The Value of Observing the Buyer Arrival Time in Dynamic Pricing

José Correa
Departamento de Ingeniería Industrial, Universidad de Chile, Santiago, Chile, correa@uchile.cl

Dana Pizarro
Instituto de Ciencias de la Ingeniería, Universidad de O'Higgins, Rancagua, Chile, dana.pizarro@uoh.cl

Gustavo Vulcano
School of Business, Universidad Torcuato di Tella, and CONICET,
Buenos Aires, 1428, Argentina, gvulcano@utdt.edu

We consider a dynamic pricing problem where a firm sells one item to a single buyer to maximize expected revenues. 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 his 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 his arrival time. On the contrary, when she 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 where 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.

Key words: dynamic pricing, strategic consumer behavior, price commitment, price discrimination, worst-case analysis.

History: Original version: August 2020. Revisions: July 2021, October 2022. A preliminary version of this paper appeared in the Proceedings of the EC'20 Conference.

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 B2C and C2C 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 have been rooted in the increasing availability of data from internal and external sources, sophisticated and rapid advances in machine learning and artificial intelligence, and tremendous advances 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, since 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– has accounted 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 a consumer arrival time and adjust the pricing policy accordingly as opposed to the case where this capability is absent. To our knowledge, this *value of observability* has been 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, while 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 due to the seller ignoring the consumer arrival time or being unable to accurately register it. The latter may happen in online marketplaces, where 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 him. 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 case and that in the unobservable case. 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 *value of observability* (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 have described. The model could be alternatively interpreted as having a continuum of buyers with total mass normalized to 1. 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 which 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, where 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 above. 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 some important business settings. For instance, in the video game industry, the connection between the producer or seller, and the buyer, could be modeled as a one-to-one relationship where implementing personalized dynamic pricing is indeed feasible. For example, for purchasing a video game, some consumers with a high discount rate will purchase the game immediately upon release. Some other consumers may anticipate an eventual price decrease and wait, while others may not have heard of the game for a while, and will 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.

In 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. Since the solution of the observable case already takes an infinite time to implement (which would imply 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, an additional difficulty stems from some arrival distributions which might be biased towards regions where 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, 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 would be 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 monotone hazard rate, and even less important when the seller has a level of patience similar to the customer or is very much patient than the customer.

Roadmap. The remainder of this paper is organized as follows: We start with the precise model description in Section 3, spanning both the buyer (Section 3.1) and the seller (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 is 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 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 shows that price discrimination is not profitable compared to 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), where consumers have an exponentially discounted utility function and are more impatient than the seller (which has become a standard setup in the literature and that 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 has been 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 which may imply departing from basic pricing schemes (e.g. Board and Skrzypacz (2016), Pavan et al. (2014)), and in computer science the interest is in designing simple mechanisms which 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 et al. (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 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 econ journal) was 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 payed 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 above, 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 while guaranteeing 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 on their valuation and their willingness to wait before purchasing or leaving. There is no inventory limitation. They show that the firm may restrict attention to cyclic pricing policies, which have length, at most, twice the maximum willingness to wait of the customer population. 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. Customers' types differ along willingness-to-pay and arrival time during the selling season. They assume that the seller only knows the support of the customers' valuations and do not make any other distributional assumptions about customers' willingness-to-pay or arrival times. They further assume that there is no inventory limitation, and that 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 new features: (i) the private valuations of these customers decay over time and (ii) the customers incur monitoring costs. Both the rates of decay and the monitoring costs are private information. The authors propose a "robust pricing" mechanism which 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 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 where 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 1), all present at time zero. Although we consider a similar setting, his model imposes more technical structure, assuming that both the price function and the purchasing functions are monotone decreasing. He also considers an extension with buyers arriving according to a Poisson process and where the seller bargains with one buyer at a time, upon his arrival. In our model, we do not discuss bargaining but allow a general arrival distribution. Wang (2001) also considers the non observable 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 that the price function may not be decreasing. More fundamentally, he did not address the VO bound.

As mentioned above, our result can also be interpreted as *the price of discrimination*. A recent work by Elmachtoub 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 costumer his 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 fixing the pricing policy, with the one earned by each of the two strategies described above. However, they do not consider a dynamic problem but 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 a revenue management (RM) point of view and assume that the seller cannot replenish this unit throughout the selling horizon. On the demand side, a single consumer will arrive at a time that follows a cumulative distribution function (cdf) $G:[0,\infty]\to[0,1]$ 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 $F:[0,\bar{v}]\to[0,1]$ and density f. We assume that the arrival time and the valuation for the buyer are independent and that 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 $p:[\tau,\infty]\to[0,\bar{v}]$. 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 $p:[0,\infty]\to[0,\bar{v}]$. Note that given the bounded support of the valuation distribution, the price function is lower bounded by 0 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 the seller's problems, as well as some preliminary definitions and results. The discussion of the model assumptions is deferred to 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 the consumer 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 non-purchase 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 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 arrival time τ, then the utility function $U(t,v)$ would only be 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:

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

Note that the maximum in the problem above can be attained at $t=\infty$, and then we are considering that the buyer has an outside option with utility equals to zero. That is, a buyer with valuation v who arrives at time τ will purchase (at a finite time) if and only if there exists $\tau_p\in[\tau,\infty)$ such that $U(\tau_p,v)\ge 0$, and will not purchase (or will purchase at time $\tau_p=\infty$) otherwise, obtaining zero utility. We are then assuming, as it 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 will 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, $\phi : [0,\infty) \to [0,\overline{v}] \cup \{\infty\}$ 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 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 where customers strategize over the timing of their purchase, we assume here that the seller is more patient than the buyer and hence her discount rate δ verifies $\delta < \mu$. This setup is also the interesting one to consider, since the problem becomes easy when $\delta \geq \mu$ (see 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 will be shown the pricing function p_τ. In the unobservable case, the seller can only choose $p_0(t)$ and the customer will face that price curve irrespective of his arrival time. Both cases are formally presented below.

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 will pretend that the buyer arrives at time zero, i.e., we initially assume that $\tau = 0$ (w.l.o.g.). To simplify notation, we will 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 will purchase at the first time $t \geq 0$ satisfying $v \geq \phi(t)$ and will not purchase if $v < \phi(\infty)$. His 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 will buy at the first time t at which ψ takes the value v (or will buy immediately if $v \geq \psi(0)$). Due to the lower semi-continuity of p, we have that ϕ is also lower semi-continuous and therefore, ψ is well defined (see Proposition A1 in Appendix A1 for a proof). The purchasing function $\psi(t)$ is the unique non increasing function that supports $\phi(t)$ from below (see Figure 1(a)). The instantaneous probability of selling at time t is given by $\mathrm{d}(1 - F(\psi(t)))$. With this observation, we may write the seller's problem conditioned on the event that the buyer arrives at time 0:

$$[SPO_0] \quad \max_{p,\psi}\ p(0)(1 - F(\psi(0))) + \int_0^\infty e^{-\delta t} p(t)\ \mathrm{d}(1 - F(\psi(t))).$$
$$\text{s.t.} \quad t \in \arg\max_{s \geq 0}\ U(s, \psi(t)) \ \text{ for all } t \geq 0.$$

The first term in the objective function stands for the event where the customer buys immediately at time 0, and the second term accounts for his forward looking behavior. The incentive compatible constraint specifies that a consumer arriving at time zero with valuation $\psi(t)$ maximizes his utility at time $s = t$. We remark here that the individual rationality constraint is implicitly included in the equilibrium constraints due to the argmax being taken over $[0, \infty]$ and, by definition, $U(\infty, v) = 0$ for all valuation v.

We can 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 $\tau > 0$, i.e., $R_\tau = e^{-\delta\tau} R_0$, with R_0 being the objective function value of problem $[SPO_0]$. Finally, the maximum expected revenue of the seller can be written as $R = \int_0^\infty R_\tau g(\tau)\mathrm{d}\tau = R_0 \int_0^\infty e^{-\delta\tau} g(\tau)\mathrm{d}\tau$, so that our assumption above on writing the seller's problem when the customer arrives at time zero is without loss of generality in terms of characterizing the structure of the optimal pricing policy.

The formulation $[SPO_0]$ 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, where there is a continuum of buyers with mass $\mathrm{d}(1-F(\psi(t)))$ who buy at time t, for a total mass of 1 over the infinite horizon; and (ii) price of discrimination, where 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 will 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 her 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 where ϕ takes the same value as $\phi(t)$ (or $s_t = 0$ if such time does not exist, see Figure 1(b)), i.e.,

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

the seller's problem can be described as follows:

$$\begin{aligned}
[SPN]\ \max\nolimits_{p,\phi} &\int_0^\infty e^{-\delta t} p(t) \left[(1-F(\phi(t)))g(t) + 1_{\{\phi'(t)\leq 0\}}(G(t)-G(s_t))\mathrm{d}(1-F(\phi(t)))\right] \mathrm{d}t. \\
\text{s.t.} \quad & t \in \arg\max_{s\geq t} \ U(s,\phi(t)) \ \text{ for all } t.
\end{aligned}$$

The term in brackets stands for the probability of purchasing between times t and $t+\mathrm{d}t$. 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 are accounting for a customer arriving at t_1 with valuation $v \geq v_1$. The second term, $(G(t) - G(s_t))\mathrm{d}(1 - F(\phi(t)))$, is the probability of arriving during the interval $(s_t, t]$ with valuation between $\phi(t)$ and $\phi(t+\mathrm{d}t)$. 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, he will not buy if ϕ is increasing at t, and thus this second term only holds at points where 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} the objective function value of $[SPN]$.

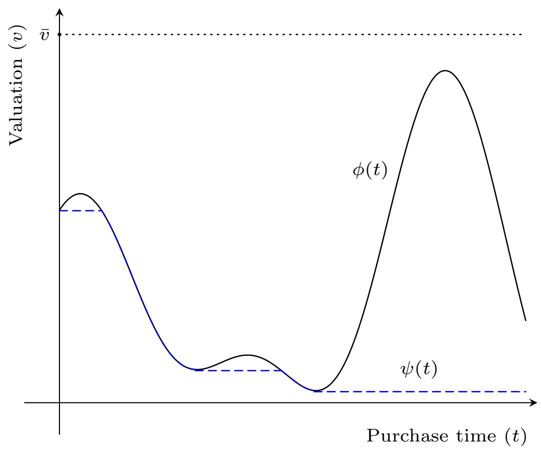

(a) Observable case. Definition of the function $\psi(t)$. For a given function $\phi(t)$, a customer with valuation $\psi(t)$ arriving at $\tau = 0$ will buy at time t.

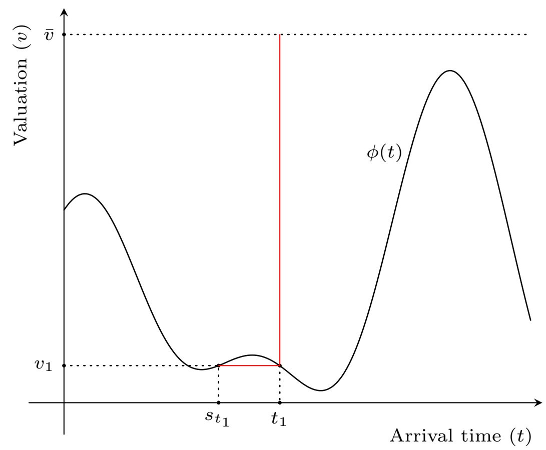

(b) Unobservable case. Characterization of a buyer purchasing at time t_1 including the one arriving exactly at t_1 with valuation $v \geq v_1$, and those arriving between s_{t_1} and t_1 with valuation v_1.

