

Online Posted Pricing with Unknown Time-Discounted Valuations

Giulia Romano, Gianluca Tartaglia, Alberto Marchesi, Nicola Gatti

Politecnico di Milano, Piazza Leonardo da Vinci 32, I-20133, Milan, Italy

giulia.romano@polimi.it, gianluca.tartaglia@mail.polimi.it, alberto.marchesi@polimi.it, nicola.gatti@polimi.it

Abstract

We study the problem of designing *posted-price mechanisms* in order to sell a single unit of a single item within a finite period of time. Motivated by real-world problems, such as, *e.g.*, long-term rental of rooms and apartments, we assume that customers arrive *online* according to a Poisson process, and their valuations are drawn from an *unknown* distribution and *discounted* over time. We evaluate our mechanisms in terms of competitive ratio, measuring the worst-case ratio between their revenue and that of an optimal mechanism that knows the distribution of valuations. First, we focus on the *identical valuation* setting, where all the customers value the item for the same amount. In this setting, we provide a mechanism \mathcal{M}_c that achieves the best possible competitive ratio, discussing its dependency on the parameters in the case of linear discount. Then, we switch to the *random valuation* setting. We show that, if we restrict the attention to distributions of valuations with a monotone hazard rate, then the competitive ratio of \mathcal{M}_c is lower bounded by a strictly positive constant that does not depend on the distribution. Moreover, we provide another mechanism, called \mathcal{M}_{PC} , which is defined by a piecewise constant pricing strategy and reaches performances comparable to those obtained with \mathcal{M}_c . This mechanism is useful when the seller cannot change the posted price too often. Finally, we empirically evaluate the performances of our mechanisms in a number of experimental settings.

Introduction

Posted-price mechanisms try to sell an item by proposing a *take-it-or-leave-it* price to each arriving agent, who then decides whether to buy the item or not (Chawla et al. 2010). If an agent opts for purchasing the item, then the mechanism terminates; otherwise, the agent leaves without any further possibility of buying the item, and the mechanism goes on by proposing prices to upcoming agents. Over the last years, growing attention has been devoted to the analysis of posted-price mechanisms, both in the classical economic literature (Seifert 2006) and in computer science (Babaioff et al. 2015, 2017; Adamczyk et al. 2017; Correa et al. 2017), within artificial intelligence and machine learning in particular (Kleinberg and Leighton 2003; Shah, Johari, and Blanchet 2019). This is mainly motivated by the overwhelming number of online economic transactions carried out by

posted-price mechanisms. This happens, for example, in online travel agencies (*e.g.*, *Expedia*), accommodation websites (*e.g.*, *Booking.com*), and e-commerce platforms (*e.g.*, *Amazon*, *eBay*). As studied by Einav et al. (2018), an increasing number of *eBay* users prefer buying goods via posted prices rather than participating in auctions.

Posted-price mechanisms provide many advantages over traditional auction-style mechanisms. From the designer’s perspective, posting prices requires a much lower effort than running an auction, since it avoids the burden of first *eliciting information* (the bids) from the agents, and then *collecting the payments*. At the same time, posted-price mechanisms retain most of the desirable properties of classical auctions, such as truthfulness. Indeed, even though the agents are not required to report their valuations for the item, they are always better off deciding whether to buy the item or not on the basis of their true valuations, without acting strategically (Babaioff et al. 2017). From the agents’ perspective, participating in a posted-price mechanism is preferable over competing in an auction, for several reasons. For instance, agents may prefer revealing minimal information about their true preferences if they plan to participate in similar markets in the future. Moreover, in some real-world settings, requiring the agents to figure out their true valuations for the item might need some additional efforts on their behalf, while answering a take-it-or-leave-it offer is usually much easier.

In this work, we study posted-price mechanisms for selling a single unit of a single item within a finite period of time, when the value of the item is discounted over time according to an arbitrary continuous and non-increasing discount function. Discounting is common in many real-world applications and widely studied for a number of economic situations, such, *e.g.*, bargaining (Rubinstein 1982; Gatti, Di Giunta, and Marino 2008) and auctions (Mao et al. 2018). We tackle settings in which agents arrive sequentially—a common assumptions in *online* mechanism design (Lavi and Nisan 2004; Parkes 2007)—and the number of agents is unknown *a priori*. In particular, following a mainstream approach in economics (see, *e.g.*, (Mason and Välimäki 2011; Rosenthal 2011)), we assume that agents’ arrivals are governed by a Poisson process. Remarkably, posted pricing with Poisson arrivals has been previously investigated by Wang (1993) and Rong, Qin, and An (2018) for undiscounted settings, though without providing any theoretical result.

We assume that each agent arriving at the mechanism has a different initial (*i.e.*, undiscounted) valuation for the item, which is independently drawn according to a common probability distribution. This leads to a fundamental trade-off between setting high prices so as to achieve high revenue and, on the other side, progressively lowering posted prices so as to increase the probability of selling the item. Our assumption is that the mechanism is only aware of the range of valuations, while it does not know anything about the shape of the distribution. This is reasonable since, differently from the actual distribution, the range of valuations can be estimated from previous data or market surveys.

Lavi and Nisan (2004) and Babaioff et al. (2017) provide the main state-of-the-art results on posted-price mechanisms for single-item single-unit scenarios. However, their models do not fit to our setting, since the agents' valuations are *not* discounted over time and the number of agents is known *a priori*. As a result, these models do not embed an explicit time representation and the proposed pricing strategies are only driven by the number of agents arrived.

Our model encompasses many real-world scenarios, such as, *e.g.*, long-term rental of rooms and apartments. Think of a website renting rooms to students for fixed periods of one year. The value of a room naturally decreases over time, reflecting the fact that a future tenant will benefit from the room for a period shorter than one year. Moreover, the potential customers arrive at the renting website according to a stochastic process, which can be reasonably modeled by a Poisson process whose rate parameter can be easily estimated by looking at traffic logs of the website.

