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

MANAGEMENT SCIENCE
Vol. 66, No. 10, October 2020, pp. 4417–4432
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Prior-Independent Optimal Auctions

Amine Allouah,[a] Omar Besbes[a]

[a] Graduate School of Business, Columbia University, New York, New York 10027

Contact: mallouah19@gsb.columbia.edu, https://orcid.org/0000-0001-9919-0606 (AA); ob2105@gsb.columbia.edu, https://orcid.org/0000-0002-2982-3794 (OB)

Received: March 10, 2018
Revised: May 2, 2019
Accepted: July 19, 2019
Published Online in Articles in Advance: May 12, 2020

https://doi.org/10.1287/mnsc.2019.3459

Abstract. Auctions are widely used in practice. Although auctions are also extensively studied in the literature, most of the developments rely on the significant common prior assumption. We study the design of optimal prior-independent selling mechanisms: buyers do not have any information about their competitors, and the seller does not know the distribution of values but only knows a general class to which it belongs. Anchored on the canonical model of buyers with independent and identically distributed values, we analyze a competitive ratio objective in which the seller attempts to optimize the worst-case fraction of revenues garnered compared with those of an oracle with knowledge of the distribution. We characterize properties of optimal mechanisms and in turn establish fundamental impossibility results through upper bounds on the maximin ratio. By also deriving lower bounds on the maximin ratio, we are able to crisply characterize the optimal performance for a spectrum of families of distributions. In particular, our results imply that a second price auction is an optimal mechanism when the seller only knows that the distribution of buyers has a monotone nondecreasing hazard rate, and it guarantees at least 71.53% of oracle revenues against any distribution within this class. Furthermore, a second price auction is near optimal when the class of admissible distributions is that of those with nondecreasing virtual value function (a.k.a. regular). Under this class, it guarantees a fraction of 50% of oracle revenues, and no mechanism can guarantee more than 55.6%.

History: Accepted by Kalyan Talluri, revenue management and market analytics.
Supplemental Material: The electronic companion is available at https://doi.org/10.1287/mnsc.2019.3459.

Keywords: optimal auctions • limited information • prior-independent • robust mechanism design • competitive ratio • dominant-strategy incentive compatibility • regular distributions • monotone hazard rate distributions

1. Introduction

Auctions have been run for many centuries and play today a prominent role in applications as diverse as e-commerce, spectrum allocation, antique sales, online advertising, and procurement. In turn, auction design has been a central topic of research at the intersection of operations research, computer science, and economics. The monograph of Krishna (2009) provides an overview of auction theory, and Talluri and Van Ryzin (2006) detail many revenue management applications. Although there is an elegant theory of auction design dating back to the seminal works of, for example, Vickrey (1961) and Myerson (1981), the classical theory of auctions is anchored on a fundamental assumption: that of a common prior. This assumption stipulates that the seller as well as the buyers share the same common prior on the process generating the values for the object. In turn, this assumption leads naturally to the buyers using this common prior to play equilibrium bidding strategies, forming a Bayesian Nash equilibrium; the seller, anticipating such equilibrium behavior, can optimize the selling mechanism based on this prior. This poses a challenge in practice as such a prior is not available, and it is not clear how the seller's belief and the buyers' beliefs about values should coincide or how they would be formed correctly. In turn, a fundamental question from practical and theoretical perspectives pertains to how to relax such an assumption and what performance can one expect in its absence. This fundamental need to move beyond mechanisms that rely on priors is often referred to as the "Wilson doctrine" (Wilson 1987). Relaxing the assumption on common priors leads to a trade-off between information about the distribution of values and performance, which motivates the following questions: What is the maximum fraction of revenues that one can guarantee compared with an oracle that would have access to the underlying distribution of values? How does this fraction vary as a function of the information available about the underlying distribution? These are the central questions that this paper aims to address.

In the present paper, we aim to address the above in the canonical private value model of a seller trying to sell a good to buyers with independent and identically distributed values.[1] Although mechanism design is very well understood for this classical model under the common prior assumption, it remains challenging in prior-independent environments. (We review shortly in detail related work in Section 1.3.) As soon as one relaxes the common prior assumption, a first question is how to formulate the problem. On the one hand, the common prior affected bidding behavior of buyers. On the other hand, it also affects the seller's mechanism optimization problem. We maintain the fact that values are drawn from an underlying distribution (the true distribution of values), as in the classical framework, but we do not assume knowledge of this distribution by the buyers or the seller. In turn, one needs to specify the information available to the buyers and the resulting equilibrium, as well as the seller's knowledge and feasible mechanisms, and these two are tightly interconnected. For the buyers' side, we will adopt a detail-free approach and assume that buyers' optimal decisions are independent of any information about the other buyers' values. For that, we will restrict attention to mechanisms for which truth-telling is a dominant strategy, so-called dominant-strategy incentive compatible (DSIC) mechanisms. Against such a mechanism, buyers bidding their values represents a dominant-strategy Nash equilibrium. On the seller's side, we will assume that the seller is free to select among such mechanisms. Given that the seller does not know the true distribution of values, we will adopt a maximin ratio approach. We model our problem as a game between nature and the seller. The seller first selects a mechanism in the class of DSIC mechanisms. Then, nature may counter such a mechanism with *any* distribution for buyers' values from a given class of admissible distributions. In particular, the resulting equilibrium induced by the mechanism is dominant-strategy incentive compatible, and the only knowledge the seller is endowed with is the class of admissible distributions. For any distribution and mechanism, we measure the performance of the seller through the ratio of the revenue she garners using this mechanism over the optimal revenue she would have obtained with access to the exact knowledge of the distribution. We refer to the latter as the *oracle revenues*. The ratio is always between 0 and 1, and the higher the ratio, the better the performance. We focus on a maximin setting in which the seller attempts to maximize the worst-case performance ratio (or competitive ratio) over the class of admissible distributions.

Our results provide a characterization of the maximin ratio across a spectrum of distribution classes. In particular, we consider three main classes of distributions. It is possible to show that against the general class of distributions, no DSIC mechanism can guarantee a positive fraction of oracle revenues, and hence there is a need to study how different structures of the underlying distributions affect the type of performance that can be achieved. Beyond the general class of distributions, we will consider a class that is central to mechanism design (including under the common prior assumption), that of so-called regular distributions. These are distributions that admit increasing virtual value function. In addition to the class of regular distributions, we will also analyze the subclass of monotone increasing hazard rate (MHR) distributions (also often referred to as increasing failure rate distributions), which contains many distributions often assumed in practice and in the literature (e.g., uniform, exponential).

1.1. Summary of Contributions

Before laying out our main results, it is important to highlight the nature of the problem we study. On the one hand, given a particular mechanism, nature selects the worst possible distributions in the nonparametric classes above. So nature, when minimizing the ratio of revenues compared with oracle performance, is solving a nonconvex infinite dimensional optimization problem. In turn, fully understanding the worst-case performance of a specific mechanism is highly nontrivial and not necessarily tractable. On the other hand, the seller, when optimizing over DSIC mechanisms, is also solving an infinite dimensional problem (over allocation and payment mappings). An important contribution of the present paper is to propose an approach to tackle this class of problems and characterize optimal or near-optimal performance.

For regular distributions, it is known that a second price auction[2] guarantees, in the worst-case scenario, 50% of the oracle revenues, as articulated in Dhangwatnotai et al. (2015) through a reinterpretation of the results in Bulow and Klemperer (1996). Notably, Fu et al. (2015) recently establish that a second price auction is not prior-independent optimal. In particular, they exhibit a mechanism that randomizes between a second price auction and an auction that inflates the second value, and they establish that it ensures a competitive ratio of at least 51.2%. Table 1 summarizes the best-known lower bounds on the maximin ratios as well as implications of our results. Although there is a lower bound on the maximin ratio against regular distributions, there is no notion of what performance one should aim at and how good are the prior-independent auctions previously proposed. In the popular subclass of MHR distributions, to the best of our knowledge, no lower or upper bounds are available in the literature.

A first significant layer of contribution pertains to the methodological domain and allows for obtaining

Table 1. Maximin Performance

	Maximin ratio (%)			
	Lower bounds		Upper bounds	
Distribution class	Best known	This paper	This paper	Best known
Regular	51.2	51.9	55.6	X
MHR	X	71.53	71.53	X

Notes. The table contrasts known results in the existing literature with the bounds derived on the maximin ration through the analysis in the present paper. X, no known bound.

the first impossibility results for *any* mechanism in a broad class of DSIC mechanisms. We mainly focus on the case with two buyers, which intuitively is the case with most tension, while relaxing the common prior assumption, and then we establish that the bounds obtained for the case of two buyers also apply to the case where the number of buyers is adversarially selected (in Section 7.1).

We first develop families of tractable upper bounds on the maximin ratio. These are obtained through successive dimensionality reductions on the space of mechanisms and the space of distributions. We show that, under some mild regularity assumption on the mechanisms, an optimal mechanism is scale-free (see Theorem 1). In other words, it is sufficient to focus on mechanisms that rely only on the ratio of values of buyers. In turn, leveraging properties of the allocations, we are able to "discretize" the mechanisms without loss of optimality and reduce the description of mechanisms to a countable set (Proposition 1).

Given the result above, we then introduce general subsets of distributions. These abstract subsets are developed in order to, on the one hand, be "hard" for any mechanism and, on the other hand, allow for further reduction of the complexity of the set of mechanisms under consideration, leading to a new generic upper bound (Theorem 2). By customizing this bound through appropriate concrete classes and leveraging additional properties of the classes, we obtain parametric upper bounds for the maximin ratio against regular distributions (Theorem 3) and MHR distributions (Theorem 4). In turn, these upper bounds lead to the first impossibility results for general randomized mechanisms against these two central classes of distributions. No DSIC mechanism considered can guarantee more than 55.6% of oracle performance against all regular distributions, and no DSIC mechanism considered can guarantee more than 71.53% of oracle performance against all regular distributions.

These results have a significant implication for regular distributions. They imply that the mechanisms proposed to date in the literature are, in fact, near optimal. A second price auction is within 5.6% of optimal, and the mechanism proposed in Fu et al. (2015) is within 4.4% of optimal. These impossibility results allow for quantifying the quality of *any* mechanism compared with *optimal* performance in the class of DSIC mechanisms.

