Dynamic Pricing with Heterogeneous Patience Levels

Ilan Lobel

Abstract

We consider the problem of dynamic pricing in the presence of patient consumers. We call a consumer patient if she is willing to wait a certain number of periods for a lower price and will purchase as soon as the price is equal to or below her valuation. We allow for arbitrary joint distributions of patience levels and valuations. We propose an efficient dynamic programming algorithm for finding optimal pricing policies. We find numerically that optimal policies can take the form of incomplete cyclic policies, mixing features of nested sales policies and of decreasing cyclic policies.

1 Introduction

When a consumer encounters a price she considers too high, she sometimes waits for a given amount of time to see if the price drops below a given target level. This is referred to in the literature as *patient consumer behavior*. To be more precise, a patient consumer is characterized by her arrival time, willingness-to-wait and valuation. If the price of the product on offer drops below her valuation within her patience window, she will purchase the product. This paper is a study of optimal pricing policies in the presence of patient consumers who are heterogeneous in their arrival times, patience levels and valuations.

There is a growing body of evidence that retail firms engage in dynamic pricing as a tool for intertemporal price discrimination (Hendel and Nevo [2013], Li et al. [2014] and Moon et al. [2018]). This is true of brick-and-mortar and is especially true of e-commerce retailers, where firms often use sophisticated dynamic pricing algorithms to sort patient consumers from impatient ones. Our work aims to understand how a firm should dynamically change its prices in order to maximize revenues given that consumers are quite diverse in their willingness-to-wait for lower prices and that willingness-to-wait is often correlated with willingness-to-pay.

Our paper builds on earlier work on dynamic pricing with patient consumers by Ahn et al. [2007] and Liu and Cooper [2015]. Our model is essentially the same as the one considered in Liu and Cooper [2015], with the only difference being that they study an infinite horizon model while we focus on a finite horizon formulation. At the heart of Liu and Cooper [2015] lies a fascinating claim. The authors of that paper argue that pricing for patient consumers is a more challenging task than pricing for strategic consumers, a claim that is at odds with the received wisdom of the revenue management research community. To be more precise, Liu and Cooper [2015] study a model that is identical to Besbes and Lobel [2015], except for the assumed consumer behavior (Besbes and Lobel [2015] assume consumers are strategic). While Besbes and Lobel [2015] are able to construct a polynomial-time algorithm for finding optimal prices for strategic consumers, Liu and Cooper [2015] are able to do the same for patient consumers only for the special case where all patient consumers have the same patience level (there are also impatient consumers in this version of their model). For the case of heterogeneous patience levels, they offer only an algorithm that runs in exponential time on the maximum willingness-to-wait.

Our contributions. Our main result is the construction of an algorithm for computing optimal dynamic pricing policies in the presence of patient consumers with heterogeneous patience levels. The algorithm we propose works for any "patient demand model." That is, it works for any joint distribution of valuations and patience levels. The algorithm runs in polynomial time in both the time horizon and the number of available prices. Our algorithm can also be used in an infinite horizon model by combining our results with a bound on the maximum length of any policy from Liu and Cooper [2015]. In the infinite horizon model, our algorithm runs in polynomial time on the maximum willingness-to-wait and the number of prices. The algorithm is loosely inspired by the dynamic program proposed in Besbes and

Lobel [2015] for strategic consumers, but it relies on a different recursion and state space.

At first glance, this result would seem to invalidate Liu and Cooper [2015]'s claim regarding prices for patient consumers being harder to compute than prices for strategic consumers. However, this would be an incorrect conclusion since the algorithm we propose requires a bigger state space than the one for strategic consumers from Besbes and Lobel [2015]. Our algorithm for patient consumers runs in quadratic time in the number of available prices, while the algorithm for strategic consumers requires only linear time in the number of prices.

At the heart of the increased difficulty of the patient consumers problem lies a simple idea. Let us call a group of consumers who arrive at a given period with a given patience level a cohort. In a strategic consumers model, all consumers within a given cohort either purchase at a given price or do not purchase at all. That is, we can understand revenues by assigning each cohort to a single price available to its consumers. With patient consumers, this is no longer the case. Different consumers within a cohort purchase at different prices. We therefore need to assign different prices to different intervals of valuation within a cohort, which necessitates a more complex state space.

We also take advantage of our algorithmic result to perform a numerical exploration of the structure of optimal dynamic pricing policies under patient consumers. Besbes and Lobel [2015] show that, with strategic consumers, pricing policies often take a form they call nested sales. Liu and Cooper [2015] show that with one patient consumer class, the optimal policies are cyclic decreasing. With heterogeneous patience levels, we find that optimal policies can take a form that combines properties of both nested sales and cyclic decreasing. Specifically, we call this kind of policy *incomplete cyclic*. An incomplete cyclic policy looks like a cyclic decreasing policy, but its cycles are sometimes incomplete, with the price increasing sharply back to its high point before reaching its low point. We also argue that the seller is able to better price discriminate in a patient consumers model than in a strategic consumers model.

Other related works. The first paper to consider a model where consumers stay in the system until they encounter a price below their valuation appears to be Kalish [1983], a study of pricing under learning-by-doing and word-of-mouth effects. Another early paper to consider a model with patient consumers is Besanko and Winston [1990], though the main focus of that paper is a model with strategic consumers. Patience is assumed to be infinite in both of these papers, in the sense that consumers only depart the system if they purchase or the selling season ends. Some recent papers, such as Caldentey et al. [2017], also compare and contrast strategic consumer models with infinite patience models.

In a very recent contribution that advances the field in a different direction, Araman and Fayad [2017] consider a model with patient consumers who have stochastic valuations. Consumer valuations change over time according to a Markov chain, and one of the states of the Markov chain is assumed to represent departure from the system. The consumers are patient, since they purchase immediately if their valuation is above the current price. The authors show that cyclic policies are near-optimal for this stochastic problem. Cohen et al. [2017] study a model of consumer demand that depends on past prices. Their model allows for demand to accumulate over time and, thus, has the flavor of a patient consumers model.

