

The Value of Observing the Buyers' Arrival Time in Dynamic Pricing

José Correa,^{a,*} Dana Pizarro,^b Gustavo J. Vulcano^c

^aDepartamento de Ingeniería Industrial, Universidad de Chile, 1058 Santiago, Chile; ^bInstituto de Ciencias de la Ingeniería, Universidad de O'Higgins, 2820000 Rancagua, Chile; ^cSchool of Business, Universidad Torcuato Di Tella, and CONICET, Buenos Aires 1428, Argentina

*Corresponding author

Contact: correa@uchile.cl, <https://orcid.org/0000-0002-3012-7622> (JC); dana.pizarro@uoh.cl (DP); gvulcano@utdt.edu, <https://orcid.org/0000-0003-0083-4737> (GJV)

Received: July 26, 2021

Revised: November 9, 2022; March 8, 2023

Accepted: April 4, 2023

Published Online in Articles in Advance:
June 1, 2023

<https://doi.org/10.1287/mnsc.2023.4794>

Copyright: © 2023 INFORMS

Abstract. We consider a dynamic pricing problem in which a firm sells one item to a single buyer to maximize expected revenue. The firm commits to a price function over an infinite horizon. The buyer arrives at some random time with a private value for the item. He is more impatient than the seller and strategizes over the timing of the purchase in order to maximize his expected utility, which implies either buying immediately, waiting to benefit from a lower price, or not buying. We study the value of the seller's ability to observe the buyer's arrival time in terms of her expected revenue. When the seller can observe the buyer's arrival, she can make the price function contingent on the buyer's arrival time. On the contrary, when the seller can't, her price function is fixed at time zero for the whole horizon. The value of observability (VO) is defined as the worst-case ratio between the expected revenue of the seller when she observes the buyer's arrival and that when she does not. First, we show that, for the particular case in which the buyer's valuation follows a monotone hazard rate distribution, the upper bound of VO is $\exp(1)$. Next, we show our main result: in a setting very general on valuation and arrival time distributions: VO is at most 4.911. To obtain this bound, we fully characterize the solution to the observable arrival problem and use this solution to construct a random and periodic price function for the unobservable case. Finally, we show by solving a particular example to optimality that VO has a lower bound of 1.136.

History: Accepted by Itai Ashlagi, revenue management and market analytics.

Funding: This work was partially supported by the Agencia Nacional de Investigación y Desarrollo Chile [Grants ACT210005 and FONDECYT 1220054], and by the Center for Mathematical Modeling at the University of Chile [Grant FB210005]. The work of D. Pizarro has benefited from the Artificial and Natural Intelligence Toulouse Institute, which is funded by the French "Investing for the Future - PIA3" program [Grant ANR-19-P3IA-0004].

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/mnsc.2023.4794>.

Keywords: dynamic pricing • strategic consumer behavior • price commitment • price discrimination • worst-case analysis

1. Introduction

The dynamic pricing practice has been around for decades in different industries, ranging from airlines and hotels to supermarkets and clothing retailers. Yet the rise of online business platforms since the early 2000s has accelerated its presence. In particular, dynamic pricing has expanded from traditional domains to almost any business-to-consumer and consumer-to-consumer environment with ride-sharing probably being one of the current most prominent examples.

The widespread use of dynamic pricing has multiple drivers, including the need to liquidate excess inventories within a limited time frame (e.g., apparel retailers), the varying opportunity cost of scarce capacity (e.g., car rental), the chance of gathering information about the underlying demand (e.g., learning about elasticity), and

the possibility of extracting most of the surplus from a heterogeneous customer base via a form of intertemporal price discrimination. Its feasibility and successful implementation are rooted in the increasing availability of data from internal and external sources, sophisticated and rapid advances in machine learning and artificial intelligence, and tremendous progress in computational speed. In fact, digital technology has made it possible to continuously adjust prices to changing environments with minimal efforts and costs. On the downside, its extended use, concurrent with the price transparency occurring in online platforms, has raised some concerns, particularly among online retailers, because consumers have learned to strategize over the timing of their purchases and wait to take advantage of a potential future lower price.

The academic literature in operations management (OM)—and from a different perspective, also in economics and computer science—accounts for the rapid evolution of the dynamic pricing practice, acknowledging in several articles the threat represented by the forward-looking behavior of consumers. Our paper contributes to this stream by providing an assessment of the incremental revenue attainable by being able to observe consumer arrival time and adjust the pricing policy accordingly as opposed to the case in which this capability is absent. To our knowledge, this value of observability (VO) is overlooked in the literature so far.

1.1. Problem Description

We consider a stylized yet fundamental model in which one seller interacts with one buyer. The seller holds a single item whose value is normalized to zero, whereas the buyer has a random private valuation for it. The buyer's arrival time follows an arbitrary distribution over the nonnegative reals. Both the buyer and the seller discount the future, but they do it at different rates, the buyer being more impatient than the seller. The seller's goal is to set up a price function to maximize her expected discounted revenue. On the buyer's side, upon arrival, he observes the price function and decides to buy at the time that is most profitable for him or not buy at all.

The ability of the seller to observe the buyer's arrival (or not) determines two scenarios. In one case, the seller sets the price curve from the very beginning of the time horizon regardless of the effective arrival time of the consumer. This could be simply a result of the seller ignoring the consumer arrival time or being unable to accurately register it. The latter may happen in online marketplaces, in which it may be challenging to distinguish interested buyers from other traffic on the website (e.g., robots searching competitors' prices). Therefore, assuming that the seller can observe the arrival of the interested buyers may not be realistic. In the other scenario, the seller is indeed able to track the arrival time of the interested customer and, therefore, may internalize this piece of information in the price curve she proposes to the buyer. The extent to which this observational ability produces an additional rent to the seller is the main subject of this paper.

In the observable case, the seller designs a menu of price functions indexed by time and shows the buyer the specific price curve tailored to his arrival time. In the unobservable case, this information is absent, and the seller has to set a price curve from the beginning of the selling horizon, only knowing the arrival distribution. These two scenarios naturally lead to define the value of observability of a given instance of the problem as the ratio between the revenue of the seller in the observable and unobservable cases. Here, an instance of the problem is defined by the buyer's arrival time and valuation

distributions and the discount rates for both the buyer and the seller. Then, the more general VO is defined as the supremum of the corresponding instance-specific value of observability taken over all possible distributions and discount rates. This VO corresponds to the worst-case ratio between the revenue of the seller in the observable and unobservable cases. The focus of our work is to bound this worst-case ratio.

There are two equivalent interpretations for our model that are worth highlighting. The first concerns the single buyer and single unit situation we describe. The model could be alternatively interpreted as having a continuum of buyers with total mass normalized to one. In the observable case, the infinitesimal mass of each of these buyers is represented by the probability density function (pdf) of the buyer's private valuation. This interpretation is extended by also accounting for the pdf of the buyers' arrival distribution in the unobservable case and considering an infinitesimal mass of buyers described by the joint pdf between the valuation and the arrival time distributions. On the supply side, we assume a unit supply that is infinitesimally partitioned so that it can be assumed unlimited.

The second interpretation connects our VO results to the notion of price of discrimination. To see this, consider the continuum-of-buyers view of our model, in which the seller sets a customized price curve for each possible arrival time of buyers so that the total expected discounted revenue she obtains is the same as that one achieved in the observable case. On the other hand, if the seller does not have this power, she should offer the same price curve to all customers since the beginning of the selling horizon. The latter problem is exactly the same as the unobservable case described earlier. Therefore, if we define the price of discrimination as the additional rent the seller can obtain by posting a customized price curve, it becomes equivalent to VO. In other words, we are also providing a bound for the price of discrimination.

Before proceeding, we point out here that our model, though stylized, captures some key features of several important business settings. For instance, in the video game industry, the connection between the producer or seller and the buyer can be modeled as a one-to-one relationship in which implementing personalized dynamic pricing is indeed feasible. For example, for purchasing a video game, some consumers with a high discount rate purchase the game immediately upon release. Some other consumers may anticipate an eventual price decrease and wait, whereas others may not have heard of the game for a while and, consequently, have later arrival times. These latter two types of consumers may or may not be distinguishable.

1.2. Our Results

Our main contribution is to establish that, for arbitrary arrival and valuation distributions of the buyer and

arbitrary discount rates of both the seller and the buyer, the value of observability is bounded above by a small constant. This result is somewhat surprising because of three key factors: (i) the setup of the model is very general; (ii) the bound is totally independent of the model primitives; and (iii) simple pricing strategies, such as fixed pricing, fail to guarantee a constant bound.

En route to this result, we first analyze the observable arrival case. In this context, we take a pricing approach that, as usually different from the mechanism design approach, allows us to write the seller's problem as an optimal control problem and fully characterize its solution. In particular, we can prove a key result (Lemma 1), establishing that, under optimal pricing, the seller extracts a constant fraction of the total revenue within a short time period that solely depends on the seller's discount rate.

