The Value of Observability in Dynamic Pricing

JOSÉ R. CORREA, Department of Industrial Engineering, Universidad de Chile
DANA PIZARRO, Department of Industrial Engineering, Universidad de Chile
GUSTAVO VULCANO, School of Business, Universidad Torcuato di Tella, and CONICET

Research on dynamic pricing has been growing during the last four decades due to its use in practice by a variety of companies as well as the several model variants that can be considered. In this work, we consider the particular pricing problem where a firm wants to sell one item to a single buyer in order to maximize expected revenues. The firm commits to a price function over an infinite horizon. The buyer has a private value for the item and purchases at the time when his utility is maximized. In our model, the buyer is more impatient than the seller and we study how important is to observe the buyer time arrival in terms of the seller's expected revenue. When the seller can observe the arrival of the buyer, she can make the price function contingent on the buyer's arrival time. On the contrary, when the seller cannot observe the arrival, her price function is fixed at time zero for the whole horizon. The *value of observability* 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. Our main result is to prove that in a very general setting, the value of observability is at most 4.911. To obtain this result we fully characterize the observable setting and use this solution to construct a random and periodic price function for the unobservable case.

CCS Concepts: • **Social and professional topics** → *Pricing and resource allocation*; • **Theory of computation** → **Algorithmic mechanism design**; **Computational pricing and auctions**.

Additional Key Words and Phrases: Pricing, Mechanism Design, Value of Observability

ACM Reference Format:
José R. Correa, Dana Pizarro, and Gustavo Vulcano. 2020. The Value of Observability in Dynamic Pricing. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC '20), July 13–17, 2020, Virtual Event, Hungary.* ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3391403.3399489

1 INTRODUCTION

In recent years we have witnessed an enormous amount of work in dynamic pricing and dynamic mechanism design. Driven by the increasingly important online marketplaces, the area has been particularly active in Economics, Operations Management and Computer Science.

Typically, the literature studies a game between a seller and one or more buyers. The seller owns an item or a set of items, and the buyers have private valuations for them [7, 8, 16]. The game takes place over a time interval since either the buyers arrive over time [1, 2, 4, 5, 9, 11, 13–15, 17, 21–23], or since the buyers and the seller discount the future differently which implies that there is delay on trade [18, 24, 25]. Other assumptions to explain dynamic trading are that valuations may change over time [14, 23] or that buyers may be short-lived [22]. The goal of the seller is to set up a pricing mechanism so as to maximize revenue (or welfare) while the buyers respond to this pricing strategically and decide to buy at the time maximizing their (discounted) utility. Although the borders are blurred, often research in operations management deals with finding optimal or

EC '20, July 13–17, 2020, Virtual Event, Hungary

ACM ISBN 978-1-4503-7975-5/20/07…$15.00
https://doi.org/10.1145/3391403.3399489

approximately optimal dynamic pricing mechanisms (e.g. [2, 6, 15]), in economics the central interest is to find optimal dynamic mechanisms (e.g. [4, 23]) which may involve departing from basic pricing schemes, while in computer science the interest is in designing simple mechanisms which are approximately optimal (e.g. [3, 8, 10]).

One drawback of part of the literature is the underlying assumption that the seller is informed of the buyers' arrivals. This assumption allows the seller to update the pricing/mechanism when observing a new arrival. In some contexts, such as in online marketplaces, it may be difficult for the seller to distinguish interested buyers from other traffic on the website and therefore assuming that the seller observes the buyers' arrivals may not be realistic. The extent to which this observational ability produces additional rents to the seller is the main subject of this paper.