There is a substantial related literature on dynamic pricing with strategic consumers. Conlisk et al. [1984] were the first to consider the problem of how sellers should price their products in settings where consumers arrive over time and are intertemporal utility maximizers. They assumed consumers had two possible valuations, low and high, and showed that this would lead to optimal policies being cyclic. The two papers that introduced the problem of dynamic pricing with strategic consumers to the operations community are Su [2007] and Aviv and Pazgal [2008]. Some of the recent contributions to this topic include Chen and Shi [2016], Chen and Wang [2016], Correa et al. [2016], Briceno et al. [2017] and Chen and Farias [2018]. As yet, there is no empirical work that attempts to disentangle patient behavior from strategic behavior within the context of dynamic pricing. There is, however, experimental work that supports the idea that humans often make decisions via satisficing rather than utility-maximizing (Caplin et al. [2011], Reutskaja et al. [2011] and Stüttgen et al. [2012]). There also exists a quickly growing family of papers studying pricing in the presence of more complex behavioral effects. In particular, there are several papers that study dynamic pricing with reference price effects (Kopalle et al. [1996], Popescu and

Wu [2007], Nasiry and Popescu [2011], Wang [2016] and Hu et al. [2016]). Our model is a behavioral one with a demand accumulation effect but no reference price effect.

2 Model

We consider a monopolist facing a multi-period single-product pricing problem. Consumers arrive with unit demand and are characterized by their arrival time t, their valuation for the product $v \in \mathbb{R}^+$ and their willingness-to-wait $w \in \{0, 1, ..., S\}$, where $S \in \mathbb{N}$. Consumers are infinitesimal and arrive deterministically. The mass of consumers arriving in each period with patience level w is denoted by γ_w. For each value of w, the cumulative distribution of valuations is denoted by $\mathcal{F}_w(\cdot)$. For each w and v, we let $F_w(v) = \lim_{v' \uparrow v} \mathcal{F}(v')$ be the left limit of $\mathcal{F}_w(\cdot)$ at v. We assume a finite horizon equal to $T \in \mathbb{N}$ and that the demand is stationary, so that in every period $t \in \{1, 2, ..., T\}$, the mass of consumers arriving with patience w and valuation greater than or equal to v is given by $\gamma_w(1 - F_w(v))$. We do not impose any assumptions on the demand model $\{\gamma_w, F_w(\cdot)\}_{w=0,...,S}$.

We let $\mathcal{D}$ be the set of feasible prices available to the seller, which we assume to be a finite set with cardinality D. The seller selects a sequence of prices $\mathbf{p}_T = (p_1, p_2, ..., p_T)$. We denote the set of available pricing policies by $\mathcal{P}_T = \mathcal{D}^T$. Consumers are assumed to be patient but not strategic. That is, a consumer will buy the product as soon as the price is equal to or below her valuation. If a consumer with patience w encounters a price above her valuation, she will wait for the price to drop for up to w periods. A consumer arriving at time t with patience w and valuation v will exit the system without purchasing if $v < \min\{p_t, p_{t+1}, ..., p_{t+w}\}$.

Consider a consumer arriving at period t' with patience w and valuation v. For any period $t \in \{t' + 1, t' + 2, ..., t' + w\}$, the consumer will still be in the system at time t if $v < \min\{p_{t'}, ..., p_{t-1}\}$ and will purchase at time t if $v \in [p_t, \min\{p_{t'}, ..., p_{t-1}\})$. Therefore, the fraction of consumers with patience w who arrived at period t' who purchase at time $t \in \{t' + 1, t' + 2, ..., t' + w\}$ is given by $(F_w(\min\{p_{t'}, ..., p_{t-1}\}) - F_w(p_t))^+$, where we use the notation $x^+ = \max\{x, 0\}$. In order to calculate the seller's overall revenue, we define

$R_{t,w}(\mathbf{p}_T)$ as the revenue earned by the seller at period t by selling to consumers with patience w, which is equal to

$$R_{t,w}(\mathbf{p}_T) = \gamma_w p_t \left[(1 - F_w(p_t)) + \sum_{i=1}^{w} (F_w(\min\{p_{t-i}, ..., p_{t-1}\}) - F_w(p_t))^+ \right], \tag{1}$$

where for any $t \leq 0$, the value of p_t is defined to be zero. The first term inside the brackets above, $(1 - F_w(p_t))$, represents the fraction of the consumer population arriving at time t with patience w that purchases at time t itself, while the summation captures purchases by consumers who arrived with patience w strictly before time t.

Without loss of generality, we assume the seller's product has a marginal cost of zero. We also assume the seller does not face production capacity constraints, that there are no resale markets, and that there are no costs to adjusting prices. The seller's goal is to maximize his revenue over the entire time horizon $t = 1, 2, ..., T$. The total revenue is nothing but the sum of all revenues earned by the seller in each period t by selling to the consumers with each patience level w:

$$Z_T(\mathbf{p}_T) = \sum_{t=1}^{T} \sum_{w=0}^{S} R_{t,w}(\mathbf{p}_T). \tag{2}$$

The seller's problem is to choose a policy $\mathbf{p}_T \in \mathcal{P}_T$ to maximize the total revenue:

$$\max_{\mathbf{p}_T \in \mathcal{P}_T} Z_T(\mathbf{p}_T). \tag{3}$$

3 A Dynamic Programming Formulation

In this section, we develop a dynamic programming solution to the problem in Eq. (3). We begin by introducing the notion of cutoff prices. For a given pricing policy $\mathbf{p}_T \in \mathcal{P}_T$, we represent the period t cutoff price for consumers arriving at period $t - i$ by $c_{t,i}(\mathbf{p}_T)$. The cutoff price $c_{t,i}(\mathbf{p}_T)$ represents the least upper bound of consumer valuations that is present in the system at the beginning of period t among consumers who arrived in the system at period $t - i$. Formally,

$$c_{t,i}(\mathbf{p}_T) = \begin{cases} \min\{p_{t-i}, p_{t-i+1}, ..., p_{t-1}\} & \text{if } i \geq 1; \\ \infty & \text{if } i = 0. \end{cases}$$