Figure 1 Consumer purchasing behavior.

The formulation $[SPN]$ allows us to revisit the connection with the two alternative model interpretations described in Section 1: (i) infinite supply and demand setup, where here there is a continuum of infinitesimal buyers with point mass $(1 - F(\phi(t)))g(t) + 1_{\{\phi'(t) \leq 0\}}(G(t) - G(s_t))\mathrm{d}(1 - F(\phi(t)))$, who buy at time t (and who have arrived before or at t), and (ii) price of discrimination, where here the model does not allow to price discriminate since all buyers face the same price curve posted at time zero.

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 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 above, 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 1. Also assume the seller discount rate is 1 while that of the buyer is extremely large (so that in fact the buyer behaves myopically: he will buy as soon as the price is below his valuation[4]). Then, if the seller can observe the buyer's arrival, she will start pricing at 1 and then suddenly decrease the price in a continuous fashion until hitting the customer valuation, when the transaction is executed. In this way, she will be extracting all the consumer surplus, with expected value 1/2. Thus, in expectation, the seller gets $R=\int_0^\infty (e^{-t}/2)e^{-t}\mathrm{d}t = 1/4$. 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 non increasing price function, then the problem is relatively easy to solve. Indeed, the seller would need to maximize, over all non increasing functions p, the quantity $R^{uo}=\int_0^\infty [e^{-t}(1-p(t))-(1-e^{-t})p'(t)]\, e^{-t}p(t)\,\mathrm{d}t$. Note that for a non increasing $p(t)$, trade occurs between t and $t+dt$ if either the buyer arrives in that interval and his valuation is above $p(t)$ (hence the term $e^{-t}(1-p(t))$), or the buyer arrived before t and his valuation is between $p(t)$ and $p(t+\mathrm{d}t)$ (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 non increasing 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 non increasing price functions. While we have argued that for the observable case the optimal price curve will be non increasing, for the unobservable case this characterization is unclear. Suppose that she 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 1 to 0) in the last ϵ^2 time units of each interval. As the buyer is myopic, he will buy at the first point in time in which the price is below his valuation, and since ϵ is very small, the probability that the buyer arrives when the price is $p(t)=1$ is close to 1. 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 1. 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, 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 1. On the other side of the spectrum, when $\mu \to \infty$, the buyer essentially behaves myopically. Similarly to the example above 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 $\epsilon \to 0$, the buyer is myopic, so he will buy at the first point in time in which the price is below his valuation. Because the probability that the buyer arrives when the price is $\bar{v}$ is close to 1, 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 1 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 like 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. Finally, 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 $[SPO_0]$, where the buyer arrives at time 0.

The problem $[SPO_0]$ is difficult to solve because of its equilibrium constraint. Our approach will be 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 will 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 with, 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:

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

$$\text{s.t.}\quad \psi(t)=p(t)-\frac{p'(t)}{\mu},\quad \forall t\geq 0. \tag{1}$$

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 $p(t)$:

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

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:

$$\max_p \int_0^\infty I(t,p(t),p'(t),p''(t))\ \mathrm{d}t+r_0.$$

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

$$\frac{\mathrm{d}^2}{\mathrm{d}t^2}\frac{\partial I}{\partial p''}-\frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial I}{\partial p'}+\frac{\partial I}{\partial p}=0.$$

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

$$\int_0^\infty I(t,p(t),p'(t),p''(t))\ \mathrm{d}t.$$

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

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

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 ODEs (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 Appendix A4 we solve the problem for the special case where the valuation is uniformly distributed in $[0,1]$.

Let us highlight that though we know that $\psi(t)$ is non increasing 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 non increasing 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 non increasing.*

The non increasing 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 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 to the relaxed problem $[SPO_0^r]$, and solve it. Its solution will depend 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 her 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 1, 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* $(1-c)R_0$; *i.e.,*

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

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 where the seller is not able to observe the arrival time of the buyer. Differently from the previous observable case, where 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 she commits to a price function at time zero without knowing the precise 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 will 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 will be $2T$ where T is such that, until time T, the seller's expected revenue in the observable case when the buyer arrives at time zero 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

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

where $I_{2k-1} = (2(k-1)T, (2k-1)T]$ and $I_{2k} = ((2k-1)T, 2kT]$ for $k \in \mathbb{N}$, and where 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 will consider along the rest of the section, with origin at a value $t_0 \geq 0$.

We explain below 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 of all, 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 were a

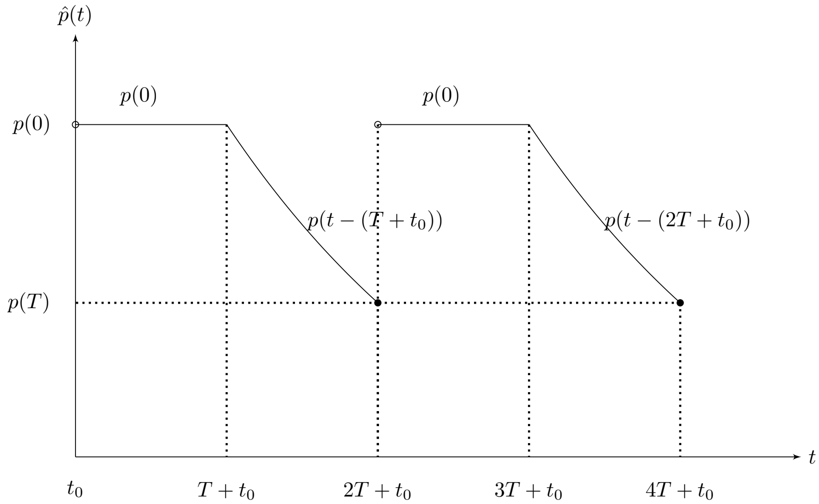

Figure 2 Periodic pricing policy $\hat{p}$ after performing a random shift and setting the origin at time t_0.

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 on the buyer’s purchasing behavior: he is strategic and would not face a similar curve as in the observable case; hence, his behavior would be different. This argument sustains the setting of the high price $p(0)$ in the odd intervals of length T: It allows to accumulate 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 will further reduce the seller’s revenue, which is acceptable for the sake of computing a lower bound. Indeed, the buyers arriving during the odd intervals behave like 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 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 will 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 $t_0 \sim$ 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 Unif$[0,2T]$.*

Therefore, by applying the proposition above, we can assume that the buyer's arrival time, conditional on the fact that the arrival belongs to a specific interval, is Unif$[0,2T]$, and that 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 will denote $\hat{p}$ the feasible pricing policy after performing the random shift. We will also relabel the intervals of the function $\hat{p}$ and denote by $\tilde{I}_{2k-1}$ the range where $\hat{p}$ is constant, and will denote by $\tilde{I}_{2k}$ the range where $\hat{p}$ is the translation of the function p after performing the random shift.

For illustration purposes and as a supplement to this section, in 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 Appendix A6, we characterize the optimal solution of a particular unobservable arrival problem instance with a two-point arrival time distribution and where the valuation of the buyer is Unif$[0,1]$. 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 value of observability 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 monotone hazard rate (MHR), and show that it can be upper bounded by $\exp(1)$, and that 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 upper bound for VO. In Appendix A7 we argue by resorting to an example that VO is lower bounded away from the trivial 1, 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 a Monotone Hazard Rate (MHR) Valuation Distribution.

The case of monotone hazard rate (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 $R = R_0 \int_0^\infty e^{-\delta t} g(t)\mathrm{d}t$, where R_0 is the objective function value of problem $[SPO_0]$ and therefore verifies $R_0 \le \mathbb{E}(v)$, the expected value of the valuation drawn from F. Hence, the seller's expected revenue in the observable case verifies

$$R \le \mathbb{E}(v) \int_0^\infty e^{-\delta t} g(t)\mathrm{d}t.$$

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

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

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 $p^*(1-F(p^*)) \ge \frac{1}{e}\mathbb{E}(v)$, and thus $VO(G,F,\delta,\mu) \le e$, 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 1 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 will announce 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 will buy 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 will set 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

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

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 the argument above does not imply that VO is lower bounded by e, since for the unobservable case we are considering a feasible but not necessarily optimal strategy. What we proved 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 non monotone hazard rate), we start by noting that using a fixed pricing policy in the unobservable arrival case does not necessarily lead to an upper bound defined by a constant. For instance, consider the game where the buyer's valuation is distributed according to a truncated Pareto distribution with parameter 1 and support $[1, M]$, denoted by TruncPareto(1,1,M). That is, $v \sim$ 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 argument above, in the observable case, the seller will announce a price curve that spans all the valuation support (e.g., $p(t) = 1/t$), and the consumer will buy 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 will 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* $-\frac{2W_{-1}(-1/(2\sqrt{e}))+1}{\left(e^{W_{-1}(-1/(2\sqrt{e}))+1/2}-1\right)^2} \approx 4.911$.

Proof. Recalling that R_τ and R_τ^{uo} are the expected values of the seller's revenue in the observable case 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 $-\frac{2W_{-1}(-1/(2\sqrt{e}))+1}{\left(e^{W_{-1}(-1/(2\sqrt{e}))+1/2}-1\right)^2}$, which is roughly 4.911. □

It is worth noting that our result is *robust* to the specification of the arrival distribution. That is, even in the case where 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 her 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 pre-commits to a price curve with the objective of maximizing revenues. 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 *value of observability* (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*, i.e., 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 since 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 where 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 monotone hazard rate, and it is even less important in the extreme cases where 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, e.g., an online seller facing robots searching for competitors' prices.

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

Endnotes

1. Due to this assumption the seller could potentially lose at most a negligible extra revenue and therefore it does not affect our results. Moreover, the lower semi-continuity 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 $U(t,v') = U(t,v) + \epsilon e^{-\mu t} > U(t',v) + \epsilon e^{-\mu t'} = U(t',v')$, i.e., $U(t,v') \geq U(t',v'), \forall t' \geq t$. Thus, the purchasing time of buyer v' cannot be later than t.

3. If $p(t)$ is an arbitrary pricing function we could consider $p'(t)$ the largest nonincreasing function that is below p. Since the buyer is forward-looking, the possible purchasing times will coincide under p and p'.

4. Readers should not confuse the notion of *myopic* buyer with that of *impatient* buyer in the sense of waiting only an infinitesimal amount of time before purchasing. The notion of myopic buyer we will use throughout the paper is the one where the consumer buys as soon as the price is below his valuation. It may then happen that he 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)), where 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. *Operations Research* 68(4):1116–1131.

Araman V, Fayad B (2021) Intertemporal price discrimination with time-varying valuations. *Operations Research* 69(1):245–265.

