Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Optimal Continuous Pricing with Strategic Consumers

Luis Briceño-Arias, José R. Correa, Andrés Perlroth

To cite this article:
Luis Briceño-Arias, José R. Correa, Andrés Perlroth (2017) Optimal Continuous Pricing with Strategic Consumers. Management Science 63(8):2741-2755. https://doi.org/10.1287/mnsc.2016.2462

informs
http://pubsonline.informs.org/journal/mnsc/

MANAGEMENT SCIENCE
Vol. 63, No. 8, August 2017, pp. 2741–2755
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Optimal Continuous Pricing with Strategic Consumers

Luis Briceño-Arias,[a] José R. Correa,[b] Andrés Perlroth[b]

[a] Department of Mathematics, Universidad Técnica Federico Santa María, Santiago, Chile; [b] Department of Industrial Engineering, Universidad de Chile, Santiago, Chile

Contact: luis.briceno@usm.cl (LB-A); correa@uchile.cl (JRC); andresperlrothv@gmail.com (AP)

Received: May 20, 2014
Revised: June 18, 2015; November 2, 2015
Accepted: December 16, 2015
Published Online in Articles in Advance: June 21, 2016

https://doi.org/10.1287/mnsc.2016.2462

Abstract. An important economic problem is that of finding optimal pricing mechanisms to sell a single item when there are a random number of buyers who arrive over time. In this paper, we combine ideas from auction theory and recent work on pricing with strategic consumers to derive the optimal continuous time pricing scheme in this situation. Under the assumption that buyers are split among those who have a high valuation and those who have a low valuation for the item, we obtain the price path that maximizes the seller's revenue. We conclude that, depending on the specific instance, it is optimal to either use a fixed price strategy or to use steep markdowns by the end of the selling season. As a complement to this optimality result, we prove that under a large family of price functions there is an equilibrium for the buyers. Finally, we derive an approach to tackle the case in which buyers' valuations follow a general distribution. The approach is based on optimal control theory and is well suited for numerical computations.

History: Accepted by Yinyu Ye, optimization.

Funding: This work was supported by the Millennium Nucleus Information and Coordination in Networks [Grant ICM/FIC RC130003] and by the Chilean National Commission for Science and Technology (CONICYT) [Grants FONDECYT 11140360, FONDECYT 1160079, and ACT1106].

Keywords: revenue management • strategic consumers • optimal pricing

1. Introduction

In many practical situations, particularly when selling items online, the precise number of potential consumers is unknown. Therefore, studying auction mechanisms with a random number of bidders has been an important question in economic theory since the work of McAfee and McMillan (1987). Significant effort has been put in understanding these type of auctions under different assumptions (see, e.g., Levin and Smith 1994, Levin and Ozdenoren 2004, Haviv and Milchtaich 2012). However, these works assume a static situation in which the action takes place at one time (or in two rounds) and potential bidders are always present. Recent work in economic theory, including that of Gershkov et al. (2014) and Board and Skrzypacz (2016), considers a more general setting in which buyers arrive over time and fully strategize their decisions. However, this work assumes that the seller has full flexibility and can design any type of mechanism, rather than just the posted price mechanisms more common in the pricing literature.

The issues of intertemporal price discrimination and that of limited pricing flexibility have been central in the area of revenue management, which is concerned with price discrimination over time when selling perishable goods (see Talluri and van Ryzin 2005 for an in-depth treatment). Typically, revenue management theory and practice firms use posted price mechanisms with a price that may vary over time. On the other hand, a frequent assumption in early revenue management work is that customers are not forward-looking. This assumption is usually violated because customers anticipate the pricing policy and incorporate such knowledge in their purchase decision, thus influencing the firm's pricing decision. The evidence that consumers act strategically (Li et al. 2014) has opened a new line of research that analyzes the conditions under which different pricing policies optimize the firm's profitability (Caldentey and Vulcano 2007, Aviv and Pazgal 2008, Elmaghraby et al. 2009, Yin et al. 2009, Jerath et al. 2010, Osadchiy and Vulcano 2010, Surasvadi and Vulcano 2013, Cachon and Feldman 2015, Caldentey et al. 2015, Correa et al. 2016). In particular, two types of pricing policies have been studied: ones in which the price depends on the number of remaining items at the end of each period and ones that are fully preannounced. A drawback of this literature is that it assumes that the price can only be changed at discrete time steps, usually limited to two periods.

In this paper we address the problem of pricing an item over time when fully informed and forward-looking rational consumers arrive according to a random process. This situation is common in several economic activities, including real estate markets and electronic commerce. For instance, Mercado Minero in Chile offers second-hand mining machinery (mining is the largest industry in Chile) through a continuous pricing scheme.[1] In their business model, they announce at time 0 the full price curve until the last day, say, 180. Then, a consumer arriving at any time may decide to buy upon arrival or wait until the price goes down enough. However, if a consumer decides to wait, she risks not getting the item because another consumer may purchase it. Other examples of companies using such announced pricing schemes include Lands End overstocks and Dress for Less.

Our model differs from most of those in the revenue management literature in that the pricing can be adjusted continuously. In this sense, our model is similar to that of Su (2007), although consumers' arrivals are structured by a deterministic continuous flow. On the other hand, our work also differs from the recent literature in dynamic mechanism design, particularly the work of Gershkov et al. (2014) and Board and Skrzypacz (2016), in two key aspects. First, we only allow the seller to use posted prices rather than general mechanisms. Second, and probably more important, is the fact that in these papers both the buyers and the seller discount the future at the same rate, which is key for a Myersonian approach. As argued by Pai and Vohra (2013), among others, it is probably more realistic to assume that buyers are more impatient than the seller. Therefore, we follow the classic economic modeling of impatience and consider that buyers possess a temporal discount rate, while the seller does not. It is worth noting that there are alternative models available in the literature, including those proposed by Su (2007), Pai and Vohra (2013), or Mierendorff (2016). In these models, buyers have an arrival time and either a time when they leave the system or an explicit waiting cost.

1.1. Our Results

We consider the situation in which a single item is to be sold over a period of time. Two types of strategic buyers arrive over time: high-value and low-value consumers. Consumers arrive according to independent nonhomogeneous Poisson processes and upon arrival they decide when is most convenient for them to buy. To this end, they balance the price to be paid against the probability of getting the item, and discount the future at rate $\mu \geq 0$. The seller, who does not discount the future, chooses an arbitrary price function over the time period in order to maximize his expected profit. Intuitively, one may think that the optimal policy for the seller is to wait until the end of the time period and then engage in a first price auction. We prove that this intuition is correct when $\mu = 0$, but fails when $\mu > 0$. In this case, we explicitly find the optimal pricing policy for the seller and observe that offering steep discounts by the end of the period is optimal. Interestingly, we observe that the dynamic pricing mechanism obtains more revenue than that of an optimal auction. This additional revenue arises from exploiting the buyers' impatience.

To complement these results, we show that for any fixed continuous price function, possibly with a discontinuity at the end of the season, a mixed strategy equilibrium exists for the buyers. We do this by constructing a symmetric equilibrium, since the standard tools from fixed point theory do not apply to our setting. In this equilibrium, some buyers will buy upon arrival, whereas others will use a mixed strategy over future times. Although not central to our paper, this result may be of practical relevance since it allows sellers to evaluate the expected revenue across a wide array of price functions—an option that may be necessary if, for instance, regulations or practical constraints impede the application of the optimal pricing.

Finally, we extend our approach to a more general setting in which buyers' valuations are arbitrary. Under an assumption about the buyers behavior, we formulate the seller's problem and apply optimal control theory to distill it down to solving a system of ordinary differential equations, which we are able to solve numerically.

From a methodological perspective, our work may be of interest since we depart from the Myersonian approach and design a method based on optimal control formulations that explicitly consider the time. This seems to be crucial to account for the difference in degree of impatience between the buyers and the seller.

1.2. Assumptions

Let us briefly discuss the main assumptions we make in the subsequent analysis. First, we assume that only two types of buyers' valuation are present. Although this makes some of the analysis simpler, we believe that the main conclusions are not altered. Furthermore, this is a common simplifying assumption in the pricing literature (Caldentey and Vulcano 2007, Su 2007, Yin et al. 2009). In Section 4 we relax this two-valuation assumption. Our second assumption on the consumers' side is that they arrive according to independent nonhomogeneous Poisson processes. This is more general than the homogeneous Poisson arrivals imposed in part of the literature, and moreover there is strong empirical evidence that this is a good modeling assumption in electronic commerce (Russo et. al 2010). The third assumption is that consumers discount the future at rate $\mu \geq 0$. Note that under this impatience measure, the buyers' behavior is straightforward: A higher value of μ make buyers more prone to buy earlier, increasing their chances of actually getting the item. We assume that the seller does not discount the future, given that it is reasonable to assume that the vendor is more patient than the buyers. Finally, we assume that a single item is on sale. This is the case in many applications, and quite common in the economic theory literature and some of the revenue management work (see, e.g., Caldentey et al. 2015).

2. Model

Consider a risk neutral seller who wants to sell a single item over a season described by the time interval $[0, T]$. At time 0 the seller commits to a price function for the item $p(t)$, in case the item is available at time t. Consumers, who arrive according to a random process, take the price function as given and make purchasing decisions strategically; namely, they decide to buy at a time (which is at least their arrival time) maximizing their expected utility, which balances the price to pay with the probability of actually obtaining the item. Naturally, we model the game as a Stackelberg game (two stage dynamic game), where in the last stage strategic buyers seek to maximize their profit under the imposed pricing, whereas in the first stage the seller selects a price function to maximize his own profit. We now describe the seller's and buyers' problems.

2.1. The Seller

The vendor's problem is to determine a price function $p(t)$. Since $p(t)$ is public information and consumers are strategic, any choice of $p(t)$ induces a function $f(t)$, denoting the probability that the item is available at time t. Thus, the seller, who is risk neutral and has the distributional knowledge of the buyers, needs to anticipate this function $f(t)$ and selects $p(t)$ to maximize his expected revenue.

2.2. The Buyers

There are two classes of buyers: *high-value* consumers, whose valuation for the item is V, and *low-value* consumers, whose valuation for the item is $v < V$. We consider the buyers' discount rate to be $\mu \geq 0$. This discount rate models impatience on the buyers' side and also the fact that buyers may be risk averse. Both the high-value and low-value buyers arrive to the store according to independent stochastic processes (counting process) with strictly positive interarrival times. We assume for simplicity that the arrival processes are nonhomogeneous Poisson processes with rates $\Lambda(t)$ for high-value buyers and $\lambda(t)$ for low-value buyers, though this assumption may be relaxed for some of our results. We then denote by $Q_k(t)$ (respectively, $q_k(t)$) the probability that exactly k high-value (respectively, low-value) buyers arrive in $[0, t]$. Note that, given that a buyer arrived at time $s < t$, these probabilities also represent the probability that k other buyers arrive in $[0, t]$. The informational assumption we make is that buyers know the arrival process and that other buyers are strategic, but do not have information on how many other buyers have arrived at any point in time.

