also a difficult task to compute an optimal policy. Note that the computational approach from Section 6.1 relies on the fact that optimal cycles are decreasing, which need not be the case here. That approach can readily be modified to cover the variable patience levels now under consideration, but doing so yields a dynamic program with K-dimensional state variable. The main idea is to recursively compute $\widetilde{\pi}_t(x_{t-K}, \ldots, x_{t-1})$, the maximum revenue from time t onward in an L-period problem given that the prices in periods $t - K, \ldots, t - 1$ are $x_{t-K}, \ldots, x_{t-1}$. This can be done via the relation

$$\begin{aligned}\widetilde{\pi}_t(x_{t-K}, \ldots, x_{t-1}) = \max_{z: z\in\mathcal{P}} \Big\{ z\Big[1 - G(z) + \sum_{k=1}^{K}\sum_{i=k}^{K} \alpha_i [F(\min\{x_{t-k}, \ldots, x_{t-1}\}) - F(z)]^+\Big] \\ + \widetilde{\pi}_{t+1}(x_{t-K+1}, \ldots, x_{t-1}, z)\Big\}\end{aligned} \tag{9}$$

with the convention that $x_j = 0$ for $j \leq 0$. Such dynamic programs are intractable except when K is small. Hence, we will confine ourselves to developing bounds and heuristics.

To do so, let $\rho_t^k(p)$ denote the revenue function defined in (1) with $\alpha = 1 - \alpha_0$ and fixed patience level k. Likewise, let $H^k(p) = \liminf_{T\to\infty} T^{-1}\sum_{t=1}^T \rho_t^k(p)$ be the value of a policy p and let $H^k = \sup\{H^k(p) : p \in \mathcal{P}^\infty\}$ denote the optimal value in (2). A little algebra shows that

$$\widetilde{\rho}_t(p) = \sum_{k=1}^K \frac{\alpha_k}{1-\alpha_0}\rho_t^k(p).$$

Therefore, for the optimal cyclic policy $\widetilde{p}$ whose existence was established in Proposition 8, we have

$$\begin{aligned}\widetilde{H} = \widetilde{H}(\widetilde{p}) &= \liminf_{T\to\infty}\frac{1}{T}\sum_{t=1}^T \widetilde{\rho}_t(\widetilde{p}) = \lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^T \widetilde{\rho}_t(\widetilde{p}) = \lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^T\sum_{k=1}^K \frac{\alpha_k}{1-\alpha_0}\rho_t^k(\widetilde{p})\\ &= \sum_{k=1}^K \frac{\alpha_k}{1-\alpha_0}\liminf_{T\to\infty}\frac{1}{T}\sum_{t=1}^T \rho_t^k(\widetilde{p}) \le \sum_{k=1}^K \frac{\alpha_k}{1-\alpha_0}\sup_p \liminf_{T\to\infty}\frac{1}{T}\sum_{t=1}^T \rho_t^k(p).\end{aligned}$$

Hence, we have established the following result, which states that the optimal value in (8) for the problem with variable patience levels is bounded above by a convex combination of optimal objective values of problems with fixed patience levels.

Proposition 9 *The optimal value in* (8) *satisfies*

$$\widetilde{H} \le \bar{H} := \sum_{k=1}^K \frac{\alpha_k}{1-\alpha_0} H^k.$$

We now turn our attention to finding a computationally tractable heuristic approach for problems with variable patience levels. To do so, observe that it is straightforward to maximize $\widetilde{H}(\cdot)$ over decreasing cyclic policies using a slight variation of the algorithm from Section 6.1. Specifically, in (9) we can simply maximize over $\{z \in \mathcal{P} : z \le x_{t-1}\}$. We can also drop all arguments except x_{t-1} from the function $\widetilde{\pi}(\cdot)$ because $\min\{x_{t-k}, \dots, x_{t-1}\} = x_{t-1}$ for a decreasing policy. Such an optimization is tractable because the state variable is one-dimensional. The policy (say $\widetilde{q} = (\widetilde{d}, \widetilde{d}, \dots)$ where $\widetilde{d}$ is decreasing) produced by this procedure will generally not be an optimal policy — even though it is the best decreasing cyclic policy — because in this setting of variable patience levels, it may be the case that the optimal policy does not have decreasing cycles (i.e., $\widetilde{H}(\widetilde{q}) < \widetilde{H}$). Nevertheless, using the best decreasing cyclic policy is a workable heuristic. Proposition 9 above can help us evaluate this approach. For instance, the proposition yields computable bounds on the fraction of the optimal revenue that can be obtained by implementing the best decreasing cyclic policy. Specifically, we have $\widetilde{H}(\widetilde{q})/\bar{H} \le \widetilde{H}(\widetilde{q})/\widetilde{H} \le 1$. In the next section we consider the best decreasing policy heuristic for some example problems with variable patience levels.