Original Contributions We adopt the perspective of competitive analysis (Borodin and El-Yaniv 2005) and evaluate our mechanisms in terms of *competitive ratio*, measuring the worst-case ratio between their revenue and that of an optimal mechanism that knows the distribution of valuations. As it is customary in the literature (see, *e.g.*, (Babaioff et al. 2017; Kleinberg and Leighton 2003)), we first focus on the identical valuation setting in which all the agents share the same initial valuation for the item. Then, we extend our results to the random valuation setting where the agents' valuations are drawn i.i.d. from the same distribution satisfying the monotone hazard rate condition (when the distributions of valuations are unrestricted, Lavi and Nisan (2004) and Babaioff et al. (2017) show that then there is no algorithm with good performances). In the identical valuation setting, we design a posted-price mechanism \mathcal{M}_C and prove that it is optimal, *i.e.*, it provides the best possible competitive ratio. In order to derive the ratio, we first identify two crucial properties that characterize optimal mechanisms: their undiscounted price is non-increasing in time and they always guarantee the same fraction of the expected revenue of an optimal mechanism that knows the agents' valuation, independently of its actual value. For the specific case of linear discount, we discuss how the competitive ratio depends on the parameters. In the random valuation setting, we first show that mechanism \mathcal{M}_C still provides good performances by proving that its competitive ratio is lower

bounded by a constant, which does not depend on the distribution of agents' valuations. Then, motivated by real-world scenarios in which the seller is constrained not to change the posted prices too often, we propose a new mechanism \mathcal{M}_{PC} defined by a piecewise constant pricing strategy and prove that its performances in terms of competitive ratio are comparable with those obtained by \mathcal{M}_C . In conclusion, we empirically compare \mathcal{M}_C with a natural adaption of the mechanism proposed by Babaioff et al. (2017) to our setting, showing that the latter is inefficient even without time discounting. We also empirically evaluate the performances of \mathcal{M}_C and \mathcal{M}_{PC} as the frequency with which prices are allowed to change decreases, showing that, when this is not too low, then the performances of \mathcal{M}_{PC} and \mathcal{M}_C are comparable.

Other Related Works As showed by Hajiaghayi, Kleinberg, and Sandholm (2007) for single-item settings, posted pricing is strictly related to the *secretary problem* and to *prophet inequalities*; see also the work by Babaioff et al. (2009) for single-item settings and that of Lucier (2017) for multi-item scenarios. Differently from our model, these works assume that the mechanism knows the probability distribution of agents' valuations and that the agents reveal their actual valuation for the item upon arrival. When multiple units of the same item are available, learning approaches based on *bandit* techniques are customarily adopted. In particular, Kleinberg and Leighton (2003) study an unlimited-supply setting where the number of buyers is fixed, and derive upper bounds on the regret. Several recent works extend the results in (Kleinberg and Leighton 2003). Shah, Johari, and Blanchet (2019) study a contextual setting, providing a semi-parametric model that learns from the observation of a binary outcome which stands for acceptance or rejection of the offered price. Mohri and Munoz (2014) study revenue-maximizing learning algorithms for posted pricing with strategic buyers. They consider a repeated game in which, at each round, the seller offers the item at a certain price and a strategic buyer accepts or rejects it. In that work, the goal is to learn the buyers' valuation for the item by minimizing the *strategic-regret* of the algorithm.

Preliminaries

We study a model in which a seller is interested in selling a single unit of an item within a finite time period of length T . The seller implements a *posted-price mechanism* by setting a take-it-or-leave-it price at each time $t \in [0, T]$. We denote by $p : [0, T] \rightarrow \mathbb{R}_+$ the *pricing strategy* adopted by the seller, with $p(t)$ being the price offered at time $t \in [0, T]$. The agents (*i.e.*, the buyers) arrive sequentially over time, according to a Poisson process with rate parameter $\lambda > 0$.

We label agents according to their order of arrival (*i.e.*, agent i is the i -th agent arriving in $[0, T]$). Each agent i has a private valuation V_i for the item, drawn from a distribution F with finite support $[v_{\min}, v_{\max}]$, where $v_{\max} > v_{\min} > 0$ denote the maximum and minimum valuation, respectively. In the following, for the ease of presentation, we normalize agents' valuations in the range $[1, h]$, where $h := \frac{v_{\max}}{v_{\min}}$. Accordingly, we scale the support of F to $[1, h]$. Then, we

denote by f the probability density function of F .

The value of the item for sale decreases over time. In particular, V_i is agent i 's initial valuation at time $t = 0$. We model decreasing values by introducing a continuous non-increasing *discount function* $\xi : [0, T] \rightarrow [0, 1]$ such that $\xi(0) = 1$ and $\xi(T) = 0$. By letting W_i be the random variable representing the arrival time of agent i , we define the agent i 's *discounted valuation* as $D_i := V_i \xi(W_i)$, which represents how much agent i is willing to pay upon her arrival. As a result, whenever agent i arrives, she buys the item if and only if $D_i \geq p(W_i)$, *i.e.*, her discounted valuation is at least as large as the price offered by the mechanism.

We introduce the following additional notation. We denote by $I_{s,\tau} := [s, s+\tau] \subseteq [0, T]$ the time interval of length $\tau \in [0, T]$ starting from time $s \in [0, T - \tau]$. The number of agents arriving in $I_{s,\tau}$ is a random variable denoted by $N_{s,s+\tau}$. Given $\tau \in [0, T]$, the random variables $N_{s,s+\tau}$ are equally distributed for all $s \in [0, T - \tau]$, as the arrivals are generated by a Poisson process. For the sake of presentation, we omit s in $N_{s,s+\tau}$, denoting by N_τ the random variable of the number of agents arriving in any time interval of length τ , which follows a Poisson distribution with parameter $\lambda\tau$.¹ Thus, N_T is the random variable of the total number of agents arriving in the overall time period. In the following, we sometimes focus on the *linear* discount function, denoted as $\xi_{\text{lin}} : [0, T] \rightarrow [0, 1]$ with $\xi_{\text{lin}}(t) := 1 - \frac{t}{T}$. In this case, each agent i 's discounted valuation is $D_i := V_i \left(1 - \frac{W_i}{T}\right)$.

Performances of Posted-Price Mechanisms

Given a deterministic posted-price mechanism \mathcal{M} defined by a price function $p_{\mathcal{M}} : [0, T] \rightarrow \mathbb{R}_+$, we denote by $\mathbb{E}_F[\mathcal{R}(\mathcal{M})]$ the expected revenue that the mechanism provides to the seller. The expectation is calculated with respect to both the Poisson arrivals and the distribution F of agents' initial valuations. We made explicit the dependence on F , as we will frequently refer to it along the paper.

We adopt the perspective of competitive analysis and measure the performances of a mechanism \mathcal{M} by comparing the seller's expected revenue with that of a benchmark mechanism \mathcal{M}^* , which is optimal having knowledge of the distribution F . Notice that the benchmark has no information on the actual realizations of agents' initial valuations, but only on their distribution, whereas the mechanisms we propose operate having knowledge of their range only.

Our goal is to bound the performances of our mechanisms w.r.t. those of the benchmark \mathcal{M}^* by looking at the worst case over the set \mathcal{F} of possible distributions F , *i.e.*, all those with support $[1, h]$. This is captured by the following:

Definition 1. *The competitive ratio of a deterministic posted-price mechanism \mathcal{M} is defined as:*

$$\rho(\mathcal{M}) := \min_{F \in \mathcal{F}} \rho_F(\mathcal{M}), \quad \text{where } \rho_F(\mathcal{M}) := \frac{\mathbb{E}_F[\mathcal{R}(\mathcal{M})]}{\mathbb{E}_F[\mathcal{R}(\mathcal{M}^*)]}.$$

Moreover, we say that a mechanism is optimal when its competitive ratio is the highest possible among all the deterministic posted-price mechanisms.

