

Robust Bidding in First-Price Auctions: How to Bid without Knowing what Others are Doing*

Bernhard Kasberger* Karl H. Schlag[◇]

March 27, 2023

Abstract

We propose how to bid in first-price auctions when a bidder knows her value but not how others will bid. To do this, we introduce a methodology to show how to make choices in strategic settings without assuming common knowledge or equilibrium behavior. Accordingly, we first eliminate environments that are believed not to occur and then find a robust rule that performs well in the remaining environments. We test our bids using data from laboratory experiments and the field and find that our bids outperform those made by real bidders.

JEL Classification: D44, D81, C57.

Keywords: robust decision-making, first-price auction, minimax loss, empirical loss, non-equilibrium model.

1 Introduction

In this paper, we specify how to bid in real-life first-price sealed-bid auctions. To understand current bidding practices, we conduct a survey among 44 expert academics and practitioners involved in auctions. This survey reveals that classic

*We thank Paul Heidhues, Fabian Herweg, Jacob Goeree, Paul Klemperer, Itzhak Rasooly, József Sákovics, Kyle Woodward and the participants at the 2016 Auction and Market Design Workshop in Vienna, the 2017 ZEW Workshop on Market Design, the 2017 Stony Brook International Conference on Game Theory, the 2017 EEA-ESEM meeting in Lisbon, the 2018 NOeG annual meeting, the 2018 EARIE conference, the Nuffield Economic Theory Workshop, the 2019 VfS annual conference, the 2019 INFORMS annual meeting, and the 2019 NOeG winter meeting for their comments.

*Düsseldorf Institute for Competition Economics (DICE), Heinrich Heine University Düsseldorf, bernhard.kasberger@hhu.de

[◇]Department of Economics, University of Vienna, karl.schlag@univie.ac.at

game-theoretic analysis provides minimal guidance. One obstacle is that the literature typically assumes that the setting is common knowledge and that everyone is bidding in equilibrium. A further issue is that the theory typically quantifies uncertainty using probabilities, an approach that is hardly applicable in real-life settings.

Uncertainty is at the heart of auctions; for example, it is not clear how other bidders value the good, what they might bid, and what information and budgets they have at their disposal. In fact, it may not even be clear how many other bidders there are.

The way we deal with uncertainty is based on the understanding that it can be easier to focus on what *will not* happen when one is uncertain about what *will* happen. Based on this principle, we rule out as much as possible and then find a bid that performs best in a well-defined sense, without committing to what will occur. Note that we need a pragmatic approach to implement this principle, as it is not feasible or sensible to confront the bidder with each possible environment and ask whether or not it is conceivable to the bidder. Therefore, we introduce a methodology to show how to describe certain features of the environment to then ask whether the bidder can rule out environments with these features. As a result, all environments that have the same features are dealt with simultaneously.

The best way to describe the features of the environment depends on the mindset or perspective of the bidder and on the bidder's information. Closest to the classic approach, one might first consider the values of others, however, some survey respondents have noted that it is easier to think about the bids of others. In fact, it is simplest to only focus on the maximal bid of the other bidders. This motivates three different approaches.

The first is referred to as the “outcome-based” approach, which focuses on the maximal bid of the other bidders without incorporating how it results from individual bidding. The aim is to rule out features of the maximal bid, for instance, by specifying a lower bound under which the bid is unlikely to fall. The second is called the “choice-based” approach, which is quantified in terms of the bids of the others, without asking why these bids are made. This approach allows us to include bounds on the degree of independence between bidders, as well as on the variation of the bids. Historical bidding data can be incorporated when modeling the variation. Information on individual bidding is then aggregated so as to gather insights on the maximal bid. The “intention-based” approach is the third approach, which connects the values of others to their bids. This yields restrictions on which individual bids are expected. Only this approach allows us to reason about the val-

ues of others, such as conjecturing independent private values or affiliated values. A bid may be ruled out, for instance, because a bidder is not expected to bid above her value or because bidders will not bid too far below their own value. Information on values and bidding strategies is used to deduce information on the maximal bid of the other bidders.

Each of these approaches uses a different way to describe an environment. Common among these approaches is the objective to rule out the environments that are believed not to occur. The set of all environments that cannot be ruled out is called the *ambiguity set*. The next step is to choose a bid, given the ambiguity set. Specifically, for a given ambiguity set, we search for a bid that achieves a payoff that never falls too far below the payoff that could be achieved if more were known. Such a bid comes with a performance guarantee. Formally, we implement the minimax loss criterion (Savage, 1951).

In practice, it may not be straightforward to construct the ambiguity set as one faces the following dilemma. On the one hand, we wish to rule out as few environments as possible; on the other hand, the performance (guarantee) of the bid improves as the ambiguity set shrinks. To help resolve this dilemma, we propose the following iterative process. We start with the ambiguity set that contains all environments. Given the ambiguity set, we then compute the bid with the best performance guarantee and identify the worst-case environments. These are the environments that prevent the guarantee from being better. Next, we consider whether these worst-case environments can be ruled out: if yes, then we construct a new ambiguity set and continue as above; if none can be ruled out, then the process stops. Sometimes we may add an additional step to investigate how the final bid performs when additional features are specified. We do this when it is not clear how to quantify these features. For instance, it may be clear that there is a lower bound on some aspects, but it may be hard to assess where this lower bound is.

We illustrate our general methodology by presenting a specific implementation for each of the three approaches. In these illustrations, we incorporate only minimal information about the other bidders. Our implementations lead to three bids, which we later assess empirically. In the outcome-based approach, we find a bid that depends on very few assumptions. In the choice-based approach, we create a model that assumes some independence and variation in opponent bidding. In the intention-based approach, we assume independent values and put bounds on the value distribution and bidding functions. Each of the resulting bids come with a few free parameters, that are either determined by the known properties of the auction (such as the reserve price and the number of bidders) or are set equal to

salient values.

We make some comments on these implementations. The outcome-based approach makes no assumptions on how the maximal bid of the other bidders is generated. Hence, it allows for heterogeneous bidding of any form; independent or correlated. It is simple and straightforward, as everything is formulated in terms of the maximal bid, and thus it is easy to justify in front of others. However, it does not delve into the connection between individual and maximal bids. The choice-based approach involves making assessments about the individual bids of others. However, there is no connection between the bids of others and their valuations. The intention-based approach connects bids to values. Putting bounds on bidding functions allows us to incorporate intentions without making specific assumptions on the objectives of others. This approach can be useful for those who like the qualitative features of the classical equilibrium model.

Next, we investigate how these three bids perform in real-life first-price auctions. There are several reasons for doing this: first, we are interested in how our bids perform outside the worst-case scenario; second, we want to demonstrate that our bids can be used off-the-shelf to bid in real auctions; and third, we are interested in how our bids compare to those placed by real bidders. This is the ultimate test as, after all, our bids rely on minimal information, while real bidders are often more experienced and tend to have more information. Finally, we wish to compare our methodology to the state-of-the-art in the literature, where bidders are assumed to bid according to a Bayes-Nash equilibrium (BNE).

To assess the empirical performance of our bids, we use data from auctions in the field and from a laboratory experiment. For the field data, we analyze data on timber auctions. The advantage of this data set is that it is well studied, earnings can be substantial, and bidders are experienced. The disadvantage is that values are not observable. This is overcome by estimating bidders' values from English auctions for similar lots. The data from the laboratory experiment is taken from Güth and Ivanova-Stenzel (2003). One advantage of laboratory experiments is that we know the bidders' values. Two particular disadvantages are that earnings are limited, and that bidders are students who may lack experience in bidding in first-price auctions.

We evaluate bids by their empirical loss: the empirical loss measures how close a bid's payoff comes to the best possible payoff given the empirical distribution of the opponent bids. A small loss means that the bid is essentially a best response given the actual bidding behavior. To achieve better comparability across bidders and data sets, we compute a normalized loss that is given by reporting the loss of

a bidder as a percentage of the largest payoff attainable by this bidder.

We find that the bids of our three implementations perform well; the choice-based approach performs best. Its average loss is 1% in the field and 3.1% in the laboratory experiment across the two treatments. Bids resulting from the outcome-based and intention-based implementations perform slightly worse. Note how this compares to our benchmarks. When a BNE can be computed, the corresponding bid yields an average loss of 1.2% in the field and 4.2% in the laboratory experiment. The bids placed by the actual bidders yield an average loss of 6.6% in the field and 10.5% in the laboratory. Hence, we outperform actual bidders and perform similarly to the BNE wherever it can be calculated.

We now review the related literature. The theoretical auction literature assumes that all bidders reason identically or within the same framework. This is true for the BNE approach (Vickrey, 1961), rationalizability (e.g., Battigalli and Siniscalchi, 2003), iterative regret (Halpern and Pass, 2012), level-k reasoning (Crawford and Iriberri, 2007), and papers using ambiguity (Lo, 1998; Gretschko and Mass, 2019; Mass, 2018). Rothkopf (2007) proposes analyzing auctions with decision-theoretic methods but does not provide any detail on how this should be done. Note that while we take a decision-theoretic approach, we also incorporate some understanding of the behavior of others. Even outside the theoretical auction literature, there are only a few approaches to playing games that do not assume that everyone reasons identically and do not treat the game as a decision problem. Renou and Schlag (2011) assume that others reason identically with some probability. Schlag (2018) proposes a solution concept on the basis of outperforming others where players do not reason about others. However, the objective therein for evaluating an action is far from that used in classic economic models, unlike the proposal brought forward in this paper.

The way in which we select a bid relates to the literature in different ways. It is more in line with Huber (1965), who qualifies a choice as robust if it is almost optimal whenever the model is slightly misspecified. It can be interpreted as looking for a particular ϵ -optimal action (Radner, 1980). It has axiomatic foundations within the framework of minimax regret (Milnor, 1954; Hayashi, 2008) when environments are identified with states. We do not use the term “regret,” as most of the literature on regret is concerned with loss that accrues ex-post (e.g., Loomes and Sugden, 1982; Linhard and Radner, 1989; Guo and Shmaya, 2019); although an exception is Schlag and Zapechelnuyk (2016a,b).