8 Numerical Experiments

In this section we present numerical results for problems with both fixed and variable patience levels. In the first set of experiments, we consider settings with fixed patience levels and we suppose that valuations for both patient and impatient customers are drawn from the beta distribution with parameters (a, b). The density of the beta distribution is given by $f(x) = F'(x) = x^{a-1}(1 - x)^{b-1}/B(a, b)$ for $x \in (0, 1)$ where $B(a, b) = \int_0^1 t^{a-1}(1 - t)^{b-1}dt$. Taking $a = b = 1$ yields the uniform distribution on $(0, 1)$. The mean and variance of the beta distribution with parameters (a, b) are given by $\mu = a/(a + b)$ and $\sigma^2 = (ab)/[(a + b)^2(a + b + 1)]$.

In Table 1, we consider examples with $a = b$ so that the mean valuation is fixed at $\mu = 1/2$, the variance is $\sigma^2 = 1/(8a + 4)$, the coefficient of variation is $\text{cv} = \sigma/\mu = 1/\sqrt{2a + 1}$, and the valuation density is symmetric about 1/2. We set $m = 10$ with prices evenly spaced on $(0, 1]$; that is, the price set is $\mathcal{P} = \{1/10, 2/10, \ldots, 9/10, 1\}$. The pairs of numbers in the table denote the optimal cycle length and the optimal average revenue for particular values of α, k, and a. For example, the optimal cycle length is 4 and optimal average revenue is 0.2736 for $\alpha = 0.5$, $k = 2$, and $a = b = 2$. The variance of the valuations decreases as a (and therefore b) increases. Hence, as we move to the right in the table, we get examples with lower variances and lower coefficients of variation. When $a = b > 1$ the mode (the maximum value of the density) occurs at 1/2, and as a gets bigger the distribution becomes more concentrated around 1/2. For $a = b < 1$ the density has a "U shape" and is minimized at 1/2 and grows without bound as $x \downarrow 0$ or $x \uparrow 1$. As a becomes smaller, the distribution becomes more concentrated just above 0 and just below 1.

As expected, the revenue increases in k and in α (it is simple to prove this). For fixed k and α the revenue is lowest when $a = 1$, which is the case of the uniform distribution. Revenue increases as a function of the variance of the valuation distribution for $a < 1$ and decreases in the variance when $a > 1$. This is perhaps counterintuitive at first because one might expect revenue to decrease as a function of the variance, given a fixed mean. This phenomenon can be explained as follows. Although the variance increases as a goes down to 0, the distribution becomes more concentrated in a sense as noted in the previous paragraph, and therefore one may view the "variability" to be low even though the variance is itself high (of course, variance is only a summary measure of variability). In the limit as $a \downarrow 0$, one may view the distribution as essentially placing mass 1/2 at both 0 and 1, and indeed it turns out that the revenue converges to 1/2 as $a \downarrow 0$. (Note that if the valuation distribution places mass 1/2 at both 0 and 1, then it is optimal to always price at 1, yielding a revenue of 1/2.) As $a \uparrow \infty$, the valuation distribution becomes essentially a unit mass at 1/2. In that case the revenue also approaches 1/2. (Note that if the valuation distribution places mass 1 at 1/2, then it is optimal to always price at 1/2, yielding a revenue of 1/2.)

We also considered a variety of examples in which valuations follow the gamma distribution. In those examples (not shown) we found that for a given mean valuation, the optimal average revenue decreases in the variance of the valuation. This is consistent with intuition.

