

WORKING

PAPERS

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

Bernhard Kasberger
Karl H. Schlag

September 2017

Working Paper No: 1707



DEPARTMENT OF ECONOMICS

UNIVERSITY OF VIENNA

All our working papers are available at: <http://mailbox.univie.ac.at/papers.econ>

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

Bernhard Kasberger* Karl Schlag[◇]

September 28, 2017

Abstract

Bidding optimally in first-price auctions is complicated. In the classical equilibrium framework, optimal bidding relies on detailed beliefs about other bidders' value distributions and bidding functions. This article shows how to find a robust bidding rule that does well with minimal information and thus achieves good performance in many situations. Robust bidding means to minimize the maximal difference between the payoff and the payoff that could be achieved if one knew the other bidders' value distributions and bidding functions. We derive robust bidding rules under different scenarios, including complete uncertainty. Our bid recommendations are evaluated with experimental data.

Keywords– Auction, Robust, Uncertainty

JEL Codes: C72, D44, D81

1 Introduction

Consider a bidder in a first-price auction who knows her value for the auctioned good.¹ The bidder is aware of the basic trade-off. On the one hand, bidding close to value raises her chances of winning the auction, but it also leads to a

*We thank 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, and the 2017 VfS annual conference for their comments.

*Kasberger: Vienna Graduate School of Economics and Department of Economics, University of Vienna, bernhard.kasberger@univie.ac.at

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

¹In a first-price auction every bidder submits a sealed bid. The bidder with the highest bid wins the object and pays his or her bid. Ties are resolved randomly.

small payoff conditional on winning. Bidding low, on the other hand, leads to a large payoff conditional on winning, but goes along with a smaller likelihood of winning. The bidder is unsure about the probability that a bid becomes winning and wishes to hire an economic consultant to help her figure out the optimal bid. A consultant will ideally be guided by the current economic literature and hence will seek to compute a Bayesian Nash equilibrium. For this he needs to know the joint value distribution, the information structure, the number of bidders, other bidders' utility functions, their beliefs, and needs to assume common knowledge of all of these ingredients. He would then (try to) calculate an equilibrium. Our bidder has only a vague conception of these ingredients and concludes that she cannot hire such an economic consultant. In particular, she does not believe that everyone is playing according to equilibrium as she is unsure and others might be unsure too. Alternative methods to make recommendations exist but rely on the bidder knowing how others make their choices (rationalizability (e.g. Battigalli and Siniscalchi, 2003) and bidding under ambiguity (e.g. Lo, 1998)).

In this paper we provide a novel methodology that a consultant can use to help the bidder formalize her vague conception and place a bid in the first-price auction. The methodology requires minimal information about other bidders and even allows misunderstandings about the faced circumstances. Our general idea is the following. There is a true but unknown bid distribution. This bid distribution is generated by a joint value distribution, an information structure, and other bidders' bidding behavior. If the bid distribution was known, one could maximize the payoff; this generates the highest possible payoff. Our bid recommendations aim to get a payoff close to the highest possible payoff in the following way. Consider as the *loss* of a bid the difference between the maximized payoff if the bid distribution was known, and the payoff generated by the bid under the true bid distribution. Bidding behavior is chosen to minimize the maximal loss (the error bound) for the bidders' conceivable bid distributions. Note that loss is never experienced; it is a number that measures how well the highest possible payoff is approximated given the true but unknown bid distribution.

We suggest the consultant to formalize the bidder's vague conception with a set of conceivable environments. An environment generates a bid distribution and is defined as a joint value distribution, an information structure, and the bidding functions of the other bidders. Different perceptions lead to different sets of conceivable environments. In the text, we discuss conceptions that reach from complete uncertainty to bounds on the value distribution and others' bidding functions. Maximal loss is minimized with respect to the bidder's set of conceivable

environments.

We start by making a recommendation for the case of complete uncertainty, so when the bidder deems all value distributions and bidding functions possible. Other bidders might, for example, collude or have interdependent values. Our recommendation under complete uncertainty is to bid half of the own value. This bid insures that loss is at most half of the own value. In contrast, the loss of a slightly misspecified Nash equilibrium bidding function can be 100% of the own value. Nash equilibrium play is designed to perform optimally in a specific environment, but can perform poorly in another. Our bidding functions balance the risks of bidding too low and too high across all conceivable environments. One scenario we have to protect against, as any bids of others is deemed possible, is that all other bidders bid 0 and hence the robust bidder incurs a large loss due to bidding above 0. The bidder might be willing to rule out such extreme environments. We address this in two different ways. First, we consider restrictions on the joint bid distribution by imposing lower bounds and variability of bids. Later we assume others bid independently and limit the possible environments by bounding value distributions and bidding functions.

Our next recommendation applies when the bidder can assess a lower bound on the maximal bid of others. The bidder might be willing to rule out that the maximal bid of the others will be below some threshold $L < v$, where v is the private value of the bidder. For these environments, bidding $(v + L)/2$ guarantees a loss below $(v - L)/2$. This means, for example, that if one does not expect that the maximal bid is below 80% of one's own value, so $L = 0.8 \cdot v$, then the error bound will be at most 10% of v . However, a threshold L for which the bidder is willing to rule out with certainty that the maximal bid will not be below L may be very small. Hence, we also offer a bidding strategy for a bidder who can bound the probability of others bidding below L . We find that the above result for the case where the maximal bid has no mass below $L = 0.80 \cdot v$ continues to hold if at most mass 0.11 of the maximal bid is below this value of L .

The following recommendation can be used when the bidder is willing to rule out that others bid the same bid for sure. The bidding functions recommended so far do not depend on the number of bidders, as it is conceivable that the other bidders bid the same bid for sure. We introduce some degree of independence in a reduce form. There are n bidders and the maximal bid of others is not below L . We assume that independently for each bidder, with probability ε , this bidder is believed to bid between L and v with each bid being equally likely. The remaining bidders are believed to bid above L . The optimal bidding function of the ε -uniform

model is increasing in the number of bidders and a convex combination of the value v and L . The weight on v is at least one half. For instance, if ε is at least 0.15, there are 10 bidders, and $L = 0$, then bidding $0.88 \cdot v$ insures that loss is at most 13% of value v . The same minimax loss is attained by bidding $0.66 \cdot v$ when there are five bidders and $L = v/2$.

We finally derive recommendations for more refined conceptions in which the bidder is willing to place bounds on the joint value distribution and others' bidding behavior. Others' values and bids are conceived to be independent and the bounds rule out too much mass of others' bids close to L . There are upper bounds on the likelihood of others having low values as well as bounds on how low a bid can be for each possible value. We find, for instance, if there are five bidders, other bidders are believed to bid less than half of their value, i.e. $b(v) \geq v/2$, and the values are independently distributed according to F with $F(x) \leq \sqrt{x}$, then a robust bidder who believes that the own value is not in the upper 0.29 quantile of the value distribution can guarantee a loss below 10% of the own value by using a linear bidding strategy. In this example the optimal bid is $0.7 \cdot v$.

Our bid recommendations are designed to minimize the difference of the payoff and the highest possible payoff in the worst case. Supposedly, the worst case does not always occur. Hence, we want to evaluate our recommendations in actual situations. Experimental data provides a good setting for the evaluation as the private values are known. We data from three experiments. In all treatments of the three experiments, we find that our methods lead to much smaller loss on average than the actually submitted bids. The best recommendation across all treatments is from the ε -uniform model. The average loss of this recommendation tends to be around 20% of the actual loss. Depending on the treatment, knowing the empirical bid distributions allows to increase the payoff by 0.6% to 3.8% of the own value on average. Our methods tend to outperform Nash equilibrium play in terms of average and maximal loss.

The outline of the paper is as follows. We next discuss the related literature. The formal methodology used for this study is described in Section 2. Section 3 presents simple examples of the non-robustness of bidding functions that are optimal in specific environments under expected and maximin expected utility. The core of the paper begins in Section 4 where we derive deterministic bid recommendations for different restrictions on the possible bid distributions. In Section 5 we derive bids under restrictions on value distributions and bidding functions. In Section 6 we evaluate our bid recommendations with experimental data. Section 7 concludes. Proofs are given in Appendix A. Appendix B shows that appro-

priate randomization can further decrease the minimax loss. For example, the error bound can be improved by appropriate randomization to 36% of $(v - L)$ under complete uncertainty. Appendix C supplements the empirical section. In Appendix D we consider risk averse bidders.

Related Literature

First-price auctions are prevalently analyzed by deriving Bayesian Nash equilibria (e.g. Vickrey, 1961).² As this practice requires many common knowledge assumptions, it is subject to the critique of Wilson (1987) who suggests to successively weaken common knowledge assumptions in game theory.³ Our work dispenses with all common knowledge assumptions and is therefore related to studies that weaken these assumptions. Common with the literature on rationalizability is dispensing with equilibrium play (e.g. Battigalli and Siniscalchi, 2003; Dekel and Wolinsky, 2003; Cho, 2005; Robles and Shimoji, 2012). In contrast to this approach, we neither assume a commonly known value distribution, nor common knowledge of rationality. Common with the literature on ambiguity is uncertainty about the value distribution (e.g. Lo, 1998; Levin and Ozdenoren, 2004; Chen et al., 2007).⁴ This approach, however, assumes symmetric ambiguity preferences, a common prior over priors, and equilibrium play, so that behavior is common knowledge up to private information.

Compte and Postlewaite (2013) suggest to reduce the sophistication needed in fine tuning bids to the environment by restricting attention to a specific functional form. Thus, the bidding functions are simple by assumption while in our paper they are simple as result of the optimization. Most importantly, they assume common knowledge of the restriction. In our approach, we do not impose that all bidders think about the problem in the same way. We only consult one bidder and try to put as little restrictions on what others do as possible.

²Experimental research seems to have rejected the hypothesis of risk-neutral equilibrium play in first-price auctions. There does not seem to be one accepted model that explains observed behavior in the lab; risk aversion, behavioral motives, and confusion have been suggested as explanations (e.g. Kagel and Levin, 2016). In principle, our theory can be seen as an alternative to existing models to rationalize observed behavior. The primary objective of the paper is, however, deriving bid recommendations. Interestingly, Harrison (1989) suggests considering the payoff space and not the bid space in evaluating experimental bids. Using a certain notion of loss and equilibrium beliefs, he concludes that observed bids in experiments might be close to optimal bids in the payoff space.

³Rothkopf (2007) informally favors decision theoretic methods over game theoretic models for deriving bid recommendations, but has no explicit suggestion of how to address bidding as a decision problem.

⁴Interestingly, experimental results by Güth and Ivanova-Stenzel (2003) and Chen et al. (2007) indicate that bidding behavior is very similar with and without a known value distribution.

Our approach is minimizing the maximal loss; a concept introduced by Savage (1951) for decision problems that was subsequently used in the literature on minimax regret and robust statistics. To start with the latter, Huber (1965, 1981) introduces a loss function to derive robust test statistics in slightly misspecified environments. In this paper, we do not restrict ourselves to slightly misspecified environments, but consider more arbitrary sets of conceivable environments, that is, we look at globally robust bidding functions. Note that the term “robust” has also been used differently in denoting expected maximin utility (e.g. Carroll, 2015). The literature on minimax regret (e.g. Milnor, 1954; Hayashi, 2008) has, for example, looked at the news-vendor problem under partial information (Perakis and Roels, 2008), the pricing problem of a monopolist (Bergemann and Schlag, 2008, 2011; Caldentey et al., 2017), and dynamically consistent robust search rules (Schlag and Zapechelnyuk, 2016). In strategic settings, Linhard and Radner (1989) consider bargaining and Sošić (2007) develops a specific collusive scheme in auctions. Renou and Schlag (2011) and Halpern and Pass (2012) develop solution concepts for games where it is common knowledge all players follow minimax regret.

A methodological innovation is that we fully exploit the advantages of *ex-ante* over *ex-post* loss. The difference between ex-ante and ex-post loss can be easily seen in the auction context. From an ex-ante perspective environments are bid distributions, whereas from an ex-post perspective environments are realized bids. Put differently, ex-ante loss occurs because one does not know the bid distribution, but ex-post loss occurs because the realized bids are unknown. While some studies have also considered ex-ante loss (e.g. Perakis and Roels, 2008; Jiang et al., 2011; Schlag and Zapechelnyuk, 2016), we explicitly vary the information available.

We talk about loss and not about regret because, first, our evaluation of performance has no behavioral context as the term “regret” might suggest, and second, because the term regret is used differently by the literature that follows Loomes and Sugden (1982). This approach has been applied to auctions (e.g. Engelbrecht-Wiggans, 1989; Filiz-Özbay and Özbay, 2007; Engelbrecht-Wiggans and Katok, 2008), assumes a commonly known value distribution, and deals with anticipated behavioral effects of learning other bidders’ bids. We do not have a behavioral motive in mind.

2 Methodology

There are n bidders who participate in a first-price sealed bid auction for an indivisible good. We consider the bidding behavior of a single bidder among them. Let 1 be the index of this bidder. Bidder 1 is risk neutral with a utility function quasilinear in her bid.⁵ She knows that winning the good with bid b yields her utility $v_1 - b$ and losing the auction gives her utility 0. Let b_1 be her bidding function, so $b_1 : \mathbb{R}_+ \rightarrow \Delta\mathbb{R}_+$ maps the own value v_1 into (a distribution of) bids, as $\Delta\mathbb{R}_+$ denotes the set of probability distributions over positive reals.

When choosing their bids, bidders in first-price auctions are interested in the bid distribution of the other bidders. The bid distribution is generated from three inputs. First, there is a true and exogenous joint value distribution $F \in \Delta\mathbb{R}_+^n$. Second, there is an information structure that maps the profile of values into the information available to the bidders. Third, each participant's bidding behavior translates the information into bids. Traditional game theoretic analysis of bidding behavior in auctions assumes common knowledge of the value distribution, information structure, and bidders' preferences and searches for an equilibrium in which each bidder best responds to the bidding behavior of the other bidders. In equilibrium every bidder knows how the other bidders behavior up to private information and therefore knows the bid distribution.

We depart from the classic game theoretic setting and analyze the auction as a decision problem. We investigate how bidder 1 should bid if she is uncertain about the value distribution, the information structure, and the bidding behavior of the others. No assumptions about the utility functions of the other bidders are made. Bidder 1 might not know the true joint value distribution F , but conceive that F belongs to the class of joint distributions \mathcal{F} with $\mathcal{F} \subseteq \bigcup_{n \in \mathbb{N}} \Delta\mathbb{R}_+^n$. One can incorporate uncertainty over the number of bidders n by including distributions with different numbers of bidders in \mathcal{F} . The bidder might be uncertain about the bidding functions of the other bidders. They might bid independently, collude, communicate, or know not only their own value. We abstract from the information structure and directly specify which bidding functions are deemed possible. The set B_F collects the profiles of other bidders' bidding functions $b_{-1} = (b_2, \dots, b_n) \in B_F$ that bidder 1 conceives under the joint distribution F . Any bidder $i > 1$ uses the bidding function b_i that maps the available information into a distribution of bids.