Specifically, we consider a simple, yet fundamental, model in which one seller interacts with a single buyer. The seller holds a single item whose value is normalized to zero, while the buyer has a private random valuation for the item. The buyer arrives according to an arbitrary distribution over the nonnegative reals. As usual in the literature, 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 goal of the seller is to set up a price function so as to maximize her expected discounted revenue. On the other hand, upon arrival, the buyer observes the price function and decides to buy at the time that is most profitable for him. We distinguish two basic situations. In the *observable* case, the buyer has the ability to actually observe the buyer's arrival and thus the price function she sets may be dependent on the buyer's arrival time. In the *unobservable* case this ability is absent and therefore the seller has to set a price curve from the beginning only knowing the arrival distribution. These two scenarios naturally lead to define the *value of observability*, $\mathrm{VO}(G, F, \delta, \mu)$, of an instance of the problem with arrival distribution G, valuation distribution F, and discount rates δ, μ for the seller and buyer, respectively, as the ratio between the revenue of the seller in the observable case and that in the unobservable case. Accordingly, the *value of observability* (VO) is defined as the supremum of the latter quantity over all possibles G, F, δ, μ, namely $\mathrm{VO} = \sup_{G,F,\delta,\mu} \mathrm{VO}(G, F, \delta, \mu)$. Of course the definition can immediately be extended to other pricing problems.

Motivating Example

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

To better grasp this difficulty and the difference between the observable and unobservable cases let us describe 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 the end the buyer is myopic, he will buy as long as the price is below his valuation). Then, if the seller can observe the buyer's arrival in our dynamic pricing setting, she will start pricing at 1 and then decrease the price suddenly in a continuous fashion until hitting the customer valuation, where 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 $\int_0^\infty (e^{-t}/2)e^{-t}\mathrm{d}t = 1/4$ (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 decreasing price function then the problem is relatively easy to solve. Indeed, the seller would need to maximize, over all decreasing functions $p(\cdot)$, the quantity $\int_0^\infty (e^{-t}(1-p(t)) - (1-e^{-t})p'(t))e^{-t}p(t)\mathrm{d}t$. Note that for a decreasing $p(t)$, trade occurs between t and $t + dt$ if either the buyer arrived 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 $p(t+dt)$ (hence the term $-(1-e^{-t})p'(t)$). In both cases the discounted revenue for the seller is $e^{-t}p(t)$. The solution of this problem turns out to be $p(t) = e^{-t}$, which results in an expected revenue of $1/6$. Overall the ratio of the revenues between the observable and non observable cases is 3/2, and therefore this suggests that $VO \geq 3/2$. However, the seller's strategy space is richer than that of decreasing price functions. Suppose that she splits the time horizon into short intervals of length ϵ and considers 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 1 is close to 1. Thus, even in the unobservable case the seller is able to obtain a revenue arbitrarily close to 1/4.

Furthermore, when the discount rate of the buyer is not too large strategic behavior comes into play, which adds an additional layer of difficulty in formulating the problem, as we discuss in Section 2.2. It should be noted that if both the buyer and seller discount at equal rates then the optimal pricing function is simply constant in both the observable and unobservable cases, therefore strategic behavior vanishes, there is no delay on trade, and the resulting expected revenues for the seller are equal. Thus we assume throughout that the buyer's discount rate is strictly larger than that of the seller.

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 several factors: (i) the generality of the model; (ii) the bound is totally independent of the model primitives; and (iii) simple pricing strategies, such as fixed pricing, fail to guarantee a constant bound.

En route to this result we first revisit the observable case. It is worth mentioning that this problem is far from new and indeed already Stokey [25] notes that intertemporal price discrimination happens only due to the difference in discount rates. Later Landsberger and Meilijson [18] precisely show that this price discrimination through time is optimal, while Shneyerov [24] considers the situation in which there are multiple units to sell. In this context we take a pricing approach (rather than a mechanism design) which, as usual in this literature, allows us to write the seller's problem as an optimal control problem and furthermore to fully characterize its solution. In particular, we can prove a key result (Lemma 3.3) establishing that in the optimal pricing the seller extracts a constant fraction of the total revenue within a short time, 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, in order to prove that the value of observability is bounded by a constant it is enough to exhibit a pricing policy that can recover a constant fraction of the revenue of the optimal solution in the observable case. To this end we use the solution of the observable case and try to repeat it over time to contract a periodic price function. Of course this is not possible since already that solution takes infinite amount of time to implement. Thus the aforementioned key result comes into play and allows us to do this repeated pricing within small time windows. The second obstacle is that we should be careful with the buyer's strategic behavior. To avoid this issue, 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, a difficulty arises since the arrival distribution might now impose a lot of weight in regions where our price 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 are able to show that the proposed pricing scheme obtains an expected revenue of at least a fraction $1/4.911 = 0.203$ of the optimal revenue in the observable case.