Then, we turn to study the unobservable case. Unfortunately, this problem is much harder to analyze, and obtaining an explicit solution seems hopeless. However, to prove that the value of observability is bounded by a constant, it is enough to exhibit a feasible pricing policy that can recover a constant fraction of the revenue of the optimal solution in the observable case. There are three main obstacles that we must circumvent to get our main result. First, we use part of the structure of the solution of the observable case and repeat it over time to construct a periodic price function. Because the solution of the observable case already takes an infinite time to implement (which implies an arbitrarily long period when plugged in as a feasible solution to the unobservable arrival case), the aforementioned key lemma comes into play and allows us to implement this repeated pricing within small time windows. The second obstacle is that we should be careful with the buyer's forward-looking behavior. To account for this strategic wait, we simply introduce empty space, say, by using a very high price, before each application of the optimal observable pricing so as to make a buyer, arriving within this empty space, behave as in the observable case. Again, this comes at a loss of a constant fraction of the revenue. Finally, the third difficulty stems from some arrival distributions that might be biased toward regions in which the feasible pricing policy we consider is too low. To overcome this, we apply a random shift to our price curve, which allows us to treat the buyer's arrival time as if it were uniform on a given interval. Ultimately, by carefully dealing with these three obstacles, we can state our main result (Theorem 1): the proposed pricing scheme for the unobservable case attains an expected revenue of at least a fraction $1/4.911 = 0.203$ of the optimal revenue in the observable case. Along the way, we characterize an explicit pricing policy to approximately solve the hard unobservable case.

We show that the situation is much simpler for the special and relevant case of valuation distributions having a monotone hazard rate (MHR), which includes several standard distributions, such as the normal, uniform, logistic, exponential, and double exponential. Indeed, it is enough to consider a fixed price curve in the unobservable arrival case (i.e., the price is constant over the whole period) to recover a fraction $1/e$ of the optimal revenue in the observable case. We further note that fixed pricing cannot guarantee a constant fraction in the general valuation distribution case.

Interestingly, we also observe that our results are robust to the distribution of arrivals. Even if the arrival time of the buyer was chosen by an adversary that knows the price function of the seller (but does not know the realization of the random shift) then our bound on the VO still applies.

Beyond the specific bound we are able to characterize, our result has important managerial implications. The seller may wonder how much she is leaving on the table by not being able to track the arrival time of the customers or, in other words, how much she is willing to pay for introducing this capability. Our conclusion is that this value could be significant from a business perspective, but it is not unbounded and does not depend on the problem parameters. It is not very significant when consumers' valuation distribution has a monotone hazard rate and even less important when the seller has a level of patience similar to the customer or is much more patient than him.

1.3. Road Map

The remainder of this paper is organized as follows: We start with a literature review in Section 2, followed by the precise model description in Section 3, spanning both the buyer's (Section 3.1) and seller's (Section 3.2) problems. The seller problem description includes the formulations of both the standard observable case and the more challenging unobservable case. Both cases are later analyzed in detail in Sections 4 and 5, respectively. Finally, the bounds for the VO are established in Section 6. We close the paper with our concluding remarks in Section 7. The proofs of the results stated in the main body of the paper are relegated to Online Appendix A2.

2. Literature Review

The literature on intertemporal price discrimination under forward-looking consumers was pioneered by Stokey (1979), who considers a monopolist selling an unlimited inventory of a product by committing to a continuously declining price scheme over a finite horizon. All consumers are present at time zero and stay until either purchasing a unit or the end of the season, whichever occurs first. Stokey (1979) shows that price discrimination is not profitable compared with a fixed-price strategy when the seller and the consumers discount the

future at the same rate. Landsberger and Meilijson (1985) study a particular case of Stokey (1979) in which consumers have an exponentially discounted utility function and are more impatient than the seller (which has become a standard setup in the literature and which we also follow in our model). The seller announces a price function that is continuous and differentiable. They show that, in this setup, intertemporal price discrimination strictly dominates the fixed-price policy.

Over the last four decades, there is a vast literature in economics and OM on the topic of dynamic pricing and strategic consumer behavior. More recently, the topic caught the interest of part of the computer science community. Although the borders are blurred, often, current research in OM deals with finding optimal or approximately optimal dynamic pricing mechanisms (e.g., Besbes and Lobel 2015, Caldentey et al. 2017, Gershkov et al. 2018), whereas in economics, the central interest is to find optimal dynamic mechanisms that may imply departing from basic pricing schemes (e.g., Pavan et al. 2014, Board and Skrzypacz 2016), and in computer science, the interest is in designing simple mechanisms that are approximately optimal (e.g., Blumrosen and Holenstein 2008, Chawla et al. 2010, Correa et al. 2019, Kessel et al. 2022). We refer the reader to the books by Talluri and van Ryzin (2004) and Gallego and Topaloglu (2019) for a detailed technical presentation of models on pricing. The book chapter by Aviv and Vulcano (2012) surveys the literature on dynamic list pricing until the 2000s with emphasis on operational applications.

One of the early papers to address a dynamic pricing problem under strategic consumer behavior in operational contexts (although published in an economics journal) is Conlisk et al. (1984). The authors analyze the problem of a monopolist facing an arriving stream of customers over time who are, in turn, intertemporal utility maximizers. They assume that consumer valuations could be either low or high and characterize optimal cyclic policies. Later, since the 2000s, the OM community has paid attention to the design of pricing mechanisms to mitigate the adverse impact of strategic consumer behavior on firms' revenues. These mechanisms exceed traditional list pricing and include capacity rationing (e.g., Su 2007, Liu and Van Ryzin 2008), quick response production (e.g., Cachon and Swinney 2009), changing inventory display formats (e.g., Yin et al. 2009), making price and capacity commitments (e.g., Aviv and Pazgal 2008, Su and Zhang 2008, Mersereau and Zhang 2012, Correa et al. 2016), internal price matching (e.g., Lai et al. 2010), and binding reservations (e.g., Elmaghraby et al. 2009, Osadchiy and Vulcano 2010). A comprehensive reference on this topic is the book chapter by Aviv et al. (2009). Despite the common wisdom about the existence of such forward-looking consumer behavior and the need to incorporate it in the

decision-making process within operational applications, it was not until the mid 2010s when Li et al. (2014) show that between 5.2% and 19.2% of the consumer base they study within the air-travel industry strategized the timing of their bookings.

As mentioned, an important body of recent research on dynamic pricing with an algorithmic twist has emerged within the OM community. Borgs et al. (2014), motivated by the selling strategy of online services, analyze a multiperiod pricing problem of a firm with capacity levels that vary over time. Customers are heterogeneous in their arrival and departure periods as well as valuations and are fully strategic with respect to their purchasing decisions. The firm's problem is to set a sequence of prices that maximizes its revenue and guarantees service to all paying customers. Besbes and Lobel (2015) study a fluid model in which customers arrive over time, are strategic in timing their purchases, and are heterogeneous in their valuation and willingness to wait before purchasing or leaving. There is no inventory limitation. They show that the firm may restrict attention to cyclic pricing policies that have length, at most, twice the maximum willingness to wait of the customer population. Caldentey et al. (2017) take a robust approach for the intertemporal pricing problem based on the minimization of the seller's worst case regret over a finite horizon. Customer types differ along willingness to pay and arrival time during the selling season, and the seller only knows the support of the customers' valuations. They further assume that there is no inventory limitation and the seller and the consumers discount the future at the same rate. For markets with either myopic or strategic customers, they characterize optimal price paths. Chen and Farias (2018) study the typical dynamic pricing problem under forward-looking consumers with two particular features: (i) the private valuations of these customers decay over time, and (ii) the customers incur monitoring costs. Both the rates of decay and monitoring costs are private information. The authors propose a "robust pricing" mechanism that is guaranteed to achieve expected revenues that are at least within 29% of those under an optimal (not necessarily posted price) dynamic mechanism. In Chen et al. (2019), a paper with a focus on multiproduct, network revenue management (RM), the authors show for the single-product case that an optimally set fixed price guarantees the seller revenues that are within at least 63.2% of that under an optimal dynamic mechanism.

A different line of models in which customers are "patient" rather than "strategic" has recently caught the attention of the RM community. Liu and Cooper (2015) and Lobel (2020) belong to this stream and show the structural optimality of cyclic pricing policies. In Araman and Fayad (2021), consumers are not only patient

but also have time-varying stochastic valuations. The authors show that cyclic policies are near-optimal in this case.

We have discussed so far some of the most relevant papers belonging to the prolific literature on dynamic pricing under strategic consumer behavior. Different from the existing literature, the center of our analysis is the distinction between the consumers' arrivals being observable or not. Regarding observable arrivals, the paper possibly closest to ours is Wang (2001), who also resorts to Euler–Lagrange optimality conditions to solve the pricing problem though his focus is related to the impact of different relative magnitudes between the discount rates of the seller and the infinitesimal buyers (adding up to a mass of one), all present at time zero. Although we consider a similar setting, his model imposes more technical structure, assuming that both the price and purchasing functions are monotone decreasing. He also considers an extension with buyers arriving according to a Poisson process and in which the seller bargains with one buyer at a time upon the buyer's arrival. In our model, we do not discuss bargaining, but allow a general arrival distribution. Wang (2001) also considers the nonobservable arrival case in his work. In particular, he considers buyers arriving according to a Poisson process facing a decreasing price function. The main modeling difference with our work is that we assume only one buyer with an arrival time following a general distribution and the price function may not be decreasing. More fundamentally, Wang (2001) did not address the VO bound.

