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

Bernhard Kasberger Karl Schlag

November 8, 2016

Very Preliminary and Incomplete

Abstract

Finding optimal bids in first-price auctions in the classical framework requires detailed information and specific assumptions about the other bidders' value distributions and bidding behavior. This paper shows how to bid with less information. A bidding rule is evaluated by comparing the payoff of the rule to the payoff that could be achieved if one knew the other bidders' value distributions and bidding functions. Robust bidding approximates the payoff under more information by minimizing the highest payoff difference. We derive robust bidding rules under complete uncertainty and for cases in which one imposes bounds on the bid or value distributions of the other bidders.

1 Introduction

We investigate how to bid in a first-price auction. The rules of the first-price auction are simple and commonly known.[1] Yet, it is very hard to bid optimally as lots of other aspects are unknown. A bidder might be uncertain about the number of other bidders, their value distribution, the way

*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

[1]In 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.

they translate values into bids, and their beliefs about other bidders' values and behavior. First, we propose how to bid if no further assumptions on the value distribution and bidding behavior are made. Then we introduce restrictions on the anticipated bid distribution. Finally, we formulate constraints on other bidders' value distributions and on their bidding behavior. The specific application will identify whether it is simpler to incorporate knowledge and perceptions of the environment by specifying possible bid distributions or by specifying possible values and bidding behavior.

This is the first paper that makes theory based explicit recommendations when the value distribution is unknown and strategic uncertainty is not resolved in equilibrium. Each of our recommendations comes together with an error bound that specifies an upper bound on how much more surplus one could get if one knew the joint value distribution and bidding behavior. Interestingly, we can make a recommendation even for the case in which one does not wish to make any assumptions about values and/or bidding behavior of others. In this case, it is as if we are comparing the payoff to the case where we would know the bids of all others. The error bound of our proposed randomized bidding strategy is 36% of the own value v. In contrast, the error bound of a slightly misspecified Nash equilibrium bidding function can be 100% of the own value.

Our objective is to derive bidding functions that minimize the error bound for a given set of conceivable environments. An environment is a pair that generates the faced bid distribution and consists of the joint value distribution and the other bidders' bidding functions, that is their rules that transform values into bids. The loss for a given environment is the difference between the hypothetical oracle payoff, that is the highest possible surplus if the environment is known, and the surplus generated by the bidding function. The error bound is the maximal loss; bidding behavior is chosen to minimize this maximal loss.

We develop two approaches to reduce the error bound obtained under complete uncertainty. A first step is to think about the driving forces of the stated result. If all environments are possible, then one has to deal with the case in which all others bid 0 for sure, which generates a large loss for bidding high. An immediate reaction is that this situation is not very likely. We address this in two different ways, by imposing lower bounds and by imposing variability of bids.

First, one can imagine that one has enough information about the context that one is willing to rule out that the maximal bid of the others will be below some threshold L. In this case we provide a bidding rule that guarantees a loss below 36% of $(v - L)$. In particular this means that if one does not expect that the maximal bid is below 73% of one's own value then the error bound will be at most 10% of v. However, a threshold L for which one is willing to rule out with certainty that the maximal bid will not be below L may be very small. Clearly, one could choose such a threshold larger if one is allowed to assign some maximal probability that the maximal bid is below L. We find that the 10% of v error bound can be sustained if one assigns at most 12% to the possibility that the maximal bid will be below 73% of v. The reason why the error bound is smaller is because the constraint on the probability of the maximal bid being below L must also hold for the true environment.

Second, we consider bidding when one adds some assumptions about the bids above L, maintaining the assumption that the maximal bid will not be below L. These assumptions will concern both independence of bidding and the variability of bids. Independently for each other bidder, with probability ε, this bidder is believed to bid between L and v with each bid being equally likely. With probability $1 - \varepsilon$ that bidder can do anything as long as she bids above L. So some bidders are believed to bid independently and uniformly on $[L, v]$, while no assumptions are made on what the other bidders do. Hence, it is not possible that all bidders bid L with certainty. We find, for instance, if ε is at least 0.15 and there are 10 bidders and $L = 0$ (or there are 5 bidders and $L = v/2$) that loss is bounded above by 10% of value v. The parameter ε measures the weight on the particular beliefs and influences the variability of bids that are expected.

Bounds on the value distribution and on bidding behavior can further improve the error bound. A bound either on the value distribution or on bidding behavior cannot reduce the uncertainty sufficiently much to reduce maximal loss. The combination, however, can produce sharp results. We investigate the case in which the value distribution can be bounded above so that very low values cannot be drawn with certainty, and the other bidders' bidding functions can be bounded from below so that not all bidders bid very low with certainty. In this case, a linear bidding function can bound the error from above. Remarkably, this bidding function only depends on the parameter used to bound the value distribution and on the number of bidders. What is more, if other bidders are perceived to use the same linear bidding function, this bidding function remains to perform well. Thus, behavior forms an *ϵ-loss-equilibrium*, that is similar in spirit to ϵ-Nash equilibrium. The advantage of this solution concept is that it is applicable to very complicated settings in which (strategic) uncertainty is present.

Related Literature

Minimizing the maximal loss is a concept introduced by Savage (1951) for decision problems and 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. The literature on minimax regret (e.g. Hayashi, 2008) has, for example, looked at the pricing problem of a monopolist (Bergemann and Schlag, 2008, 2011), and dynamically consistent robust search rules (Schlag and Zapechelnyuk, 2016). In strategic settings, Linhard and Radner (1989) consider bargaining, and Renou and Schlag (2011) and Halpern and Pass (2012) develop solution concepts for games where it is common knowledge all players follow minimax regret. In the last section of our paper we also make recommendations in an equilibrium style framework, but remain less precise in terms of anticipated beliefs and bidding behavior of others. Some of the papers are discussed in more depth in the text.

We talk about loss and not regret because, first, our evaluation of performance has no behavioral context as the term "regret" might suggest, and second, because in the context of auctions the term regret is used differently by Engelbrecht-Wiggans (1989). In this literature (e.g. Filiz-Özbay and Özbay, 2007; Engelbrecht-Wiggans and Katok, 2008) the value distributions is known and the objective is the maximization of expected utility plus additive regret terms. Regret arises from learning other bids.

In our setting, the benchmark of learning the environment of value distribution and bidding function is purely hypothetical and there is no additive "behavioral" motive. This benchmark is introduced to measure the compromise that a bidder who believes to be better informed has to undergo when following our recommendations.

Our approach should not be confused with recent literature on robust decision making (e.g. Carroll, 2015) that borrowed the term "robust" (Huber, 1981) for a different context. Common is the objective to find a rule for making decisions without a prior. However, in this alternative strand of literature the benchmark to be close to the optimal policy is dropped. Instead, preferences are introduced for how to deal with multiple priors without connecting this to the optimal policy. The plausibility of the rule depends also on how plausible these alternative preferences are. Strategic uncertainty is resolved in equilibrium. In the context of the first-price auction, Lo (1998) and Chen et al. (2007) present models in which symmetric bidders have a set of priors over the value distribution. In the first paper, bidders have maxmin preferences. In the second paper, bidders' preferences are a generalization of maximin expected utility that allows ambiguity loving. Levin and Ozdenoren (2004) study auctions with an uncertain number of bidders who have maximin preferences. Interestingly, experimental results by Güth and Ivanova-Stenzel (2003) and Chen et al. (2007) indicate that bidding behavior is very similar with and without priors.

There are papers that study rationalizable bids in first-price auctions in which bidders have a prior over the value distribution. Battigalli and Siniscalchi (2003) show that all nonzero bids below, and some bids above equilibrium are rationalizable. For sufficiently many bidders, rationalizable bids are very close to value (Dekel and Wolinsky, 2003; Cho, 2005). When bidders believe that others' values are only slightly lower than the own value, then bidding close to value is rationalizable (Robles and Shimoji, 2012). We neither assume common knowledge of rationality, nor common knowledge of the beliefs about the value distribution.