In the pricing problem of a monopolist, Bergemann and Schlag (2011) consider as the ambiguity set the Prokhorov neighborhood of a particular demand function.

Analogously, in our choice-based approach, it is possible to consider the neighborhood of a particular bid distribution. However, in the absence of sufficient data, it is not clear how to choose this distribution. Note that our bids do not require information about bidding in other auctions. When there is abundant data, there are many alternative approaches, such as those outlined in the learning literature; for example, fictitious play (Hon-Snir et al., 1998) or no-regret learning (Blum and Mansour, 2007; Han et al., 2020).

We propose strategies for a given game form (first-price sealed-bid auction) for which choosing a strategy is non-trivial as there is no dominant strategy. An alternative research agenda in mechanism design deals with game forms where players do not have to be strategically sophisticated as there is a dominant strategy (Azar and Micali, 2012; Azar et al., 2012; Giannakopoulos et al., 2020).

Some of our bids involve linear bid shading. Behavior of this form has been documented in laboratory experiments (Bajari and Hortacısu, 2005; Filiz-Özbay and Özbay, 2007) and has been exogenously imposed as a model of simple bidding (Compte and Postlewaite, 2013).

Section 2 contains the results of our survey. In Section 3 we introduce the basic setup. Section 4 proposes how to bid for a given ambiguity set, which we refer to as “robust bidding.” In Section 5, we present the general principles for constructing an ambiguity set and Section 6 provides a specific implementation for each of the three perspectives that results in three different bids. In Section 7, we bring these bids to the data, and Section 8 concludes. The Online Appendix contains the proofs, theoretical extensions, a detailed description of the data, the analysis of the laboratory experiment, robustness checks to the empirical analysis, and alternatives to the parameter selection.

2 Survey

In an anonymous survey conducted in Fall 2021, we asked 35 leading academic economists who work on auctions and nine professional auction consultants about their approach to bidding in first-price auctions. We received 27 responses. Nineteen of the 27 respondents had previously been asked to consult or bid for someone else in such an auction. The median number of auctions they had participated in, conditional on participating in at least one auction, is around four. A few respondents had participated in many auctions, the maximum being around 100. In Online Appendix A, we explain the format of the survey in more detail.

In our first question, which was verbal, we asked whether the theoretical and

empirical auction literatures offer guidance when bidding. In terms of the empirical literature, one respondent suggested finding the bid by estimating the equilibrium bidding behavior using historical data. Five respondents suggested using historical bids (and covariates) to estimate the other bidders' bid distribution. One of them noted that they had never had access to data when bidding. None of these respondents suggested how to proceed when no satisfactory data is available. A large majority of the respondents (about 70%) mentioned the qualitative guidance of the theoretical literature. These respondents found the literature helpful in guiding strategic thinking, as it raised issues, such as the necessity to bid below value (bid shading), the possibility of the winner's curse, and the effects of risk attitudes, budget constraints, the number of bidders, asymmetries, and behavioral biases. Other respondents (around 20%) viewed the theoretical literature as "fairly useless," as it did not "adequately reflect situations found in practice." They pointed out that the theoretical models rely on common knowledge assumptions that do not hold in practice. They mentioned that clients found it difficult to assign probabilities to opponents' values and bids and to then trust these subjective assessments.

In the next questions, we asked the bidders to quantify the main difficulties when bidding in first-price auctions. For this we used a Likert scale with the ordered categories, "not a problem at all," "can be a problem," "a problem," "a serious problem," and "a substantial problem." The vast majority of the respondents regarded assessing other bidders' beliefs, understanding how others bid, assessing opponents' values and their distribution, and identifying the correlation between bidders' values as a serious or substantial problem, the two highest categories on the Likert scale. Most of the respondents indicated that it can be a problem for a bidder to find their own value, and that formulating an objective is less of an issue. Our interpretation of these findings is that most respondents face serious and substantial problems when wishing to apply the existing equilibrium-based framework.

In the verbal part of the questionnaire, we also received various comments on other practical issues. A very experienced participant said that clients find it easier to reason about bids as opposed to values. Another experienced respondent said that their clients could not specify the distribution of opponent bids but tried to estimate the number of bidders and their expected bids. Two others said that the clients' values and objectives were often determined by within-firm negotiations that change over time. There was also concern that top managers might face ex-post complaints through a public discussion of the auction outcome. One participant noted that auctioneers often lack commitment to the auction outcome and try to negotiate after observing the bids. Finally, some participants described some clients

as being reluctant to reason strategically.

We learn the following lessons from the survey. While a handful of respondents provided an explicit recommendation for how to find bids when there is data, none identified a method for making an explicit recommendation *without* using data. Broadly speaking, the answers in the survey reveal the following hierarchy of challenges. The objectives of the bidder appear to be clear; understanding their own value is already more difficult. A quantification of the remaining ingredients (beliefs, bids, and values of others) seems unrealistic. In particular, all but one of the respondents neither explicitly nor implicitly refer to bidders playing an equilibrium. We note that these respondents use very different forms of reasoning, as some focus more on values, others on bids, and others on general features. In summary, there seems to be demand for a quantitative approach that can be applied to real-life first-price auctions.

3 The Basics

We consider a first-price, sealed-bid auction for a single good. There are n bidders, where $n \geq 2$. Bidders simultaneously make non-negative bids. The bidder with the highest bid wins the good and pays her bid. Ties are broken randomly. In our analysis, we allow for a reserve price.

We ask how bidder 1 should bid. Bidder 1 knows her private value v_1 for the good, where $v_1 \in \mathbb{R}_+$. We formally look for a bidding function that leads to a bid recommendation b_1 for each value v_1 . Our analysis also applies if v_1 is an unbiased estimate of bidder 1's value, as long as the distribution of this value is independent of the bids of the other bidders.

Bidder 1 is risk-neutral. Alternative preferences can be incorporated (see Online Appendix F for risk aversion). Bidder 1's payoff when bidder j bids b_j , $j = 1, \dots, n$, is denoted by $u_1(b_1, b_{-1})$ and equal to

$$u_1(b_1, b_{-1}) = \begin{cases} \frac{1}{m} (v_1 - b_1) & \text{if } b_1 = \max_j \{b_j\} \text{ and } |\{k: b_k = b_1\}| = m, \\ 0 & \text{if } b_1 < \max_j \{b_j\}. \end{cases}$$

The bids of the other bidders are influenced by many factors, including their values, information, preferences, budgets, and the characteristics of the good. To capture this variation, we model bidder 1 as facing some bid distribution B . We also refer to the bid distribution faced by the bidder as the "environment." Formally, the bid distribution B belongs to $\Delta(\mathbb{R}_+^{n-1})$, where ΔX denotes the set of all cumulative

distribution functions (cdf) on X . Note that we do not assume that others bid independently. The payoff of a bid b_1 when facing a bid distribution B is measured by the expected utility. Specifically, we extend u_1 to be the expected utility of bidding b_1 when facing the bid distribution B , so

$$u_1(b_1, B) = \int u_1(b_1, b_{-1}) dB(b_{-1}).$$

If bidder 1 knows the environment, we assume that she would choose the bid that maximizes her expected utility. However, we are considering a bidder who is uncertain about the environment she is facing, as she is unable to assess the exact distribution of the opponent bids. Our survey highlights this as an important element in real-life auctions. In contrast, note that the conceptual framework of subjective utility maximization implicitly assumes that the decision-maker is able to pin down the precise probabilities.

4 Dealing with Uncertainty: Robust Bidding

We now present how to bid when uncertain about the environment, where the environment is given by the distribution of bids made by the other bidders. An environment is called “conceivable” if the bidder cannot rule out that this is the environment she is facing. We call the set of conceivable environments the *ambiguity set*. As the bidder is uncertain about the environment, the ambiguity set contains at least two elements. In the following, we propose how to bid for a given ambiguity set, then, in the next section, we show how to construct the ambiguity set.

When evaluating a bid, we do not pay particular attention to any specific environment. While there are typically many conceivable environments, the bidder can only place a single bid, and it is unlikely that this bid maximizes the expected utility in each conceivable environment; thus, it will often lead to a loss in payoff relative to the best possible payoff. Desiring close-to-optimal payoffs without paying particular attention to a specific environment motivates us to select the bid that attains the smallest maximal loss among the conceivable environments. The value of the smallest maximal loss is called the “minimax loss.” The selected bid comes with the performance guarantee that loss is bounded above by the minimax loss. The performance guarantee depends on the ambiguity set. Smaller sets (in terms of set-wise inclusion) represent less uncertainty and yield weakly superior performance guarantees. This approach to decision-making under uncertainty is founded on axioms (Milnor, 1954; Hayashi, 2008; Stoye, 2011). It is important to keep in

mind that the selected bid will lead to a loss that is smaller than its performance guarantee and will often be much smaller.

We introduce the model formally. Let $\mathcal{B} \subseteq \Delta(\mathbb{R}_+^{n-1})$ be the set of bid distributions that bidder 1 thinks she may face. Bid distributions in \mathcal{B} are the conceivable environments and \mathcal{B} is the ambiguity set. Importantly, the bidder is not willing or able to assess the likelihood of each conceivable bid distribution. This stands in contrast to the Bayesian approach, in which the bidder would have a probabilistic prior belief over \mathcal{B} .

We suggest how to bid given \mathcal{B} . Intuitively, a bid performs well when facing a given bid distribution if it attains a payoff that is close to the highest possible payoff for this distribution. We measure this “closeness” with the following loss function l . For a given bid distribution B and an own bid b_1 , let l be the difference between the best possible payoff given this bid distribution and the own payoff, that is:

$$l(b_1, B|v_1) = \sup_{\bar{b}_1} u_1(\bar{b}_1, B) - u_1(b_1, B).$$

The value of the maximal loss among all conceivable distributions, given by $\sup_{B \in \mathcal{B}} l(b_1, B)$, is a performance guarantee of the bid b_1 . The maximal loss measures the highest possible loss in utility of bidding b_1 due to a lack of information about the true distribution.