The overall uncertainty is modeled in the form of bidder 1 identifying a set

⁵In Appendix D we consider risk-averse bidders.

of conceivable environments \mathcal{E} . The set of conceivable environments combines the uncertainty about the joint value distribution and the bidding functions, i.e. $\mathcal{E} = \bigcup_{F \in \mathcal{F}} \{F\} \times B_F$. An environment $E \in \mathcal{E}$ generates the bid distribution faced by bidder 1 and is a pair (F, b_{-1}) , where F is a joint value distribution and $b_{-1} = (b_2, \dots, b_n)$ specifies the bidding behavior of other bidders. Two extreme cases of conceivable environments are complete uncertainty and Nash equilibrium. Under complete uncertainty the set of conceivable environments is the universe of all possible value distributions and bidding functions $\mathcal{U} = \bigcup_{F \in \mathcal{F}} \{F\} \times B_F$, where $\mathcal{F} = \bigcup_{n \in \mathbb{N}} \Delta \mathbb{R}_+^n$ is the set of all joint distributions and B_F the set of all positive functions b_{-1} , where n given by F . Another extreme case is Nash equilibrium. Simply set $\mathcal{E} = \{(F, b_{-1}^*)\}$ where b^* is a Bayesian Nash equilibrium under joint value distribution F . Similarly, one can model the environment in which some but not all other bidders choose Bayesian Nash bidding strategies.

Bidder 1 ideally selects the best bid given the true environment. We consider, however, a bidder who does not know the environment and hence cannot perform this task (i.e. $|\mathcal{E}| > 1$). In the following we present our model of how bidder 1 bids without having a subjective probabilistic belief. The performance of a bidding function in a given environment is measured using a loss function. The loss of bidder 1 conditional on her value v_1 is defined as the difference between what she could get if she knew the environment what she gets. Formally, loss is given by

$$l(b_1, F, b_{-1} | v_1) = \sup_y \{(v_1 - y)Q(y, b_{-1}, F)\} - \int (v_1 - x)Q(x, b_{-1}, F) db_1(x),$$

where Q is the probability that bidder 1 wins the object when bidder i uses bidding function b_i and values are drawn according to F . Note that $\sup_y \{(v_1 - y)Q(y, b_{-1}, F)\}$ describes the payoff bidder 1 could (approximately) achieve if she knew F and the bidding behavior of the others. We call this purely hypothetical situation in which the environment is known the *oracle*. In general loss is zero if the optimal bidding function for the true environment is chosen and bounded above by v_1 if no bids above value are placed.

What remains to be specified is the description of how bidder 1 solves the problem of deciding how to bid without knowing the environment. A bidding function is evaluated by the maximal loss it can generate among the conceivable environments \mathcal{E} . For a given set of conceivable environments \mathcal{E} , bidder 1 prefers bidding functions that generate smaller maximal loss. The best bidding function bidder 1 can choose according to this criterion is the one at which minimax loss is attained. Minimax loss is defined relative to a set of conceivable environments

\mathcal{E} . We distinguish minimax loss and deterministic minimax loss. In the former all bidding strategies can be used to minimize loss, whereas in the latter only deterministic (pure) strategies are allowed.

Definition 0. Call ξ the value of *minimax loss* for the conceivable environments \mathcal{E} if, for all environments in \mathcal{E} , (i) loss is guaranteed to be at most ξ , and (ii) there is no bidding function that guarantees a loss strictly lower than ξ .

Call ξ the value of *deterministic minimax loss* for the conceivable environments \mathcal{E} if, for all environments in \mathcal{E} , (i) there exists a deterministic bidding function that guarantees loss to be at most ξ , and (ii) there is no deterministic bidding function that guarantees a loss strictly lower than ξ .

In the main part of the article we consider deterministic minimax loss. In Appendix B we show that minimax loss can be further decreased with appropriate mixed strategies.

3 Sensitivity of Expected and Maximin Utility

In this section we provide two simple examples that illustrate potential loss of game theoretic solution concepts due to slightly misspecified environments. The first example illustrates loss in the standard Bayesian Nash equilibrium framework, while the second example deals with equilibrium among maximin expected utility maximizers. Maximal loss is (approximately) 100% of the value in both examples. This is the largest possible maximal loss if one does not bid above value. In both examples there are two risk-neutral bidders participating in a first-price auction. Their respective values are drawn independently from a parameterized value distribution. Loss will depend on the parameter of the value distribution. The two examples differ in how the other bidder's conceived bidding function depends on the parameter. In the first example bidder 1 conceives that the other bidder's bidding function depends on the parameter, so it is as if the other bidder behaves as if the true parameter was common knowledge. In the second example the other bidder is conceived to use a bidding function independent of the parameter. This might occur if there is a common understanding about behavior, but not about the value distribution.

In the first example any conceivable value distribution distributes mass $\epsilon \in (0, 1)$ uniformly on $(\delta, 1]$ and mass $1 - \epsilon$ on δ , so $F^\delta(x) = 0$ for $x < \delta$ and $F^\delta(x) = \min \left\{ 1 - \epsilon + \epsilon \frac{x - \delta}{1 - \delta}, 1 \right\}$ for $x \geq \delta$. Bidder 1 expects the other person to

play the Bayesian Nash equilibrium strategy

$$b^\delta(x) = x - \frac{\int_\delta^x F^\delta(\tilde{x}) d\tilde{x}}{F^\delta(x)} = \frac{x^2\epsilon - \delta^2(2 - \epsilon) + 2\delta(1 - \epsilon)}{2(1 - \delta + x\epsilon - \epsilon)}$$

for $x \in [\delta, 1]$. Note bidder 2 has value δ with probability $1 - \epsilon$, in which case she bids her value. The set of conceivable environments is given by $\mathcal{E} = \{(F^\delta, b^\delta) | \delta \in [0, 1]\}$. In a standard textbook model bidder 1 has a subjective probabilistic belief that a certain environment (F^δ, b^δ) happens with probability 1. Consider bidder 1 with type $v_1 = 1$ and the belief that (F^0, b^0) occurs with probability 1. Her optimal bid is $b^0(1) = \epsilon/2$. This bid is never winning if the true $\delta > \epsilon/2$. The loss of the bid $b^0(1) = \epsilon/2$ is maximized if the true δ is slightly above $\epsilon/2$ and equal to

$$\sup_{\delta > \epsilon/2} l(b_1, F^\delta, b_2^\delta | v) = \sup_{\delta > \epsilon/2} 1 - b^\delta(1) = \sup_{\delta > \epsilon/2} \frac{(1 - \delta)(2 - \epsilon)}{2} = \frac{(2 - \epsilon)^2}{4}.$$

For ϵ close to 0, the maximal loss is approximately equal to the value. To summarize, the Bayesian Nash equilibrium performs optimally for the true value distribution, but it is very sensitive to slight chances in the environment.

In the second example any of the conceivable value distributions distributes mass γ uniformly on $[0, 1 - \gamma)$ and mass $1 - \gamma$ uniformly on $[1 - \gamma, 1]$. Let $0 < \gamma_1 < \gamma_2 < 1$, and $\mathcal{F} = \{F^\gamma | \gamma \in [\gamma_1, \gamma_2]\}$, where $F^\gamma(x) = \gamma x / (1 - \gamma)$ for $0 \leq x \leq 1 - \gamma$ and $F^\gamma(x) = (2\gamma - \gamma x + x - 1) / \gamma$ for $1 - \gamma < x \leq 1$. The set of conceivable environments is $\mathcal{E} = \{(F^\gamma, b^{\min}) | \gamma \in [\gamma_1, \gamma_2]\}$, where

$$b^{\min}(v) = v - \frac{\int_0^v F^{\gamma_1}(x) dx}{F^{\gamma_1}(v)} = \begin{cases} \frac{v}{2} & \text{for } v \in [0, 1 - \gamma_1) \\ \frac{(1 - \gamma_1)(v^2 + 2\gamma_1 - 1)}{4\gamma_1 + 2v(1 - \gamma_1)v - 2} & \text{for } v \in [1 - \gamma_1, 1] \end{cases}.$$

This behavior is consistent with the equilibrium model of maximin expected utility maximizers. Lo (1998) shows that bidders with identical \mathcal{F} and maximin preferences select the worst-case prior F^{\min} as the lower envelope of conceivable value distributions in \mathcal{F} . In our example this corresponds to $F^{\min} = F^{\gamma_1}$. Subsequently, bidders behave as if F^{\min} is the true value distribution and strategic uncertainty is resolved in equilibrium, that is the bidding function is used by both players. Consider bidder 1 with value 1 who bids $b^{\min}(1) = 1 - \gamma_1$. Let γ_1 be close to 0 and the true $\gamma = \gamma_2$ and γ_2 close to 1. In this case most types are close to 0 and bidder 1 should bid very low to maximize her payoff. Bidder 1 bids, however, almost 1 as γ_1 is close to 0. Consequently, her loss of not bidding as she would if she knew the true γ can be almost 100% of the own value.

In both equilibrium models the loss can be as high as 100% of the value. Although equilibrium play depends on many restrictive assumptions, these assumptions are not the source of the high loss in our examples. The loss comes from the underlying decision theory. A bidder with a subjective probability assessment of the highest bid among the other bidders might bid very well if the assessment is accurate, but very badly if it is slightly wrong. Likewise, bidder 1 with maximin expected utility preferences might bid value, because bidding value is a best response to the worst-case in which all other bidders bid bidder 1's value. Bidding value is not a sensible bid recommendation in a first-price auction as it leads to zero expected utility for sure.

4 Conceiving Bid Distributions

We come to the main part of the paper. A bidder who is uncertain about the value distribution, the information structure, and other's bidding functions is in fact uncertain about the bid distribution. In this section we consider a bidder who narrows down the possible environments by putting restrictions directly on the conceivable bid distributions. First, she imposes a lower bound on the possible maximal bid, then she allows for some mass below this threshold, and finally, she considers a model with some independent bidding and bid dispersion. In Section 5 the robust bidder thinks more explicitly about value distribution and other bidders' behavior.

4.1 Imposing a Lower Bound on the Maximal Bid of Others

We start by looking at the case in which bidder 1 is completely uncertain about the other bidders' types (value distributions, risk preferences, higher-order beliefs, etc.) and their bidding behavior. This means, for example, that bidder 1 does not insist that the other bidders bid independently, but also deems colluding behavior possible. More formally, we allow the set of conceivable environments to be the set of all possible environments \mathcal{U} .

The model of complete uncertainty can readily be generalized to the following situation. Bidder 1 believes that she needs to bid at least $L \geq 0$ so that her bid becomes winning. The value of L can be a known reserve price,⁶ or the perception that the maximal bid of other bidders is at least L . In the following we consider the case where $v_1 > L$ as bidding under $v_1 \leq L$ is simple; all bids less than or

⁶The value L can also be the endogenously determined reserve price in a unit-demand Anglo-Dutch auction (Binmore and Klemperer, 2002).

equal to v_1 are optimal. The perception that the other bidders' highest bid is at least L is a restriction on the set of all possible environments.

Definition 1. Let \mathcal{E}_L be the set of environments belonging to \mathcal{U} in which the highest bid of other bidders is almost surely at least L for $L \geq 0$.

The objective is to find the deterministic bid b^* that minimizes the maximal loss for the set of conceivable environments \mathcal{E}_L . Loss associated with bid b is the difference between the oracle payoff, the maximized payoff if the environment was known, and the payoff generated by bid b . First, observe that loss cannot be minimized by bids above v_1 , because these bid results in a non-positive payoff.⁷ Bids below L are losing for sure and cannot minimize loss. Loss cannot be maximized by environments in which the highest bid among the other bidders M is above v , because then loss is zero. Therefore, loss can only be maximized by environments with $M \in [L, v]$. Note that for a given environment E bidder 1 wins only if her bid is higher than M . Hence, loss is maximized by environments in which M is revealed in the oracle, because in these environments the oracle payoff is as high as possible. A simple type of environment in which this is the case is when all other bidders bid M with certainty. If bidder 1 bids b and the oracle reveals M , then loss equals

$$l(b, M|v) = \sup_{x > M} \{v - x\} - \mathbb{1}_{b > M}(v - b) = v - M - \mathbb{1}_{b > M}(v - b),$$

where $\mathbb{1}_{b > M} = 1$ if $b > M$ and 0 otherwise. The oracle payoff is $\sup_{x > M} \{v - x\}$, because bidder 1 knows the bid M she has to match.⁸

Loss can come from bidding too low and from bidding too high. Bidder 1 bids too low if the bid b does not become winning. In this case, loss is $v - M$, but not higher than $v - b$, as loss is maximized when b is slightly outbid. The bid b is too high when $b > M \geq L$, so when bidder 1 could raise her payoff by decreasing her bid. The loss of bidding too high is not more than $v - L - v + b = b - L$. Thus, the maximal loss is $\max \{v - b, b - L\}$. Maximal loss is minimized by the bid that balances the loss from bidding too low and too high, so by the bid that equates the two expressions. The proposition gives the resulting deterministic minimax bid and the corresponding loss. Minimax loss relative to the distance $v - L$ is $\frac{1}{2}$.

⁷We will often drop the index if we think this causes no confusion.

⁸Formally, without discrete bids and with a non-degenerate tie-breaking rule, bidder 1 has no best response when he knows that the maximal bid among the other bidders is $M < v$. Hence, we consider the supremum because we are interested in the payoff and not in the specific bid.

Proposition 1. *For the set of conceivable environments \mathcal{E}_L deterministic minimax loss is equal to $\frac{v-L}{2}$ and attained by bidding*

$$b^*(v) = \frac{v}{2} + \frac{L}{2}. \quad (1)$$

The bidding function in Proposition 1 is independent of the number of bidders. Note that no assumption on the number of bidders is made. Even if this number was known, the true value distributions could assign the same value to all other bidders, or the bidding function could be such that all submit the same bid, making the number of bidders irrelevant.

Bergemann and Schlag (2008) look at a related problem—the optimal pricing scheme of a monopolist who does not know the value distribution of the buyer. The monopolist minimizes maximal regret (loss), where regret is the difference in profit when the value is known and when it is not known. It turns out that their solution of the monopolist’s optimal pricing strategy resembles our solution. Bidding in a first-price auction with no assumptions is like pricing in markets with no information on demand. Apart from directions where higher payoffs can be achieved, in auctions one wishes as bidder to have a low winning bid, in markets as seller a high sale prices. A methodological difference is that Bergemann and Schlag (2008) consider ex-post loss, while we consider ex-ante loss. The difference of those two concepts is decision maker’s knowledge used for the computation of the oracle payoff. In Bergemann and Schlag (2008) the monopolist knows the strategy of the potential buyer and uses the buyer’s value in the oracle. In this article the bidder uses the distribution and bidding function of the other bidders only when computing the oracle payoff. Halpern and Pass (2012) introduce iterated elimination of strategies that do not attain minimax regret in normal form games (with a known prior). For this approach it is crucial that all players are known to minimize maximal regret. They provide a simple example of a first-price auction in which they essentially look at ex-post minimax regret, limit attention to deterministic strategies, and iteration is not needed. They find that the bidding function $b(v) = v/2$ minimizes maximal regret. We do not assume that all bidders minimize loss and we consider ex-ante and not ex-post loss. The deterministic bidding rule $b(v) = v/2$ was also found by Sošić (2007).