As a second layer of contribution, we also develop lower bounds on the maximin ratio. We develop a series of generic parametric lower bounds (Propositions 4 and 5) and in turn obtain lower bounds on the worst-case performance of specific mechanisms. For the case of regular distributions, we establish that there exists a mechanism that guarantees at least 51.9%, improving the best-known lower bound and further closing the gap with the upper bound we have developed. For the case of MHR distributions, we establish that a particular mechanism, a second price auction, guarantees at least 71.53% of oracle performance.

Whereas we improve the lower bound on regular distributions, the significant implication of the lower bounds is for the MHR class. The first implication stems from comparing it to the novel upper bound we derive for regular distributions. In particular, our results show how refined class information (from regular to MHR) translates into improved performance. Against MHR distributions, even with only two buyers, a seller is guaranteed 71.53% of oracle performance. The second implication is even more notable. The conjunction of our upper and lower bounds imply that a second price auction is actually *optimal* against MHR distributions and that we have exactly characterized the maximin ratio for that class. Overall, the results above provide a crisp characterization of the maximin ratio as information regarding distributions is refined.

In addition, the results shed light on the trade-off that an auctioneer might face between running an auction with limited information and the cost of collecting additional information to approach the oracle optimal revenue. Our results highlight how this trade-off might be affected by the nature of distributions that a decision maker might face, for example, if distributions are more "concentrated" (as is the case for MHR).

From a different angle, in practice, there is also often a trade-off between revenue maximization and social efficiency. In the canonical class studied, our results highlight that, in a prior-independent environment, a second price auction is near optimal for the wide class of regular distributions and optimal for the large subclass of MHR distributions. As such, when limited information about the underlying distribution of values is available, a simple, practical, and socially efficient mechanism appears "sufficient" from a revenue maximization perspective. Hence, there is a weak trade-off between revenue maximization and social efficiency when facing regular distributions and no trade-off when facing MHR distributions.

1.2. The Remainder of the Paper

After relating our paper to the existing literature, we formulate our problem and set up our framework for two buyers. In Section 3, we establish that one may restrict attention to scale-free mechanisms and characterize the maximin ratio for general distributions. In Section 4, we derive a family of upper bounds on the maximin ratio against subsets of regular distributions. In Section 5, we investigate the case of regular distributions, whereas the subset of MHR distributions is the focus of Section 6. Then, in Section 7, we extend our results to the case in which the number of buyers is arbitrary and adversarially selected, and we discuss future directions. All proofs are presented in the electronic companion.

1.3. Literature Review

Our work relates to a rich literature on auction design. Since the seminal work Myerson (1981) that characterized the structure of an optimal revenue-maximizing mechanism when the seller has access to the exact distributions of values of buyers, the research community has raised early on the need of designing auctions that do not rely on such informational assumptions, often referred to as the "Wilson doctrine" (Wilson 1987). Our work belongs to the stream that aims to relax such assumptions. There are different layers of informational assumptions that have been analyzed in the literature. Some layers relate to the seller's knowledge about the distributions of values of buyers or the number of participating buyers. Other layers relate to the knowledge of buyers about their own values as well as the values or number of competitors.

When relaxing informational assumptions in auction design, there are two implications. On the one hand, the information affects the type of mechanisms that the seller can adopt. On the other hand, the information also affects the type of equilibrium played by the competing buyers.

In terms of the assumption that each buyer makes on the value-generating process of his competitors, various alternatives have been analyzed. One extreme is to assume that the buyers know their competitors' distributions of values. In this case, Caillaud and Robert (2005) show that the seller could exploit this and recover the optimal oracle revenue even if she does not have access to the distributions of values through a dynamic mechanism. A first relaxation is to assume that the buyers know some ambiguity sets characterizing the distributions; see Bose et al. (2006), Chiesa et al. (2015), and Koçyiğit et al. (2020). A further relaxation is to assume that buyers do not have access to any information about values of other buyers; this is typically done by assuming DSIC mechanisms. We refer the reader to Chung and Ely (2007), who give a formal foundation of such an assumption by showing that a dominant-strategy mechanism always dominates in terms of revenue any other mechanisms, when the buyers' beliefs about the distribution of their competitor are selected adversarially. Our work aims to make minimum assumptions on both the seller's and the buyer's side, and in turn, we focus on DSIC mechanisms. Furthermore, we do not make any assumption on the buyer's knowledge on the number of competitors.

Another line of work relaxes the knowledge of the buyers regarding her true value, by assuming that the buyer observes some signal related to the true value. We refer the reader to, for example, Bergemann et al. (2016), who aim to characterize optimal auctions when there is uncertainty on the information structure of the buyers. See also Bergemann and Morris (2013) for a broader overview. We would like to note that in this line of work, it is typically assumed that the seller knows the distribution of values of buyers. Compared with our work, we assume that the seller does not know the distribution of values of buyers but knows the information structure of the buyers. Furthermore, the DSIC assumption also implies that the equilibrium of buyers does not depend on the underlying distributions of values. The buyers' strategies also does not depend on the number of buyers. In that regard, we also note that another dimension of information on the side of buyers pertains to the number of buyers. Harstad et al. (1990) and Levin and Ozdenoren (2004) relax this while maintaining knowledge of the distribution of values.

Once information on the buyers' side is formulated, the next dimension relates to the layer of information that the seller has. In that regard, there are at a high level three main classes of information structures assumed on the knowledge of distributions of values of buyers: nonparametric canonical classes of distributions (Dhangwatnotai et al. 2015), statistics of the distributions (Azar et al. 2013), and uncertainty sets on the distributions (Koçyiğit et al. 2020) or the values (Bandi and Bertsimas 2014).

Finally, a fundamental other dimension pertains to how performance of a mechanism is measured in such an environment. One approach, typically referred to as "robust," is to use the *absolute* worst-case performance based on the information available; see, for example, Carrasco et al. (2015), Bandi and Bertsimas (2014), and Koçyiğit et al. (2020). Another approach is to measure worst-case performance *relative* to a full information benchmark; see, for example, Neeman (2003) and Dhangwatnotai et al. (2015). A more detailed discussion on various candidate objectives can be found in Borodin and El-Yaniv (1998). Our work relates to the last branch of literature because we characterize the optimal competitive ratio when the seller has only access to the class of distributions of buyers. The ratio we analyze is unitless and has a physical interpretation in terms of the fraction of oracle performance one can obtain compared with an oracle.[3]

In this stream, an important set of results pertain to "existence" of mechanisms with good guarantees. Looking at different classes of distributions, Neeman (2003) derives an early result and establishes a guarantee for the English auction, compared with the social optimum. In particular, the author characterizes tight lower bounds as a function of some summary statistics on the performance of an English auction with or without reserve price. The setting we focus on is the independently and identically distributed values case. In this setting, if the seller knows that the distribution of values of buyers belongs to the regular class of distributions, then an implication of classical results of Bulow and Klemperer (1996), based on the interpretation of Dhangwatnotai et al. (2015), is that there exists a particular mechanism—namely, a second price auction—that extracts 50% of the oracle revenue had one known the true distribution, against *any* regular distribution. Recently, Fu et al. (2015) show that a second price auction is suboptimal against regular distributions by exhibiting a randomized mechanism that has a higher guarantee than a second price auction. In the present work, we focus on optimizing over a very broad class of DSIC mechanisms and in turn establish fundamental impossibility results for any such mechanism. The results complement the literature by not only characterizing what is achievable by a particular mechanism but also characterizing *optimal* performance through upper bounds on the maximin ratio. Furthermore, by focusing on the widely considered subclass of MHR distributions, we establish that a second price auction is actually the exact optimal mechanism in that case. This also sheds light on the role of randomization and its relationship to the class of distributions one faces.

In the case of multiple goods, Goldberg et al. (2006) introduce and analyze the competitive ratio, where the worst case is taken with respect to any possible inputs, and then establish that some auctions are competitive compared with a fixed pricing benchmark. In more general environments, Dhangwatnotai et al. (2015) leverage the connection to Bulow and Klemperer (1996) to propose a mechanism that has a nontrivial performance even in general allocation environments. Relatedly, Sivan and Syrgkanis (2013) extend a result of Bulow and Klemperer (1996) to the case in which the distributions of values of buyers are a convex combination of regular distributions.

A related stream of literature focuses on alternative information about the distribution. For instance, Azar and Micali (2012) and Azar et al. (2013) propose mechanisms in cases in which the seller has access to some summary statistics of the distributions of values of buyers (mean or median). They exhibit mechanisms that have performance guarantees compared with an oracle using these. In the present paper, we do not assume that the seller has access to some summary statistics, and we focus on the optimal mechanism among a broad set of randomized mechanisms.

In our paper, we focus on a static model with limited information. Other examples of directions analyzed pertain to the amount information available or the dynamics. Cole and Roughgarden (2014) analyze the size of the sample that the seller needs to observe from past data in order to design a near optimal mechanism. Dynamic models have also been considered in the literature; see, for example, Bose and Daripa (2009) for a dynamic model under ambiguity. We refer the reader to review papers of Hartline and Roughgarden (2009), Hartline (2013), and Carroll (2019) for a broader overview.

Another information assumption from seller's perspective that the literature has tried to relax is the knowledge of the exact number of buyers. For instance, while maintaining the common prior assumption, McAfee and McMillan (1987) characterize the optimal auctions when the seller has some prior on the number of buyers, and Levin and Ozdenoren (2004) study the seller's best response when the number of buyers is picked adversarially from some ambiguity set.

Our work also relates to pricing under limited information. Monopoly pricing with unknown demand information was analyzed with various considerations in Bergemann and Schlag (2008) for a minimax regret objective and in Eren and Maglaras (2010) for the competitive ratio. Caldentey et al. (2016) extend this line of work to account for the presence of strategic customers. Cohen et al. (2016) derive performance guarantees for pricing heuristics when the firm has some knowledge about the demand shutdown price. More recently, Chen et al. (2017) study robust single-item and bundle pricing based on summary statistics of buyers' values distribution. Leveraging existing data, Huang et al. (2015) focus on pricing based on a finite sample of values. There is also an extensive body of work on joint learning and pricing with various informational structures. We refer the reader to Kleinberg and Leighton (2003), Keskin and Zeevi (2014), and Besbes and Zeevi (2015) for various informational structures, as well as to Besbes and Zeevi (2009), Araman and Caldentey (2009), Farias and Van Roy (2010), and Wang et al. (2014) for inventory considerations in such pricing problems. den Boer (2015) provides a survey of this line of work.