Other approaches to complete-information games with non-rational bidders are Eliaz (2002) and de Clippel (2014). Work on robust mechanism design, initiated by Bergemann and Morris (2005), loosens the prior while maintaining rationality assumptions. Bergemann et al. (2016) analyze properties of equilibrium play in the first-price auction that hold for all information structures for a common prior over the joint value distribution.

The formal methodology used for this study is described in the next section. Section 3 presents simple examples of the non-robustness of bidding functions that are optimal in specific environments under expected and maximin utility. The core of the paper begins in Section 4. Various beliefs about the bid distribution are used to derive bid recommendations. In Section 5 more specific beliefs about the bidding behavior and the value distribution are made.

2 Methodology

We consider the bidding behavior of bidder 1 in a first-price seal bid auction for a single indivisible good. Bidder 1 has a value $v_1 \in \mathbb{R}_+$ for the auctioned object and quasilinear preferences that have a von Neumann–Morgenstern expected utility representation. The utility of losing the auction is normalized to zero. Winning the auction with bid b_1 yields utility $u_1(v_1 - b_1)$. Bidder 1 employs the (possibly mixed) bidding function $b_1 : \mathbb{R}_+ \to \Delta\mathbb{R}_+$ that maps the own value v_1 into (a distribution of) bids. Likewise, every other bidder $i > 1$ has a value v_i for the object and bids according to a certain bidding function. We will be more specific about the other bidders.

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 two inputs. First, there is a true and exogenous joint value distribution $F \in \Delta\mathbb{R}^n_+$, with n denoting the number of bidders. Second, each participant's bidding behavior translates values into bids. The bidding behavior can depend on the own value and on information about other bidders. Traditional analysis of bidding behavior in auctions assumes that all bidders know the value distribution and other bidders' risk attitudes and that there is common knowledge of rationality. One then searches for an equilibrium in which each bidder best responds to the bidding behavior of the other bidders. In short, every bidder knows the value distribution and, in equilibrium, the other bidders' bidding strategies.

We depart from the classic setting and investigate how bidder 1 should bid if she is uncertain both about the value distribution and the bidding behavior of the others. This uncertainty is modeled in the form of bidder 1 identifying a set of conceivable environments $\mathcal{E}$. An environment $E \in \mathcal{E}$ is a pair (F, b_{-1}), where F is a joint value distribution and b_{-1} specifies the bidding behavior of other bidders. Hence, an environment generates the bid distribution faced by bidder 1. What follows is the formal description of the two components of environments and the set of conceivable environments.

Bidder 1 does not know the true F, but she conceives that F belongs to the class of joint distributions $\mathcal{F}$ with $\mathcal{F} \subseteq \bigcup_{n \in \mathbb{N}} \Delta\mathbb{R}_+^n$. In particular, one can incorporate uncertainty over the number of bidders n by including distributions with different number of bidders in $\mathcal{F}$. Complete uncertainty about the other's values can be modeled by $\mathcal{F}$ being the set of all joint distributions. Alternatively, the bidder might assert that there are n bidders with iid values. The set of possible joint value distributions $\mathcal{F}$ would then be the set of all joint distributions that are generated from some iid distributions and denoted by $\mathcal{F}_{iid}$. The set $\mathcal{F}$ plays a role in the evaluation of bidding functions.

Bidders might be uncertain about the bidding functions of the others, just as they are uncertain about the value distribution. First, we formally introduce the bidding functions of other bidders. Then we introduce sets of bidding functions that are used to model bidder 1's uncertainty about the other participants' bidding behavior. Any bidder $i > 1$ uses the (possibly mixed) bidding function $b_i : \mathbb{R}_+^{k_i} \to \Delta\mathbb{R}_+$ for some $k_i \in \mathbb{N}$, where $\Delta\mathbb{R}_+$ denotes the set of probability distributions over positive reals. Deterministic bidding functions are $b_i : \mathbb{R}_+^{k_i} \to \mathbb{R}_+$. Let $b_{-1} = (b_2, \ldots, b_n)$ denote the other bidders' profile of bidding functions. For example, the other bidders can bid independently, in which case $k_i = 1$ for all bidders. Another possibility is that bidder 1 thinks that the others are communicating or colluding, so $k_i > 1$.

Let B_F be the set of bidding functions that bidder 1 conceives that the other bidders use under the joint distribution F. A generic element of B_F is denoted by b_{-1}. The bidder may no nothing about the bidding behavior of others, in which case $B_F = \{(b_2, \ldots, b_n) \,|b_i : \mathbb{R}_+^{k_i} \to \Delta\mathbb{R}_+, \text{ for all } k_i \text{ and } 1 < i \leq n\}$, with n given by F. She might conceive that the other bidders use identical linear bidding functions such that they do not submit negative bids and never bid above value, i.e. $B_F = \{(b_2, \ldots, b_n) \,|b_i(v) = b_j(v) = \tau v$, for all $\tau \in [0,1]$, $i, j \neq 1\}$.

One can view environments as having a purely exogenous part F, and a potentially endogenous part b_{-1} that might depend on the behavior of other bidders. The bidder conceives that the environment she faces belongs to $\mathcal{E}' \subseteq \mathcal{E}$, where $\mathcal{E} = \bigcup_{F \in \mathcal{F}} \{F\} \times B_F$ is the universe of all environments. The set $\mathcal{E}'$ is varied throughout the text.

Note that our model about the bidder's belief about the faced environment is rich enough to include the Bayes-Nash framework. Let F be a joint distribution such that a Bayes-Nash equilibrium exits and $\mathcal{F} = \{F\}$. Let $b = (b_1 \ldots, b_n)$ be the equilibrium strategy profile, i.e. a strategy profile such that for every bidder $i \in \{1, \ldots, n\}$ the bidding function b_i is a best-response to the bidding profile b_{-i}. Then $\mathcal{E}' = \{(F, b_{-1})\}$, so the only conceivable environment within the Bayes-Nash equilibrium framework is (F, b_{-1}). Similarly, one can model the environment in which some other bidders are in a Bayes-Nash framework. This approach is included in Section 5.3.

Ideally, bidder 1 selects the best bid given the 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 knowing the environment. 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 gets and what she could get if she knew the environment. Formally, loss is given by

$$l\left(b_1, F, b_{-1} | v_1\right) = \sup_y \left\{u_1(v_1 - y) Q\left(y, b_{-1}, F\right)\right\} - \int u_1(v_1 - x) Q\left(x, b_{-1}, F\right) db_1\left(x\right),$$

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 $\sup_y \left\{u_1(v_1 - y) Q\left(y, b_{-1}, F\right)\right\}$ describes the payoff bidder 1 could (approximately) achieve if she knew F and the bidding behavior of the others. In general, loss is zero if the optimal bidding function for the true environment is chosen and bounded above by $u_1(v_1)$ 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 always defined relative to a set of conceivable environments $\mathcal{E}' \subseteq \mathcal{E}$. 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 1. *Call ϵ the value of* minimax loss *for the conceivable environments $\mathcal{E}' \subseteq \mathcal{E}$ if, for all environments in $\mathcal{E}'$, (i) loss is guaranteed to be at most ϵ, and (ii) there is no bidding function that guarantees a loss strictly lower than ϵ.*

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

3 Insensitivity of Expected and Maximin Utility

In this section we illustrate the notion of loss in a simple example and discuss the robustness of equilibrium bidding functions that are derived under expected utility maximization with a (slightly) misspecified prior. In particular, we show that a bidding rule that maximizes expected utility under a certain prior can perform very badly in another setting. Additionally, we comment on shortcomings of bidding functions that are optimal under maximin utility.