Loss can come from bidding too low and from bidding too high. Loss from bidding too low is maximized when the mass point of the bid distribution is just marginally larger than the own bid. Loss from bidding too high, however, is maximized by a bid distribution that puts all the mass on the lowest possible bid

L . There are certainly situations in which such extreme situations are deemed implausible. In the following sections we show that minimax loss is smaller if very low bids are not conceived to be likely.

4.2 Allowing some Mass Below the Threshold

In the analysis above bidder 1 restricted her bids to be above L , because she deemed that her bids below L are never winning. When the likelihood of bids below L is sufficiently small, we show that it is best to ignore possible bids below L and to bid as in Proposition 1.⁹

Definition 2. Let $L \geq 0$ and $\bar{p} \in [0, 1]$. Define $\mathcal{E}_{L\bar{p}}$ to be the set of all environments such that the probability that the highest bid among the other bidders is below L is bounded above by \bar{p} .

The maximal probability that the highest bid among the other bidders is below L is \bar{p} for environments in $\mathcal{E}_{L\bar{p}}$. Each environment specifies a number of bidders n with $n \geq 2$. Suppose the other bidders bid independently. For every bidder $i > 1$ there is a $p_i \in [0, 1]$ such that at most mass p_i of i 's bids can be below L . Then the maximal probability that the highest bid among the other bidders is below L is $\prod_{1 < i \leq n} p_i \leq \bar{p}$. Moreover, if p_i does not depend on i , then $p^{n-1} \leq \bar{p}$. In the analysis above we had $\bar{p} = 0$. In this section the maximal bid M can be in $[L, 1]$ with probability 1, but maximal bids below L can only be induced by distributions in which the highest bid among the other bidders is above L with probability at least $1 - \bar{p}$.

Consider bidder 1 having a relatively high value $v > L$ and suppose she uses the bidding strategy $b^*(v) = (v + L)/2$ of Proposition 1. Above we saw that if the highest bid among the other bidders is always above L , loss is at most $(v - L)/2$. Therefore, loss of not bidding below L can only be made larger if the highest bid among the others is below L . Potentially, loss can be made largest by all other bidders bidding 0, which can, under the set of conceivable environments $\mathcal{E}_{L\bar{p}}$, only happen with probability \bar{p} . This insight is associated with a loss that depends on \bar{p} . The following proposition states that if \bar{p} is sufficiently small, then maximal loss is minimized by ignoring potential bids below L .

Proposition 2. *Let $v > L > 0$ and $0 < \bar{p} < 1$. For the set of conceivable environments $\mathcal{E}_{L\bar{p}}$ with $\bar{p} \leq \frac{v-L}{v+L}$ deterministic minimax loss is equal to $\frac{v-L}{2}$ and attained by the deterministic bidding strategy stated in Proposition 1.*

⁹In Appendix B we discuss a related setting in which one knows that the own value is relatively small and the implications on loss.

The proof is in Appendix A. Note that for any $L \in (0, v)$ and $p_i \in (0, 1)$, the upper bound on \bar{p} in Proposition 2 is satisfied for large enough n . In order to get a feeling for the result, in Example 1 we fix the bound on loss to be at most 10% of v and ask which L and \bar{p} give rise to this loss.

Example 1. Let $L = 0.8v$. Proposition 2 implies that one does not need to bid below L to minimize maximal loss if $\bar{p} \leq \frac{v-L}{v+L} = \frac{0.2}{1.8} = 0.11$. In this case minimax loss is equal to 10% of the value. Assuming independent bidding, so $\bar{p} = p^{n-1}$, loss is less than one tenth of v when $p \leq 0.57$ and $n = 5$.

The proposition illustrates the difference between an ex-ante and an ex-post perspective. In an ex-post perspective the oracle reveals the bids of the other bidders. If it happens that all bids are below L , the optimal bid is below L . One can incorporate constraints on the bid distribution in the ex-ante approach. The oracle reveals the bid distribution, but not the specific bids, and therefore it is optimal not to bid below L . In the next model the ex-ante approach is further developed by assuming a certain independence of other bidders' behavior.

4.3 The ε -Uniform Model

We now return to our original model in which the bidder believes that all bids are above L . In Subsection 4.1 it was conceivable that all other bidders bid the same bid. This had the consequence that the optimal bidding function was independent of the number of bidders. Here we assume that the bidder expects a certain number of bidders and some heterogeneity among the other bidders. We model this in a reduced form by assuming that the bidder believes that any other given bidder puts a minimal weight of ε on bids above L . Thus, no relevant bid can be ruled out and, in particular, it cannot be the true environment that all other bidders bid some bid M for sure. Formally, bidder 1 conceives that the bid distribution of a given bidder can be written as ε times the uniform distribution on $[L, v]$ plus $1 - \varepsilon$ times some arbitrary distribution.¹⁰

Definition 3. Let $L \geq 0$ and $\varepsilon \in (0, 1)$. Define $\mathcal{E}_{L, \varepsilon, n}$ to be the set of all environments belonging to \mathcal{E}_L such that (i) there are n bidders, (ii) for any bidder $i > 1$ it is as if the bid is independently drawn uniformly from the interval $[L, v]$ with probability ε .

¹⁰We assume the uniform distribution for simplicity. With the uniform distribution one gets quite far in terms of closed form solutions. It might be that one has to rely entirely on numerical calculations for other continuous distributions.

For $\mathcal{E}_{L,\varepsilon,n}$ it is again a simple form of environment that potentially maximizes loss. These environments generate bid distributions such that for every bidder $i > 1$ the bid is drawn uniformly from $[L, v]$ with probability ε and equal to $M \in [L, v]$ with probability $1 - \varepsilon$. It is enough to restrict bids to the interval $[L, v]$, as loss is made smaller if bids are above value with positive probability. In these simple environments it is as if bidder 1 learns the highest bid M among the other bidders whose bid is not drawn uniformly in the oracle. The bid distribution has a mass point at M . If $b > M$, then bidder 1 wins against the $n - 1 - k$ bidders who bid M , so the payoff is equal to

$$\pi(b) = \sum_{k=0}^{n-1} \binom{n-1}{k} \varepsilon^k (1 - \varepsilon)^{n-1-k} (v - b)p(b)^k, \quad (2)$$

with $p(b) = (b - L)/(v - L)$ being the probability that bid b is larger than a uniformly drawn bid. This function is maximized by bidding $\tilde{b} = \frac{L + \varepsilon n v - v}{\varepsilon n}$. Note that $\tilde{b} > L$ if and only if $\varepsilon > 1/n$. In principle, the oracle bid recommendation can be above, equal to, or below M . Bidding M is never optimal, because slightly bidding above M avoids the tie-breaking. If bidder 1 bids below M , then bidder 1 only wins if all other bidders' bids are drawn uniformly. It turns out that loss is not maximized by environments in which $M < \tilde{b}$, or in which M is so large such that one wants bid below M in the oracle.

Loss can come again from bidding too low and from bidding too high. Bidding too low means that the bid b is below the mass point M . The difference to complete uncertainty is that here bidder 1 might win against the k bidders who bid uniformly. The bidder bids too high when $b > M$ so that a lower bid would also have been higher than the mass point of the bid distribution. The minimax bid balances the maximal loss from bidding too low and the maximal loss from bidding too high. It depends on ε , the number of bidders n , and the lowest possible winning bid L . A closed form solution is not available—it needs to be computed numerically. As a result, also the value of minimax loss can only be stated implicitly. As ε tends to zero, the model and the results converge to the previously stated $1/2$ bound.

Proposition 3. *Let there be $n \geq 2$ bidders, $v > L \geq 0$, and $\varepsilon \in (0, 1)$. For the set of conceivable environments $\mathcal{E}_{L,\varepsilon,n}$ deterministic minimax loss is attained by b^* such that*

$$\pi(\max\{L, \tilde{b}\}) - \pi(b^*) = \sum_{k=0}^{n-2} \binom{n-1}{k} \varepsilon^k (1 - \varepsilon)^{n-1-k} (v - b^*)p(b^*)^k \quad (3)$$

and equal to the value on either side of the equation, where $\tilde{b} = \frac{L + \varepsilon n v - v}{\varepsilon n}$.

The left-hand side of Equation (3) is the maximal loss of bidding too high. The bid generates a payoff of $\pi(b)$, but a higher payoff could be achieved by bidding $\max\{L, \tilde{b}\}$. The right-hand side is the maximal loss of bidding too low. In this case the bid can only be winning if all other bidders' bid is drawn uniformly. The maximal loss is when bidder 1 is just slightly outbid.

To get a feeling for the magnitude of minimax loss, we look for a numerical approximation when we vary the number of bidders. Hence, we fix $\varepsilon = 0.15$ and consider $n \leq 10$. Two cases need to be distinguished, as $\tilde{b} > L$ when $n > 6$ and $\tilde{b} < L$ otherwise. For every case, we look for a linear bidding function b such that the maximum of the difference between the maximal loss of the linear bidding function and the minimax loss is as small as possible. The best fit is given by

$$b(v|n) = \begin{cases} \lambda_1(n) \cdot v + (1 - \lambda_1(n)) \cdot L & \text{for } n \leq 6 \\ \lambda_2(n) \cdot v + (1 - \lambda_2(n)) \cdot \tilde{b} & \text{for } 6 < n \leq 10, \end{cases}$$

where

$$\lambda_1(n) = 0.46 + 0.04 \cdot n$$

$$\lambda_2(n) = 0.38 + 0.05 \cdot n.$$

The approximated bidding function leads to a higher maximal loss than the minimax bids. Proportional to $v - L$, the maximal loss of the approximated bidding function is at most 0.017 larger for $n \leq 6$ and of at most 0.027 for $6 < n \leq 10$. A linear fit of maximal loss proportional to $v - L$ of the approximated bidding function is

$$\frac{\max l(b)}{v - L} = \begin{cases} 0.51 - 0.05 \cdot n & \text{for } n \leq 6 \\ 0.33 - 0.02 \cdot n & \text{for } 6 < n \leq 10. \end{cases}$$

One can see that loss is decreasing in the number of bidders and that the optimal bid is increasing in the number of bidders. For $n = 10$ loss relative to $v - L$ is only about 13%. This is substantially lower than the 50% bound under complete uncertainty.

5 Conceiving Value Distributions and Bidding Functions

In Section 4 we described the uncertainty directly in terms of the conceivable bid distributions. Now we consider the other bidders more explicitly. We first

reexamine the case of complete uncertainty and look for justifications of the lower bound L . Next, we discuss the impact of beliefs about the value distribution and the bidding behavior of other bidders. Beliefs only about the value distribution or the bidding behavior do not have much bite, as one cannot improve the $(v - L)/2$ bound of minimax loss. A combination of beliefs about the value distribution and bidding behavior, however, does lead to a lower upper bound on loss. In particular, we present a set-up in which minimax loss is attained by a bidding function linear in the own value.

In the setting of complete uncertainty (Definition 1) the robust bidder is completely uncertain about the value distribution and bidding functions. She only conceives that the maximal bid among the other bidders is above L . Two simple scenarios lead to the conclusion of the maximal bid being above L . In the first scenario there are at least two other bidders, who are rational, bid independently, and believe that the value of the other is above L . In this case, neither of them will bid below L . In the second scenario, the maximal bid is above L if one believes that there is another robust bidder who applies the results of this paper and who believes that the maximal bid of the bidders he faces is certainly above L . Hence, it is not necessarily the own beliefs that lead to the bound L , but it can be anticipating beliefs and behavior of others.

One might wonder whether the error bound $(v - L)/2$ becomes smaller if one makes more assumptions on the bidding behavior of the others while maintaining complete uncertainty about the value distributions. We explore this in the framework closest to the classic model of independent private values. Assume that all other bidders know the value distribution, that there is common knowledge of rationality among them and that they know the strategy of the robust bidder. Moreover, assume that the robust bidder knows the above assumptions, but does not know the value distribution. The error bound does not change if there are at least two other bidders. To see this, let the value distribution F be iid and equal to $\epsilon[v_1] + (1 - \epsilon)[M]$, $M \geq L$ and ϵ sufficiently small.¹¹ Bayesian bidders basically know that all bidders have the same value M . Consequently, bidding (almost) value is a best response to each other. This behavior generates the same conditions as in Subsection 4.1 and Proposition 1 applies. Loss cannot be decreased by simply restricting uncertainty to uncertainty over values.

One cannot improve the $1/2$ error bound if one knows the value distribution, but not the bidding function of other bidders. Suppose the value distribution was

¹¹Alternatively, consider an asymmetric value distribution in which the Bayesian bidders have some prior over v_1 , and know that all rational bidders have the same value.

known, but any bidding behavior was deemed possible. In this case it cannot be ruled out that all other bidders bid M irrespective of their value. Any bid above L can happen with probability 1, so loss can only be bounded by $(v - L)/2$. One can imagine that the bidder's perception of the environment leads to constraints that rule out this extreme case. We have seen that constraints would need to apply both to the value distribution, so that not all bidders can have value L with certainty, and to bidding behavior to imply that not all bidders bid L irrespective of their type. In the next subsection we make restrictions that rule out these cases.

5.1 Behavioral Beliefs: A Linear Lower Bound on Bidding Functions

There are $n \geq 2$ bidders. Values are distributed independently and identically. The lowest possible value is $K \in [0, v)$. All conceivable value distributions can be bounded from above. In particular, the true value distribution F is such that $F(v) \leq \eta(v - K)^\alpha$, where $\eta > 0$ and $\alpha > 0$. The larger the parameter α , the less likely are values around K . The smaller η the less mass can be put on low types. Other bidders use deterministic bidding functions monotone in their value that can be bounded from below by linear functions, i.e. bidder $i > 1$ uses a bidding function $b_i(v) \geq \sigma v$, with $0 < \sigma < 1$. The set of conceivable environments is formally defined as follows.

Definition 4. Given $0 < \alpha$, $0 < \eta$, $0 < \sigma < 1$ and $0 \leq K$, let $\mathcal{E}_{\alpha\eta}^\sigma$ be the set of environments belonging to \mathcal{U} in which for all conceivable value distributions F , (i) there are n bidders, (ii) values are identically and independently distributed, (iii) $F(v) \leq \min\{1, \eta(v - K)^\alpha\}$, and for all conceivable bidding functions b_{-1} we have that $b_i(v) \geq \sigma v$ for $i > 1$.