¹By definition of Poisson distribution, $\mathbb{P}\{N_\tau = j\} = \frac{(\lambda\tau)^j e^{-\lambda\tau}}{j!}$.

Notice that $\rho(\mathcal{M}) \in [0, 1]$ and, for every possible distribution $F \in \mathcal{F}$, we are guaranteed that the seller's expected revenue $\mathbb{E}_F[\mathcal{R}(\mathcal{M})]$ provided by mechanisms \mathcal{M} is at least a fraction $\rho(\mathcal{M})$ of that achieved by \mathcal{M}^* , *i.e.*, it holds $\mathbb{E}_F[\mathcal{R}(\mathcal{M})] \geq \rho(\mathcal{M}) \mathbb{E}_F[\mathcal{R}(\mathcal{M}^*)]$.

As previously showed by Babaioff et al. (2017) in similar settings, we can safely restrict our analysis to mechanisms maintaining the bottom price for a non-negligible period of time. Indeed, in the case in which F places all the probability mass on 1, then any mechanism providing a non-null seller's expected revenue must set the minimum price during some time interval, otherwise no agent would buy the item.

Proposition 1. *Every deterministic posted-price mechanism \mathcal{M} such that $\rho(\mathcal{M}) > 0$ must set the minimum price $p_{\mathcal{M}}(t) = \xi(t)$ for every t in a time interval of length $\tau > 0$.*

Identical Valuation Setting

We start studying the *identical valuation* (IV) setting, where all the agents share the same initial valuation $v \in [1, h]$ for the item, *i.e.*, it holds $V_i = v$ and $D_i = v \xi(W_i)$ for every agent i . The IV setting is a special case of the general random valuation model where one restricts the attention to distributions F placing all the probability mass on a single valuation in $[1, h]$. In the following, we adjust notation for expected revenues and competitive ratios accordingly, writing $\mathbb{E}_v[\mathcal{R}(\mathcal{M})]$ and $\rho_v(\mathcal{M})$ instead of $\mathbb{E}_F[\mathcal{R}(\mathcal{M})]$ and $\rho_F(\mathcal{M})$.

Our main result (Theorem 1) is to provide a deterministic posted-price mechanism, called \mathcal{M}_c , which is optimal for the IV setting for every discount function ξ . We also study the specific case of a linear discount function ξ_{lin} , where we design an optimal mechanism $\mathcal{M}_{c,\text{lin}}$ (Theorem 2) that enjoys an easily interpretable analytical description. All the proofs are in (Romano et al. 2020).

First, we describe the shape of the benchmark mechanism \mathcal{M}^* for the IV setting. Indeed, since \mathcal{M}^* knows the actual initial valuation v , its price function $p_{\mathcal{M}^*} : [0, T] \rightarrow \mathbb{R}_+$ is such that $p_{\mathcal{M}^*}(t) = v \xi(t)$ for $t \in [0, T]$. Therefore, we can compute the expected revenue of \mathcal{M}^* as follows:

$$\mathbb{E}_v[\mathcal{R}(\mathcal{M}^*)] := \int_0^T p_{\mathcal{M}^*}(t) \lambda e^{-\lambda t} dt = \int_0^T v \xi(t) \lambda e^{-\lambda t} dt = v k^*, \quad (1)$$

where $k^* := \int_0^T \xi(t) \lambda e^{-\lambda t} dt$ does not depend on v , but only on the problem parameters T , λ , and the discount function ξ . Let us remark that the expected revenue of the benchmark \mathcal{M}^* defined in Equation (1) is expressed as a *linear* function of v .

Optimal Mechanism for a General Discount

We start proving two lemmas that highlight two crucial properties which characterize optimal posted-price mechanisms for the IV setting. Lemma 1 implies that the pricing strategy of an optimal mechanism must be such that the undiscounted price defined as $\frac{p(t)}{\xi(t)}$ is non-increasing in t , whereas Lemma 2 shows that any mechanism which always provides

a constant fraction of the expected revenue of the benchmark, independently of the agents' initial valuation v , is an optimal mechanism.

Lemma 1. *In the IV setting, given any deterministic posted-price mechanism \mathcal{M} , there always exists a deterministic posted-price mechanism \mathcal{M}' with undiscounted price $\frac{p_{\mathcal{M}'(t)}}{\xi(t)}$ non-increasing in t such that $\mathbb{E}_v[\mathcal{R}(\mathcal{M})] \leq \mathbb{E}_v[\mathcal{R}(\mathcal{M}')] for every possible agents' initial valuation $v \in [1, h]$.$*

Notice that, since ξ is continuous and non-increasing by definition, Lemma 1 also shows that there is always an optimal mechanism whose pricing strategy is non-increasing. Moreover, by recalling Proposition 1, we can conclude that any optimal mechanism must set the minimum price at the end of the overall time period, *i.e.*, during a time interval $[t_0, T] \subseteq [0, T]$ defined for some $t_0 \in [0, T]$. This result is exploited to prove the following lemma.

Lemma 2. *In the IV setting, let \mathcal{M} be a deterministic posted-price mechanism whose pricing strategy $p_{\mathcal{M}}$ satisfies $p_{\mathcal{M}}(t) = \xi(t)$ for $t \in [t_0, T]$ with $t_0 \in [0, T]$. If the ratio $\rho_v(\mathcal{M}) = \frac{\mathbb{E}_v[\mathcal{R}(\mathcal{M})]}{\mathbb{E}_v[\mathcal{R}(\mathcal{M}^*)]}$ for \mathcal{M} does not depend on the agents' initial valuation v , then \mathcal{M} is an optimal mechanism.*

By Lemma 2, in order to find an optimal mechanism for the IV setting, we can restrict the attention to mechanisms \mathcal{M} whose ratios $\rho_v(\mathcal{M})$ do not depend on the initial valuation v . Therefore, since the expected revenue of the benchmark \mathcal{M}^* is a linear function of v (see Equation (1)), we can search for an optimal mechanism among those having an expected revenue which linearly depends on v . This crucial observation allows us to design the optimal mechanism \mathcal{M}_c in Theorem 1 by leveraging the condition $\mathbb{E}_v[\mathcal{R}(\mathcal{M}_c)] = kv$ for every $v \in [1, h]$, with k being a suitably defined constant independent of v . The key insight that allows us to derive an expression for \mathcal{M}_c is that we can always find the desired pricing strategy $p_{\mathcal{M}_c}$ among the continuous price functions such that $\frac{p_{\mathcal{M}_c}(t)}{\xi(t)}$ is non-increasing in $t \in [0, T]$. Intuitively, using Lemma 1, we can always express the expected revenue $\mathbb{E}_v[\mathcal{R}(\mathcal{M}_c)]$ as a function of the time $t^* := \sup\{t \in [0, t_0] \mid p_{\mathcal{M}_c}(t) > v\xi(t)\}$, which is the first time in which $p_{\mathcal{M}_c}$ intersects $v\xi(t)$. The reason is that it holds $p_{\mathcal{M}_c}(t) \leq v\xi(t)$ if and only if $t \geq t^*$, and, thus, only agents arriving after t^* are willing to buy the item. By using the relation among $\mathbb{E}_v[\mathcal{R}(\mathcal{M}_c)]$ and t^* , we can find the desired pricing strategy $p_{\mathcal{M}_c}$ as a solution to a suitably defined differential equation. This leads to the following theorem.