Suppose there are two risk-neutral bidders participating in a first-price auction. Their respective values are drawn independently for the value distribution F^δ parameterized by $\delta \in [0,1)$. Below we compute loss of having a slightly wrong prior belief about δ, but first we describe the set-up. The value distribution puts mass $\epsilon \in (0,1)$ uniformly on $(\delta, 1]$ and mass $1-\epsilon$ on δ. Thus, $F^\delta(x) = 0$ for $x < \delta$ and $F^\delta(x) = \min\left\{1-\epsilon+\epsilon\frac{x-\delta}{1-\delta}, 1\right\}$

for $x \geq \delta$. For a given δ, bidder 2, acting as if both bidders know F^δ, uses the bidding function

$$b_2^\delta(x) = x - \frac{\int_\delta^x F^\delta(\tilde{x})\, d\tilde{x}}{F^\delta(x)} = \frac{x^2\epsilon - \delta^2(2-\epsilon) + 2\delta(1-\epsilon)}{2(1-\delta+x\epsilon-\epsilon)}$$

that corresponds to the Bayes-Nash equilibrium bidding function. Note bidder 2 has value δ with probability $1-\epsilon$, in which case she bids her value. Let b_1^δ denote the best-response to bidding function b_2^δ. Since b_2^δ is the equilibrium bidding function, $b_1^\delta \equiv b_2^\delta$.

Consider bidder 1 with type $v_1 = 1$. If bidder 1 (wrongly) believes that $\delta = 0$, then the optimal bid is $b_1^0(1) = \epsilon/2$. This bid is never winning if the true $\delta > \epsilon/2$. The loss of the bidding strategy $b_1(1) = \epsilon/2$ is maximized if the true δ is slightly above $\epsilon/2$ and equal to

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

For ϵ close to 0, the maximal loss is approximately equal to the value, which is is the upper bound on loss, conditional on not bidding higher than v_1. Hence, for a slightly misspecified prior about the environment, the loss can be as large as possible.

Under maximin utility a uncertainty averse bidder maximizes expected utility given the worst possible prior. Strategic uncertainty is resolved in equilibrium. We show that with these preferences loss can also be as high as 100% of the own value. This is done by the construction of a simple two-bidder example parameterized by γ. 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 $1-\gamma < x \leq 1$. 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

$$b^{min}(v) = v - \frac{\int_0^v F^{min}(x)\,dx}{F^{min}(v)} = \begin{cases} \frac{v}{2} & \text{for } v < 1-\gamma_1 \\ \frac{(\gamma_1-1)\left(2\gamma_1+v^2-1\right)}{-4\gamma_1+2(\gamma_1-1)v+2} & \text{for } v \geq 1-\gamma_1 \end{cases}$$

is used by both players. Note that in particular $b^{min}(1) = 1-\gamma_1$. Consider bidder with value 1. If the true γ equals γ_2 and γ_2 is large (close to 1) then most bidders have value close to 0 and bidder 1 should bid very low to maximize her surplus. However, since her bidding follows the maximin solution b^{min} and γ_1 is small (close to 0), she bids very high. Consequently, her loss of not bidding as she would if she knew the true γ can be as large as it can get, namely it can be approximately 1, that is her value.

4 Conceiving the Bids of Others

When many value distributions and bidding functions are conceivable then it can be hard to understand how different restrictions influence bidding behavior. Hence, it can be simpler and more insightful to work directly with beliefs over bid distributions, instead of deriving these from value distributions and bidding functions.

In this section we consider a bidder who narrows down the possible environment she is facing by putting restrictions on the bids of others. First she imposes a lower bound on the possible maximal bid, then she allows for some mass below this threshold, and finally, we consider a model with independent bidding and bid dispersion.

4.1 Imposing a Lower Bound on the Maximal Bid of Others

We consider the bidding behavior of risk-neutral bidder 1 in a first-price auction for a single good.[2] In this section, 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 environments $\mathcal{E}$ as defined

[2] In Appendix A we consider risk-averse bidders.

in Section 2. This set is the set of all possible value distributions and all possible bidding functions.

Suppose bidder 1 believes that she needs to bid at least $L \geq 0$ so that her bid becomes winning. The value of L can come from a known reserve price,[3] or from the belief that the maximal bid of other bidders is at least L.

Definition 2. *Given $L \geq 0$, let $\mathcal{E}_L \subseteq \mathcal{E}$ be such that $\mathcal{E}_L$ is the set of environments belonging to $\mathcal{E}$ in which the highest bid is almost surely at least L.*

In the following we consider the case where $v_1 > L$ as bidding under $v1 \leq L$ is simple; all bids less than or equal to v_1 are optimal.

We briefly repeat the considered problem and solution concept. The bidder does not know the environment that consists of the joint value distribution and bidding behavior of the other bidders. All she knows is that the maximal bid of the others will almost surely be above L. If she knew the joint value distribution and their bidding functions, she could maximize expected utility. The objective is to get a payoff that is close to the one she would achieve if she knew the environment. The closer her payoff is to the best possible payoff, the smaller the loss. Hence, the objective is to find a bidding function that minimizes the maximal loss.

In order to maximize loss we do not need to consider all conceivable environments, since it suffices to restrict attention to a particularly simple environments. Loss associated with bidding function b_1 is the difference between the "oracle payoff", the maximized expected utility if the environment is known, and the expected utility generated by bidding function b_1. First, observe that loss is zero if one of the other bidders bids higher than v_1. Therefore, in order to maximize loss we only have to consider environments in which the maximal bid of other bidders is in the interval $[L, v_1]$. Second, notice that for given environment E bidder 1 wins if her bid is higher than the maximal bid among the other bidders, denoted by M. Loss is increased if bidder 1 learns this bid M in the oracle. As a result, loss is maximized by simple environments in which all other bidders bid M

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

with certainty.[4]

First we consider the performance of deterministic bids. Observe that bidder 1 never bids above v_1 because this results in a negative surplus and never below L, because these bids are losing for sure. Bidder 1 wins the auction if her bid is above M and loses it otherwise. Winning the auction with bid b_1 yields utility $v_1 - b_1$ and losing gives zero utility.[5] If bidder 1 bids b, then loss equals

$$l(v, M) = \sup_{\tilde{b}>M} \left\{ v - \tilde{b} \right\} - \mathbb{1}_{b>M}(v - b) = v - M - \mathbb{1}_{b>M}(v - b),$$

where $\mathbb{1}_{b>M} = 1$ if $b > M$ and 0 otherwise. The oracle payoff is $\sup_{\tilde{b}>M}\{v - \tilde{b}\}$, since bidder 1 knows the bid M she has to match.[6] Loss can occur from bidding too low and from bidding too high. Bidder 1 bids too low if the bid does not become winning. In this case, loss is highest if b is slightly outbid and not higher than $v - b$. In the other case, the bid b is too high. If $b > M \geq L$, then bidder 1 could raise her expected surplus by decreasing her bid. The loss of bidding too high is not more than $v - L - v + b = b - L$.

The maximal loss is $\max\{v - b, b - L\}$. Maximal loss is minimized by the bid that equates the two expressions. This bid is equal to $b = \frac{v+L}{2}$ and guarantees that absolute loss is below $\frac{v-L}{2}$. Loss relative to the distance $v - L$ is $\frac{1}{2}$.

Next we show how appropriate randomized bidding can reduce loss further. Assume that bidder 1 uses 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 surplus. 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 M and loses it otherwise. Loss is equal to the following difference when M

[4] Environments of this form can be rationalized by the following behavior. Suppose there are at least two other bidders and the value distribution puts all the mass on M. Let the $n - 1$ bidders know this. Then bidding value is an equilibrium for them.

[5] We will often drop the index if we think this causes no confusion.

[6] Formally, 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.

is known and when it is not known, i.e.

$$
\begin{aligned}
l(v, M) &= \max\left\{\sup_{\tilde{b}>M}\left\{v-\tilde{b}, 0\right\}\right\} - \int_M^v v - b\, dG(b|v) \\
&= \max\{v-M, 0\} - \int_M^v v - b\, dG(b|v) \qquad (1)
\end{aligned}
$$

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 expected utility of using the randomized bidding function G.

Proposition 1. *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*

$$
g(b|v) = \frac{1}{v-b} \text{ on } \left[L, v - \frac{v-L}{e}\right]. \qquad (2)
$$

The deterministic minimax loss is equal to $\frac{v-L}{2}$ and attained by bidding $\frac{v+L}{2}$.