2. Problem Formulation

We consider a seller offering an indivisible object for sale to two buyers. For now, we focus on the two-buyers case because it is the case with minimum competition, and it isolates the impact of relaxing the common prior assumption. We return to the case of more than two buyers in Section 7. The two buyers have values identically and independently distributed according to a distribution F with support S_F in $[0, \infty)$. We will denote by $\overline{F}(\cdot) := 1 - F(\cdot)$ the complementary cumulative distribution function (ccdf) of values.

We assume that the seller does not know exactly the distribution of values of buyers; however, she knows that it belongs to a particular class. The goal of the seller is to design a mechanism that maximizes her revenue given the limited information about the underlying distribution of values of buyers.

2.1. Seller's Problem

We model our problem as a game between the seller and nature, in which the seller selects a prior-independent selling mechanism and then nature may counter such a mechanism with *any* distribution of buyers' values from an admissible class.

A selling mechanism $m = (\boldsymbol{x}, \boldsymbol{t})$ is characterized by an allocation mapping $\boldsymbol{x}$ and a payment mapping $\boldsymbol{t}$, where $\boldsymbol{x} : \mathbb{R}^2 \to [0,1]^2$ and $\boldsymbol{t} : \mathbb{R}^2 \to \mathbb{R}$. In particular, given reports b_1, b_2 by buyers 1 and 2, a mechanism would allocate the good to buyer i with probability $x_i(b_i, b_{-i})$ and the expected payment of buyer i is $t_i(b_i, b_{-i})$. Here, and in all that follows, the notation (v_i, v_{-i}) is the vector that has value v_i at position i and v_{-i} at the other position.

We do not make any assumption on the buyer's knowledge of the distribution. Given this, we will restrict attention to DSIC mechanisms. For such mechanisms, buyers need not make any assumptions about the underlying distribution of values and will find it optimal to report their true value, independent of the realization of value of the other buyers.[4]

More formally, we focus on the class of mechanisms $m = (\boldsymbol{x}, \boldsymbol{t})$ that satisfy the following constraints:

$$v_i\, x_i(v_i, v_{-i}) - t_i(v_i, v_{-i}) \geq 0, \quad \text{for all } i \text{ and } v_i, v_{-i} \text{ in } \mathbb{R}_+^2 \tag{IR}$$

$$v_i\, x_i(v_i, v_{-i}) - t_i(v_i, v_{-i}) \geq v_i\, x_i(\hat{v}_i, v_{-i}) - t_i(\hat{v}_i, v_{-i}), \quad \text{for all } i \text{ and } v_i,\ v_{-i},\ \hat{v}_i \text{ in } \mathbb{R}_+^3, \tag{IC}$$

$$\sum_{i=1,2} x_i(v_i, v_{-i}) \leq 1, \quad \text{for all } \boldsymbol{v} \text{ in } \mathbb{R}_+^2. \tag{AC}$$

The first constraint (IR) captures ex post *individual rationality* and states that buyer i should be willing to participate compared with his outside option, normalized to 0. The second constraint (IC) captures ex post *incentive compatibility* and imposes that a buyer should always find it optimal to report his true value, independent of the value of the other buyer. Finally, (AC) is a constraint on the allocation probabilities that captures that the seller can allocate at most one good. Note here that we allow for randomized mechanisms by the seller. In addition, we will introduce a regularity assumption on mechanisms. We denote by $TV(x_i, [a,b] \times [c,d])$ the Arzelà total variation of the allocations on the set $[a,b] \times [c,d]$.[5] We assume that the allocations around 0 have finite Arzelà total variation. In particular, we will be focusing on the following set of mechanisms:

$$\mathcal{M} = \left\{ (\boldsymbol{x}, \boldsymbol{t}) \colon \text{(IR)}, \text{(IC)}, \text{(AC)} \text{ and } \max_{i=1,2} \left\{ TV\left(x_i, [0, \varepsilon]^2\right) \right\} < \infty \text{ for some } \varepsilon > 0 \right\}. \tag{1}$$

This class of mechanisms is a rich one, containing, for example, the second price auction with a deterministic reserve price and most mechanisms typically considered in the literature. The assumption on the boundedness of the total variation of allocations around 0 is technical in nature[6] but could also be seen as a way to avoid potentially overly complex mechanisms that might be hard to implement in practice, given the high burden this would put on the buyers.[7]

The revenue of the seller using a feasible mechanism m in $\mathcal{M}$, if nature is selecting a distribution F, is given by

$$\mathbb{E}_F\left[\sum_{i=1}^{2} t_i(v_i, v_{-i})\right].$$

We will use the subscript F to emphasize that the expectation is taken with respect to that distribution.

The challenge in the present paper is that the seller does not know the distribution F and as a result cannot evaluate the objective above to select a "good" or optimal mechanism. We next introduce a performance benchmark and pose a proper objective for the seller for this environment with an unknown distribution of values.

2.1.1. Oracle Benchmark. The benchmark we will use, opt(F), is the maximal performance one could achieve *with knowledge of the exact distribution of buyers' values* when selecting mechanisms in $\mathcal{M}$. More formally,

$$\text{opt}(F) := \sup_{m\in\mathcal{M}} \mathbb{E}_F\left[\sum_{i=1}^{2} t_i(v_i, v_{-i})\right]. \tag{2}$$

2.1.2. Seller's Objective. For an arbitrary mechanism in $\mathcal{M}$, we define its performance against a distribution F such that opt(F) > 0 as follows:

$$R(m,F) = \frac{\mathbb{E}_F\left[\sum_{i=1}^{2} t_i(v_i, v_{-i})\right]}{\text{opt}(F)}.$$

In other words, $R(m,F)$ represents the fraction of the oracle benchmark performance the mechanism is able to achieve. The ratio $R(m,F)$ always lies in $[0,1]$, and the closest the ratio is to 1, the better the performance of the mechanism.

Let $\mathcal{G}$ denote the set of distributions with support included in $[0,\infty)$ with finite and nonzero expectation; that is,

$$\mathcal{G} = \{F : [0,\infty) \to [0,1] : F \text{ is a cdf and } 0 < \mathbb{E}_F[v] < \infty\}. \tag{3}$$

Note that $\mathbb{E}_F[v] > 0$ if only if opt(F) > 0. Hence the ratio $R(m,F)$ is well defined for any element of the class $\mathcal{G}$.

The objective of the present paper is to characterize for classes $\mathcal{F} \subseteq \mathcal{G}$ the maximin ratio:

$$\mathcal{R}(\mathcal{M},\mathcal{F}) = \sup_{m\in\mathcal{M}} \inf_{F\in\mathcal{F}} R(m,F). \tag{4}$$

In other words, we are interested in designing mechanisms that admit "good" performance independently of the underlying distribution of values. In particular, the value $\mathcal{R}(\mathcal{M},\mathcal{F})$ represents the maximal fraction of oracle revenues (obtained with knowledge of the distribution of values) that can be recovered when nature may select any distribution in $\mathcal{F}$.

Definition 1. *A cdf F is said to be regular on its support S_F if it admits a density f and if the corresponding virtual value function $\phi_F : v \mapsto v - (1 - F(v))/f(v)$ is nondecreasing over S_F. We will further say that the distribution has MHR if $v \mapsto f(v)/(1 - F(v))$ is nondecreasing over S_F.*

The class of regular distributions is very widely used and plays a central role in mechanism design (with knowledge of the distribution of buyers), and the class of monotone hazard rate distributions is a wide subclass of the set of regular distribution that encompasses all distributions with log-concave densities (e.g., uniform, exponential). In particular, beyond $\mathcal{G}$, we will analyze the two subclasses of distributions:

$$\mathcal{F}_{reg} = \{F \in \mathcal{G} : F \text{ is regular}\},$$
$$\mathcal{F}_{mhr} = \{F \in \mathcal{G} : F \text{ has a monotone nondecreasing hazard rate}\}$$

It is clear that we have $\mathcal{F}_{mhr} \subset \mathcal{F}_{reg} \subset \mathcal{G}$, and hence

$$\mathcal{R}(\mathcal{M},\mathcal{G}) \le \mathcal{R}(\mathcal{M},\mathcal{F}_{reg}) \le \mathcal{R}(\mathcal{M},\mathcal{F}_{mhr}).$$

In the coming sections, we will be interested in quantifying the three quantities above and characterizing optimal or near-optimal mechanisms.

2.1.3. Review of Some Known Results. Although, to the best of our knowledge, the problem above has not been addressed in the literature, some mechanisms m have been exhibited and their performance characterized. A classical mechanism in $\mathcal{M}$ is the second price auction m_{spa}, defined by

$$x_i(v_i, v_{-i}) = \mathbb{1}\{v_i > v_{-i}\} + .5\ \mathbb{1}\{v_i = v_{-i}\},$$
$$t_i(v_i, v_{-i}) = v_{-i}\mathbb{1}\{v_i > v_{-i}\} + .5\ v_{-i}\ \mathbb{1}\{v_i = v_{-i}\}.$$

The results of Bulow and Klemperer (1996) and their reinterpretation for the performance of the second price auction (see, e.g., Dhangwatnotai et al. 2015) imply that

$$\inf_{F\in\mathcal{F}_{reg}} R(m_{spa}, F) = 50\%.$$

Recently, Fu et al. (2015) exhibited a mechanism m that randomizes between the identity and a mapping that inflates the second-highest value and established that

$$\inf_{F\in\mathcal{F}_{reg}} R(m,F) \ge 51.2\%.$$

The results above imply a lower bound on $\mathcal{R}(\mathcal{M},\mathcal{F}_{reg})$ through specific mechanisms but leave open the question of *optimal* performance. In the present paper, we aim at characterizing the maximin ratio (4) and corresponding near-optimal solutions not only for $\mathcal{F}_{reg}$ but also for $\mathcal{G}$ and $\mathcal{F}_{mhr}$.