Whenever $i = 0$, the cutoff price is infinity, since all consumers that arrived at period t are still present at period t; that is, $F_w(c_{t,0}(\mathbf{p}_T)) = 1$ for all t and w. As before, p_t is defined to be equal to zero for all $t \le 0$. Thus, $c_{t,i}(\mathbf{p}_T) = 0$ whenever $i \ge t$, since no consumers arrive in the system before period $t = 1$. Using the definition of cutoff prices, we can simplify the expression in Eq. (1) that computes the revenue earned by the seller at period t by selling to consumers with patience w:

$$R_{t,w}(\mathbf{p}_T) = \gamma_w \; p_t \; \sum_{i=0}^{w} (F_w(c_{t,i}(\mathbf{p}_T)) - F_w(p_t))^+.$$

To simplify this expression further, we define $H_w(q, r)$ to be the revenue from sales to a group of consumers with patience w with valuations at most q when the price on offer is r:

$$H_w(q, r) = \gamma_w \; r \; (F_w(q) - F_w(r))^+.$$

Thus, Eq. (1) can be written more succinctly as

$$R_{t,w}(\mathbf{p}_T) = \sum_{i=0}^{w} H_w(c_{t,i}(\mathbf{p}_T), p_t), \tag{4}$$

and the overall revenue function in Eq. (2) reduces to

$$Z_T(\mathbf{p}_T) = \sum_{t=1}^{T} \sum_{w=0}^{S} \sum_{i=0}^{w} H_w(c_{t,i}(\mathbf{p}_T), p_t). \tag{5}$$

The key to understanding the revenue function, therefore, is understanding the cutoff prices. We now develop a method for decomposing the revenue function that takes advantage of the structure of the cutoff prices.

Revenue decomposition. We now propose a methodology for analyzing the revenue generated by a given pricing policy $\mathbf{p}_T \in \mathcal{P}_T$. Eq. (5) suggests that the revenue of a pricing policy can be computed by adding all the elements in a $T \times (S+1) \times (S+1)$ tensor. We call this object the *cutoff prices tensor*. We can represent this tensor via a $T \times (S+1)$ matrix where every entry (t, w) includes the entire sequence of cutoff prices $(c_{t,0}(\mathbf{p}), c_{t,1}(\mathbf{p}), ..., c_{t,w}(\mathbf{p}))$.

Let us consider the case of policies where the lowest price is offered at time T; that is, $p_T = \min\{p_1, p_2, ..., p_T\}$.[1] Let k be a period where the lowest price is offered other than p_T,

[1] We note that the optimal policy that we eventually find might not belong to this class.

i.e., $k \in \arg\min_{i\in\{1,2,...,T-1\}} p_i$.[2] We will now propose a decomposition of the revenue of a policy in this class into the revenues of two subpolicies applied to two smaller subproblems, one encompassing periods 1 through k and the other encompassing periods $k+1$ through T, plus an additional revenue term. This additional revenue term is given by a new function $Y_{k,T}(q,r)$, which accounts for the profits from sales that occur in period T to consumers that arrived in periods $1,...,k$, assuming that $p_k = q$ and $p_T = r$. It takes the form

$$Y_{k,T}(q,r) = \sum_{w=0}^{S} (\min\{w+k+1-T, k\})^+ \cdot H_w(q,r). \tag{6}$$

Lemma 1 (Decomposition). *Consider a policy $\mathbf{p}_T \in \mathcal{P}_T$ such that $p_T = \min_{i\in\{1,2,...,T\}} p_i$. Let $k \in \arg\min_{i\in\{1,2,...,T-1\}} p_i$. Let $Z(\cdot)$ be as defined in Eq.* (5) *and $Y(\cdot,\cdot)$ be as defined in Eq.* (6)*. Then, the revenue of policy $\mathbf{p}_T$ satisfies*

$$Z_T(\mathbf{p}_T) = Z_k(p_1,...,p_k) + Z_{T-k}(p_{k+1},...,p_T) + Y_{k,T}(p_k,p_T). \tag{7}$$

Proof. We first decompose the seller's revenue function from Eq. (5) into three parts. The first part represents revenues obtained in periods 1 to k, the second one is for revenues obtained in periods $k+1$ through $T-1$ and the third is for revenues obtained in period T:

$$Z_T(\mathbf{p}_T) = \sum_{t=1}^{k}\sum_{w=0}^{S}\sum_{i=0}^{w} H_w(c_{t,i}(\mathbf{p}_T),p_t) + \sum_{t=k+1}^{T-1}\sum_{w=0}^{S}\sum_{i=0}^{w} H_w(c_{t,i}(\mathbf{p}_T),p_t) + \sum_{w=0}^{S}\sum_{i=0}^{w} H_w(c_{T,i}(\mathbf{p}_T),p_T). \tag{8}$$

We note that the decomposition in Eq. (8) is not the same as the decomposition in Eq. (7).

The first term on the right-hand side of Eq. (8) depends only on prices $p_1, p_2, ..., p_k$. Therefore, we obtain

$$\sum_{t=1}^{k}\sum_{w=0}^{S}\sum_{i=0}^{w} H_w(c_{t,i}(\mathbf{p}_T),p_t) = Z_k(p_1,...,p_k). \tag{9}$$

We cannot immediately make the same statement for the second term on the right-hand side of Eq. (8), since prices $p_1,...,p_k$ do appear inside the cutoff prices $c_{t,i}(\mathbf{p}_T)$ whenever $t-i \leq k$. In fact, whenever $t-i \leq k$, $c_{t,i}(\mathbf{p}_T) = p_k$, since $k \in \arg\min_{i\in\{1,2,...,T-1\}} p_i$. Since

[2] If there are multiple minimizers, any of them can be chosen as k. It is also valid to have $p_T = p_k$.