Loss comes from bidding too high or too low. Under complete uncertainty the highest bid of the other bidders M can be any bid with probability 1. Hence, the bid $b \in [K, v)$ can be too low with probability 1 and it can be too high with probability 1. Under the constraints of this section low bids of other bidders can only occur with low probability. For example, the maximal bid of others cannot be K for sure. In particular, the probability that the bid b is larger than bidder i 's bid can be bounded by

$$\mathbb{P}(b_i(v) \leq b) \leq \mathbb{P}(\sigma v \leq b) \leq \eta \left(\frac{b}{\sigma} - K \right)^\alpha.$$

The robust bidder faces an unknown bid distribution. Let $B(b)$ be the probability that the bid b becomes winning. Given the constraints, we have $B(b) \leq$

$\gamma (b - K\sigma)^\beta$, with $\gamma = \left(\frac{\eta}{\sigma^\alpha}\right)^{n-1}$ and $\beta = \alpha(n - 1)$. The minimax bid balances the maximal loss from bidding too low and bidding too high and depends on the parameters of the model.

The following proposition says that the deterministic minimax bid is an affine function. The optimal bid is a convex combination of the own value and the lowest possible bid of other bidders $L = \sigma K$. A bid below L is, as above, never winning. Weight ρ is put on v and weight $1 - \rho$ on L , where $\rho = (1 + \beta)^{1+\beta} / ((1 + \beta)^{1+\beta} + \beta^\beta)$. The parameter $\rho \in (1/2, 1)$ is strictly increasing in β . As β is increasing in the number of bidders, ρ increases in n and reaches 1 in the limit as n grows large.

Proposition 4. *Let $0 < \alpha$, $0 < \eta$, $0 < \sigma < 1$, $0 \leq K$ and $v \leq L + \frac{\sigma}{\rho\eta^{1/\alpha}}$. For the set of conceivable environments $\mathcal{E}_{\alpha\eta}^\sigma$, deterministic minimax loss is equal to*

$$\gamma \frac{\beta^\beta (1 + \beta)^{\beta(1+\beta)}}{\left((1 + \beta)^{1+\beta} + \beta^\beta\right)^{1+\beta}} (v - L)^{1+\beta}$$

and attained by bidding function

$$b^*(v) = \rho v + (1 - \rho) L. \tag{4}$$

The bound on the bid distribution is only useful if the value v is not too high, i.e. if $\eta(\rho v / \sigma - K)^\alpha \leq 1$, which is equivalent to $v \leq L + \frac{\sigma}{\rho\eta^{1/\alpha}}$.

We look at two examples to get a feeling for how the results can be applied and for the magnitude of minimax loss. In both examples we set $K = 0$. In this case b^* only depends on the total number of bidders and the bound on the value distribution. It is independent of the beliefs about the other bidders bidding behavior (σ). When low types are relatively likely (small α) and there are few bidders, then the optimal bid is just above $v/2$. For a linear bound ($\alpha = 1$) and two bidders, the optimal bid equals $0.8v$. In the first example we ask how much mass can be at most below the own type. In the second example we fix the highest level of loss and find an upper bound for the quantile in which the own value can lie. For the examples it is convenient to introduce a new variable. Let $\mu = \eta^{1/\alpha}$, so $F(v) \leq \eta v^\alpha = (\mu v)^\alpha$. The bound on loss is tight for $v \leq \frac{\sigma}{\rho\eta^{1/\alpha}} = \frac{\sigma}{\rho\mu}$, so if $\mu v \leq \frac{\sigma}{\rho}$. Loss is then $l \leq \left(\frac{\mu}{\sigma}\right)^\beta (1 - \rho) \rho^\beta v^{1+\beta} = \left(\frac{\rho}{\sigma}\right)^\beta (1 - \rho) (\mu v)^\beta v \leq (1 - \rho) v$, where we use the fact that $\mu v \leq \frac{\sigma}{\rho}$.

Example 2. There are five bidders and the lowest possible value is $K = 0$. All other bidders bid at least half of their value, i.e. $\sigma = 1/2$. Let $\alpha = 1$. The robust bidder bids quite aggressively as $\rho = \frac{3125}{3381} \approx 0.92$. Proposition 4 applies for

$v \leq \frac{\sigma}{\rho\mu} = \frac{3381}{6250} \frac{1}{\mu} \approx 0.54 \cdot \frac{1}{\mu}$. Hence, when bidder 1 puts at most mass 0.54 below the own value v_1 , then loss is at most $(1 - \rho)v = 0.0757 \cdot v$.

Now let $\alpha = \frac{1}{2}$, so intuitively there can be more low types. In this case we have $\rho = \frac{27}{31} = 0.87$. The bound on v_1 is tight if $v \leq 0.57 \cdot \frac{1}{\mu}$. This condition translates to the case in which the robust bidder puts at most mass 0.76 below the own value, which is true whenever $\eta \leq \sqrt{\sigma/(\rho v_1)}$. Loss can then be bounded by $(1 - \rho)v = 0.13 \cdot v$, which is worse than above as more mass is allowed below value. If the bidder puts less mass below the own type, then loss can be bounded further. For example, suppose at most mass 0.54 (this is the number from above) is put below the own value v_1 (e.g. $\eta = 0.54/\sqrt{v_1}$). Loss can be bounded by $(\frac{\rho}{\sigma})^\beta (1 - \rho) ((\mu v)^\alpha)^{n-1} v \leq (\frac{\rho}{\sigma})^\beta (1 - \rho) (0.54)^{n-1} v = 0.03 \cdot v$. This bound on loss is tighter than in the case with $\alpha = 1$.

In the second example we ask how much mass can be at most below the own type such that loss is not more than 10% of the own value.

Example 3. Let $K = 0$, $n = 5$, $\sigma = 1/2$ and $\alpha = 1/2$. Minimax loss is less than 10% of the own value if $(\frac{\rho}{\sigma})^\beta (1 - \rho) (\mu v)^\beta v \leq 0.1 v$. This is true if $\mu v \leq 0.51$. Note that in this case the constraint on loss is tight, as $0.51 < \sigma/\rho = 0.57$. We have that $\mu v \leq 0.51$, when $\eta \leq \sqrt{\frac{0.71}{v}}$, so if the maximal mass below the own value is 0.71. The restrictions on the value distribution leads to a loss below 10% of the own value if one believes that the own value is not in the upper 0.29 quantile of the value distribution.

5.2 Behavioral Beliefs: An Affine Lower Bound on Bidding Functions

Above we analyzed the case in which other bidders' bidding can be bounded below by a linear function. There might be cases in which one wants to use an affine lower bound on bidding functions, i.e. $b_i(v) \geq L + \sigma(v - L)$ for all other bidders, where L denotes the lowest possible value for other bidders. One case is when one expects other robust bidders who have similar perceptions of the environment. In this case the lower bound on the bidding function is linear, but the robust bidding function from above is affine. Both $\rho > \sigma$ and $\rho < \sigma$ are possible. In particular, when α is sufficiently small and $\sigma > \frac{1}{2}$, then $b^*(v) < \sigma v$ for large v . So it can be that $b^*(v) \geq \sigma v$ is inconsistent with some other bidder being robust and having identical conceptions about environment. With an affine lower bound this inconsistency does not occur.

Let $B(b)$ be the probability that all other bidders bid below b . The probability that a bid b is winning is maximized by taking both the lower bound on the bidding

function and the upper bound on the value distribution as binding, that is,

$$\mathbb{P}(b_i(v) \leq b) \leq \mathbb{P}(L + \sigma(v - L) \leq b) \leq \eta \left(\frac{b - L}{\sigma} \right)^\alpha.$$

The upper bound on the bid distribution is therefore $B(b) \leq \min \left\{ \eta \left(\frac{b-L}{\sigma} \right)^\alpha, 1 \right\}^{(n-1)}$. Thus, everything is as in the linear case, except replacing σK by L . This leads to the following proposition, where the set of conceivable environments $\mathcal{E}_{\alpha\eta}^{\sigma'}$ is as in Definition 4, but for the different bound on others' bidding functions.

Proposition 5. *Let $0 < \alpha$, $0 < \eta$, $0 < \sigma < 1$, $0 \leq L$ and $v \leq L + \frac{\sigma}{\rho\eta^{1/\alpha}}$. For the set of conceivable environments $\mathcal{E}_{\alpha\eta}^{\sigma'}$, deterministic minimax loss is equal to*

$$\gamma \frac{\beta^\beta (1 + \beta)^{\beta(1+\beta)}}{\left((1 + \beta)^{1+\beta} + \beta^\beta \right)^{1+\beta}} (v - L)^{1+\beta} \quad (5)$$

and attained by bidding function

$$b^*(v) = \rho v + (1 - \rho) L. \quad (6)$$

In particular, the robust bidder can be more or less aggressive than the boundary bidder as there are no restriction on how ρ relates to σ . If the robust bidder conceives that there may be other robust bidders like her then it may be good to choose σ such that $\rho \geq \sigma$.

6 Comparison with Experimental Data

The objective of this section is twofold. First, we want to empirically assess our bid recommendations with experimental data. Experimental data has the advantage that the bidder's value is known. Second, we demonstrate the flexibility of our methods in using available and relevant information. We use the data of the following three studies. Filiz-Özbay and Özbay (2007) run three treatments to test the effect of post auction feedback rules on bidding behavior. Güth and Ivanova-Stenzel (2003) and Chen et al. (2007) investigate the difference in bidding behavior with a known and an unknown value distribution in an auction.¹²

We compare the performance of our bidding strategies to those used, evaluating performance based on the true value and the empirical distribution of bids. Loss

¹²Appendix C contains a more detailed description of the data and the methods.

associated with a bidding strategy is the difference between the maximal payoff when the bid distribution is known and the payoff of the bid strategy. We compute the empirical bid distribution for each treatment. The results are presented in a table for every experiment. In every table the columns refer to the treatments of the experiment. For each treatment and every bid strategy we present the average and the maximal loss. Every table reports the loss of actually observed bids, loss of deterministic bid recommendations, and loss of Nash equilibrium. We assume risk neutrality and give loss in per cent of the own value.

We first consider the data of Filiz-Özbay and Özbay (2007) (FÖÖ). In their experiment there are four bidders who know that the true value distribution is iid and uniform on $[0, 100]$. We compare three deterministic bid recommendations. First, the model of complete uncertainty (Proposition 1) is evaluated with $L = 0$ and gives the bidding function $b(v) = v/2$. The second recommendation is the ε -uniform model of Proposition 3. The lower bound is set to 0 and the information on the number of bidders is used. The value of $\varepsilon = 0.15$ was chosen ex-ante, so we get $b(v) = 0.6365 \cdot v$. The last recommendation is a compound model, as it is a combination of Propositions 2 and 4. One can set Proposition 4's parameters $\alpha = 1$ and $\eta = 1/100$, as it is known that the true value distribution is uniform on $[0, 100]$. Moreover, we assume that no other bidder bids below half of the own value, so we set $\sigma = 0.5$. The proposition only holds for relatively small types. For larger types we use Proposition 2. The proposition requires that the probability that the maximal bid of the other bidders is below a certain threshold is not too high. This probability can be computed, because the true value distribution is uniform and we assume that others bid at least half their value. Hence, we choose the parameter L of Proposition 2 such that the bound of the proposition is satisfied, i.e. that

$$\mathbb{P}(M < L) \leq \mathbb{P}((\sigma\tilde{v})^3 < L) = \mathbb{P}\left(\tilde{v} < \frac{L^{\frac{1}{3}}}{\sigma}\right) = \frac{L^{\frac{1}{3}}}{\sigma} \leq \frac{\sigma\eta^{-1/\alpha}/\rho - L}{\sigma\eta^{-1/\alpha}/\rho + L}$$

and that the bidding function is continuous. The bidding function that combines the two propositions is then given by

$$b(v) = \begin{cases} \rho \cdot v & \text{for } v \leq \frac{\sigma}{\rho\eta^{1/\alpha}} \\ \frac{v+L}{2} & \text{for } v > \frac{\sigma}{\rho\eta^{1/\alpha}} \end{cases} \approx \begin{cases} 0.90 \cdot v & \text{for } v \leq 55.27 \\ \frac{v}{2} + 22.4 & \text{for } v > 55.27 \end{cases}. \quad (7)$$

Table 1 shows that the mean loss of the actual bids is around 2.5% of the value, while the mean loss is less than 2% for all deterministic bid recommendations.

	Loser		Winner		No		
	L	Feedback		Feedback		Feedback	
		Mean	Max	Mean	Max	Mean	Max
Observed Bids	-	2.36	11.82	2.93	10.53	2.44	8.36
Nash	-	0.51	2.44	0.88	4.74	0.69	3.31
Complete Uncertainty	0	1.70	4.74	1.81	4.59	1.97	5.91
ε -uniform	0	0.63	2.28	0.66	2.59	0.60	2.50
Compound Model	-	0.88	3.46	1.59	5.20	1.31	4.34

Table 1: Performance of different bidding strategies according to the Filiz-Özbay and Özbay (2007) data.




Figure 1: The distributions of Chen et al. (2007) and our bound

What is more, the ε -uniform model leads to a mean loss well below 1% of value. Across treatments, the mean loss of the ε -uniform model is just 25% of the loss of real bids. Our bid recommendations also perform better in terms of the maximal loss. The maximal loss of the ε -uniform recommendation is around 2.5% of the value, which is the number of the average loss of the real bids. The bid strategy associated with Nash equilibrium does not outperform other bid recommendations.

Chen et al. (2007) test how information about the true value distribution influences bidding in a first-price auction. They have two treatments and there are two bidders in every auction. Values are drawn independently. Specifically there are two known and piece-wise linear distributions F^1 and F^2 (see Figure 1) where a value is drawn from F^1 with probability δ . In one treatment the parameter δ is known (it was chosen equal to 0.7), and hence the distribution of values is known. In the other treatment the parameter δ is not known and hence there is ambiguity about the distribution.

	Known		Unknown		
	L	Distribution		Distribution	
		Mean	Max	Mean	Max
Observed Bids	-	17.07	143.60	12.14	190.40
Nash	-	1.97	4.62	3.08	8.60
Complete Uncertainty	0	1.98	4.62	2.00	6.14
ε -uniform	0	2.10	4.91	1.89	6.07
Compound Model	-	4.33	7.72	4.33	10.02

Table 2: Performance of different bidding strategies according to the Chen et al. (2007) data.