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*, volume 131 of *International Series in Operations Research & Management Science*, chapter 12, 323–352 (Springer).

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

Aviv Y, Vulcano G (2012) Dynamic list pricing. Ozer O, Phillips R, eds., *The Oxford Handbook of Pricing Management*, chapter 23, 522–584, Oxford Handbooks (Oxford University Press).

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

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

Blumrosen L, Holenstein T (2008) Posted prices vs. negotiations: an asymptotic analysis. *Proceedings of the 9th ACM conference on Electronic commerce*, 49, EC'08 (New York, NY: ACM).

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

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

Briceño-Arias L, Correa J, Perlroth A (2017) Optimal continuous pricing with strategic consumers. *Management Science* 63(8):2741–2755.

Cachon GP, Swinney R (2009) Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Science* 55(3):497–511.

Cachon GP, Swinney R (2011) The value of fast fashion: Quick response, enhanced design, and strategic consumer behavior. *Management science* 57(4):778–795.

Caldentey R, Liu Y, Lobel I (2017) Intertemporal pricing under minimax regret. *Operations Research* 65(1):104–129.

Chawla S, Hartline JD, Malec DL, Sivan B (2010) Multi-parameter mechanism design and sequential posted pricing. *Proceedings of the forty-second ACM symposium on Theory of computing*, 311–320.

Chen Y, Farias V (2018) Robust dynamic pricing with strategic customers. *Mathematics of Operations Research* 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 Science* 65(12):5535–5555.

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

Correa J, Foncea P, Pizarro D, Verdugo V (2019) From pricing to prophets, and back! *Operations Research Letters* 47(1):25–29.

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

Dhangwatnotai P, Roughgarden T, Yan Q (2015) Revenue maximization with a single sample. *Games and Economic Behavior* 91:318–333.

Elmachtoub AN, Gupta V, Hamilton ML (2021) The value of personalized pricing. *Management Science* 67(10):6055–6070.

Elmaghraby W, Lippman S, Tang C, Yin R (2009) Will more purchasing options benefit customers? *Production and Operations Management* 18:381–401.

Gallego G, Topaloglu H, et al. (2019) *Revenue management and pricing analytics*, volume 209 (Springer).

Gershkov A, Moldovanu B, Strack P (2018) Revenue-maximizing mechanisms with strategic customers and unknown, markovian demand. *Management Science* 64(5):2031–2046.

Golrezaei N, Javanmard A, Mirrokni V (2021) Dynamic incentive-aware learning: Robust pricing in contextual auctions. *Operations Research* 69(1):297–314.

Gönsch J, Klein R, Neugebauer M, Steinhardt C (2013) Dynamic pricing with strategic customers. *Journal of Business Economics* 83(5):505–549.

Kessel K, Saberi A, Shameli A, Wajc D (2022) The stationary prophet inequality problem. *Proceedings of the 23rd ACM conference on Electronic commerce*, EC'22 (New York, NY: ACM).

Lai G, Debo L, Sycara K (2010) Buy now and match later: Impact of posterior price matching on profit with strategic consumers. *MSOM* 12:33–55.

Landsberger M, Meilijson I (1985) Intertemporal price discrimination and sales strategy under incomplete information. *The RAND Journal of Economics* 424–430.

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

Liu Q, Van Ryzin GJ (2008) Strategic capacity rationing to induce early purchases. *Management Science* 54(6):1115–1131.

Liu Y, Cooper W (2015) Optimal dynamic pricing with patient customers. *Operations Research* 63(6):1307–1319.

Lobel I (2020) Dynamic pricing with heterogeneous patience levels. *Operations Research* 68(4):1038–1046.

Mas-Colell A, Whinston M, Green J (1995) *Microeconomic Theory* (New York: Oxford University Press).

Mersereau A, Zhang D (2012) Markdown pricing with unknown fraction of strategic consumers. *MSOM* 14:355–370.

Myerson RB (1981) Optimal auction design. *Mathematics of operations research* 6(1):58–73.

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

Papanastasiou Y, Savva N (2017) Dynamic pricing in the presence of social learning and strategic consumers. *Management Science* 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).

Stokey NL (1979) Intertemporal price discrimination. *The Quarterly Journal of Economics* 355–371.

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

Su X, Zhang F (2008) Strategic customer behavior, commitment and supply chain performance. *Management Science* 54:1759–1773.

Swinney R (2011) Selling to strategic consumers when product value is uncertain: The value of matching supply and demand. *Management Science* 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) (New York, NY: Springer).

Wang R (2001) Optimal pricing strategy for durable-goods monopoly. *Journal of Economic Dynamics and 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 Science* 55:1391–1408.

The Value of Observing the Buyer Arrival Time in Dynamic Pricing

APPENDIX

José Correa

Departamento de Ingeniería Industrial, Universidad de Chile. Email: correa@uchile.cl

Dana Pizarro

Instituto de Ciencias de la Ingeniería, Universidad de O'Higgins, Rancagua, Chile, Email: dana.pizarro@uoh.cl

Gustavo Vulcano

School of Business, Universidad Torcuato di Tella, and CONICET, Buenos Aires, Argentina. Email: gvulcano@utdt.edu

A1. Complementary results

PROPOSITION A1. *The function $\phi(t)$ is lower semi-continuous.*

Proof. We need to show that for all $t_0 \geq 0$, it holds that

$$\liminf_{t\to t_0} \phi(t) \geq \phi(t_0). \tag{A1}$$

First, note that from the definition of $\phi(t)$, we are looking for a value v to set $\phi(t) = v$, i.e., v must be the smallest valuation verifying: $U(t,v) \geq U(t',v)$, for all $t' > t$[1], or equivalently,

$$e^{-\mu t}(v - p(t)) \geq e^{-\mu t'}(v - p(t')),$$

where by isolating v we get

$$v \geq \frac{p(t) - e^{-\mu(t'-t)}p(t')}{1 - e^{-\mu(t'-t)}}.$$

Observe that due to the lower semi-continuity of $p(t)$, the definition of the threshold function ϕ is equivalent to

$$\phi(t) = \sup_{t'>t} \left\{ \frac{p(t) - e^{-\mu(t'-t)}p(t')}{1 - e^{-\mu(t'-t)}} \right\}. \tag{A2}$$

To prove that ϕ is lower semi-continuous we need the following auxiliary result, whose proof follows from the definition of liminf for functions.

Auxiliary lemma. If f and g are functions such that for all $y \geq x$ it holds that $f(x) \geq g(x,y)$, then $\liminf_{x\to x_0} f(x) \geq \liminf_{x\to x_0} g(x,y)$, for all $y \geq x_0$.

[1] Note that the inequality holds for all v setting $t' = t$ and therefore we can restrict the condition for t' strictly greater than t.

From (A2) we have $\phi(t) \geq \frac{p(t)-e^{-\mu(t'-t)}p(t')}{1-e^{-\mu(t'-t)}}$ for all $t' \geq t$, and using the auxiliary lemma it follows that, for all $t' \geq t_0$,

$$\liminf_{t\to t_0} \phi(t) \geq \liminf_{t\to t_0} \frac{p(t)-e^{-\mu(t'-t)}p(t')}{1-e^{-\mu(t'-t)}}.$$

Due to the lower semi-continuity of $p(t)$ and the continuity of the exponential function, and in view of (A1), the right side of this inequality is at least $\frac{p(t_0)-e^{-\mu(t'-t_0)}p(t')}{1-e^{-\mu(t'-t_0)}}$, and therefore,

$$\liminf_{t\to t_0} \phi(t) \geq \frac{p(t_0)-e^{-\mu(t'-t_0)}p(t')}{1-e^{-\mu(t'-t_0)}} \quad \forall t' \geq t_0.$$

Then, $\liminf_{t\to t_0} \phi(t)$ is at least the maximum, over all $t' \geq t_0$, of $\frac{p(t_0)-e^{-\mu(t'-t_0)}p(t')}{1-e^{-\mu(t'-t_0)}}$, which is equal to $\phi(t_0)$. Thus, ϕ is lower semi-continuous in $\mathbb{R}_0^+$. □

PROPOSITION A2. *The Euler-Lagrange equation associated to the problem*

$$\max_p \int_0^{+\infty} I(t, p(t), p'(t), p''(t))\mathrm{d}t$$

is given by

$$f'\left(p(t)-\frac{p'(t)}{\mu}\right)\left(-\frac{p''(t)}{\mu}+p'(t)\right)(-\delta p(t)+p'(t)) + f\left(p(t)-\frac{p'(t)}{\mu}\right)[\delta(\delta-\mu)p(t)-2\delta p'(t)+2p''(t)] = 0. \tag{A3}$$

Proof. Recall that $I(t,p(t),p'(t),p''(t)) = e^{-\delta t}p(t)\left(-p'(t)+\frac{p''(t)}{\mu}\right)f\left(p(t)-\frac{p'(t)}{\mu}\right)$. We have to check that

$$\frac{d^2}{dt^2}\frac{\partial I}{\partial p''} - \frac{d}{dt}\frac{\partial I}{\partial p\prime} + \frac{\partial I}{\partial p} = 0 \tag{A4}$$

is equivalent to equation (A3).

The first term of the RHS of (A4) is given by

$$\begin{aligned}\frac{d^2}{dt^2}\frac{\partial I}{\partial p''} = {} & \frac{e^{-\delta t}}{\mu} f\left(p(t)-\frac{p'(t)}{\mu}\right)\left(p''(t)-2\delta p'(t)+\delta^2 p(t)\right) + \\ & \frac{e^{-\delta t}}{\mu} f'\left(p(t)-\frac{p'(t)}{\mu}\right)\left[2(p'(t)-\delta p(t))\left(p'(t)-\frac{p''(t)}{\mu}\right)+p(t)\left(p''(t)-\frac{p'''(t)}{\mu}\right)\right] + \\ & \frac{e^{-\delta t}}{\mu} f''\left(p(t)-\frac{p'(t)}{\mu}\right)p(t)\left(p'(t)-\frac{p''(t)}{\mu}\right)^2.\end{aligned}$$

On the other hand, computing the second term we obtain

$$\begin{aligned}\frac{d}{dt}\frac{\partial I}{\partial p'} = {} & \frac{e^{-\delta t}}{\mu} f\left(p(t)-\frac{p'(t)}{\mu}\right)(\delta p(t)-p'(t)) + \\ & \frac{e^{-\delta t}}{\mu} f'\left(p(t)-\frac{p'(t)}{\mu}\right)\left[\left(\frac{p''(t)}{\mu}-p'(t)\right)\left(\frac{\delta p(t)-p'(t)}{\mu}+p(t)\right)+\frac{p(t)}{\mu}\left(p''(t)-\frac{p'''(t)}{\mu}\right)\right] + \\ & \frac{e^{-\delta t}}{\mu} f''\left(p(t)-\frac{p'(t)}{\mu}\right)p(t)\left(\frac{p''(t)}{\mu}-p(t)\right)^2.\end{aligned}$$

Finally, the partial derivative of I with respect to p is the following

$$\frac{\partial I}{\partial p} = e^{-\delta t}\left(\frac{p''(t)}{\mu} - p'(t)\right)\left(f\left(p(t) - \frac{p'(t)}{\mu}\right) + p(t) f'\left(p(t) - \frac{p'(t)}{\mu}\right)\right).$$

Hence, equation (A3) comes from using the expressions above and equalizing the LHS of (A4) to zero. □

A2. Proofs of results in the main body

A2.1. Proof of Proposition 1

Let $p(t)$ be an optimal solution of the relaxed problem $[SPO_0^r]$ and suppose that there exists t such that $\psi(t)$ is an inner local maximum. Then, it must hold that

$$\psi'(t) = p'(t) - \frac{p''(t)}{\mu} = 0. \tag{A5}$$

Recalling that the valuation density f is positive, observe that, at t, the Euler-Lagrange equation (3) becomes

$$\delta(\delta-\mu)p(t) - 2\delta p'(t) + 2p''(t) = 0,$$

and therefore, together with (A5), $p'(t) = \frac{\delta p(t)}{2}$, and thus, $p'(t) - \delta p(t) < 0$.