As mentioned, our result can also be interpreted as the price of discrimination. A recent work by Elmachoub et al. (2021) studies when implementing price discrimination is indeed convenient and when it is not. Specifically, they provide lower and upper bounds (that depend on some parameters of the model) on the ratio between the revenue achievable from charging each customer the customer's own valuation and the revenue obtainable through a fixed price policy. They also compare the profit obtained when the seller observes some information (but not the buyer valuation) before setting the pricing policy, with the one earned by each of the two strategies described earlier. However, they do not consider a dynamic problem; only a static one, which is a substantial difference with our work.

3. Model Description

We study the problem faced by a firm (seller) endowed with a single unit for sale over an infinite time horizon. The value of the item for the seller is normalized to zero. We take an RM point of view and assume that the seller cannot replenish this unit throughout the selling horizon. On the demand side, a single consumer arrives at a time that follows a cumulative distribution function

(cdf) $G : [0, \infty] \rightarrow [0, 1]$ and density g . The consumer may never arrive, and we model this option by considering the arrival time $t = \infty$. The buyer has a private valuation v for the item with cdf $F : [0, \bar{v}] \rightarrow [0, 1]$ and density f . We assume that the arrival time and the valuation for the buyer are independent and both G and F are common knowledge.

The interaction between the seller and the buyer is formalized as a Stackelberg game in which the seller is the leader and precommits to a price function $p(t)$ over time in order to maximize her expected revenue. The buyer is the follower and has to decide whether and when to purchase the item given the price function set up by the seller. The seller and the buyer discount the future at rates δ and μ , respectively.

We discuss two possible variants of this problem. In the observable case, the seller is able to track the buyer's arrival time τ , and from that moment onward, she commits to a price function $p : [\tau, \infty] \rightarrow [0, \bar{v}]$. In the unobservable case, the seller does not see the buyer's arrival time (although she does know the arrival time distribution G), and since time 0, she commits to a price function $p : [0, \infty] \rightarrow [0, \bar{v}]$. Note that, given the bounded support of the valuation distribution, the price function is lower bounded by zero and upper bounded by \bar{v} .

For technical reasons, in both cases, we impose the mild condition that the price function p is lower semi-continuous and twice differentiable almost everywhere.¹ In what follows, we introduce the buyer's and seller's problems as well as some preliminary definitions and results. The discussion of the model assumptions is deferred to Online Appendix A3.

3.1. The Buyer's Problem

When the buyer arrives, he observes the price function for the rest of the horizon and decides whether and when to buy in order to maximize his utility. We assume that she is forward-looking and sensitive to delay, and denote by $U(t, v)$ the quasilinear discounted utility function of a consumer with valuation v purchasing at time t . When $t = \infty$, we interpret it as a nonpurchase decision of the buyer, and we define $U(\infty, v) = 0$. In particular, we consider an exponentially discounted utility function: $U(t, v) = e^{-\mu t}(v - p(t))$, where $\mu > 0$ is the discount rate. This intertemporal utility function discounts the buyer's payoff from time zero, and it is without loss of generality (w.l.o.g.) for the sake of characterizing an optimal policy. That is, if the buyer purchasing at time t only incurs the disutility for waiting from his own arrival time τ , then the utility function $U(t, v)$ is only affected by a fixed multiplicative constant: $U(\tau, t, v) = e^{-\mu(t-\tau)}(v - p(t)) = e^{\mu\tau}U(t, v)$. Note that, in our definition, the discount rate affects both the valuation and the price paid, which is a standard assumption in OM (see, e.g., Swinney 2011, Caldentey et al. 2017, Papanastasiou and Savva 2017, Golrezaei et al. 2021).

Given a price function $p(t)$, a forward-looking buyer arriving at time τ with valuation v solves

$$[BP] \quad \max_{\tau \leq t \leq \infty} U(t, v).$$

Note that the maximum in this problem can be attained at $t = \infty$, and then, we consider that the buyer has an outside option with utility equal to zero. That is, a buyer with valuation v who arrives at time τ purchases (at a finite time) if and only if there exists $\tau_p \in [\tau, \infty)$ such that $U(\tau_p, v) \geq 0$ and does not purchase (or purchases at time $\tau_p = \infty$) otherwise, obtaining zero utility. We are then assuming, as is common in the mechanism design literature, individual rationality (or voluntary participation) of the buyer.

It may be possible that the buyer’s problem has multiple solutions, and to avoid ambiguity, we further assume for convenience that the buyer purchases the item at the earliest time maximizing his utility. Next, we introduce an auxiliary function, namely, ϕ , which, for any given purchasing time t , returns the minimum valuation the buyer must have in order to buy at time t and no later. In other words, the function ϕ defines a threshold in the sense that, if a buyer with valuation v buys at time t , then a buyer with valuation $v' > v$ buys no later than time t .² More formally, $\phi : [0, \infty) \rightarrow [0, \bar{v}] \cup \{\infty\}$ is defined as

$$\phi(t) = \inf\{v : U(t, v) \geq U(t', v), \forall t' \geq t\},$$

and we set $\phi(t) = \infty$ if there exists $t' > t$ such that $U(t, v) < U(t', v)$ for all $v \in [0, \bar{v}]$.

We also extend the domain of ϕ by setting $\phi(\infty)$ as the minimum valuation the buyer must have to buy at some time $t \in [0, \infty)$ (or, equivalently, the maximum valuation he must have to not buy). That is,

$$\phi(\infty) = \inf\{\phi(t) : t \in [0, \infty)\}.$$

Based on ϕ , we are able to describe the equilibrium conditions for the buyer’s purchasing behavior and use them to formulate the seller’s problem.

3.2. The Seller’s Problem

The seller’s problem is to post a price function to maximize her expected revenue, taking into account the forward-looking behavior of the buyer. Following a standard assumption in the OM literature (see, e.g., Cachon and Swinney 2011, Briceño-Arias et al. 2017, Aflaki et al. 2020, Golrezaei et al. 2021), particularly motivated by retail settings in which customers strategize over the timing of their purchase, we assume here that the seller is more patient than the buyer, and hence, the seller’s discount rate δ verifies $\delta < \mu$. This setup is also the interesting one to consider because the problem becomes easy when $\delta \geq \mu$ (see Online Appendix A3).

The seller’s problem can be stated based on her ability to observe arrivals and flexibility to set prices. In the observable case, she can choose a menu of pricing functions $\{p_\tau(t)\}_\tau$, indexed by τ and defined over $[\tau, \infty)$, so that a buyer arriving at time τ is shown the pricing function p_τ . In the unobservable case, she can only choose $p_0(t)$, and the customer faces that price curve irrespective of his arrival time. Both cases are formally presented herein.

3.2.1. Observable Arrival Case. Even though the seller can design a menu of pricing functions $\{p_\tau(t)\}_\tau$ and pick the pricing function $p_\tau(t)$ if the buyer arrives at time τ , for now, we pretend that the buyer arrives at time zero, that is, we initially assume $\tau = 0$. To simplify notation, we drop the index from $p_0(t)$. Observe that we can assume that $p(t)$ is nonincreasing. Otherwise, we could easily find an alternative nonincreasing pricing returning the same revenue to the seller.³

Given the threshold function ϕ induced by the price function p , a buyer with valuation v purchases at the first time $t \geq 0$ satisfying $v \geq \phi(t)$ and does not purchase if $v < \phi(\infty)$. The buyer’s purchasing behavior could be better represented by resorting to the auxiliary function $\psi(t)$ defined as

$$\psi(t) = \min\{\phi(s) : s \leq t\}.$$

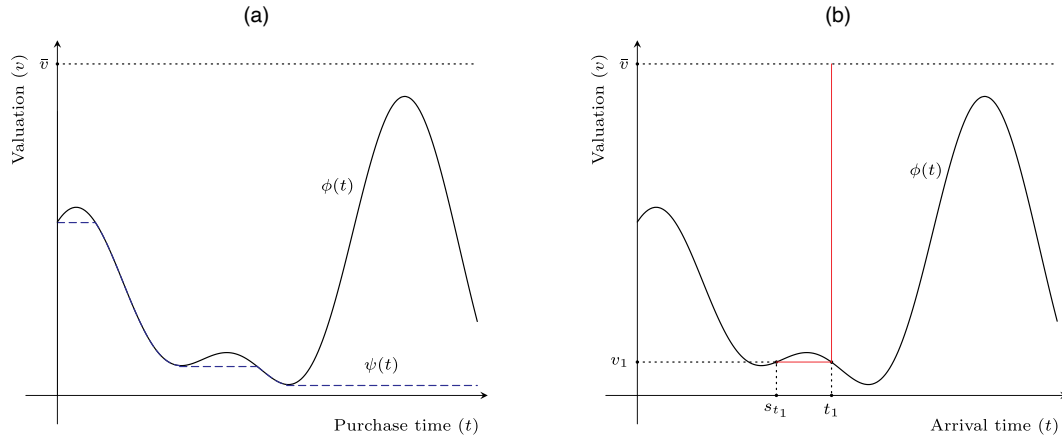
In other words, a customer arriving at $\tau = 0$ with valuation v buys at the first time t at which ψ takes the value v (or buys immediately if $v \geq \psi(0)$). Because of the lower semicontinuity of p , we have that ϕ is also lower semicontinuous, and therefore, ψ is well-defined (see Proposition A1 in Online Appendix A1 for a proof). The purchasing function $\psi(t)$ is the unique nonincreasing function that supports $\phi(t)$ from below (see Figure 1(a)). The instantaneous probability of selling at time t is given by $d(1 - F(\psi(t)))$. With this observation, we may write the seller’s problem conditioned on the event that the buyer arrives at time 0:

$$[SPO_0] \quad \max_{p, \psi} p(0)(1 - F(\psi(0))) + \int_0^\infty e^{-\delta t} p(t) d(1 - F(\psi(t))).$$

s.t. $t \in \operatorname{argmax}_{s \geq 0} U(s, \psi(t))$ for all $t \geq 0$.