We basically evaluate the same bidding functions for the Chen et al. (2007) (CKO) data, but we have to make some changes due to different parameters. A first difference is that there are only two bidders, so the bidding function $b(v) = 0.5220 \cdot v$ is used for the ε -uniform model. A second difference is the construction of the compound bidding function that uses Propositions 2 and 4. In the calculation of the bidding function we use the following parameters: $n = 2$, $\alpha = 0.4$, $\eta = 100^{-\alpha}$, and $\sigma = 0.5$. The parameters α and η are chosen such that for any δ the distribution of values first order stochastically dominates ηv^α (*cf.* Figure 1). Proposition 4 can be applied for values below 71.64. The threshold L of Proposition 2 cannot be chosen so that the bidding function is continuous. Hence, we set it as large as possible so that Proposition 2 can be applied to the borderline value 71.64. For $L = 16.03$ the inequality $\eta(L/\sigma)^\alpha \leq (71.64 - L)/(71.64 + L)$ holds as an equality. The bidding function is then

$$b(v) = \begin{cases} \rho \cdot v & \text{for } v \leq \frac{\sigma}{\rho\eta^{1/\alpha}} \\ \frac{v+L'}{2} & \text{for } v > \frac{\sigma}{\rho\eta^{1/\alpha}} \end{cases} \approx \begin{cases} 0.70 \cdot v & \text{for } v \leq 71.64 \\ \frac{v}{2} + 8.01 & \text{for } v > 71.64 \end{cases}. \quad (8)$$

The last difference relates to Nash equilibrium play. The value distribution is only common knowledge in the known valuation treatment. In this case we compute the risk-neutral Nash equilibrium bidding function. In the unknown value distribution treatment we compute the maximin expected utility Nash equilibrium, that is, the Nash equilibrium for the case in which the worst parameter of δ is the true one. This corresponds to F_2 being the true value distribution. The row “Nash” in Table 2 displays the finding for the risk-neutral Nash equilibrium in the known treatment and the maximin expected utility equilibrium in the unknown treatment.

Table 2 shows the performance of the bid strategies using the Chen et al.

(2007) data. The results are similar to the findings in Table 1, however, some bids substantially above value contribute to the high actual loss. The median loss in per cent of the value of the observed bid is 9.78% for the known and 5.50% for the unknown distribution; both numbers are much lower than the respective means. Our bid recommendations perform again much better than the observed bids. The worst candidate is the compound model. The two other functions perform quite well. The Nash recommendation performs well in the known treatment, whereas the maximin recommendation performs less well in the unknown treatment.

In the experiment of Güth and Ivanova-Stenzel (2003) (GIS) there are two asymmetric bidders. For the weak bidder, bidder 1, the value distribution is uniform on $[50, 150]$ and for the strong bidder 2 it is uniform on $[50, 200]$. The supports and the distribution are common knowledge only in one treatment. In the other treatment nothing is known about the value distribution. In the known value distribution treatment, bidders know that the lowest possible value is 50. We consider a robust bidder $i = 1, 2$ who conceives that the other bidder does not bid below 50% of the own value, so $\sigma = 0.5$ and $b_{3-i}(v) \geq \sigma v$. This gives us a lower bound of 25 for the highest bid of the other bidder. We compute the loss associated with the bid strategy $b(v) = (v + L)/2$, where we use $L = 25$. The deterministic minimax bid of the ε -uniform model is also computed with $L = 25$ and equal to $b(v) = 0.5220 \cdot v + 11.95$. The weak bidder 1 knows the strong bidder's value distribution. This translates into $\alpha = 1$, $\eta_1 = 1/150$. The strong bidder 2 has $\eta_2 = 1/100$. In both cases we have $F(v) \leq \eta(v - 50)$ for the relevant v . Proposition 4 can be applied for values below 118.75 for the weak bidder and for values below 87.5 for the strong bidder. As above, for larger values we choose L so that Proposition 2 can be applied. This gives rise to bidder $i = 1, 2$'s compound bidding function

$$b_i(v) = \begin{cases} 0.8 \cdot v + 5 & \text{for } v \leq 25 + \frac{\sigma}{\rho\eta_i^{1/\alpha}} \\ \frac{v+L_i}{2} & \text{for } v > 25 + \frac{\sigma}{\rho\eta_i^{1/\alpha}}. \end{cases} \quad (9)$$

For the weak bidder 1 $L_1 = 53.44$ and for strong bidder 2 $L_2 = 42.37$. The Nash equilibrium functions are computed for the weak and the strong bidder, respectively. In the unknown value distribution treatment, bidders do not know that there are asymmetric value distributions and they do not know the support of the value distribution. The two basic bidding functions are evaluated with $L = 0$. For the compound model, we choose $\alpha = 1/3$, $\sigma = 0.5$ at discretion and assume that no bidder believes that her value is in the top 0.09 quantile. Hence, the

	Known Distribution					Unknown Distribution		
	L	Weak Bidder		Strong Bidder		L	Mean	Max
		Mean	Max	Mean	Max			
Observed Bids	-	6.17	24.64	14.02	83.43	-	7.85	26.86
Nash	-	2.57	3.91	3.47	5.79	-	-	-
Complete Uncertainty	25	1.70	3.77	2.59	6.14	0	4.02	7.92
ε -uniform	25	1.55	3.90	1.86	4.85	0	3.79	8.13
Compound Model	-	3.26	6.10	3.46	6.17	-	3.13	9.92

Table 3: Performance of different bidding strategies according to the Güth and Ivanova-Stenzel (2003) data.

bidding function $b(v) = 0.68 \cdot v$ is used.

Table 3 summarizes the findings for the data of Güth and Ivanova-Stenzel (2003). The overall picture is similar to other two data sets. The actual loss is higher on average than for our recommendations. The ε -uniform recommendation of Proposition 3 performs well in both treatments. The combination of Propositions 2 and 4 achieves the lowest loss on average in the unknown treatment.

We conclude this section by looking at the statistical difference between our bid recommendations in order to find the best recommendation. Eyeballing suggests that the ε -uniform model performs best across treatments, while contenders are $b(v) = v/2$ and Nash.¹³ Statistical tests summarized in Appendix C confirm that the recommendation is either statistically indistinguishable, or better than any of the alternative suggestions. Hence, the bidding function derived in the ε -uniform model seems like the best choice across treatments and experiments. It introduces the important dependency on the number of bidders in a simple and reduced way.

7 Conclusion

One of the major obstacles and challenges to bidding in first-price auctions is limited information. In many instances it is difficult to assess other bidder’s value distributions and bidding functions (i.e. the environment) and to specify beliefs and higher-order beliefs. Misspecification can lead to substantial loss. This is the first paper that derives robust bidding rules in first-price auctions. We deal with the uncertainty by searching for a compromise that performs well for a wide

¹³In the “known” treatment of Güth and Ivanova-Stenzel (2003) we use the respective equivalents of the two bidding functions with different lower bounds.

variety of situations. The methodology based on compromises is easy to explain and justify. We evaluate bidding functions based on loss, where loss compares the payoff in an environment to the payoff of the best bidding rule if the true environment were known (that is, in the truly hypothetical and unrealistic oracle).

Our methodology has been designed to aid bidding in real auctions. A bidder can choose the recommendation best suited for the faced situation. First, one needs to decide whether one wants to think directly about the faced bid distribution, or if one wants to think about bounds on the value distribution and bounds on other bidders' behavior. The first step in the former case is finding a suitable threshold L . The threshold L can be the lowest maximal bid among the other bidders, it can be a known reserve price, or it can be a value chosen at discretion. For example, a bidder might never want to bid below half of the own value due to some reason. Note that one does not need to say that the maximal bid of other bidders is above L with certainty. We provide a bound on the probability with which this must be true. An alternative is to use the ε -uniform model of Subsection 4.3. One simply has to choose an ε and decide on a number of bidders in addition to L . This model has the advantage that bids are increasing in the number of bidders and that knowing that there are more bidders reduces loss. One can also change the uniform distribution in the ε -uniform model to some other preferred bid distribution.

The alternative to thinking about the bid distribution is specifying bounds on the value distribution and linear or affine bounds on other bidders' bidding behavior. In this case we develop a model with quite some flexibility. Bidding is linear (or affine) in the own value and depends on only few parameters, the number of bidders and the parameters used to bound the value distribution. This model can be used if one is willing to place enough mass above the own value, where what is enough depends on the chosen parameters.

It is interesting that the same simple functional form can be optimal in different situations. The minimax bid is always a convex combination between the own value and the reserve price or lowest maximal bid of the other bidders. In the environments we have considered, the weight on the own value is at least $1/2$. Bidding half of the value remains approximately optimal when there are few bidders. The weight on the own value increases if independent bidding is assumed and there are more bidders.

A Proofs

Proposition 2. *Let $v > L > 0$ and $0 < \bar{p} < 1$. For the set of conceivable environments $\mathcal{E}_{L\bar{p}}$ with $\bar{p} \leq \frac{v-L}{v+L}$ deterministic minimax loss is equal to $\frac{v-L}{2}$ and attained by the deterministic bidding strategy stated in Proposition 1.*

Proof. Let $\bar{p} \leq \frac{v-L}{v+L}$. We show that loss cannot be higher than $\frac{v-L}{2}$ when strategy $\frac{v+L}{2}$ is played. Loss is potentially maximized by environments of the form $\bar{p}[0] + (1-\bar{p})[M]$ with $M \geq L$. Note that in the oracle one either bids slightly above zero or slightly above M . We distinguish two cases. First, it can be that one bids too low relative to M , i.e. $M \geq \frac{v+L}{2}$ and second, that one bids too high. We start with the first case, i.e. $M \geq \frac{v+L}{2}$. In the oracle, payoff is maximized by either bidding zero or M . If one bids zero, then loss equals

$$\bar{p}v - \bar{p} \left(v - \frac{v+L}{2} \right) = \bar{p} \frac{v+L}{2}.$$

This loss is smaller than $\frac{v-L}{2}$, as $\bar{p} \leq \frac{v-L}{v+L}$. The other sub-case occurs when one bids M in the oracle. It is straightforward to verify that then maximal loss at most $\frac{v-L}{2}$.

The second case is when $L \leq M \leq \frac{v+L}{2}$. If M is the payoff maximizing bid in the oracle, then loss is smaller than $\frac{v-L}{2}$. The other sub-case is when zero is chosen in the oracle. Loss is then equal to $\bar{p}v - \frac{v-L}{2}$. This is less than $\frac{v-L}{2}$, because $\bar{p} \leq \frac{v-L}{v+L} < \frac{v-L}{v}$. To summarize, there exists no environment in $\mathcal{E}_{L\bar{p}}$ with $\bar{p} \leq \frac{v-L}{v+L}$ that increases loss above $\frac{v-L}{2}$ when bidder 1 bids $b^*(v) = \frac{v+L}{2}$. \square

Proposition 3. *Let there be $n \geq 2$ bidders, $v > L \geq 0$, and $\varepsilon \in (0, 1)$. For the set of conceivable environments $\mathcal{E}_{L,\varepsilon,n}$ deterministic minimax loss is attained by b^* such that*

$$\pi(\max\{L, \tilde{b}\}) - \pi(b^*) = \sum_{k=0}^{n-2} \binom{n-1}{k} \varepsilon^k (1-\varepsilon)^{n-1-k} (v-b^*)^k p(b^*)^k \quad (3)$$

and equal to the value on either side of the equation, where $\tilde{b} = \frac{L+\varepsilon nv-v}{\varepsilon n}$.

Proof. We first analyze the oracle payoff and then deterministic minimax loss. The following lemma considers the maximization of the payoff in the oracle with known M . It says that payoff in the oracle is maximized by either bidding (slightly above) the highest bid among the other bidders M or by bidding \tilde{b} , which is independent of M . Intuitively, when ε is small relative to n , then one basically has to outbid

only those who bid M . On the other hand, when ε is relatively large relative to n , then it can be the case that the optimal bid is independent of M , as one needs to outbid the uniform bids.

Lemma 1. *Let $M \in [L, v)$. The payoff in the oracle is maximized by $\tilde{b} = \frac{v(n\varepsilon-1)+L}{n\varepsilon}$ whenever $M < \tilde{b}$, or by bidding (slightly above) M , or by bidding $\tilde{y} = \frac{(n-1)v+L}{n}$.*

Proof. The bid distribution has a mass point at M . Therefore, we have to distinguish the case from bidding below M and bidding above M . It is never optimal to bid M in the oracle, as bidding slightly above M avoids tie-breaking and leads to a higher payoff. We consider the two relevant cases separately. Conditional upon bidding above M , the payoff in the oracle is equal to

$$\pi(b|b > M) = \sum_{k=0}^{n-1} \binom{n-1}{k} \varepsilon^k (1-\varepsilon)^{n-1-k} (v-b)p(b)^k. \quad (10)$$

Conditional on bidding below M , payoff is equal to

$$\pi(b|b < M) = \varepsilon^{n-1}(v-b)p(b)^{n-1}. \quad (11)$$

We start by maximizing $\pi(b|b > M)$. To this end we first discuss the roots of the first order condition of Equation (10) and identify \tilde{b} as the payoff maximizing root. Depending on the parameters, the bid \tilde{b} can be smaller than L and consequently smaller than M . In this case, we show that M is the argument at which the supremum of payoff is attained.

The first order condition of Equation (10) for $b > M$ with respect to b is

$$\frac{\partial \pi(b|b > M)}{\partial b} = \frac{(v-L)^{1-n}(b\varepsilon - L - \varepsilon v + v)^n(-b\varepsilon n + L + v(\varepsilon n - 1))}{(-b\varepsilon + L + (\varepsilon - 1)v)^2} = 0. \quad (12)$$

Its distinct roots are $b' = \frac{L+\varepsilon v-v}{\varepsilon}$ and \tilde{b} . The root \tilde{b} is the relevant root, as $b' < L$. The payoff is maximized by bidding \tilde{b} if $\tilde{b} > M$, because $\pi(b|b > M)$ is decreasing for $b \in (\tilde{b}, v)$. This can be seen from Equation (12). Therefore, whenever $\tilde{b} \geq L$ and $\tilde{b} > M$, then \tilde{b} is the unique maximizer of $\pi(b|b > M)$. A necessary condition for $\tilde{b} > M$ is $\tilde{b} > L$ and this is true whenever $\varepsilon > \frac{1}{n}$. However, if $\tilde{b} \leq M$, then bidding M is optimal, i.e. $M \in \arg \sup_b \pi(b|b > M)$, because the payoff is decreasing in b and therefore has its supremum in the smallest possible b .

The payoff $\pi(b|b < M)$ is maximized by bidding $\tilde{y} = \frac{(n-1)v+L}{n}$. This follows directly from the respective first order condition.