Let $\epsilon > 0$ and $\rho > 0$ be such that $t_1 = t - \epsilon$ and $t_2 = t + \rho$ satisfying $\psi(t_1) = \psi(t_2)$ and $p'(t_i) - \delta p(t_i) < 0$ for $i = 1, 2$. Furthermore, since ψ has a maximum at t, it must hold that $\psi'(t_1) > 0$ and $\psi'(t_2) < 0$.

Let us first suppose that $f'(\psi(t_1)) = f'(\psi(t_2)) \geq 0$. In this case, considering the first term in (3),

$$\underbrace{f'\left(p(t_2) - \frac{p'(t_2)}{\mu}\right)}_{f'(\psi(t_2)\geq 0}\underbrace{\left(-\frac{p''(t_2)}{\mu} + p'(t_2)\right)}_{\psi'(t_2)<0}\underbrace{(-\delta p(t_2) + p'(t_2))}_{<0} \geq 0,$$

and therefore, since $p(t)$ satisfies the Euler-Lagrange equation (3) for all t –and, in particular, for t_2–, we must have

$$\delta(\delta-\mu)p(t_2) - 2\delta p'(t_2) + 2p''(t_2) \leq 0. \tag{A6}$$

Since by construction $p'(t_2) - \delta p(t_2) < 0$, we have $\frac{p'(t_2)}{\delta} < p(t_2)$, and bounding from below the first term in the LHS of (A6), we obtain

$$-(\delta+\mu)p'(t_2) + 2p''(t_2) < 0. \tag{A7}$$

Recalling that $t_2 = t + \rho$, taking the liminf in the LHS of (A7) when $\rho \to 0$, by the lower semi-continuity of the price function $p(t)$, we obtain

$$-(\delta+\mu)p'(t) + 2p''(t) \leq 0,$$

which is equivalent to $2p''(t) \leq (\delta+\mu)p'(t)$. But from (A5), $\mu p'(t) = p''(t)$, so $2\mu p'(t) \leq (\delta+\mu)p'(t)$ and therefore $\mu \leq \delta$ (because $p'(t) = \frac{\delta p(t)}{2} > 0$), which is a contradiction.

Now, consider the case where $f'(\psi(t_1)) = f'(\psi(t_2)) < 0$. Then, it must hold that

$$\underbrace{f'\left(p(t_1) - \frac{p'(t_1)}{\mu}\right)}_{f'(\psi(t_1))<0} \underbrace{\left(-\frac{p''(t_1)}{\mu} + p'(t_1)\right)}_{\psi'(t_1)>0} \underbrace{(-\delta p(t_1) + p'(t_1))}_{<0} > 0$$

and now we can proceed analogously to the argument above.

Therefore $\psi(t)$ cannot have an inner local maximum, and with a similar argument, neither an inner local minimum. Hence, $\psi(t)$ has to be monotone.

We are now left with showing that the function $\psi(t)$ is indeed non increasing. By contradiction, suppose that ψ is increasing. We will see that if so we could improve the expected revenue, contradicting that ψ corresponds to the optimal solution of the relaxed problem $[SPO_0^r]$. To this end, let us consider the constant function $\hat{p}(t) = p(0)$ for all t. Then, $\hat{\psi}(t) = p(0)$ and therefore the value of the objective function of $[SPO_0^r]$ by considering the feasible pricing policy $\hat{p}$ is given by

$$\hat{p}(0)(1 - F(\hat{\psi}(0))) + \int_0^\infty e^{-\delta t}\hat{p}(t)(-\hat{\psi}'(t))f(\hat{\psi}(t))\mathrm{d}t = p(0)(1 - F(p(0))).$$

On the other hand, the expected revenue of the seller under the pricing policy p can be computed as

$$p(0)(1 - F(\psi(0))) + \int_0^\infty e^{-\delta t}p(t)(-\psi'(t))f(\psi(t))\ \mathrm{d}t.$$

Note that the second term is negative and therefore the expression above is upper bounded by the expected revenue obtained by selling at time 0. That is,

$$p(0)(1 - F(\psi(0))) + \int_0^\infty e^{-\delta t}p(t)(-\psi'(t))f(\psi(t))\ \mathrm{d}t < p(0)(1 - F(\psi(0))).$$

Note that $1 - F(\psi(0)) > 1 - F(p(0))$, and therefore the expected revenue under the price function $\hat{p}$ is greater than the expected revenue under the price function p, which contradicts the optimality of p. Thus, we can conclude that ψ is a non increasing function.

A2.2. Proof of Proposition 2

Given a pair $(p(t), \psi(t))$ solution of $[SPO_0^r]$, with $\psi(t) = p(t) - \frac{p'(t)}{\mu}$ for all t, we must show that it meets the equilibrium constraint of $[SPO_0]$, that is:

$$t \in \arg\max_{s \geq 0}\ e^{-\mu s}(\psi(t) - p(s)) \quad \forall t. \tag{A8}$$

Let $h(s) = e^{-\mu s}(\psi(t) - p(s))$, leading to

$$h'(s) = e^{-\mu s}(-\mu(\psi(t) - p(s)) - p'(s)),$$

and

$$h''(s) = -\mu e^{-\mu s}(-\mu(\psi(t) - p(s)) - p'(s)) + e^{-\mu s}(\mu p'(s) - p''(s)).$$

Given an interior solution t of (A8), it must verify $h'(t) = 0$ and

$$h''(t) = \mu e^{-\mu t}\left(p'(t) - \frac{p''(t)}{\mu}\right) \leq 0.$$

Since $(p(t), \psi(t))$ is solution of $[SPO_0^r]$, then from Proposition 1 we know that $\psi'(t) \leq 0$, and therefore, $h''(t) \leq 0$. Hence, $t \in \arg\max_{s \geq 0}\ e^{-\mu s}(\psi(t) - p(s))$, for any pair of functions $(p(t), \psi(t))$ solution of $[SPO_0^r]$. Recalling that the solution of $[SPO_0^r]$ defines an upper bound of $[SPO_0]$, we have that such pair $(p(t), \psi(t))$ indeed defines a solution to $[SPO_0]$.

A2.3. Proof of Lemma 1

Note that an equivalent inequality would be

$$\begin{aligned} cR_0 &\geq R_0 - \left[p(0)(1 - F(\psi(0))) + \int_0^T e^{-\delta t} p(t) \mathrm{d}(1 - F(\psi(t)))\right] \\ &= \int_T^\infty e^{-\delta t} p(t) \mathrm{d}(1 - F(\psi(t))) \end{aligned}$$

By contradiction, suppose that for $T = \ln(1/c)/\delta$, we have that:

$$\int_T^\infty e^{-\delta t} p(t) \mathrm{d}(1 - F(\psi(t))) > c\left[p(0)(1 - F(\psi(0))) + \int_0^\infty e^{-\delta t} p(t) \mathrm{d}(1 - F(\psi(t)))\right]. \tag{A9}$$

Consider the price function $\hat{p}(t) = p(t + T)$ and its associated purchasing function $\hat{\psi}$. The seller's expected revenue can be computed as:

$$R_{\hat{p}} = \hat{p}(0)(1 - F(\hat{\psi}(0))) + \int_0^\infty e^{-\delta t} \hat{p}(t) \mathrm{d}(1 - F(\hat{\psi}(t))).$$

By the definition of $\hat{p}$ and doing the change of variable $u = t + T$, it follows that the seller's expected revenue is given by:

$$R_{\hat{p}} = p(T)(1 - F(\psi(T))) + e^{\delta T} \int_T^\infty e^{-\delta t} p(t) \mathrm{d}(1 - F(\psi(t))).$$

Applying (A9), it follows that this expression verifies

$$R_{\hat{p}} > p(T)(1-F(\psi(T))) + e^{\delta T} c \left[p(0)(1-F(\psi(0))) + \int_0^{\infty} e^{-\delta t} p(t) \mathrm{d}(1-F(\psi(t))) \right].$$

Note that $p(T)(1-F(\psi(T)))$ is non negative, and that $T = \ln(1/c)/\delta$ implies $e^{\delta T} c = 1$. Thus, the seller's expected revenue for the pricing policy $\hat{p}$ is bigger than the seller's expected revenue for the pricing policy p, which contradicts the optimality of the price function p.

A2.4. Proof of Proposition 3

Suppose that we have the periodic function $\hat{p}$ with period $2T$ given by (4) and consider a random shift, that is, for a random variable $t_0 \sim \mathrm{Unif}[0, 2T]$, consider the function $\hat{p}_{t_0}(t) = \hat{p}(t + t_0)$. Then, given that the buyer arrives in the interval $I_{2k-1} \cup I_{2k}$ of length $2T$, for some $k \in \mathbb{N}$, and denoting by X the random variable *arrival time*, we have the following:

$$\begin{aligned}
\mathbb{P}(X \leq t | X \in (I_{2k-1} \cup I_{2k})) &= \mathbb{P}(X \leq t | X \in (2(k-1)T - t_0, 2kT - t_0]) \\
&= \mathbb{P}(X \in (2(k-1)T - t_0, t]).
\end{aligned}$$

Letting s be the length of the interval $[2(k-1)T - t_0, t]$, i.e., $s = t - (2(k-1)T - t_0)$, the expression above verifies

$$\begin{aligned}
\mathbb{P}(X \leq t | X \in (2(k-1)T - t_0, 2kT - t_0]) &= \mathbb{P}(X \in (2(k-1)T - t_0, 2(k-1)T - t_0 + s]) \\
&= \mathbb{P}(2(k-1)T - X < t_0 \leq 2(k-1)T - X + s) \\
&= \frac{s}{2T} \quad (\text{because } t_0 \sim \mathrm{Unif}[0, 2T]) \\
&= \frac{t - (2(k-1)T - t_0)}{2T},
\end{aligned}$$

which proves that X is uniformly distributed in $I_{2k-1} \cup I_{2k}$, and the proof is completed.

A2.5. Proof of Lemma 2

Without loss of generality let us suppose that $k = 1$, that is, the buyer arrives at time $t_0 + t$ belonging to $\tilde{I}_1$ with valuation $v \geq p(T)$, and further assume that he will not purchase before time $T + t_0$ so that the seller is making less revenue than she could really make.