$p_k \le p_t$ for all $t \le T-1$, we have $H_w(p_k, p_t) = 0$ for all $t \le T-1$ by the definition of $H_w(\cdot,\cdot)$. Let us create a new pricing policy $\hat{\mathbf{p}}_T$ with the first k prices of $\mathbf{p}_T$ replaced by zero, i.e., $\hat{\mathbf{p}}_T = (0, ..., 0, p_{k+1}, ..., p_T)$. This new policy $\hat{\mathbf{p}}_T$ will also have $H_w(c_{t,i}(\hat{\mathbf{p}}_T), \hat{p}_t) = 0$ whenever $t-i \le k$ and $t \ge k+1$ and $H_w(c_{t,i}(\hat{\mathbf{p}}_T), p_t) = H_w(c_{t,i}(\mathbf{p}_T), p_t)$ whenever $t-i > k$. Therefore,

$$\sum_{t=k+1}^{T-1} \sum_{w=0}^{S} \sum_{i=0}^{w} H_w(c_{t,i}(\mathbf{p}_T), p_t) = \sum_{t=k+1}^{T-1} \sum_{w=0}^{S} \sum_{i=0}^{w} H_w(c_{t,i}(\hat{\mathbf{p}}_T), \hat{p}_t).$$

Since the first k prices are equal to zero in $\hat{\mathbf{p}}_T$, we can also add those first k periods to our summation without changing our results:

$$\sum_{t=k+1}^{T-1} \sum_{w=0}^{S} \sum_{i=0}^{w} H_w(c_{t,i}(\mathbf{p}_T), p_t) = \sum_{t=1}^{T-1} \sum_{w=0}^{S} \sum_{i=0}^{w} H_w(c_{t,i}(\hat{\mathbf{p}}_T), \hat{p}_t). \tag{10}$$

We now focus on the third term from the right-hand side of Eq. (8). Along the same line as the analysis in the paragraph above, we will separate our analysis into two groups of terms depending on whether $c_{T,i}(\mathbf{p}_T) = p_k$ for a given i:

$$\sum_{w=0}^{S} \sum_{i=0}^{w} H_w(c_{T,i}(\mathbf{p}_T), p_T) = \sum_{w=0}^{S} \sum_{\substack{i=0 \\ T-i>k}}^{w} H_w(c_{T,i}(\mathbf{p}_T), p_T) + \sum_{w=0}^{S} \sum_{\substack{i=0 \\ T-i\le k}}^{w} H_w(c_{T,i}(\mathbf{p}_T), p_T). \tag{11}$$

For the values of i where $T-i > k$, $c_{T,i}(\mathbf{p}_T) = c_{T,i}(\hat{\mathbf{p}}_T)$, since the these cutoff prices do not depend on the first k prices. Therefore,

$$\sum_{w=0}^{S} \sum_{\substack{i=0 \\ T-i>k}}^{w} H_w(c_{T,i}(\mathbf{p}_T), p_T) = \sum_{w=0}^{S} \sum_{\substack{i=0 \\ T-i>k}}^{w} H_w(c_{T,i}(\hat{\mathbf{p}}_T), \hat{p}_T).$$

Note that $c_{T,i}(\mathbf{p}_T) = 0$ whenever $T-i \le k$. Therefore, $H_w(c_{T,i}(\hat{\mathbf{p}}_T), p_T) = 0$ whenever $T-i \le k$. We can therefore include these terms in the summation above as well:

$$\sum_{w=0}^{S} \sum_{\substack{i=0 \\ T-i>k}}^{w} H_w(c_{T,i}(\mathbf{p}_T), p_T) = \sum_{w=0}^{S} \sum_{\substack{i=0 \\ T-i>k}}^{w} H_w(c_{T,i}(\hat{\mathbf{p}}_T), \hat{p}_T) = \sum_{w=0}^{S} \sum_{i=0}^{w} H_w(c_{T,i}(\hat{\mathbf{p}}_T), \hat{p}_T). \tag{12}$$

We now consider the summation of terms in Eq. (11) with $T-i \le k$. For the terms where $T-i \le 0$, we have $c_{T,i}(\mathbf{p}_T) = 0$ and $H_w(c_{T,i}(\mathbf{p}_T), p_T) = 0$. For the other terms, where $1 \le T-i \le k$, $c_{T,i}(\mathbf{p}_T) = p_k$, since $k \in \arg\min_{i\in\{1,2,...,T-1\}} p_i$. Therefore,

$$\sum_{w=0}^{S} \sum_{\substack{i=0 \\ T-i\le k}}^{w} H_w(c_{T,i}(\mathbf{p}_T), p_T) = \sum_{w=0}^{S} \sum_{\substack{i=0 \\ 1\le T-i\le k}}^{w} H_w(p_k, p_T) = \sum_{w=0}^{S} H_w(p_k, p_T) \sum_{\substack{i=0 \\ 1\le T-i\le k}}^{w} 1.$$

Computing the value of that final summation in the equation above, we obtain

$$\sum_{\substack{i=0\\1\le T-i\le k}}^{w} 1 = \sum_{i=T-k}^{\min\{w,T-1\}} 1 = (\min\{w+k+1-T,k\})^+,$$

where the $(\cdot)^+$ operator accounts for the fact that $\min\{w, T-1\}$ could be strictly smaller than $T-k$, in which case, the summation should yield zero. Thus,

$$\sum_{w=0}^{S}\sum_{\substack{i=0\\T-i\le k}}^{w} H_w(c_{T,i}(\mathbf{p}_T), p_T) = \sum_{w=0}^{S}(\min\{w+k+1-T,k\})^+ \cdot H_w(p_k, p_T) = Y_{k,T}(p_k, p_T), \quad (13)$$

where the second equality follows from the definition of $Y_{k,t}(\cdot,\cdot)$ in Eq. (6).