Now we come back to maximizing payoff in the oracle. Note that $\tilde{b} < \tilde{y}$, so if $M \leq \tilde{b}$, bidding \tilde{b} yields higher payoff utility than \tilde{y} . For a medium high M , i.e. $\tilde{b} < M < \tilde{y}$, it is clear that bidding M is optimal. Whenever M is high, so when $\tilde{y} < M$, it might be best to “ignore” M and bid \tilde{y} . \square

Now we derive optimal deterministic bids. Any bid b can be too low or too high for the true environment. First, we consider a too low bid, i.e. $b < M$. Loss cannot be maximized by environments with $M < \tilde{b}$, because then loss $\pi(\tilde{b}|b > M) - \pi(b|b < M)$ is the same for all such M . Hence, it is enough to consider only $M \geq \tilde{b}$. A similar argument shows that loss is also not maximized by environments with M such that \tilde{y} maximizes the oracle payoff. Thus, the optimal action in the oracle is to bid slightly above M , leading to a loss of $\pi(M) - \pi(b|b < M)$.¹⁴ The proof of Lemma 1 showed that $\pi(b)$ is decreasing in b for $b > \tilde{b}$, thus the highest loss of bidder 1 by being slightly outbid by one of the non-uniform bidders by bidding $M > b$ is equal to

$$\sup_{x > M > b} \left\{ \sum_{k=0}^{n-2} \binom{n-1}{k} \varepsilon^k (1-\varepsilon)^{n-1-k} (v-x) p(x)^k \right\} = \sum_{k=0}^{n-2} \binom{n-1}{k} \varepsilon^k (1-\varepsilon)^{n-1-k} (v-b) p(b)^k.$$

We still have to consider loss that results from bidding too high (i.e. $b > M$). It is again enough to only consider environments such that the oracle payoff is $\pi(M)$. Hence, loss is equal to $\pi(M) - \pi(b)$. This loss is maximized by $M = \max\{L, \tilde{b}\}$. Conditional on b , the highest loss is then

$$\max \left\{ \pi(\max\{L, \tilde{b}\}) - \pi(b), \sum_{k=0}^{n-2} \binom{n-1}{k} \varepsilon^k (1-\varepsilon)^{n-1-k} (v-b) p(b)^k \right\}$$

The optimal bid b^* equates these two expressions. \square

Proposition 4. *Let $0 < \alpha, 0 < \eta, 0 < \sigma < 1, 0 \leq K$ and $v \leq L + \frac{\sigma}{\rho\eta^{1/\alpha}}$. For the set of conceivable environments $\mathcal{E}_{\alpha\eta}^\sigma$, deterministic minimax loss is equal to*

$$\gamma \frac{\beta^\beta (1+\beta)^{\beta(1+\beta)}}{\left((1+\beta)^{1+\beta} + \beta^\beta \right)^{1+\beta}} (v-L)^{1+\beta}$$

and attained by bidding function

$$b^*(v) = \rho v + (1-\rho)L. \quad (4)$$

¹⁴Note that we slightly abuse notation, as the function $\pi(b)$ is equal to $\pi(b|b > M)$.

Proof. The probability that bid b is winning is bounded by $B(b) \leq \min \left\{ \eta \left(\frac{b}{\sigma} - K \right)^\alpha, 1 \right\}^{n-1}$. Let us first ignore the constraint that probabilities are bounded above by 1, so we obtain $B(b) \leq \gamma(b-L)^\beta$. We consider the loss of the robust bidder who faces some bid distribution B with $B(b) \leq \gamma(b-L)^\beta$. Loss is maximized by bid distributions Q_x that put mass $\gamma(x-L)^\beta$ on x and the rest of mass above v . Bid distributions of the form Q_x amplify what it means to bid too low and too high, because a marginal change in the bid can lead to a substantial change in payoff. To see this, consider the case in which a bid b is too low. A bid is too low when a higher bid would improve payoff. In the worst-case there is no mass of other bidders' bids below the own bid, so the payoff of bidding b is zero. Loss is maximized if as much mass as possible is on x with $b < x < v$. Conversely, the bid b is too high if lowering the bid increases payoff. The problem of bidding too high is most severe if slightly lowering the bid does not increase payoff. In the worst-case the bid distribution has a point mass at the bid $x < b$.

Loss associated with bid b is given by

$$l(b, Q_x) = \gamma(x-L)^\beta ((v-x) - 1_{b>x}(v-b)).$$

The loss of bidding too low is maximized by $\max\{b, \frac{L+\beta v}{\beta+1}\}$. It holds that $b^*(v) > \frac{L+\beta v}{\beta+1}$, so the relevant maximizer of loss is an x slightly above b . The inequality $b^*(v) > \frac{L+\beta v}{\beta+1}$ is equivalent to the true inequality $\beta^{\beta+1} < (\beta+1)^{\beta+1}$. Hence, loss cannot be raised by setting $x = \frac{L+\beta v}{\beta+1}$. The loss of bidding too high is maximized by $x = \frac{\beta b+L}{\beta+1}$ for $L < b$, as $\left. \frac{d}{dx} \left(\gamma(x-L)^\beta (b-x) \right) \right|_{x=\frac{\beta b+L}{\beta+1}} = 0$. Maximal loss is therefore

$$\max \left\{ \gamma(b-L)^\beta (v-b), \gamma \beta^\beta \left(\frac{b-L}{1+\beta} \right)^{1+\beta} \right\}.$$

Maximal loss is attained by the bid that equalizes the two expressions, so choose b such that $\gamma(b-L)^\beta (v-b) = \gamma \beta^\beta \left(\frac{b-L}{1+\beta} \right)^{1+\beta}$. The bidding function $b^*(v) = \rho v + (1-\rho)L$ satisfies this equation. Plugging in the bidding function into maximal loss gives the minimax loss.

$$\begin{aligned} \gamma \beta^\beta \left(\frac{b-L}{1+\beta} \right)^{1+\beta} \Big|_{b=\rho v+(1-\rho)L} &= \gamma \rho^\beta (1-\rho)(v-L)^{1+\beta} \\ &= \gamma \frac{\beta^\beta (1+\beta)^{\beta(1+\beta)}}{\left((1+\beta)^{1+\beta} + \beta^\beta \right)^{1+\beta}} (v-L)^{1+\beta}. \end{aligned}$$

The bound on loss is tight if $\eta \left(\frac{1}{\sigma} b^*(v) - K \right)^\alpha \leq 1$, i.e. if $b^*(v) \leq L + \frac{\sigma}{\eta^{1/\alpha}}$, which is equivalent to $v \leq L + \frac{\sigma}{\rho\eta^{1/\alpha}}$. □

B Randomized Bidding

Appropriate randomization can further reduce minimax loss relative to deterministic minimax loss. We will consider the cases of complete uncertainty, mass below L , and the congestion ε -uniform model. The section on mass below L also considers the case with $v < L$. The last subsection evaluates the bidding functions with the experimental data.

B.1 Conceiving the Lower Threshold

We show how appropriate randomized bidding can reduce loss compared to the deterministic minimax bidding function of Proposition 1. Let bidder 1 use a mixed strategy with probability density function (pdf) $g(b|v)$ on some support, which is a subset of $[L, v]$. Bids below L are never winning and bids above v yield non-positive payoff. The corresponding cumulative distribution function (cdf) of the mixed bidding function is denoted by $G(b|v)$. Bidder 1 wins the auction if her bid is above the highest bid of the other bidders M and loses it otherwise. Loss is equal to the following difference when M is known and when it is not known, i.e.

$$\begin{aligned} l(G, M|v) &= \max \left\{ \sup_{x > M} \{v - x\}, 0 \right\} - \int_M^v v - b dG(b|v) \\ &= \max \{v - M, 0\} - \int_M^v v - b dG(b|v) \end{aligned} \quad (13)$$

If M is known then bidder 1 gets either (approximate) utility of $v - M$ by bidding (slightly above) M , or 0 if $M \geq v$. All bids above M are winning and bidder 1 computes the payoff of using the randomized bidding function G .

Proposition 6. *For the set of conceivable environments \mathcal{E}_L minimax loss is $\frac{v-L}{e}$ and attained by the randomized bidding strategy with density*

$$g(b|v) = \frac{1}{v-b} \text{ on } \left[L, v - \frac{v-L}{e} \right]. \quad (14)$$

The proof specifies the details how the bidding function is derived. The mean bid of bidding function (14) is $(v + L(e - 1))/e$ and less than the median, which

is equal to $(v(\sqrt{e} - 1) + L)/\sqrt{e}$. The median and the mean are both less than the deterministic bid. This shows that one needs to bid relatively low in order to minimize maximal loss.

Syrngkanis and Tardos (2013) show that in any equilibrium of the first-price auction with independent private values the ratio of the realized social welfare and the highest possible social welfare is at least $(1 - \frac{1}{e})$. Interestingly, in the proof they use the same randomized bidding function that we identify as the minimax bidding function under complete uncertainty (Equation (14)).

Proof. We show two ways to derive the optimal randomized bidding function for bidder 1. Then we consider environments, that is, bid distributions such that minimax loss is attained.

We derive the optimal randomized bidding function for bidder 1. Let M_1 and M_2 be two highest bids such that maximal loss is attained at M_1 and M_2 . Clearly, the loss needs to be the same for these two bids. Without loss of generality, let $M_1 > M_2$ and observe that $l(G, M_1) = l(G, M_2)$ is equivalent to

$$v - M_2 - v + M_1 = \int_{M_2}^{M_1} (v - b)g(b) db.$$

This equation is satisfied by $g(b|v) = \frac{1}{v-b}$ with support $[L, \bar{b}]$. The upper bound of the support is determined by $\bar{b} \leq v$ that solves $\int_L^{\bar{b}} g(b|v)db = 1$ and equal to $\bar{b} = v - \frac{v-L}{e}$.

Plugging in the bidding function and the support in Equation (13) gives loss

$$v - M - \int_M^{v - \frac{v-L}{e}} db = \frac{v - L}{e}.$$

An alternative derivation of the bidding function g is to take the first derivative of loss as specified in Equation (13) with respect to M and solve the first order condition

$$g(M|v)(v - M) - 1 = 0$$

for $g(b|v)$. This leads to the same random bidding function as specified in Equation (14). We will mostly use the FOC approach.

Now we derive environments in which the bound on loss is tight. One can model the minimization of the maximal loss as a zero-sum game between bidder 1 and nature. Nature knows v and chooses the highest bid among other bidders M . The objective of the bidder is to minimize loss, while the nature's objective is the maximization of loss. Nature is indifferent between all M if bidder 1 uses an

optimal bidding function. One obtains this bidding function by setting the first derivative of (13) equal to zero, as shown above.

Nature chooses the cumulative distribution function $H(M|v) = \frac{v-M}{e}$ to make bidder 1 indifferent between all bids $b_1, b'_1 \in [L, v - \frac{v-L}{e}]$. Loss must be equal for both bids, i.e.

$$\int (v-M)dH(M|v) - (v-b_1)H(b_1|v) = \int (v-M)dH(M|v) - (v-b'_1)H(b'_1|v)$$

must hold. Plugging in $b'_1 = L$ and simplifying gives

$$H(b_1|v) = \frac{(v-L)H(L|v)}{v-b_1}.$$

Observe that nature does not want to place any bids above $v - \frac{v-L}{e}$, because this only decreases loss, thus $H(v - \frac{v-L}{e}|v) = 1$. Solving for $H(L|v)$ gives $H(L|v) = \frac{1}{e}$. Nature puts mass $1/e$ on L .

To summarize, an environment $E = (F, B_F)$ in which minimax loss is attained is given by the value distribution F such that $v_2 = \dots = v_n$, $v_2 \sim F$, $F(v_2|v_1) = \frac{v_1}{e(v_1-v_2)}$ on $[L, v - \frac{v-L}{e}]$ and $B_F = \{b_{-1}|b_i(x) = x \text{ for } 1 < i \leq n\}$. There are other environments that generate the same loss. In any of these environments, the distribution of the maximal bid among other bidders is given by H . \square

B.2 Conceiving Mass Below the Threshold

We come back to the model of Subsection 4.2. There is potentially some mass below the threshold L .

Proposition 7. *Let $v > L > 0$ and $0 < \bar{p} < 1$. For the set of conceivable environments $\mathcal{E}_{L\bar{p}}$ with $\bar{p} \leq \frac{v-L}{v-L+eL}$ minimax loss is equal to $\frac{v-L}{e}$ and attained by the randomized bidding strategy stated in Proposition 6.*

Proof. From Proposition 6 we know that if no bids are below L , then minimax loss is attained by the randomized bidding strategy and equal to $\frac{v-L}{e}$. Therefore, we have to show that loss is maximized if there are no relevant bids of the other bidders below L , i.e. that nature does not want to put mass below L in the fictitious zero-sum game. Subsequently, we show that loss is less than $\frac{v-L}{e}$.

Let bidder 1 use the random bidding function with density $g(b) = \frac{1}{v-b}$ on support $[L, v - \frac{v-L}{e}]$. Loss can potentially be increased if as much mass as possible is below L . Hence, consider bid distributions of the form $\bar{p}[M_1] + (1-\bar{p})[M_2]$, with $0 \leq M_1 < L$ and $L \leq M_2$. The bidding function g performs badly if $M_1 = 0$

and $M_2 > v - \frac{v-L}{e}$. To see this, note that all bids of bidder 1 beat M_1 , but they are too high. Conversely, all bids are lower than M_2 . Loosely speaking, \bar{p} times of the cases the bids are too high and $1 - \bar{p}$ times too low. Loss of environment $E = \bar{p}[0] + (1 - \bar{p})[M_2]$, $M_2 > v - \frac{v-L}{e}$ equals

$$\begin{aligned} l(G, E) &= \max \left\{ \sup_{x>0} \{\bar{p}v - x\}, \sup_{x>M_2} \{v - x\}, 0 \right\} - \bar{p} \int_L^{v - \frac{v-L}{e}} db \\ &= \max \{\bar{p}v, v - M_2, 0\} - \bar{p} \int_L^{v - \frac{v-L}{e}} db. \end{aligned}$$

If $\bar{p}v \geq v - M_2$, then loss equals $\frac{\bar{p}(v+(e-1)L)}{e}$. This loss is less than $\frac{v-L}{e}$, as $\bar{p} \leq \frac{v-L}{v-L+eL}$. On the other hand, if $\bar{p}v < v - M_2$, the inequality must hold in particular for $M_2 = v - \frac{v-L}{e}$, in which case loss is $\frac{(v-L)(1-\bar{p}(e-1))}{e}$, which is less than $\frac{v-L}{e}$. Loss cannot be made larger through other bid distributions. \square

In Example 4 we fix L and \bar{p} and ask for which v the inequality is satisfied so that Proposition 2 can be used.

Example 4. Suppose the true value distribution is uniform on $[0, 1]$ and the other $n-1$ bidders are risk-neutral and play according to the risk neutral Bayesian Nash equilibrium $\beta(v) = \frac{n-1}{n}v$. Bidder 1, however, only knows the median bid $L = \frac{n-1}{2n}$ and $p = \frac{1}{2}$. The inequality $p^{n-1} \leq \frac{v-L}{v-L+eL}$ gives a bound on v . If v is higher than the upper bound, then loss is bounded by $\frac{v-L}{e}$ for the set of conceivable environments $\mathcal{E}_{L,\bar{p}}$. If $n = 2$, then $L = 0.25$ and $p^{n-1} = \frac{1}{2} \leq \frac{v-L}{v-L+eL}$ for $v \geq (1+e)/4 (\approx 0.93)$. For $n = 5$ the median bid is $L = 0.4$ and $v \geq (30+2e)/75 (\approx 0.47)$ is necessary. If $n = 10$, then $L = 0.45$ and $v \geq (4599+9e)/10220 (\approx 0.4524)$ is required.

So far we have considered relatively large v , but now we turn attention to smaller v . In particular, we look at $v < \frac{e}{e-1}L$. The next proposition says that bidder 1 can minimize maximal loss by using the randomized bidding function of Equation (14) on $[0, v - \frac{v}{e}]$. This bidding functions ensures that all bids are below L , as $v < \frac{e}{e-1}L$.