In Appendix B, the proof specifies the details how the bidding function is derived. The mean bid of bidding function (2) 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. In the following sections we show that loss can be made smaller if very low bids are not conceived to be likely.

So we find that a bidder who is not willing to narrow down the environments further than $\mathcal{E}_L$ can only guarantee loss to be below $(v-L)/e$. This guarantee is achieved by appropriate randomized bidding and is 74% or less of the guarantee that can be achieved by a deterministic bidding strategy. The mixed strategy reduces loss because it allows some kind of hedging against unknown bids.

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, where regret is the difference in profit when the value is known and when it is not known. It turns out that the pdf of the monopolist's optimal pricing strategy resembles Equation (2). 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 interim loss. The difference of those two concepts is decision maker's knowledge used for the computation of the oracle payoff. On the one hand, in Bergemann and Schlag (2008) the monopolist knows the strategy of the potential buyer and uses the buyer's value in the benchmark. On the other hand, in this paper the bidder uses only the distribution and bidding function of the other bidders 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 minimax loss and we consider interim and not ex post loss.

4.2 Allowing some Mass Below the Threshold

In the analysis above bidder 1 restricted her bids to be above L, since she deemed that her bids below L are never winning. In this subsection we first show that it is best to ignore possible bids below L and to bid as in Proposition 1, as long as the likelihood of bids below L is sufficiently small. Then we discuss a related setting in which one knows that the own value is relatively small and the implications on loss.

Definition 3. *Let $L \geq 0$ and $\bar{p} \in [0, 1]$. Define $\mathcal{E}_{L\bar{p}} \subseteq \mathcal{E}$ 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 definition describes the following. Consider environments in $\mathcal{E}_{L\bar{p}}$. The maximal probability that the highest bid among the other bidders is below L is $\bar{p}$. 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, bidder 1 thought that the minimal maximal bid is L, thus $\bar{p} = 0$. 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 randomized bidding strategy of Proposition 1 with support $[L, v - \frac{v-L}{e}]$. Above we saw that if the highest bid among the other bidders is always above L, loss is at most $\frac{v-L}{e}$. 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 $\bar{p}$. If bidder 1 learns that all other bidders bid 0, then she knows that she bids too high. 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 $v > L \geq 0$ and $\bar{p} \leq \frac{v-L}{v-L+eL}$. For the set of conceivable environments $\mathcal{E}_{L\bar{p}}$ minimax loss is equal to $\frac{v-L}{e}$ and attained by the randomized bidding strategy stated in Proposition 1.*

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. One can, of course, also fix the bound on loss for type v and ask which L and $\bar{p}$ give rise to this loss (Example 1), or fix L and $\bar{p}$ and ask for which v the inequality is satisfied (Example 2).

Example 1. *Let* $v = 1$ *and* $L = 1 - \frac{e}{10}\,(\approx 0.73)$. *We need that* $\bar{p} = p^{n-1} \leq \frac{v-L}{v-L+eL} = \frac{1}{11-e}$. *Loss is less than one tenth if* $n = 2$ *and* $p \leq \frac{1}{11-e}\,(\approx 0.12)$ *or if* $n = 4$ *and* $p \leq \left(\frac{1}{11-e}\right)^{\frac{1}{3}}$ (≈ 0.49).

Example 2. *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 Bayes-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 (2) on $[0, v - \frac{v}{e}]$. This bidding functions ensures that all bids are below L, since $v < \frac{e}{e-1}L$.

Proposition 3. *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* $\bar{p}\frac{v}{e}$ *and attained by the randomized bidding strategy stated in Proposition 1 evaluated as if* $L = 0$.

The condition $\bar{p} \geq e\frac{v-L}{v}$ implies $\bar{p} \geq \frac{v-L}{v-L+eL}$. Hence, for v such that $L < v \leq \frac{e}{e-1}L$ Propositions 2 and 3 cannot hold at the same time.

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 our previous analysis (Subsection 4.1) the bidder thought it is conceivable that all other bidders bid the same value. 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 by assuming that the bidder believes that any given other 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 L for sure. Formally, bidder 1 conceives that the bidding distribution of a given bidder can be written as ε times the uniform distribution on $[L, v]$ plus $1 - \varepsilon$ times some arbitrary distribution.[7]

Definition 4. *Let $L \geq 0$ and $\varepsilon \in (0, 1)$. Define $\mathcal{E}_{L,\varepsilon,n} \subseteq \mathcal{E}_L$ to be the set of all environments such that (i) there are n bidders, (ii) for any bidder $i > 1$ it is as if the bid is drawn uniformly from the interval $[L, v]$ with probability at least ε.*

Within $\mathcal{E}_{L,\varepsilon,n} \subseteq \mathcal{E}_L$ it is again a simple form of environments that potentially maximize loss. These environments generate bid distributions such that for every bidder $i > 1$ the bid is drawn uniformly from $[L, v_1]$ with probability ε and equal to $M \in [L, v_1]$ with probability $1 - \varepsilon$. It is enough to restrict bids to the interval $[L, v_1]$, since we know from above that loss is made smaller if bids are above value with positive probability. Moreover, the more mass is distributed uniformly, the lower loss, hence we consider bid distributions in which the minimum of ε is distributed uniformly independently for every bidder $i > 1$.

In these simple environments it is as if bidder 1 learns the highest bid among the other bidders M whose bid is not drawn uniformly in the benchmark. Conditional upon bidding above M, the expected utility in the benchmark is equal to

$$EU(b|b > M) = \sum_{k=0}^{n-1} \binom{n-1}{k} \varepsilon^k (1-\varepsilon)^{n-1-k} (v - b) \left(\frac{b - L}{v - L}\right)^k. \tag{3}$$

If $b > M$, then bidder 1 wins against the $n - 1 - k$ bidders who bid M. The probability that b is the highest among k uniformly distributed bids is $\left(\frac{b-L}{v-L}\right)^k$. We sometime abbreviate $EU(b|b > M)$ simply with $EU(b)$. When we use $EU(M)$ then we mean that bidder 1 bids slightly above M. If bidder 1 bids below M, then the expected utility is

$$EU(b|b < M) = \varepsilon^{n-1}(v - b) \left(\frac{b - L}{v - L}\right)^{n-1},$$

[7] We 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.

since bidder 1 only wins if all other bidders' bids are drawn uniformly.

The following lemma says that the surplus in the benchmark is maximized by bidding (slightly above) the highest bid among the other bidders, or by bidding b^*, which is independent of M. Whenever $\varepsilon \leq 1/n$, then b^* is smaller than L and therefore smaller than M. In this case the bid b^* cannot become winning and bidder 1 needs to bid M to maximize utility. In the case of a relatively large $\varepsilon > 1/n$, the bid b^* is larger than L. Hence, there can be bids of the other bidders $M \in [L, b^*)$ such that bidder 1 ignores this information and bids b^*, which is independent of M.

Lemma 1. *Whenever $M < b^* = \frac{v(n\varepsilon-1)+L}{n\varepsilon}$, expected utility $EU(b|b > M)$ is maximized by bidding b^*. Otherwise $M \in \arg\sup_b EU(b|b > M)$.*

This lemma has direct consequences on the maximization of loss. Suppose $\varepsilon > 1/n$ so that $b^* > L$ and that bidder 1 uses a randomized bidding strategy with support $[\underline{b}, \bar{b}]$, where $\underline{b} < b^*$. If it turns out that the highest bid among the other bidders is $M \in [L, b^*)$, then bidder 1's optimal bid in the benchmark is b^*. Loss is non-decreasing for other bidders' highest bids $M' \in (M, b^*)$. This can be seen by observing that for such an M' the oracle payoff is unchanged since bidder 1's optimal bid is still b^*. The higher the other bidders bid, the less likely it is that a bid drawn from $[\underline{b}, \bar{b}]$ is winning. Therefore, loss cannot be maximized if the highest bid among the other bidders is below $\max\{L, b^*\}$.