For the special and relevant case of valuation distributions having monotone hazard rate, which includes several of the standard distributions, we show that the situation is much simpler. Indeed it is enough to consider a fixed price curve (i.e., the price is constant over the whole period) to recover a fraction $1/e$ of the revenue in the observable case. We further note that fixed pricing cannot guarantee a constant in general.

We also note that our result is robust to the distribution of arrivals. Indeed, 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.

Roadmap. We start with the precise model description in Section 2, including describing the buyer's problem in Section 2.1 and the seller's problem in Section 2.2. This latter section includes the formulations of the standard observable case and the more challenging unobservable case. Both cases are later analyzed in detail in Sections 3 and 4, respectively. Finally, the bound for the VO is established in Section 5.

2 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 buyer has a private valuation v for the item with cdf $F : [0, \bar{v}] \to [0, 1]$ and density f. Both G and F are common knowledge[1].

The interaction between the seller and the buyer is formalized as a Stackelberg game in which the seller is the leader and pre-commits 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.

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 onwards 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 $p : [0, \infty] \to [0, \bar{v}]$.

Even though for the ease of exposition the game between the seller and the buyer is presented as if the seller were to announce the price function in a first stage, and the buyer were going to decide *if* and *when* to purchase in the second stage, strictly speaking, the game can also be described as a simultaneous game with no need of precommitment since the calculation of the price function and the timing of the buyer's purchase decision are based on common knowledge information.

For technical reasons, in both cases we impose the mild condition that the price function p is lower semi-continuous and differentiable almost everywhere[2]. In what follows, we introduce the buyer's and the seller's problems, as well as some preliminary definitions and results.

[1] As it is standard in the literature we can equivalently think that the seller has unlimited supply, and that on the consumer side we have a mass of consumers with arrivals distributed as G and valuation distributed as F.

[2] 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.

2.1 The buyer's problem.

When the buyer arrives, he observes the price function for all future times 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. In particular, we consider an exponentially discounted utility function: $U(t, v) = e^{-\mu t}(v - p(t))$, where $\mu > 0$ is the discount factor. Following the standard assumption in the literature, the buyer is more impatient than the seller and hence $\mu > \delta$, with δ being the discount factor of the seller. Note that as $t \to \infty$, $U(t, v) \to 0$, so that the buyer eventually purchases the item as long as $v > p(t)$, for some t. Given a price function $p(t)$, a forward-looking buyer arriving at time τ with valuation v solves:

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

It may be possible that this 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. We define the auxiliary function $\phi : [0, \infty) \to [0, \overline{v}]$ as:

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

that represents the minimum valuation that the buyer must have in order to buy at time t and not later, and it is defined irrespective of the buyer's arrival time $\tau \leq t$. 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 at the same time and not later[3].

After defining ϕ we are able to describe the equilibrium conditions for the buyer purchasing behavior and used them to formulate the seller's problem.

2.2 The seller's problem

The seller's problem is to select a price function to maximize her expected revenue, taking into account the forward-looking behavior of the buyer.

2.2.1 Observable arrival case. In this situation the seller observes the arrival time τ of the buyer and therefore sets a price function p defined over $[\tau, \infty)$. For now, we will pretend that the buyer arrives at time zero, i.e., we initially assume that $\tau = 0$.

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 $v \geq \phi(t)$. In this observable case, the buyer's purchasing behavior could be better represented by resorting to the auxiliary function $\psi(t)$, defined as

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

In other words, a customer arriving at time zero with valuation $\psi(t)$ will buy at time t. Due to the lower semi-continuity of p we have that ϕ is also lower semi-continuous and therefore, ψ is well defined (as we prove in the full version of the paper). 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 $d(1 - F(\psi(t)))$. With this observation we may write the

[3]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.

seller's problem conditioned on the event that the buyer arrives at time 0:

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

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 $s = t$.

Note that every p feasible solution of the problem $[SPO_0]$ must be non increasing. Otherwise, there would exist $t > s > 0$ such that $p(t) > p(s) > 0$. Thus, $\psi(t) - p(s) > \psi(t) - p(t)$ and $e^{-\delta s} > e^{-\delta t} > 0$, and therefore, $U(s, \psi(t)) > U(t, \psi(t))$, which contradicts the definition of t in the constraint of $[SPO_0]$.