3. Optimality of Scale-Free Mechanisms

The goal of this section is to establish that one may reduce the space of mechanisms to a simpler class, without loss of optimality. In particular, we will establish that one may restrict attention to scale-free mechanisms (as defined later in Equation (5)).

We first state a classical result from the mechanism design literature (see Myerson 1981) that links payments and allocations for any incentive compatible mechanism.

Lemma 1. *A mechanism* $(\boldsymbol{x}, \boldsymbol{t})$ *verifies* (IC) *if and only if* $x_i(\cdot, v_{-i})$ *is nondecreasing for any* $v_{-i} \geq 0$ *and the payment mapping satisfies*

$$t_i(v_i, v_{-i}) = v_i \; x_i(v_i, v_{-i}) - \int_0^{v_i} x_i(l, v_{-i}) dl + t_i(0, v_{-i}), \quad \textit{for all } v_i, v_{-i} \geq 0.$$

Note that by the constraint (IR), $t_i(0, v_{-i}) \leq 0$. Hence, we can restrict attention to mechanisms that set $t_i(0, v_{-i}) = 0$ without loss of optimality. With some abuse of notation, we impose this additional constraint in the class of mechanisms $\mathcal{M}$. In other words, given (IC), we can restrict attention to allocations that are monotone in own values, and payments are fully determined by the allocations.

Before stating the main result of this section, let us now introduce some definitions pertaining to scaled distributions as well as a scale-invariant classes of distributions.

For any distribution F in $\mathcal{G}$, and $\theta > 0$, we define $F_\theta(\cdot) := F(\theta \, \cdot)$ to be the θ-scaled distribution.

Definition 2 (Scale Invariance). *A class of distributions* $\mathcal{F} \subseteq \mathcal{G}$ *is said to be invariant under scaling if for any element* F *in* $\mathcal{F}$, *the distribution* F_θ *also belongs to* $\mathcal{F}$ *for any* $\theta > 0$.

Note that $\mathcal{G}$, $\mathcal{F}_{reg}$, and $\mathcal{F}_{mhr}$ are all scale invariant. The scale invariance of $\mathcal{G}$ follows from the fact that $\mathbb{E}_{F_\theta}[v] = \theta^{-1}\mathbb{E}_F[v]$. For $\mathcal{F}_{mhr}$ and $\mathcal{F}_{reg}$, note that for any F, we have, for all v in its support,

$$\frac{f_\theta(v)}{1 - F_\theta(v)} = \theta \frac{f(\theta v)}{1 - F(\theta v)} \quad \text{and}$$
$$v - \frac{1 - F_\theta(v)}{f_\theta(v)} = \frac{1}{\theta}\left(\theta v - \frac{1 - F(\theta v)}{f(\theta v)}\right).$$

Hence, the MHR and regularity properties of any distributions F_θ are inherited from the original distribution F.

3.1. Scale-Free Mechanisms

Recall the class of mechanisms $\mathcal{M}$ introduced in (1). We next introduce the subclass of scale-free mechanisms $\mathcal{M}_{sf} \subset \mathcal{M}$, defined as follows:

$$\mathcal{M}_{sf} = \{m \in \mathcal{M} : x_i(\theta v_i, \theta v_{-i}) = x_i(v_i, v_{-i}) \quad \text{for all } v_1, v_2 \geq 0, \theta > 0, i = 1, 2\}. \tag{5}$$

This subclass of mechanisms have the property that the allocations do not depend on the scale of values. With these definitions in place, we may now state the main result of this section.

Theorem 1. *For any class* $\mathcal{F} \subseteq \mathcal{G}$ *that is invariant under scaling, when solving* (4), *it is sufficient to consider scale-free mechanisms. Namely, we have*

$$\mathcal{R}(\mathcal{M}_{sf}, \mathcal{F}) = \mathcal{R}(\mathcal{M}, \mathcal{F}).$$

This result establishes that we can restrict attention to the scale-free mechanisms without loss of optimality. Intuitively, an optimal prior-independent mechanism should not depend on the scale of buyers' values. If that were the case, then nature could leverage it to significantly affect the performance of the seller. The proof builds on this idea by evaluating a mechanism in $\mathcal{M}$ against a particular distribution and noting that the performance of this mechanism against any scaled version of the distribution serves as an upper bound on the worst-case performance of this mechanism. (For this step, we leverage the boundedness of the total variation of feasible mechanisms around 0.) In turn, by "swapping" the scale from the distribution to the mechanism, we establish that the limiting performance of the mechanism against a scaled version of the distribution as the scale goes to ∞ can be reinterpreted as the performance of a scale-free mechanism against the original distribution. In other words, we obtain that there exists a scale-free mechanism that performs at least as well (in the worst case) as the original mechanism.

The reduction to scale-free mechanisms significantly simplifies the set of mechanisms under consideration, and we will leverage this property to further reduce the space of mechanisms in upcoming sections when we consider regular distributions and its subsets. Before that, we directly leverage Theorem 1 to characterize the maximin ratio under arbitrary distributions $\mathcal{G}$, defined in (3).

In the previous literature, it was alluded to that without restrictions, the seller cannot have any guarantee (Dhangwatnotai et al. 2015). For completeness, we formalize this here in our specific context.

Lemma 2. *No mechanism in* $\mathcal{M}$ *can achieve a positive maxmin ratio against the general class* $\mathcal{G}$*—namely,*

$$\mathcal{R}(\mathcal{M}, \mathcal{G}) = 0.$$

Lemma 2 shows that it is *impossible* for the seller to design a mechanism that achieves positive worst-case performance against arbitrary distributions. The proof relies on two main ideas. Given Theorem 1, one may restrict attention to scale-free mechanisms. In turn, we establish that if the value of a buyer is 0, then necessarily, a scale-free mechanism charges 0 to the other buyer, independent of its value. Given this, we establish that the performance of any scale-free mechanism when facing the family of the Bernoulli distribution of values can be arbitrarily small.

In the rest of the paper, we focus on characterizing the maximin ratio for the set of regular distributions $\mathcal{F}_{reg}$ and the set of monotone hazard rate distributions $\mathcal{F}_{mhr}$.

4. Maximin Ratio for Subsets of Regular Distributions

In this section, we focus on the development of a family of upper bounds on $\mathcal{R}(\mathcal{M}, \mathcal{F})$ for any $\mathcal{F}$ that is a subset of the class of regular distributions $\mathcal{F}_{reg}$. In particular, the analysis of this section applies to both $\mathcal{F}_{reg}$ and $\mathcal{F}_{mhr}$, and we will leverage these results in Sections 5 and 6, when we specialize the analysis to those classes.

In Section 4.1, we establish that one may, without loss of optimality, restrict attention to a simpler set of mechanisms that are characterized by a sequence of thresholds. In Section 4.2, we focus on a simplification of the set of distributions against which one competes, which leads to a further simplification of the set of mechanisms one needs to consider. The conjunction of results leads to a generic family of upper bounds on $\mathcal{R}(\mathcal{M}, \mathcal{F})$ presented in Theorem 2.

Regarding the oracle performance for regular distributions, note that when the distribution of values F is known and is regular, it is a standard result (see Myerson 1981) that an optimal mechanism is given by a second price auction with the reserve price given by $r_F := \phi_F^{-1}(0)$, and in turn,

$$\text{opt}(F) = \mathbb{E}_F\left[\phi_F(\max\{v_1, v_2\})\mathbb{1}\{\max\{v_1, v_2\} \geq r_F\}\right].$$

In particular, the optimal oracle mechanism depends on the knowledge of the distribution through the reserve price. In what follows, we denote by $q_F = 1 - F(r_F)$ the quantile associated with r_F.

4.1. From General Mechanisms to Discrete Threshold Mechanisms

Our first result consists of a reduction of the set of mechanisms that one needs to focus on when the seller faces a subset of regular distributions. To that end, we introduce the subset of mechanisms $\mathcal{M}'_{sf}$ defined by

$$\mathcal{M}'_{sf} = \Bigg\{m \in \mathcal{M}_{sf} : \text{for } i = 1, 2,\ x_i(v_i, v_{-i}) = \sum_{n=1}^{N} \frac{1}{N}\ \mathbb{1}\{v_i > \gamma_n v_{-i}\}\mathbb{1}\{v_i \neq v_{-i}\} + c\ \mathbb{1}\{v_i = v_{-i}\}, \text{ for some } N \geq 1,\ \gamma \in \mathbb{R}^N \text{ and } c \in [0, 1/2]\Bigg\}.$$

Note first that this set $\mathcal{M}'_{sf}$ is nonempty. For example, the second price auction (without reserve price) belongs to this set. (To see that, one can take $N = 1$, $\gamma_1 = 1$, and $c = 1/2$.) This set represents a subset of the scale-free mechanisms $\mathcal{M}_{sf}$; it consists of mechanisms that are constructed using a randomization over prices to be paid by the buyer that is a linear transformation of the value of the competitor.[8] The next result characterizes the performance of mechanisms in $\mathcal{M}'_{sf}$.

Proposition 1. *For any subclass $\mathcal{F}$ of the set of regular distributions $\mathcal{F}_{reg}$, it is sufficient to focus on mechanisms in $\mathcal{M}'_{sf}$; that is,*

$$\mathcal{R}(\mathcal{M}'_{sf}, \mathcal{F}) = \mathcal{R}(\mathcal{M}_{sf}, \mathcal{F}).$$

Proposition 1 shows that, without loss of optimality, we can focus on mechanisms that belong to $\mathcal{M}'_{sf}$. Furthermore, note that this result allows one to move from a (potentially intractable) functional space of mechanisms, $\mathcal{M}_{sf}$, to the union of finite dimensional vector spaces, $\mathcal{M}'_{sf}$.

The result relies on three key ingredients. We first leverage the monotonicity of the allocations (see Lemma 1) to establish that one may approximate those from below by a combination of step functions, where the steps are chosen so that the new allocation stays appropriately close to the original allocation. This leads to a new mechanism in $\mathcal{M}'_{sf}$. Then, leveraging the scale-free property of mechanisms and the fact that the distributions are regular, we can establish that, necessarily, the performance (in terms of the ratio of revenues achieved compared the optimal oracle revenues) of the new mechanism is necessarily appropriately close to that of the original mechanism.