The first term in the objective function stands for the event when the customer buys immediately at time 0, and the second term accounts for the customer’s forward-looking behavior. The incentive-compatible constraint specifies that a consumer arriving at time zero with valuation $\psi(t)$ maximizes utility at time $s = t$. We remark here that the individual rationality constraint is implicitly included in the equilibrium constraints because of the argmax being taken over $[0, \infty]$ and, by definition, $U(\infty, v) = 0$ for all valuation v .

Figure 1. (Color online) Consumer Purchasing Behavior



Notes. (a) Observable case: Definition of the function $\psi(t)$. For a given function $\phi(t)$, a customer with valuation $\psi(t)$ arriving at $\tau = 0$ buys at time t . (b) Unobservable case: Characterization of a buyer purchasing at time t_1 including the one arriving exactly at t_1 with valuation $v \geq v_1$ and those arriving between s_{t_1} and t_1 with valuation v_1 .

We now extend the seller's revenue optimization problem to the case when the buyer arrives at time $\tau > 0$. Let R_τ be the seller's maximum expected revenue conditioned on the event that the buyer arrives at time τ . This corresponds to shifting the seller's revenue from $\tau = 0$ to $\tau > 0$, that is, $R_\tau = e^{-\delta\tau}R_0$, with R_0 being the objective function value of problem [SPO₀]. Finally, the maximum expected revenue of the seller can be written as $R = \int_0^\infty R_\tau g(\tau) d\tau = R_0 \int_0^\infty e^{-\delta\tau} g(\tau) d\tau$ so that our assumption on writing the seller's problem when the customer arrives at time 0 is without loss of generality in terms of characterizing the structure of the optimal pricing policy.

The formulation [SPO₀] and the related expected revenue R allow us to make a clear connection to the two alternative interpretations of our model discussed in Section 1: (i) infinite supply and demand setup, for which there is a continuum of buyers with mass $d(1 - F(\psi(t)))$ who buy at time t , for a total mass of one over the infinite horizon, and (ii) price of discrimination, for which, here, the seller is indeed able to keep track of the arrival time τ of each of these buyers and post a personalized price curve p upon each arrival.

3.2.2. Unobservable Arrival Case. When the seller does not observe the buyer's arrival time, the price function $p_0(t)$ that she has to set from time 0 can only depend on the arrival time distribution G . To simplify notation, we also drop the index from p_0 .

Although it is possible to formulate the seller's problem without any assumption over the threshold function ϕ , it is necessary to be careful on how to express the seller's expected revenue when ϕ is not continuous. Thus, just for simplicity and because it does not affect the analysis in what follows, we describe the seller's problem under the assumption of p being continuous, which, in turn, implies ϕ being continuous.

Defining the point of time s_t as the last time previous to t when ϕ takes the same value as $\phi(t)$ (or $s_t = 0$ if such time does not exist; see Figure 1(b)), that is,

$$s_t = \sup\{l < t : \phi(l) = \phi(t)\} \vee 0,$$

the seller's problem can be described as follows:

$$[SPN] \max_{p, \phi} \int_0^\infty e^{-\delta t} p(t) [(1 - F(\phi(t)))g(t) + 1_{\{\phi'(t) \leq 0\}} (G(t) - G(s_t))d(1 - F(\phi(t)))] dt.$$

s.t. $t \in \arg \max_{s \geq t} U(s, \phi(t))$ for all t .

The term in brackets stands for the probability of purchasing between times t and $t + dt$. Within it, the first term $(1 - F(\phi(t)))g(t)$ represents the probability of arriving in that small time interval with valuation $v \geq \phi(t)$ and, hence, of purchasing immediately. This corresponds to the points in the vertical line at t_1 in Figure 1(b); that is, we account for a customer arriving at t_1 with valuation $v \geq v_1$. The second term, $(G(t) - G(s_t))d(1 - F(\phi(t)))$, is the probability of arriving during the interval $(s_t, t]$ with valuation between $\phi(t)$ and $\phi(t + dt)$. This is the probability of being in the line connecting $\phi(s_{t_1})$ and $\phi(t_1)$ for a valuation v_1 in Figure 1(b). Note that, if the buyer has arrived before t and is still present at t , the buyer does not buy if ϕ is increasing at t , and thus, this second term only holds at points at which ϕ is decreasing, which is captured by the indicator function. In both arrival situations, the discounted revenue for the seller is $e^{-\delta t}p(t)$.

For future reference, we denote R^{uo} as the objective function value of [SPN].

The formulation [SPN] allows us to revisit the connection with the two alternative model interpretations described in Section 1: (i) an infinite supply and demand

setup, in which, here, there is a continuum of infinitesimal buyers with point mass $(1 - F(\phi(t)))g(t) + 1_{\{\phi'(t) \leq 0\}}(G(t) - G(s_t))d(1 - F(\phi(t)))$ who buy at time t (and who arrive before or at t), and (ii) price of discrimination, for which, here, the model does not allow price discrimination because all buyers face the same price curve posted at time 0.

3.3. Value of Observability: Overview

We start this section with the formal definition of the value of observability, followed by an example of its calculation that also illustrates the associated challenges.

3.3.1. Definition. After describing the seller's problem in both the observable and unobservable cases, we can formally define the value of observability for a specific problem instance characterized by distributions F , G and discount rates δ and μ with $\delta < \mu$. Recalling that, for a particular problem instance, R is the objective function value of problem $[SPO_0]$ and R^{uo} is the objective function value of problem $[SPN]$, we define

$$VO(G, F, \delta, \mu) = \frac{R}{R^{uo}}.$$

Our objective is to provide a bound for the instance-independent value of observability, when $\delta < \mu$, which we denote by VO:

$$VO = \sup_{G, F, \delta, \mu} VO(G, F, \delta, \mu).$$

A key difficulty in evaluating the value of observability is that, as mentioned, the unobservable case is typically very hard to solve, and standard approaches based on optimal control to tackle dynamic pricing and mechanism design problems fail.

3.3.2. Preview Example. To better grasp this difficulty and the difference between the observable and unobservable cases, let us describe here a quick example. Take a buyer with valuation uniformly distributed in $[0, 1]$ and arrival time distributed as an exponential with mean one. Also assume the seller discount rate is one, whereas that of the buyer is extremely large (so that, in fact, the buyer behaves myopically: she buys as soon as the price is below the buyer's valuation).⁴ Then, if the seller can observe the buyer's arrival, she starts pricing at one and then suddenly decreases the price in a continuous fashion until hitting the customer valuation when the transaction is executed. In this way, the seller extracts all the consumer surplus with expected value $1/2$. Thus, in expectation, the seller gets $R = \int_0^\infty (e^{-t}/2)e^{-t} dt = 1/4$. Here, the first e^{-t} represents the discounting, and the second e^{-t} represents the density of the exponential.

On the other hand, in the unobservable case, if we assume that the seller needs to set a nonincreasing price function, then the problem is relatively easy to solve.

Indeed, the seller needs to maximize, over all nonincreasing functions p , the quantity $R^{uo} = \int_0^\infty [e^{-t}(1 - p(t)) - (1 - e^{-t})p'(t)]e^{-t}p(t)dt$. Note that, for a nonincreasing $p(t)$, trade occurs between t and $t + dt$ if either the buyer arrives in that interval and the buyer's valuation is above $p(t)$ (hence, the term $e^{-t}(1 - p(t))$) or the buyer arrives before t and the buyer's valuation is between $p(t)$ and $p(t + dt)$ (hence, the term $-(1 - e^{-t})p'(t)$). In both cases, the discounted revenue for the seller is $e^{-t}p(t)$. The solution of this problem turns out to be $p(t) = e^{-t}$, which results in an expected revenue of $1/6$. Overall, the ratio of the revenues between the observable and unobservable cases in this example and when restricting the seller to use a nonincreasing price function for the unobservable case is $3/2$.