Given a price function p, a buyer arriving at time $t \in [0, T]$ with valuation $u \in \{v, V\}$ who finds out that the item is still available will buy at a time maximizing her own utility (provided that it is eventually nonnegative). Her utility at time $s \geq t$ is given by

$$(u - p(s))e^{-\mu s}\,\mathbb{P}(\text{getting the item at time } s \mid \text{item is available at time } t).$$

Of course, for a buyer to get the item at time s, the item had to be available at time t. Thus, the event in which the item is available at time t is already included in the event in which the buyer gets the item at s. Therefore, by Bayes' rule, the buyer's utility maximization problem is equivalent to

$$\begin{aligned}\max_{s\geq t} U(s) &:= \max_{s\geq t}(u - p(s))e^{-\mu s}\,\mathbb{P}(\text{getting the item at time } s)\\ &= \max_{s\geq t}(u - p(s))e^{-\mu s}\alpha(s)f(s).\end{aligned}$$

Recall that $f(s)$ denotes the probability that the item is available at time s, so that $\alpha(s)$ represents the probability of actually getting the item given that it is available in a random allocation model. In other words, if $r_k(s)$ represents the probability that there are k other buyers who chose to buy at time s, then $\alpha(s) = \sum_{k=0}^{\infty} r_k(s) \cdot (1/(k+1))$. If the chances that two different buyers decide to buy in the exact same moment are zero, then this term is just $\alpha(s) = 1$. It is easy to observe that the latter happens, for instance, when the price function $p(t)$ is continuous, since in such a case a buyer may slightly anticipate the purchasing decision, thus paying only infinitesimally more, but saving a significant amount in the term $\alpha(s)$.

2.3. Equilibrium of the Second Stage

Naturally, an equilibrium may be defined as a set of strategies for all potential buyers such that they cannot strictly improve their own profit by unilaterally deviating from the current situation; that is, the potential buyers $I = \{1, \ldots, n, \ldots\}$ that may arrive over $[0, T]$ must have a plan of action that is optimal given the plans of the other buyers. Even though our results apply to this general equilibrium concept, because players are ex ante equal, we restrict our attention to symmetric equilibria.

Since the seller does not have incentives to lower the price below v, the game actually occurs only among high-value buyers. Consider, thus, a price function p such that $p(t) > v$ for all $t \in [0, T)$ and $p(T) = v$. We define a mixed strategy profile of the game as a family of distributions H_t over $[t, T]$, such that if a high-value buyer arrives at time t, she buys at a random time chosen according to H_t; that is, we forget about defining a full plan of actions for every possible buyer and just index the strategies by the time when buyers arrive. We define equilibrium in terms of $f(t)$, accounting for the probability that the item is available at time t. Specifically, given a price function p such that $p(t) > v$ for all

$t \in [0, T)$ and $p(T) = v$, we say that a probability function f *is an equilibrium for* p if there exists a family of distributions H_s over $[s, T]$, such that if a high-value buyer arriving at time s buys at a random time chosen according to H_s, it holds that

$$f(t) = \text{probability that the item is available at } t, \text{ given buyers behavior} \tag{1}$$
$$= \sum_{k=0}^{\infty} Q_k(t) \int \prod_{i=1}^{k} (1 - H_{x_i}(t))\, dF_t(x_1, \ldots, x_k),$$
$$\operatorname{supp}(H_t) \subseteq \underset{t \le s \le T}{\arg\max} (V - p(s)) e^{-\mu s} \alpha(s) f(s). \tag{2}$$

Here $F_t(x_1, \ldots, x_k)$ denotes the conditional distribution of k high-value buyer arrivals in $[0, t]$, subject to one buyer having arrived at time 0.

This definition appears to be more natural and simpler: creating a full possible contingent of actions seems much harder to understand. Also, it is easy to see that defining for buyer $i \in I$ with arrival time t its contingent plan as H_t leads to a strategy in the classic Bayesian setting. Therefore, every equilibrium in our definition generates a Bayes–Nash equilibrium almost everywhere. Also, our definition only in terms of $f(t)$ is more robust in the sense that condition (2) could be imposed almost everywhere, so that if we change the distributions H_t for a negligible subset of $[0, T]$, $f(t)$ remains unaltered.

Note that considering mixed strategies is key since equilibrium in pure strategies may fail to exist. For the intuition behind this, suppose the price function p is continuous and that a buyer arriving at time t maximizes her utility by buying at time $s > t$ (i.e., the buyer chooses s deterministically). Then, since all buyers have equal valuation for the item, it is clear that every buyer arriving between t and s will also maximize her utility by buying at time s. But since there is a positive probability that someone arrives in $[t, s]$ and p is continuous, a sufficient small ε exists such that $s - \varepsilon$ is a better option for the buyer arriving at time t, and therefore buying at s is not a best response for this buyer.

2.4. Equilibrium of the First Stage

Given that the seller does not discount the future, it is quite evident that his optimal strategy belongs to one of the following two families:

Constant pricing: Select the price function $p(t) \equiv V$ (i.e., a no markdown strategy).

Markdown: Select a price function such that $p(t) > v$ for all $t \in [0, T)$ and $p(T) = v$.

Indeed, if $p(t) > v$ throughout the season, low-valuation buyers will not buy, and thus setting the price to V is optimal. Otherwise, the price function will reach the value v at some point, but because the seller does not discount the future, it is in his best interest to mark down at the end since low-valuation buyers will buy anyhow. In the former case, there is no strategic waiting: high-value buyers buy upon arrival, whereas low-value buyers never buy. Thus, the expected revenue of the seller is easily computed as

$$\pi(V) = V\,\mathbb{P}(\text{one or more high-value buyers arrive}) = V(1 - Q_0(T)). \tag{3}$$

To express the expected revenue in the latter case observe that a price function induces an equilibrium, which determines an availability probability function $f(t)$. Thus, for $t < T$, $G(t) = 1 - f(t)$ can be seen as the distribution of the random variable expressing the time at which the item is sold considering only high-value buyers, and for convenience we set $G(T) = 1 - Q_0(T)$. This holds since low-value buyers only buy at time T. Thus, the seller's expected revenue under the markdown pricing scheme can be expressed as the integral of the price function with respect to the measure induced by G plus a term accounting for the probability of selling the item at price v when no high-value buyer arrives:

$$\pi(p, G) := \int_0^T p(t)\, dG(t) + v Q_0(T)(1 - q_0(T)). \tag{4}$$

Note that in this expression the revenue corresponding to selling the item at price v is exactly v times the probability of selling the item at that price. Indeed, this probability is expressed as $Q_0(T)(1 - q_0(T))$, in the second term of (4), plus the jump $G(T) - G(T^-)$ in the integral term (since $p(T) = v$). Overall, this adds up to $1 - Q_0(T) - G(T^-) + Q_0(T) - Q_0(T) q_0(T) = 1 - Q_0(T) q_0(T) - G(T^-)$, i.e., the probability of selling the item minus the probability of not selling it before T, as desired.

In light of this observation, to determine the optimal pricing strategy, the seller needs to solve first the subproblem in which there are only buyers with valuation V and the price function has to be chosen among those belonging to the markdown family, i.e., satisfying $p(t) > v$ for all $t \in [0, T)$ and $p(T) = v$. In this subproblem, buyers with valuation v only interfere as a threat to buyers of high valuation if these decide to postpone their purchase until time T. In summary, the seller needs to find a price function maximizing $\tilde{\pi}(p, G) = \int_0^T p(t)\, dG(t)$. With this, the seller evaluates Equation (4), then compares the quantities (3) and (4), and selects the price function inducing the largest profit among these two.

3. Optimal Pricing

In this section we explicitly obtain the optimal price function and its corresponding equilibrium for the seller. To this end, we consider that only high-value buyers arrive according to a nonhomogeneous Poisson process. In Section 3.3 we plug this back into the seller's revenue to obtain the optimal profit. The space of functions over which the seller needs to make her selection, called $\mathcal{F}$, is the set of price functions $p\colon [0,T] \to [v,V]$ such that $p(T)=v$ and $p(t)>v$ for every $t \in [0,T)$. Note that we do not impose any continuity or regularity of the price functions; we only require that an equilibrium for the buyers exists. By defining $\mathcal{D}$ as the set of non decreasing functions $G\colon [0,T] \to [0,1]$ such that $G(0)=0$ and setting

$$E = \{(p,G) \in \mathcal{F} \times \mathcal{D}\colon 1-G \text{ is an equilibrium for } p\},$$

the vendor's subproblem may be written as

$$\max_{(p,G)\in E} \tilde{\pi}(p,G) = \int_0^T p(t)\,dG(t).$$

To solve this problem, we first compute an upper bound on the expected revenue of any price function in $\mathcal{F}$ and then find a particular pricing scheme whose expected revenue matches this upper bound. The tightness of this upper bound is based on two guesses that hold at the optimal price path: first, that high-valuation buyers may buy as long as the item is available, and second, that these buyers are indifferent between different purchasing times. Thus, our upper bounds essentially comes by lower bounding the utility of a buyer with what she would get if she buys at the end of the season. This translates into an upper bound on the seller's profit that we may match with a price function satisfying the guesses.

3.1. Upper Bound on Vendor's Profit

Recall that given a function $p \in \mathcal{F}$, buyers choose a distribution function for deciding when they will buy so as to maximize their utility $U(s) = \alpha(s)(V - p(s)) \cdot f(s)e^{-\mu s}$, where f is an equilibrium for p. So we let $(p,G) \in E$ and for convenience let $f = 1-G$. Thus, a buyer arriving at time t will get an expected utility of $u_t = \max_{s\geq t} U(s)$ and may buy at any time belonging to the set $S_t := \arg\max_{s\in[t,T]} U(s)$. Note that the value u_t is nonincreasing in t. In particular,

$$u_t \geq Q_0(T)\alpha(T)e^{-\mu T}(V-v) \quad \text{for all } t \in [0,T]. \tag{5}$$

Indeed, the quantity on the right-hand side is the expected utility the buyer would get if no other high-value buyer arrives in $[0,T]$ and she gets the item in the random assignment at time T, so $\alpha(T) = \sum_{k=0}^{\infty} q_k(T)/(1+k)$. Let us define $S = \bigcup_{t\in[0,T]} S_t$, the set of times at which the item can be sold. Now, it follows from (5) that for all $t \in S$, i.e., a time at which there is interest in buying, we have that

$$U(t) = \alpha(t)f(t)e^{-\mu t}(V-p(t)) \geq Q_0(T)\alpha(T)e^{-\mu T}(V-v).$$

Using that $\alpha(t) \leq 1$, since it is a probability, we deduce

$$\begin{aligned} p(t) &\leq V - \frac{Q_0(T)\alpha(T)}{f(t)} e^{-\mu(T-t)}(V-v) \\ &= V - \frac{Q_0(T)\alpha(T)}{1-G(t)} e^{-\mu(T-t)}(V-v) \quad \text{for all } t \in S. \end{aligned}$$