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 ex-ante maximum expected revenue of the seller can be written as $R = R_0 \int_0^\infty e^{-\delta\tau} g(\tau) \mathrm{d}\tau$, 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 seller's optimal pricing policy.

2.2.2 Unobservable arrival case. When the seller does not observe the buyer's arrival time, the price function that she has to set can only depend on the arrival time distribution G.

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 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:

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

The term in brackets stands for the probability of purchasing at time t. Within it, the first term $(1 - F(\phi(t)))g(t)$ represents the probability of arriving at time t with valuation $v \geq \phi(t)$ and hence purchasing immediately. This corresponds to the points in the vertical line in Figure 1(b); that is, we are accounting for a customer arriving in t_1 with valuation $v \geq v_1$.

The second term, $(G(t) - G(s_t))(1 - F(\phi(t)))'$, is the probability of purchasing at time t when arriving at any time between s_t and t with valuation $\phi(t)$, that is, the probability of being in the line connecting $\phi(s_{t_1})$ and $\phi(t_1)$ 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 the latter term only holds at points where ϕ is decreasing. The description of this optimization problem is included for completeness, but strictly speaking we will not solve it in our forthcoming development, but rather we would

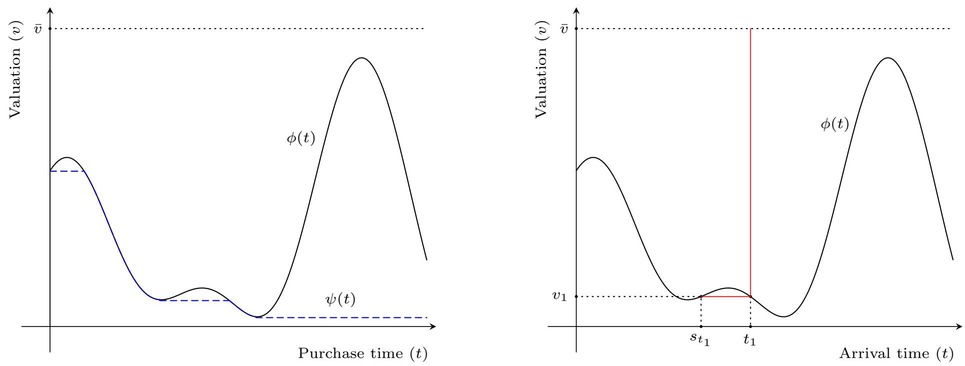

(a) Observable case. Definition of the function $\psi(t)$. For a given function $\phi(t)$, a customer with valuation $\psi(t)$ arriving at $\tau = 0$ will buy at time t.

(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.

Fig. 1. Consumer purchasing behavior

focus in a feasible pricing policy that would allow us to bound the ratio between the revenue from $[SPO_0]$ and $[SPN]$.

3 ANALYSIS OF THE MODEL WITH AN OBSERVABLE ARRIVAL

Given the argument stated in Section 2.2.1, to analyze the observable case it is sufficient to focus on the solution of $[SPO_0]$, where the buyer arrives at time 0. This problem has already been studied in the economics literature from both a mechanism design approach [18, 24] and a pricing approach [26]. In particular Shneyerov [24] considers a very similar situation but the characterization of the optimal price function is done through the maximum principle and therefore it involves a rather complicated hamiltonian. Our approach is simpler, and based on the Euler-Lagrange optimality conditions we can derive a simpler ODE which we prove has a unique solution.

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 for at least for F being a uniform distribution.

To begin with, consider the incentive compatible constraint in problem $[SPO_0]$. If $t^* > 0$ 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:

$$[SPO_0^r] \quad \max_{p,\psi} \quad \int_0^\infty e^{-\delta t} p(t)(-\psi'(t)) f(\psi(t))\, dt + p(0)(1 - F(\psi(0)))$$

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

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 $p(t)$:

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

Letting the integrand function be $G(t, p(t), p'(t), p''(t))$ and the expected revenue at time zero be r_0, problem (1) is equivalent to:

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

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

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

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

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

After some algebra (detailed in the full version of the paper) the Euler-Lagrange equation becomes:

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

Of course, 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 (e.g. Theorem 20.9 in [20]) 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. These conditions hold if we for instance assume that $p(0) > 0$ and $p'(0) < 0$. While the former is natural to assume, the latter makes sense in the context of this observable case with price commitment, where a forward-looking consumer will never buy within an ϵ-interval starting at zero if the price is non decreasing at zero. Therefore, for a large set of valuation distributions, we have that the relaxed problem has exactly one solution.