4.2. Family of Upper Bounds on $\mathcal{R}(\mathcal{M}, \mathcal{F})$

Having reduced the strategies of the seller to a more tractable space by discretizing the allocation function, we next reduce the complexity of the space of distribution functions under consideration $\mathcal{F} \subset \mathcal{F}_{reg}$. To that end, we introduce the subclass of distributions,

$$\mathcal{W} := \Bigg\{F \in \mathcal{G} : \overline{v}_F < \infty,\ F \text{ admits a density on } [\underline{v}_F, \overline{v}_F) \text{ and } \sup_{v \in [\underline{v}_F, \overline{v}_F)} \phi_F(v) \leq 0\Bigg\},$$

where for any distribution $F \in \mathcal{G}$, we let $\underline{v}_F = \inf\{x : x \in S_F\}$ and $\overline{v}_F = \sup\{x : x \in S_F\}$. In particular, $\mathcal{W}$ denotes the class of distributions with bounded support that have nonpositive virtual value function on the interior of the support and a potential mass at the upper limit of the support. Note that this set is clearly nonempty, and we will consider explicit examples in Sections 5 and 6. Moreover, note also that for each element of $\mathcal{W}$, the expectation of the virtual value function is not necessarily equal to the expected revenue. The expected revenue is given by

$$\mathbb{E}_F[t_i(v_i, v_{-i})] = \int_0^{\bar{v}} \phi_i(v_i)\bar{x}_i(v_i)f(v_i)dv_i + \bar{F}(\bar{v})\left(\bar{v}\bar{x}_i(\bar{v}) - \int_0^{\bar{v}} \bar{x}_i(s)ds\right),$$

where $\bar{x}_i(v_i) = \int_0^\infty x_i(v_i, u)f(u)du$ is the interim allocation to buyer i. In addition to the "classical" first term on the right-hand side, a second term, driven by the mass at $\bar{v}$, is also present.

Note that the payment of any mechanism in $\mathcal{M}'_{sf}$ takes the form (see Lemma B-5 in the electronic companion)

$$t_i(v_i, v_{-i}) = \sum_{k=1}^{N} \frac{1}{N}\ \gamma_k v_{-i}\ \mathbb{1}\{v_i > \gamma_k v_{-i}\}\mathbb{1}\{v_i \neq v_{-i}\} + c'\ v_{-i}\ \mathbb{1}\{v_i = v_{-i}\},$$

for some appropriate c'. In particular, when evaluating the expected revenues of a mechanism in $\mathcal{M}'_{sf}$, one needs to consider terms of the form $\mathbb{E}_F[v_2\mathbb{1}\{v_1 > \alpha v_2\}]$. The next result establish that one may characterize the performance of terms of $\mathbb{E}_F[v_2\mathbb{1}\{v_1 > \alpha v_2\}]$, not only for elements of $\mathcal{F}_{reg}$ but also for limits of such elements.

Lemma 3. *Suppose that a sequence $\{F_n : n \geq 1\}$ in $\mathcal{F}_{reg}$, with $\sup_{n\geq 1}\{\bar{v}_{F_n}\} < \infty$, converges weakly to a distribution F, where the latter has at most a discontinuity at $\bar{v}_F < \infty$. Then, for any $\alpha \geq 0$,*

$$\lim_{n\uparrow\infty} \mathbb{E}_{F_n}[v_2\mathbb{1}\{v_1 > \alpha v_2\}] = \begin{cases} \mathbb{E}_F[v_2\mathbb{1}\{v_1 > \alpha v_2\}], & \text{if } \alpha \neq 1, \\ \frac{1}{2}\ \mathbb{E}_F[\min(v_1, v_2)] & \text{if } \alpha = 1. \end{cases} \tag{6}$$

This result is established by leveraging the weak convergence in conjunction with the regularity of the distributions F_n's. This result is a key step in linking the performance against elements of $\mathcal{W}$ to that against $\mathcal{F}$.

Proposition 2. *Fix a nonempty subset $\mathcal{F}$ of $\mathcal{F}_{reg}$ and a nonempty subset $\mathcal{W}'$ of $\mathcal{W}$. Suppose that for any element of $\mathcal{W}'$, there exists a sequence of distributions in $\mathcal{F}$ that weakly converges to that element. Then we have*

$$\mathcal{R}(\mathcal{M}'_{sf}, \mathcal{F}) \leq \mathcal{R}(\mathcal{M}'_{sf}, \mathcal{W}').$$

In other words, although $\mathcal{W}'$ is not a subset of $\mathcal{F}_{reg}$, the result states that the maximin ratio against the class of distributions $\mathcal{W}'$ upper bounds the maximin ratio against the class $\mathcal{F}_{reg}$. The proof of this result leverages the fact that we are working under the tractable space of mechanisms $\mathcal{M}'_{sf}$ in conjunction with the limits established in Lemma 3. Indeed, the worst-case performance of any mechanism in $\mathcal{M}'_{sf}$ against $\mathcal{F}$ is upper bounded by that against any element of a sequence F_n that converges weakly to an element F of $\mathcal{W}'$. In the proof, we characterize an asymptotic upper bound on the performance of any mechanism in $\mathcal{M}'_{sf}$ against F_n. Then, we establish that the asymptotic upper bound may be expressed as the performance of a new mechanism in $\mathcal{M}'_{sf}$ when facing the distribution corresponding to the weak limit F.

4.2.1. Subclass of Optimal Mechanisms Against $\mathcal{W}$. Next, we exploit the structure of the distributions in $\mathcal{W}$ to further simplify the maximin ratio against subclasses $\mathcal{W}'$ of $\mathcal{W}$, $\mathcal{R}(\mathcal{M}'_{sf}, \mathcal{W}')$. Let us introduce the following subset of mechanisms of $\mathcal{M}_{sf}$:

$$\mathcal{M}_{sf}^{max} = \Big\{m \in \mathcal{M}_{sf} : \text{for } i = 1, 2,\ x_i(v_i, v_{-i}) = \sum_{n=1}^{N} \frac{1}{N}\ \mathbb{1}\{v_i > \gamma_n v_{-i}\}\mathbb{1}\{v_i \neq v_{-i}\} + c\ \mathbb{1}\{v_i = v_{-i}\}, \text{ for some } N \geq 1, \gamma \in ([1, \infty))^N \text{ and } c \in [0, 1/2]\Big\}.$$

Note that $\mathcal{M}_{sf}^{max}$ is a subset of $\mathcal{M}'_{sf}$ and is the set of mechanisms in $\mathcal{M}'_{sf}$ that never allocate to the minimum value of buyers (when both values are different).

Proposition 3. *For any subset of distributions $\mathcal{W}'$ of $\mathcal{W}$,*

$$\mathcal{R}(\mathcal{M}'_{sf}, \mathcal{W}') = \mathcal{R}(\mathcal{M}_{sf}^{max}, \mathcal{W}').$$

This proposition shows that, without loss of optimality, when facing distributions in $\mathcal{W}$, one can focus on mechanisms that never allocate to the minimum value (if the latter is different from the maximum value). The intuition behind the result is that under the class of distributions $\mathcal{W}$, the seller would like to set a reserve price equal to the upper bound of the support if she would know the distribution (see Lemma B-4 in the electronic companion). In addition, allocating to a buyer with value strictly below this reserve price yields a negative contribution to the revenue of the seller (cf. Myerson 1981). When the seller sees two values, although she does not know the distribution, she knows that it belongs to $\mathcal{W}$, and hence she still knows that both values are weakly below the optimal oracle reserve price. In turn, the seller never wants to allocate to the minimum value (if it is different from the maximum value).

We are now ready to put together all earlier results and state the main result of this section.

Theorem 2. *Fix a nonempty scale-invariant subset $\mathcal{F}$ of $\mathcal{F}_{reg}$ and a nonempty subset $\mathcal{W}'$ of $\mathcal{W}$. Suppose that for any element of $\mathcal{W}'$, there exists a sequence of distributions in $\mathcal{F}$ that weakly converges to that element. Then we have*

$$\mathcal{R}(\mathcal{M}, \mathcal{F}) \leq \mathcal{R}(\mathcal{M}_{sf}^{max}, \mathcal{W}').$$

This result provides a family of upper bounds on the maximin ratio associated with any subset of the set of regular distributions and in particular applies to $\mathcal{F}_{mhr}$ and $\mathcal{F}_{reg}$. In Section 5, we apply this upper bound to $\mathcal{F} = \mathcal{F}_{reg}$, and in Section 6, we apply it to $\mathcal{F} = \mathcal{F}_{mhr}$, where for each, we select a suitable set $\mathcal{W}'$.

5. Maximin Ratio for Regular Distributions

In this section, we develop upper and lower bounds on $\mathcal{R}(\mathcal{M}, \mathcal{F}_{reg})$, leading to a narrow interval to which $\mathcal{R}(\mathcal{M}, \mathcal{F}_{reg})$ belongs.

5.1. Upper Bound

Theorem 3 (Upper Bound for Regular Distributions). *The maximin ratio $\mathcal{R}(\mathcal{M}, \mathcal{F}_{reg})$ is upper bounded as follows:*

$$\mathcal{R}(\mathcal{M}, \mathcal{F}_{reg}) \leq \sup_{N\geq 1} \sup_{\gamma\in[1,+\infty)^N} \inf_{q\in(0,1)} \frac{N-|\mathcal{I}^+|}{N}\frac{1}{2-q} + \frac{|\mathcal{I}^+|}{N}\frac{q}{2-q} + \sum_{k\in\mathcal{I}^+}\frac{1}{N}\overline{\psi}(\gamma_k, q),$$

where $\mathcal{I}^+ = \{k \in [1, N] : \gamma_k > 1\}$ and

$$\overline{\psi}(\gamma_k, q) := 2\frac{\gamma_k}{\gamma_k - 1}\frac{1}{1-q}\frac{1}{2-q}\left[\frac{1-q}{1-q+\gamma_k q} - \frac{1}{\gamma_k - 1}\ln\left(\frac{\gamma_k}{1-q+\gamma_k q}\right)\right].$$

Theorem 3 provides a fundamental limit on the performance of *any* mechanism in $\mathcal{M}$. At a high level, the upper bound captures the complexity of the space of mechanisms through a vector $\gamma \in [1, +\infty)^N$, and the space of distributions has been distilled down to a scalar $q \in [0, 1]$. This is in stark contrast with the initial space of mechanisms $\mathcal{M}$ and the space of regular distributions. The sharpness of this upper bound will be apparent in the coming subsections, when we evaluate it and compare it to a lower bound.