Theorem 1. *In the IV setting, there exists an optimal deterministic posted-price mechanism \mathcal{M}_c whose pricing strategy $p_{\mathcal{M}_c}$ is defined as follows:*

$$p_{\mathcal{M}_c}(t) := \begin{cases} a e^{\int b(t)dt} & \text{if } t \in [0, t_0] \\ \xi(t) & \text{if } t \in [t_0, T] \end{cases},$$

where b is a function such that $b(t) := \lambda - \frac{\lambda}{k\zeta(t)} - \frac{\zeta'(t)}{\zeta(t)}$ with $\zeta(t) := \frac{1}{\xi(t)}$, whereas a, k, t_0 are suitably defined constants that do not depend on the agents' initial valuation v .

As a byproduct of the proof of Theorem 1, we also get an expression for the competitive ratio of the mechanism \mathcal{M}_c , as stated by the following corollary.

	$T \rightarrow \infty$	$\lambda \rightarrow \infty$	$h \rightarrow \infty$
t_0	$\Theta(\sqrt{T})$	$\Theta(\sqrt{T/\lambda})$	$\Theta(T)$
$\rho(\mathcal{M}_{c,\text{lin}})$	$\Theta\left(1 - \frac{1}{\sqrt{T}}\right)$	$\Theta\left(1 - \frac{1}{\sqrt{\lambda}}\right)$	$\Theta\left(\frac{1}{\log^2(h)}\right)$
$\lim \rho(\mathcal{M}_{c,\text{lin}})$	1	1	0

Table 1: Values of t_0 and $\rho(\mathcal{M}_{c,\text{lin}})$ as T, λ, h go to infinity.

Corollary 1. *In the IV setting, mechanism \mathcal{M}_c achieves:*

$$\rho(\mathcal{M}_c) = \frac{\int_0^{T-t_0} \xi(t) \lambda e^{-\lambda t} dt}{\int_0^T \xi(t) \lambda e^{-\lambda t} dt}.$$

Optimal Mechanism for a Linear Discount

The pricing strategy $p_{\mathcal{M}_c}$ of the optimal mechanism defined in Theorem 1 still depends on some parameters, namely a, k , and t_0 , which do not admit an easy analytical formula for a general discount function ξ . Nevertheless, they can be expressed analytically if we restrict the attention to functions ξ having a particular shape. In the following Theorem 2 and Corollary 2, we analyze the case of a linear discount function ξ_{lin} , defining an optimal mechanism $\mathcal{M}_{c,\text{lin}}$ for such setting.

Theorem 2. *In the IV setting with linear discount function ξ_{lin} , there exists an optimal deterministic posted-price mechanism $\mathcal{M}_{c,\text{lin}}$ whose pricing strategy $p_{\mathcal{M}_{c,\text{lin}}}$ is defined as:*

$$p_{\mathcal{M}_{c,\text{lin}}}(t) := \begin{cases} h \left(1 - \frac{t}{T}\right) e^{\lambda(1-\frac{1}{k})t + \frac{\lambda}{2kT}t^2} & \text{if } t \in [0, t_0] \\ 1 - \frac{t}{T} & \text{if } t \in [t_0, T] \end{cases},$$

where $k := \lambda t_0 \frac{2T-t_0}{2T(\lambda t_0 + \ln h)}$ and the time $t_0 \in [0, T]$ is defined as the unique positive real root of the following equation: $1 - \frac{1}{\lambda T} (1 + \lambda t_0 - e^{-\lambda(T-t_0)}) = k$.

The prices posted by $\mathcal{M}_{c,\text{lin}}$ decrease as a linearly discounted exponential function until $t = t_0$, starting, at time $t = 0$, by setting the price equal to the maximum agents' initial valuations h . Then, during the time interval $[t_0, T]$, the price function linearly decreases and equals zero in $t = T$.

Corollary 2. *In the IV setting with linear discount function ξ_{lin} , $\mathcal{M}_{c,\text{lin}}$ achieves a competitive ratio:*

$$\rho(\mathcal{M}_{c,\text{lin}}) = \frac{1 - \frac{1}{\lambda T} (1 + \lambda t_0 - e^{-\lambda(T-t_0)})}{1 - \frac{1}{\lambda T} (1 - e^{-\lambda T})}.$$

The asymptotic values of t_0 and $\rho(\mathcal{M}_{c,\text{lin}})$ as T, λ, h go to infinity are in Table 1 (see (Romano et al. 2020) for more details). In particular, $\rho(\mathcal{M}_{c,\text{lin}})$ goes asymptotically to 1 as λ or T increases. This corresponds to having an infinite number of agents and, thus, selling the item with certainty. Instead, $\rho(\mathcal{M}_{c,\text{lin}})$ decreases as h increases, going asymptotically to 0 as $\frac{1}{\log^2(h)}$. The range $[1, h]$ represents the degree of uncertainty that the mechanism has on the agents' valuation. Therefore, $\rho(\mathcal{M}_{c,\text{lin}})$ decreases as the uncertainty increases and it cannot be lower bounded by any strictly positive constant if no finite upper bound on h is known (*i.e.*, when $h \rightarrow +\infty$). However, the dependency of $\rho(\mathcal{M}_{c,\text{lin}})$ on the degree of uncertainty is logarithmic. Instead, notice that a trivial mechanism setting the price equal to $\xi_{\text{lin}}(t)$ for $t \in [0, T]$ would have a competitive ration of $\frac{1}{h}$, which depends linearly on the degree of uncertainty.