Combining Eqs. (9), (10), (11), (12) and (13), we obtain

$$\begin{aligned}
Z_T(\mathbf{p}_T) = \ & Z_k(p_1, ..., p_k) + \sum_{t=1}^{T-1}\sum_{w=0}^{S}\sum_{i=0}^{w} H_w(c_{t,i}(\hat{\mathbf{p}}_T), \hat{p}_t) \\
& + \sum_{w=0}^{S}\sum_{i=0}^{w} H_w(c_{T,i}(\hat{\mathbf{p}}_T), \hat{p}_T) + Y_{k,T}(p_k, p_T) \\
= \ & Z_k(p_1, ..., p_k) + \sum_{t=1}^{T}\sum_{w=0}^{S}\sum_{i=0}^{w} H_w(c_{t,i}(\hat{\mathbf{p}}_T), \hat{p}_t) + Y_{k,T}(p_k, p_T) \\
= \ & Z_k(p_1, ..., p_k) + Z_T(\hat{\mathbf{p}}_T) + Y_{k,T}(p_k, p_T),
\end{aligned}$$

where the second equality follows from the definition of $Z_T(\cdot)$ from Eq. (5). Since the first k prices in $\hat{\mathbf{p}}_T$ are equal to zero, the total revenue obtained from the pricing policy $\hat{\mathbf{p}}_T$ must be identical to the revenue obtained by using a policy $(p_{k+1}, p_{k+2}, ..., p_T)$ in a $T-k$ horizon problem. Thus, $Z_T(\hat{\mathbf{p}}_T) = Z_{T-k}(p_{k+1}, ..., p_T)$, completing our proof. □

Bellman equation. We have proved that the revenue of a pricing policy $\mathbf{p}_T$ could be decomposed into three parts, as long as the lowest price was used in the period T. We now show how to leverage this decomposition to prove a Bellman equation for our problem.

Before we define our value function, we construct a set of pricing policies over which the value function will optimize. For any $q, r \in \mathcal{D}$ such that $q \ge r$, we define $\mathcal{P}_T(q, r) = \{\mathbf{p}_T \in \mathcal{P}_T \mid p_t \ge q \text{ for all } t \in \{1, ..., T-1\} \text{ and } p_T = r\}$. We can now define the value function $V_T(q,r)$ as the maximum revenue possible for a horizon of length T where only pricing policies within $\mathcal{P}_T(q,r)$ are allowed. Formally, for any $q,r \in \mathcal{D}$ such that $q \geq r$, we define

$$V_T(q,r) = \max_{\mathbf{p}_T \in \mathcal{P}_T(q,r)} Z_T(\mathbf{p}_T). \tag{14}$$

If we design an efficient algorithm for computing $V_T(q,r)$ for all T and $q \geq r$, we will have solved our problem, since we can add a price equal to zero to our set of available prices $\mathcal{D}$ and then compute $\max_{\mathbf{p}_T \in P_T} Z_T(\mathbf{p}_T) = V_{T+1}(0,0)$. The additional period at time $T+1$ with price zero does not lead to any additional revenues. Furthermore, the price zero will not be used in periods before $T+1$, since it generates no revenues and clears the system. The value $V_{T+1}(0,0)$ corresponds to the value of an optimal pricing policy since it imposes no constraints on prices $p_1, ..., p_T$.

We now prove the value function $V_T(q,r)$ satisfies a Bellman equation. The Bellman equation builds on the revenue decomposition from Lemma 1.

Lemma 2 (Bellman)**.** *For any $T \in \mathbb{N}$ and $q,r \in \mathcal{D}$ such that $q \geq r$, the value function $V_T(q,r)$ satisfies the following Bellman equation:*

$$V_T(q,r) = \max_{\substack{k \in \{1,2,\dots,T-1\} \\ x \in \mathcal{D}:\ x \geq q}} \left\{ V_k(x,x) + V_{T-k}(x,r) + Y_{k,T}(x,r) \right\}.$$

Proof. Let us define the set of prices $\tilde{\mathcal{P}}_T(q,r,k,x)$ as the subset of prices within $\mathcal{P}_T(q,r)$ where k is the period with the lowest price among periods 1 through $T-1$ and $p_k = x$:

$$\tilde{\mathcal{P}}_T(q,r,k,x) = \left\{ \mathbf{p}_T \in \mathcal{P}_T(q,r) \;\middle|\; k \in \operatorname*{arg\,min}_{i \in \{1,2,\dots,T-1\}} p_i \ \text{ and } \ p_k = x \right\}.$$

Using the definition above, we can decompose the problem in Eq. (14) of optimizing prices over $\mathcal{P}_T(q,r)$ into a sequence of two problems: we first optimize over the period of the lowest price k among $1, ..., T-1$ and the price p_k, and then we optimize over prices that satisfy these two properties. That is,

$$V_T(q,r) = \max_{\substack{k \in \{1,2,\dots,T-1\} \\ x \in \mathcal{D}:\ x \geq q}} \max_{\mathbf{p}_T \in \tilde{\mathcal{P}}_T(q,r,k,x)} Z_T(\mathbf{p}_T).$$

Since the lowest price of policy $p_T \in \tilde{\mathcal{P}}_T(q,r,k,x)$ is used in period T, Lemma 1 applies:

$$V_T(q,r) = \max_{\substack{k\in\{1,2,...,T-1\} \\ x\in\mathcal{D}:\ x\geq q}} \max_{\mathbf{p}_T\in\tilde{\mathcal{P}}_T(q,r,k,x)} \{Z_k(p_1,...,p_k) + Z_{T-k}(p_{k+1},...,p_T) + Y_{k,T}(p_k,p_T)\}.$$

The values of k, $p_k = x$ and $p_T = r$ are constants within all policies in $\tilde{\mathcal{P}}_T(q,r,k,x)$. We can therefore move the last term outside the price maximization problem:

$$V_T(q,r) = \max_{\substack{k\in\{1,2,...,T-1\} \\ x\in\mathcal{D}:\ x\geq q}} \left\{ Y_{k,T}(x,r) + \max_{\mathbf{p}_T\in\tilde{\mathcal{P}}_T(q,r,k,x)} \{Z_k(p_1,...,p_k) + Z_{T-k}(p_{k+1},...,p_T)\} \right\}.$$