Let us highlight that though we know that in the observable case $\psi(t)$ 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 $\psi(t)$. However, the following result, which we prove in the full version of the paper, 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 decreasing function.

Proposition 3.1. *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.*

Proposition 3.1, along with the upper bound defined by the solution to $[SPO_0^r]$, allow us to show that any solution of $[SPO_0^r]$ also solves the seller's problem $[SPO_0]$.

THEOREM 3.2. *Any solution of the relaxed problem* $[SPO_0^r]$ *such that* p *is differentiable with continuous derivative also solves the seller's problem* $[SPO_0]$.

PROOF. 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:

$$t \in \arg\max_{s\geq 0}\ e^{-\mu s}(\psi(t) - p(s))\ \ \forall t. \tag{3}$$

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 (3), it must verify $h'(t) = 0$ and

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

Since $(p(t), \psi(t))$ is solution of $[SPO_0^r]$, then from Proposition 3.1 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 $[SPO_0]$. □

Theorem 3.2 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 $[SPO_0]$, first we formulate the Euler-Lagrange equation (2) and solve it. Its solution will depend on the initial values $p(0) > 0$ and $p'(0) < 0$. Then, we replace that solution in problem (1) 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 $[SPO_0]$.

To conclude this subsection, we present a following technical result that states that if for a given parameter $c \in (0, 1)$, 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 $T = \ln(1/c)/\delta$.

LEMMA 3.3. *For a given parameter* $c \in (0, 1)$*, up to time* $T = \ln(1/c)/\delta$*, the seller's expected revenue in the observable arrival case is at least* $(1 - c)R_0$*; i.e.,*

$$\int_0^T e^{-\delta t} p(t) \mathrm{d}(1 - F(\psi(t))) \geq (1 - c)R_0,$$

where $p(t)$ *is the solution from* (2) *to the observable case problem.*

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

$$\int_T^\infty e^{-\delta t} p(t) \mathrm{d}(1 - F(\psi(t))) > c\left[p(0)(1 - F(\psi(0))) + \int_0^\infty e^{-\delta t} p(t) \mathrm{d}(1 - F(\psi(t)))\right]. \tag{4}$$

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:

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

Applying (4), 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. □

For instance, if we want to reach at least half of R_0 and we normalize the seller's rate discount to 1, from this result we conclude that it is enough to consider the problem until $T = \ln(2)$. This implies that the time needed to get a big fraction of R_0 is relatively small and, moreover, it does not depend on the valuation distribution.

4 ANALYSIS OF THE MODEL WITH UNOBSERVABLE ARRIVAL

Consider now the problem stated in Section 2.2.2 where the seller is not able to observe the arrival time of the buyer. Different from the previous observable case, where the seller knows the arrival time τ of the buyer and sets the price function $p(t)$ 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.

This problem turns out to be very difficult in the general case. To partially overcome, we will focus on analyzing the seller's problem under a feasible pricing policy, with the objective of bounding the *value of observability*; that is, the ratio between the expected revenues under the observable case $[SPO_\tau]$ and the unobservable case $[SPN]$.

Our main result states that under a general valuation distribution, the *value of observability* is upper bounded by 4.911. However, in the case where the valuation distribution is monotone hazard rate, the bound is improved to $e \approx 2.718$. To ease the exposition we first prove the latter result.

4.1 Monotone hazard rate valuation distribution

The case of monotone hazard rate valuation distribution turns out to be quite simple. We start this section by reviewing some basic concepts on the theory of optimal auctions introduced in the seminal work of Myerson [19]. A key building blocks 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 $p^* = J^{-1}(0)$.[4]

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 distributions 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.