The upper bound in Theorem 3 also explicitly highlights the tension associated with the design of a prior-free mechanism. On the one hand, one may want to put weight on values $\gamma_k = 1$ to guarantee performance in line with a second price auction, which hedges against deterministic values. This corresponds to the first term in the upper bound (i.e., $1/(2-q)$). On the other hand, putting weight on terms $\gamma_k > 1$ may yield higher performance if nature selects a distribution with a heavy tail.

5.1.1. Key Ideas Underlying the Proof of Theorem 3. The first step in the proof is to derive an upper bound on $\mathcal{R}(\mathcal{M}, \mathcal{F}_{reg})$ through Theorem 2. Given the latter, the key then is to identify an appropriate subset of distribution $\mathcal{W}_{reg}$ that verifies the conditions of Theorem 2, and the rest of the proof is organized around identifying such a subset and explicitly deriving an upper bound on the worst-case performance of any mechanism in $\mathcal{M}_{sf}^{max}$ against $\mathcal{W}_{reg}$.

The family of distributions for which the revenue curve in the quantile space is a triangle has the following expression:

$$F_a(v) = \begin{cases} 1 - \frac{1}{v+1}, & \text{if } v < a, \\ 1, & \text{if } v \geq a, \end{cases} \tag{7}$$

for some $a \geq 0$, and it has received attention in the literature in various contexts. If we introduce the following class of distribution $\mathcal{W}_{reg} := \{F_a : a > 0\}$, then one can show that each element in this class of distribution $\mathcal{W}_{reg}$ can be approached by a sequence of elements of $\mathcal{F}_{reg}$ (see Lemma C-1 in the electronic companion). As a result, $\mathcal{R}(\mathcal{M}_{sf}^{max}, \mathcal{W}_{reg})$ is a valid upper bound for $\mathcal{R}(\mathcal{M}, \mathcal{F}_{reg})$. The proof then relies on deriving an analytical expression for $\mathcal{R}(\mathcal{M}_{sf}^{max}, \mathcal{W}_{reg})$.

5.2. Lower Bound

We have just established an upper bound on $\mathcal{R}(\mathcal{M}, \mathcal{F}_{reg})$. We next focus on deriving a lower bound.

Proposition 4 (Lower Bound for Regular Distributions). *Consider any mechanism $m = (\boldsymbol{x}, \boldsymbol{t})$ in $\mathcal{M}_{sf}^{max}$ and the corresponding parameters $N \geq 1$, $\gamma \in [1, \infty)^N$, and $c \in [0, 1/2]$. Let $\mathcal{I}^+ = \{k \in [1, N] : \gamma_k > 1\}$. If $|\mathcal{I}^+|/N \leq 1/3$, then the performance of such a mechanism in the presence of two buyers against a distribution F with optimal quantile q_F is lower bounded as follows:*

$$R(m, F) \geq \frac{N-|\mathcal{I}^+|}{N}\frac{1}{2-q_F} + \sum_{k\in\mathcal{I}^+}\frac{1}{N}\underline{\psi}(\gamma_k, q_F),$$

where

$$\underline{\psi}(\gamma_k, q_F) := \frac{\gamma_k}{\gamma_k - 1}\left(1 - q_F - \frac{1}{\gamma_k - 1}\ln\left[\frac{\gamma_k}{1 + (\gamma_k - 1)q_F}\right]\right)\frac{1}{1-q_F} - \frac{2\,\gamma_k\, q_F}{1 + (\gamma_k - 1)q_F}\frac{1}{2-q_F}.$$

The proposition above gives an explicit lower bound for any mechanism in $\mathcal{M}_{sf}^{max}$ that satisfies $|\mathcal{I}^+|/N \leq 1/3$ (i.e., which does not inflate the second price more than a third of the time). In particular, the lower bound admits the same structure as the function characterizing the upper bound up to a correction factor. In particular, it is possible to see that the difference between the upper and lower bounds goes to 0 as q approaches 0.

5.2.1. Comparison with the Lower Bound Obtained in Fu et al. (2015). The authors study a mechanism that randomizes between a second price auction and an inflation factor of γ, which can be viewed as a special instance of the mechanisms in $\mathcal{M}_{sf}^{max}$.

For $\gamma = 2$, and using the second price auction with probability $1 - p$ and inflation γ with probability p, one may establish that the lower bound obtained in Proposition 4 is tighter and higher by a factor of

$$p\frac{2\ q_F^2\ (1-q_F)}{(1+q_F)\ (2-q_F)}.$$

The key drivers of the improvement are dual. A first improvement stems from bounding in a dependent fashion the contributions of the second price auction ($\gamma_k = 1$) and that of the inflation mechanisms ($\gamma_k > 1$). A second improvement stems from obtaining a tighter bound on the contributions of high γ_k terms.

5.3. Characterization of $\mathcal{R}(\mathcal{M}, \mathcal{F}_{reg})$

We next evaluate numerically values for upper and lower bounds on $\mathcal{R}(\mathcal{M}, \mathcal{F}_{reg})$. Using Theorem 3, we derive an upper bound on the maximin ratio. To that end, we fix $q = 0.17$. For such a value, we have $1/(2-q) = 54.64\%$. Furthermore, the function $\zeta : (1, +\infty) \to \mathbb{R}$ defined by

$$\zeta(\gamma) = 2\ \gamma\ \frac{1}{1-q}\ \frac{1}{2-q}\left[\frac{1}{\gamma-1}\ \frac{1-q}{1-q+\gamma q} - \frac{1}{(\gamma-1)^2}\ \ln\left[\frac{\gamma}{1+(\gamma-1)q}\right]\right] + \frac{q}{2-q}$$

reaches its maximum at about $\gamma = 1.5$, and its maximal value is 55.59%. From the above, we deduce that maximin ratio is upper bounded by 55.59%.

Applying Proposition 4, we evaluate numerically the lower bound by taking $\gamma = (1, 1, 1, 1, 2)$ and a vector q of values from 0 to 1 with a step 0.001. We find that the lower bound is 51.9%. We conclude that

$$51.9\% \le \mathcal{R}(\mathcal{M}, \mathcal{F}_{reg}) \le 55.59\%.$$

In other words, we have characterized the maximin ratio up to less than 4%. There is an important implication of the results above. In the face of regular distributions, although randomization is helpful compared with a second price auction (that guarantees 50% of oracle revenues), the extent to which one may improve performance is limited to at most 5.59%. An interpretation of our results is that the second price auction is near optimal in environments with unknown regular distributions.

6. Maximin Ratio for MHR Distributions

In this section, we focus on the maximin ratio when nature can only select distributions in $\mathcal{F}_{mhr}$, which is a subset of the regular class of distributions $\mathcal{F}_{reg}$. In other words, the seller now has more information about the distribution of buyers compared with the setting analyzed in Section 5.

6.1. Upper Bound

Theorem 4 (Upper Bound for MHR Distributions). *The maximin ratio $\mathcal{R}(\mathcal{M}, \mathcal{F}_{mhr})$ is upper bounded as follows:*

$$\mathcal{R}(\mathcal{M}, \mathcal{F}_{mhr}) \le \inf_{q \in [e^{-1}, 1]} \frac{1-q^2}{2\ q\ (2-q)\ \ln(1/q)}.$$

Theorem 4 provides a fundamental limit on the performance of any mechanism against distributions in $\mathcal{F}_{mhr}$. Quite notably, this upper bound comes in quasi-closed form and takes a significantly much simpler form than for the broader class of regular distributions. We next highlight the main ideas in the proof and highlight the role of the MHR knowledge in the derivation of this upper bound.

The proof of this result follows initially the same structure as that of Theorem 3. As earlier, we leverage Theorem 2, but now, we use a different family, $\mathcal{W}_{mhr}$, suited to the increasing hazard rate family of distributions $\mathcal{F}_{mhr}$. In particular, we define $\mathcal{W}_{mhr}$ to be the set of distribution F parametrized by $a \ge b > 0$ such that

$$F_{a,b}(v) = \begin{cases} 1 - \exp\left(-\frac{v}{a}\right), & \text{if } v < b, \\ 1, & \text{if } v \ge b. \end{cases}$$

This family is constructed by truncating the exponential family distribution. This family is rich enough to cover the range of all possible optimal oracle quantiles (q_F) of MHR distributions. We establish that any such element can be "approached" by a sequence in $\mathcal{F}_{mhr}$, and in turn, $\mathcal{R}(\mathcal{M}_{sf}^{max}, \mathcal{W}_{mhr})$ is an upper bound on $\mathcal{R}(\mathcal{M}, \mathcal{F}_{mhr})$. The role of the MHR assumption comes into play when we evaluate the performance of any mechanism in $\mathcal{M}_{sf}^{max}$ against $\mathcal{W}_{mhr}$. In this context, we are able to establish that the optimal performance against $\mathcal{W}_{mhr}$ is given by that of a second price auction. In particular, it is suboptimal to randomize the allocation when facing the family $\mathcal{W}_{mhr}$.

6.2. Lower Bound

We next establish a lower bound on $\mathcal{R}(\mathcal{M}, \mathcal{F}_{mhr})$ by lower bounding the performance of a second price auction.

It is worthwhile to note that a first coarse lower bound may be readily obtained from existing results by simply noting that the oracle optimal quantile q_F cannot be less than e^{-1} for MHR distributions; see, for example, Hartline et al. (2008). Combining this with the lower bound on the performance of a second price auction of $1/(2-q_F)$ obtained in Fu et al. (2015) for regular distributions, one readily obtains that

$$R(m_{spa}, F) \ge \frac{1}{2-e^{-1}} \approx 61.2\%.$$

One can already see that a significantly higher performance is possible with the additional knowledge that the distributions belong to the MHR class. Next, we establish a sharp lower bound on $R(m_{spa}, F)$.