The term $Z_k(p_1,...,p_k)$ depends only on p_1 through p_k, and the term $Z_{T-k}(p_{k+1},...,p_T)$ depends only on p_{k+1} through p_T. Furthermore, the choice of p_1 through p_k imposes no restrictions on the choice of p_{k+1} through p_T, given that $\mathbf{p}_T \in \tilde{\mathcal{P}}_T(q,r,k,x)$, or vice versa. Therefore, we are free to optimize the prices separately. Formally, we do so by constructing two complete pricing policies $\mathbf{p}_T, \tilde{\mathbf{p}}_T \in \tilde{\mathcal{P}}_T(q,r,k,x)$:

$$V_T(q,r) = \max_{\substack{k\in\{1,...,T-1\} \\ x\in\mathcal{D}:\ x\geq q}} \left\{ Y_{k,T}(x,r) + \max_{\tilde{\mathbf{p}}_T\in\tilde{\mathcal{P}}_T(q,r,k,x)} Z_k(\tilde{p}_1,...,\tilde{p}_k) + \max_{\mathbf{p}_T\in\tilde{\mathcal{P}}_T(q,r,k,x)} Z_{T-k}(p_{k+1},...,p_T) \right\}. \tag{15}$$

The maximization over $\tilde{\mathbf{p}}_T$ does not depend in any way on the values of T and r. It also satisfies the property that p_k is the lowest price within periods 1 through k. All prices above $p_k = x$ can be used in periods 1 through $k-1$. Therefore,

$$\max_{\tilde{\mathbf{p}}_T\in\tilde{\mathcal{P}}_T(q,r,k,x)} Z_k(\tilde{p}_1,...,\tilde{p}_k) = \max_{\tilde{\mathbf{p}}_k\in\mathcal{P}_k(x,x)} Z_k(\tilde{\mathbf{p}}_k) = V_k(x,x). \tag{16}$$

For the maximization over $\mathbf{p}_T$ in Eq. (15), the problem can be recast as an optimization over $T-k$ periods, with the lowest price being used in the final period, $p_{T-k} = r$. However, only prices above x can be used in periods $k+1$ through $T-1$. Otherwise, the original pricing policy $\mathbf{p}_T$ would not correspond to a policy in $\tilde{\mathcal{P}}_T(q,r,k,x)$. Thus,

$$\max_{\mathbf{p}_T\in\tilde{\mathcal{P}}_T(q,r,k,x)} Z_{T-k}(p_{k+1},...,p_T) = \max_{\mathbf{p}_{T-k}\in\mathcal{P}_{T-k}(x,r)} Z_{T-k}(\tilde{\mathbf{p}}_{T-k}) = V_{T-k}(x,r). \tag{17}$$

We obtain the desired result by combining Eqs. (15), (16) and (17). $\square$

A polynomial-time algorithm. Using Lemma 2, we can write a dynamic program to compute an optimal pricing policy. This algorithm requires $O(D^2T^2)$ steps, where a step is an arithmetic operation, such as summing, multiplying or comparing two numbers.

Theorem 1. *There exists an algorithm that finds an optimal* $\mathbf{p}_T \in \mathcal{P}_T$ *in* $O(D^2T^2)$ *steps.*

Proof. The algorithm begins by precomputing the value of $Y_{k,T}(q,r)$ for all k, T, q and r, where $Y_{k,T}(q,r)$ is defined in Eq. (6). A naïve algorithm for computing the values of $Y_{k,T}(q,r)$ would require a relatively large number of steps. We can speed up the precomputation phase by taking advantage of the structure of $Y_{k,T}(q,r)$.

For any $w \in \{0,1,...,S\}$ and $q,r \in \mathcal{D}$ with $q \geq r$, define $G_w(q,r)$ to be equal to $\sum_{i=w}^{S} H_i(q,r)$. Note that the partial derivative of $(\min\{w+k+1-T,k\})^+$ with respect to w is equal to one whenever $(\min\{w+k+1-T,k\})^+ = w+k+1-T$ and zero when $w > T-1$ or $w < k$. Therefore, for any k and T, there exists two numbers $a(k,T)$ and $b(k,T)$ such that $Y_{k,T}(q,r) = \sum_{w=0}^{S}(\min\{w+k+1-T,k\})^+ \cdot H_w(q,r) = \sum_{w=a(k,T)}^{b(k,T)} G_w(q,r)$. For some values of k and T, we have $Y_{k,T}(q,r) = 0$, a case we can represent by a sum where $a(k,T) > b(k,T)$. To compute the values of $Y_{k,T}(q,r)$, we first compute all values of $G_w(q,r)$ in $O(D^2S)$ steps. We then compute the values of $\sum_{w=a}^{b} G_w(q,r)$ for all $q \geq r$ and all integer values of $0 \leq a \leq b \leq S$. This task can be accomplished in $O(D^2S^2)$ steps by first computing all sums where $b-a=1$, then using those results to compute all sums where $b-a=2$, and so on. We then compute all values of $a(k,T)$ and $b(k,T)$ and assign $Y_{k,T}(q,r) = \sum_{w=a(k,T)}^{b(k,T)} G_w(q,r)$ in $O(D^2T^2)$ steps.