To prove the lemma we analyze the consumer behavior in the unobservable case under the pricing policy $\hat{p}$ depending on his valuation. More specifically we will prove the followings three statements:

1. If $v \in [p(T), \psi(T))$, then the buyer buys at time $2T + t_0$.
2. If $v \in [\psi(T), \psi(0))$, then the buyer waits and buys at time $\tau \in (T + t_0, 2T + t_0]$ satisfying $\psi(\tau) = v$.

3. If $v \geq \psi(0)$ the buyer purchases at time $t_0 + T$.

First, consider a buyer with valuation $v \in [p(T), \psi(T))$. Knowing that he will purchase to gain some positive utility (eventually at time $2T + t_0$), if he decides to buy at time $\tau < 2T + t_0$, then by the monotonicity of the purchasing function ψ in the observable case, we have that $\psi(\tau - (T + t_0)) > \psi(2T + t_0 - (T + t_0)) = \psi(T)$ and it means that the buyer must have valuation greater than $\psi(T)$ to be optimum to purchase at time τ, which is not the case. We then conclude that in this case he will buy at time $2T + t_0$.

Secondly, if the buyer has valuation $v \in [\psi(T), \psi(0))$, then by using the calculation of the purchasing function for the observable arrival case –conducted under the assumption that the buyer arrives at time 0–, we have that for some $t \in [0, T]$, it holds that $v = \psi(t)$, i.e.,

$$t \in \arg\max_{s \geq 0} U(s, \psi(t)),$$

which means that

$$e^{-\mu t}(\psi(t) - p(t)) \geq e^{-\mu s}(\psi(t) - p(s)), \ \forall s \geq 0.$$

This is equivalent to

$$e^{-\mu(T+t_0)} e^{-\mu t}(\psi(t) - p(t)) \geq e^{-\mu(T+t_0)} e^{-\mu s}(\psi(t) - p(s)), \ \forall s \geq 0.$$

Hence, the buyer will buy at time $\tau = T + t_0 + t$ satisfying $\psi(t) = v$.

Finally, the third statement follows directly from the definition of the threshold function ψ.

The lemma follows by observing that if the buyer has valuation at least $\psi(T)$, the seller's revenue is the same as in the observable case with the buyer arriving at time $T + t_0$ and accumulating revenue up to time $2T + t_0$ (cases (2) and (3)). But if the buyer has valuation between $p(T)$ and $\psi(T)$ (case (1)), then he will buy before time $2T + t_0$ in the unobservable setting under the price function $\hat{p}$, but he will buy after that time in the observable case with arrival time $T + t_0$.

Therefore, we conclude that, conditioned on the event that the buyer with valuation greater than $p(T)$ arrives at time $T + t_0$ –which is equivalent to looking at the problem in the interval $[T + t_0, 2T + t_0]$ in the observable case–, the seller's expected revenue under the policy $\hat{p}$ in the unobservable case is at least the expected revenue earned up to time $2T + t_0$ in the observable case with arrival time $T + t_0$.

A2.6. Proof of Lemma 3

Consider the pricing policy $\hat{p}$ described in Figure 2 and fix the buyer arrival time τ. Recall that t_0 is the uniform random variable involved in the random shift applied over the original price function p to get $\hat{p}$. Recall that these functions have period $2T$, where we are setting $T = \ln(1/c)/\delta$.

Suppose, without loss of generality, that the buyer arrives during the first cycle of the policy; i.e., $\tau \in [t_0, t_0 + 2T]$. Thus, $t_0 \sim \text{Unif}[\tau - 2T, \tau]$. In order to have intervals defined around t_0, we denote $\tilde{\mathcal{I}}_1 := [\tau - T, \tau]$ and $\tilde{\mathcal{I}}_2 := [\tau - 2T, \tau - T]$. With this definition, we have that $\tau \in \tilde{I}_i$ if and only if $t_0 \in \tilde{\mathcal{I}}_i$, for $i = 1, 2$.

In our analysis we will only consider the buyer's arrival if it belongs to the interval $\tilde{I}_1$, otherwise, we simply bound the revenue by 0.

Note that if $\tau \in \tilde{I}_1$, we can lower bound R^{uo}_τ by the expected revenue obtained by considering that the buyer has valuation at least $p(T)$ and that he purchases after time $t_0 + T$. This is because the buyer does not purchase if $v < p(T)$, and his wait until $t_0 + T$ to buy when he could have bought earlier would only hurt the seller's revenue. Then, from Lemma 2, R^{uo}_τ is at least the expected revenue earned up to time $2T + t_0$ in the observable case with arrival time $T + t_0$, i.e., $R^{uo}_\tau \geq R_{[t_0+T, t_0+2T]}$.

Let R_{t_0+T} be the expected revenue in the observable case if the buyer arrives at time $t_0 + T$, for a given value t_0. After applying Lemma 1, we have $R_{[t_0+T,t_0+2T]} \geq (1-c)R_{t_0+T}$, so that $R^{uo}_\tau \geq (1-c)R_{t_0+T}$.

We now use the analysis above to compute a bound for the expected value of the seller's revenue in the unobservable case conditioned on the event that the buyer arrives at time τ. To this end, we define the sample path-based revenue S^{uo}_τ for the unobservable case from time τ onward, i.e., $R^{uo}_\tau = \mathbb{E}(S^{uo}_\tau)$, and conditioning on the random, shifted origin time t_0, we get:

$$
\begin{aligned}
\mathbb{E}(S^{uo}_\tau) &= \mathbb{E}_{t_0}(\mathbb{E}(S^{uo}_\tau \mid t_0)) \\
&= \mathbb{E}(S^{uo}_\tau \mid t_0 \in \tilde{\mathcal{I}}_1)\mathbb{P}(t_0 \in \tilde{\mathcal{I}}_1) + \mathbb{E}(S^{uo}_\tau \mid t_0 \in \tilde{\mathcal{I}}_2)\mathbb{P}(t_0 \in \tilde{\mathcal{I}}_2) \\
&= \frac{1}{2}\mathbb{E}(S^{uo}_\tau \mid t_0 \in \tilde{\mathcal{I}}_1) + \frac{1}{2}\mathbb{E}(S^{uo}_\tau \mid t_0 \in \tilde{\mathcal{I}}_2) \\
&\geq \frac{1}{2}(1-c)\mathbb{E}_{t_0}(R_{t_0+T} \mid t_0 \in \tilde{\mathcal{I}}_1),
\end{aligned}
$$

where the last equality holds because $t_0 \sim \text{Unif}[\tau - 2T, \tau]$, and the inequality follows from the observation above.

Note that $R_{t_0+T} = e^{-\delta(t_0+T-\tau)}R_\tau = ce^{-\delta(t_0-\tau)}R_\tau$, where the second equality holds from $e^{-\delta T} = c$. Therefore, it is enough to compute $\mathbb{E}_{t_0}\left(e^{-\delta(t_0-\tau)} \mid t_0 \in \tilde{\mathcal{I}}_1\right)$. In fact, now for $t_0 \sim \text{Unif}[\tau - T, \tau]$, we have

$$
\begin{aligned}
\mathbb{E}_{t_0}\left(e^{-\delta(t_0-\tau)} \mid t_0 \in \tilde{\mathcal{I}}_1\right) &= \int_{\tau-T}^{\tau} e^{-\delta(t_0-\tau)}\frac{1}{T}\mathrm{d}t_0 \\
&= \frac{e^{\delta T} - 1}{\delta T}.
\end{aligned}
$$

By the definition of T, we know that $T\delta = \ln(1/c)$ and $e^{\delta T} = 1/c$, and therefore we have

$$
\mathbb{E}_{t_0}\left(e^{-\delta(t_0-\tau)} \mid t_0 \in \tilde{\mathcal{I}}_1\right) = \frac{1-c}{c\ln(1/c)}.
$$

We then obtain the following lower bound for the expectation of the seller's revenue in the unobservable case that depends on c:

$$R_\tau^{uo} = \mathbb{E}(S_\tau^{uo}) \geq \frac{(1-c)^2}{2\ln(1/c)} R_\tau,$$

which completes the proof.

A3. Discussion of model assumptions

Our model is quite general and follows the standard setup of the OM literature. Yet, there a few features that are worth discussing.

First, we consider a consumer's utility function of the intertemporal type (e.g., see chapter 20 in Mas-Colell et al. (1995)), where the buyer discounts the net payoff $v - p(t)$ if he decides to buy at a later time t, following the classic econ approach (e.g., see Landsberger and Meilijson (1985), Besanko and Winston (1990)). It has also been the prevailing utility function within the OM literature on strategic consumer behavior (Gönsch et al. (2013)). An alternative model would be one where only the valuation of the buyer declines, obtaining the utility function $U(t,v) = e^{-\mu t}v - p(t)$ (e.g. Aviv and Pazgal (2008), Cachon and Swinney (2009)). However, this model is equivalent to one where a utility function like ours is considered and the seller is more impatient than the buyer. Indeed, by considering the utility function $U(t,v) = e^{-\mu t}v - p(t)$ and a seller discount rate $\delta \geq 0$, the setup is equivalent to defining the utility function $\bar{U}(t,v) = e^{-\mu t}(v - \bar{p}(t))$ and a seller's discount rate $\mu + \delta$, with $\bar{p}(t) = e^{\mu t}p(t)$. In this case, $U(t,v) = e^{-\mu t}v - p(t) = e^{-\mu t}(v - e^{\mu t}p(t)) = e^{-\mu t}(v - \bar{p}(t)) = \bar{U}(t,v)$, and in both cases the seller gets $e^{-\delta t}p(t) = e^{-(\delta+\mu)t}\bar{p}(t)$.

Second, in our model we assume that the seller is more patient than the consumers, i.e., $\delta < \mu$. Otherwise, when $\delta \geq \mu$, the optimal pricing policy is to fix a constant price equal to the monopoly price, obtaining that VO is trivially 1. It is enough to prove it for the observable case because if the optimal price function is constant (and independent of the parameters of the problem), then the advantage of being able to observe the arrival vanishes and both problems are equivalent.

To see this, consider first the case where both discount rates are equal ($\delta = \mu$). This is exactly the setting considered by Stokey (1979) in which she proved that no price discrimination occurs and the optimal pricing policy is to charge the monopoly price, namely p_m, during the whole selling horizon. If the seller is more impatient than the buyer, i.e. $\mu < \delta$, it is enough to note that her optimal expected revenue is upper bounded by the one in the case $\mu = \delta$ since a bigger seller's discount rate can only lead to a worse revenue for her, which is achieved by taking $p(t) = p_m, \forall t$. At the same time, a fixed pricing policy is a feasible solution for the unobservable case, and the revenue derived from it provides a lower bound for R^{uo}. All in all, in this case VO is upper bounded by 1, and hence it is exactly 1.

A4. Solving the Observable Case.

In the particular case when the buyer's valuation has density function $f(x)=kx^{\alpha}$, for some integer number α and positive number k, we can explicitly calculate the pricing function $p(t)$ and the purchasing function $\psi(t)$ for the observable case.