The optimal cycle length is either 1 (fixed price), $k+1$, or $k+2$ for all the examples shown on Table 1. We have seen the same in other examples, and in fact, we have not been able to find any examples in which the optimal cycle length exceeds $k+2$. We conjecture that there are broad conditions under which there is an optimal policy with cycle length of either $k+1$ or $k+2$. Recall that for uniform valuations and continuous price set, Ahn et al. (2007) have previously established that there is an optimal policy with cycles of length 2 for problems with $k=1$. As further support for the conjecture, we have also proved through a different argument (not presented here) that for uniform valuations and continuous price set, there is an optimal policy with cycles of length 3 for problems with $k=2$. However, neither of these arguments appears to readily extend to more general settings. Therefore, it is an open problem to prove the conjecture.

The next set of examples is shown in Table 2, where we consider variable patience levels as described in Section 7. We take $K=2$ so that patient customers are willing to wait either 1 or 2 periods (and impatient customers will not wait). In the examples we assume that valuations are uniformly distributed on $(0,1)$ and we assume that the fraction of customers that is impatient is $\alpha_0 = 0.2$. We take $m=50$ prices evenly spaced on $(0,1]$. For $K=2$, we can use the recursion (9) to compute an optimal policy and the associated optimal revenue. Hence, these examples allow us to make some assessment of the quality of the upper bound $\bar{H}$ from Proposition 9 and also to evaluate the performance $\widetilde{H}(\widetilde{q})$ of the best decreasing cyclic policy $\widetilde{q}$ for problems with variable patience levels. The columns in the table correspond to different combinations of (α_1, α_2), which are the fractions of customers with patience levels 1 and 2. As we move rightward in the table, customers become less patient. In five of the seven examples, a decreasing cyclic policy turns out to be optimal. In the other two, the best decreasing policy achieves 99.9% of the optimal revenue. The upper bound on the optimal revenue is also quite tight. Overall, Table 2 suggests that the best decreasing policy works well for $K=2$ when customer valuations are uniform. Is this still the case with larger K and different valuation distribution? We take this up next.

In Table 3, we consider variable patience levels but with larger values of K, namely $K=4$ and $K=10$. For these values, it is not practical to compute an optimal policy with (9), so we instead obtain the best decreasing cyclic policy and compare it against the upper bound $\bar{H}$ from Proposition 9. In these examples we suppose that valuations are drawn from the gamma distribution with parameters $n=\lambda=1/2$ so that the mean valuation is $\mu=1$ and the variance is $\sigma^2=2$.

Table 1: Optimal Cycle Lengths and Average Revenues for Beta-Distributed Valuations.

	$a = b \longrightarrow$	1/8	1/2	1	2	8
α	$k \downarrow$	cv = 0.89	cv = 0.71	cv = 0.58	cv = 0.45	cv = 0.24
0.2	1	(1, 0.3484)	(2, 0.2639)	(1, 0.2500)	(2, 0.2605)	(1, 0.3148)
	2	(3, 0.3504)	(3, 0.2647)	(3, 0.2520)	(3, 0.2610)	(1, 0.3148)
	5	(6, 0.3526)	(6, 0.2682)	(6, 0.2550)	(6, 0.2638)	(6, 0.3180)
	10	(11, 0.3543)	(11, 0.2715)	(12, 0.2575)	(11, 0.2666)	(11, 0.3209)
0.5	1	(2, 0.3529)	(2, 0.2705)	(2, 0.2600)	(2, 0.2694)	(1, 0.3148)
	2	(3, 0.3567)	(3, 0.2786)	(3, 0.2667)	(4, 0.2736)	(3, 0.3211)
	5	(6, 0.3632)	(6, 0.2900)	(6, 0.2783)	(6, 0.2858)	(6, 0.3302)
	10	(11, 0.3692)	(11, 0.2995)	(11, 0.2864)	(11, 0.2923)	(11, 0.3343)
0.8	1	(2, 0.3577)	(2, 0.2840)	(2, 0.2730)	(2, 0.2826)	(2, 0.3283)
	2	(3, 0.3637)	(3, 0.2974)	(3, 0.2887)	(3, 0.2973)	(4, 0.3404)
	5	(6, 0.3750)	(6, 0.3209)	(6, 0.3143)	(6, 0.3206)	(6, 0.3604)
	10	(11, 0.3863)	(11, 0.3405)	(12, 0.3330)	(12, 0.3380)	(11, 0.3698)
1	1	(2, 0.3609)	(2, 0.2943)	(2, 0.2850)	(2, 0.2944)	(2, 0.3398)
	2	(3, 0.3700)	(3, 0.3159)	(3, 0.3067)	(3, 0.3163)	(4, 0.3621)
	5	(6, 0.3839)	(6, 0.3479)	(6, 0.3467)	(6, 0.3567)	(6, 0.3963)
	10	(11, 0.3981)	(11, 0.3751)	(12, 0.3750)	(11, 0.3842)	(11, 0.4164)