One may think that the latter example implies that, in general, $VO \geq 3/2$. However, the seller's strategy space is richer than that of nonincreasing price functions. Whereas we argue that, for the observable case, the optimal price curve is nonincreasing, for the unobservable case, this characterization is unclear. Suppose that the seller now splits the time horizon into short intervals of length ϵ and posts a periodic price function that sets price 1 for the first $\epsilon - \epsilon^2$ time units of each interval and a quickly decreasing price (from one to zero) in the last ϵ^2 time units of each interval. As the buyer is myopic, he buys at the first point in time in which the price is below his valuation, and because ϵ is very small, the probability that the buyer arrives when the price is $p(t) = 1$ is close to one. Thus, even in the unobservable case, the seller is able to obtain a revenue arbitrarily close to $1/4$ —higher than the one under the decreasing pricing policy—bringing the value of observability down to one. This observation illustrates the difficulty of obtaining a general upper bound for VO that is independent of the instance-specific parameters of the problem.

More generally, and beyond the uniform valuation distribution, it can easily be shown that the worst VO is not attained when the buyer's discount rate μ is in an extreme. Indeed, when $\mu = \delta$ (or, even more generally, when $\mu \leq \delta$), the seller cannot extract extra revenue by using any type of dynamic screening, and the optimal pricing policy is simply to fix a constant price equal to the monopoly price. Therefore, the VO equals one. On the other side of the spectrum, when $\mu \rightarrow \infty$, the buyer essentially behaves myopically. Similar to the preceding example, the seller may split the time horizon into short intervals of length ϵ . She then posts a periodic price function that sets a very high price (say, equal to \bar{v} , the largest possible valuation) for the first $\epsilon - \epsilon^2$ time units of each interval and a continuously (and quickly) decreasing to zero price function in the last ϵ^2 time units of each interval. In the limit as $\epsilon \rightarrow 0$, the buyer is myopic, so the buyer buys at the first point in time in which the price is below the buyer's valuation. Because the probability

that the buyer arrives when the price is \bar{v} is close to one, even in the unobservable case, the seller is able to obtain a revenue arbitrarily close to the buyer's valuation. This implies that the value of observability approaches one when μ grows large, and therefore, the most interesting cases occur when μ is in the middle of the range. However, finding this worst possible μ in terms of VO seems to be a very challenging problem.

4. Analysis of the Model with an Observable Arrival

In this section, we study in detail the observable arrival case. We start by deriving some structural properties of the solution to this problem, spanning both the optimal price and purchasing functions. Then, we present a technical result that provides a guarantee for a fraction of the revenue to be attainable over a finite time window.

4.1. Structural Characterization of the Optimal Solution

Given the argument stated in Section 3.2.1, to analyze the observable case, it is sufficient to focus on the solution of $[SPO_0]$, in which the buyer arrives at time 0.

Problem $[SPO_0]$ is difficult to solve because of its equilibrium constraint. Our approach is to formulate a relaxed version of the problem by computing the first order condition of the equilibrium constraint. Then, by applying the Euler–Lagrange equation, we show that any solution of the relaxed problem also solves $[SPO_0]$. Moreover, we provide a characterization of the optimal price function as a solution of an ordinary differential equation, which turns out to have a unique solution for a large set of valuation distributions, and furthermore, it can be solved explicitly at least for F being a uniform distribution.

To begin, consider the incentive-compatible constraint in problem $[SPO_0]$. If the optimal solution to the optimization problem in the constraint of problem $[SPO_0]$, namely, t^* , is in the interior of the feasible region, then it must satisfy the first order condition $h(t)=0$, where $h(s) = U_s(s, \psi(t))$ or, equivalently, $\psi(t) = p(t) - \frac{p'(t)}{\mu}$. Now, consider the relaxed formulation:

$$\begin{aligned}
 [SPO_0^r] \max_{p, \psi} & \int_0^\infty e^{-\delta t} p(t) d(1 - F(\psi(t))) + p(0)(1 - F(\psi(0))) \\
 \text{s.t. } & \psi(t) = p(t) - \frac{p'(t)}{\mu}, \quad \forall t \geq 0.
 \end{aligned} \tag{1}$$

The feasible region of this constrained problem is larger than the one of $[SPO_0]$, and therefore, the objective function value of $[SPO_0^r]$ provides an upper bound of $[SPO_0]$.

Note that the problem $[SPO_0^r]$ can be written as the following unconstrained maximization problem on the

price function $p(t)$:

$$\begin{aligned}
 \max_p \int_0^\infty e^{-\delta t} p(t) & \left(-p'(t) + \frac{p''(t)}{\mu} \right) f \left(p(t) - \frac{p'(t)}{\mu} \right) dt \\
 & + p(0) \left(1 - F \left(p(0) - \frac{p'(0)}{\mu} \right) \right).
 \end{aligned} \tag{2}$$

Letting the integrand function be $I(t, p(t), p'(t), p''(t))$ and the expected revenue at time zero be r_0 , problem (2) is equivalent to

$$\max_p \int_0^\infty I(t, p(t), p'(t), p''(t)) dt + r_0.$$

Focusing on the first term, the associated Euler–Lagrange equation that must be satisfied by an optimal price function $p(t)$ states that

$$\frac{d^2}{dt^2} \frac{\partial I}{\partial p''} - \frac{d}{dt} \frac{\partial I}{\partial p'} + \frac{\partial I}{\partial p} = 0.$$

Such function $p(t)$ is a stationary point of the functional

$$\int_0^\infty I(t, p(t), p'(t), p''(t)) dt.$$

After some algebra (detailed in Proposition A2 in Online Appendix A1), the Euler–Lagrange equation becomes

$$\begin{aligned}
 f' \left(p(t) - \frac{p'(t)}{\mu} \right) & \left(-\frac{p''(t)}{\mu} + p'(t) \right) (-\delta p(t) + p'(t)) \\
 & + f \left(p(t) - \frac{p'(t)}{\mu} \right) [\delta(\delta - \mu)p(t) - 2\delta p'(t) + 2p''(t)] = 0.
 \end{aligned} \tag{3}$$

This equation can be written as a system of two first order differential equations by defining the auxiliary variable $u(t) = p'(t)$. Thus, by standard results on ordinary differential equations (see, e.g., theorem 54.A in Simmons 2016) we can show that there exists one and only one solution to the initial value problem given $p(0)$ and $p'(0)$ under mild continuity and differentiability conditions, which are satisfied for a large set of valuation distributions. To illustrate this, in Online Appendix A4, we solve the problem for the special case in which the valuation is uniformly distributed in $[0, 1]$.

Let us highlight that, though we know $\psi(t)$ is nonincreasing by construction—and, indeed, we use this fact to formulate the seller's problem— $[SPO_0^r]$ could potentially have an optimal solution with a generic function $\psi(t)$ not meeting this monotonicity. However, the following result establishes that this does not happen. In other words, if $\psi(t)$ corresponds to an optimal solution of the seller's relaxed problem, then it must be a nonincreasing function.

Proposition 1. Assume that the density function f is strictly positive. If the price function $p(t)$ is a continuously differentiable optimal solution of the relaxed problem $[SPO_0^r]$, then the optimal purchasing function $\psi(t) = p(t) - \frac{p'(t)}{\mu}$ is nonincreasing.

The nonincreasing structure of the purchasing function $\psi(t)$ for the relaxed problem along with the upper bound defined by the solution to $[SPO_0^r]$ allow us to show the following result:

Proposition 2. Any solution of the relaxed problem $[SPO_0^r]$ such that p is differentiable with continuous derivative, also solves the seller’s problem $[SPO_0]$.

This result allows us to simplify the solution of the seller’s problem $[SPO_0]$. Furthermore, we show that the solution of the relaxed problem is a solution of an autonomous system of ordinary differential equations.

Thus, to solve the seller’s problem $[SPO_0]$, we first formulate the Euler–Lagrange Equation (3) associated with the relaxed problem $[SPO_0^r]$ and solve it. Its solution depends on the initial values $p(0) > 0$ and $p'(0) < 0$. Then, we replace that solution form in problem (2) and solve it in terms of the scalar variables $p(0)$ and $p'(0)$. Finally, using these optimal initial values, we can recover the optimal price function $p(t)$ and purchasing function $\psi(t)$, which are the optimal solutions of the original seller’s problem $[SPO_0]$.

4.2. Bounding a Fraction of Revenue over Time

In preparation to bound VO in Section 6, we present here the following technical result: for a given parameter $c \in (0, 1)$, if we need to ensure that the seller earns a fraction $(1 - c)$ of the expected revenue in problem $[SPO_0]$, it is enough to look at the problem until time $T = \ln(1/c)/\delta$. For instance, if we want to reach at least half of R_0 and we normalize the seller’s discount rate to one, we conclude that it is enough to consider the problem until time $T = \ln(2)$. This implies that the time needed to get a big fraction of R_0 is relatively short, and moreover, it does not depend on the valuation distribution. The result is formally stated as follows.

Lemma 1. For a given parameter $c \in (0, 1)$, up to time $T = \ln(1/c)/\delta$, the seller’s expected revenue $R_{[0,T]}$ in the observable arrival case is at least $(1 - c)R_0$; that is,

$$R_{[0,T]} = p(0)(1 - F(\psi(0))) + \int_0^T e^{-\delta t} p(t) d(1 - F(\psi(t))) \geq (1 - c)R_0,$$

where $p(t)$ is the solution from (3) to the observable case problem.