To start with, the problem $[SPO_0^r]$ becomes:

$$\max_{p}\int_{0}^{\infty} I(t,p(t),p'(t),p''(t))\mathrm{d}t + p(0)(1-F(\psi(0))), \tag{A10}$$

where $I(t,p(t),p'(t),p''(t)) = e^{-\delta t}p(t)k\left(p(t)-\frac{p'(t)}{\mu}\right)^{\alpha}\left(-p'(t)+\frac{p''(t)}{\mu}\right)$.

By formulating the Euler-Lagrange equation (3) in this case, we obtain the following ODE in the function $p(t)$ [2]

$$\alpha\left(p'(t)-\frac{p''(t)}{\mu}\right)(p'(t)-\delta p(t)) + \left(p(t)-\frac{p'(t)}{\mu}\right)(\delta(\delta-\mu)p(t)-2\delta p'(t)+2p''(t))=0,$$

which is equivalent to

$$\delta(\delta-\mu)-\frac{p'(t)}{p(t)}\delta\left(\alpha+1+\frac{\delta}{\mu}\right)+\frac{p'(t)^2}{p(t)^2}\left(\alpha+\frac{2\delta}{\mu}\right)-\frac{p'(t)p''(t)}{p(t)^2}\frac{1}{\mu}(2+\alpha)+\frac{p''(t)}{p(t)}\left(\frac{\alpha\delta}{\mu}+2\right). \tag{A11}$$

Setting $y(t)=(\log(p(t)))'$, we have that $p'(t)/p(t)=y(t)$, $p'(t)^2/p(t)^2=y^2(t)$, $p'(t)p''(t)/p(t)^2 = y(t)(y'(t)+y^2(t))$, and $p''(t)/p(t)=(y'(t)+y^2(t))$, which allows to rewrite equation (A11) as:

$$\delta(\delta-\mu)-y(t)\delta\left(\alpha+1+\frac{\delta}{\mu}\right)+y^2(t)\left(\alpha+\frac{2\delta}{\mu}\right)-(y(t)y'(t)+y^3(t))\frac{1}{\mu}(2+\alpha)+(y'(t)+y^2(t))\left(\frac{\alpha\delta}{\mu}+2\right)=0.$$

Rearranging terms, this equation can be written as

$$\begin{aligned}&\delta(\delta-\mu)-y(t)\delta\left(\alpha+1+\frac{\delta}{\mu}\right)+y^2(t)(2+\alpha)\left(1+\frac{\delta}{\mu}\right)\\&-y^3(t)\frac{1}{\mu}(2+\alpha)-y(t)y'(t)\frac{1}{\mu}(2+\alpha)+y'(t)\left(\frac{\alpha\delta}{\mu}+2\right)=0.\end{aligned} \tag{A12}$$

Note that (A12) is a separable first-order nonlinear ODE, that is, can be written as $y'=H(y)$, for some function H. Therefore it can be solved by integration. For our case the solution turns out to be quite contrived (and obtainable only in impliocit form), but for two special cases we can give clean closed form solutions. We thus show the optimal pricing policy for the cases where the buyer's valuation distribution is uniform and truncated Pareto.

[2] We highlight here that the equation gives the optimal solution of (A10) for $p(t)-p'(t)/\mu$ in the support of the valuation distribution. Otherwise, the equation is reduced to $0=0$.

Case $\alpha = 0$. We first fix $k = 1$ and $\alpha = 0$, i.e., we consider the buyer's valuation Unif[0,1]. In this case, we obtain a second order ODE in the function $p(t)$ with constant coefficients, expressed by

$$p''(t) - \delta p'(t) + \frac{\delta^2 - \delta\mu}{2} p(t) = 0.$$

Its solution is given by:

$$p(t) = c_1 e^{\frac{1}{2}t(\delta - \sqrt{-\delta(\delta - 2\mu)})} + c_2 e^{\frac{1}{2}t(\delta + \sqrt{-\delta(\delta-2\mu)})},$$

where c_1, c_2 are constants to be determined.

Note that $\delta + \sqrt{-\delta(\delta - 2\mu)} > 0$ and $\delta - \sqrt{-\delta(\delta - 2\mu)} < 0$ due to $\mu > \delta$. Therefore, the optimal pricing function is a sum of a negative exponential function and a positive exponential function and $p(t)$ could in principle go to infinity when t goes to infinity. However, $p(t) \in [0, 1]$ for all t, and then, it must be the case that $c_2 = 0$. Thus, the optimal price function is a negative exponential function of the form:

$$p(t) = c_1 e^{\frac{1}{2}t(\delta - \sqrt{-\delta(\delta - 2\mu)})},$$

with $c_1 = p(0)$.

In order to simplify the notation, define the positive constant $A = -\delta + \sqrt{-\delta(\delta - 2\mu)}$. We are left with finding the value $p(0)$. Replacing the function $p(t)$ in the unconstrained problem (A10), we can rewrite it as a maximization problem over $p(0)$ as follows:

$$\max_{p(0)} \quad \left\{ p^2(0) \int_0^\infty e^{-(\delta + A)t} \left(\frac{A}{2} + \frac{A^2}{4\mu} \right) \mathrm{d}t + p(0) \left(1 - p(0) - p(0)\frac{A}{2\mu} \right) \right\}.$$

Solving this problem, we obtain $p(0) = \frac{2\mu(\delta+A)}{(A+2\mu)(A+2\delta)}$. Noting that $p'(t) = -\frac{1}{2}Ap(0)e^{-\frac{1}{2}At}$, we also obtain $p'(0) = -\frac{A\mu(\delta+A)}{(A+2\mu)(A+2\delta)}$. Therefore, the pricing function that solves the seller's problem is given by

$$p(t) = \frac{2\mu(\delta + A)}{(A + 2\mu)(A + 2\delta)} e^{-\frac{A}{2}t},$$

with corresponding purchasing function (derived from (1))

$$\psi(t) = \frac{\delta + A}{A + 2\delta} e^{-\frac{A}{2}t}.$$

Note that in this uniform valuation case, the purchasing function turns out to be a positive multiplicative scaling of the pricing function. The optimal expected revenue of the seller in this case can be easily computed obtaining

$$R = \frac{\mu(A + \delta)}{(A + 2\mu)(A + 2\delta)}.$$

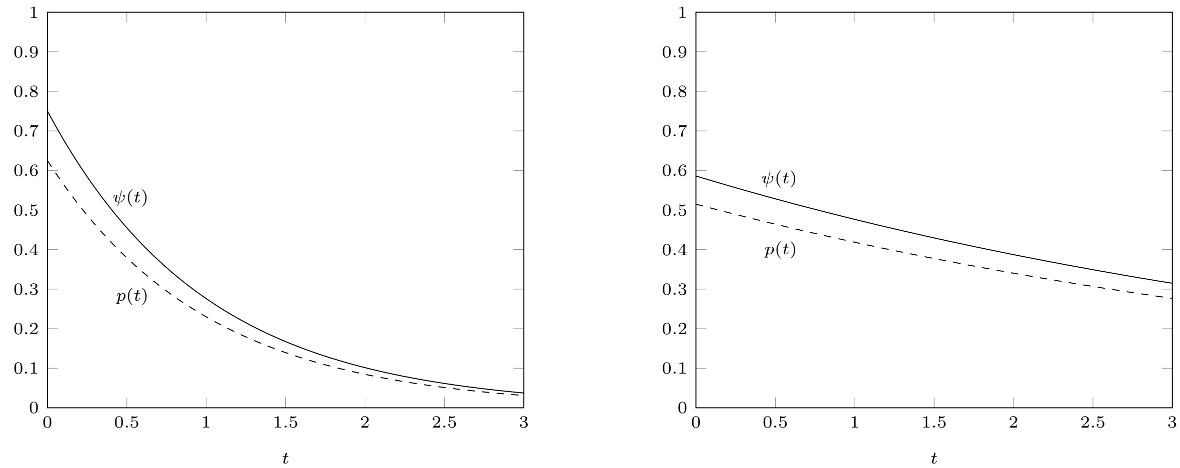

(a) Very impatient buyer ($\mu = 5$). Expected revenue: 0.3125.

(b) Slightly impatient buyer ($\mu = 1.5$). Expected revenue: 0.2574.

Figure A1 Optimal purchasing and price functions for different levels of asymmetry in the patience of the seller and the buyer. In both panels we normalize the discount rate of the seller at $\delta = 1$.

In what follows, we analyze the optimal curves obtained for some specific values of the discount rates μ and δ, corresponding to different levels of asymmetry in the patience of the seller and the buyer. Without loss of generality, we normalize the discount rate of the seller by setting $\delta = 1$.

In Figure A1, the left panel captures the case where the buyer is five times more impatient than the seller, whereas the right panel illustrates the scenario where he is only 50% more impatient. In panel (a), when the buyer is noticeably more impatient, we can observe that the optimal initial values of $p(0)$ and $p'(0)$ are greater than in panel (b), and that both price and purchasing optimal functions decrease faster. These curves reflect the fact that when facing a more impatient consumer (panel (a)), the seller will price more aggressively early in the horizon but will also drop the price relatively fast. Noting that the decreasing price pattern plays the role of a valuation discovery mechanism, the wider span of the pricing in (a) attempts to keep in the market a low valuation consumer by offering an attractive enough price relatively soon. On the contrary, when the buyer is more patient (panel (b)), the seller can offer a slow decaying price curve so that a consumer with mid to low valuation will buy later (compared to (a)) but at a higher price. The fact that the seller takes advantage of the buyer's impatience is confirmed when computing the ex-ante expected revenue by solving $[SPO_0]$ in both cases, leading to values 0.3125 and 0.2574, respectively.

Figure A2 illustrates two limit scenarios for a normalized seller's discount rate $\delta = 1$. In panel (a) we consider the case in which the buyer is extremely impatient (with $\mu = 1000$). Here, the seller drops the price very quickly from 1 to 0, charging almost instantaneously the valuation of the buyer and extracting his whole surplus. In panel (b) we present the case in which the buyer's discount rate

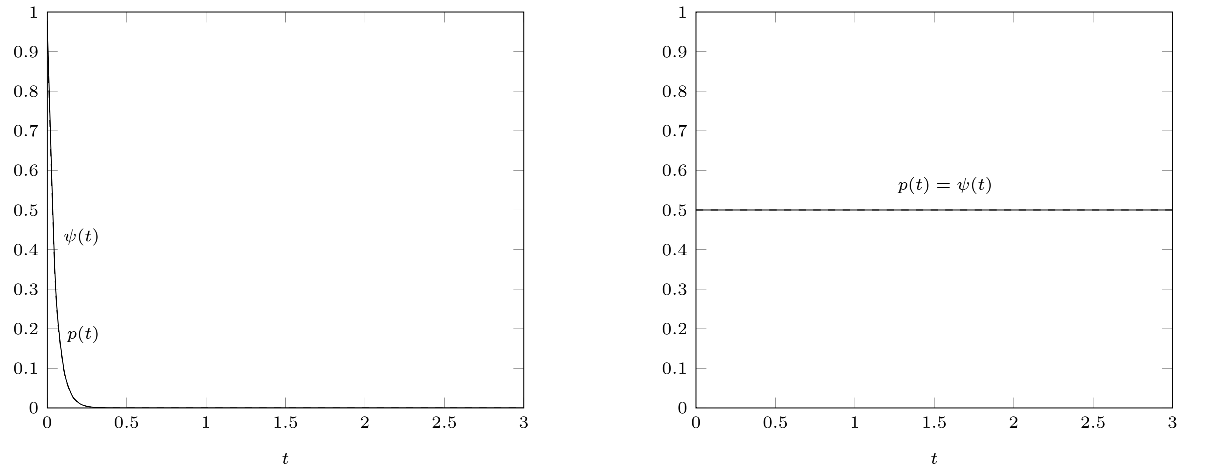

(a) Extremely impatient buyer ($\mu = 1000$). Expected revenue: 0.4786.