Indeed, we know from Section 2.2.1 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 $R_0 \leq \mathbb{E}(v)$, the expected value of the valuation drawn from F. Hence, the seller's expected revenue in the observable case is upper bounded by

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

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 is at least

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

Finally, by Lemma 3.10 (p.325) of Dhangwatnotai et al. [12], it follows that $p^*(1 - F(p^*)) \geq \frac{1}{e}\mathbb{E}(v)$, and this the claimed bound follows.

The bound is tight in the case of exponentially distributed valuation ($F(v) = 1 - e^{-v}$) and a myopic buyer (with $\mu = \infty$), and when the seller does not discount revenues (i.e., $\delta = 0$). In this setting, in the observable case, the seller will announce a price curve that spans all the support $[0, \bar{v}]$ (e.g., $p(t) = 1/t$), and the consumer will buy immediately when his valuation $v = p(t)$. In this case, the ex-ante expected revenue is $\mathbb{E}(v) = 1$. In order to get the revenue for the unobservable case, the seller will offer a fixed price p maximizing $p(1 - F(p)) = pe^{-p}$. This function is maximized at $p = 1$ with optimal revenue e^{-1}.

4.2 A feasible periodic price function

We start by noting that using fixed pricing does not work in general. For instance, consider the game where the buyer's valuation is distributed according to a truncated Pareto distribution with parameter 1, that is, with cdf $F(x) = (1 - 1/x)M/(M - 1)$ for $x \in [1, M]$, and again $\mu = \infty$ and $\delta = 0$. Here, we have that the expected value of the buyer's valuation is $M \ln M/(M - 1)$ whereas $p^*(1 - F(p^*)) = M/(M - 1)$, leading to the ratio $\mathbb{E}(v)/r^*(1 - F(r^*)) = \ln M$ growing with M. Note then that the ratio grows arbitrarily large independent on the arrival distribution.

[4] 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.

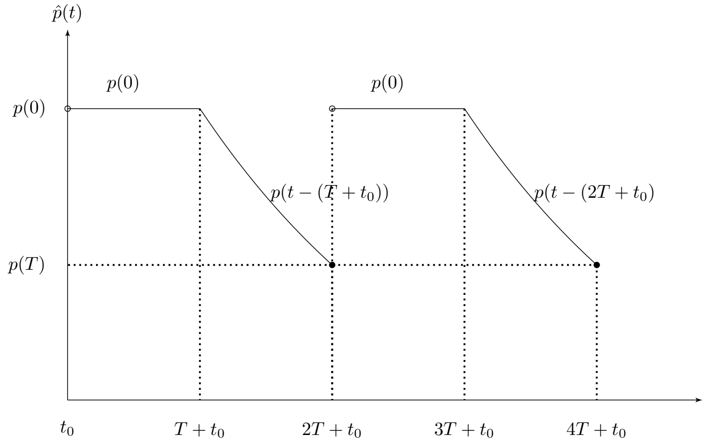

Fig. 2. Periodic pricing policy $\hat{p}$ after performing a random shift and setting the origin at time t_0

Thus, to bound the value of observability in the general case we need to consider a pricing policy that allows us to compare the expected revenue in the observable and unobservable case. We define it in Section 4.2 and present our main result in Section 4.3.

The feasible pricing policy $\hat{p}$ we consider is periodic and depends on the optimal pricing policy p of $[SPO_0]$. 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 big enough. In particular, the price function we use to bound the seller's expected revenue in the unobservable case is defined by

$$\hat{p}(t) = \begin{cases} p(0) & \text{if } t \in I_{2k-1},\ k \in \mathbb{N} \\ p(t-(2k-1)T) & \text{if } t \in I_{2k},\ k \in \mathbb{N} \end{cases} \tag{5}$$

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 constant price $p(0)$ 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 $t_0 \geq 0$. The fact of having a time origin set at t_0 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 $[SPN]$. In order to perform the analysis independently of the specific function g, let us first observe that by doing a random shift on the price function we can assume without loss of generality that the buyer's arrival time is uniformly distributed within a period of length $2T$.

More formally, suppose that we have a periodic function h with period $2T$ and consider a random shift, that is, for a random variable $t_0 \sim \text{Unif}[0, 2T]$, consider the function $\hat{h}(t) = h(t + t_0)$. 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

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

which proves that X is uniformly distributed in $I_{2k-1} \cup I_{2k}$. Therefore, by applying a random shift over the function p to obtain $\hat{p}$, we can assume that buyer's arrival, conditional on the arrival 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 $2T$.