The table shows various different values of α_0. For each, we considered three different problem instances. In instance *i*, we take $\alpha_i = (1 - \alpha_0)/K$ for $i = 1, \ldots, K$, which means the fraction of customers with each patience level is the same. In instance *ii*, we take $\alpha_1 = \alpha_K = (1 - \alpha_0)/2$, which means half of the patient customers have patience level 1 and the other half have patience level K. In instance *iii*, we take $\alpha_{K/2} = \alpha_{K/2+1} = (1 - \alpha_0)/2$, which means half of the patient customers have patience level $K/2$ and the other half have patience level $K/2 + 1$. Instances *i* and *ii* represent situations where the problem in question is dissimilar to a problem with fixed patience levels, whereas instance *iii* represents a case where the problem is "almost" a problem with a fixed patience level. Not surprisingly, the best decreasing policy does quite well in instance *iii*, attaining at least 99% of the optimal revenue in each case. However, in instances *i* and *ii* the best decreasing policy may not perform as well. For example, when $\alpha_0 = 0$ and $K = 10$, the best decreasing policy attains roughly 89% of the upper bound. Note that the comparison is made against the upper

Table 2: Variable Patience Levels with $K = 2$

$(\alpha_1, \alpha_2) \longrightarrow$	(0.1,0.7)	(0.2,0.6)	(0.3,0.5)	(0.4,0.4)	(0.5,0.3)	(0.6,0.2)	(0.7,0.1)
optimal revenue $\widetilde{H}$	0.2869	0.2841	0.2816	0.2791	0.2769	0.2749	0.2745
upper bound $\bar{H}$	0.2878	0.2859	0.2840	0.2821	0.2801	0.2782	0.2763
best decreasing $\widetilde{H}(\widetilde{q})$	0.2869	0.2841	0.2816	0.2791	0.2769	0.2747	0.2744
ratio $\widetilde{H}(\widetilde{q})/\widetilde{H}$	1.00	1.00	1.00	1.00	1.00	0.9993	0.9996
ratio $\widetilde{H}(\widetilde{q})/\bar{H}$	0.9969	0.9937	0.9915	0.9894	0.9886	0.9874	0.9931

bound, so it is not possible to tell if the gap arises from a shortcoming of decreasing cyclic policies or because the upper bound is loose. The table suggests that the ratio $\widetilde{H}(\widetilde{q})/\bar{H}$ decreases as K gets bigger, which is not surprising. Even for $K = 10$, the best decreasing policy performs reasonably well in all instances.