We select a bid that attains the best performance guarantee for \mathcal{B} . This bid minimizes the maximal loss and is therefore called the *minimax loss bid*. The minimum is calculated among all bids and the maximum is computed from among all environments in the ambiguity set. In our analysis, the minimax loss bid is typically unique. Formally, b_1^* is a minimax loss bid if

$$b_1^* \in \arg \min_{b_1 \in \mathbb{R}_+} \sup_{B \in \mathcal{B}} l(b_1, B|v_1).$$

We denote the value of the minimax loss as l_1^* , where

$$l_1^* = \min_{b_1 \in \mathbb{R}_+} \sup_{B \in \mathcal{B}} l(b_1, B|v_1).$$

We refer to the conceivable bid distributions that maximize loss as the “worst-case” bid distributions. Maximal loss is sometimes obtained by the limit of a sequence of converging distributions (where we consider convergence in probability). We then refer to the limit point as a worst-case distribution.

We emphasize that loss is computed conditional on a bid distribution. Thus, we differ from the classic minimax regret literature (Savage, 1951; Hayashi, 2008;

Linhard and Radner, 1989), where minimax regret is computed in expectation based on the realized (ex-post) outcomes; that is, the bids. We choose to condition on the bid distribution to better capture uncertainty. Specifically, if there is only a single conceivable bid distribution, and thus no uncertainty, then the minimax loss is equal to 0. In contrast, minimax regret will typically not be zero in such cases. Note that we do not allow for randomized bidding, although it is known from the theoretical literature on minimax regret that superior performance can be attained by allowing for randomization in choices. However, we want to propose a simple bid, but arguably, randomization is not simple. Moreover, to implement a randomized bid, we have to be able to commit to the randomization device. Such a commitment might be difficult, and even more so when a bidder has to justify the actual bid in front of others.

Finally, we believe that minimax loss bids are easy to explain. Our survey reveals the importance of being able to justify bids in front of others. Minimax loss can loosely be explained as follows. Bidding according to the minimax loss methodology means computing a flexible bid that is good in many environments. The cost of flexibility is measured by the maximal loss. The alternative would be to guess the environment and best respond to it. Guesses come with the risk of not being correct.

We qualify the minimax loss bids as being *robust*, which is in the spirit of Huber (1965), who calls a policy robust if it is almost optimal whenever the model is (slightly) misspecified.

5 Conceptualizing Uncertainty: General Principles

In the previous section, we show how to select a bid for a given understanding of uncertainty as formally captured by the ambiguity set. In this section, we show how to construct this ambiguity set, thereby presenting a method for conceptualizing uncertainty. The ambiguity set is constructed in two steps. In the first step, we choose the mindset or perspective of the bidder that will be used when describing the uncertainty. In the second step, we use an iterative procedure to conceptualize uncertainty from the perspective selected in the first step. We illustrate the process by presenting specific implementations in Section 6.

The first step consists of choosing one of three perspectives the bidder may use to reason about the environment she is facing. We introduce the three perspectives in order of increasing detail about what leads to the maximal bid of the opponents. The first perspective involves only reasoning about the highest bid of the other

bidders, which is the object most cared about in first-price auctions. The choices and underlying motives of others are ignored. We call this perspective *outcome-based*. The second perspective is called *choice-based* and arises when reasoning about the individual bids of others and their joint distribution without thinking about what motivates these bids. The third perspective is called *intention-based* and includes reasoning about opponent bidders' intentions and values.

The selection of the perspective of the bidder is influenced by many factors and include her individual circumstances and the information available to her. An important role is also played by the simplicity of the resulting model and how well it can be justified.

We discuss some general considerations that can assist in selecting the most appropriate perspective for the application in mind. A bidder who has a good feeling for the intentions of the other bidders, and with some understanding of their values, might choose the intention-based approach. An advantage of this approach is that the degree of similarity of others can be well modeled using constraints on their values and on how they bid given these values. Moreover, this approach is close to the classic framework, as values and strategies are the primitives. The disadvantage is that constraints on bid distributions have to be computed from constraints on values and bidding strategies. This can be intricate and opaque, particularly if values can be correlated. The approach can be less straightforward in the mindset of a pragmatic bidder who has observed past bids and wishes to adjust to these bids, in which case, the choice-based approach might be more appropriate. One advantage of this approach is that it is very effective in modeling the correlation between bids. Both the intention- and choice-based approach build on specific models of individual behavior that can be avoided by the outcome-based approach. Under the outcome-based approach, the bidder only has to make some conjectures about the highest opponent bid. Focusing on the object that determines the payoffs can lead to this approach naturally. Moreover, it is particularly useful when the bidder has little information.

We now come to the second step in which we determine the ambiguity set. Our leading principle is the understanding that it can be straightforward to establish what *will not* happen, even when there is uncertainty about what *will* happen. Following this principle, the ambiguity set will be obtained by sequentially eliminating environments that can be ruled out. Those that are left over are declared to be conceivable. The elimination occurs by describing a class or feature of the environment, using the perspective selected in the first step, and then asking the bidder to confirm that (she thinks) this feature will *not* occur. In the outcome-based approach,

the feature is described in terms of properties of the distribution of the maximal opponent bid. In the choice-based approach, the feature is formulated directly in terms of the joint bid distribution. In the intention-based approach, features of the values and their correlation, together with bidding behavior, are described.

The elimination of environments within the second step follows an iterative process. We propose starting with a large ambiguity set, which is potentially the set containing all environments. We then refine the ambiguity set by eliminating environments that we think are unlikely to be true. These environments are described in terms of the chosen perspective. After each elimination round, we compute all worst-case environments; that is, the environments that determine why the performance guarantee is not better. The features of these environment are considered and determined, whether or not they are conceivable. The underlying idea is that environments should only be ruled out if they lead to bids with better performance. A necessary condition for obtaining better performance from one round of elimination to the next is, therefore, that at least some worst-case distributions can be eliminated. For instance, in the choice-based approach when all environments are conceivable, it is observed that a worst-case distribution involves all others placing the same bid. This can motivate the introduction of a lower bound on variation in the bidding of others (as we do in Section 6). The iterative process stops when none of the worst-case distributions can be ruled out or when the performance guarantee is sufficient, as assessed subjectively.

Note that it makes little sense to eliminate environments one by one, particularly as there are infinitely many. The idea is therefore to use the perspective chosen in the first step to formulate general features that are considered to be inconceivable and then to rule out all environments that have this feature.

There are two trade-offs in constructing the ambiguity set. The first is to find the balance between performance and the risk that the ambiguity set does not contain the true environment. The more environments are eliminated, the more likely it is that the true environment is ruled out, and the better the performance guarantee when facing one of the remaining environments. The second trade-off concerns the bounds between simplicity and precision. A simple model is one that can be described with a few features that involve few numerical parameters. Ideally, only a small number of these features impact the optimal bidding function and the corresponding parameter values are easy to identify. On the other hand, identifying many different features might lead to a high precision and good performance, which may be at the expense of simplicity and tractability.

Sometimes the iterative procedure might stop, because it is difficult to quantify

the bound that underlies some intuitive features. For instance, it may be difficult to quantify a lower bound on the variation, even if it is known that there will be some variation. In such cases, we propose not to change the bid but only to investigate how the performance changes in the bound. In other words, we obtain a parametric family of ambiguity sets. Instead of committing to specific parameter values, we take the bid from the last iteration and investigate how its performance changes in these parameters. This can be a powerful way to understand performance without needing to commit to specific features of the environment. An alternative method of proceeding, when the performance guarantee is unsatisfactory, is to gather more information prior to placing a bid.

We list a few questions that may assist the bidder to eliminate environments from her ambiguity set. In all three approaches, it can be useful to start by asking vague questions, as these can help rule out many environments at the same time. These questions include: Do you have any information on bidding in similar auctions? What range of bids do you expect? Do you expect others to be very aggressive in their bids? Next, we list some questions that are specific to the outcome-based approach: What is the possible range of the highest opponent bid? How likely is it that the highest opponent bid is very high or very low? How dispersed can the distribution of the highest opponent bid be, or how likely is it that the highest opponent bid will take a specific value? The answers can be used to specify bounds on the distribution of the highest opponent bid. Additional specific questions can be useful for the choice-based approach and the intention-based approach: Do you know how many bidders are taking part in the current auction? How similar are these bidders to each other and to yourself? Are the other bidders incumbents or entrants? Are there aggregate factors that might cause others to bid in a similar way? (This will lead to correlated bid distributions.) Or is it highly likely that they will bid independently, given the information you have? What do you know about how others will bid? What kind of budget constraints do the other bidders face? For the intention-based approach, the following questions can also be asked: Do you know anything about the values of the other bidders? What do you know about their risk preferences? Are they experienced? Are the bidders sophisticated or are they naive?

Data on past auctions can be included in each of the three approaches. On an intuitive level, data can help understand what is conceivable. On a more formal level, data in the form of empirical distributions can help quantify bounds on the possible outcomes, and the behavior and values of others. Information on past highest opponent bids is most straightforwardly included in the outcome-based approach. Data

on individual bids in the past can be incorporated in the choice-based approach. Data on past valuations and the bidding of others is useful in the intention-based approach, however, this type of information is rare.

A final comment regards the value of strategizing and incorporating beliefs that other bidders might also be following the ideas outlined in this paper. Such considerations do not enter the outcome or choice-based approaches. Strategic behavior is only modeled in the intention-based approach. Note, however, that our methodology is not about strategizing in the sense of forming higher-order beliefs. Rather, we focus on what others do not do and not about what they believe. In the intention-based approach, we suggest confirming whether or not the solution is contained in the ambiguity set to deal with the possibility that others follow the methodology of this paper. We further discuss this in the light of our implementations at the end of Section 6.

6 Constructing Three Explicit Bids

We now illustrate how the general principles introduced in Section 5 can be used to obtain specific bid recommendations in a low-information setting. Specifically, we consider a bidder who knows her own value and the number of bidders but otherwise has very little information. We implement each of the three approaches outlined in Section 5 and thus generate three bids.