Proposition 8. *Let $v \leq \frac{e}{e-1}L$ and $\bar{p} \geq \frac{v-L}{v}e$ if $L < v$. For the set of conceivable environments $\mathcal{E}_{L,\bar{p}}$ minimax loss is equal to $\bar{p}v/e$ and attained by the randomized bidding strategy stated in Proposition 6 evaluated as if $L = 0$.*

Proof. Loss is maximized by bid distributions of the form $\bar{p}[M_1] + (1 - \bar{p})[M_2]$, where M_1 is the highest bid below L and M_2 the highest bid above L . Bidder 1 bidding M_1 in the oracle yields payoff $\bar{p}(v - M_1)$ and bidding M_2 gives $v - M_2$.

Bidder 1 always bids M_1 if $\bar{p}(v - M_1) \geq v - M_2$ for all $M_1 < L$ and $M_2 \geq L$. The inequality holds for all such M_1, M_2 if it holds for the largest M_1 and smallest M_2 , that is, for $M_1 = v - v/e$ and $M_2 = L$. Bidder 1 always bids M_1 in the oracle if $\bar{p} \geq e(v - L)/v$ and M_2 is irrelevant for loss. Under the proposed bidding function and the restriction on \bar{p} and v , loss is

$$\bar{p}(v - M_1) - \bar{p} \int_{M_1}^{v - \frac{v}{e}} db = \bar{p} \frac{v}{e}.$$

Nature chooses bid distributions of the form $\bar{p}[M_1] + (1 - \bar{p}[M_2])$, where $M_1 \in [0, L]$ and distributed according to the cdf specified in the proof of Proposition 6 with support $[0, \frac{v}{e}]$ and $M_2 \geq L$ arbitrary. \square

Example 5. Suppose one knows that one has a value below the median bid L , that other bidders bid independently and that there are n bidders. Then the proposition says that minimax loss is $(1/2)^{n-1}v/e$. Then for two bidders minimax loss is approximately $0.18v$, with five bidders it is $0.02v$ and loss is at most $0.0007v$ with ten bidders.

Note that the condition $\bar{p} \geq \frac{v-L}{v}e$ implies $\bar{p} \geq \frac{v-L}{v-L+eL}$. Hence, for v such that $L < v \leq \frac{e}{e-1}L$ Propositions 7 and 8 cannot hold at the same time.

B.3 The ε -Uniform Model

Minimax loss is also lower in the ε -uniform model of Section 4.3.

Proposition 9. *Let there be $n \geq 2$ bidders, $v > L \geq 0$, and $\varepsilon \in (0, 1)$. Let $\alpha(b) = v(1-\varepsilon) + b\varepsilon - L$ and $\beta(b) = \varepsilon(b-L)$. For the set of conceivable environments $\mathcal{E}_{L,\varepsilon,n}$ minimax loss is attained by the randomized bidding strategy with density conditional on v given by*

$$g(b|v) = \frac{\alpha(b)^{n-1}\beta(b)(v(1-\varepsilon n) + b\varepsilon n - L)}{(v-b)((\varepsilon-1)v\beta(b)^n + b\varepsilon(\alpha(b)^n - \beta(b)^n) + L(\beta(b)^n - \varepsilon\alpha(b)^n))} \quad (15)$$

for $b \in [\underline{b}, \bar{b}]$, where $\underline{b} = \max\{L, \tilde{b}\}$ and \bar{b} solves $\int_{\underline{b}}^{\bar{b}} g(b|v) db = 1$. Minimax loss equals

$$\pi(\bar{b}) - \varepsilon^{n-1} \int_{\underline{b}}^{\bar{b}} g(b|v) \left(\frac{b-L}{v-L}\right)^{n-1} (v-b) db. \quad (16)$$

Proof. The oracle is as in Proposition 3. In particular, Lemma 1 has direct consequences on the maximization of loss. Suppose $\varepsilon > 1/n$ so that $\tilde{b} > L$ and that

bidder 1 uses a randomized bidding strategy with support $[\underline{b}, \bar{b}]$, where $\underline{b} < \tilde{b}$. If it turns out that the highest bid among the other bidders is $M \in [L, \tilde{b})$, then bidder 1's optimal bid in the oracle is \tilde{b} , so loss equals

$$\pi(\tilde{b}) - \varepsilon^{n-1} \int_L^v (v-b)p(b)^{n-1} dG(b|v) - \sum_{k=0}^{n-2} \binom{n-1}{k} \varepsilon^k (1-\varepsilon)^{n-1-k} \int_M^v (v-b)p(b)^k dG(b|v).$$

Loss is increased by $M' \in (M, \tilde{b})$, because this does not change the oracle payoff, but decreases the chance of winning under unknown M' . Hence, maximal loss is attained by $M \geq \tilde{b}$ and given by

$$\pi(M) - \varepsilon^{n-1} \int_L^v (v-b)p(b)^{n-1} dG(b|v) - \sum_{k=0}^{n-2} \binom{n-1}{k} \varepsilon^k (1-\varepsilon)^{n-1-k} \int_M^v (v-b)p(b)^k dG(b|v). \quad (17)$$

This is also the maximal loss when $\varepsilon \leq 1/n$. To summarize, loss cannot be maximized if the highest bid among the other bidders is below $\max\{L, \tilde{b}\}$.

Taking the first derivative with respect to M of Equation (17) and solving the first order condition for $g(M|v)$ leads to bidder 1 using the density

$$g(b|v) = \frac{\alpha(b)^{n-1} \beta(b) (v(1-\varepsilon n) + b\varepsilon n - L)}{(v-b) ((\varepsilon-1)v\beta(b)^n + b\varepsilon(\alpha(b)^n - \beta(b)^n) + L(\beta(b)^n - \varepsilon\alpha(b)^n))},$$

where $\alpha(b) = v(1-\varepsilon) + b\varepsilon - L$ and $\beta(b) = \varepsilon(b-L)$.

In order to determine the support of the random bidding function one needs to distinguish between $\varepsilon < \frac{1}{n}$ and its converse. Let $\underline{b} = \max\{\tilde{b}, L\}$. Then the support of g is $[\underline{b}, \bar{b}]$, where \bar{b} solves $\int_{\underline{b}}^{\bar{b}} g(b|v) db = 1$. Under this bidding function, loss is guaranteed to be below

$$\pi(\bar{b}) - \varepsilon^{n-1} \int_{\underline{b}}^{\bar{b}} g(b|v) \left(\frac{b-L}{v-L}\right)^{n-1} (v-b) db,$$

as loss must be the same for all $M \in [\underline{b}, \bar{b}]$. Evaluate loss at $M = \bar{b}$. The function g never selects a winning bid, hence one only wins if all other bids are uniformly drawn.

Loss is maximized if nature can choose the distribution of M and has preferences for higher losses, i.e. if bidder 1 played a zero-sum game against nature, where bidder 1 wants to minimize loss and nature wants to maximize bidder 1's loss. Nature's strategy H is missing for the proof of the Proposition. In the equilibrium of a zero-sum game, a player must be indifferent between two actions,

therefore, bidder 1 must be indifferent between b and b' . Loss for bid b is

$$l(b|v) = \int \pi(M) dH(M|v) - \varepsilon^{n-1}(v-b)p(b)^{n-1} - (v-b)H(b|v) \sum_{k=0}^{n-2} \varepsilon^k (1-\varepsilon)^{n-1-k} p(b)^k.$$

In equilibrium, it must hold that $l(b|v) = l(b'|v) = l(\underline{b}|v)$. Solving for $H(b|v)$ yields

$$H(b|v) = \frac{(v - \underline{b}) (H(\underline{b}|v) \gamma(\underline{b})^{n-1} + (1 - H(\underline{b}|v)) \delta(\underline{b})^{n-1}) - (v - b) \delta(b)^{n-1}}{(v - b) (\gamma(b)^{n-1} - \delta(b)^{n-1})},$$

where $\gamma(b) = \frac{v(1-\varepsilon)+\varepsilon b-L}{v-L}$ and $\delta(b) = \varepsilon p(b)$. Loss is decreased for $b > \bar{b}$, thus $H(\bar{b}|v) = 1$. From the last equation one can solve for $H(\underline{b}|v)$. \square

Example 6. Table 4 provides numerical calculations for $v = 1$, $L = 0$ and different values of ε and n . For every ε and n the table reports the support of the random bidding function, the mean bid, and the upper bound on loss, all rounded to two decimals. One interesting feature of the model is that the expected bid is increasing in the number of bidders. As one might expect, loss is decreasing in ε and n . The maximal loss under random bidding is approximately 74% of the maximal loss under deterministic bidding.

Figure 2 shows the density of the bidding function of Equation (15) for $n = 2$ (dotted), $n = 5$ (dashed), and $n = 10$ (solid), where $\varepsilon = 0.15$, $L = 0$, and $v = 1$. One can see that for $n = 2$ and $n = 5$ the lower bound of the bidding function is 0, but not for $n = 10$. For $n = 10$, $\tilde{b} > 0$ and therefore \tilde{b} is the lowest possible bid. As the number of other bidders increases, more mass is put on higher bids.

B.4 Comparison with Experimental Data

Table 5 summarizes the mean and max loss of the randomized bidding functions of the model with complete uncertainty and the ε -uniform model. The randomized bid recommendations tend to lead to a higher loss on average than the corresponding deterministic bid, but sometimes achieve a lower maximal loss. For the same L , the randomized bidding functions has a mean bid lower than the deterministic minimax bid. This reduces maximal loss, but might not be ideal in other situations. For the Filiz-Özbay and Özbay (2007) data, mean loss is about the same as of the observed bids. In the experiments of Chen et al. (2007) and Güth and Ivanova-Stenzel (2003), the randomized bid recommendations would have led to a lower loss than real bids.

Randomized Bidding						
ε	$n = 2$			$n = 5$		
	Support	Mean Bid	Loss	Support	Mean Bid	Loss
0.10	[0, 0.64]	0.38	0.32	[0, 73]	0.47	0.24
0.15	[0, 0.65]	0.39	0.30	[0, 78]	0.54	0.19
0.20	[0, 0.65]	0.40	0.28	[0, 0.82]	0.62	0.15
0.25	[0, 0.66]	0.41	0.26	[0.2, 0.86]	0.70	0.12
0.40	[0, 0.68]	0.46	0.19	[0.5, 0.91]	0.81	0.07
0.50	[0, 0.70]	0.51	0.14	[0.6, 0.93]	0.85	0.06
	$n = 10$					
	Support	Mean Bid	Loss			
0.10	[0, 0.83]	0.64	0.14			
0.15	[0.33, 0.89]	0.76	0.10			
0.20	[0.5, 0.92]	0.82	0.07			
0.25	[0.6, 0.93]	0.85	0.06			
0.40	[0.75, 0.96]	0.91	0.04			
0.50	[0.80, 0.97]	0.93	0.03			

Table 4: Minimax Bids and Loss for different values of ε and n with $L = 0$ and $v = 1$




Figure 2: Probability density function of Equation (15) for $n = 2$ (dotted), $n = 5$ (dashed), and $n = 10$ (solid), where $\varepsilon = 0.15$, $L = 0$, and $v = 1$

	L	Loser Feedback		Winner Feedback		No Feedback	
		Mean	Max	Mean	Max	Mean	Max
		Complete Uncertainty	0	2.95	9.47	3.66	10.37
ε -uniform	0	2.02	6.22	2.50	7.18	2.46	6.72

(a) Filiz-Özbay and Özbay (2007) data

	L	Known Distribution		Unknown Distribution	
		Mean	Max	Mean	Max
		Complete Uncertainty	0	5.14	9.17
ε -uniform	0	4.88	8.67	5.56	9.25

(b) Chen et al. (2007) data

	L	Known Distribution				Unknown Distribution		
		Weak Bidder		Strong Bidder		L	Mean	Max
		Mean	Max	Mean	Max			
Complete Uncertainty	25	4.12	6.44	6.30	11.01	0	8.32	14.27
ε -uniform	25	3.84	5.85	5.94	10.53	0	7.75	13.19

(c) Güth and Ivanova-Stenzel (2003) data

Table 5: Performance of randomized bidding strategies

C Data

This appendix supplements Section 6 in the description of the data and the methods, and provides statistical tests. In all three experiments every subject participated only in one treatment. Filiz-Özbay and Özbay (2007) test anticipated regret in the first-price auction. For this purpose they run three treatments that differ in the pre-auction announcement of the post-auction feedback about the winner’s or loser’s bid. The authors find that bidders who know that they will learn the winning bid (loser feedback treatment) significantly higher than in the other two treatments. The strategy method is used to elicit bidding functions. Hence, each bidder receives a list of ten values and has to indicate a bid for each value. The list for bidder $i = 1, \dots, 4$ is the same in every auction, hence there are 40 different values. We generate the empirical bid distribution B_i by pooling all bids submitted by a bidder with value list $i = 1, \dots, 4$. For a bidder with value list i the probability of bid b to become winning is a function of B_{-i} and tie-breaking with equal probabilities of winning. There are $m \in \{7, 8, 9\}$ markets in a treatment, so there are m bidders with the same value. The empirical bid distribution contains $30 \cdot m$ observations. For a player with value v who bid b in the experiment we simply compute the actual loss $l(b|v)$. We compute the mean loss for the m bidders with the same value. The loss in per cent of the value is then $l(b|v)/v \cdot 100$.

In Chen et al. (2007) bidders repeatedly receive values and bid. Values and bids are integers from 0 to 100. We only consider the data from the first round of bidding in order to ensure independent observations. So our data only consists of one observation (value and bid) per subject and treatment. There are 40 independently drawn values per treatment. We compute for bidder i the empirical bid distribution B_j that she faces, that is, we calculate for every bidder a cdf that excludes the bidder’s bid and use it to compute the probability of winning. The bid distribution faced by a bidder therefore contains 39 bids.

Güth and Ivanova-Stenzel (2003) compare bidding in a first-price auction with and without common knowledge of the value distribution. In one treatment the value distributions are known, in the other treatment the value distributions are unknown. They find very similar behavior in both treatments. We only use the data from the first round of bidding to have independent observations. We have 35 independent values for the weak and the strong bidder, respectively. For every treatment, we use all the bids of bidder 2 to compute the empirical bid distribution faced by bidder 1 and all the bids of bidder 1 to get the empirical bid distribution

	FÖÖ (2007)			CKO (2007)		GIS (2003)		
	L	W	No	K	UNK	W	S	UNK
WMW	0.01	0.00	0.00	0.70	0.81	0.43	0.07	0.56
S-IEQ at 5%	=	<	<	=	=	=	=	=
S-IEQ at 10%	<	<	<	=	=	=	=	=

(a) Comparison eps and L0

	FÖÖ (2007)			CKO (2007)		GIS (2003)		
	L	W	No	K	UNK	W	S	UNK
WMW	0.37	0.99	0.65	0.76	0.01	0.00	0.00	-
S-IEQ at 5%	=	=	=	=	=	<	<	-

(b) Comparison eps and Nash

Table 6: Significance Tests

faced by bidder 2. The used empirical bid distribution contains 35 independent bids.