Proposition 5 (Lower Bound for MHR Distributions). *For any F in $\mathcal{F}_{mhr}$, the performance of the second price auction in the presence of two buyers is bounded below as follows:*

$$R(m_{spa}, F) \geq \frac{1}{2} \frac{1 - q_F^2}{q_F(2 - q_F)(-\ln(q_F))},$$

where $q_F = 1 - F(r_F)$ is the oracle optimal quantile.

The key idea underlying this result is to leverage the structural properties that the MHR distribution imposes on the structure of the revenue curve in the quantile space. In particular, leveraging a single crossing property between the ccdf of *any* MHR distribution and any exponential tail developed in the reliability theory literature (Barlow and Proschan 1975), we establish a lower bound on the ccdf of the distribution of any MHR distribution through that of a particular exponential distribution. This leads to a lower bound on the revenue curve in the quantile space, ultimately leading to the bound above.

We discuss the implications of this result next.

6.3. Optimality of Second Price Auction and Characterization of $\mathcal{R}(\mathcal{M}, \mathcal{F}_{mhr})$

We are now ready to state the main result of Section 6, which follows from the two earlier results.

Theorem 5 (Optimality of Second Price Auction). *The second price auction is optimal in $\mathcal{M}$ when facing two buyers with MHR distributions—namely,*

$$\inf_{F \in \mathcal{F}_{mhr}} R(m_{spa}, F) = \mathcal{R}(\mathcal{M}, \mathcal{F}_{mhr}).$$

Furthermore,

$$\mathcal{R}(\mathcal{M}, \mathcal{F}_{mhr}) = \inf_{q \in [e^{-1}, 1]} \frac{1 - q^2}{2\, q\, (2 - q)\, \ln(1/q)} \approx 71.53\%.$$

We conclude that the second price auction and an element in $\mathcal{W}_{mhr}$ represent a (quasi) saddle point for the maximin ratio $\mathcal{R}(\mathcal{M}, \mathcal{F}_{mhr})$.[9] Theorem 5 provides an exact characterization of the maximin ratio and the corresponding optimal prior-free auction.

Interestingly, whereas randomization of the allocation helped the seller counter nature when facing regular distributions, such randomization does not help anymore when facing the subclass of monotone hazard rate distributions. It is quite notable that this simple mechanism, a second price auction, which is also efficient, is actually optimal in this environment.

The result above also quantifies the value of additional knowledge of the distributions. If a seller knows that the distribution is MHR, then she gains at least 71.53% − 55.59% = 15.94% in guaranteed performance (compared with an oracle). Indeed, MHR distributions have limited variability as measured (e.g., through the coefficient of variation). The latter is bounded by 1 (Barlow and Proschan 1975), whereas it is unbounded for regular distributions. With such limited variability, a second price auction appears "sufficient."

7. Extensions and Concluding Remarks

We have analyzed the problem of *optimally* selling one indivisible good to two symmetric and independent buyers when one relaxes the common prior assumption. For that, we look at the model where the buyers are not assumed to know any information about the other buyers and the seller does not know the exact distribution. We characterize the maximin ratio for a broad subclass of DSIC mechanisms against the classes of regular and MHR distributions. We refer the reader back to Table 1 for a summary of some implications of our results. Whereas we have done so while focusing on the case of two buyers, we establish next that the bounds we have derived apply to the case when the number of buyers is selected adversarially.

7.1. Extension to the Case of an Adversarially Selected Number of Buyers

In this section, we will show that our bounds apply to the case in which the number of buyers is arbitrary but adversarially selected.

We assume as earlier that the seller does not know exactly the distribution of values of buyers but knows it belongs to some class of distribution $\mathcal{F}$ in the general class of distributions $\mathcal{G}$. Moreover, we assume also that the seller does not know the exact the number of buyers $K \geq 2$. We model the seller's problem as a game between the seller and nature, where the seller will first pick a collection of prior-independent mechanisms contingent on the number of bids K, and then nature picks both the number of buyers and their distribution of values from some class.

A seller's mechanism is now a set of allocations and payment functions contingent on the number of bids received $K \geq 2$. The seller will apply a mechanism characterized by an allocation mapping $\boldsymbol{x}^K$ and a payment mapping $\boldsymbol{t}^K$, where $\boldsymbol{x}^K : \mathbb{R}^K \to [0,1]^K$ and $\boldsymbol{t}^K : \mathbb{R}^K \to \mathbb{R}^K$. We focus on DSIC mechanisms that verify for any $K \geq 2$

$$v_i\, x_i^K(v_i, \boldsymbol{v}_{-i}) - t_i^K(v_i, \boldsymbol{v}_{-i}) \geq 0, \quad \text{for all } i \text{ and } v_i, \boldsymbol{v}_{-i} \text{ in } \mathbb{R}_+^K, \tag{IR-K}$$

$$v_i\, x_i^K(v_i, \boldsymbol{v}_{-i}) - t_i^K(v_i, \boldsymbol{v}_{-i}) \geq v_i\, x_i^K(\hat{v}_i, \boldsymbol{v}_{-i}) - t_i^K(\hat{v}_i, \boldsymbol{v}_{-i}), \quad \text{for all } i \text{ and } v_i,\ \boldsymbol{v}_{-i},\ \hat{v}_i \text{ in } \mathbb{R}_+^{K+2}, \tag{IC-K}$$

$$\boldsymbol{x}^K(v_i, \boldsymbol{v}_{-i}) \text{ belongs to } \Delta^K, \quad \text{for all } v_i, \boldsymbol{v}_{-i} \text{ in } \mathbb{R}_+^K, \tag{AC-K}$$

where Δ^K is the probability simplex of $\mathbb{R}^K$. These constraints are similar to those introduced earlier (see (IR), (IC). and (AC)).

More formally, the seller's strategy is a collection of mechanisms from the set $\widetilde{\mathcal{M}}$, where

$$\widetilde{\mathcal{M}} = \Big\{(\boldsymbol{x}^K, \boldsymbol{t}^K)_{K\geq 2}\colon (\boldsymbol{x}^K, \boldsymbol{t}^K) \text{ satisfies (IR-K), (IC-K), (AC-K) and } \max_{i=1,2}\{TV(x_i^2, [0,\varepsilon]^2)\} < \infty \text{ for some } \varepsilon > 0\Big\}.$$

Similarly, we define the oracle benchmark, as well as the performance of each mechanism contingent on having K buyers, by

$$\text{opt}^K(F) := \sup_{m\in\mathcal{M}} \mathbb{E}_F\left[\sum_{i=1}^K t_i^K(v_i, \boldsymbol{v}_{-i})\right] \text{ and}$$
$$R^K(m,F) = \frac{\mathbb{E}_F\left[\sum_{i=1}^K t_i^K(v_i, \boldsymbol{v}_{-i})\right]}{\text{opt}^K(F)}.$$

In the case of an arbitrary but adversarially selected number of buyers, the objective of the seller is now given by

$$\widetilde{\mathcal{R}}(\widetilde{\mathcal{M}}, \mathcal{F}) = \sup_{m\in\widetilde{\mathcal{M}}} \inf_{K\geq 2} \inf_{F\in\mathcal{F}} R^K(m,F).$$

Next, we state the main result of this section.

Proposition 6. (1) *The maximin ratio for the regular class of distributions verifies*

$$51.9\% \leq \widetilde{\mathcal{R}}(\widetilde{\mathcal{M}}, \mathcal{F}_{ref}) \leq 55.6\%.$$

(2) *A second price auction is an optimal mechanism when facing MHR distributions. Furthermore,*

$$\widetilde{\mathcal{R}}(\widetilde{\mathcal{M}}, \mathcal{F}_{mhr}) = \inf_{F\in\mathcal{F}_{mhr}} R^2(m_{spa}, F) \approx 71.53\%.$$

Note that a priori it is not clear that the case of two buyers is the worst case, because the oracle benchmark also varies with the number of buyers. In the proof of Proposition 6, we show formally that the smallest maximin ratio is achieved when only two buyers participate in the auction. The proof of these results relies fundamentally on the case of two buyers studied earlier in the paper in conjunction with some known results in the literature. Hence, when nature can pick adversarially any number of buyers $K \geq 2$, a second price auction is still near optimal against regular distributions and is actually optimal against MHR distributions.

7.2. Future Directions

A direct and complementary direction would be to characterize the maximin ratio as a function of the number of buyers, when this number cannot be selected adversarially by nature.

Also, in our analysis we have mainly focused on the regular and MHR classes, which are subsets of the α-strongly regular class of distributions; see Ewerhart (2013), Cole and Roughgarden (2014), Cole and Rao (2015), and Schweizer and Szech (2016) for more details about this class of distributions. The results developed in Section 4 for the upper bounds have the potential to be applied to any subclass of the α-strongly regular class where α would be a parameter that would capture the degree of knowledge of the seller. As such, an interesting direction is to characterize the maximin ratio as a function of the degree of knowledge of the seller.

Another way to incorporate the knowledge of the seller is to assume that she has access to extra information such as the moments, and a potential research question is how one could leverage such additional information to improve the performance and what the structure of optimal mechanisms is in such cases. (We refer the reader to, e.g., Azar et al. (2013), who study deterministic mechanisms that incorporate such information.)

More generally, our work tries to relax the common prior assumption and we have focused on the canonical setting of one indivisible good and symmetric buyers with independent values that are regular. There are various generalizations that naturally emerge. For example, it would be interesting to see if one can develop results of a similar nature when the class of distribution is a "structured" irregular class (see, e.g., Sivan and Syrgkanis 2013 for an example of such a subclass). Similarly, developing parallel lower and upper bounds on the maximin ratio for general environments that would allow, for example, for correlation among values or asymmetric buyers is a promising direction.

Acknowledgments

The authors are grateful to the review team for the feedback that allowed to improve the initial version of the paper.