6.1 An Outcome-Based Bid: Maximal Uncertainty

We first adopt the outcome-based perspective and reason directly about the distribution of the highest opponent bid without thinking about where it may come from. As there are only very few restrictions, we refer to this particular implementation as bidding under *maximal uncertainty*. Following Section 5, we adopt an iterative procedure to determine the ambiguity set.

We start with the set of all bid distributions as the ambiguity set. So, \mathcal{B} equals $\Delta(\mathbb{R}_+^{n-1})$. We need to identify the minimax loss bid and the corresponding worst-case distributions. The minimax loss bid is $v_1/2$. The reasoning is as follows. Consider the payoffs of bidder 1 who bids b_1 . Assume first that the highest opponent bid is 0 (in which case all others bid 0). Then, bidder 1 wins and has utility $v_1 - b_1$. However, bidder 1 could have attained a utility close to v_1 by bidding slightly above 0. Hence, her loss is $(v_1 - 0) - (v_1 - b_1) = b_1$. Now consider a bid distribution in which the highest opponent bid is slightly above b_1 . Then, bidder 1 does not win

the object but could have won by bidding slightly higher. In this case, her loss is given by $v_1 - b_1 - 0 = v_1 - b_1$. In Online Appendix C, we show that all other distributions yield a weakly lower loss than either of these. The minimax loss bid b_1^* minimizes the maximum of these two losses by equating them, therefore, b_1^* solves $b_1^* = v_1 - b_1^*$, and hence $b_1^* = v_1/2$. The level of minimax loss is $v_1/2$.

The minimax loss is not small, so we proceed to the next step of the iteration. We need to investigate whether the features of the worst-case distributions can be ruled out from the perspective of the bidder. In the outcome-based approach, the perspective is to reason only about the highest opponent bid, therefore we must determine whether we can rule out that the maximal bid of the others is almost certainly equal to 0 or almost certainly slightly above $v_1/2$. In many auctions, it is not plausible that such extreme bids occur almost certainly. Thus, we introduce a lower threshold L , below which the highest opponent bid is not believed to fall. While it is easy to assume that such a threshold exists, a bidder might find it difficult to pin down an exact number. Therefore, we introduce \bar{p} , a parameter that specifies the maximal probability under which it is conceivable that the maximal opponent bid falls below L . Let $\mathcal{B}_{L,\bar{p}}$ be the resulting set of distributions, formally given by

$$\mathcal{B}_{L,\bar{p}} = \{B \in \Delta(\mathbb{R}_+^{n-1}) : B(b_{-1}) \leq \bar{p} \text{ if } b_j < L \text{ for all } j \neq 1\}.$$

In the case where the auction has a reserve price R , the reserve price bounds the threshold L from below, so $L \geq R$. Moreover, \bar{p} must be 0 if $L = R$.

We present the minimax loss bid and the minimax loss for the set of conceivable distributions $\mathcal{B}_{L,\bar{p}}$. The proof of the result can be found in Online Appendix C.

Proposition 1. *Assume that $0 \leq L < v_1$ and $\bar{p} \leq (v_1 - L)/(v_1 + L)$. The unique minimax loss bid for the set of conceivable environments $\mathcal{B}_{L,\bar{p}}$ is given by*

$$b_1^* = \frac{v_1 + L}{2}. \tag{1}$$

The corresponding minimax loss equals

$$l_1^* = \frac{v_1 - L}{2}.$$

The above formulae for minimax loss and the corresponding bid are computed almost analogously as when $L = 0$. The findings do not depend on the number of

bidders n . This is because there are always worst-case distributions in which all other bidders make the same bid. In particular, the proposition is also applicable when the number of bidders is unknown or when the other bidders collude perfectly. Note also that the parameter \bar{p} does not enter the above result. The intuition for this is as follows. When the likelihood of the highest opponent bid being below L is sufficiently small, then the best response is to bid at least L . However, once \bar{p} is too large given L , the minimax loss bid will depend on \bar{p} . Recall from Section 5 that we argue in favor of simple implementations that depend on as few parameters as possible. Note that L is decreasing in \bar{p} . So, we suggest to choose L so small that the maximal mass below L falls below the threshold for \bar{p} , as identified in Proposition 1. This yields a minimax loss bid that does not depend on \bar{p} and, as such, is simple. Given this reasoning, we do not present the minimax loss bid when \bar{p} is above the threshold.

The bidding function presented in Equation (1), evaluated for the special case where $L = 0$, is also a solution to the minimization of maximal regret found in Sošić (2007) (Example 1) and Halpern and Pass (2012). These papers follow the minimax regret literature (Savage, 1951) in computing loss ex post; that is, conditional on the actual bids of others. As mentioned in Section 4, we choose a different approach and compute loss conditional on the bid distribution. Note that the results obtained in Proposition 1 remain unchanged when minimax regret is used instead of minimax loss (and $\bar{p} = 0$). This equivalence of the two decision criteria does not hold for the other results in this paper.

Adding the lower bound L on bids in the second step of the iteration has reduced the minimax loss, yet the minimax loss remains large. Following Proposition 1, the minimax loss is 50% of the maximal possible payoff of bidder 1, which is equal to $v_1 - L$. Thus, we proceed with the iteration.

The worst-case bid distributions feature the highest opponent bid being either almost surely equal to L or almost surely slightly above b_1^* . We proceed to rule out that the maximal bid is deterministic as in both of these cases; that is, we expect some variation in the maximal bid. Specifically, we impose lower and upper bounds on the density of the highest opponent bid. This rules out point masses and holes in the support of the distribution of the highest opponent bid. For analytic simplicity, we limit attention to the case where $\bar{p} = 0$, so the maximal bid is believed to be certainly above L .

We formally introduce these bounds. Let α_1 and α_2 be such that $0 \leq \alpha_1 \leq 1 \leq$

α_2 . A bid distribution is qualified as being conceivable if

$$\alpha_1 \cdot \frac{x_2 - x_1}{v_1 - L} \leq \mathbb{P}(\max b_{-1} \in [x_1, x_2]) \leq \alpha_2 \cdot \frac{x_2 - x_1}{v_1 - L} \quad (2)$$

for all $L \leq x_1 \leq x_2 \leq v_1$. The resulting set of conceivable bid distributions given such parameters is denoted by $\mathcal{B}_{L, \alpha_1, \alpha_2}$.

We hesitate to specify values for the bounds α_1 and α_2 as there are no salient values. Therefore, instead of committing to specific values for α_1 and α_2 , we adopt the flexible approach (as presented in Section 5) and investigate how the minimax loss bid from Proposition 1 performs as a function of these parameters. Surprisingly, we find that the optimal bid from Proposition 1 often still attains minimax loss. The proof is in Online Appendix C.

Proposition 2. *Let $b_1^* = (v_1 + L)/2$ and*

$$\bar{l}_1 = \frac{(\alpha_2 - \alpha_1)^2}{16\alpha_2} (v_1 - L).$$

(i) *The maximal loss of bidding b_1^* when $B \in \mathcal{B}_{L, \alpha_1, \alpha_2}$ equals*

$$\sup_{B \in \mathcal{B}_{L, \alpha_1, \alpha_2}} l(b_1^*, B) = \begin{cases} \bar{l}_1 & \text{if } \alpha_1 \leq -\alpha_2 + 2\sqrt{\alpha_2} \text{ and } \alpha_2 \leq 4 \\ (1 - \alpha_1) \frac{\alpha_1^2 + \alpha_2^2 - 2\alpha_2}{(\alpha_2 - \alpha_1)^2} \frac{v_1 - L}{2} & \text{if } \alpha_1 > -\alpha_2 + 2\sqrt{\alpha_2} \text{ or } \alpha_2 > 4. \end{cases}$$

(ii) *If $0 \leq \alpha_1 \leq 2\alpha_2 - \sqrt{5\alpha_2^2 - 4\alpha_2}$, then b_1^* is the minimax loss bid and the minimax loss is \bar{l}_1 .*

The above result shows that the minimax loss bid from Proposition 1 performs well when the difference between the two density bounds α_1 and α_2 is not too large. For example, if $\alpha_1 \geq 1/3$ and $\alpha_2 \leq 3$, then the minimax loss is below 15% of the maximal possible payoff $v_1 - L$. Note, also, that the model maintains a low level of specificity as the bidding behavior only depends on a single parameter, L .

6.2 A Choice-Based Bid: Variation in Bidding

In this section, we apply the choice-based approach and reason about the other bidders' choices without taking their motivation into account. We refer to this particular implementation as *variation in bidding*.

We start with the set of all bid distributions, as in the outcome-based approach. In the first iteration step, we eliminate all bids below L . As shown in Subsection

6.1, this leaves us with worst-case distributions in which everyone bids L , or where they coordinate so that at least one of them bids just above $(v_1 + L)/2$ and all others bid weakly less. Note the lack of variation in the individual bidding in the first case and the perfect correlation in the second case. We proceed to rule out such environments as there is too little variation and independence in the bidding. In the outcome-based approach, we can only impose conditions on the maximal bid. Here, in the choice-based approach, we can put restrictions directly on the bids.

We introduce some independence by assuming that with probability τ a bidder bids independently. We call such a bidder “constrained.” This allows us to capture the correlation structure with a single variable. We introduce variation in individual bidding by assuming that the bids of the constrained bidders are independent draws from a distribution G with a support that contains $[L, v_1]$. The bids of the unconstrained bidders are arbitrary; in particular, any degree of correlation among the bids of the unconstrained bidders is allowed. It follows that bidder j ’s marginal bid distribution must be at least τG . We discuss the choices of τ and G below.