4.3 Revenue analysis

Along the rest of the paper we will 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 a translation of the function p after performing the random shift.

We start by providing a simple lower bound for the seller's revenue within a limited time frame in the unobservable arrival case.

Lemma 4.1. *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 $v \ge p(T)$, 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 $(2k-1)T + t_0$.*

Proof. 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 $v \ge p(T)$, and further assume that he will not purchase before time $T + t_0$.

To prove the lemma we analyze the consumer behaviour 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 $v \ge \psi(0)$ 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.,

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

which means that

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

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 $T + t_0$. □

5 BOUNDING THE VALUE OF OBSERVABILITY

We are now able to bound the value of observing arrivals by considering the particular pricing policy $\hat{p}$ in (5) to give a lower bound of the seller's expected revenue in the unobservable case. To do so, we will only consider the buyer with valuation $v \geq p(T)$ when he arrives in an interval where the price is constant. We then obtain our main result which states that the value of observing the arrival is at most roughly 4.91.

This bound can be written as a function of W_{-1}, the negative branch of the well known Lambert function[5].

Theorem 5.1. *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. We will compute a lower bound of the seller's expected revenue for the unobservable case. For that purpose, 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 $\hat{p}$. By defining $T = \ln(1/c)/\delta$, the price function has period $2T$.

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

Let us denote by R^{uo}_τ the seller's revenue in the unobservable case if the arrival time is τ. We 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 the expected value of 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 by waiting up to $t_0 + T$ to buy when he could buy would hurt the seller's revenue given her discount factor.

[5]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.*

Then, by Lemma 4.1, R_τ^{uo} is at least the expected revenue earned up to time $2T+t_0$ in the observable case with arrival time $T+t_0$. Applying Lemma 3.3, we have $R_\tau^{uo} \geq (1-c)R_{t_0+T}$, where R_{t_0+T} denotes the expected revenue in the observable case if the buyer arrives at time 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.

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

where the last equality holds because $t_0 \sim \mathrm{Unif}[\tau - 2T, \tau]$ and the inequality follows from the analysis above. Note that $R_{t_0+T} = ce^{-\delta(t_0-\tau)}R_\tau$, with $e^{-\delta T} = c$, and therefore it is enough to compute $\mathbb{E}\left(e^{-\delta(t_0-\tau)} \mid t_0 \in \mathcal{I}_1\right)$. In fact,

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

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

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

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

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

Noting that R_τ is the expected value of the seller's revenue in the observable case with buyer's time arrival τ, follows that for each time arrival τ, the ratio between the expected revenue in the observable and the unobservable case is at most

$$\frac{\mathbb{E}(\text{Rev. Obs} \mid \tau)}{\mathbb{E}(\text{Rev. Unobs} \mid \tau)} \leq \frac{2\ln(1/c)}{(1-c)^2}.$$

The latter expression is minimized at $c = e^{W_{-1}(-1/(2\sqrt{e}))+1/2} \approx 0.284$ and the minimum is $-\frac{2W_{-1}(-1/(2\sqrt{e}))+1}{\left(e^{W_{-1}(-1/(2\sqrt{e}))+1/2}-1\right)^2}$, which is roughly 4.911. □

It is worth noting that our result is robust in the sense that it holds independently of the arrival distribution. That is, even in the case where the buyer arrival is adversarial-the worst possible for the seller-, we prove that the seller's expected revenue in the observable case is at most 4.911 times the seller's expected revenue if she does not observe the buyer's arrival.

Unfortunately it is not so easy to obtain lower bounds for the value of observability. The difficulty relies in that it is notoriously difficult to solve, even numerically, the optimal price function in the unobservable case. Consider a buyer with valuation distributed uniformly in [0, 1] and two possible time arrivals, 0 or T. We can solve the unobservable case using a dynamic programming approach which is detailed in the full version of the paper. Thus we can maximize the value of observability over T and μ, fixing the seller's discount rate to 1. This leads to a lower bound on the value of observability of 1.042, which is attained by taking $\mu = 1.8$ and $T = 0.9$.