This result is used to construct the feasible pricing policy for the unobservable case that we present in Section 5,

becoming a key ingredient to bound the value of observability.

5. Analysis of the Model with an Unobservable Arrival

Consider now the problem stated in Section 3.2.2 in which the seller is not able to observe the arrival time of the buyer. Different from the previous observable case, in which the seller knows the arrival time τ of the buyer and, resorting to the menu $\{p_\tau(t)\}_\tau$ of price curves, posts $p_\tau(t)$ over the horizon $[\tau, \infty)$ (even though, as explained before, the analysis was conducted without loss of generality by assuming $\tau = 0$), in this case, the seller commits to a price function at time 0 without knowing the precise buyer’s arrival time.


This problem turns out to be notoriously hard in the general case, and obtaining an explicit solution seems hopeless. To partially overcome this, we focus on analyzing the seller’s problem under a feasible (and suboptimal) pricing policy with the objective of bounding the value of observability later.

The feasible pricing policy \hat{p} we consider is periodic and depends on the policy p that solves the observable case formulation, $[SPO_0]$. The length of the period is $2T$, where T is such that, until time T , the seller’s expected revenue in the observable case when the buyer arrives at time 0 is of significant magnitude. In particular, the price function we use to bound from below the seller’s expected revenue in the unobservable case is defined by

$$\hat{p}(t) = \begin{cases} p(0) & \text{if } t \in I_{2k-1}, k \in \mathbb{N} \\ p(t - (2k - 1)T) & \text{if } t \in I_{2k}, k \in \mathbb{N}, \end{cases} \quad (4)$$

where $I_{2k-1} = (2(k - 1)T, (2k - 1)T]$ and $I_{2k} = ((2k - 1)T, 2kT]$ for $k \in \mathbb{N}$ and the price function p comes from the solution of $[SPO_0]$. Note that the function \hat{p} is continuous at the points kT for odd values of k . Figure 2 shows the structure of the periodic pricing policy we consider

Figure 2. Periodic Pricing Policy \hat{p} After Performing a Random Shift and Setting the Origin at Time t_0



along the rest of the section with its origin at a value $t_0 \geq 0$.

We subsequently explain the reason for shifting the origin by t_0 units of time. Here, we provide a discussion for the intuition behind this feasible pricing policy. In fact, its structure is motivated by several concurrent factors. First, in view of the definition of value of observability introduced in Section 3.3, we are seeking an expected revenue R^{uo} for the unobservable case that is a constant fraction of the expected revenue R of the observable case. The optimal pricing policy of the observable case is not periodic and spans the whole infinite horizon, but the result in Lemma 1 is helpful in providing a bound for the revenue over a limited interval of length T , where T is the threshold defined therein that warrants a “big enough” revenue. The pricing policy borrowed from the observable case is embedded in the even intervals of the proposed periodic policy.

The next question is why it would not be sufficient then to just replicate this policy with period T again and again. The answer is rooted in the buyer’s purchasing behavior: the buyer is strategic and would not face a similar curve as in the observable case; hence, the buyer’s behavior would be different. This argument sustains the setting of the high price $p(0)$ in the odd intervals of length T : it allows the accumulation of a mass of buyers before launching the observable case price policy. Even though some of the mass of these buyers could buy before time T (i.e., the ones with valuation higher than $\psi(0)$), we assume that they still wait until time T to execute the purchase. This extra delay further reduces the seller’s revenue, which is acceptable for the sake of computing a lower bound. Indeed, the buyers arriving during the odd intervals behave similarly to the buyers in the observable case except for the fact that they wait a bit extra.

In summary, within a period of length $2T$, say, between t_0 and $t_0 + 2T$, the mass of buyers with valuation below $p(T)$ is priced out of the market, and for those with valuation above $p(T)$, we only account for the ones arriving between t_0 and $t_0 + T$. Note that, in this scenario, a buyer with valuation between $p(T)$ and $\psi(T)$ who would buy after time T in the regular observable case, here would also buy at time T , pay $p(T)$, and get a positive utility (higher than the zero utility of the no purchase). Those with valuation above $\psi(T)$ are the ones playing the strategy of the observable case except for the small extra wait until $t_0 + T$.

Coming back to the fact of having a time origin set at t_0 , it is justified as follows. One element that makes the unobservable arrival case particularly challenging to analyze from a revenue computation perspective is the presence of the density $g(t)$ in the formulation of [SPN]. In order to perform the analysis independently of the specific function g , we use the next proposition, stating

that, by doing a random shift on the price function \hat{p} defined in (4), we can assume without loss of generality that the buyer’s arrival time is uniformly distributed within a period of length $2T$.

Proposition 3. *Let us consider the function \hat{p}_{t_0} obtained by performing a random shift over the function \hat{p} defined in (4), that is, for a random variable $t_0 \sim \text{Unif}[0, 2T]$, consider the function $\hat{p}_{t_0}(t) = \hat{p}(t + t_0)$. It holds that the buyer’s arrival, conditional in the arrival interval, is $\text{Unif}[0, 2T]$.*

Therefore, by applying the preceding proposition, we can assume that the buyer’s arrival time, conditional on the fact that the arrival belongs to a specific interval, is $\text{Unif}[0, 2T]$ and the function’s new origin is t_0 ; that is, t_0 is the starting point of a period of length $2T$.

Along the rest of the paper, we denote as \hat{p} the feasible pricing policy after performing the random shift. We also relabel the intervals of the function \hat{p} and denote by \tilde{I}_{2k-1} the range in which \hat{p} is constant and denote by \tilde{I}_{2k} the range in which \hat{p} is the translation of the function p after performing the random shift.

For illustration purposes and as a supplement to this section, in Online Appendix A5, we compare the performance of our heuristic policy to that of a fixed price, showing the advantage of the former when the buyer is moderately to significantly more impatient than the seller. Then, in Online Appendix A6, we characterize the optimal solution of a particular unobservable arrival problem instance with a two-point arrival time distribution and in which the valuation of the buyer is truncated Pareto. Despite the unobservable arrival case being hard to solve in general, we can fully solve this particular problem instance by simplifying its formulation to a sequence of observable arrival cases.

6. Bounding the Value of Observability

Recall that following the definition in Section 3.3, the value of observability $VO(G, F, \delta, \mu)$ is the instance-specific ratio between the expected revenue in the corresponding observable and unobservable cases. Accordingly, the VO is defined as the supremum of instance-specific parameters: $VO = \sup_{G, F, \delta, \mu} VO(G, F, \delta, \mu)$. In what follows, we first study VO under valuation distributions with MHR and show that it can be upper bounded by $\exp(1)$ and the bound is tight within the space of fixed price policies for the unobservable case. Next, we analyze the more challenging case beyond MHR valuations that may arise when fitting real data. For this very general case, we characterize the main result of our paper: the constant 4.911 as an upper bound for VO. In Online Appendix A7, we argue by resorting to an example that VO is lower bounded away from the trivial one, obtaining 1.136 based on a two-point distribution for the arrival and a truncated Pareto distribution for the valuations.

6.1. An Upper Bound for an MHR Valuation Distribution

The case of MHR valuation distribution turns out to be relatively simple. For completeness, we review here some basic concepts on the theory of optimal auctions introduced in the seminal work of Myerson (1981). A key building block to state the seller’s optimal expected revenue in a general single unit auction is the so-called virtual value of the bidder, defined as

$$J(v) := v - \frac{1 - F(v)}{f(v)} = v - \frac{1}{\rho(v)},$$

where $\rho(v) = f(v)/(1 - F(v))$ is the hazard rate function associated with the distribution F . The value $J(v)$ represents the expected value of the revenue that the seller may intend to collect from a bidder with valuation v , which naturally verifies $v > J(v)$. Alternatively, when considering the static price-optimization problem of a seller trying to maximize the revenue function $r(p) = p(1 - F(p))$, the first order condition states that $J(p) = 0$. In other words, $J(p)$ stands for the marginal revenue function. As a consequence, an optimal monopoly reserve price p^* is defined as $p^* = J^{-1}(0)$.⁵

A distribution F is said to be regular if the virtual value function $J(v)$ is strictly increasing in v . This assumption is not overly restrictive and is satisfied by distributions with increasing hazard rate $\rho(v)$, including standard ones, such as the normal, uniform, logistic, exponential, and extreme value distributions.

In what follows, we assume that the buyer’s valuation is distributed according to a monotone (increasing) hazard rate distribution F and prove that the value of observability is upper bounded by e . Moreover, this bound is tight if we restrict the space of feasible policies to the set of fixed price policies for the unobservable case.