(b) Very patient buyer ($\mu \to \delta$). Expected revenue: 0.25.

Figure A2 Optimal purchasing and price functions for limiting asymmetries in the patience level of the seller and the buyer for different levels of asymmetry in the patience of the seller and the buyer. In both panels we normalize the discount rate of the seller at $\delta = 1$.

tends to 1. The optimal price and purchasing functions are the same and equal to 0.5 throughout the selling horizon. In this case, we recover the optimal auction of Myerson (1981), with reservation price 0.5 and the buyer purchasing at time zero if and only if his valuation is at least 0.5. In this case, he pays the reservation price for the item. The seller's advantage revenue-wise is even more emphasized, with values 0.4786 and 0.25, respectively.

Case $\alpha = -2$. We now take $k = M/(M-1)$ and $\alpha = -2$, which corresponds to the truncated Pareto distribution with parameter 1 and support $[1, M]$.

After some algebra and taking $\delta = 1$, we obtain a second order ODE in the function $p(t)$, expressed by

$$2p''(t)p(t) + p'(t)p(t) - \mu p(t)^2 - 2p'(t)^2 = 0,$$

whose solution is given by

$$p(t) = c_2 e^{\mu t + c_1 e^{-t/2}},$$

where c_1, c_2 are constants to be determined.

Therefore, the optimal pricing policy is

$$p^*(t) = \begin{cases} c_2 e^{\mu t + c_1 e^{-t/2}} & \text{if } t \leq \tilde{t} \\ 1 & \text{if } t > \tilde{t}, \end{cases}$$

where $\tilde{t}$ is such that $c_2 e^{\mu t + c_1 e^{-t/2}} = 1$. On the other hand, it must hold that $p^{*\prime}(t) = 0$, and from both conditions we obtain $c_1 = 2\mu e^{\tilde{t}/2}$ and $c_2 = e^{-\mu(\tilde{t}+2)}$.

Replacing the function $p(t)$ in the unconstrained problem (A10), we can rewrite it as a maximization problem over $\tilde{t}$ as follows:

$$\max_{\tilde{t}} \quad \left\{ \frac{1}{M-1} \left(M e^{-\tilde{t}/2} \left(1 + (2\mu - 1)(e^{-\tilde{t}/2} - 1) + \mu e^{\tilde{t}/2}(1 - e^{-\tilde{t}}) \right) - e^{-\mu(\tilde{t}+2)+2\mu e^{\tilde{t}/2}} \right) \right\}. \tag{A13}$$

We conclude that the optimal price function of the observable case when the buyer's valuation is TruncPareto(1,1,M) is given by

$$p^*(t) = \begin{cases} e^{-\mu(\tilde{t}+2-t)-2\mu e^{\tilde{t}/2-t/2}} & \text{if } t \le \tilde{t} \\ 1 & \text{if } t > \tilde{t}, \end{cases}$$

where $\tilde{t}$ is the optimal solution of problem A13.

A5. Performance of Two Heuristic Solutions of the Unobservable Case

In this section we evaluate the performance of two heuristics that may be implemented in the context of the difficult unobservable arrival case, for which an exact solution is hard to characterize: the periodic pricing policy presented in Section 5, and the simple fixed price policy.

In particular, we will assume that the buyer arrives according to an exp(1) distribution, and consider two valuation scenarios: Unif[0,1] and TruncPareto(1,1,100) –as defined in Section 6.2. We normalize the seller's discount rate to 1, and vary the buyer's discount rate μ starting from values close to 1.

In Table A1 we present the results for the uniform valuation case. More specifically, for different values of μ listed in the first column, we use the analysis in Appendix A4 to compute the optimal expected revenue of the seller in the observable case (second column). In the third column, we present the expected revenue of our periodic pricing policy in the unobservable case. Finally, we compute the ratio between the latter values and the expected revenue of the best fixed price policy, 0.125, and present this values in the fourth column of the table. [3]

From the values in Table A1 we can see that, although our policy is better than fixed price, this advantage is small when the valuation is uniformly distributed.

In Table A2 we present a similar analysis for the TruncPareto(1,1,100) distribution. We remark here that this is possible because, in this particular case, we are able to explicitly solve the observable case, as shown in Appendix A4. Here, in contrast to the uniform case, we report that our policy performs significantly better than the fixed price policy, and the gap increases as the buyer's discount rate increases. Note that in this case the optimal fixed price policy gives revenue 0.5.[4]

[3] Note that the best fixed price policy for the uniform valuation, unobservable case is given by the price p maximizing $p(1-F(p)) = p(1-p)$, which is $p = 1/2$, and therefore the total expected revenue is $\frac{1}{4}\int_0^\infty e^{-t}e^{-t}\mathrm{d}t = \frac{1}{8}$, where one of the exponential factors in the integral comes from seller's discount rate and the other one from the density function of the exp(1) valuation distribution.

[4] Note that the best fixed price policy for the unobservable case under TruncPareto(1,1,M) valuations is given by the price p maximizing $p(1-F(p)) = \frac{1}{M-1}(M-p)$, which is $p = 1$, and thus the total expected revenue is $\int_0^\infty e^{-t}e^{-t} = \frac{1}{2}$.

Discount μ	Revenue observable case optimal policy	Revenue unobservable case periodic policy	Revenue ratio periodic policy vs. fixed price
1.1	0.1257	0.1232	0.9859
1.5	0.1287	0.1252	1.0016
2.0	0.1340	0.1274	1.0192
5.0	0.1563	0.1398	1.1184
7.0	0.1650	0.1439	1.1512
10.0	0.1741	0.1477	1.1816
100.0	0.2191	0.1463	1.1704

Table A1 Comparison among expected revenues under different buyer's discount rates μ, for Unif$[0,1]$ valuation distribution and exp(1) arrival distribution. Columns: Revenue from the optimal policy in the observable case (column 2), revenue from the periodic policy in the unobservable case (column 3), and ratio between the latter and that from the fixed price policy, i.e. 0.125, in the unobservable case (column 4). The seller discount rate is normalized to $\delta = 1$.

Discount μ	Revenue observable case optimal policy	Revenue unobservable case periodic policy	Revenue ratio periodic policy vs. fixed price
1.1	0.5079	0.4669	0.9338
1.5	0.5669	0.4932	0.9864
2.0	0.6358	0.5352	1.0704
5.0	0.9017	0.7328	1.4656
7.0	1.0098	0.8048	1.6096
10.0	1.1265	0.8667	1.7334
100.0	1.7841	1.1488	2.2976

Table A2 Comparison among expected revenues under different buyer's discount rates μ, for TruncPareto(1,1,M) valuation distribution and exp(1) arrival distribution. Columns: Revenue from the optimal policy in the observable case (column 2), revenue from the periodic policy in the unobservable case (column 3), and ratio between the latter and that from the fixed price policy, i.e. 0.5, in the unobservable case (column 4). The seller discount rate is normalized to $\delta = 1$.

A6. Unobservable case with Truncated Pareto Valuation and Two Possible Arrival Times.

With the objective of characterizing the optimal solution of an unobservable arrival problem instance, suppose that the valuation of the buyer is distributed according to a TruncPareto(1,1,M), and that he arrives at one of two possible times: either at time 0 with probability β, or at time T with probability $1-\beta$, for some predetermined value $T > 0$.

We define the threshold valuation α as the value so that if the buyer arrives at time 0 with valuation $v \geq \alpha$, then he would buy before time T. This implies that, at time T, conditioned on the event that the seller has not sold the item yet, the buyer's valuation is the mixture of two truncated Pareto distributions: (i) a truncated Pareto in $[1,\alpha]$ accounting for the mass of buyers who arrived at 0 and decided to wait for a good price to be offered after T, with weight β, and (ii) a truncated Pareto in $[1,M]$ accounting for the buyer arriving at time T, with weight $1-\beta$.

Assume also that the seller's discount rate is normalized to $\delta=1$, and that the buyer discounts the future at rate $\mu>1$.

The general approach to compute the optimal expected revenue would be to decouple the problem in two independent subproblems that occur sequentially over time, by conditioning on the purchasing time of the buyer: either before or after T. Assume for now that the value of α is given. We will argue that if the buyer arrives at time zero with valuation above α, he does not have an incentive to delay his purchase beyond time T. Otherwise (i.e., he arrives at time zero with valuation below α or he arrives at time T), he will buy after time T. Hence, we can solve these two subproblems separately and then link them through the threshold α occurring at time T. Furthermore, each subproblem, once we condition on the information available at times 0 and T, corresponds to the observable case. In this regard, the threshold function ϕ of the unobservable arrival case is defined by parts combining three different purchasing functions ψ for the observable arrival case.

The first subproblem, defined over the time window $[0,T)$, is solved as in the observable case with TruncPareto(1,1,M) valuation but in a finite horizon, by adding a terminal condition for the purchasing function: $\psi(T)=\alpha$. For the second subproblem, defined over $[T,\infty)$, we first guess the time τ by which the buyer arriving at T would have purchased the item if and only if his valuation were above α, and then we solve two "observable arrival" problems assuming a truncated Pareto valuation distribution for each of them. More specifically, the problem in $[\tau,\infty)$ is solved with valuations TruncPareto(1,1,α); whereas the problem in the interval $[T,\tau)$ is solved with valuations TruncPareto(1,1,M) and with two terminal conditions: (i) the value of the purchasing time function at the boundary has to verify: $\psi(\tau)=\alpha$, and (ii) the price function has to be continuous at τ.

This whole procedure gives a price function and a purchasing function that depend on α and τ –see Figure A3–, which are then optimized to maximize the seller's revenue. Note that as we can derive explicit solutions for the truncated Pareto valuation distribution in the observable arrival case, the numerical part of the optimization to solve the unobservable case is only over these two parameters.

We provide below more details of the subproblems we need to solve to obtain the optimal price and threshold functions: one over the time interval $[0,T)$, and other over the time interval $[T,\infty)$, which in turn could be divided into two problems, with splitting time at τ, where in principle τ is assumed to be fixed. In particular, these three problems lead to three pricing policies, namely p_1, p_2 and p_3, with corresponding purchasing functions ψ_1, ψ_2 and ψ_3, as can be seen in Figure A3. After that, we prove that the problem can indeed be decoupled into these sub-problems, by showing that there are no purchasing deviations (see Propositions A3 and A4 below).