Hence, we have that

$$\begin{aligned} &\int_0^T p(t)\,dG(t) \\ &\quad = \int_S p(t)\,dG(t) \\ &\quad \leq \int_S \left(V - \frac{Q_0(T)\alpha(T)}{1-G(t)} e^{-\mu(T-t)}(V-v)\right) dG(t) \\ &\quad \leq \int_0^T \left(V - \frac{Q_0(T)\alpha(T)}{1-G(t)} e^{-\mu(T-t)}(V-v)\right) dG(t) \\ &\quad = V\int_0^T dG(t) \\ &\qquad + Q_0(T)\alpha(T)(V-v)e^{-\mu T}\int_0^T e^{\mu t}\frac{-dG(t)}{1-G(t)}. \end{aligned} \tag{6}$$

Here, the first equality follows from the fact that outside S the function $G(t)$ is constant, since no buyer buys outside S. (The support of the measure induced by G is exactly S.) The first inequality follows from the bound on $p(t)$ obtained earlier, whereas the second inequality is direct because $S \subseteq [0,T]$ and the integrand is nonnegative. Hence, noting that $\int_0^T dG(t) = G(T) - G(0) = 1 - Q_0(T)$ for the first integral and using integration by parts for the second, we obtain that, for all $(p,G) \in E$,

$$\begin{aligned} &\tilde{\pi}(p,G) \\ &\quad \leq V(1-Q_0(T)) + Q_0(T)\alpha(T)(V-v) \\ &\qquad \cdot \left(\ln(Q_0(T)) - e^{-\mu T}\mu\int_0^T e^{\mu t}\ln(1-G(t))\,dt\right). \end{aligned} \tag{7}$$

3.2. Matching Upper Bound

Now we provide a pricing scheme, together with a very natural equilibrium for it, that attains the latter upper bound for the revenue of the vendor. We then conclude that the seller's revenue under this price function is best possible for the subproblem in which only high-value consumers arrive. Define

$$p^*(t) = \begin{cases} V - \dfrac{Q_0(T)\alpha(T)}{Q_0(t)} e^{-\mu(T-t)}(V-v) & \text{for } t \in [0,T), \\ v & \text{for } t = T, \end{cases}$$

and observe that, for every $t \in [0,T)$, $p^*(t) > v$ and $p^*(T) = v$. Note also that this function is discontinuous at time T. Under this pricing scheme, a buyer arriving at time t will seek to maximize

$$\max_{s\in[t,T)} f(s)\alpha(s)e^{-\mu s}\frac{Q_0(T)\alpha(T)}{Q_0(s)} e^{-\mu(T-s)}(V-v),$$

or may prefer to buy at time T for a profit of $f(T)\alpha(T)e^{-\mu T}(V-v)$. Then, if all buyers buy upon arrival $Q_0(s)=f(s)$ and $\alpha(s)=1$, the previous quantity is actually the constant $Q_0(T)\alpha(T)e^{-\mu T}(V-v)$, which in particular is maximized at time $s=t$. Therefore, $f(s)=Q_0(s)$ is an equilibrium for the proposed pricing p^*. By denoting $G^*=1-Q_0$, we conclude that $(p^*,G^*)\in E$.

Observe that by subtracting an arbitrarily small quantity $\varepsilon>0$ to the price function above, we can force that the utility of each buyer is strictly maximized at time $s=t$. Indeed this extra term ε will be multiplied by $e^{-\mu t}Q_0(t)$, so the expected utility of a buyer deciding to buy at time s will be the constant above plus $\varepsilon e^{-\mu s}Q_0(s)$, which is strictly decreasing. This implies that buying upon arrival (i.e., $f(t)=Q_0(t)$), is an equilibrium in which all buyers strictly prefer their choice. Of course, because ε is arbitrarily small, this price function achieves a profit for the seller that is arbitrarily close to that obtained with p^*.

Now we compute the expected revenue for the seller of the pricing policy p^*. This revenue calculation follows exactly as in (6) and doing integration by parts after this. It follows that the revenue of the seller is given by

$$\begin{aligned}\tilde{\pi}(p^*,G^*)&\\ &=V(1-Q_0(T))+Q_0(T)\alpha(T)(V-v)\\ &\quad\cdot\left(\ln(Q_0(T))-e^{-\mu T}\mu\int_0^T e^{\mu t}\ln(1-G^*(t))\,dt\right). \end{aligned}\tag{8}$$

Finally, to connect the quantities in (7) and (8), we need to establish a relation between the last integrals in both terms. Note that since for every $(p,G)\in E$ we have that $1-G\geq Q_0$, we can bound any such distribution as $G\leq G^*$. Therefore, using the monotonicity of the function $\ln(\cdot)$, we deduce from (7) that for all equilibrium $(p,G)\in E$,

$$\begin{aligned}\tilde{\pi}(p,G)&\leq V(1-Q_0(T))+Q_0(T)\alpha(T)(V-v)\\ &\quad\cdot\left(\ln(Q_0(T))-e^{-\mu T}\mu\int_0^T e^{\mu t}\ln(1-G(t))\,dt\right).\\ &\leq V(1-Q_0(T))+Q_0(T)\alpha(T)(V-v)\\ &\quad\cdot\left(\ln(Q_0(T))-e^{-\mu T}\mu\int_0^T e^{\mu t}\ln(1-G^*(t))\,dt\right).\\ &=\tilde{\pi}(p^*,G^*).\end{aligned}$$

Hence, (p^*,G^*) is the optimal pricing function and associated equilibrium.

3.3. Best Pricing Strategy

We are now ready to compare the revenues obtained by the candidates to the best pricing policy. Recall that this is either *constant pricing* or the best possible *markdown*. To make the calculation explicit, note that we are assuming nonhomogeneous Poisson arrivals so that if we let $m(t)=\int_0^t\Lambda(s)\,ds$ and $\ell(t)=\int_0^t\lambda(s)\,ds$, we have $Q_0(T)=e^{-m(T)}$, $Q_1(T)=m(T)e^{-m(T)}$, $q_0(T)=e^{-\ell(T)}$, $q_1(T)=\ell(T)e^{-\ell(T)}$, and $\alpha(T)=\sum_{k=0}^{\infty}q_k(T)/(1+k)=(1-e^{-\ell(T)})\ell(T)$.

Constant pricing: Observe that in the constant price strategy, the seller obtains value V if and only if at least one high-value buyer arrives. Thus, his expected revenue is

$$V(1-Q_0(T))=V(1-e^{-m(T)}).\tag{9}$$

Markdown: To evaluate the revenue of the optimal markdown pricing strategy, we use Equations (4) and (8), which lead to an expected revenue of

$$\begin{aligned}&V(1-Q_0(T))+Q_0(T)\alpha(T)(V-v)\\ &\quad\cdot\left(\ln(Q_0(T))-e^{-\mu T}\mu\int_0^T e^{\mu t}\ln(1-G^*(t))\,dt\right)\\ &\quad+vQ_0(T)(1-q_0(T))\\ &=V(1-Q_0(T))+e^{-m(T)}(1-e^{-\ell(T)})\\ &\quad\cdot\left(v-\frac{V-v}{\ell(T)}\left(m(T)-e^{-\mu T}\mu\int_0^T e^{\mu t}m(t)\,dt\right)\right).\end{aligned}\tag{10}$$

In summary, the markdown strategy is better if and only if

$$\begin{aligned}&v\left(\ell(T)+m(T)-\mu\int_0^T e^{-\mu(T-t)}m(t)\,dt\right)\\ &\qquad>V\left(m(T)-\mu\int_0^T e^{-\mu(T-t)}m(t)\,dt\right).\end{aligned}$$

This condition may be rewritten as

$$\frac{V}{v}<\left(1+\frac{\ell(T)}{m(T)-\mu\int_0^T e^{-\mu(T-t)}m(t)\,dt}\right).\tag{11}$$

Observe that, as one may intuitively expect, this condition is invariant under a rescaling of the valuations. What is probably less intuitive is that this condition is also invariant under a rescaling of the arrival rates. This happens since, under such a rescaling, the price path p^* gets closer to the value V, because the threat of low-valuation buyers becomes more powerful.

Interestingly, in the situation without a discount rate, i.e., $\mu=0$, our price path obtains the same revenue as an optimal mechanism where all bidders arrive at time 0. This constitutes a stronger version of revenue equivalence that applies to a random number of bidders (similar to the results of Levin and Ozdenoren 2004 and Haviv and Milchtaich 2012). Indeed, as pointed out by Skreta (2006), one can argue that the optimal auction in our setting can be derived as follows. Consider that a random number of bidders, distributed as a Poisson random variable of parameter $m(T)+\ell(T)$, participate in the auction. The valuation of the bidders are independent and identically distributed (i.i.d.), equal to v

with probability $p=\ell(T)/(m(T)+\ell(T))$, and equal to V with probability $1-p$. In this situation, the virtual valuation is not monotone, and therefore Myerson's ironing is needed. It turns out that the optimal mechanism can be implemented as follows:

Case $v>V(1-p)$. Use a fixed price equal to V. Note that the condition of this case is the opposite of (11) (for $\mu=0$) and, furthermore, the revenue is exactly that given by (9).

Case $v<V(1-p)$. Here the situation is more involved. The optimal mechanism is implemented by running a second price auction with reservation price equal to $b=V-(V-v)((1-e^{-\ell(T)})/\ell(T))$, but if all bids are below b, the item is randomly allocated to any of the participating bidders. The auction is incentive compatible, and thus the seller's revenue in this case equals to V times the probability of having two or more high-valuation bidders plus b times the probability of having exactly one high-valuation bidder plus v times the probability of having no high-valuation bidder and at least one with low valuation; that is,

$$\begin{aligned}&V(1-e^{-m(T)}-m(T)e^{-m(T)})+bm(T)e^{-m(T)}\\&\quad+ve^{-m(T)}(1-e^{-\ell(T)})\\&=V(1-e^{-m(T)})+e^{-m(T)}(1-e^{-\ell(T)})\\&\quad\cdot\left(v-(V-v)\frac{m(T)}{\ell(T)}\right),\end{aligned}$$

exactly as is (10).

Therefore, the nonnegative term $((1-e^{-\ell(T)})/\ell(T))\cdot e^{-m(T)}(V-v)e^{-\mu T}\mu\int_0^T e^{\mu t}m(t)\,dt$ may be seen as the additional revenue obtained by the seller by exploiting the impatience of the consumers. It is worth mentioning that our pricing scheme can get up to twice the revenue of the optimal static mechanism and not more than that. To see this, note that the best case for our pricing occurs when μ is very large. In this case, p^* is essentially a constant equal to V and drops to v in the very last minute, implying a revenue of $V(1-e^{-m(T)})+ve^{-m(T)}(1-e^{-\ell(T)})$. On the other hand, the optimal mechanism gets the revenue expressed in the previous cases. A tedious but straightforward calculation shows that this ratio is at most 2, and this can be attained (in the limit) using, for instance, $V=1, v=\sqrt{\varepsilon}$, $m(T)=\varepsilon$, $\ell(T)=\sqrt{\varepsilon}$. Indeed, in this case, the revenue of the optimal mechanism is $1-e^{-\varepsilon}\approx\varepsilon$, whereas our pricing mechanism obtains $1-e^{-\varepsilon}+e^{-\varepsilon}(1-e^{-\sqrt{\varepsilon}})\sqrt{\varepsilon}\approx 2\varepsilon$.