Indeed, we know from Section 3.2.1 that the optimal seller’s expected revenue in the observable case is given by $R = R_0 \int_0^\infty e^{-\delta t} g(t) dt$, where R_0 is the objective function value of problem $[SPO_0]$ and, therefore, verifies $R_0 \leq \mathbb{E}(v)$, the expected value of the valuation drawn from F . Hence, the seller’s expected revenue in the observable case verifies

$$R \leq \mathbb{E}(v) \int_0^\infty e^{-\delta t} g(t) dt.$$

For the unobservable case, consider the feasible, fixed pricing policy $p(t) = p^*$ for all t , where $p^* = J^{-1}(0)$ is the optimal monopoly price. Then, the seller’s expected revenue verifies

$$R^{uo} \geq \int_0^\infty e^{-\delta t} p^*(1 - F(p^*)) g(t) dt = p^*(1 - F(p^*)) \int_0^\infty e^{-\delta t} g(t) dt.$$

From this lower bound, we can now establish an upper bound for the value of observability. By lemma 3.10 of Dhangwatnotai et al. (2015), it follows that $p^*(1 - F(p^*)) \geq \frac{1}{e} \mathbb{E}(v)$, and thus, $VO(G, F, \delta, \mu) \leq e$, when F is regular.

This upper bound is tight if, in the unobservable case, we restrict the feasible space of price functions to the set of fixed price strategies. To see this, consider a myopic buyer (i.e., a buyer with discount rate μ extremely large so that he buys as soon as the price drops below his valuation) arriving according to a general distribution G and with valuation distributed according to the truncated exponential random variable with parameter one and support $[0, M]$, that is, with cdf $F(x) = (1 - e^{-x})/(1 - e^{-M})$. Assume that the seller does not discount revenues (i.e., $\delta = 0$). In this setting, in the observable case, the seller announces a price curve for a consumer arriving at time τ that spans all the valuation support (e.g., $p_\tau(t) = Me^{-(t-\tau)}$), and the consumer buys immediately when his valuation v verifies $v = p(t)$ with expected revenue for the seller $R = \mathbb{E}(v) = (1 - (M - 1)e^{-M})/(1 - e^{-M})$. In the unobservable case, if the seller offers the best possible fixed price, she sets a price p^* maximizing $\bar{R}^{uo}(p) = p(1 - F(p))$. This function is maximized at $p^* = 1 - W(e^{1-M})$, where W denotes the positive branch of the well-known Lambert function,⁶ obtaining a revenue \bar{R}^{uo} given by

$$\bar{R}^{uo} = \frac{W^{-1}(e^{1-M}) + W(e^{1-M})}{(1 - e^{-M})e^M}.$$

That leads to a gap between the observable and unobservable case that converges to e when M grows large, obtaining the tightness.

We highlight here that this argument does not imply that VO is lower bounded by e because, for the unobservable case, we are considering a feasible but not necessarily optimal strategy. What we prove is that, by considering a fixed price policy in the unobservable case, the upper bound of e can be attained.

6.2. An Upper Bound for a General Valuation Distribution

For a general valuation distribution (with a nonmonotone hazard rate), we start by noting that using a fixed price policy in the unobservable arrival case does not necessarily lead to an upper bound defined by a constant. For instance, consider the game in which the buyer’s valuation is distributed according to a truncated Pareto distribution with parameter one and support $[1, M]$, denoted by $\text{TruncPareto}(1, 1, M)$. That is, $v \sim \text{TruncPareto}(1, 1, M)$ has cdf

$$F(x) = \left(1 - \frac{1}{x}\right) \frac{M}{M-1}, \text{ for } x \in [1, M].$$

We further assume that μ is extremely large and $\delta = 0$. Following the earlier argument, in the observable case,

the seller announces a price curve that spans all the valuation support (e.g., $p(t) = 1/t$), and the consumer buys immediately when his valuation v verifies $v = p(t)$, with expected revenue for the seller equal to $R = \mathbb{E}(v) = M \ln M / (M - 1)$. The revenue for the unobservable case under a fixed price p^* is $R^{uo} = p^*(1 - F(p^*)) < M / (M - 1)$, leading to the ratio $VO(G, \text{TruncPareto}(1, 1, M), 0, \mu) = \mathbb{E}(v) / (p^*(1 - F(p^*))) > \ln M$, which grows with M unboundedly and independently of the arrival distribution.

Despite the difficulty imposed by the ineffectiveness of a fixed price policy, we are able to bound from above the value of observability by resorting to the particular pricing policy \hat{p} introduced in Section 5 (and defined in (4)).

Before presenting our main result, we need to introduce two preliminary lemmas. In the first one, we provide a simple lower bound for the seller's revenue within a limited time frame in the unobservable arrival case. In the second one, we give a lower bound for R_τ^{uo} , which we define as the seller's expected revenue in the unobservable case when the buyer's arrival time is τ . To characterize this bound, we focus on arrivals during the odd intervals.

Lemma 2. *If the buyer is present at time τ being the beginning of a period \tilde{I}_{2k} for some $k \in \mathbb{N}$, and has valuation $v \geq p(T)$, then the seller's expected revenue by offering the price function \hat{p} in the unobservable case is at least the expected revenue earned up to time $2kT + t_0$ in the observable case, with arrival time $(2k - 1)T + t_0$.*

Lemma 3. *For a given parameter $c \in (0, 1)$, and denoting by R_τ the expected revenue in the observable case when the buyer arrives at time τ , it holds that*

$$R_\tau^{uo} \geq \frac{(1 - c)^2}{2 \ln(1/c)} R_\tau.$$

We are now ready to provide an upper bound for the value of observability that can be written as a function of W_{-1} , the negative branch of the Lambert function.

Theorem 1. *For any valuation distribution and arrival time distribution, the value of observability is at most*

$$-\frac{2W_{-1}(-1/(2\sqrt{e}))+1}{(e^{W_{-1}(-1/(2\sqrt{e}))+1/2}-1)^2} \approx 4.911.$$

Proof. Recalling that R_τ and R_τ^{uo} are the expected values of the seller's revenue in the observable and unobservable cases, respectively, when the buyer's arrival time is τ , it follows that, for each arrival time τ , the ratio between them verifies

$$\frac{R_\tau}{R_\tau^{uo}} \leq \frac{R_\tau}{\frac{(1-c)^2}{2 \ln(1/c)} R_\tau} = \frac{2 \ln(1/c)}{(1-c)^2}.$$

This bound ratio holds for any constant $c \in (0, 1)$, and it is minimized at $c = e^{W_{-1}(-1/(2\sqrt{e}))+1/2} \approx 0.284$, giving

a minimum of $-\frac{2W_{-1}(-1/(2\sqrt{e}))+1}{(e^{W_{-1}(-1/(2\sqrt{e}))+1/2}-1)^2}$, which is roughly 4.911. \square

It is worth noting that our result is robust to the specification of the arrival distribution. That is, even in the case in which the arrival time of the buyer is adversarial (i.e., the arrival distribution is chosen by an adversary that knows the price function of the seller but does not know the realization of the random shift), we prove that the seller's expected revenue in the observable case is at most 4.911 times his expected revenue if she does not observe the buyer's arrival.⁷

7. Conclusions

In this paper, we revisit a standard formulation for the dynamic pricing problem when a monopolistic seller faces the arrival of a single buyer over an infinite time horizon and precommits to a price curve with the objective of maximizing revenue. The buyer observes the price curve and strategically purchases at a time when his utility is maximized. Both players discount the future at different rates with the buyer being more impatient than the seller in the most realistic and technically challenging situation. The model feature that we analyze in this paper is the ability of the seller to observe (or not) the arrival time of the buyer.

We define the VO as the worst case ratio between the revenue attainable in the observable and unobservable cases taken over all model parameters, namely, the distribution of the arrival time of the buyer, the distribution of the buyer's valuation, and the discount rates. Our main result is that VO can be upper bounded by the constant 4.911 irrespective of the model primitives. In the particular and relevant case of monotone hazard rate valuation distribution, the upper bound can be improved to $e \approx 2.718$. We also show by an example that a lower bound for VO can be set at 1.136.

It is worth pointing out again that VO can be interpreted as the price of discrimination, that is, the additional rent that the seller can obtain from being able to post a customized price curve for (infinitesimal) buyers as opposed to setting the same price curve for everyone.

We highlight that, because our upper bound of 4.911 for the general valuation distribution setup relies on the implementation of a modified optimal policy for the observable case as a feasible policy for the unobservable case, and knowing that the upper bound for VO in the monotone hazard rate valuation distribution is 2.718, there is room for a potential improvement of the former bound. Closing the gap between them or showing that our bound is tight would be interesting venues for further research.

Our analysis also carries important managerial insights by characterizing business contexts in which gathering information about the buyer arrival time is particularly

valuable. We observe that it is not very significant when the consumer's valuation distribution has a monotone hazard rate, and it is even less important in the extreme cases in which the seller has a level of patience similar to the customer or is much more patient than the customer.

Finally, as future work, it would be interesting to conduct a case study (using synthetic or real data) demonstrating the value of tracking customer arrivals in different contexts, for example, an online seller facing robots searching for competitors' prices.

Acknowledgments

The authors thank the review team for many comments and for raising several issues that substantially improved the presentation of the paper.

Endnotes

¹ Because of this assumption, the seller could potentially lose at most negligible extra revenue, and therefore, it does not affect our results. Moreover, the lower semicontinuity is necessary to ensure that the buyer's problem can always be solved.