Random Valuation Setting

We now switch to the *random valuation* (RV) setting, where agents' initial valuations V_i are i.i.d. random variables defined by a cumulative distribution function F with support $[1, h]$. We focus on distributions F satisfying the *monotone hazard rate* (MHR) condition. Formally, a distribution F is MHR if the hazard rate $H(x) := \frac{f(x)}{1-F(x)}$ is non-decreasing in x . This assumption is common when studying posted-price mechanisms that operate without knowing the shape of the distribution of valuations (see (Babaioff et al. 2015, 2017)) and many distributions used in practice satisfy it (such as, e.g., uniform, normal, and exponential distributions). Moreover, the MHR condition is necessary for proving our main results (Theorems 3 and 4). Indeed, when the family of possible distributions is unrestricted, one cannot design posted-price mechanisms guaranteeing a constant fraction of the revenue of \mathcal{M}^* independently of the distribution F , as shown by Babaioff et al. (2017) for the easier setting in which agents do not arrive stochastically. All the proofs are provided in (Romano et al. 2020).

Auxiliary Definitions and Results We introduce the random variable $X_{\lambda\tau}$ as the maximum initial valuation of agents arriving in an interval of length $\tau \in (0, T]$. Formally:

$$X_{\lambda\tau} := \max_{i \in \{1, \dots, N_\tau\}} V_i.^2$$

$X_{\lambda\tau}$ is the first order statistic of N_τ samples drawn from F and, since agents' arrivals are governed by a Poisson process, its cumulative distribution function $F_{X_{\lambda\tau}}$ is defined as:

$$\begin{aligned} F_{X_{\lambda\tau}}(x) &:= \sum_{j=1}^{\infty} \mathbb{P}\{N_\tau = j\} F_{X_{\lambda\tau}|N_\tau=j}(x) = \\ &\sum_{j=1}^{\infty} \frac{(\lambda\tau)^j e^{-\lambda\tau}}{j!} [F(x)]^j = e^{-\lambda\tau(1-F(x))}. \end{aligned}$$

We also define $Y_{s,\lambda\tau}$ as the random variable representing the maximum discounted valuation of agents arriving in an interval $I_{s,\tau}$ of length $\tau \in (0, T]$ starting at $s \in [0, T - \tau]$:

$$Y_{s,\lambda\tau} := \max_{i \in \{1, \dots, N_\tau\}} D_i.$$

The cumulative distribution function $F_{Y_{s,\lambda\tau}}$ of $Y_{s,\lambda\tau}$ is:

$$\begin{aligned} F_{Y_{s,\lambda\tau}}(x) &:= \sum_{j=1}^{\infty} \mathbb{P}\{N_\tau = j\} F_{Y_{s,\lambda\tau}|N_\tau=j}(x) = \\ &\sum_{j=1}^{\infty} \frac{(\lambda\tau)^j e^{-\lambda\tau}}{j!} F_{Y_{s,\lambda\tau}|N_\tau=j}(x), \end{aligned}$$

where $F_{Y_{s,\lambda\tau}|N_\tau=j}$ is the cumulative distribution function of $Y_{s,\lambda\tau}$ conditioned on the event $N_\tau = j$. Let us remark that,

²In the definition of $X_{\lambda\tau}$ and $Y_{s,\lambda\tau}$, overloading the notation, we assume that the agents arriving in the considered time interval of length τ are labeled from 1 to N_τ according to their order. Their actual labels referred to the overall period $[0, T]$ may be different.

by definition, $F_{Y_{s,\lambda\tau}}$ depends on distribution F . In the following, we also let $Y_{\lambda T} := Y_{0,\lambda T}$ be the random variable representing the maximum discounted valuation of agents arriving in the overall time period $[0, T]$. In the Supplemental Material, for the specific case of a linear discount function, we show how to exploit some useful properties of Poisson processes so as to find an analytical expression for $F_{Y_{\lambda T}}$. In particular, by letting $Z := VU$, where V and U are independent random variables distributed according to F and $\mathcal{U}(0, 1)$, respectively, we obtain:

$$F_{Y_{\lambda T}}(x) := \sum_{j=1}^{\infty} \frac{(\lambda T)^j e^{-\lambda T}}{j!} [F_Z(x)]^j,$$

where

$$F_Z(x) := \begin{cases} x \int_1^h \frac{1}{z} f(z) dz & \text{if } x \in [0, 1) \\ F(x) + x \int_x^h \frac{1}{z} f(z) dz & \text{if } x \in [1, h] \end{cases}.$$

Bounding $\rho(\mathcal{M}_C)$ in the RV Setting

We show that mechanism \mathcal{M}_C (see definition in Theorem 1), which is optimal in the IV setting, provides good performances also in the RV setting. Our main result (Theorem 3) is a lower bound on the competitive ratio of the mechanism, which is obtained by showing that \mathcal{M}_C always provides at least a constant fraction of the seller's expected revenue achieved by the benchmark \mathcal{M}^* , independently of the distribution of agents' initial valuations F .³ This is surprising since, differently from \mathcal{M}^* , our mechanism works without having knowledge about F (except for its range).

We first need some definitions and lemmas.

Definition 2. Let $I_{s,\tau}$ be any interval of length $\tau \in (0, T]$ starting at $s \in [0, T - \tau]$. Then, the ratio between the prices posted by \mathcal{M}_C at the endpoints of $I_{s,\tau}$ is defined as:

$$\kappa_\tau(s) := \frac{p_{\mathcal{M}_C}(s)}{p_{\mathcal{M}_C}(s + \tau)}.$$

Intuitively, $\kappa_\tau(s)$ bounds the slope of the price function of \mathcal{M}_C in the time interval $I_{s,\tau}$, which depends on both the starting time s and the length τ of the interval. Moreover, notice that $\kappa_\tau(s) \geq 1$ since $p_{\mathcal{M}_C}$ is non-increasing by Lemma 1. Next, we introduce an upper bound on the price ratios of all the time intervals of length τ , which is useful in deriving our main result.

Definition 3. The maximum price ratio of \mathcal{M}_C over intervals of length $\tau \in (0, T]$ is denoted by:

$$\kappa_\tau := \max_{s \in [0, T - \tau]} \kappa_\tau(s).$$

The following lemma establishes a relation between the price function $p_{\mathcal{M}_C}$ of \mathcal{M}_C and the expected value of the random variable $X_{\lambda T}$ representing the maximum initial valuation of agents arriving in the overall time period. This is crucial to prove Theorem 3.

³To prove the lower bound, we follow an approach similar to that used by Babaioff et al. (2017) to bound the competitive ratio of their *Equal-Sample-of-Every-Scale* mechanism. However, our setting introduces additional challenges, since the agents' arrivals are stochastic and the valuations are discounted. Thus, our proofs require different techniques w.r.t. those of Babaioff et al. (2017).