The following proposition shows that minimax loss is attained by a randomized bidding function that depends on ε, the number of bidders n, and the lowest possible winning bid L. A closed form solution for the upper bound of the bidding function is not available–it needs to be computed numerically in applications. 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/e bound.

Proposition 4. *Let $v > L \geq 0$, $\varepsilon \in (0, 1)$, and n an integer and $n \geq 2$. For the set of conceivable environments $\mathcal{E}_{L,\varepsilon,n}$ minimax loss is attained by the randomized bidding strategy*

$$g(b|v) = \frac{\alpha(b)^{n-1}\beta(b)(v(1-\varepsilon n) + b\varepsilon n - L)}{(v-b)\left((\varepsilon-1)v\beta(b)^n + b\varepsilon\left(\alpha(b)^n - \beta(b)^n\right) + L\left(\beta(b)^n - \varepsilon\alpha(b)^n\right)\right)}, \tag{4}$$

with $\alpha(b) = v(1-\varepsilon) + b\varepsilon - L$ *and* $\beta(b) = \varepsilon(b-L)$*, on support* $\left[\underline{b}, \bar{b}\right]$ *, where* $\underline{b} = \max\{L, b^*\}$ *and* $\bar{b}$ *solves*

$$\int_{\underline{b}}^{\bar{b}} g(b|v)\, db = 1. \tag{5}$$

Minimax loss equals

$$EU\left(\bar{b}\right) - \varepsilon^{n-1} \int_{\underline{b}}^{\bar{b}} g(b|v) \left(\frac{b-L}{v-L}\right)^{n-1} (v-b)\, db. \tag{6}$$

Deterministic minimax loss is attained by $\hat{b}$ *such that*

$$EU(\max\{M, b^*\}) - EU(\hat{b}) = \sum_{k=1}^{n-1} \binom{n-1}{k} \varepsilon^k (1-\varepsilon)^{n-1-k} (v-\hat{b}) p(\hat{b})^k$$

and equal to the value on either side of the equation.

Example 3. *Table 1 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 roughly 74% of the maximal loss under deterministic bidding.*

Figure 1 shows the density of the bidding function of Equation (4) 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* $n = 10$*,* $b^* > 0$ *and therefore* b^* *is the lowest possible bid. As the number of other bidders increases, more mass is put on higher bids.*

5 Bidding with Beliefs about Behavior

In Section 4 we described uncertainty in terms of the bid distributions the robust bidder thinks she could be facing. In this section we focus on where these bids come from and formulate bidding in terms of uncertainty about value distributions and bidding behavior of others.

Randomized Bidding						
	$n = 2$			$n = 5$		
ε	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			

Deterministic Bidding						
	$n = 2$		$n = 5$		$n = 10$	
ε	Bid	Loss	Bid	Loss	Bid	Loss
0.10	0.51	0.44	0.62	0.33	0.76	0.19
0.15	0.52	0.41	0.68	0.26	0.84	0.13
0.20	0.53	0.38	0.75	0.20	0.88	0.10
0.25	0.54	0.34	0.80	0.16	0.90	0.08
0.40	0.58	0.25	0.87	0.10	0.94	0.05
0.50	0.62	0.19	0.90	0.08	0.95	0.04

Table 1: Loss for different values of ε and n, where $L = 0$ and $v = 1$

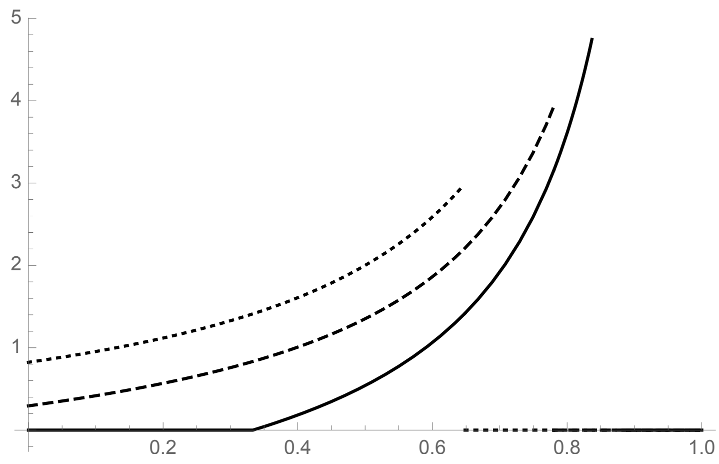

Figure 1: Probability density function of Equation (4) for $n = 2$ (dotted), $n = 5$ (dashed), and $n = 10$ (solid), where $\varepsilon = 0.15$, $L = 0$, and $v = 1$

5.1 Complete Uncertainty about Values

In this section we present two simple scenarios regarding beliefs about value distributions and behavior that imply that the maximal bid of the others will certainly be above L. This then means that we can apply our results of Section 4.1 where we showed that the maximal loss can be bounded above by $(v - L)/e$.

One might think that 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. Similarly, 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 $(v - L)/e$ 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 assumption, but does not know the value distribution. We show that 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 the limit of $\epsilon[v_1] + (1-\epsilon)[M]$, $M \geq L$ where ϵ converges to zero from above.[8] Bayesian bidders basically know that all bidders have the same value M. Consequently, bidding value is a best response to each other. This behavior generates the same conditions as in Subsection 4.1 and therefore Proposition 1 applies. Loss cannot be decreased by simply restricting uncertainty to uncertain over values. However, as we show in the next sections, loss can be decreased through bounds on the value distribution.

5.2 Behavioral Beliefs

The fact that the highest bid can be L with probability 1 contributes substantially to the bound on maximal loss provided in Section 4.1. One can imagine that the bidder's perception of the environment leads to constraints that rule out this extreme case. Such 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 this section we make plausible or intuitive restrictions on bidding behavior. In the next section we make assumptions on the other bidders' objective. In both sections we distinguish between robust and non-robust bidders with the understanding that all robust bidders are facing the same uncertainty. If there are at least two robust bidders then this means that we no longer have a decision problem and consequently have to introduce an equilibrium concept.

The robust bidder conceives the following environments possible. There are k other robust out of n total bidders, with $0 \leq k \leq n-1$. Values are distributed independently. All conceivable value distributions can be bounded from above. Bidding strategies of robust and non-robust bidders can be different. It is believed that the bidding strategies of non-robust bidders can be bounded from below. Robust bidders are believed to use a linear bidding function.[9] All bidders use deterministic (pure) strategies

[8] Alternatively, consider an asymmetric value distribution in which the Bayesian bidders have some prior over v_1, and know that all $i > 1$ have the same value.

[9] A similar model appears in Eliaz (2002). In his model, there are up to k "faulty"

and their bids are increasing in their value.

Definition 5. *Given* $0 < \alpha, \eta$ *and* $0 < \sigma, \tau < 1$, *let* $\mathcal{E}_{\alpha\eta}^{\sigma\tau} \subseteq \mathcal{E}$ *be such that* $\mathcal{E}_{\alpha\eta}^{\sigma\tau}$ *is the set of environments belonging to* $\mathcal{E}$ *in which for all conceivable value distributions* F, *(i) values are distributed independently, (ii)* $F(v) \leq \min\{\eta v^{\alpha}, 1\}$, *and for all conceivable bidding functions* b_{-1} *there are* k *bidders such that for every bidder* $b(v) = \tau v$, *and there are* $n-k$ *bidders with* $b(v) \geq \sigma v$.

In this section it is believed that all robust bidders use a linear bidding function. This belief is confirmed–we solve for the robust bidding strategy and find that a linear strategy generates suitable loss. This gives rise to a certain equilibrium notion. Let $O \subseteq \{1, \ldots, n\}$ be the set of robust bidders, where $|O| \geq 1$. In this section there are $k+1$ robust bidders, so $|O| = k+1$. Let $\bar{b}$ be the profile of bidding strategies of the robust bidders, $\bar{b} : O \to \mathbb{R}_{+}^{\Delta\mathbb{R}_{+}}$. The profile $\bar{b}$ is indexed by $i \in O$, hence $\bar{b}(i)$ denotes the bidding strategy of bidder i. Loss depends on the set of conceivable environments and we allow bidders to have heterogeneous perceptions $\mathcal{E}_i \subseteq \mathcal{E}$ for $i \in O$. Hence, for every bidder i there exists an individual bound on loss $\epsilon(i)$ that might depend on v_i. This is captured by defining $\epsilon : O \to \mathbb{R}_{+}^{\mathbb{R}_{+}}$.