Appendix. Auxiliary Definition

Here, we recall the definition of the Arzelà total variation definition for functions of two variables; see, for example, Clarkson and Adams (1933).

Definition A.1. *The Arzelà total variation of a mapping* $h : [a,b]\times[c,d] \to \mathbb{R}$ *is given by*

$$TV(h, [a,b]\times[c,d]) := \sup_{N\geq 1} \sup_{\substack{u\in[a,b]^N \\ u_1\leq\ldots\leq u_N}} \sup_{\substack{v\in[c,d]^N \\ v_1\leq\ldots\leq v_N}} \sum_{j=1}^N \left|h(u_{j+1}, v_{j+1}) - h(u_j, v_j)\right|.$$

Furthermore, h is said to have finite total variation in $[a,b]\times[c,d]$ *if* $TV(h,[a,b]\times[c,d]) < \infty$.

Endnotes

[1] We initially focus on two buyers and then return to the case of an arbitrary and adversarially selected number of buyers in Section 7.1.

[2] Here and throughout the paper, whenever we refer to a "second price auction," unless otherwise noted, it is implicitly assumed that there is no reserve price.

[3] It is worthwhile noting here that against the classes we consider (regular and mhr), a worst-case absolute performance analysis would lead to a value of 0, and all feasible mechanisms would be optimal. The relative benchmark approach allows for controlling the environment and deriving guarantees on broader sets of distributions.

[4] We also refer the reader to Chung and Ely (2007) for an in-depth discussion of DSIC mechanisms.

[5] We recall the definition of Arzelà total variation in the appendix.

[6] Although it is needed for the proofs, we conjecture that it does not imply a loss of optimality.

[7] In recent years, there has been growing literature advocating for simple mechanisms (see, e.g., Hartline and Roughgarden 2009 and Daskalakis and Pierrakos 2011). In that sense, the mechanisms in $\mathcal{M}$ could be thought as a formalization of some broad class of "simple" mechanisms.

[8] Note also that this set captures explicitly the probability of allocation to a buyer when the value of buyers are equal. Although seemingly unimportant in the class $\mathcal{F}_{reg}$, because ties happen with probability 0, this explicit inclusion of the case of ties will play an important role when we will be dealing with the limiting performance of a mechanism against an appropriate sequence of distributions that converges weakly to a point outside of $\mathcal{F}_{reg}$ (see Proposition 3).

[9] Technically speaking, it is not exactly a saddle point given that the elements of $\mathcal{W}_{mhr}$ do not belong to $\mathcal{F}_{mhr}$.

References

Araman VF, Caldentey RA (2009) Dynamic pricing for nonperishable products with demand learning. *Oper. Res.* 57(5):1169–1188.

Azar P, Micali S (2012) Optimal parametric auctions. CSAIL Technical Report MIT-CSAIL-TR-2012-015, Massachusetts Institute of Technology, Cambridge.

Azar P, Daskalakis C, Micali S, Weinberg SM (2013) Optimal and efficient parametric auctions. Khanna S, ed. *Proc. 24th Annual ACM-SIAM Sympos. Discrete Algorithms* (ACM, New York), 596–604.

Bandi C, Bertsimas D (2014) Optimal design for multi-item auctions: A robust optimization approach. *Math. Oper. Res.* 39(4):1012–1038.

Barlow RE, Proschan F (1975) *Statistical theory of reliability and life testing: Probability models*. Technical report, Florida State University, Tallahassee.

Bergemann D, Morris S (2013) An introduction to robust mechanism design. *Foundations Trends Microeconom.* 8(3):169–230.

Bergemann D, Schlag KH (2008) Pricing without priors. *J. Eur. Econom. Assoc.* 6(2–3):560–569.

Bergemann D, Brooks, BA, Morris S (2016) Informationally robust optimal auction design. Cowles Foundation Discussion Paper 2065, Yale University, New Haven, CT.

Besbes O, Zeevi A (2009) Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Oper. Res.* 57(6):1407–1420.

Besbes O, Zeevi A (2015) On the (surprising) sufficiency of linear models for dynamic pricing with demand learning. *Management Sci.* 61(4):723–739.

Borodin A, El-Yaniv R (1998) *Online Computation and Competitive Analysis* (Cambridge University Press, New York).

Bose S, Daripa A (2009) A dynamic mechanism and surplus extraction under ambiguity. *J. Econom. Theory* 144(5):2084–2114.

Bose S, Ozdenoren E, Pape A (2006) Optimal auctions with ambiguity. *Theoret. Econom.* 1(4):411–438.

Bulow J, Klemperer P (1996) Auctions vs. negotiations. *Amer. Econom. Rev.* 86(1):180–194.

Caillaud B, Robert J (2005) Implementation of the revenue-maximizing auction by an ignorant seller. *Rev. Econom. Design* 9(2):127–143.

Caldentey R, Liu Y, Lobel I (2016) Intertemporal pricing under minimax regret. *Oper. Res.* 65(1):104–129.

Carrasco V, Farinha Luz V, Monteiro P, Moreira H (2015) Robust selling mechanisms. Discussion Paper 641, Pontifical Catholic University of Rio de Janeiro, Rio de Janeiro.

Carroll G (2019) Robustness in mechanism design and contracting. *Annual Rev. Econom.* 11(1):139–166.

Chen H, Hu M, Perakis G (2017) Distribution-free pricing. Working paper, Chinese Academy of Sciences, Beijing.

Chiesa A, Micali S, Zhu ZA (2015) Knightian analysis of the Vickrey mechanism. *Econometrica* 83(5):1727–1754.

Chung K-S, Ely JC (2007) Foundations of dominant-strategy mechanisms. *Rev. Econom. Stud.* 74(2):447–476.

Clarkson JA, Adams CR (1933) On definitions of bounded variation for functions of two variables. *Trans. Amer. Math. Soc.* 35(4):824–854.

Cohen MC, Perakis G, Pindyck RS (2016) Pricing with limited knowledge of demand. *Proc. 2016 ACM Conf. Econom. Comput.* (ACM, New York), 657.

Cole R, Rao S (2015) Applications of α-strongly regular distributions to Bayesian auctions. Markakis E, Schäfer G, eds. *Proc. 11th Internat. Conf. Web Internet Econom.* (Springer, Berlin), 244–257.

Cole R, Roughgarden T (2014) The sample complexity of revenue maximization. *Proc. 46th Annual ACM Sympos. Theory Comput.* (ACM, New York), 243–252.

Daskalakis C, Pierrakos G (2011) Simple, optimal and efficient auctions. Chen N, Elkind E, Koutsoupias E, eds. *Proc. 7th Internat. Workshop Internet Network Econom.* (Springer, Berlin), 109–121.

den Boer AV (2015) Dynamic pricing and learning: Historical origins, current research, and new directions. *Surveys Oper. Res. Management Sci.* 20(1):1–18.

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

Eren SS, Maglaras C (2010) Monopoly pricing with limited demand information. *J. Revenue Pricing Management* 9(1–2):23–48.

Ewerhart C (2013) Regular type distributions in mechanism design and ρ-concavity. *Econom. Theory* 53(3):591–603.

Farias VF, Van Roy B (2010) Dynamic pricing with a prior on market response. *Oper. Res.* 58(1):16–29.

Fu H, Immorlica N, Lucier B, Strack P (2015) Randomization beats second price as a prior-independent auction. *Proc. 16th ACM Conf. Econom. Comput.* (ACM, New York), 323.

Goldberg AV, Hartline JD, Karlin AR, Saks M, Wright A (2006) Competitive auctions. *Games Econom. Behav.* 55(2):242–269.

Harstad RM, Kagel JH, Levin D (1990) Equilibrium bid functions for auctions with an uncertain number of bidders. *Econom. Lett.* 33(1):35–40.

Hartline JD (2013) Bayesian mechanism design. *Foundations Trends® Theoret. Comput. Sci.* 8(3):143–263.

Hartline JD, Roughgarden T (2009) Simple vs. optimal mechanisms. *Proc. 10th ACM Conf. Electronic Commerce* (ACM, New York), 225–234.

Hartline J, Mirrokni V, Sundararajan M (2008) Optimal marketing strategies over social networks. *Proc. 17th Internat. Conf. World Wide Web* (ACM, New York), 189–198.

Huang Z, Mansour Y, Roughgarden T (2015) Making the most of your samples. *Proc. 16th ACM Conf. Econom. Comput.* (ACM, New York), 45–60.

Keskin NB, Zeevi A (2014) Dynamic pricing with an unknown demand model: Asymptotically optimal semi-myopic policies. *Oper. Res.* 62(5):1142–1167.

Kleinberg RD, Leighton T (2003) The value of knowing a demand curve: Bounds on regret for online posted-price auctions. *Proc. 44th Annual IEEE Sympos. Foundations Comput. Sci.* (IEEE, Piscataway, NJ), 594–605.

Koçyiğit Ç, Iyengar G, Kuhn D, Wiesemann W (2020) Distributionally robust mechanism design. *Management Sci.* 66(1):159–189.

Krishna V (2009) *Auction Theory* (Academic Press, New York).

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

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

Myerson RB (1981) Optimal auction design. *Math. Oper. Res.* 6(1): 58–73.

Neeman Z (2003) The effectiveness of English auctions. *Games Econom. Behav.* 43(2):214–238.

Schweizer N, Szech N (2016) The quantitative view of Myerson regularity. Working paper, Tilburg School of Economics, Tilburg, Netherlands.

Sivan B, Syrgkanis V (2013) Vickrey auctions for irregular distributions. Chen Y, Immorlica N, eds. *Internat. Conf. Web Internet Econom.* (Springer, Berlin), 422–435.

Talluri KT, Van Ryzin GJ (2006) *The Theory and Practice of Revenue Management*, International Series in Operations Research and Management Science, vol. 68 (Springer Science & Business Media, New York).

Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* 16(1):8–37.

Wang Z, Deng S, Ye Y (2014) Close the gaps: A learning-while-doing algorithm for single-product revenue management problems. *Oper. Res.* 62(2):318–331.

Wilson R (1987) Game theoretic analyses of trading processes. Bewley TF, ed. *Adv. Econom. Theory: Fifth World Congress* (Cambridge University Press, Cambridge, UK), 33–70.