We run two types of tests to test for statistically significant differences in loss. The results are summarized in Table 6. First, we use the Wilcoxon-Mann-Whitney (WMW) test to examine whether the distribution of mean loss is the same for the ε -uniform model (eps) and the model of complete uncertainty (L0) and Nash in a treatment.¹⁵ The table reports p values. First, we compare eps and L0. The null that eps and L0 are identically distributed is rejected for the treatments of Filiz-Özbay and Özbay (2007). For the other treatments the data is found to be indistinguishable. As a next step, we use an exact stochastic inequality test (Schlag, 2008, S-IEQ) to verify that mean loss is indeed statistically less under eps than L0. This is true for the Winner and No Feedback treatment when the level of a type I error is fixed at 5% and true for the Loser treatment when the significance level is 10%. Now we compare eps and Nash. The WMW test suggests that mean loss is not identically distributed in the treatments of Chen et al. (2007) and in the known treatment of Güth and Ivanova-Stenzel (2003). The stochastic inequality test finds that in the latter case the mean loss is indeed significantly less under eps than Nash.

¹⁵For the tests one needs independent observations. As a result, we have restricted ourselves to observations of the first round only. Nevertheless, we basically use the same empirical bid distribution in the computation for mean loss for eps and L0 (Nash) and for every different value. Strictly speaking, we have to assume independence. One could develop an exact test for dependent variables, but this goes beyond the scope of this exercise.

D Risk-Aversion

D.1 Complete Uncertainty

In this appendix, we derive robust bid recommendations for risk averse bidders. Let u denote bidder 1's non-decreasing Bernoulli utility function with the normalization that the utility of losing is equal to 0. Under complete uncertainty, with $L \geq 0$, loss is

$$l(b, M|v) = \max\{u(v - M), 0\} - \int_L^M u(v - b) dG(b|v) - \int_M^v 0 dG(b|v).$$

Taking the first derivative with respect to M yields the density of the optimal bidding function

$$g(b|v) = \frac{u'(v - b)}{u(v - b)} \text{ on } [L, \bar{b}].$$

The upper bound of the support is determined by $\bar{b} \leq v$ that solves $\int_L^{\bar{b}} g(b|v) db = 1$. In the following we focus on two popular parameterized utility functions.

D.1.1 CRRA

Let $u(x) = x^{1-\rho}$, $0 \leq \rho < 1$. Risk neutrality corresponds to $\rho = 0$. The density of the bidding function is

$$g(b|v) = \frac{1 - \rho}{v - b} \text{ on } [0, v - e^{\frac{1}{\rho-1}}(v - L)].$$

Loss is bounded by

$$(v - M)^{1-\rho} - \int_M^{v - e^{\frac{1}{\rho-1}}(v-L)} (v - b)^{1-\rho} g(b) db = \frac{(v - L)^{1-\rho}}{e}.$$

The mean bid

$$\int_L^{v - e^{\frac{1}{\rho-1}}(v-L)} b dG(b|v) = v - \left(1 - e^{\frac{1}{\rho-1}}\right) (1 - \rho)(v - L)$$

is increasing in ρ . The more risk averse, the higher the bid on average, because larger bids reduce the risk of not winning. When $L = 0$, then for $\rho \geq 0.7$ we have that the expected bid is approximately equal to $\rho \cdot v$.

Deterministic minimax loss is equal to $(v - L)^{1-\rho}/2$ and attained by the bidding strategy $b^*(v) = v - (v - L)2^{\frac{1}{\rho-1}}$, because this bid balances the loss of bidding too

high and too low, i.e.

$$b^*(v) \in \arg \min_b \max\{(v-b)^{1-\rho}, (v-L)^{1-\rho} - (v-b)^{1-\rho}\}.$$

Minimax loss can be increasing or decreasing in the level of risk aversion. It is strictly increasing in the risk aversion parameter ρ when $v-L < 1$, so when v is small relative to L . For v large relative to L minimax loss is strictly decreasing. Minimax loss normalized by $u(v-L)$ is $1/e$ and $1/2$ for the randomized and deterministic case, respectively.

D.1.2 CARA

Under constant absolute risk aversion, $u(x) = \frac{1-e^{-\alpha x}}{1-e^{-\alpha}}$, $\alpha > 0$, the probability density function of the bidding function is

$$g(b|v) = \frac{\alpha e^{\alpha(-(v-b))}}{1-e^{\alpha(-(v-b))}} \text{ on } \left[L, \frac{\log(e^{\alpha L-1} + e^{\alpha v} - e^{\alpha v-1})}{\alpha} \right].$$

Loss is not higher than $\frac{e^{\alpha(1-e^{-\alpha(v-L)})}}{e^{\alpha(e^{\alpha}-1)}}$. The deterministic minimax bid is $b^*(v) = \log\left(\frac{1}{2}(e^{\alpha L} + e^{\alpha v})\right)/\alpha$ and deterministic minimax loss equals $\frac{e^{\alpha(1-e^{-\alpha(v-L)})}}{2(e^{\alpha}-1)}$.

Loss is increasing in α when v is relatively small in relation to L , i.e. when $e^{\alpha L}(e^{\alpha(v-L)} + L - v + 1) > e^{\alpha v}$. Otherwise minimax loss is decreasing. Minimax loss normalized by $u(v-L)$ is $1/e$ under randomization and $1/2$ under deterministic bidding.

D.2 The ε -Uniform Model

D.2.1 CRRA

The ε -uniform model can also be solved for a CRRA Bernoulli utility function $u(x) = x^{1-\rho}$. The derivation of minimax loss is analogous to the risk-neutral case. One only has to change the Bernoulli utility function. Payoff conditional on bidding higher than M equals

$$\pi(b|b > M) = \sum_{k=0}^{n-1} \binom{n-1}{k} \varepsilon^k (1-\varepsilon)^{n-1-k} u(v-b)p(b)^k,$$

and is maximized by $\max\{M, \tilde{b}\}$, where $\tilde{b} = \frac{v(n\varepsilon - \varepsilon\rho + \rho - 1) + L(1-\rho)}{\varepsilon(n-\rho)} \geq L$ if and only if $n \geq (1 - \rho(1 - \varepsilon))/\varepsilon$. The optimal bid \tilde{b} is strictly increasing in ρ , so the more

risk averse, the higher the optimal bid. Let $\delta(b) = v(1 - \varepsilon) + b\varepsilon - L$.

Minimax loss is equal to

$$\pi(\bar{b}|\bar{b} > M) - \varepsilon^{n-1} \int_{\underline{b}}^{\bar{b}} g(b|v)p(b)^{n-1}u(v-b) db \quad (18)$$

and attained by the mixed strategy with density

$$g(b|v) = \frac{\varepsilon(b-L)p(v(1-\varepsilon)+b\varepsilon)^n((v-L)(1-\rho)-(v-b)\varepsilon(n-\rho))}{(v-b)\delta(b)(\varepsilon(b-L)p(v(1-\varepsilon)+b\varepsilon)^n - \varepsilon\delta(b)p(b)^n)}$$

on $[\underline{b}, \bar{b}]$, with $\underline{b} = \max\{L, \tilde{b}\}$ and \bar{b} such that $\int_{\underline{b}}^{\bar{b}} g(b|v) db = 1$.

D.2.2 CARA

Finding closed form solutions for the ε -uniform model with CARA utility function $u(x) = \frac{1-e^{-\alpha x}}{1-e^{-\alpha}}$ is more challenging. The payoff conditioned on bidding higher than M is maximized by $\max\{\tilde{b}, M\}$, with $\tilde{b} = \frac{\varepsilon - \varepsilon n + \varepsilon W((n-1)e^{n+\frac{\alpha v}{\varepsilon}-1}) + \alpha \varepsilon v + \alpha(-v)}{\alpha \varepsilon}$, where $W(\cdot)$ is the Lambert-W function.

The density of the mixed strategy at which minimax loss is attained is equal to

$$g(b|v) = \frac{\varepsilon(b-L)p(v(1-\varepsilon)+b\varepsilon)^n (\varepsilon((n-1)e^{\alpha(v-b)} - n + 1) - \alpha(v(1-\varepsilon) + b\varepsilon - L))}{(1 - e^{\alpha(v-b)})\delta(b)(\varepsilon(b-L)p(v(1-\varepsilon)+b\varepsilon)^n - \varepsilon\delta(b)p(b)^n)}$$

and has support $[\underline{b}, \bar{b}]$, with $\underline{b} = \max\{L, \tilde{b}\}$ and \bar{b} such that $\int_{\underline{b}}^{\bar{b}} g(b|v) db = 1$. Minimax loss is given by Equation (18).

References

- Battigalli, Pierpaolo and Marciano Siniscalchi**, “Rationalizable Bidding in First Price Auctions,” *Games and Economic Behavior*, October 2003, 45 (1), 38–72.
- Bergemann, Dirk and Karl Schlag**, “Pricing without Priors,” *Journal of the European Economic Association*, 2008, 6 (2-3), 560–569.
- **and –**, “Robust Monopoly Pricing,” *Journal of Economic Theory*, 2011, 146 (6), 2527–2543.

- Binmore, Ken and Paul Klemperer**, “The Biggest Auction Ever: The Sale Of The British 3G Telecom Licences,” *The Economic Journal*, March 2002, *112* (478), C74 –C96.
- Caldentey, Rene, Ying Liu, and Ilan Lobel**, “Intertemporal Pricing Under Minimax Regret,” *Operations Research*, January-February 2017, *65* (1), 104–129.
- Carroll, Gabriel**, “Robustness and Linear Contracts,” *American Economic Review*, February 2015, *105* (2), 536–563.
- Chen, Yan, Peter Katuscak, and Emre Ozdenoren**, “Sealed Bid Auctions with Ambiguity: Theory and Experiments,” *Journal of Economic Theory*, 2007, *136*, 516–535.
- Cho, In-Koo**, “Monotonicity and Rationalizability in a Large First Price Auction,” *Review of Economic Studies*, 2005, *72* (4), 1031–1055.
- Compte, Olivier and Andrew Postlewaite**, “Simple Auctions,” *Working paper*, 2013.
- Dekel, Eddie and Asher Wolinsky**, “Rationalizable Outcomes of Large Private-Value First-Price Discrete Auctions,” *Games and Economic Behavior*, May 2003, *43* (2), 175–188.
- Engelbrecht-Wiggans, Richard**, “The Effect of Regret on Optimal Bidding in Auctions,” *Management Science*, June 1989, *35* (6), 685–692.
- **and Elena Katok**, “Regret and Feedback Information in First-Price Sealed-Bid Auctions,” *Management Science*, April 2008, *54* (4), 808–819.
- Filiz-Özbay, Emel and Erkut Y. Özbay**, “Auctions with Anticipated Regret: Theory and Experiment,” *American Economic Review*, September 2007, *97* (4), 1407–1418.
- Güth, Werner and Radosveta Ivanova-Stenzel**, “Asymmetric Auction Experiments With(out) Commonly Known Beliefs,” *Economics Letters*, August 2003, *80* (2), 195–199.
- Halpern, Joseph Y. and Rafael Pass**, “Iterated Regret Minimization: A New Solution Concept,” *Games and Economic Behavior*, January 2012, *74* (1), 184–207.

- Harrison, Glenn W.**, “Theory and misbehavior of first-price auctions,” *The American Economic Review*, September 1989, 79 (4), 749–762.
- Hayashi, Takashi**, “Regret Aversion and Opportunity Dependence,” *Journal of Economic Theory*, March 2008, 139 (1), 360–379.
- Huber, Peter J.**, “A Robust Version of the Probability Ratio Test,” *The Annals of Mathematical Statistics*, December 1965, 36 (6), 1753–1758.
- , *Robust Statistics*, New York: John Wiley and Sons, 1981.
- Jiang, Houyuan, Serguei Netessine, and Sergei Savin**, “Robust Newsvendor Competition Under Asymmetric Information,” *Operations Research*, January-February 2011, 59 (1), 254–261.
- Kagel, John H. and Dan Levin**, “Auctions: A Survey of Experimental Research,” in John H. Kagel and Alvin E. Roth, eds., *The Handbook of Experimental Economics*, Vol. 2, Princeton University Press, 2016, chapter 9.
- Levin, Dan and Emre Ozdenoren**, “Auctions with Uncertain Numbers of Bidders,” *Journal of Economic Theory*, October 2004, 118 (2), 229–251.
- Linhard, P. B. and Roy Radner**, “Minimax-Regret Strategies for Bargaining over Several Variables,” *Journal of Economic Theory*, June 1989, 48 (1), 152–178.
- Lo, Kin Chung**, “Sealed Bid Auctions with Uncertainty Averse Bidders,” *Economic Theory*, July 1998, 12 (1), 1–20.
- Loomes, Graham and Robert Sugden**, “Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty,” *The Economic Journal*, December 1982, 92 (368), 805–824.
- Milnor, John**, “Games Against Nature,” in R. H. Thrall, C. H. Coombs, and R. L. Davis, eds., *Decision Processes*, Wiley, 1954, chapter 4.
- Perakis, Georgia and Guillaume Roels**, “Regret in the Newsvendor Model with Partial Information,” *Operations Research*, 2008, 56 (1), 188–203.
- Renou, Ludovic and Karl Schlag**, “Implementation in Minimax-Regret Equilibrium,” *Games and Economic Behavior*, March 2011, 71 (2), 527–533.

- Robles, Jack and Makoto Shimoji**, “On Rationalizability and Beliefs in Discrete Private-Value First-Price Auctions,” *The B.E. Journal of Theoretical Economics*, 2012, 12 (1).
- Rothkopf, Michael H.**, “Decision Analysis: The Right Tool for Auctions,” *Decision Analysis*, 2007, 4 (3), 167–172.
- Savage, Leonard J.**, “The Theory of Statistical Decision,” *Journal of the American Statistical Association*, March 1951, 46 (253), 55–67.
- Schlag, Karl and Andriy Zapechelnjuk**, “Robust sequential search,” *Working paper*, 2016.
- Schlag, Karl H.**, “A New Method for Constructing Exact Tests without Making any Assumptions,” *Universitat Pompeu Fabra, Department of Economics and Business Working Paper 1109* 2008.
- Sošić, Greys**, “Collusion in Second-Price Auctions under Minimax Regret Criterion,” *Production and Operations Management*, 2007, 16 (4), 471–482.
- Syrkkanis, Vasilis and Eva Tardos**, “Composable and Efficient Mechanisms,” *44th Symposium on Theory of Computing Conference, STOC’13*, 2013, pp. 211–220.
- Vickrey, William**, “Counterspeculation, auctions, and competitive sealed tenders,” *The Journal of Finance*, March 1961, 16 (1), 8–37.
- Wilson, Robert B.**, “Game-Theoretic Analyses of Trading Processes,” in Truman Bewley, ed., *Advances in Economic Theory: Fifth World Congress*, Econometric Society Monographs, Cambridge University Press, 1987, chapter 2, pp. 33–70.