Lemma 3. *In the RV setting with agents' initial valuations drawn from a distribution F , given $\tau \in (0, T]$ and $0 < \epsilon < 1$, there exists at least an interval $I_{s, \tau}$ of length $\tau \in (0, T]$ starting at $s \in [0, T - \tau]$ such that the prices $p_{\mathcal{M}_c}(t)$ posted by mechanism \mathcal{M}_c during the time instants $t \in I_{s, \tau}$ lie in the range $\left[\frac{\mathbb{E}[X_{\lambda T}] \xi(s + \tau)(1 - \epsilon)}{\kappa_\tau}, \mathbb{E}[X_{\lambda T}] \xi(s + \tau)(1 - \epsilon) \right]$.*

The following two lemmas are the final pieces that we need to prove Theorem 3. Lemma 4 establishes that, if the distribution F is MHR, then the same holds for the distribution $F_{X_{\lambda \tau}}$ of $X_{\lambda \tau}$. Lemma 5, given two intervals of length τ and τ' with $\tau \leq \tau'$, provides a lower bound on the expected value of $X_{\lambda \tau}$ which depends on the expected value of $X_{\lambda \tau'}$ and the logarithms of the expected number of agents' arrivals in the two intervals, respectively $\lambda \tau$ and $\lambda \tau'$.

Lemma 4. $F_{X_{\lambda \tau}}$ has non-decreasing monotone hazard rate.

Lemma 5. For every $\tau, \tau' \in (0, T]$ with $\tau \leq \tau'$, it holds:

$$\frac{\mathbb{E}[X_{\lambda \tau}]}{\mathbb{E}[X_{\lambda \tau'}]} \geq \frac{\ln(\lambda \tau)}{\ln(\lambda \tau')}.$$

Theorem 3. Consider the RV setting with $\lambda \tau = (\lambda T)^{1 - \epsilon} \geq 1 - \ln(e - 1)$ for some $\tau \in (0, T]$ and $0 < \epsilon < 1$. Then, restricted to the set \mathcal{F} of distributions F satisfying the MHR condition, mechanism \mathcal{M}_c has a competitive ratio that can be lower bounded as follows:

$$\rho(\mathcal{M}_c) \geq \frac{\xi(t_0 + T^{1 - \epsilon} \lambda^{-\epsilon})(1 - \epsilon)}{\kappa_\tau e}.$$

The idea of the proof is to use $\rho_F(\mathcal{M}_c) \geq \frac{\mathbb{E}_F[\mathcal{R}(\mathcal{M}_c)]}{\mathbb{E}[Y_{\lambda T}]}$, following from the fact that $\mathbb{E}_F[\mathcal{R}(\mathcal{M}^*)]$ cannot be larger than $\mathbb{E}[Y_{\lambda T}]$, which is the expected revenue achieved by an optimal mechanism that knows the realization of agents' initial valuations and arrivals. Then, $\mathbb{E}_F[\mathcal{R}(\mathcal{M}_c)]$ is lower bounded by the revenue that \mathcal{M}_c achieves in a suitably defined interval $I_{s, \tau}$, whose existence is guaranteed by Lemma 3. Moreover, Lemmas 3, 4, and 5, together with the properties of MHR distributions, allow us to write $\mathbb{E}_F[\mathcal{R}(\mathcal{M}_c)] \geq \frac{\mathbb{E}[X_{\lambda T}] \xi(t_0 + \tau)(1 - \epsilon)}{\kappa_\tau e}$, giving the result as $\mathbb{E}[Y_{\lambda T}] \leq \mathbb{E}[X_{\lambda T}]$.

A Mechanism with a Piecewise Constant Price

We introduce a new mechanism \mathcal{M}_{pc} whose pricing strategy $p_{\mathcal{M}_{pc}}$ is a piecewise constant function. This turns out to be useful in all the situations in which the seller is constrained not to change the posted price too often, *e.g.*, when the mechanism is required to set prices for time intervals having a given minimum length. Our main result (Theorem 4) is a lower bound on the competitive ratio of \mathcal{M}_{pc} in the RV setting, which is comparable to that obtained for \mathcal{M}_c in Theorem 3. Thus, we show that, even in presence of constraints on the allowed pricing strategies, we are still able to design mechanisms with good performances in terms of competitive ratio. Clearly, \mathcal{M}_{pc} depends on the minimum length requirement, which influences the resulting lower bound. In particular, \mathcal{M}_{pc} is tuned by a parameter δ related to the number of time intervals in which the price must be constant.

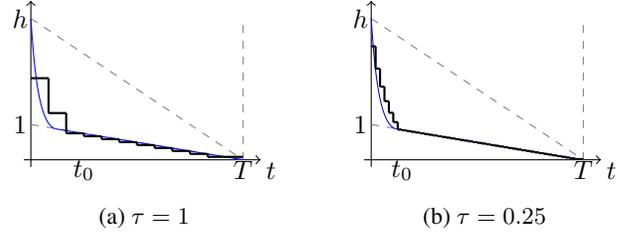


Figure 1: Prices of mechanisms $\mathcal{M}_{c,lin}$ (blue) and $\mathcal{M}_{pc,lin}$ (black) when $h = 2.8$, $\lambda = 10$, and $T = 12$.

Mechanism \mathcal{M}_{pc} works by evenly partitioning the time interval $[0, t_0]$ into $\lceil \log_\delta h \rceil$ sub-intervals of length τ , where $\delta \in (1, h)$ and $t_0 \in [0, T]$ are suitably defined parameters. Then, the remaining time $[t_0, T]$ is organized in other sub-intervals of length τ . As a result, $[0, T]$ is partitioned into $\lceil \frac{T}{\tau} \rceil$ sub-intervals, which, overloading notation, we denote by $I_i := [(i - 1)\tau, \min\{i\tau, T\}]$ for $i = 1, \dots, \lceil \frac{T}{\tau} \rceil$. Notice that $\tau = \frac{t_0}{\lceil \log_\delta h \rceil}$, and, thus, parameters t_0 and δ can be tuned to match the required minimum length τ . The pricing strategy $p_{\mathcal{M}_{pc}}$ of \mathcal{M}_{pc} is defined in such a way that the price is constant in each interval I_i . By letting $p_{\mathcal{M}_{pc}}(I_i)$ be the price posted during I_i , we define the function $p_{\mathcal{M}_{pc}}$ as follows:⁴

$$p_{\mathcal{M}_{pc}}(I_i) := \begin{cases} \frac{h}{\delta^i} \xi(i\tau) & \text{if } i = 1, \dots, \lceil \log_\delta h \rceil \\ \xi(i\tau) & \text{if } i = \lceil \log_\delta h \rceil, \dots, \lceil \frac{T}{\tau} \rceil - 1 \\ \xi((i - 1)\tau) & \text{if } i = \lceil \frac{T}{\tau} \rceil \end{cases}$$

We compare in Figure 1 the prices of $\mathcal{M}_{c,lin}$ and $\mathcal{M}_{pc,lin}$ (*i.e.*, \mathcal{M}_{pc} with a linear discount) in a specific setting for two values of τ . Notice that \mathcal{M}_{pc} can be thought of as an extension of the *Equal-Sample-of-Every-Scale* (ESoS) mechanism by Babaioff et al. (2017) to the more general setting in which agents arrive stochastically according to a Poisson process and agents' valuations are discounted over time.