To finish this section, we plot in Figure 1 the optimal pricing path p^* when the rate of the arrival process is constant. Interestingly, the price remains relatively constant until close to the end of the season, where it drops steeply. This seems consistent with the common practice in retail and other industries where aggressive markdown strategies are used. In our setting, the optimal price is even discontinuous at time T. This is partially because of the valuations we consider; however, the phenomenon of steep discounts is prevalent in continuous distributions of valuations, as we show in the next section.

Figure 1. Optimal Markdown Price Function for Homogeneous Poisson Processes of Rates $\Lambda=1$ and $\lambda=0.2$ and Parameters $V=3$, $v=1$, $\mu=0.5$, and $T=5$

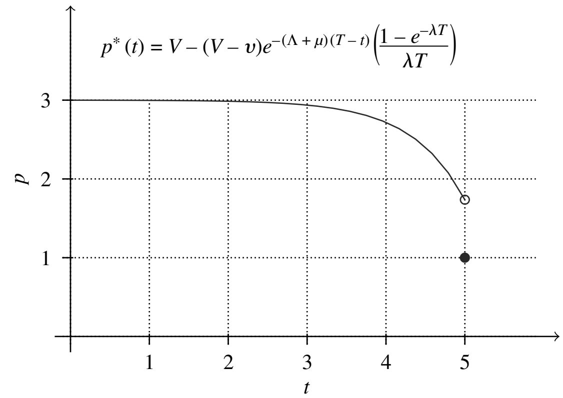

4. Continuous Valuation

In this section we relax the assumption of having just two possible valuations for the item and consider the general case in which the buyers' valuations for the item are i.i.d. according to a continuous distribution $\Phi\colon[v,V]\to[0,1]$ with associated density ϕ. Therefore, we assume that buyers arrive according to a nonhomogeneous Poisson process of rate $\Lambda(\cdot)$, and the probability that a buyer arriving at time t has a valuation less than or equal to v is $\Phi(v)$.

The main result of this section is that, under an assumption over the equilibrium strategies, we reduce the seller's problem to solving a system of ordinary differential equations. To this end, we first show that we can reduce to equilibria taking the form of a threshold $\varphi(\cdot)$, implying that a buyer arriving a time t with valuation u will buy upon arrival if $u>\varphi(t)$, will buy at time $s\in\varphi^{-1}(u)$ (with $s\geq t$) if $u\leq\varphi(t)$, and will not buy if such a time does not exist. Second, we prove that a first-order approach is sufficient to write the seller's optimization problem. Finally, using optimal control theory, we write down the seller's problem in a way that can be dealt with numerically.

4.1. Threshold Strategies

Following the notation in Section 2, given a price function p, we let f be a corresponding equilibrium with associated distributions $H=(H_t^u)_{t\in[0,T]}^{u\in[v,V]}$. Here, H_t^u is a probability distribution over $[t,T]$ corresponding to the (mixed) strategy of a buyer arriving at time t with valuation u. Observe that, for a continuous price function, the random allocation probability $\alpha(t)$ will always be 1, and the probability that the item is available at time t, $f(t)$, is a continuous function. Thus, we have that $\mathrm{supp}(H_t^u) \subseteq \arg\max_{t\le s\le T}(u-p(s))e^{-\mu s}f(s)$. A key *monotonicity property* is that if we consider an equilibrium and two valuations $u < u'$, then for all $s \in \mathrm{supp}(H_t^u)$ with $s > t$ and $s' \in \mathrm{supp}(H_t^{u'})$, we have that $s > s'$. Indeed, assume s is the smallest element in $\mathrm{supp}(H_t^u)$, and write the utility of a buyer with utility u' buying at time τ as

$$\begin{aligned}(u'-p(\tau))&e^{-\mu\tau}f(\tau)\\ &=(u-p(\tau))e^{-\mu\tau}f(\tau)+(u'-u)e^{-\mu\tau}f(\tau). \quad (12)\end{aligned}$$

Since the second term is decreasing and the first is maximized at $\tau = s$, the whole utility is maximized at a point $s' \in [0, s)$.[2] This monotonicity property implies that if a buyer arriving at t buys upon arrival, then any higher-valuation buyer arriving at t will also buy upon arrival.

With the monotonicity property at hand, it follows that for an equilibrium (p, f) with associated distributions H, the function $\varphi(t) = \inf\{u \mid t \in \mathrm{supp}(H_t^u)\}$[3] defines a threshold with the desired property. Indeed, note that φ may be defined as

$$\varphi(t)=\inf\left\{u \mid t\in \arg\max_{t\le s\le T}(u-p(s))e^{-\mu s}f(s)\right\},$$

so that if for a buyer arriving at t her valuation u is greater than $\varphi(t)$, then she will buy immediately, whereas if her valuation is $u \le \varphi(t)$, she will wait until a time s for which $u = \varphi(s)$.

As the reader could realize, the monotonicity property imposes a certain order in the equilibrium strategies, which is summarized by the threshold characterization. From a mechanism design perspective, the threshold function is inherently connected to the allocation rule associated with the mechanism of a posted price p. In fact, if the threshold function turns out to be nonincreasing, the allocation rule consists of giving the item to the player with minimum $\tau = \max\{t, \varphi^{-1}(u)\}$, where (t, u) is the respective type of the player.

Moreover, a violation of this nonincreasing threshold property would imply scenarios (with positive probability of occurrence) where a player u arriving at t receives the item, but if she arrives shortly afterward she waits to purchase at time $t + c$. Thus, the chances of obtaining the item depend on whether no one with higher valuation arrives between $[t, t + c]$. This strange situation leads us to conjecture that in the optimal pricing, the threshold must be nonincreasing. Unfortunately, we have been unable to formally prove the latter. Nevertheless, under some conditions, e.g., if the sets $S_u := \arg\max_{s\ge 0}(u-p(s))e^{-\mu s}f(s)$ are connected, one can indeed prove that the threshold equilibrium is nonincreasing. For the rest of this section, we make the following assumption.

Assumption 1. *In the revenue maximizing pricing policy, the buyers' equilibrium is characterized by a nonincreasing threshold function.*

Thus, for sorting out the seller's problem, we restrict our attention to this class of equilibria. In what follows, we exploit this conjecture to simplify the seller's optimization problem. An important implication of our assumption is that if φ is nonincreasing, there are at most countably many valuations u for which the preimage $\varphi^{-1}(u)$ is not a singleton. Therefore, for almost all valuations u and arrival times t, a buyer with valuation u and arriving at time t will buy at time $\max\{t, \varphi^{-1}(u)\}$. In conclusion, almost all players are playing pure strategies.

4.2. First-Order Approach

We now consider a price function with a nonincreasing threshold equilibrium, which we denote by (p, φ), and assume that both p and φ are differentiable. The goal of this section is to show that the first-order approach is sufficient to deal with the seller's problem.

Let us first write down the seller's problem. Recalling that $m(t) = \int_0^t \Lambda(s)\,ds$ is the average arrival rate until time t, denoting by E the set of equilibrium pairs[4] (p, φ) such that φ is nonincreasing, and noting that $f(t) = e^{-m(t)(1-\Phi(\varphi(t)))}$ expresses the probability of having no arrivals in $[0, t]$ with valuation above $\varphi(t)$, the seller's problem can be written as

$$\max_{p:(p,\varphi)\in E}\int_0^T -\frac{d}{ds}\left(e^{-m(s)(1-\Phi(\varphi(s)))}\right)p(s)\,ds. \quad (13)$$

Let $U(u,t) := (u-p(t))e^{-\mu t}f(t) = (u-p(t))e^{-\mu t}\cdot e^{-m(t)(1-\Phi(\varphi(t)))}$ denote the utility (at equilibrium) of a buyer with valuation u buying at time t. Then, $(p, \varphi) \in E$ if and only if $U(u, t)$ is maximized at $t = \varphi^{-1}(u)$. On the other hand, the first-order optimality condition is $\partial_2 U(u, \varphi^{-1}(u)) = 0$. Similarly to the monotonicity property, we compute

$$\begin{aligned}\partial_2U(u,t)&=-p'(t)e^{-\mu t}f(t)+p(t)\mu e^{-\mu t}f(t)-p(t)e^{-\mu t}f'(t)\\ &\quad -ue^{-\mu t}(\mu f(t)-f'(t)) \quad (14)\end{aligned}$$

and note that, evaluating at $t = \varphi^{-1}(u)$, the previous quantity is zero. Then, since the term $\mu f(t) - f'(t) > 0$, $\partial_2 U(u', \varphi^{-1}(u)) < 0$ whenever $u' > u$, and $\partial_2 U(u', \varphi^{-1}(u)) > 0$ whenever $u' < u$. Equivalently, since φ is nonincreasing, we have that

$$\partial_2U(u,t)\begin{cases}>0 & \text{if } t<\varphi^{-1}(u),\\ =0 & \text{if } t=\varphi^{-1}(u),\\ <0 & \text{if } t>\varphi^{-1}(u).\end{cases}$$

Thus, the first-order optimality condition is enough to guarantee that, as a function of t, $U(u, t)$ increases until $t = \varphi^{-1}(u)$, and then it decreases, implying that

$\varphi^{-1}(u)$ is a global maximizer. We have thus established that (13) is equivalent to

$$\max_{p,\varphi:\partial_2 U(u,\varphi^{-1}(u))=0,\varphi'\le 0} \int_0^T -\frac{d}{ds}(e^{-m(s)(1-\Phi(\varphi(s)))})p(s)\,ds. \tag{15}$$

4.3. Solving the Seller's Problem Using Optimal Control

Under Assumption 1, we now apply optimal control theory to solve problem (15). It is worth mentioning that although we must restrict our attention to a differentiable setting, smooth functions are dense in the continuous functions space, so our method obtains numerical solution to the seller's problem in this more general setting.