Definition 6. *The strategy profile* $\bar{b}$ *is an* ϵ-loss-equilibrium *if for all* $i \in O$, *(i)* $\bar{b}(i)$ *guarantees a loss below* $\epsilon(i)$ *for environments in* $\mathcal{E}_i$, *and (ii) if* $(F, b_{-i}) \in \mathcal{E}_i$ *and* $j \in O \backslash \{i\}$ *then* $b_{-i}(j) = \bar{b}(j)$.

Without any equilibrium reasoning, the following proposition shows that a linear bidding functions allows to bound loss substantially. The beliefs about the value distribution and the bidding functions allow to bound the probability of a winning bid from above. The higher α, the less likely low values for other bidders. The higher σ and τ, the higher other bidders are expected to bid, the lower the probability of winning with a certain bid. It turns out that loss can be bounded from above by environments, that is, bid distributions, that put as much mass as possible on a unique bid and the rest of the mass above value v_1. Thus it is as if

players who choose any action. We assume that there are $n-k$ non-robust bidders bid can be bounded from below. Also see Gradwohl and Reingold (2014) for fault tolerance in large games.

bidder 1 learns the maximal bid of the other bidders, given this bid is below value. Note however, that these distributions are not always conceived as possible. If $\sigma > \tau$ and $0 < k < n-2$, robust and non-robust bidders submit different bids even if they have the same value. Since values are iid, it cannot be the case that robust and non-robust bidders have different value distributions. As a result, we do not talk about minimizing maximal loss, but rather that we bound loss.

Proposition 5. *Let* $0 < \alpha,\ \eta,\ 0 < \sigma,\ \tau < 1$ *and* $0 \leq v \leq \frac{1}{\eta^{1/\alpha}} \left(\frac{\sigma}{\tau}\right)^{\frac{n-k}{n-1}}$*. For the set of conceivable environments* $\mathcal{E}_{\alpha\eta}^{\sigma\tau}$*, loss can be guaranteed to be below*

$$\gamma \frac{\beta^\beta \left(1+\beta\right)^{\beta(1+\beta)}}{\left(\left(1+\beta\right)^{\beta+1}+\beta^\beta\right)^{1+\beta}} v^{1+\beta}, \tag{7}$$

by using the deterministic bidding function

$$b_1^*(v_1) = \frac{\left(1+\beta\right)^{1+\beta}}{\left(1+\beta\right)^{1+\beta}+\beta^\beta} v, \tag{8}$$

where

$$\gamma = \eta^{n-1} \left(\left(\frac{\sigma}{\tau}\right)^k \frac{\tau}{\sigma^n}\right)^\alpha \text{ and } \beta = \alpha\left(n-1\right).$$

Bidding function b^* is remarkable in its independence of the beliefs about the other bidders bidding behavior (σ and τ) and about the number of robust bidder k. The function only depends on the number of other bidders and the bound on the value distribution. 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 $0.8\,v$. On the other hand, the value of loss is dependent on all parameters of the model.

The bound on loss increases in the number of robust bidders k if and only if $\sigma \geq \tau$, that is, if robust bidders bid less aggressive than the other bidders. The intuition can be seen from the bound on the winning probability. Loss is increasing the more likely low bids are, that is, the lower the probability that a low bid is winning. The winning probability of a bid can be bounded from above by bids of non-robust bidders equal to σv.

Whenever $\sigma > \tau$, then it is the robust bidders that make low bids possible and therefore loss increases in their number.

The robust bidder chooses a linear bidding strategy if she anticipates that the other robust bidders choose a linear strategy. This gives rise to an ϵ-loss equilibrium if the other robust bidders have the same perception about the environments, i.e. use the same α and n. The notion of ϵ-loss equilibrium is comparable to the notion of ϵ Nash equilibrium in the sense that a lower loss might be possible, like a higher expected surplus might be possible in ϵ Nash, but players are content with giving up ϵ. Note that the bound on loss is a non-linear function in v.

Remark 1. *For $k+1$ robust bidders with conceivable environments $\mathcal{E}_{\alpha\eta}^{\sigma\tau}$ with*

$$\tau = \frac{(1+\beta)^{1+\beta}}{(1+\beta)^{1+\beta} + \beta^{\beta}},$$

bidding function (8) forms an ϵ-loss equilibrium with $\epsilon(v)$ equal to loss given in Equation (7).

Example 4. *Let there be five bidders ($n = 5$), $\alpha = \frac{1}{2}$, $\eta = 1$, and $v = 1$. Let $b_1^*(v) = \tau v = 27/31/, \approx 0.87$, and $\sigma = 0.9$. For $k = 1$ loss is guaranteed to be below* 0.12. *In the case of $k = 5$, loss is bounded by* 0.13.

Alternatively, for $\sigma = 0.8$ we need that $\eta \leq \frac{2}{45}\sqrt{465} \approx 0.96$ so that the bound on v is satisfied, i.e.

$$v \leq \frac{1}{\eta^{1/\alpha}} \left(\frac{\sigma}{\tau}\right)^{\frac{n-k}{n-1}}.$$

As $\sigma < \tau$, the robust bidder is more aggressive and loss decreases in the number of robust bidders. If $\eta = \frac{2}{45}\sqrt{465}$, then loss is smaller than 0.13 *for $k = 1$ and less than* 0.11 *if $k = 5$, that is, all bidders bid robustly.*

5.3 Explicit Rational Bidding

6 Conclusion

One of the major obstacles and challenges to bidding in first-price auctions is limited information. It is difficult, in many instances, 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 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 benchmark). Our method identifies how important detailed information is. We show that there are instances in which more detailed information is not needed to guarantee loss to be below 10% of the own value. In Subsection 4.1, the robust bidder knows very little about the other bidders. However, loss can be low if she knows that she has a relatively low value compared to other bidders. More formally, it is needed that the robust bidder knows that the highest bid of the other bidders is above the own value with probability at least 0.728. Another case is with the likely existence of at least one competitor who bids close to the own value. This happens if the maximal rivaling bid is below 72.8% of the own value with probability at most 0.12.

In Subsection 4.3, any open set of bids below v occurs with ex ante positive probability. This is modeled by drawing the bid for very other bidder uniformly with probability $\varepsilon \in (0, 1)$ and arbitrarily with probability $1 - \varepsilon$. Loss is decreasing in the number of bidders n and in the uncertainty ε. For example, if $\varepsilon = 0.15$ and $n = 10$, then loss is below 10% of the value.

In Section 5, the robust bidder bounds the other bidders' value distribution and bidding functions. Among other things, we solve a model in which there are only robust bidders. If all robust bidders use a linear bidding function, then a linear bidding function is robust. Loss is decreasing in the number of bidders. The less likely low types are, the lower loss.

A Risk-Aversion

A.1 Complete Uncertainty

Let u denotes bidder 1's von-Neumann-Morgenstern increasing utility function with $u(0) = 0$. Loss is

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

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

$$g(b|v) = -\frac{u'(v - b)}{u(0) - u(v - b)}$$

with support $[L, \bar{b}]$. The upper bound of the support is determined by $\bar{b} \leq v$ that solves $\int_L^{\bar{b}} g(b|v)db = 1$. One can get analytic solutions with two popular parameterized utility functions.