Before proving our main result, we need the following lemma, which is the analog of Lemma 3 working for mechanism \mathcal{M}_{pc} instead of \mathcal{M}_c .

Lemma 6. *In the RV setting with agents' initial valuations drawn from a distribution F , given $0 < \epsilon < 1$, there exists $i = 1, \dots, \lceil \log_\delta h \rceil$ such that the price $p_{\mathcal{M}_{pc}}(I_i)$ posted by \mathcal{M}_{pc} during the interval I_i lies in the range $[\frac{\nu}{\delta} \xi(i\tau), \nu \xi(i\tau)]$, where $\nu := \max\{1, \mathbb{E}[X_{\lambda T}](1 - \epsilon)\}$.*

Now, we provide our main result. The idea behind its proof is similar to the one used for Theorem 3.

Theorem 4. Consider the RV setting with $\lambda \tau = (\lambda T)^{1 - \epsilon} \geq 1 - \ln(e - 1)$ for some $\tau \in (0, T]$ and $0 < \epsilon < 1$. Then, restricted to the set \mathcal{F} of distributions F satisfying the MHR condition, mechanism \mathcal{M}_{pc} has a competitive ratio that can be lower bounded as follows:

$$\rho(\mathcal{M}_{pc}) \geq \frac{\xi((\lceil \log_\delta h \rceil + 1)T^{1 - \epsilon} \lambda^{-\epsilon})(1 - \epsilon)}{\delta e}.$$

⁴Whenever T is not divisible by τ , then the last time interval is shorter than τ . Thus, in order to satisfy the minimum length constraint, we set its price equal to the one in the preceding interval.

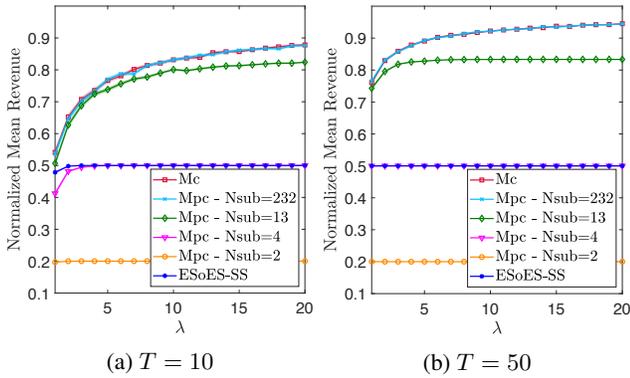


Figure 2: Average normalized revenue of \mathcal{M}_C , \mathcal{M}_{PC} , ESoES-SS in a RV setting w. uniform distribution ($h = 10$).

Empirical Evaluation

We evaluate mechanisms \mathcal{M}_C , \mathcal{M}_{PC} , and a natural adaptation of the ESoES mechanism by Babaiouf et al. (2017) to *stochastic settings* with no time discounting (called ESoES-SS). The pricing strategy of ESoES-SS is defined as follows. First, we compute the prices of ESoES by setting the number of agents equal to the expected number λT of agents arriving in $[0, T]$ according to a Poisson process of parameter λ . Then, ESoES-SS proposes the price that ESoES would propose to the i -th agent arrived if $i \leq \lambda T$ and 1 otherwise.

We use the following parameters values for the experiments: $\lambda \in \{1, \dots, 20\}$, $T \in \{10, 20, 50, 100\}$, and $h \in \{2, \dots, 20\}$. The following results do not consider time discounting so as to have a fair comparison between our mechanisms and ESoES-SS. Further results with a linear discount function are provided in (Romano et al. 2020).

Result #1 We study a RV setting with a uniform probability distribution over $[1, h]$. For every combination of values of λ, T, h , we run 1000 Monte Carlo simulations, evaluating the revenue provided by mechanisms ESoES-SS, \mathcal{M}_C , and \mathcal{M}_{PC} . In particular, we analyze some variants of mechanisms \mathcal{M}_{PC} differing for the number of subintervals (*i.e.*, N_{sub}) in which $[0, t_0]$ is partitioned. Furthermore, we normalize the revenue provided by the mechanisms in each simulation with respect to h . We report the results in Figure 2 for $T = 10$ and $T = 50$, when $h = 10$. The results obtained for different values of h are similar. \mathcal{M}_C and \mathcal{M}_{PC} with $N_{sub} = 232$ have overlapping performances that beat those of the other mechanisms. \mathcal{M}_{PC} with $N_{sub} = 13$ has a performance close to that of the previous two mechanisms, showing that mechanism \mathcal{M}_{PC} provides good performances even with few subintervals. \mathcal{M}_{PC} with $N_{sub} = 4$ and ESoES-SS have almost overlapping performances, showing that very few subintervals are sufficient to \mathcal{M}_{PC} to match the performances of ESoES-SS. The worst mechanism is \mathcal{M}_{PC} with $N_{sub} = 2$. The loss of ESoES-SS w.r.t. \mathcal{M}_C averaged over the values of λ is about $0.3h$ when $T = 10$, and $0.4h$ when $T = 50$. Surprisingly, the performances of ESoES-SS seem to do not strictly depend on λ and T .

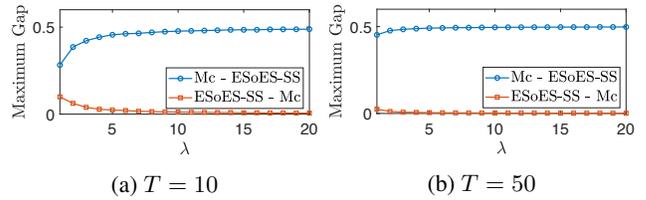


Figure 3: Maximum difference between the normalized revenues of \mathcal{M}_C and ESoES-SS in an IV setting ($h = 10$).

Result #2 We study an IV setting. For every combination of values of λ, T, h , and for every $v \in \{1.0, 1.5, 2.0, \dots, h\}$, we run 1000 Monte Carlo simulations, evaluating the normalized revenue provided by mechanisms ESoES-SS and \mathcal{M}_C . For every combination of values of λ, T, h , we calculate $\max_v \frac{\mathbb{E}_v[\mathcal{R}(\mathcal{M}_C)] - \mathbb{E}_v[\mathcal{R}(\text{ESoES-SS})]}{h}$, corresponding to the maximum normalized loss of ESoES-SS w.r.t. \mathcal{M}_C over all valuations v , and $\max_v \frac{\mathbb{E}_v[\mathcal{R}(\text{ESoES-SS})] - \mathbb{E}_v[\mathcal{R}(\mathcal{M}_C)]}{h}$, corresponding to the maximum normalized loss of \mathcal{M}_C w.r.t. ESoES-SS over all valuations v . These two indexes are shown in Figure 3 for $T = 10$ and $T = 50$, when $h = 10$. The results obtained for different values of h are similar. The loss of ESoES-SS w.r.t. \mathcal{M}_C is always larger than $0.5h$ except when both λ and T assume small values, while the loss of \mathcal{M}_C w.r.t. ESoES-SS is negligible. Furthermore, the two losses converge to two constants as λ and T increase. This shows that, even if there are some special settings where ESoES-SS performs better than \mathcal{M}_C , the improvement is negligible. Instead, mechanism \mathcal{M}_C , which is designed to deal with stochastic arrivals, provides a very significant improvement. In particular, we observe that the difference between the revenue provided by ESoES-SS and that provided by \mathcal{M}_C is maximized for small values of v close to 1, while between \mathcal{M}_C and ESoES-SS for large values of v close to h .