To transform (15) into the classic optimal control framework, we observe that at an equilibrium (p,φ) we must have that $p(T)=\varphi(T)$. Then, writing f explicitly in (14), we have that (15) becomes

$$\max_{\alpha,p,\varphi} \int_0^T -\frac{d}{ds}(e^{-m(s)(1-\Phi(\varphi(s)))})p(s)\,ds$$
$$\text{s.t.}\ \begin{cases} -p'=(\varphi-p)(\mu+m'(s)(1-\Phi(\varphi)) \\ \qquad -m(s)\phi(\varphi(s))\varphi'(s)) \quad \text{for } s\in(0,T), \\ p(T)=\varphi(T)=\alpha, \\ \varphi'\le 0. \end{cases}$$

To solve this problem, we introduce the auxiliary functions $q(s)=p(T-s)$ and $\psi(s)=\varphi(T-s)$, for every $s\in[0,T]$. By using the change of variables $\tau=T-s$, the problem becomes

$$\max_{\alpha,q,\psi} I:=\int_0^T \frac{d}{d\tau}(e^{-m(T-\tau)(1-\Phi(\psi(\tau)))})q(\tau)\,d\tau$$
$$\text{s.t.}\ \begin{cases} q'=(\psi-q)(\mu+m'(T-\tau)(1-\Phi(\psi)) \\ \qquad +m(T-\tau)\phi(\psi)\psi') \quad \text{for } \tau\in(0,T), \\ q(0)=\psi(0)=\alpha, \\ \psi'\ge 0. \end{cases} \tag{16}$$

Note that from the differential equation on q, we obtain

$$q(\tau)=\psi(\tau)-e^{\mu(T-\tau)+m(T-\tau)(1-\Phi(\psi(\tau)))} \cdot\int_0^\tau \psi'(s)e^{-\mu(T-s)-m(T-s)(1-\Phi(\psi(s)))}\,ds.$$

Using integration by parts, this yields

$$\begin{aligned} I &= \int_0^T \frac{d}{d\tau}(e^{-m(T-\tau)(1-\Phi(\psi(\tau)))})\psi(\tau)\,d\tau \\ &\quad+\int_0^T \frac{d}{d\tau}(m(T-\tau)(1-\Phi(\psi(\tau))))e^{\mu(T-\tau)} \\ &\quad\cdot\int_0^\tau \psi'(s)e^{-\mu(T-s)-m(T-s)(1-\Phi(\psi(s)))}\,ds\,d\tau \\ &= \psi(T)-\alpha e^{-m(T)(1-\Phi(\alpha))} \\ &\quad-\int_0^T e^{-m(T-\tau)(1-\Phi(\psi(\tau)))}\psi'(\tau)\,d\tau \\ &\quad-\int_0^T m(T-\tau)(1-\Phi(\psi(\tau)))e^{\mu(T-\tau)} \\ &\quad\cdot(\psi'(\tau)e^{-\mu(T-\tau)-m(T-\tau)(1-\Phi(\psi(\tau)))}-\mu r(\tau))\,d\tau, \end{aligned}$$

where r is the solution to the ordinary differential equation

$$\begin{cases} r'=\psi' e^{-\mu(T-\tau)-m(T-\tau)(1-\Phi(\psi))} & \text{for } \tau\in(0,T), \\ r(0)=0. \end{cases}$$

Therefore, by setting $u=\psi'$, (16) is equivalent to

$$\min_{\alpha,q,\psi} \int_0^T \ell(\tau,\psi(\tau),r(\tau),u(\tau))d\tau+R(\psi(T),r(T))$$
$$\text{s.t.}\ \begin{cases} r'=ue^{-\mu(T-\tau)-m(T-\tau)(1-\Phi(\psi))} & \text{for } \tau\in(0,T), \\ r(0)=0, \\ \psi'=u\ge 0 & \text{for } \tau\in(0,T), \\ \psi(0)=\alpha, \end{cases} \tag{17}$$

where

$$\begin{cases} \ell(\tau,\psi,r,u):=e^{-m(T-\tau)(1-\Phi(\psi))}u(1+m(T-\tau)(1-\Phi(\psi))) \\ \qquad\qquad -\mu r m(T-\tau)(1-\Phi(\psi))e^{\mu(T-\tau)}, \\ R(\psi,r):=\alpha e^{-m(T)(1-\Phi(\alpha))}-\psi. \end{cases}$$

Note that (17) is a classical optimal control problem, where u is the control and (ψ,r) is the state. Hence, we deduce that any solution of this problem must satisfy the first-order necessary conditions (Pontryagin's minimum principle; see, e.g., Vinter 2000, Section 6.2):

$$(\forall\tau\in\,]0,T[)$$
$$\cdot\begin{cases} u(\tau)\in\arg\min_{u\in\mathbb{R}_+} H(\tau,\psi(\tau),r(\tau),u,w(\tau),\eta(\tau)), \\ -w'(\tau)=\dfrac{\partial H}{\partial r}(\tau,\psi(\tau),r(\tau),u(\tau),w(\tau),\eta(\tau)), \\ w(T)=\dfrac{\partial R}{\partial r}(\psi(T),r(T))=0, \\ -\eta'(\tau)=\dfrac{\partial H}{\partial \psi}(\tau,\psi(\tau),r(\tau),u(\tau),w(\tau),\eta(\tau)), \\ \eta(T)=\dfrac{\partial R}{\partial \psi}(\psi(T),r(T))=-1, \end{cases} \tag{18}$$

where H is the Hamiltonian of the system

$$\begin{aligned} &H(\tau,\psi,r,u,w,\eta) \\ &\quad=\ell(\tau,\psi,r,u)+wue^{-\mu(T-\tau)-m(T-\tau)(1-\Phi(\psi))}+\eta u. \end{aligned} \tag{19}$$

Thus, we have transformed the seller's problem to solving (17)–(18), a system of four ordinary differential equations with initial value, coupled with a Hamiltonian equation.

Figure 2. Optimal Price and Threshold for Homogeneous Poisson Arrivals of Rate $\lambda = 3$

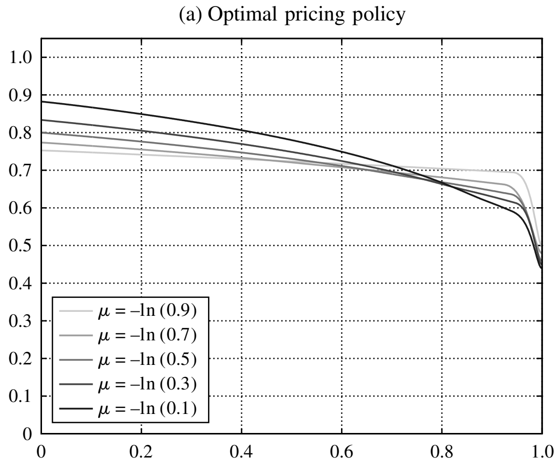

4.4. Numerical Experiments

For solving system (17)–(18) numerically, we discretize the interval $[0, T]$ into $N_h = T/h$ subintervals of length h, starting from a given piecewise linear control function u_h^0. For every $n \in \mathbb{N}$, we find piecewise linear functions r_h^n and ψ_h^n by solving the differential equations in (17) via a forward Euler's method, and we find w_h^n and η_h^n by solving the differential equations in (18) via a backward Euler's method. Then we obtain u_h^{n+1} by computing the projected gradient step

$$u_h^{n+1}(\tau_i) = P_{\mathbb{R}_+}\Bigg(u_h^n(\tau_i) - \gamma \frac{\partial H}{\partial u}(\tau_i, \psi_h^n(\tau_i), r_h^n(\tau_i), u_h^n(\tau_i), w_h^n(\tau_i), \eta_h^n(\tau_i))\Bigg), \quad (20)$$

where, for $i = 0, \ldots, N_h$, we let $\tau_i = i \cdot h$, the parameter $\gamma > 0$ is chosen appropriately, and $P_{\mathbb{R}_+}(x) = \max\{0, x\}$. The algorithm stops when $\max_{0 \le i \le N_h} |u_h^{n+1}(\tau_i) - u_h^n(\tau_i)| < \varepsilon$, for $\varepsilon > 0$ small enough. All our computations consider that Φ is the uniform distribution in $[0, 1]$ and $T = 1$. The parameters are set to be $\varepsilon = 0.005$, $h = 0.001$, and $\gamma = 0.3$.

In Figure 2, we vary the discount rate, whereas the arrivals are modeled via an homogeneous Poisson process of fixed rate $\lambda = 3$. On the other hand, in Figure 3, we fix $\mu = -\ln(0.7)$ and vary the arrival rate of the buyers. Tables 1 and 2 exhibit the profits obtained in each case and compare them to those of an optimal auction (that takes place at the end of the season). We verify that the profit obtained with the markdown strategy is always better than that of the optimal auction and that, quite naturally, the difference increases when μ increases. Maybe not so naturally, note that when μ is large, it is more profitable to decrease the reserve price. The situation with fixed discount rate and varying λ is different. There, it is not clear that a higher arrival rate impacts the profit ratio. It is natural, however, that the reservation price is not significantly affected by the number of buyers as this is also the case in an optimal auction.

Additionally, the numerical results show that in the continuous valuation model, the reservation price $p(T)$ is affected by the temporal discount rate, in contrast to

Figure 3. Optimal Price and Threshold Functions for a Discount Factor $\mu = -\ln(0.7)$

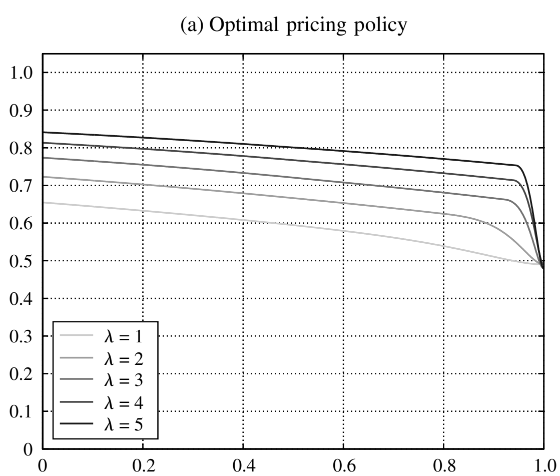

Table 1. Optimal Profits When $\lambda = 3$ Compared with the Revenue of an Optimal Auction at the End of the Season

μ	Profit	α	Opt. auction	Ratio
$-\ln(0.9) = 0.105$	0.4877	0.50	0.4821	1.011
$-\ln(0.7) = 0.357$	0.4900	0.48	0.4821	1.016
$-\ln(0.5) = 0.693$	0.4973	0.46	0.4821	1.031
$-\ln(0.3) = 1.204$	0.5072	0.45	0.4821	1.052
$-\ln(0.1) = 2.303$	0.5230	0.44	0.4821	1.085

Table 2. Optimal Profits When $\mu = -\ln(0.7)$ Compared with the Revenue of an Optimal Auction at the End of the Season

λ	Profit	α	Opt. auction	Ratio
1	0.2166	0.49	0.2131	1.016
2	0.3733	0.49	0.3679	1.015
3	0.4900	0.48	0.4821	1.016
4	0.5777	0.49	0.5677	1.018
5	0.6450	0.48	0.6328	1.019

the discrete valuation model, where we proved that the reservation price is unaffected by the discount factor. The intuition behind this is that in the discrete valuation model, the reservation price is just used to avoid that high-valuation buyers decide to go to the lottery at the end the season. In contrast, in the continuous setting, the reservation price is used to split the bidders that are ex ante interesting (for the seller) to trade with. Since the valuation is affected by time with the discount rate, it is quite natural that here the reservation price depends on the discount factor.