² To see this, knowing that $v = \phi(t)$, we have $U(t, v) \geq U(t', v)$, $\forall t' \geq t$. Now, consider a buyer with valuation $v' = v + \epsilon$, $\epsilon > 0$. By simple algebra, we have $U(t, v') = U(t, v) + \epsilon e^{-\mu t} > U(t', v) + \epsilon e^{-\mu t'} = U(t', v')$, that is, $U(t, v') \geq U(t', v')$, $\forall t' \geq t$. Thus, the purchasing time of buyer v' cannot be later than t .

³ If $p(t)$ is an arbitrary pricing function, we could consider $\tilde{p}(t)$ the largest nonincreasing function that is below p . Because the buyer is forward-looking, the possible purchasing times coincide under p and \tilde{p} .

⁴ Readers should not confuse the notion of a myopic buyer with that of an impatient buyer in the sense of waiting only an infinitesimal amount of time before purchasing. The notion of a myopic buyer we use throughout the paper is the one in which the consumer buys as soon as the price is below the buyer's valuation. It may then happen that the buyer waits a long time before buying or even does not buy at all.

⁵ More generally, the optimal reserve price is defined as $p^* = \max\{v : J(v) = 0\}$, and by convention, $p^* = \infty$ if $J(v) < 0$ for all v .

⁶ The Lambert W function is defined as the multivalued function that satisfies $z = W(z)\exp(W(z))$ for any complex number z . If x is real, then for $-1/e \leq x < 0$, there are two possible real values of $W(x)$. We denote the branch satisfying $-1 \leq W(x)$ by $W_0(x)$ —namely, the principal branch—and the branch satisfying $W(x) \leq -1$ by $W_{-1}(x)$, referred to as the negative branch.

⁷ Note that this notion of robustness is different from the common one in the pricing literature (e.g., Caldentey et al. 2017), in which it usually indicates that the price function does not depend on any of the parameters of the model. In our case, the feasible pricing policy we consider in the unobservable case depends on the optimal price function of the observable case, and therefore, it depends on the valuation distribution and discount rates. However, it does not depend on the arrival distribution.

References

- Aflaki A, Feldman P, Swinney R (2020) Becoming strategic: Endogenous consumer time preferences and multiperiod pricing. *Oper. Res.* 68(4):1116–1131.
- Araman V, Fayad B (2021) Intertemporal price discrimination with time-varying valuations. *Oper. Res.* 69(1):245–265.
- Aviv Y, Pazgal A (2008) Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing Service Oper. Management* 10(3):339–359.
- Aviv Y, Vulcano G (2012) Dynamic list pricing. Ozer O, Phillips R, eds. *The Oxford Handbook of Pricing Management*, Oxford Handbooks (Oxford University Press, Oxford, UK), 522–584.
- Aviv Y, Levin Y, Nediak M (2009) Counteracting strategic consumer behavior in dynamic pricing systems. Netessine S, Tang C, eds. *Consumer-Driven Demand and Operations Management Models*, International Series in Operations Research & Management Science, vol. 131 (Springer, New York), 323–352.
- Besbes O, Lobel I (2015) Intertemporal price discrimination: Structure and computation of optimal policies. *Management Sci.* 61(1):92–110.
- Blumrosen L, Holenstein T (2008) Posted prices vs. negotiations: An asymptotic analysis. *Proc. Ninth ACM Conf. Electronic Commerce* (ACM, New York), 49.
- Board S, Skrzypacz A (2016) Revenue management with forward-looking buyers. *J. Political Econom.* 124(4):1046–1087.
- Borgs C, Candogan O, Chayes J, Lobel I, Nazerzadeh H (2014) Optimal multiperiod pricing with service guarantees. *Management Sci.* 60(7):1792–1811.
- Briceño-Arias L, Correa J, Perloth A (2017) Optimal continuous pricing with strategic consumers. *Management Sci.* 63(8):2741–2755.
- Cachon GP, Swinney R (2009) Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Sci.* 55(3):497–511.
- Cachon GP, Swinney R (2011) The value of fast fashion: Quick response, enhanced design, and strategic consumer behavior. *Management Sci.* 57(4):778–795.
- Caldentey R, Liu Y, Lobel I (2017) Intertemporal pricing under minimax regret. *Oper. Res.* 65(1):104–129.
- Chawla S, Hartline JD, Malec DL, Sivan B (2010) Multi-parameter mechanism design and sequential posted pricing. *Proc. 42nd ACM Sympos. Theory Comput.* (ACM, New York), 311–320.
- Chen Y, Farias V (2018) Robust dynamic pricing with strategic customers. *Math. Oper. Res.* 43(4):1119–1142.
- Chen Y, Farias V, Trichakis N (2019) On the efficacy of static prices for revenue management in the face of strategic customers. *Management Sci.* 65(12):5535–5555.
- Conlisk J, Gerstner E, Sobel J (1984) Cyclic pricing by a durable goods monopolist. *Quart. J. Econom.* 99(3):489–505.
- Correa J, Montoya R, Thraves C (2016) Contingent preannounced pricing policies with strategic consumers. *Oper. Res.* 64(1):251–272.
- Correa J, Foncea P, Pizarro D, Verdugo V (2019) From pricing to prophets, and back! *Oper. Res. Lett.* 47(1):25–29.
- Dhangwatnotai P, Roughgarden T, Yan Q (2015) Revenue maximization with a single sample. *Games Econom. Behav.* 91:318–333.
- Elmachtoub AN, Gupta V, Hamilton ML (2021) The value of personalized pricing. *Management Sci.* 67(10):6055–6070.
- Elmaghraby W, Lippman S, Tang C, Yin R (2009) Will more purchasing options benefit customers? *Production Oper. Management* 18(4):381–401.
- Gallego G, Topaloglu H (2019) *Revenue Management and Pricing Analytics*, vol. 209 (Springer, New York).
- Gershkov A, Moldovanu B, Strack P (2018) Revenue-maximizing mechanisms with strategic customers and unknown, Markovian demand. *Management Sci.* 64(5):2031–2046.
- Golrezaei N, Javanmard A, Mirrokni V (2021) Dynamic incentive-aware learning: Robust pricing in contextual auctions. *Oper. Res.* 69(1):297–314.
- Kessel K, Saberi A, Shameli A, Wajc D (2022) The stationary prophet inequality problem. *Proc. 23rd ACM Conf. Electronic Commerce* (ACM, New York), 243–244.
- Lai G, Debo L, Sycara K (2010) Buy now and match later: Impact of posterior price matching on profit with strategic consumers. *Manufacturing Service Oper. Management* 12(1):33–55.

- Landsberger M, Meilijson I (1985) Intertemporal price discrimination and sales strategy under incomplete information. *RAND J. Econom.* 16(3):424–430.
- Li J, Granados N, Netessine S (2014) Are consumers strategic? Structural estimation from the air-travel industry. *Management Sci.* 60(9):2114–2137.
- Liu Q, Van Ryzin GJ (2008) Strategic capacity rationing to induce early purchases. *Management Sci.* 54(6):1115–1131.
- Liu Y, Cooper W (2015) Optimal dynamic pricing with patient customers. *Oper. Res.* 63(6):1307–1319.
- Lobel I (2020) Dynamic pricing with heterogeneous patience levels. *Oper. Res.* 68(4):1038–1046.
- Mersereau A, Zhang D (2012) Markdown pricing with unknown fraction of strategic consumers. *Manufacturing Service Oper. Management* 14(3):355–370.
- Myerson RB (1981) Optimal auction design. *Math. Oper. Res.* 6(1):58–73.
- Osadchiy N, Vulcano G (2010) Selling with binding reservations in the presence of strategic consumers. *Management Sci.* 56(12):2173–2190.
- Papanastasiou Y, Savva N (2017) Dynamic pricing in the presence of social learning and strategic consumers. *Management Sci.* 63(4):919–939.
- Pavan A, Segal I, Toikka J (2014) Dynamic mechanism design: A Myersonian approach. *Econometrica* 82(2):601–653.
- Simmons GF (2016) *Differential Equations with Applications and Historical Notes* (CRC Press, Boca Raton, FL).
- Stokey NL (1979) Intertemporal price discrimination. *Quart. J. Econom.* 93(3):355–371.
- Su X (2007) Intertemporal pricing with strategic customer behavior. *Management Sci.* 53(5):726–741.
- Su X, Zhang F (2008) Strategic customer behavior, commitment and supply chain performance. *Management Sci.* 54(10):1759–1773.
- Swinney R (2011) Selling to strategic consumers when product value is uncertain: The value of matching supply and demand. *Management Sci.* 57(10):1737–1751.
- Talluri K, van Ryzin G (2004) *The Theory and Practice of Revenue Management*, International Series in Operations Research & Management Science, Book 68 (Springer, New York).
- Wang R (2001) Optimal pricing strategy for durable-goods monopoly. *J. Econom. Dynamic Control* 25(5):789–804.
- Yin R, Aviv Y, Pazgal A, Tang C (2009) Optimal markdown pricing: Implications of inventory display formats in the presence of strategic consumers. *Management Sci.* 55(8):1391–1408.