ACKNOWLEDGMENTS

This work was partially supported by CONYCYT-Chile through grants PIA/BASAL AFB180003 and FONDECYT 1190043.

REFERENCES

[1] Y. Aviv, A. Pazgal. 2008. Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing & Service Operations Management*, 10(3):339–359.
[2] O. Besbes, I. Lobel. 2015. Intertemporal price discrimination: Structure and computation of optimal policies. *Management Science* 61(1):92–110.
[3] L. Blumrosen, T. Holenstein. 2008. Posted prices vs. negotiations: an asymptotic analysis. EC 2008.
[4] S. Board, A. Skrzypacz. 2016. Revenue management with forward-looking buyers. *Journal of Political Economy* 124(4):1046–1087.
[5] L. Briceño-Arias, J. Correa, A. Perlroth. 2017. Optimal Continuous Pricing with Strategic Consumers. *Management Science* 63(8):2741–2755.
[6] R. Caldentey, Y. Liu, I. Lobel. 2017. Optimal Continuous Pricing with Strategic Consumers. *Operations Research* 65(1):104–129.
[7] S. Chawla, J. Hartline, R. Kleinberg. 2007. Algorithmic pricing via virtual valuations. EC 2007.
[8] S. Chawla, J. Hartline, D. Malec, B. Sivan. 2010. Multi-parameter mechanism design and sequential posted pricing. STOC 2010.
[9] J. Correa, P. Foncea, R. Hoeksma, T. Oosterwijk, T. Vredeveld. 2017. Posted price mechanisms for a random stream of customers. EC 2017.
[10] J. Correa, P. Foncea, D. Pizarro, V. Verdugo. 2019. From prophets to pricing, and back!. *Operations Research Letters* 47(1):25–29.
[11] J. Correa, R. Montoya, C. Thraves. 2016. Contingent preannounced pricing policies with strategic consumers. *Operations Research* 64(1):251–272.
[12] P. Dhangwatnotai, T. Roughgarden, Q. Yan. 2015. Revenue maximization with a single sample. *Games and Economic Behavior* 91:318–333.
[13] W. Elmaghraby, S. Lippman, C.S. Tang, R. Yin. 2009. Will more purchasing options benefit customers? *Production and Operations Management* 18(4):381–401.
[14] D.F. Garrett. 2016. Intertemporal price discrimination: Dynamic arrivals and changing values. *The American Economic Review* 106(11):3275–3299.
[15] A. Gershkov, B. Moldovanu, P. Strack. 2017. Revenue-Maximizing Mechanisms with Strategic Customers and Unknown, Markovian Demand. *Management Science* 64(5):2031–2046.
[16] M. Hajiaghayi, R. Kleinberg, T. Sandholm. 2007. Automated online mechanism design and prophet inequalities. AAAI 2007.
[17] Q. Liu, K. Mierendorff, X. Shi, W. Zhong. 2019. Auctions with limited commitment. *American Economic Review* 109(3):876–910.
[18] M. Landsberger, I. Meilijson. 1985. Intertemporal price discrimination and sales strategy under incomplete information. *RAND Journal of Economics* 16:424–430.
[19] R.B. Myerson. 1981. Optimal auction design. *Mathematics of Operations Research* 6(1):58–73.
[20] P.J. Olver. 2012.*Chapter 20:* Nonlinear Ordinary Differential Equations.
[21] N. Osadchiy, G. Vulcano. 2010. Selling with binding reservations in the presence of strategic consumers. *Management Science* 56(12): 2173–2190.
[22] M. Pai, R. Vohra. 2013. Optimal Dynamic Auctions and Simple Index Rules. *Mathematics of Operations Research* 38(3):682–697.
[23] A. Pavan, I. Segal, J. Toikka. 2014. Dynamic mechanism design: A Myersonian approach. *Econometrica* 82(2):601–653.
[24] A. Shneyerov. 2014. An optimal slow Dutch auction. *Economic Theory* 57:577-602.
[25] N.L. Stokey. 1979. Intertemporal price discrimination. *The Quarterly Journal of Economics* 93(3):355–371.
[26] R. Wang. 2001. Optimal pricing strategy for durable-goods monopoly. *Journal of Economic Dynamics and Control* 25:789–804.