Another interesting observation is that in our numerical simulations, as in Section 3 when $\mu = 0$, we recover the classic result of static optimal mechanism design when we do not consider a temporal discount factor. Indeed, as μ approaches 0, the optimal price function becomes flat and decreases very quickly to 0.5 at time 1. This is a way of simulating a first price auction at time 1 with a reserve price of 0.5, which is known to yield the optimal revenue. For instance, taking $\lambda = 3$ and $\mu = -\ln(0.99)$, we obtain that the optimal price is essentially 0.740 until time 0.95 (with an underlying threshold of value essentially 0.99). This is in almost perfect agreement with the fact that in a first price auction with a random number of bidders Poisson distributed at rate 3, the maximum possible bid, i.e., that of a bidder with valuation 1, is equal to 0.741.

5. Concluding Remarks

We have studied a two-stage dynamic game where, in the first stage, a seller proposes a markdown path price for selling a single item, and strategic consumers respond by selecting the optimal time to buy the item considering the risk of not getting it.

In particular, we have characterized the optimal price function when buyers' valuations can only take two values. Interestingly, this function satisfies an important economic property: it is incentive compatible. Indeed, even if the seller cannot observe the buyers arrival time (a common situation in practice), it is in the buyers' best interest to buy upon arrival, thus revealing their private type. Furthermore, the revenue obtained by this price function is at least as large as that of the optimal mechanism in this context. In this respect, the obtained optimal price function is discontinuous at the end of the season, which nicely mimics the implicit random allocation necessary in the optimal mechanism.

We also derive a numerical approach to tackle the general valuation case. Our numerical results show in particular that the fact that buyers discount the future faster than the seller severely affects the optimal pricing policy. Indeed, one can easily derive from the work of Gershkov et al. (2014) that if the seller and the buyers discount the future equally, then the optimal mechanism takes the form of a threshold that is constant until the end of the season, the time at which an auction is run. In our case the threshold is far from constant for small arrival rates or large discount rates.

Throughout this paper, we have assumed that the seller has commitment power and can credibly preannounce a certain price path. However, because we deal with the single unit case, this assumption is not really needed. Indeed, when there is a single unit on sale, the optimal preannounced price function and the optimal *dynamic* price function (in which the seller does not make commitments) actually coincide. Therefore, all our results apply to the case without commitment as well.

Finally, important extensions that require further investigation are to characterize the optimal path price function in more general frameworks, including when consumers have a random private value over a continuous distribution and when there are multiple units to be sold.

Acknowledgments

The authors thank Gustavo Vulcano for his valuable comments on an earlier draft of this paper. They also thank the three reviewers and the associate editor for several suggestions that greatly improved this paper.

Appendix. Existence of Equilibrium

In this section we prove that for a large family of *reasonable* price functions the seller may impose, there is an equilibrium in the buyers' subgame. Because low-valuation buyers do not behave strategically (because the seller does not have incentives to lower the price below v), our task is equivalent to show existence of equilibrium for high-value consumers. Specifically, we prove the following result:

Theorem 2. *Consider two classes of buyers arriving according to nonhomogeneous Poisson processes with continuous arrival rate; the first group has rate* $\Lambda: \mathbb{R}_{++} \to \mathbb{R}_{++}$ *and value the item at* V*, and* *the second group has rate $\lambda\colon \mathbb{R}_{++} \to \mathbb{R}_{++}$ and value the item at v. If the seller commits to a price function $p\colon [0,T) \to (v,V)$ that is continuous on $[0,T)$, satisfying $\lim_{t\to T^-} p(t) \in [v,b^*]$, where $b^* = V - ((V-v)(1-e^{-\int_0^T \lambda(t)\,dt}))/(\int_0^T \lambda(t)\,dt)$, and $p(T) = v$, then a symmetric equilibrium exists.*

Observe that we are considering the case of a markdown strategy assuming in addition that $p(t) < V$. This is without loss of generality because if $p(t) = V$ for some $t \in [0,T]$, $U(t) = 0$ for the buyers, and thus nobody buys at this time (unless $p \equiv V$, the case of constant pricing).

It is worth mentioning that the standard fixed point approaches to prove the existence of equilibrium do not seem work here since we have infinitely many players with an infinite set of available pure strategies. Furthermore, the natural fixed-point mapping is hard to analyze. Thus, to prove Theorem 2, we take a constructive approach and build an equilibrium that is in a way symmetric.

The technique for characterizing the equilibrium is innovative. Solving the trade-off between waiting for a lower price and risking loosing the item induces us to split the season into two disjoint subsets: one where the buyers' strategy is to buy upon arrival and the other where buyers use mixed strategies. Consequently, the first main idea of our construction is to divide the interval $[0,T]$ into subintervals. In some of these subintervals, buyers will simply buy upon arrival, whereas in others they will use a mixed strategy over that subinterval. In the subintervals where mixed strategies are used, the conditional distribution determining the buying time of a consumer that arrived at time t and has already waited until time s is independent of t and identical across buyers. Thus, all consumers that wait until a certain time behave identically. It is in this sense that our constructed equilibrium is symmetric.

The second main idea of the proof is the construction of this symmetric equilibrium within an interval, (t_1,t_2), in which mixed strategies are used. Here, we first *iron* the price function and show that one can assume that the function $t \mapsto (V-p(t))e^{-\mu t}$ is nondecreasing. Then, we impose that a symmetric equilibrium is molded by a distribution $H\colon [t_1,t_2] \to [0,1]$, which generates the equilibrium strategy of a consumer arriving at time t as the conditional distribution $(H_t)_{t\in[t,t_2]}$. Finally, by imposing the equilibrium conditions on this family, particularly, that the whole interval maximizes the utility of a buyer, we are able to characterize this distribution through a differential equation, whose solution somewhat surprisingly satisfies all desired properties.

Regarding the technical requirement on $p(t)$ close to T, it ensures that a high-value costumer does not prefer to wait until the end of the season and participate in the lottery with low-value costumers. Hence, the main difficulty of the proof consists in developing strategies that avoid deviations over $[0,T)$. For this reason, we first construct equilibrium strategies assuming that p is continuous over $[0,T]$, with $p(T) \in [v,b^*]$, and, at the end of the section, we show that the same strategies sustain an equilibrium for the case stated in Theorem 2.

A.1. Time Horizon Decomposition

We now decompose the interval $[0,T]$ into subintervals. The key property of these subintervals is that, in equilibrium, all consumers will actually buy in the subinterval they arrived.

Given a continuous function $p\colon [0,T) \to (v,V)$ such that $p(T) \in [v,b^*]$, and letting, for all $t \in [0,T]$, the average arrival rate until time t, $m(t) = \int_0^t \Lambda(x)dx$, we consider the set $I \subseteq [0,T]$:

$$I := \left\{t \in [0,T]\colon t \in \operatorname*{arg\,max}_{s\in[t,T]} (V-p(s))e^{-(\mu s+m(s))}\right\}. \tag{21}$$

In our constructed equilibria, buyers arriving in I will buy upon arrival. Consider now

$$t_* := \min\left\{t \in [0,T]\colon t \in \operatorname*{arg\,max}_{s\in[0,T]} (V-p(s))e^{-(\mu s+m(s))}\right\}, \tag{22}$$

$$t_0 := \min\Big\{t \in [0,t_*]\colon (V-p(t))e^{-\mu t} = (V-p(t_*))e^{-(\mu t_*+m(t_*))}\Big\}. \tag{23}$$

Lemma 3. *The quantities t_* and t_0 are well defined. Furthermore, $t_* = 0$ if and only if $t_0 = 0$.*

Proof. First note that t_* is well defined since p and m are continuous and $[0,T]$ is compact, implying that the set of maximizers of $t \mapsto (V-p(t))e^{-(\mu t+m(t))}$ is also compact. Observe that if $t_* = 0$, then $t_0 = 0$. Conversely, if $t_* > 0$, then $0 \notin \operatorname{arg\,max}_{s\in[0,T]}(V-p(s))e^{-(\mu s+m(s))}$, which yields $(V-p(0)) < (V-p(t_*))e^{-(\mu t_*+m(t_*))}$. Moreover, since $m(t_*) > 0$, we have that $(V-p(t_*))e^{-\mu t_*} > (V-p(t_*))e^{-(\mu t_*+m(t_*))}$. Altogether, the continuity of $t \mapsto (V-p(t))e^{-\mu t}$ yields the existence of t_0 and also that $t_* = 0$ if and only if $t_0 = 0$. □

Whenever $t_* > 0$, buyers arriving in $(0,t_*)$ will use a mixed strategy with support in the interval $(t_0,t_*]$. Note that it makes no sense to buy at $t < t_0$, since $U(t) \le (V-p(t))e^{-\mu t} < (V-p(t_*))e^{-(\mu t_*+m(t_*))} \le U(t_*)$.

As we will prove in Lemma 8, the remainder of the interval $[0,T]$ can be decomposed into a collection of open intervals of the form (t_1,t_2) such that

$$t_2 = \min\left\{t \in [0,T]\colon t \in \operatorname*{arg\,max}_{s\in(t_1,T]} (V-p(s))e^{-(\mu s+m(s))}\right\}, \tag{24}$$

and t_1 is the largest $t < t_2$ satisfying

$$(V-p(t_1))e^{-(\mu t_1+m(t_1))} = (V-p(t_2))e^{-(\mu t_2+m(t_2))}. \tag{25}$$

Buyers arriving in an interval of the form (t_1,t_2) will buy within the interval according to a mixed strategy defined in the next section. Note also for $t_1 < t < t_2$ that we have $(V-p(t))e^{-(\mu t+m(t))} < (V-p(t_2))e^{-(\mu t_2+m(t_2))}$. From now on we refer to these intervals (t_1,t_2) as well as the interval (t_0,t_*) as *mixing* intervals, since mixed strategies are used.

In the next section we will show that every mixing interval (t_1,t_2) has a corresponding distribution H with support on $[t_1,t_2]$ such that buyers arriving at time $t \in (t_1,t_2)$ will buy at a random time drawn according to H_t, the conditional distribution of H in $[t,t_2]$. We may summarize our constructed equilibrium as follows (see Figure A.1):

(i) Consumers arriving in I buy upon arrival.

(ii) Consumers arriving at time $t \in [0,t_0]$ buy at a random time drawn according to the distribution H, corresponding to the mixing interval (t_0,t_*).

(iii) Consumers arriving at time $t \in (t_0,t_*)$ buy at a random time drawn according to the conditional distribution H_t, corresponding to distribution H of the mixing interval (t_0,t_*).