We now proceed to the recursion. We first compute $V_1(q,r) = V_1(r,r) = \sum_{w=0}^{S} H_w(\infty,r)$ for all q and r in $O(DS)$ steps, where $V_t(q,r)$ is defined in Eq. (14). We then proceed by induction. Suppose we have computed $V_k(q,r)$ for all $q,r \in \mathcal{D}$ where $q \geq r$ and all $k \leq t-1$. We now compute $V_t(q,r)$ for all $q,r \in \mathcal{D}$ where $q \geq r$. For all $k \in \{1,...,t-1\}$ and $x,r \in \mathcal{D}$ with $x \geq r$, define $\widehat{V}_t(x,r) = \max_{k \in \{1,2,...,t-1\}} V_k(x,x) + V_{t-k}(x,r) + Y_{k,t}(x,r)$. We can compute the value of $\widehat{V}_t(x,r)$ for every x and r in $O(D^2t)$ steps by taking advantage of our precomputation of $Y_{k,t}(x,r)$. The Bellman equation from Lemma 2 gives us that the value function $V_t(q,r)$ satisfies $V_t(q,r) = \max_{\{x \in \mathcal{D}:\ x \geq q\}} \widehat{V}_t(x,r)$. Let x_1 and x_2 be consecutive elements in $\mathcal{D}$ such that $x_2 > x_1 \geq r$. Then, $V_t(x_1, r) = \max\left\{V_t(x_2, r), \widehat{V}_t(x_1, r)\right\}$. To go from having $\widehat{V}_t(q, r)$ for all $q \geq r$ to having $V_t(q, r)$ for all $q \geq r$ thus requires $O(D^2)$ steps. Therefore, we can compute $V_t(q, r)$ for all $q \geq r$ in $O(D^2 t)$ steps. Repeating this for every t until $T + 1$ requires $O(D^2 T^2)$ steps. Overall, this algorithm requires $O(D^2 T^2)$ steps. $\square$

Extension: infinite horizon model. While our work studies a finite horizon model, our results can also be used to compute optimal policies for a long-term average infinite horizon model. To do so, we can leverage Proposition 8 from Liu and Cooper [2015], which shows that optimal policies are cyclical in such a model, with a bound on the periodicity of $SD+1$, assuming S is finite. We can find an optimal policy for the infinite horizon model by considering all possible values of $T \leq SD + 1$ and then computing $\max_{t \in \{1,2,\ldots,SD+1\}} V_{t+1}(0, 0)/t$. The computational complexity of this algorithm is $O(D^4 S^2)$.

Extension: nonstationarity. Our paper assumes a stationary demand model. Our results can be extended to a nonstationary model by considering a state space $V_{t,t'}(q, r)$ where t and t' represent the initial and final periods being considered rather than merely a duration. This technique is analogous to the one used in Theorem 5 of Besbes and Lobel [2015].

4 Structure of Optimal Policies: A Numerical Study

We now consider the structure of optimal pricing policies. We do so via a numerical case study using a specific patient demand model. We consider a model where each patience level is represented by a linear demand function. We assume the maximum willingness-to-wait is $S = 11$ and the time horizon is $T = 40$. The set of prices available is $\mathcal{D} = \{0, 0.01, 0.02, ..., 1\}$. We assume $\gamma_w = 1$ for all $w \in \{0, ..., S\}$. For every $p \geq 0$ and every $w \in \{0, .., S\}$, we assume the consumer valuation distribution follows a uniform distribution with $F_w(p) = \min\{p(w + 1), 1\}$. The uniform assumption is not necessary for our numerical results to hold: similar patterns can obtained assuming valuations are drawn from light-tailed distributions such as exponential as well as from heavy-tailed ones such as Pareto.

With a single patience class (besides consumers with $w = 0$), Liu and Cooper [2015] showed that optimal policies for patient consumers are cyclic decreasing. Liu and Cooper

[2015] do not attempt to understand the structure of optimal prices in the presence of heterogenous patience levels. In contrast, Besbes and Lobel [2015] show numerically that optimal policies for strategic consumers often have a fractal-like structure they call "nested sales." Nested sales are symmetric with respect to time and include multiple promotion cycles of different depth and periodicity overlaid on top of each other (see Figure (a)).

With heterogeneous patience levels, the optimal policy often combines features of both cyclic decreasing policies and nested sales policies. Figure (b) shows the optimal policy for the linear demand system described above. The policy has some of the fractal-like structure of a nested sales policy and some of the time asymmetry of a cyclic decreasing policy. We call policies of this form *incomplete cyclic* policies.

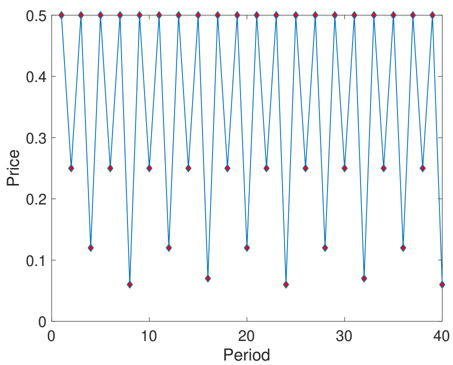

(a) Optimal prices for strategic consumers.

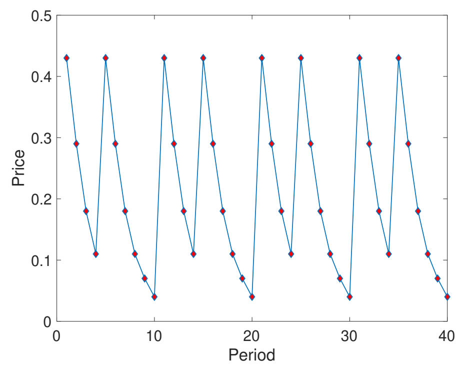

(b) Optimal prices for patient consumers.

The rationale for such policies is as follows. If the prices are relatively high for a while, a large number of low-value consumers accumulate. At that point, it becomes in the firm's best interest to use low prices. Low prices clear both low-value and high-value consumers from the market, leading the cycle of prices to restart. The exception are even-lower-value consumers, who remain in the system because the prices they find attractive haven't been offered yet. This encourages the firm to go deeper in the next round of discounts.

We now compare revenues and optimal price levels obtained in the strategic and patient consumer models. We also consider the best fixed price policy as a benchmark since the revenue of a fixed price policy does not depend on whether consumers are patient or strategic. Table 1 shows the average, minimum and maximum prices for the three different cases, assuming the linear demand system described above. It also shows the revenues of the different policies, with the revenue of the fixed price policy normalized to one.

	Fixed Price	Strategic Consumers	Patient Consumers
Average Price	0.08	0.336	0.213
Minimum Price	0.08	0.06	0.04
Maximum Price	0.08	0.50	0.43
Normalized Revenue	1	1.171	1.349

Table 1: Prices and revenues of (i) the best fixed price policy, (ii) the optimal prices for strategic consumers and (iii) the optimal prices for patient consumers.