The ambiguity set is denoted by $\mathcal{B}_{\tau, G}$. Informally, the ambiguity set is defined as the set of joint bid distributions constructed in the following way. First, we determine which of the $n - 1$ other bidders are constrained by independently making each of them constrained with probability τ . Second, the constrained bidders independently choose their bids according to distribution G . Third, depending on which bidders are constrained, there is a distribution H from which the bids of the unconstrained bidders are jointly drawn. The set of conceivable bid distributions is then formally given by

$$\mathcal{B}_{\tau, G} = \left\{ B \in \Delta(\mathbb{R}^{n-1}) : B(b_{-1}) = \sum_i \tau^{\sum_{j=1}^{n-1} i_j} (1 - \tau)^{n-1 - \sum_{j=1}^{n-1} i_j} H_i \left((b_j)_{j:i_j=0} \right) \prod_{j:i_j=1} G(b_j) \right. \\ \left. \text{for } H_i \in \Delta \left(\mathbb{R}^{n-1 - \sum_{j=1}^{n-1} i_j} \right) \text{ and } i \in \{0, 1\}^{n-1} \right\}.$$

How do we choose G ? Each of the other bidders is conceived to bid below b with probability at least $\tau \cdot G(b)$. So, G introduces some variation in the bidding of each of the other bidders. For instance, a bidder might choose G as a smooth approximation of the empirical distribution of bids in previous and similar auctions. A salient choice that we recommend in the absence of data is to choose G as the uniform distribution on $[L, v_1]$. The uniform distribution adds variation without giving prominence to any particular bid. We do not place any mass of G above v_1 as this would simply lower the loss without changing the incentives.

How do we choose the parameter τ ? The parameter τ is the minimal likelihood that any other bidder independently selects their bid from the distribution G . So τ

simultaneously captures the degree of independence and the trust in G . In this sense, τ reflects the accuracy of the estimate G of the bid distribution. This accuracy, and hence τ , may increase as the bidder gathers more data. If τ is close to 0, then the ambiguity set contains almost all distributions over bids above L . If τ is close to 1, then it is as if we are facing $n - 1$ independent bids from the distribution G . In the empirical applications reported in Section 7, we set $\tau = 0.15$, where we have committed to this value prior to looking at the data.¹ In Online Appendix E, we provide a method for selecting τ from a set of possible values and also discuss empirical approaches for determining τ and G .

Choosing G uniform would arise if the bidder thought that she was facing others who used reinforcement learning (Q-learning) with ε -greedy exploration (Sutton and Barto, 2018), where $\varepsilon = \tau$. The other bidders' Q-matrices, value distributions, and greedy actions are unknown, but the bidder knows that with probability ε they will choose a bid uniformly. A common choice for ε is 0.1.

We now add some comments on the rationale of our modeling. Modeling minimal independence with a binomial distribution is simple, rests on a single free parameter, and is easy to communicate.² The decision to let G be the uniform distribution on $[L, v_1]$ means that the set of conceivable bid distributions of bidder 1 depends on the valuation of this bidder. This reflects that the conceivable bid distributions are subjective from the viewpoint of the bidder, where v_1 is part of the bidder's characteristics. Moreover, it is the only salient choice that does not add another parameter that needs to be chosen.

We proceed to the analysis. The next lemma shows that there are always worst-case distributions that take a particularly simple form. The proof is in Online Appendix C.

Lemma 1. *For any bid b_1 of bidder 1, there is a worst-case bid distribution in which all unconstrained bidders submit the same bid.*

Given Lemma 1, we can now limit attention to investigating the maximum loss when all unconstrained bidders submit the same bid. Note that such bid distributions have a mass point at the bid of the unconstrained. Given this mass point, we have to distinguish three cases in our calculations of expected utility, depending on how b_1 compares to this mass point. If bidder 1 bids strictly above

¹Our motivation was that the constraints should not be negligible, but they should also not be substantial. Guided by our intuition (which is clearly subjective), we find a value of τ equal to 0.1 to be too small and a value of 0.2 to be too large and, hence, converge on 0.15.

²An alternative approach that we do not pursue in this paper is to model minimal independence through Bayesian networks or Markov random fields.

the mass point, then the expected utility, denoted by u_1^+ , is equal to

$$u_1^+(b_1) = \sum_{k=0}^{n-1} \binom{n-1}{k} \tau^k (1-\tau)^{n-1-k} G(b_1)^k (v_1 - b_1)$$

as bidder 1 wins only if b_1 is also higher than the bids of the k constrained bidders, $k = 0, 1, \dots, n-1$. Conversely, if bidder 1 bids strictly below the mass point, then the bidder wins only if all other bidders are constrained, in which case b_1 is higher than bidder j 's bid with probability $G(b_1)$. Therefore, the expected utility of bidder 1, denoted by u_1^- , equals

$$u_1^-(b_1) = \tau^{n-1} G(b_1)^{n-1} (v_1 - b_1).$$

We do not have to consider the third case, where the bid b_1 is equal to the mass point, as this will never be a worst-case distribution.

For our analysis of minimax loss, we distinguish between two cases. We say that bidder 1 is bidding too high in a given environment if the best response to this environment is below her bid. In the other case, we say she is bidding too low.

Consider the maximal loss of bidding too high. As the best response bid h in the given environment is smaller than b_1 , loss is maximal when the point mass of the unconstrained bidders is smaller than h . In this case, the best response payoff is given by $u^+(h)$ and the payoff of the bidder is given by $u^+(b_1)$. Thus, the maximal loss of bidding too high is given by $\sup_{\hat{h}: \hat{h} < b_1} u^+(\hat{h}) - u^+(b_1)$.

Now consider the maximal loss of bidding too low. This means that the best response h is higher than b_1 . Note that the loss of bidding too low is largest if the point mass of the unconstrained bidders lies between b_1 and h . Consequently, the maximal loss of bidding too low is given by $\sup_{\hat{h}: \hat{h} > b_1} u^+(\hat{h}) - u^-(b_1)$.

It then follows that the minimax loss bid equalizes the maximal loss of bidding too low and the maximal loss of bidding too high. As there may be multiple bids with this property, the minimax loss bid is the one with the lowest maximal loss. This leads us to the following proposition; the proof is detailed in Online Appendix C.

Proposition 3. *Let X be the set of bids such that the maximal loss of bidding too high equals the maximal loss of bidding too low, that is,*

$$X = \{b \in [L, v_1] : \sup_{\hat{h}: \hat{h} \in [L, b]} u_1^+(\hat{h}) - u_1^+(b) = \sup_{\hat{h}: \hat{h} \in (b, v_1]} u_1^+(\hat{h}) - u_1^-(b)\}.$$

The minimax loss bid b_1^ when facing $\mathcal{B}_{\tau, G}$ lies in X .*

Given the above proposition, the minimax loss bid equalizes the maximal loss from bidding too high and the maximal loss from bidding too low. The minimax loss bid is the bid for which the loss in X is minimal.

We now investigate in more detail the special case where G is uniform on $[L, v_1]$. Consider first the case of bidding too high. The uniform distribution leads to the expression

$$u_1^+(b_1) = \sum_{k=0}^{n-1} \binom{n-1}{k} \tau^k (1-\tau)^{n-1-k} \left(\frac{b_1 - L}{v_1 - L} \right)^k (v_1 - b_1).$$

This expression is maximized by $\max\{L, \tilde{b}\}$, where

$$\tilde{b} = v_1 - \frac{v_1 - L}{n\tau}.$$

In the worst case, all unconstrained bidders bid L . When τ or n is small, such that $\tau \leq 1/n$, then it is optimal to ignore the constrained bidders and to bid slightly above L . However, when $\tau > 1/n$, then it is optimal to take the constrained bidders into account and to bid above L . Consequently, the maximal loss of bidding too high is given by $u_1^+(\max\{L, \tilde{b}\}) - u_1^+(b_1)$.

Now, consider the case of bidding too low. The uniform distribution leads to

$$u_1^-(b_1) = \tau^{n-1} \left(\frac{b_1 - L}{v_1 - L} \right)^{n-1} (v_1 - b_1).$$

The worst-case distribution is attained when all unconstrained bidders bid slightly above b_1 . The best response is to bid slightly above these bids. This yields a payoff of $u_1^+(b_1)$. Consequently, the maximal loss of bidding too low is $u_1^+(b_1) - u_1^-(b_1)$.

The next proposition formally establishes the minimax loss bid and some of its properties. The proof is found in Online Appendix C.

Proposition 4. *Let $G(x) = (x - L)/(v_1 - L)$ for $x \in [L, v_1]$. The minimax loss bid b_1^* when facing $\mathcal{B}_{\tau, G}$ is unique and solves*

$$u_1^+(\max\{L, \tilde{b}\}) - u_1^+(b_1^*) = u_1^+(b_1^*) - u_1^-(b_1^*). \quad (3)$$

Minimax loss equals $l_1^ = u_1^+(b_1^*) - u_1^-(b_1^*)$. The minimax loss bid b_1^* increases in the number of bidders n .*

Evaluating (3) for $\tau = 0.15$, we obtain that the minimax loss is below 26% and 15% of the maximal possible payoff $v_1 - L$ when $n \geq 5$ and $n \geq 9$, respectively.

Note that the worst-case distributions exhibit point masses due to the identical bidding of the unconstrained bidders. This pattern of behavior does not seem plausible. Hence, analogously to Section 6.1, we investigate how the minimax loss bid performs when there is variation in the maximal bid of the unconstrained bidders. Details are provided in Online Appendix B. Once again, we do not reoptimize the bid when introducing these bounds. We find that its maximal loss is below 13% of $v_1 - L$ when $\alpha_1 \geq 1/3$ and $\alpha_2 \leq 3$ for any $n \geq 2$.

6.3 An Intention-Based Bid: Variation in Values

In our third implementation, we derive a bid by using the intention-based approach: we consider environments from the perspective of the underlying values and the bidding strategies of the other bidders. We call this implementation *variation in values*.

We begin without any constraints on values or bidding strategies. Following Subsection 6.1 the worst-case distributions that emerge exhibit either all bidding 0 or the maximal bid being almost certainly slightly above $v_1/2$. As the minimax loss is high, we proceed with the iteration. Our objective is to rule out one of the worst-case bid distributions by connecting bid distributions to conceivable value distributions and bidding strategies. As in Section 6.2, we do this by imposing variation and independence.

At this point, we assume that valuations and bidding behavior are independent across bidders. We have three reasons for doing so. First, such independence is frequently presumed. Indirect evidence can be found in our survey, where none of the respondents have referred to an explicit modeling of the values' correlation. Second, it connects our bidding to the classical literature, where independence is often assumed. Third, it leads to a particularly simple mathematical structure, which is instrumental when choosing the key parameters.