Figure A.1. Interval $[0, T]$ Is Decomposed in Three Collections of Intervals

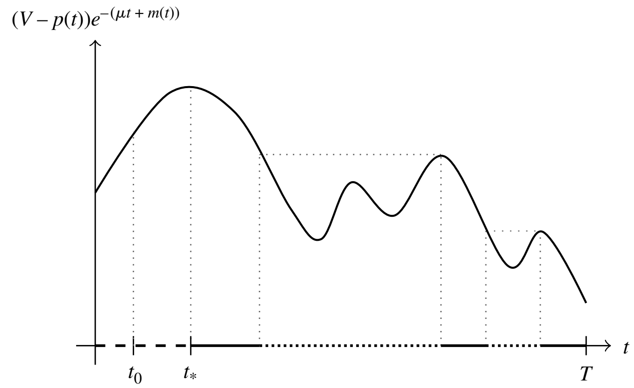

Notes. Costumers arriving in I (continuous line) buy straightaway. Costumers arriving in intervals, marked by a dotted line, play mixed strategies. Finally, costumers arriving in the first interval, marked by a dashed line, play a mixed strategy with support in (t_0, t_*).

(iv) Consumers arriving at time $t \in (t_1, t_2)$ of a generic mixing interval buy at a random time drawn according to the conditional distribution H_t, corresponding to distribution H of the mixing interval (t_1, t_2).

A.2. Strategy in a Mixing Interval

In this section we focus on a mixing interval (t_1, t_2) where t_1 and t_2 satisfy (24) and (25). We isolate (t_1, t_2) assuming that nobody arrived before t_1, which is consistent with our partitioning of the time horizon T. For simplicity we first assume that p is such that the function

$$g_p(t) := (V - p(t))e^{-\mu t}$$

is nondecreasing. At the end of this section, we consider an arbitrary continuous price function.

In the following lemma, we impose that the mixed strategy of a consumer who arrived at t should be the same as that of those who arrived earlier but did not buy before t. This gives us a closed expression of the availability probability f in terms of H.

Lemma 4. *Let H be a continuous distribution over $[t_1, t_2]$. Assume that buyers arriving at $t \in (t_1, t_2)$ buy according to the conditional distribution of H; then, the probability that the item is available at time $t \in [t_1, t_2]$ is given by*

$$f(t) = \exp\left(-m(t) + (1 - H(t)) \int_{t_1}^{t} \frac{\Lambda(x)dx}{1 - H(x)}\right).$$

Proof. Let h be the probability density function of H. Thus, the conditional density on $[t, t_2]$ is

$$h_t(s) = \frac{h(s)}{\int_t^{t_2} h(\tau)d\tau}, \quad \text{for all } t \in (t_1, t_2) \text{ and } s \in [t, t_2]$$

so that the conditional distributions is

$$H_t(s) = \frac{H(s) - H(t)}{1 - H(t)}, \quad \text{for all } t \in (t_1, t_2) \text{ and } s \in [t, t_2]. \tag{26}$$

To apply Equation (1) in the definition of equilibrium, we need an expression for the density of the arrival process. For $t \in (t_1, t_2)$, the density of the arrival time in a nonhomogeneous Poisson process between (t_1, t) is given by $dF_t(x) = (\Lambda(x)/(m(t) - m(t_1)))\,dx$. Also, for every $k \in \mathbb{N}$, we have that $Q_k(t) = e^{-m(t)+m(t_1)}(m(t) - m(t_1))^k/k!$. Hence, it follows that

$$\begin{aligned}
f(t) &= e^{-m(t_1)}\,\mathbb{P}(\text{item is available at } t \mid \text{is available at } t_1)\\
&= e^{-m(t_1)} \sum_{k\ge 0} Q_k(t) \int_{t_1}^{t} \cdots \int_{t_1}^{t} \prod_{i=1}^{k} \frac{1 - H(t)}{1 - H(x_i)}\, dF_t(x_1) \ldots dF_t(x_k)\\
&= e^{-m(t_1)} \sum_{k\ge 0} Q_k(t)(1 - H(t))^k \left(\int_{t_1}^{t} \frac{1}{1 - H(x)} dF_t(x)\right)^k\\
&= e^{-m(t_1)} \sum_{k\ge 0} Q_k(t) \left((1 - H(t)) \int_{t_1}^{t} \frac{1}{1 - H(x)} dF_t(x)\right)^k\\
&= e^{-m(t_1)} \sum_{k\ge 0} e^{-(m(t)-m(t_1))}\\
&\quad \cdot \frac{((m(t) - m(t_1))(1 - H(t)) \int_{t_1}^{t} (1/(1 - H(x)))\, dF_t(x))^k}{k!}\\
&= e^{-m(t_1)} \exp\left[(m(t) - m(t_1))\left(-1 + (1 - H(t))\right.\right.\\
&\qquad \left.\left. \cdot \int_{t_1}^{t} \frac{1}{1 - H(x)}\, dF_t(x)\right)\right]\\
&= \exp\left(-m(t) + (1 - H(t)) \int_{t_1}^{t} \frac{\Lambda(x)}{1 - H(x)}\, dx\right). \quad \square
\end{aligned}$$

We now turn to give an explicit expression for the strategies of buyers arriving in a mixing interval. For $t \in (t_1, t_2]$, consider the function

$$\begin{aligned}
H(t) &= 1 - (\ln(g_p(t_1)/g_p(t)) + m(t) - m(t_1))\\
&\quad \cdot \left(K \exp\left(-\int_t^{(t_1+t_2)/2}\right.\right.\\
&\qquad \left.\left. \cdot \frac{\Lambda(x)dx}{\ln(g_p(t_1)/g_p(x)) + m(x) - m(t_1)}\right)\right)^{-1},
\end{aligned} \tag{27}$$

where $K > 0$ is a constant to be determined later. The next result shows that this is actually a distribution and that if all consumers buy according to the conditional distribution given by (26), namely,

$$\begin{aligned}
H_t(s) &= \frac{H(s) - H(t)}{1 - H(t)}\\
&= 1 - \exp\left(-\int_t^s \frac{\Lambda(x)dx}{\ln(g_p(t_1)/g_p(x)) + m(x) - m(t_1)}\right)\\
&\quad \cdot \frac{\ln(g_p(t_1)/g_p(s)) + m(s) - m(t_1)}{\ln(g_p(t_1)/g_p(t)) + m(t) - m(t_1)},
\end{aligned} \tag{28}$$

then their utility $U(t)$ is constant in the interval $[t_1, t_2]$.

Lemma 5. *Let p be a price function such that g_p is nondecreasing. Then, there is $K > 0$ such that H, defined by (27), is nondecreasing, continuous, and satisfies that $H(t_2) = 1$. Furthermore, if all consumers buy according to $(H_t)_{t\in(t_1,t_2)}$, the family of distributions defined in (28), their utility satisfies $U(t) = U(t_1)$ for all $t \in [t_1, t_2]$.*

Proof. We proceed backward by first imposing that the utility is constant throughout the interval and study the implications of this condition. Since we assume that nobody arrived before t_1, we have $f(t_1) = e^{-m(t_1)}$, and the condition $U(t) = U(t_1)$, which is equivalent to $(V - p(t))e^{-\mu t} f(t) = (V - p(t_1))e^{-\mu t_1} f(t_1)$, yields

$$e^{m(t_1)} f(t) = \frac{(V - p(t_1))e^{-\mu t_1}}{(V - p(t))e^{-\mu t}} = \frac{g_p(t_1)}{g_p(t)} \quad \text{for all } t \in (t_1, t_2].$$

Note that, since g_p is nondecreasing, f is nonincreasing, which is consistent with the fact that f represents the probability of the item being available. From Lemma 4 we obtain the equation

$$\frac{g_p(t_1)}{g_p(t)} = \exp\left(-m(t)+m(t_1)+(1-H(t))\int_{t_1}^{t}\frac{\Lambda(x)dx}{1-H(x)}\right),$$

where the unknown is H. This can be rewritten as

$$\ln\frac{g_p(t_1)}{g_p(t)} + m(t)-m(t_1) = (1-H(t))\int_{t_1}^{t}\frac{\Lambda(x)dx}{1-H(x)}\,dx \quad \text{for all } t\in(t_1,t_2].$$

Denoting $u\colon t\mapsto \int_{t_1}^{t}(\Lambda(x)dx/(1-H(x)))$, we have that $u'(t)=\Lambda(t)/(1-H(t))$ and then we transform the previous integral equation into the differential equation

$$\frac{\Lambda(t)}{\ln(g_p(t_1)/g_p(t))+m(t)-m(t_1)} = \frac{u'(t)}{u(t)}.$$

The latter is solved by integrating from $(t_1+t_2)/2$ to t, which leads to

$$\ln(u(t))-\ln(u((t_1+t_2)/2)) = \int_{(t_1+t_2)/2}^{t}\frac{\Lambda(x)dx}{\ln(g_p(t_1)/g_p(x))+m(x)-m(t_1)}.$$

Defining $K=u((t_1+t_2)/2)=\int_{t_1}^{(t_1+t_2)/2}(\Lambda(x)dx/(1-H(x)))>0$, we obtain that the solution is

$$u(t)=K\exp\left(\int_{(t_1+t_2)/2}^{t}\frac{\Lambda(x)dx}{\ln(g_p(t_1)/g_p(x))+m(x)-m(t_1)}\right),$$

and hence,

$$\begin{aligned}(\forall\, t\in(t_1,t_2))\quad \frac{\Lambda(t)}{1-H(t)} &= u'(t)\\ &= \left(K\Lambda(t)\exp\left(\int_{(t_1+t_2)/2}^{t}\frac{\Lambda(x)dx}{\ln(g_p(t_1)/g_p(x))+m(x)-m(t_1)}\right)\right)\\ &\quad\cdot\left(\ln\left(\frac{g_p(t_1)}{g_p(t)}\right)+m(t)-m(t_1)\right)^{-1}, \qquad (29)\end{aligned}$$

which yields (27). Now, since (25) yields $g_p(t_1)e^{-m(t_1)} = g_p(t_2)e^{-m(t_2)}$, it is clear that the right-hand side of (29) goes to infinity as $t\to t_2$ so that we can set $H(t_2)=1$. Also, since p is continuous, g_p is continuous and H is continuous in $(t_1,t_2]$.

To see that H is nondecreasing, assume for simplicity that p is differentiable. In this case,

$$H'(t) = (g_p{}'(t))\cdot\left(g_p(t)K\cdot\exp\left(-\int_{t}^{(t_1+t_2)/2}\frac{\Lambda(x)dx}{\ln(g_p(t_1)/g_p(x))+m(x)-m(t_1)}\right)\right)^{-1}\geq 0.$$

In general, the monotonicity of H can be easily obtained by considering a sequence $(g_p^n)_{n\in\mathbb{N}}\in\mathscr{C}^\infty([0,T])$ of nondecreasing functions such that g_p^n converges to g_p uniformly on $[0,T]$. To conclude, note that the conditional distributions defined in (28) are indeed distributions. Since H is continuous and nondecreasing in $(t_1,t_2]$, H_t is also continuous and nondecreasing in $[t,t_2]$. Moreover, because $H(t_2)=1$, we get that $H_t(t_2)=1$ and $H_t(t)=0$. □

Remark 6. H remains constant on a subset of (t_1,t_2) if and only if g_p remains constant.