We find that the seller earns higher revenues and uses lower prices (minimum, maximum and average) when facing patient consumers compared to the strategic consumers case. Using lower prices and earning higher revenues mean that the seller makes more sales in the patient consumers model. The maximum and the minimum prices are lower under patient consumers for different reasons. The maximum price is lower under patient consumers because it is not as targeted to myopic ($w = 0$) consumers. In a strategic consumers model, only myopic consumers are affected by the maximum price, and therefore the highest price is tailored to them. In a patient consumers model, the highest price affects consumers of all patience levels arriving when that price is in effect, which has a moderating effect on the maximum price. On the other hand, the minimum price is much more targeted toward the low-value consumers in the patient model than in the strategic one. Under strategic consumers, the minimum price serves as a vortex that attracts a large number of consumers to it. In the patient consumers case, the lowest price serves only immediate arrivals and consumers who did not buy at the previous period's price. Thus, the seller is freer to price the product to attract low-value consumers in the patient model. Overall, the seller is capable of better price discrimination in the patient consumers model, and can do so via a strategy that relies on price skimming over intervals of time, a technique that would backfire in a strategic consumers setting. Therefore, the seller is able to obtain higher revenues and increase sales in a patient consumers model when compared to a strategic consumers one.

References

Hyun-soo Ahn, Mehmet Gümüş, and Philip Kaminsky. Pricing and Manufacturing Decisions When Demand Is a Function of Prices in Multiple Periods. *Operations Research*, 55(6): 1039–1057, 2007.

Victor Araman and Bassam Fayad. Intertemporal Price Discrimination with Time-Varying Valuations. *Working Paper, American University of Beirut*, 2017.

Yossi Aviv and Amit Pazgal. Optimal Pricing of Seasonal Products in the Presence of Forward-looking Consumers. *Manufacturing & Service Operations Management*, 10(3): 339–359, 2008. ISSN 1526-5498.

David Besanko and Wayne L Winston. Optimal Price Skimming by a Monopolist Facing Rational Consumers. *Management Science*, 36(5):555–567, 1990.

Omar Besbes and Ilan Lobel. Intertemporal Price Discrimination: Structure and Computation of Optimal Policies. *Management Science*, 61(1):92–110, 2015.

Luis Briceno, José Correa, and Andres Perlroth. Optimal Continuous Pricing with Strategic Consumers. *Management Science*, 63(8):2741—-2755, 2017.

René Caldentey, Ying Liu, and Ilan Lobel. Intertemporal pricing under minimax regret. *Operations Research*, 65(1):104–129, 2017.

Andrew Caplin, Mark Dean, and Daniel Martin. Search and Satisficing. *The American Economic Review*, 101(7):2899–2922, 2011.

Xi Chen and Zizhuo Wang. Bayesian Dynamic Learning and Pricing with Strategic Customers. *Working paper, New York University*, 2016.

Yiwei Chen and Vivek F Farias. Robust Dynamic Pricing With Strategic Customers. *Mathematics of Operations Research*, 43(4):1119—-1142, 2018.

Yiwei Chen and Cong Shi. Joint Pricing and Inventory Management with Strategic Customers. *Working paper, University of Michigan*, 2016.

Maxime C Cohen, Ngai-Hang Z Leung, Kiran Panchamgam, Georgia Perakis, and Anthony Smith. The Impact of Linear Optimization on Promotion Planning. *Operations Research*, 65(2):446–468, 2017.

John Conlisk, Eitan Gerstner, and Joel Sobel. Cyclic Pricing by a Durable Goods Monopolist. *The Quarterly Journal of Economics*, 99:489–505, 1984.

José Correa, Ricardo Montoya, and Charles Thraves. Contingent Preannounced Pricing Policies with Strategic Consumers. *Operations Research*, 64(1):251–272, 2016.

Igal Hendel and Aviv Nevo. Intertemporal Price Discrimination in Storable Goods Markets. *American Economic Review*, 103(7):2722–2751, 2013.

Zhenyu Hu, Xin Chen, and Peng Hu. Dynamic Pricing with Gain-Seeking Reference Price Effects. *Operations Research*, 64(1):150–157, 2016.

Shlomo Kalish. Monopolist Pricing with Dynamic Demand and Production Cost. *Marketing Science*, 2(2):135–159, 1983.

Praveen K Kopalle, Ambar G Rao, and Joao L Assuncao. Asymmetric Reference Price Effects and Dynamic Pricing Policies. *Marketing Science*, 15(1):60–85, 1996.

Jun Li, Nelson Granados, and Serguei Netessine. Are Consumers Strategic? Structural Estimation from the Air-travel Industry. *Management Science*, 60(9):2114–2137, 2014.

Yan Liu and William Cooper. Optimal Dynamic Pricing with Patient Customers. *Operations Research*, 63(6):1307–1319, 2015.

Ken Moon, Kostas Bimpikis, and Haim Mendelson. Randomized Markdowns and Online Monitoring. *Management Science*, 64(3):1271—1290, 2018.

Javad Nasiry and Ioana Popescu. Dynamic Pricing with Loss-averse Consumers and Peak-end Anchoring. *Operations Research*, 59(6):1361–1368, 2011.

Ioana Popescu and Yaozhong Wu. Dynamic Pricing Strategies with Reference Effects. *Operations Research*, 55(3):413–429, 2007.

Elena Reutskaja, Rosemarie Nagel, Colin F Camerer, and Antonio Rangel. Search Dynamics in Consumer Choice Under Time Pressure: An Eye-Tracking Study. *The American Economic Review*, 101(2):900–926, 2011.

Peter Stüttgen, Peter Boatwright, and Robert T Monroe. A Satisficing Choice Model. *Marketing Science*, 31(6):878–899, 2012.

Xuanming Su. Intertemporal Pricing with Strategic Customer Behavior. *Management Science*, 53(5):726, 2007.

Zizhuo Wang. Technical Note—Intertemporal Price Discrimination via Reference Price Effects. *Operations Research*, 64(2):290–296, 2016.