We connect bids to values, thereby indirectly incorporating the objectives of the different bidders. First, we do not think that any of the other bidders will bid above their own value. Moreover, none of them is expected to bid too far below their own value. To capture this, we introduce $\sigma \in (0, 1)$ such that none of the other bidders is believed to bid below σ times their value. Note that this assumption does not rule out that all others bid 0, as this can result from all others having a value of 0. Similar to imposing a lower bound L on bids, it is natural to impose a lower bound on values. We denote this lower bound by K . Together with our assumption about bidding behavior, this rules out bids below $\sigma \cdot K$. The worst-case distribution in which all bid 0 now turns into one in which all bid $\sigma \cdot K$. Next, we impose

some variation on the bids of other bidders. Specifically, we put bounds on how likely it is that bidders have low valuations. We do this by introducing an upper bound on the cdf of the distribution of values of any given bidder. We choose a particular parametric form of an exponential bound as it leads to simple bidding functions. This makes the resulting performance guarantee traceable to the degree of tightness of this bound. Specifically, let F_j be the distribution of bidder j 's values. We choose $\eta > 0$ and $\alpha > 0$, such that we believe that $F_j(v_j) \leq \eta(v_j - K)^\alpha$ for $v_j \geq K$. Hence, in each environment there might be bidders with values above K'' if $K'' < K' = K + \eta^{-1/\alpha}$. At the same time, there are conceivable environments in which all values are believed to be below K' . We refer to K' as the lowest possible maximal value. Consequently, the highest possible bid is at least $\sigma \cdot K'$. Note that the bound on F_j becomes tighter as α increases and η decreases. Thus, we have introduced four parameters whose values have to be chosen: σ , K , η and α .

Let $\mathcal{B}_{\eta,K,\alpha,\sigma}$ denote the set of conceivable bid distributions that arise under the above constraints. It is formally given by³

$$\mathcal{B}_{\eta,K,\alpha,\sigma} = \left\{ B \in (\Delta\mathbb{R}_+)^{n-1} : B_j(b_j) \leq \eta \left(\frac{\max\{b_j, \sigma K\}}{\sigma} - K \right)^\alpha \text{ for all } j \neq 1 \right\}.$$

We give some intuition for how the minimax loss bid is computed. The first observation is that it can be shown that loss is maximal when all other bidders bid σ times their own value and the value distribution puts the maximal mass $\eta(x - K)^\alpha$ on some value x , such that $\sigma \cdot x < v_1$ and the rest of the mass on sufficiently high values, for instance, above v_1/σ . The intuition here is that the worst case arises when the probability of bids is maximal between the bid of bidder 1 and the payoff-maximizing bid.

We then find that the maximal loss of bidding too low is attained when as many as possible bid slightly higher than bidder 1, which means that they have a value slightly above b_1/σ . This maximal loss becomes

$$\eta^{n-1} \left(\frac{b_1}{\sigma} - K \right)^{\alpha(n-1)} (v_1 - b_1).$$

The maximal loss of bidding too high is attained for some value x , such that $\sigma \cdot x$ is below the bid of bidder 1. The corresponding value of loss equals

$$\eta^{n-1} (x - K)^{\alpha(n-1)} \cdot (b_1 - \sigma \cdot x).$$

³Note that the bounds imply a single bound on the bid distribution, namely $B_j(b_j) \leq \eta(b_j/\sigma - K)^\alpha$. As this is a bound on the individual bid distribution, it can also be used in the choice-based approach.

Maximizing the loss of bidding too high with respect to x leads to

$$\sigma\eta^{n-1}\frac{\beta^\beta}{(1+\beta)^{1+\beta}}\left(\frac{b_1}{\sigma}-K\right)^{1+\beta}, \quad (4)$$

where $\beta = \alpha(n-1)$.

Equating the maximal loss of bidding too low and the maximal loss of bidding too high, we obtain $\bar{b}_1^* := \rho \cdot v_1 + (1-\rho)\sigma K$, where

$$\rho = \frac{(1+\beta)^{1+\beta}}{(1+\beta)^{1+\beta} + \beta^\beta}.$$

The parameter ρ is strictly between $1/2$ and 1 and is increasing in β . Hence, the bid \bar{b}_1^* is a weighted combination of the lowest possible bid σK and the own value v_1 , where more weight is put on the own value. As \bar{b}_1^* increases in β , it also increases in n and α . Note that the computations above rely on the bound on the mass being a true constraint. Therefore, this computation is only correct if the bound on the mass for bidders bidding slightly above \bar{b}_1^*/σ is below 1 , so if $\eta(\bar{b}_1^*/\sigma - K)^\alpha \leq 1$. This holds if and only if

$$v_1 \leq \bar{v} := \sigma \left(K + \frac{1}{\rho\eta^{1/\alpha}} \right).$$

The bid \bar{b}_1^* attains minimax loss when $v \leq \bar{v}$, where this condition means that it is not conceivable for bidder 1 that all others almost certainly have a value below \bar{b}_1^*/σ . As this condition might not hold, we also investigate the minimax loss when $v > \bar{v}$. The next proposition summarizes our findings. The formal proof is in Online Appendix C. Let \hat{v} be a cutoff value defined as

$$\hat{v} = \sigma \left(K + \frac{2+\beta}{\beta\eta^{1/\alpha}} \right).$$

Proposition 5. *The unique minimax loss bid b_1^* when facing $\mathcal{B}_{\eta,K,\alpha,\sigma}$ is given by*

$$b_1^* = \begin{cases} \rho v_1 + (1-\rho)\sigma K & \text{if } \sigma K \leq v_1 \leq \bar{v} \\ b_1 \text{ s.t. } v_1 - b_1 = \eta^{n-1} \left(\frac{\beta}{\sigma}\right)^\beta \left(\frac{b_1 - \sigma K}{1+\beta}\right)^{1+\beta} & \text{if } \bar{v} < v_1 \leq \hat{v} \\ \frac{1}{2} \left(v_1 + \sigma K + \frac{\sigma}{\eta^{1/\alpha}} \right) & \text{if } \hat{v} < v_1. \end{cases} \quad (5)$$

The minimax loss bid b_1^ is continuous, increasing in v_1 , increasing in n for $v_1 \leq \hat{v}$, and independent of n for $v_1 > \hat{v}$. The corresponding value of the minimax loss is*

given by

$$l_1^* = \begin{cases} \eta^{n-1} \left(\frac{\rho}{\sigma}\right)^\beta (1-\rho)(v_1 - \sigma K)^{1+\beta} & \text{if } \sigma K \leq v_1 \leq \bar{v} \\ v_1 - b_1^* & \text{if } \bar{v} < v_1 \leq \hat{v} \\ \frac{1}{2} \left(v_1 - \sigma K - \frac{\sigma}{\eta^{1/\alpha}}\right) & \text{if } v_1 > \hat{v}. \end{cases} \quad (6)$$

For $v_1 > \bar{v}$, the loss of bidding $\bar{b}_1^* = \rho v_1 + (1-\rho)\sigma K$ is \bar{l}_1 , where

$$\bar{l}_1 = \begin{cases} \max \left\{ v - \bar{b}_1^*, \eta^{n-1} \left(\frac{\beta}{\sigma}\right)^\beta \left(\frac{\bar{b}_1^* - \sigma K}{1+\beta}\right)^{1+\beta} \right\} & \text{for } \bar{b}_1^* \leq \sigma K + \frac{\sigma(1+\beta)}{\beta\eta^{1/\alpha}} \\ \max \left\{ v_1 - \bar{b}_1^*, \bar{b}_1^* - \sigma K - \frac{\sigma}{\eta^{1/\alpha}} \right\} & \text{otherwise.} \end{cases} \quad (7)$$

Note the simplicity of the bid when $v_1 \leq \bar{v}$. In view of our objective to present simple rules, we show in (7) also how this rule performs when it does not attain minimax loss; that is, when $v_1 > \bar{v}$.

To illustrate the magnitude of the minimax loss, consider the following example. Consider $K = 0$, $\alpha = \sigma = 1/2$, $\eta = 2/(3\sqrt{v_1 - K})$ and five bidders. This η implies that bidder 1 conceives that any other bidder has a value above v_1 with a probability of at least $1/3$. The minimax loss is less than 8% of the maximal possible payoff $v_1 - \sigma K$. If, instead, $\eta = 3/(4\sqrt{v_1 - K})$, then the probability of any other bidder having a value above v_1 is at least $1/4$. As values (and correspondingly bids) below v_1 become more likely, the upper bound on minimax loss increases to about 12%. Note that in both cases $v_1 \leq \bar{v}$.

The minimax loss bid distinguishes different regions for v_1 , depending on where it lies relative to \bar{v} and \hat{v} . The cases are determined by whether or not the implied upper bound on the opponent bid distribution is binding in the worst-case bid distributions. For low values, the worst cases associated with bidding too low and too high involve bid distributions where the bound is tight. For high values, in both cases there are worst-case distributions with a single bid in their support. In the intermediate region, the worst-case distribution puts mass on at least two bids when bidding too low; when bidding too high, all the mass can be on one bid.

We discuss the role of the different parameters and then suggest how to choose them. The parameter α influences the bound on how much mass there can be on low values. The smaller α , the more mass can be placed there. The number of bidders connects how the bounds on individual bids influence the bound on the maximal bid. The more bidders there are, the more likely it is that there will be higher bids. The auction is thus more competitive if the likelihood of low values is lower and there are more bidders. In fact, the bound on the distribution of the

maximal bid depends on n and α through $\beta = (n - 1) \cdot \alpha$. In this sense, β is interpreted as a measure of the competitiveness of the auction. This is consistent with the observation that ρ , as a monotone transformation of β , determines how much weight is put on your own value when $v \leq \bar{v}$. Note that this is the case when the bounds always bind. For a better understanding of the magnitude of α , note that facing $n - 1$ bidders for a given α is as facing $\beta = (n - 1) \cdot \alpha$ bidders under a linear bound (where $\alpha = 1$). In this sense, α can be interpreted as a conversion rate between bounds on individual values and on the highest value among the other bidders. The case where the individual value distribution is uniform is useful for the comparison, as it is a common example in the literature and identifies the situation where the bound on the density of the values does not depend on the location.