A.1.1 CRRA

Let for simplicity $L = 0$. Define $u(x) = x^{1-\rho}$, so that $u(0) = 0$. Risk neutrality corresponds to $\rho = 0$. The bidding function is

$$g(b|v) = \frac{1 - \rho}{v - b}$$

for a given v. The cumulative distribution function (cdf) is

$$G(b|v) = \int_0^b g(\tilde{b}|v)d\tilde{b} = (\rho - 1) \log\left(1 - \frac{b}{v}\right)$$

for $0 \leq b \leq v - e^{\frac{1}{\rho-1}}v$ and is equal to 1 for larger b that is, the support of g is $[0, v - e^{\frac{1}{\rho-1}}v]$. Loss is then bounded by

$$l^{CRRA}(v) = u(v - M) - \int_M^{v - e^{\frac{1}{\rho-1}}v} u'(v - b)db = \frac{v^{1-\rho}}{e}.$$

One can see that loss is increasing in ρ, that is, the more risk averse, the higher loss. Loss is increasing in risk aversion since randomization is used to trick nature. However, risk averse agents do not like to randomize, hence in their eyes randomization is less effective to reduce loss.

The mean bid

$$\int_0^{v-e^{\frac{1}{\rho-1}}} b\,dG(b|v) = \left(\rho - e^{\frac{1}{\rho-1}}(\rho-1)\right) v$$

is increasing in ρ. For $\rho \geq 0.7$ we have that the expected bid is roughly equal to ρv.

The variance of g is equal to

$$-\frac{1}{2}\left(e^{\frac{1}{\rho-1}}-1\right)(\rho-1)\left(-2\rho+e^{\frac{1}{\rho-1}}(2\rho-3)+1\right)v^2.$$

A.1.2 CARA

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

$$g(b|v) = \frac{\alpha e^{\alpha(-(v-b))}}{1-e^{\alpha(-(v-b))}}$$

and has support $\left[0, \frac{\log\left(-e^{\alpha v}+e^{\alpha v+1}+1\right)-1}{\alpha}\right]$. Loss is not higher than

$$l^{CARA}(v) = \frac{e^{\alpha+\alpha(-v)-1}\left(1-e^{\alpha v}\right)}{1-e^{\alpha}}.$$

Loss is again higher the more risk averse.

A.2 The ε-Uniform Model

A.2.1 CRRA

One can solve the model also for certain strictly concave utility functions such as CRRA $u(x) = x^{1-\rho}$. In this case the density is

$$g(b|v) = \frac{b\varepsilon(b\varepsilon-\varepsilon v+v)^{n-1}(b\varepsilon(\rho-n)+v(-\varepsilon\rho+\varepsilon n+\rho-1))}{(b-v)\left((\varepsilon-1)v(b\varepsilon)^n+b\varepsilon\left((b\varepsilon-\varepsilon v+v)^n-(b\varepsilon)^n\right)\right)}$$

and $b^* = \frac{-\varepsilon\rho+\varepsilon n+\rho-1}{\varepsilon(n-\rho)}v$. The comparative statics of loss with respect to ρ are ambiguous. The natural tendency is that loss is higher under more risk aversion because bidders do not like the randomization. But it can also be lower. A high ρ leads to a larger b^* and a large b^* reduces the exposure to randomization. Therefore, for example, if $b^* < 0$ under risk neutrality and $b^* > 0$ under risk aversion, then the loss under risk aversion can be lower.

B Proofs

Proposition 1. *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*

$$g(b|v) = \frac{1}{v-b} \text{ on } \left[L, v - \frac{v-L}{e}\right]. \tag{2}$$

The deterministic minimax loss is equal to $\frac{v-L}{2}$ and attained by bidding $\frac{v+L}{2}$.

Proof. We will first derive the optimal randomized bidding function for bidder 1. There are two ways to derive it. Then we consider environments, that is, bid distributions such that minimax loss is attained. Finally, we show which environments generate deterministic minimax loss.

We first 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 $l(v, M_1) = l(v, M_2)$ 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 (1) gives loss

$$l(v, M) = v - M - \int_M^{v-\frac{v-L}{e}} db = \frac{v-L}{e}.$$

An alternative derivation of the bidding function g is to take the first derivative of loss as specified in Equation (1) 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 (2). In most parts of the paper, we will 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 (1) equal to zero, as shown above.

Nature chooses the cumulative distribution function

$$H(M|v)=\frac{v}{e(v-M)} \tag{9}$$

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 $v_2=\cdots=v_n$, $v_2 \sim F$, $F(v_2|v_1)=\frac{v_1}{e(v_1-v_2)}$ 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 Equation (9).

Finally, deterministic loss can be equal to $\frac{v-L}{2}$ if nature selects $M^* = \frac{v+L}{2}$. To see this, suppose $M^* > b$ so that loss equals $v - M^* = \frac{v-L}{2}$. □

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

Proof. From Proposition 1 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. To do this, we consider simple environments for which the bidding function is the worst case. Then we show that the loss associated with these environments 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 $\bar{p}[0] + (1-\bar{p})[M_2]$, $M_2 > v - \frac{v-L}{e}$ equals

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

If $\bar{p}v \geq v - M_2$, then loss equals $\frac{\bar{p}(v+(e-1)L)}{e}$. This loss is less than $\frac{v-L}{e}$, since $\bar{p} \leq \frac{v-L}{v-L+eL}$. 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. □

Proposition 3. *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* $\bar{p}\frac{v}{e}$ *and attained by the randomized bidding strategy stated in Proposition 1 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 benchmark yields expected utility $\bar{p}(v - M_1)$ 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 benchmark 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 1 with support $\left[0, \frac{v}{e}\right]$ and $M_2 \geq L$ arbitrary. □

Proposition 4. *Let* $v > L \geq 0$, $\varepsilon \in (0, 1)$, *and* n *an integer and* $n \geq 2$. *For the set of conceivable environments* $\mathcal{E}_{L,\varepsilon,n}$ *minimax loss is attained by the randomized bidding strategy*

$$g(b|v) = \frac{\alpha(b)^{n-1}\beta(b)(v(1-\varepsilon n) + b\varepsilon n - L)}{(v-b)\left((\varepsilon-1)v\beta(b)^n + b\varepsilon\left(\alpha(b)^n - \beta(b)^n\right) + L\left(\beta(b)^n - \varepsilon\alpha(b)^n\right)\right)}, \tag{4}$$

with $\alpha(b) = v(1-\varepsilon) + b\varepsilon - L$ *and* $\beta(b) = \varepsilon(b - L)$, *on support* $\left[\underline{b}, \bar{b}\right]$, *where* $\underline{b} = \max\left\{L, b^*\right\}$ *and* $\bar{b}$ *solves*

$$\int_{\underline{b}}^{\bar{b}} g(b|v)\, db = 1. \tag{5}$$

Minimax loss equals

$$EU\left(\bar{b}\right) - \varepsilon^{n-1}\int_{\underline{b}}^{\bar{b}} g(b|v)\left(\frac{b-L}{v-L}\right)^{n-1}(v-b)\, db. \tag{6}$$

Deterministic minimax loss is attained by $\hat{b}$ such that

$$EU(\max\{M, b^*\}) - EU(\hat{b}) = \sum_{k=1}^{n-1} \binom{n-1}{k} \varepsilon^k (1-\varepsilon)^{n-1-k} (v - \hat{b}) p(\hat{b})^k$$

and equal to the value on either side of the equation.

Proof. We first prove the statement of the lemma and then derive the density g, loss, and the worst-case distribution. From now on we will also use the notation $p(b) = \frac{b-L}{v-L}$.

Lemma 1. *Whenever $M < b^* = \frac{v(n\varepsilon-1)+L}{n\varepsilon}$, expected utility $EU(b|b > M)$ is maximized by bidding b^*. Otherwise $M \in \arg\sup_b EU(b|b > M)$.*

Proof. In the proof we will show that b^* maximizes expected utility $EU(b|b > M)$ whenever $b^* > M \geq L$. If $b^* \leq M$, bidding arbitrarily small above M is optimal. To this end we first discuss the roots of the first order condition of Equation (3). We identify b^* as the utility maximizing root. Depending on the parameters, the bid b^* can be smaller than L and consequently smaller than M. If this is the case, then we show that M is the argument at which which the supremum of expected utility is attained.