Now we tackle the general case when the assumption on the monotonicity of g_p is dropped. For this we define an auxiliary price scheme over (t_1,t_2),

$$\bar p(t) := V+\inf_{\tau\in(t_1,t)}\{e^{\mu(t-\tau)}(p(\tau)-V)\} = V-e^{\mu t}\sup_{\tau\in(t_1,t)} g_p(\tau), \quad (30)$$

and define the distribution, $\bar H$, corresponding to the mixing interval (t_1,t_2) as in (27) but using $g_{\bar p}$ instead of g_p. Similarly, we define the strategy for $t\in(t_1,t_2)$ as $\bar H_t$, the conditional distribution of $\bar H$ obtained as in (28).

To see the intuition behind $\bar p$, note that when $\mu=0$, $\bar p(t)=\inf_{\tau\in(t_1,t)}p(\tau)$. In general, $\bar p$ is the largest function below p such that $g_{\bar p}(t)=\sup_{\tau\in(t_1,t)}g_p(\tau)$ is nondecreasing. Indeed, whenever g_p is decreasing, $g_{\bar p}$ remains constant, and furthermore, if in an interval $g_{\bar p}\neq g_p$, then $g_{\bar p}$ is constant in that interval.

Next, we assert that $\bar H$ and $\bar H_t$ are indeed distributions. In fact, observe that $p(t)=\bar p(t)$ if and only if $g_p(t)=g_{\bar p}(t)$, and therefore the set $A:=\{t\in(t_1,t_2)\mid \bar p(t)\neq p(t)\}$ is the same set as the set where $g_{\bar p}(\cdot)$ remains constant. Thus, by Remark 6, the support of $\bar H$ is actually $[t_1,t_2]\setminus A$, and therefore $\bar H$ and $\bar H_t$ are nondecreasing. Finally, invoking Lemma 4, we conclude that the subset of times over a mixing interval where $f(\cdot)$ stays constant is the set A.

Remark 7. The analysis for the mixing interval (t_0,t_*) is analogous, excepting that we have to take into account that the buyers arriving in $[0,t_0]$ are waiting to buy on (t_0,t_*). Hence, at the moment of computing f, we consider that $Q_k(t)=e^{-m(t)}m(t)^k/k!$.

A.3. Putting the Pieces Together

We are now ready to prove Theorem 2. First, recall that the strategies of buyers in the game can be summarized as follows:

(i) Consumers arriving in I buy upon arrival.

(ii) Consumers arriving at time $t\in[0,t_0]$ buy at a random time drawn according to the distribution $\bar H$, corresponding to the mixing interval (t_0,t_*) constructed using (27), (30), and Remark 7.

(iii) Consumers arriving at time $t\in(t_0,t_*)$ buy at a random time drawn according to the conditional distribution $\bar H_t$, corresponding to distribution $\bar H$ of the mixing interval (t_0,t_*).

(iv) Consumers arriving at time $t\in(t_1,t_2)$ of a generic mixing interval buy at a random time drawn according to the conditional distribution $\bar H_t$, defined by (28) and (30), corresponding to the distribution $\bar H$ of the mixing interval (t_1,t_2).

In the next lemma, we show that the time decomposition of the horizon $[0,T]$ is correct. This means that for every $t\in[0,T]$ we can associate a strategy as just described.

Lemma 8. *The interval $[0,T]$ can be partitioned into the set I, the interval $[0,t_*)$, and a collection of mixing intervals of the form (t_1,t_2), as defined by Equations (21)–(25). Hence, the constructed strategies are well defined for every $t\in[0,T]$.*

Proof. We first prove that the set I is a compact. Indeed, consider a convergent sequence in I, $x_n\to x$ and $\varepsilon>0$. By continuity of $T(s):=(V-p(s))e^{-(\mu s+m(s))}$, there is $M>0$ such that

$$\max_{s\geq x_n}T(s)+\varepsilon \geq \max_{s\geq x}T(s) \quad \text{for all } n\geq M.$$

Since $x_n \in I$ we obtain that $T(x_n) + \varepsilon \geq \max_{s\geq x} T(s)$ for all $n \geq M$. Taking $n \to \infty$, we conclude that $x \in \arg\max_{s\geq x} T(s)$, equivalently, $x \in I$. Then, I is closed and thus compact.

The compactness of I quickly implies the lemma. Let $t \in [0, T]$, and let $\bar{t} = \min\{s \geq t\colon s \in I\}$. Clearly, if $t = \bar{t}$, then $t \in I$. So assume $t < \bar{t}$. In this case, if $\bar{t} = t_*$, then $t \in [0, t_*]$. Otherwise there is $s \in I$ such that $s < t$. So, letting $t_1 = \max\{s \leq t \mid s \in I\}$ and $t_2 = \bar{t}$, they satisfy (25) and (24), respectively; therefore $t \in (t_1, t_2)$. □

Slightly abusing notation, let us call H_t the strategy defined for a buyer arriving at time $t \in [0, T]$. We are finally able to prove our main result: if all buyers behave according to H_t, we have a Bayes–Nash equilibrium of the subgame.

Proof of Theorem 2. First, let us show that if H_t is an equilibrium when p is continuous, then it will also be an equilibrium when p is only continuous over $[0, T)$ and $p(T) = v$. In fact, if all players follow H_t, the utility of any player arriving before T is greater or equal than $(V - p(T^-))e^{-(\mu T + m(T))}$. Then, since $p(T^-) \leq V - ((V - v)(1 - e^{-\int_0^T \lambda(t)\,dt}))/(\int_0^T \lambda(t)\,dt)$, we conclude that is not profitable to wait and purchase at the end of the season.

To conclude that H_t is an equilibrium for the game, it only remains to prove that there is no profitable deviation over $[0, T)$. By contradiction, assume that there is a deviation for some player, i.e., there is $t \in [0, T)$ and $z \geq t$ such that z lies outside the support of H_t and $U(z) > U(s)$ for s in the support of H_t.

If $t \in I$, we have that $H_t(t) = 1$ (buy upon arrival), and by definition of I, we obtain $(V - p(t))e^{-(\mu t + m(t))} \geq (V - p(z)) e^{-(\mu z + m(z))} \geq U(z)$, where the last inequality follows because, by construction, every buyer arriving before t will end up buying before t.

Otherwise, if t belongs to a mixing interval of the form $t \in (t_1, t_2)$, recalling that the utility of a buyer arriving at time t equals $U(t_1)$ and, by Lemma 5, for $\tilde{z} \in (t_1, t_2]$ we have that $U(\tilde{z}) \leq U(t_1)$, we conclude that $z > t_2$. But if $z > t_2$, the situation is analogous to the case $t \in I$ since, by definition of t_2, $(V - p(t_2))e^{-(\mu t_2 + m(t_2))} \geq (V - p(z))e^{-(\mu z + m(z))} \geq U(z)$.

Finally, we study the case $t \in [0, t_*]$ with $t_* > 0$. If $z \geq t_0$, the situation is analogous to the previous case. If, on the contrary, $z < t_0$ and $U(z) > U(t_0)$, since $t_* > 0$, we have that $U(0) < U(t_*) = U(t_0)$; therefore, by continuity, there is $\bar{z} < t_0$ such that $U(\bar{z}) = (V - p(\bar{z}))e^{-\mu\bar{z}} = U(t^*) = (V - p(t_*))e^{-(\mu t_* + m(t_*))}$, contradicting the minimality of t_0. □

Endnotes

[1] See http://www.mercadominero.cl/sitio/home.php?lang=1 (accessed June 9, 2016).

[2] Clearly $s' \in [0, s]$, and some basic calculus shows that actually $s' < s$.

[3] We define the infimum over the empty set as V.

[4] Note that equilibrium is actually defined as a price and a probability of availability f, but using (12), we see that φ characterizes the equilibrium as well.

References

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

Board S, Skrzypacz A (2016) Revenue management with forward-looking buyers. *J. Political Econom.* 124(4):1046–1087.

Cachon G, Feldman P (2015) Price commitments with strategic consumers: Why it can be optimal to discount more frequently . . . than optimal. *Manufacturing Service Oper. Management* 17(3):399–410.

Caldentey R, Vulcano G (2007) Online auction and list price revenue management. *Management Sci.* 53(5):795–813.

Caldentey R, Liu Y, Lobel I (2015) Intertemporal pricing under minimax regret. Working paper, University of Chicago, Chicago.

Correa JR, Montoya R, Thraves C (2016) Contingent preannounced pricing policies with strategic consumers. *Oper. Res.* 64(1): 251–272.

Elmaghraby W, Lippman S, Tang C, Yin R (2009) Pre-announced pricing strategies with reservations. *Production Oper. Management* 18(4):381–401.

Gershkov A, Moldovanu B, Strack P (2014) Revenue maximizing mechanisms with strategic customers and unknown demand: Name-your-own-price. Working paper, Hebrew University of Jerusalem, Jerusalem.

Haviv M, Milchtaich I (2012) Auctions with a random number of identical bidders. *Econom. Lett.* 114(2):143–146.

Jerath K, Netessine S, Veeraraghavan SK (2010) Revenue management with strategic customers: Last-minute selling and opaque selling. *Management Sci.* 56(3):430–448.

Levin D, Ozdenoren E (2004) Auctions with uncertain numbers of bidders. *J. Econom. Theory* 118(2):229–251.

Levin D, Smith JL (1994) Equilibrium in auctions with entry. *Amer. Econom. Rev.* 84(3):585–599.

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

McAfee RP, McMillan J (1987) Auctions with a stochastic number of bidders. *J. Econom. Theory* 43(1):1–19.

Mierendorff K (2016) Optimal dynamic mechanism design with deadlines. *J. Econom. Theory* 161:190–222.

Surasvadi N, Vulcano G (2013) Price matching under conditional discounts in the presence of strategic consumers. Working paper, National Electronics and Computer Technology Center, Pathum Thani, Thailand.

Osadchiy N, Vulcano G (2010) Selling with binding reservations in the presence of strategic consumers. *Management Sci.* 56(12): 2173–2190.

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

Russo R, Shmueli G, Jank W, Shyamalkumar N (2010) Models for bid arrivals and bidder arrivals in online auctions. Kotz S, Read CB, Balakrishnan N, Vidakovic B, eds. *Encyclopedia of Statistical Sciences*, Chap. 23 (John Wiley & Sons, New York), 293–309.

Skreta V (2006) Mechanism design for arbitrary type spaces. *Econom. Lett.* 91(2):293–299.

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

Talluri K, van Ryzin G (2005) *The Theory and Practice of Revenue Management* (Springer, New York).

Vinter R (2000) *Optimal Control* (Birkhäuser, Berlin).

Yin R, Aviv Y, Pazgal A, Tang C (2009) Optimal markdown pricing: Implications of inventory display formats in the presence of strategic customers. *Management Sci.* 55(8):1391–1408.