The parameter σ gives a bound on how low a bid can be, relative to the value of this bidder. Thus, σ also captures some of the competitiveness of the auction, as a more competitive auction is expected to generate higher bids. Note that σ also influences the range of possible bids. In particular, no single bid will lie below $\sigma \cdot K$ and the highest possible bid lies above $\sigma \cdot K'$.

There are several ways to choose the free parameters. One possibility is to pre-commit to values that appear salient or that are motivated by theory. An alternative is to estimate certain key quantities and solve for the parameters that lead to these values. In this latter case, we suggest looking at the lowest possible maximal value K' and at the minimal probability of another bidder's value being above v_1 , given by the formula $1 - \eta(v_1 - K)^\alpha$. Computing these quantities is also valuable as a "sanity check" when pre-committing to the parameters. An additional way to select the free parameters based on worst-case analysis is presented in Online Appendix E.

To further illustrate the parameter choice, we provide more detail on our empirical implementation that applies to low-information environments. For such settings, we suggest using the bidding strategy \bar{b}_1^* , because it is simple and does not depend on η . The parameter η is only needed if it is necessary to understand the performance. Only parameters K , α , and σ need to be specified for bidding. The lower bound K can be chosen equal to the reserve price r when it is conceivable that bidders with a value above r also participate.⁴ We pre-committed to using $\alpha = 0.5$: the salience of this parameter influenced this choice. Choosing $\alpha = 0.5$ gives a square root bound on individual value distribution. Facing 11 other bidders with this bound is like facing $(11 - 1)/2 = 5$ bidders under a linear bound (where $\alpha = 1$).

⁴As bids above $\sigma \cdot v_1$ have not been ruled out, the set of conceivable environments contains situations where bidders bid below the reserve price, even though this is not feasible. We have not ruled out these environments in the interest of keeping the analysis simple.

However, in our empirical application we find that α in $[0.1, 0.2]$ would have been a better choice. The auctions were less competitive than we had anticipated. We also pre-committed to choosing $\sigma = 1/2$ in our empirical analysis, influenced by its salience and our analysis of maximal uncertainty.

Finally, we discuss how to implement our methodology when others may follow this paper’s methods. We recommend choosing σ such that $\sigma \cdot v \leq \rho \cdot v + (1 - \rho)K\sigma$. This ensures that the lower bound on bids is below the minimax loss bid. Hence, the beliefs are consistent with the selected bidding function (similar to Renou and Schlag (2010)). An alternative that is not pursued in this paper is to assume some commonly known lower bound on the proportion of other bidders who follow the model of variation in values. This would lead to an equilibrium approach.

7 Empirical Performance

We now investigate how the bids we have designed perform in auctions with real bidders. In this section, we consider timber auctions. In Online Appendix D, we look at auctions in laboratory experiments. We compare the performance of our bids to those actually made in the auction and to the BNE bids. Individual valuations in timber auctions are unknown and have to be estimated. Performance is measured by the loss relative to the bid distributions encountered in the data. To emphasize that loss is computed from the data, we refer to it as the *empirical loss*.

7.1 Connecting Empirical Loss to the Theory

We start by explaining what the theory outlined in Section 6 implies for the empirical performance of each of the three approaches. In each of them, the theory identifies the maximal values of loss for any data that is consistent with the ambiguity set. The difference between these bounds and the empirical loss is then due to the particular features of the data set. As all the specific features of the data set cannot be anticipated, the empirical loss will be lower than the performance bounds. To understand the gap between the empirical and the worst-case loss, we provide insights on how additional parameters improve the bound on performance. For maximal uncertainty and variation in bidding, these parameters are bounds on the density. For variation in values, this is the scaling parameter η . In each case, we estimate the respective parameters from the data and compute the adjusted bound on loss. The final gap between these adjusted bounds and empirical performance is then due to additional unanticipated features of the data or to the misspecification

of the model.

7.2 The Data

Our next step is to bring our bids to the data. In Online Appendix D, we consider the laboratory experiment of Güth and Ivanova-Stenzel (2003). In this section, we consider particular timber auctions that are popular in the empirical study of auctions (Athey et al., 2011).⁵ In these timber auctions, the stakes are high and the bidders can be considered to be experts. Some auctions are run in the English auction format, which allows us to estimate the bidders' values. The data contains information on the bids, the number and types of bidders (loggers or mills), together with some auction characteristics, such as the size of the lot and the type of wood. The number of bidders varies from two to 12. Loggers mainly differ from mills in that they tend to be smaller companies without processing facilities. This creates an observable asymmetry between the bidders. There is likely to be additional unobserved heterogeneity as these bidders frequently bid against each other. Dependencies in bidding behavior can arise due to the repeated nature and collusive behavior. Note that all our models, except for the model of variation in values, have been constructed to allow for some collusion.

7.3 The Empirical Methodology

For each bidder we identify a value, the parameters needed for the bid, and the bid distribution they face. Throughout the process, we maintain the assumption that the bidder applying our methodology is risk neutral. In Online Appendix F, we discuss some findings for risk aversion. We assume independent bidding and use, following Athey et al. (2011), a Weibull distribution to estimate the bid distributions conditional on the observed auction covariates.

The bidders' values are unobserved and hence need to be estimated. To do this, we use data from timber auctions for similar lots that are conducted as English auctions. Our identifying assumption is that bidders bid their own value in these English auctions. We use a Weibull specification to estimate the value distribution, conditional on the auction covariates in these related auctions. We then use the expected value of the value distribution, conditional on the first-price auction's covariates and the bidder's bid, as the value of the bidder. It is as if we are repeatedly drawing a random value from the estimated value distribution, conditional on the covariates and on the value being higher than the observed bid.

⁵Online Appendix D contains a more detailed discussion of the data.

We choose the following parameters for the bids of the three different approaches. In maximal uncertainty and variation in bidding, we set L equal to the reserve price r . For variation in bidding, we set $\tau = 0.15$ following our pre-commitment. For variation in values, we set K equal to the reserve price r and use $\alpha = \sigma = 0.5$ following our pre-commitment. The bid is set equal to the maximum of $\rho \cdot v + (1 - \rho) \cdot \sigma \cdot K$ and the reserve price K .⁶ Note that the parameter η does not enter the bid.

With the above information, we compute the empirical loss of our bids and some bounds on their loss. We compare this to the empirical loss of two benchmarks: the actual bids and the BNE bidding strategies. The estimated value distributions from the related English auctions are used to compute the BNE. The BNE benchmark is not provided in asymmetric auctions that contain both loggers and millers due to the lack of a closed-form solution. To better assess the differences in the performance across the different settings and data sources, we divide the loss of a bidder with value v by the maximal payoff that she can attain. According to the approach, the maximal payoff is either $v - L$ or $v - \sigma \cdot K$.

We also compare the empirical loss to the loss predicted by our theories using two different bounds. By ‘MAX’ we refer to the minimax loss as identified by our propositions when inserting the parameters used for the bids. This value cannot be calculated for variation in values as we have not pre-committed to a value for η . By ‘AM,’ or adjusted maximum, we refer to the value of maximal loss that also uses parameters that are estimated from the data. Specifically, for maximal uncertainty and variation in bidding, we determine α_1 and α_2 from the estimated bid distribution. For variation in values, we choose η as the smallest value that respects the bounds on the estimated value distribution.

7.4 The Results

We now show how our bidding models perform against actual bidders. Table 1 contains our findings. Entries are percentages, describing the empirical loss as a percentage of the maximal possible expected payoff. A blank entry indicates that this benchmark is not available. The first five columns refer to the auctions in which there were only loggers. We compute the BNE only for these auctions. The last two columns show the mean empirical loss separately for the loggers and the mills in the entire data set. The first row shows the performance of the bids made in the data. The next three rows show the performance of the bids resulting from

⁶The constraint imposed by the reserve price is never binding.

Table 1: Loss in timber auctions

	Auctions with loggers only					All auctions	
	Empirical			Theoretical		Loggers	Mills
	Mean	90 th %ile	Max	AM	Max	Mean	Mean
Observed Bid	6.61	16.15	43.51			5.16	4.61
MaxUn	4.17	9.45	18.21	10.40	50	3.83	3.93
VarBid	0.98	2.03	14.20	9.55	25.44	0.83	0.90
VarVal	4.15	11.52	28.20	26.26		3.32	3.95
Nash	1.20	3.09	10.76				
# of auctions	159					290	134
# of bids	705					1218	205

maximal uncertainty (‘MaxUn’), variation in bidding (‘VarBid’), and variation in values (‘VarVal’). The last row contains, where it is computed, the BNE bid. The first three columns show the mean, the 90th percentile, and the maximum of the empirical loss. The fourth column shows the adjusted maximum AM. The fifth column, Max, shows the minimax loss.

Looking at Table 1, the first and most striking observation is that bidding according to each of our three approaches leads to a good performance. This is true, both for the mean and for the tails. The mean empirical loss of our three approaches lies between 1 and 4.2%. The 90th percentile lies between 2 and 12%. The maximal empirical loss ranges from 14 to 28%. The model of variation in bidding does best in terms of mean, 90th percentile, and maximum loss; for example, its mean loss is just 1%.

Our next observation is that the actual bidders perform worse than any of our approaches, both at the mean and the tails. They attain a mean loss of around 7%. The 90th percentile is 16% and the maximal empirical loss of the actual bids is 44%. The actual bids in the timber auction are typically too low (see Figure D.1 in Online Appendix D). This is possibly due to bidders underestimating the competition. This conjecture emerges when observing that the bidders’ bids do not increase sufficiently when the number of bidders increases. The bidders would have done better if they had adopted one of our bid recommendations, such as variation in bidding.

Next, we analyze how much of the good performance demonstrated by following our rules can be traced to our methodology. For this, we consider the adjusted bounds, as shown in the column AM. For maximal uncertainty and variation in bid-

ding, the adjusted maximal loss falls below 10%. Hence, we conclude that variation in bidding, as captured by the lower and upper bounds on densities, can explain why the loss of our bids is below 10%. The differences in the AM bounds provide some idea as to why maximal uncertainty and variation in bidding perform similarly and better than variation in values.