Table 3: Variable Patience Levels with $K = 4$ and $K = 10$

		$K = 4$			$K = 10$		
α_0		i	ii	iii	i	ii	iii
0	upper bound $\bar{H}$	0.4718	0.4651	0.4785	0.5455	0.5155	0.5626
	best decreasing $\widetilde{H}(\widetilde{q})$	0.4559	0.4429	0.4753	0.5094	0.4591	0.5601
	ratio $\widetilde{H}(\widetilde{q})/\bar{H}$	0.9663	0.9523	0.9933	0.9338	0.8906	0.9956
0.2	upper bound $\bar{H}$	0.43	0.4257	0.4344	0.4836	0.4635	0.496
	best decreasing $\widetilde{H}(\widetilde{q})$	0.4184	0.4091	0.4327	0.4570	0.4210	0.494
	ratio $\widetilde{H}(\widetilde{q})/\bar{H}$	0.9730	0.9610	0.9961	0.9450	0.9083	0.9960
0.5	upper bound $\bar{H}$	0.3785	0.3766	0.3803	0.4035	0.3942	0.4087
	best decreasing $\widetilde{H}(\widetilde{q})$	0.3717	0.3674	0.3790	0.3894	0.3741	0.4077
	ratio $\widetilde{H}(\widetilde{q})/\bar{H}$	0.9820	0.9756	0.9966	0.9651	0.9490	0.9976
0.8	upper bound $\bar{H}$	0.3423	0.3418	0.3427	0.3487	0.3466	0.3499
	best decreasing $\widetilde{H}(\widetilde{q})$	0.3402	0.3389	0.3420	0.3447	0.3416	0.3497
	ratio $\widetilde{H}(\widetilde{q})/\bar{H}$	0.9939	0.9915	0.9980	0.9885	0.9856	0.9994

9 Concluding Comments

In this paper, we consider dynamic pricing in the presence of patient customers. Our main result establishes that for arbitrary fixed patience levels and arbitrary valuation distributions, there is an optimal policy comprised of cycles of decreasing prices. We also provide bounds on the length of these cycles and present an algorithm for computing an optimal policy. Our work complements previous results that apply to problems with uniform valuation distributions in which patient customers are willing to wait one period to make a purchase. A direction for future research is to prove the conjecture that there are optimal policies with cycles of length $k + 1$ or $k + 2$, where k is the number of periods a patient customer is willing to wait to make a purchase. It would also be of interest to consider pricing problems with patient customers in which customer arrivals are stochastic. The models and policies discussed in this paper could potentially form the basis for heuristics in stochastic settings.

References

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

Y. Aviv, Y. Levin, and M. Nediak. Counteracting strategic consumer behavior in dynamic pricing systems. In S. Netessine and C. S. Tang, editors, *Consumer-Driven Demand and Operations Management Models*, pages 323–352. Springer, 2009.

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

O. Besbes and I. Lobel. Intertemporal price discrimination: Structure and computation of optimal policies. *Management Science*, 61(1):92–110, 2015.

A. Caplin, M. Dean, and D. Martin. Search and satisficing. *The American Economic Review*, 101 (7):2899–2922, 2011.

F. Caro and V. Martínez-de-Albéniz. Product and price competition with satiation effects. *Management Science*, 58(7):1357–1373, 2012.

J. Conlisk, E. Gerstner, and J. Sobel. Cyclic pricing by a durable goods monopolist. *The Quarterly Journal of Economics*, 99(3):489–505, 1984.

M. Gümüş, P. M. Kaminsky, and S. Mathur. The impact of product substitution and retail capacity on the timing and depth of price promotions: Theory and evidence. Working paper, McGill University, 2013.

T. Huang, G. Allon, and A. Bassamboo. Bounded rationality in service systems. *Manufacturing & Service Operations Management*, 15(2):263–279, 2013.

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

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

J. Nasiry and I. Popescu. Dynamic pricing with loss-averse consumers and peak-end anchoring. *Operations Research*, 59(6):1361–1368, 2011.

I. Popescu and Y. Wu. Dynamic pricing strategies with reference effects. *Operations Research*, 55 (3):413–429, 2007.

E. L. Porteus. *Foundations of Stochastic Inventory Theory*. Stanford University Press, 2002.

E. Reutskaja, R. Nagel, C. F. Camerer, and A. Rangel. Search dynamics in consumer choice under time pressure: An eye-tracking study. *The American Economic Review*, 101(2):900–926, 2011.

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

H. A. Simon. A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1): 99–118, 1955.

P. Stüttgen, P. Boatwright, and R. T. Monroe. A satisficing choice model. *Marketing Science*, 31 (6):878–899, 2012.

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

Author Biographies

Yan Liu is a Ph.D. candidate in the Department of Industrial and Systems Engineering at University of Minnesota. His research interests include revenue management and pricing.

William L. Cooper is a Professor in the Department of Industrial and Systems Engineering at University of Minnesota. His research interests include revenue management and applied probability.