Conclusion and Future Works

We study distribution-free posted-price mechanisms in order to sell a unique item within a finite time period. In our model, the agents arrive online according to a Poisson process, and their valuations for the item are discounted over time. Following a worst-case competitive analysis, we design a mechanism \mathcal{M}_C providing an optimal competitive ratio in the identical valuation setting. Then, as for the random valuation setting, we analyze the performances of \mathcal{M}_C and of a new mechanism \mathcal{M}_{PC} that is constrained to set constant prices during time intervals having a given minimum length. We prove that both mechanisms achieve a competitive ratio that is constant with respect to the actual valuation when the distribution of the valuations has a monotone hazard rate. This shows that our mechanisms are robust even in non-stationary markets subject to arbitrary distribution changes preserving the same support.

In future, we will investigate hybrid settings in which our robust mechanisms can be combined with machine learning tools. For instance, data could be used to learn a class of distributions, and we could design a mechanism robust with respect to all the distributions of that class.

Acknowledgments

This work has been partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR “Algorithms, Games, and Digital Market”.

Ethics Statement

Posted-price mechanisms are widely adopted in real-world economic transactions, thanks to their simplicity: a seller posts prices and buyers arrive sequentially, deciding whether to accept the offer or not. Nowadays, most e-commerce websites implement this form of interaction with their users. Our mechanisms apply to concrete scenarios where the probability distribution of buyers’ valuations is unknown, the value of item for sale may decrease over time, and buyers’ arrivals are stochastic. In these settings, our mechanisms can make economic transactions more efficient and robust, allowing agents (buyers and sellers) to find better economic agreements. As we argued in the paper, our mechanisms provide theoretical guarantees in terms of online worst-case performance. This could have an arguably positive societal impact when applied to real-world economic problems. However, further research in this direction is required to prevent scenarios with an unbalanced reward structure, where agreements may just award one side (buyers or sellers) with the largest utilities at the expense of the others.

References

- Adamczyk, M.; Borodin, A.; Ferraioli, D.; Keijzer, B. D.; and Leonardi, S. 2017. Sequential posted-price mechanisms with correlated valuations. *ACM Transactions on Economics and Computation (TEAC)* 5(4): 1–39.
- Babaioff, M.; Blumrosen, L.; Dughmi, S.; and Singer, Y. 2017. Posting prices with unknown distributions. *ACM Transactions on Economics and Computation (TEAC)* 5(2): 1–20.
- Babaioff, M.; Dinitz, M.; Gupta, A.; Immorlica, N.; and Talwar, K. 2009. Secretary problems: weights and discounts. In *Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms*, 1245–1254. SIAM.
- Babaioff, M.; Dughmi, S.; Kleinberg, R.; and Slivkins, A. 2015. Dynamic pricing with limited supply. *ACM Transactions on Economics and Computation (TEAC)* 3(1): 1–26.
- Borodin, A.; and El-Yaniv, R. 2005. *Online computation and competitive analysis*. Cambridge University Press.
- Chawla, S.; Hartline, J. D.; Malec, D. L.; and Sivan, B. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, 311–320.
- Correa, J.; Foncea, P.; Hoeksma, R.; Oosterwijk, T.; and Vredeveld, T. 2017. Posted price mechanisms for a random stream of customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 169–186.
- Einav, L.; Farronato, C.; Levin, J.; and Sundaresan, N. 2018. Auctions versus posted prices in online markets. *Journal of Political Economy* 126(1): 178–215.
- Gatti, N.; Di Giunta, F.; and Marino, S. 2008. Alternating-offers bargaining with one-sided uncertain deadlines: an efficient algorithm. *Artificial Intelligence* 172(8-9): 1119–1157.
- Hajiaghayi, M. T.; Kleinberg, R.; and Sandholm, T. 2007. Automated online mechanism design and prophet inequalities. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence*, 58–65.
- Kleinberg, R.; and Leighton, T. 2003. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, 594–605. IEEE.
- Lavi, R.; and Nisan, N. 2004. Competitive analysis of incentive compatible on-line auctions. *Theoretical Computer Science* 310(1-3): 159–180.
- Lucier, B. 2017. An economic view of prophet inequalities. *ACM SIGecom Exchanges* 16(1): 24–47.
- Mao, W.; Zheng, Z.; Wu, F.; and Chen, G. 2018. Online Pricing for Revenue Maximization with Unknown Time Discounting Valuations. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, 440–446.
- Mason, R.; and Välimäki, J. 2011. Learning about the arrival of sales. *Journal of Economic Theory* 146(4): 1699 – 1711.
- Mohri, M.; and Munoz, A. 2014. Optimal regret minimization in posted-price auctions with strategic buyers. In *Advances in Neural Information Processing Systems*, 1871–1879.
- Parkes, D. C. 2007. *Online mechanisms*. Cambridge University Press.
- Romano, G.; Tartaglia, G.; Marchesi, A.; and Gatti, N. 2020. Online Posted Pricing with Unknown Time-Discounted Valuations. *arXiv preprint arXiv:2012.05774*.
- Rong, J.; Qin, T.; and An, B. 2018. Dynamic pricing for reusable resources in competitive market with stochastic demand. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 4718–4726.
- Rosenthal, E. C. 2011. A Pricing Model for Residential Homes with Poisson Arrivals and a Sales Deadline. *The Journal of Real Estate Finance and Economics* 42(2): 143–161.
- Rubinstein, A. 1982. Perfect Equilibrium in a Bargaining Model. *Econometrica* 50(1): 97–109.
- Seifert, S. 2006. *Posted price offers in internet auction markets*, volume 580. Springer Science & Business Media.
- Shah, V.; Johari, R.; and Blanchet, J. 2019. Semi-Parametric Dynamic Contextual Pricing. In *Advances in Neural Information Processing Systems*, 2363–2373.
- Wang, R. 1993. Auctions versus posted-price selling. *The American Economic Review* 838–851.