Finally, we consider the performance of the Bayes-Nash equilibrium bidding strategies. The BNE bid performs in a similar way to variation in bidding. Indeed, Online Appendix D shows that the two bids are, on average, close to each other, although variation in bidding performs (slightly) better. Note, however, that the BNE bid is only a hypothetical benchmark. It uses information on the value distribution that is rarely available in practice. Despite its sophistication, extensive use of data, and computational complexity, the BNE bids do not outperform those under variation in bidding.

We investigate how the empirical loss changes in the number of bidders n in Online Appendix D. We find that the performance of the different models and benchmarks is similar when there are many bidders ($n \geq 10$). The exception is the bid of maximal uncertainty. Intuitively, when there are many bidders and the bid distribution exhibits some independence and variation, then any bid close to the value would perform similarly. The bid under maximal uncertainty is independent of the number of bidders, which explains the inferior performance.

In Online Appendix D, we investigate the sensitivity of the mean performance in the parameters τ and (α, σ) , respectively. In auctions in which all bidders are loggers, any $\tau \in [0.1, 0.2]$ would have performed similarly; the mean empirical loss is, at most, 1.5. In variation in values, the performance is relatively insensitive to σ , however, performance with $\alpha = 0.5$ is inferior to $\alpha \in [0.1, 0.2]$, which leads to an average loss in the range of 1 to 1.5.

In Online Appendix E, we present two alternative methods for selecting the free parameters under variation in bidding and variation in values. The first possibility is to follow a worst-case approach. Each parameter is evaluated according to how much it increases maximal loss compared to the value of minimax loss under an ideal parameter. We then choose the parameter that minimizes the maximum increase in the maximum loss across all contenders for being an ideal parameter. The second alternative is to use data from related auctions to estimate the parameters.⁷ We assess these two alternatives when there are, at most, five bidders, all of whom are loggers. When applying the worst-case approach to selecting τ from the set

⁷Note that best responding to the sample may be preferred to estimating parameters if the sample is sufficiently large and the sample generating process is understood.

[0.1, 0.2], we obtain a very similar performance to our baseline $\tau = 0.15$. The average empirical loss is 0.88. The performance is also similar when τ is selected from the set $[0.03, 1/3]$. We now summarize our findings for the empirically oriented method for selecting the parameters. To implement a low-information environment, we estimate the parameters based on small samples (60 bids) from different auctions without including their observed auction characteristics. Fixing G to be uniform, we obtain an estimate of τ equal to 0.24 on average. The corresponding mean loss is 1.91. Estimating G and fixing $\tau = 0.15$ leads to an average empirical loss of 1.23. When both τ and G are estimated from the data, the average empirical loss is 1.35. Using data leads to a slightly worse performance than the worst-case approach. These findings should be seen in light of the performance of the empirical best response to the sample, which yields a substantially higher average loss of 6.24. Why is this the case? When computing the empirical best response, the covariates are not taken into account, as we are in a low-information environment. Consequently, the empirical best response does not accommodate the auction heterogeneity. The same issues and comments apply when estimating τ and G . In contrast, the worst-case approach is designed to be able to deal with unknown heterogeneity: it is a robust design.

In Online Appendix F, we investigate the performance under two alternative ways of assigning values to bidders. The first alternative chooses values that make the observed bids optimal. In the second alternative, we assume that bidders make equilibrium bids and accommodate unobserved auction heterogeneity, as in Athey et al. (2011). The average loss of the variation in bidding bid is 1.19% and 1.53%, respectively. We also run other robustness checks for the laboratory experiment.

Based on our theoretical and empirical findings, we find that it is best to bid according to variation in bidding. The performance of the corresponding bid is close to best responding to the empirical bid distributions. Our particular choice of $\tau = 0.15$ is not critical, as a similar performance is obtained for any $\tau \in [0.1, 0.2]$. Its AM bounds suggest that this bidding rule will also perform well in other settings. Bidding under maximal uncertainty is a good alternative, as it relies on fewer assumptions and has no free parameters. It also has low AM bounds but performs slightly worse in our data sets. Its main disadvantage is that its bid does not adjust to the number of bidders and, hence, does not perform well when there are many bidders (see Figure D.1 in Online Appendix D). The performance of variation in values is similar. Choosing the right value for α is critical and benefits from a good understanding of the competitiveness of the auction. In our case, given no information about the auction, we find that our choice of $\alpha = 0.5$ was too high (see Figure

D.3 in Online Appendix D).

8 Conclusion

Our survey reveals that the classic game-theoretic tools are of little assistance when trying to determine which bid to make in real-life first-price sealed-bid auctions. Therefore, we have developed a new method on how to bid. Implementations of our method lead to bid recommendations with valuable theoretical performance guarantees and good empirical performance. Future work can extend and refine the method and investigate the performance in other data sets. Moreover, one can apply the methodology to other settings, as the inadequacy of theory for making explicit recommendations is hardly limited to first-price auctions. It may be worth considering procurement, dynamic, and multi-unit auctions, pricing, contract design or bilateral bargaining. In any case, the focus should be on bringing theories to the data and creating value for practitioners.

References

- Athey, S., J. Levin, and E. Seira (2011). Comparing open and sealed bid auctions: Evidence from timber auctions. *Quarterly Journal of Economics* 126(1).
- Azar, P., C. Daskalakis, S. Micali, and S. M. Weinberg (2012). Optimal and efficient parametric auctions. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- Azar, P. and S. Micali (2012). Optimal parametric auctions. *Technical Report MIT-CSAIL-TR- 2012-015*.
- Bajari, P. and A. Hortacısu (2005). Are structural estimates of auction models reasonable? evidence from experimental data. *Journal of Political Economy* 113(4), 703–741.
- Battigalli, P. and M. Siniscalchi (2003, October). Rationalizable bidding in first price auctions. *Games and Economic Behavior* 45(1), 38–72.
- Bergemann, D. and K. Schlag (2011). Robust monopoly pricing. *Journal of Economic Theory* 146(6), 2527–2543.
- Blum, A. and Y. Mansour (2007). From external to internal regret. *Journal of Machine Learning Research* 8, 1307–1324.

- Compte, O. and A. Postlewaite (2013). Simple auctions. *Working paper*.
- Crawford, V. P. and N. Iriberri (2007). Level-k auctions: Can boundedly rational strategic thinking explain the winner’s curse and overbidding in private-value auctions? *Econometrica* 75(6), 1721–1770.
- Filiz-Özbay, E. and E. Y. Özbay (2007, September). Auctions with anticipated regret: Theory and experiment. *American Economic Review* 97(4), 1407–1418.
- Giannakopoulos, Y., D. Poças, and A. Tsigonias-Dimitriadis (2020). Robust revenue maximization under minimal statistical information. In *16th Conference on Web and Internet Economics*.
- Gretschko, V. and H. Mass (2019). Endogenous worst-case beliefs in first-price auctions. *Working paper*.
- Guo, Y. and E. Shmaya (2019). Robust monopoly regulation. *Working paper*.
- Güth, W. and R. Ivanova-Stenzel (2003, August). Asymmetric auction experiments with(out) commonly known beliefs. *Economics Letters* 80(2), 195–199.
- Halpern, J. Y. and R. Pass (2012, January). Iterated regret minimization: A new solution concept. *Games and Economic Behavior* 74(1), 184–207.
- Han, Y., Z. Zhou, and T. Weissman (2020). Optimal no-regret learning in repeated first-price auctions. *arXiv arXiv:2003.09795*.
- Hayashi, T. (2008, March). Regret aversion and opportunity dependence. *Journal of Economic Theory* 139(1), 360–379.
- Hon-Snir, S., D. Monderer, and A. Sela (1998). A learning approach to auctions. *Journal of Economic Theory* 82(1), 65–88.
- Huber, P. J. (1965, December). A robust version of the probability ratio test. *The Annals of Mathematical Statistics* 36(6), 1753–1758.
- Linhard, P. B. and R. Radner (1989, June). Minimax-regret strategies for bargaining over several variables. *Journal of Economic Theory* 48(1), 152–178.
- Lo, K. C. (1998, July). Sealed bid auctions with uncertainty averse bidders. *Economic Theory* 12(1), 1–20.
- Loomes, G. and R. Sugden (1982, December). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal* 92(368), 805–824.

- Mass, H. (2018). Strategies under strategic uncertainty. *Working paper*.
- Milnor, J. (1954). Games against nature. In R. H. Thrall, C. H. Coombs, and R. L. Davis (Eds.), *Decision Processes*, Chapter 4. Wiley.
- Radner, R. (1980). Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives. *Journal of Economic Theory* 22(2), 136–154.
- Renou, L. and K. Schlag (2011, March). Implementation in minimax-regret equilibrium. *Games and Economic Behavior* 71(2), 527–533.
- Renou, L. and K. H. Schlag (2010). Minimax regret and strategic uncertainty. *Journal of Economic Theory* 145(1), 264–286.
- Rothkopf, M. H. (2007). Decision analysis: The right tool for auctions. *Decision Analysis* 4(3), 167–172.
- Savage, L. J. (1951, March). The theory of statistical decision. *Journal of the American Statistical Association* 46(253), 55–67.
- Schlag, K. and A. Zapechelnyuk (2016a). Robust sequential search. *Working paper*.
- Schlag, K. H. (2018). How to play out of equilibrium: Beating the average. *Working paper*.
- Schlag, K. H. and A. Zapechelnyuk (2016b). Value of information when searching for a secretary. *Working paper*.
- Sošić, G. (2007). Collusion in second-price auctions under minimax regret criterion. *Production and Operations Management* 16(4), 471–482.
- Stoye, J. (2011). Axioms for minimax regret choice correspondences. *Journal of Economic Theory* 146(6), 2226–2251.
- Sutton, R. S. and A. G. Barto (2018). *Reinforcement Learning* (2nd ed.). The MIT Press.
- Vickrey, W. (1961, March). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance* 16(1), 8–37.