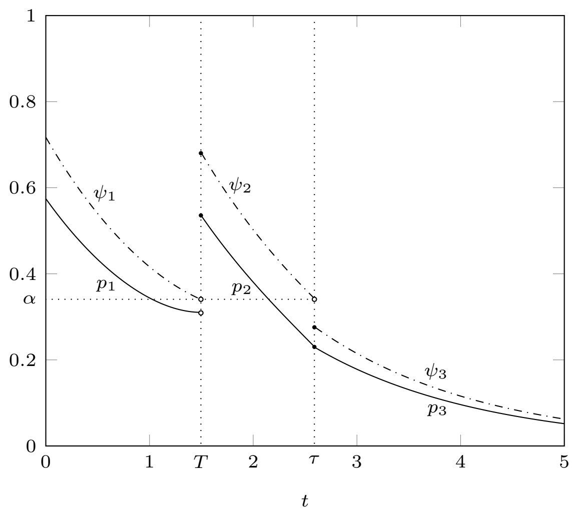

Figure A3 Price and purchasing functions for the unobservable arrival case with TruncPareto(1,0,1) and two possible arrival times: 0 and T. The threshold function ϕ is defined by three parts through respective definitions of the purchasing function ψ for the observable arrival case. In order to compel with the requirement of ϕ being lower semicontinuous, the function ψ_1 is assumed to be left continuous at some point $T - \epsilon$, for $\epsilon > 0$ arbitrarily small, and the function ψ_2 is extended to the left of T so that it is continuous at T. Note that this technical adjustment implies a negligible revenue loss from the seller's perspective.

The following steps are performed for solving a particular problem instance, defined over the parameters β, T, μ, and M. The procedure is described for fixed given values α and τ, leading to a pricing policy p and a threshold function ϕ that depend on both values. Then, both α and τ are optimized to maximize the revenues of the specific problem instance.

Step 1. Compute the pricing policy p_1 to offer in the time interval $[0, T)$ and its associated purchasing function ψ_1. To this end, we need to solve the problem $[SPO_0^r]$ affected by the probability β that the buyer indeed arrives at time 0, i.e.,

$$\max_{p} \quad \beta p_1(0)(1 - F(\psi_1(0))) + \beta \int_0^T e^{-t} p_1(t) f(\psi_1(t))(-\psi_1'(t))\,\mathrm{d}t,$$

where $\psi_1(t)$ is defined in (1), and where we further impose the boundary condition $\psi_1(T) = \alpha$.

Step 2A. Compute the pricing policy p_3 to offer in the time interval $[\tau, \infty)$ and its associated purchasing function ψ_3. To this end, we need to solve the problem $[SPO_0^r]$ but with origin at time τ, with valuations at most α, i.e.,:

$$\max_{p} \int_\tau^\infty e^{-t} p_3(t) f(\psi_3(t))(-\psi_3'(t))\,\mathrm{d}t + e^{-\tau} p_3(\tau)(F(\alpha) - F(\psi_3(\tau))),$$

where the second term represents the discounted expected revenue obtained at time τ from the mass of valuations between $\psi_3(\tau)$ and α who arrived at time 0 or at time T,[5] and the term involving the integral represents the discounted expected revenue obtained along the interval $[\tau,\infty)$.

Step 2B. Compute the pricing policy p_2 to offer in the time interval $[T,\tau)$ and its associated purchasing function ψ_2. To this end, we need to solve the problem $[SPO_0^r]$ affected by the probability $1-\beta$ that the buyer arrives at time T, i.e.,

$$\max_p \quad (1-\beta)e^{-T}p_2(T)(1-F(\psi_2(T)))+(1-\beta)\int_T^\tau e^{-t}p_2(t)f(\psi_2(t))(-\psi_2'(t))\,\mathrm{d}t,$$

where $\psi_2(t)$ is defined in (1), and where we further impose two boundary conditions: $\psi_2(\tau)=\alpha$, and $p_2(\tau)=p_3(\tau)$.[6]

With the solutions of the three steps above, we can define the pricing policy p and the threshold function ϕ for the whole horizon. To ensure that these functions solve the unobservable case for this particular instance, it only remains to prove that if the buyer arrives at time 0 with valuation $v>\alpha$, then he will indeed buy before time T. To this end, we show a preliminary result stating that if a buyer with valuation $\psi_1(\eta)$ for some $\eta\in[0,T)$ buys after T, then he will also buy after T if he has valuation belonging to the interval $(\alpha,\psi_1(\eta))$, where ψ_1 represents the solution of the problem described in Step 1.

Proposition A3. *If there exists $\eta\in[0,T)$ and $\eta'>T$ such that $U(\eta,\psi_1(\eta))<U(\eta',\psi_1(\eta))$, then for all t belonging to the interval $(\eta,T]$, there exists $t'>T$ such that $U(t,\psi_1(t))<U(t',\psi_1(t))$.*

Proof. By contradiction, suppose that there exist $t\in(\eta,T]$ such that $U(t,\psi_1(t))\geq U(t',\psi_1(t))$ for all $t'>T$. In particular, the inequality holds for $t'=\eta'$, and then we have $U(t,\psi_1(t))\geq U(\eta',\psi_1(t))$.

Define $\varepsilon=\psi_1(\eta)-\psi_1(t)$. Note that by hypothesis $\eta<t$, and thus, from the monotonicity of ψ_1 (see Proposition 1 in Section 4), we obtain $\varepsilon\geq 0$, which implies

$$U(t,\psi_1(\eta))=e^{-\mu t}(\psi_1(t)+\varepsilon-p(t))=U(t,\psi_1(t))+\varepsilon e^{-\mu t}\geq U(\eta',\psi_1(t))+\varepsilon e^{-\mu\eta'}=U(\eta',\psi_1(t)),$$

where the inequality holds due to the contradiction hypothesis for $t'=\eta'$ and because $\eta'>t$, and the last equality holds because $\psi_1(\eta)=\psi_1(t)+\varepsilon$. We conclude that $U(t,\psi_1(\eta))\geq U(\eta',\psi_1(\eta))$.

[5] This mass of valuations is composed of $\beta(F(\alpha)-F(\psi_3(\tau)))$ coming from time 0, and $(1-\beta)(F(\alpha)-F(\psi_3(\tau)))$ coming from time T. Recall that no matter the arrival time, the buyer always discounts his utility since time 0.

[6] Note that the price function has to be continuous at τ. Otherwise, knowing that p is non increasing (due to the observation in Section 3.2.1 together with Proposition 2), if it were the case that $\lim_{t\to\tau^-}p_2(\tau)>p_3(\tau)$, then a buyer with valuation $\psi_2(\tau-\epsilon)>\alpha$ would have an incentive to wait and buy at τ and take advantage of a price decreased by a non negligible amount, which would contradict the definition of α as the threshold valuation above which a buyer would purchase before τ.

On the other hand, by hypothesis we know that $U(\eta', \psi_1(\eta)) > U(\eta, \psi_1(\eta))$, and putting all together we conclude that $U(t, \psi_1(\eta)) > U(\eta, \psi_1(\eta))$ which contradicts the definition of the purchasing function ψ_1, completing the proof. □

Let us use this result to prove that if the buyer arrives at time 0 with valuation greater than α, then he does not have an incentive to buy later than T. This result justifies the decoupling of the problem that occurs at time T and that separates Step 1 from Step 2.

PROPOSITION A4. *If the buyer arrives at time* 0 *with valuation* $v > \alpha$, *then there exists an optimal pricing policy under which he purchases at some time before* T.

Proof. By contradiction, suppose that there exist $t \in (\eta, T]$ such that $U(t, \psi_1(t)) \geq U(t', \psi_1(t))$ for all $t' > T$. In particular, the inequality holds for $t' = \eta'$, and then we have $U(t, \psi_1(t)) \geq U(\eta', \psi_1(t))$.

Define $\varepsilon = \psi_1(\eta) - \psi_1(t)$. Note that by hypothesis $\eta < t$, and thus, from the monotonicity of ψ_1 (see Proposition 1 in Section 4), we obtain $\varepsilon \geq 0$, which implies

$$U(t, \psi_1(\eta)) = e^{-\mu t}(\psi_1(t) + \varepsilon - p(t)) = U(t, \psi_1(t)) + \varepsilon e^{-\mu t} \geq U(\eta', \psi_1(t)) + \varepsilon e^{-\mu \eta'} = U(\eta', \psi_1(t)),$$

where the inequality holds due to the contradiction hypothesis for $t' = \eta'$ and because $\eta' > t$, and the last equality holds because $\psi_1(\eta) = \psi_1(t) + \varepsilon$. We conclude that $U(t, \psi_1(\eta)) \geq U(\eta', \psi_1(\eta))$.

On the other hand, by hypothesis we know that $U(\eta', \psi_1(\eta)) > U(\eta, \psi_1(\eta))$, and putting all together we conclude that $U(t, \psi_1(\eta)) > U(\eta, \psi_1(\eta))$ which contradicts the definition of the purchasing function ψ_1, completing the proof. □

In summary, even though the unobservable arrival case is hard to solve in general, in the particular scenario with truncated Pareto distribution and two possible arrival times, we could solve it by simplifying its formulation to a sequence of observable arrival cases. We highlight here that the same argument can be used to solve the unobservable case with two arrival times for any setting for which the observable case can be solved, as it is the case of the uniform distribution.

A7. A Lower Bound for the Value of Observability.

In order to get a lower bound for the value of observability it is enough to get the ratio for one particular problem instance. The challenge here stems from the difficulty in solving the unobservable case, even numerically. In order to partially overcome this difficulty, we consider the particular problem instance introduced in Appendix A6: TruncPareto$(1, 1, M)$ valuations and two possible arrival times: 0 and T, with probability β and $1 - \beta$, respectively. Remember that the policy is further characterized by two additional parameters: (i) α, representing the value so that if the buyer arrives at time 0 with valuation $v \geq \alpha$, then he would buy before time T, and (ii) τ, the time by which a buyer arriving at T would have purchased the item if and only if his valuation were above α.

Setting the parameters $M = 100$ and $T = 1$, and normalizing the discount rate of the seller at $\delta = 1$, we numerically compute the value of observability for different buyer's discount rates μ and probabilities β of arriving at time 0. The results are exhibited in Table A3, from where we note that the worst case over these instances is the one determined by $\mu = 4.5$ and $\beta = 0.4$, leading to a lower bound of 1.136 for the value of observability.

μ \ β	0.2	0.4	0.6	0.8
1.0	1	1	1	1
1.5	1.036	1.017	1.005	1.001
2.0	1.074	1.047	1.020	1.001
2.5	1.111	1.069	1.038	1.004
3.0	1.128	1.089	1.048	1.013
3.5	1.124	1.107	1.058	1.020
4.0	1.115	1.124	1.066	1.025
4.5	1.103	**1.136**	1.074	1.029
5.0	1.091	1.121	1.094	1.064
5.5	1.078	1.103	1.075	1.044
6.0	1.065	1.084	1.054	1.022

Table A3 Value of Observability for a two-point arrival time distribution: 0 with probability β, and $T = 1$ with probability $1 - \beta$, assuming TruncPareto(1,1,100) valuation and seller's discount rate $\delta = 1$.