The first order condition of Equation (3) with respect to b is equal to

$$\frac{\partial EU(b|b > M)}{\partial b} = \frac{(v-L)^{1-n}(b\varepsilon - L - \varepsilon v + v)^n(-b\varepsilon n + L + v(\varepsilon n - 1))}{(-b\varepsilon + L + (\varepsilon - 1)v)^2} = 0. \tag{10}$$

Its distinct roots are $b' = \frac{L+\varepsilon v - v}{\varepsilon}$ and b^*. The root b^* is the right root to consider, since $b' < L$.

The expected utility is maximized by bidding b^* if $b^* > M$, because $EU(b|b > M)$ is decreasing for $b \in (b^*, v)$. This can be seen from Equation (10). Therefore, whenever $b^* \geq L$ and $b^* > M$, then b^* is the unique maximizer of $EU(b|b > M)$. A necessary condition for $b^* > M$ is $b^* > L$ and this is true whenever $\varepsilon > \frac{1}{n}$.

However, if $b^* \leq M$, then $M \in \arg\sup_b EU(b|b > M)$ is optimal, since the expected utility is decreasing in b and has therefore its supremum in the smallest possible b.

□

Note that he expected utility $EU(b|b < M)$ is maximized by bidding $\tilde{b} = \frac{(n-1)v+L}{n}$. This lemma follows from solving the FOC of $EU(b|b < M)$ with respect to b.

We are now fit to discuss bidding behavior in the benchmark and its implications for loss. If the probability of a uniform bid is relatively small, i.e. if $\varepsilon < \frac{1}{n}$, then $b^* < L$. Consequently, the supremum of expected utility is attained by M or $\tilde{b}$. As long as M is not too high, the expected utility is higher from bidding slightly above M. Later it will be checked (numerically) that the largest M in the support is not too high. Loss is then equal to

$$EU(M) - \varepsilon^{n-1}\int_L^v (v-b)p(b)^{n-1}\, dG(b|v) - \sum_{k=0}^{n-2}\binom{n-1}{k}\varepsilon^k(1-\varepsilon)^{n-1-k}\int_M^v (v-b)p(b)^k\, dG(b|v). \tag{11}$$

The other case is when $\varepsilon \geq \frac{1}{n}$ and consequently $b^* \geq L$. If nature puts $M \in [L, b^*)$, then the optimal benchmark bid is b^* and loss equals

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

Loss is increased by $M' \in (M, b^*)$, because this does not change the benchmark payoff, but decreases the chance of winning under unknown M'. Hence, maximal loss is attained by $M \geq b^*$ and maximal loss has the form given in Equation (11).

Taking the first derivative with respect to M of Equation (11) and solving the first order condition for $g(M|v)$ leads to bidder 1 using the density

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

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 $\underline{b} = \max\{b^*, L\}$.

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

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

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 equilibrium of a zero-sum game, a player must be indifferent between two actions, therefore, bidder 1 must be indifferent between b and b'. For a given b, loss equals

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

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

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

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 $H(\underline{b}|v)$.

Now we derive optimal deterministic bids. Conditional upon bidding b, loss can come from two various sources. First, the highest bid among the other bidders can be slightly above b, so that b is not winning, or the bid b can be too large.

The highest loss of bidder 1 by being slightly outbid by one of the non-uniform bidders by bidding $M > b$ is equal to

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

The highest loss that results from bidding too high (i.e. $b > M$) is

$$\sup_{\tilde{b}>M}\left\{EU(\tilde{b}|\tilde{b}>M)\right\}-EU(b|b>M).$$

As in the case with random bidding, the supremum of expected utility is attained by $\max\{M, b^*\}$. The same case distinction as above applies. On the one hand, if $\varepsilon \leq \frac{1}{n}$, then loss can be maximized by $M = L$. On the other hand, if $\varepsilon > \frac{1}{n}$, then loss cannot be maximized by $M = L$, because bidder 1 will always choose to bid b^* in the benchmark if $b^* > M$. Therefore, the highest loss from $M < b$ is

$$\sup_{\tilde{b}>\max\{M,b^*\}}\left\{EU(\tilde{b}|\tilde{b}>M)\right\}-EU(b|b>M)=EU(\max\{M,b^*\})-EU(b).$$

Conditional on b, the highest loss is then

$$\max\left\{EU(\max\{M,b^*\})-EU(b),\sum_{k=1}^{n-1}\binom{n-1}{k}\varepsilon^k(1-\varepsilon)^{n-1-k}(v-b)p(b)^k\right\}$$

The optimal bid equates these two expressions.

□

Proposition 5. *Let* $0 < \alpha,\ \eta,\ 0 < \sigma,\ \tau < 1$ *and* $0 \leq v \leq \frac{1}{\eta^{1/\alpha}}\left(\frac{\sigma}{\tau}\right)^{\frac{n-k}{n-1}}$. *For the set of conceivable environments* $\mathcal{E}^{\sigma\tau}_{\alpha\eta}$, *loss can be guaranteed to be below*

$$\gamma\frac{\beta^\beta\,(1+\beta)^{\beta(1+\beta)}}{\left((1+\beta)^{\beta+1}+\beta^\beta\right)^{1+\beta}}v^{1+\beta}, \tag{7}$$

by using the deterministic bidding function

$$b_1^*(v_1)=\frac{(1+\beta)^{1+\beta}}{(1+\beta)^{1+\beta}+\beta^\beta}v, \tag{8}$$

where

$$\gamma = \eta^{n-1}\left(\left(\frac{\sigma}{\tau}\right)^{k}\frac{\tau}{\sigma^{n}}\right)^{\alpha} \text{ and } \beta = \alpha\left(n-1\right).$$

Proof. We will derive the linear bidding function and the bound on loss. Let $B(b)$ be the probability that all other bidders bid below b. The bound on the bidding strategies and the value distribution allow to bound the winning probability, so

$$B(b) \le \eta^{n-1}\left(\left(\frac{b}{\tau}\right)^{\alpha}\right)^{k-1}\left(\left(\frac{b}{\sigma}\right)^{\alpha}\right)^{n-k} = \gamma b^{\beta}.$$

Consider the loss of a robust bidder who faces some bid distribution B with $B(b) \le \gamma b^{\beta}$. Loss is maximized by bid distributions of the form Q_x that puts mass γx^{β} on x and rest of mass sufficiently high above v. In this case, loss is given by

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

This loss is bounded by $l(b, Q_y) \le \max_{x<b}\left\{\gamma b^{\beta}\left(v-b\right), \gamma x^{\beta}\left(b-x\right)\right\}$. Maximizing loss with respect to the value x is done by $x = \beta\frac{b}{1+\beta}$, since $\frac{d}{dx}\left(\gamma x^{\beta}\left(b-x\right)\right)|_{x=\beta\frac{b}{1+\beta}} = 0$. Thus, $\gamma x^{\beta}\left(b-x\right) = \gamma\beta^{\beta}\left(\frac{b}{1+\beta}\right)^{1+\beta}$.

We obtain a new expression for the upper bound on loss with

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

The upper bound on loss is minimized by choosing b such that $\gamma b^{\beta}\left(v-b\right) = \gamma\beta^{\beta}\left(\frac{b}{1+\beta}\right)^{1+\beta}$. The loss minimizing bid is given in Equation (8). Using this rule leads to the bound on loss given in Equation (7).

The above equations are only tight for v such that $\gamma\left(b^{*}\left(v\right)\right)^{\beta} \le 1$. This is true for $v \le \frac{1}{\eta^{1/\alpha}}\left(\frac{\sigma}{\tau}\right)^{\frac{n-k}{n-1}}$, since $\gamma\tau^{\beta}v^{\beta} = \eta^{n-1}\left(\left(\frac{\sigma}{\tau}\right)^{k}\frac{\tau}{\sigma^{n}}\right)^{\alpha}\tau^{\alpha(n-1)}v^{\beta} = \eta^{n-1}\left(\frac{\tau}{\sigma}\right)^{\alpha(n-k)}v^{\beta} \le 1$. □

References

Battigalli, Pierpaolo and Marciano Siniscalchi, "Rationalizable Bidding in First Price Auctions," *Games and Economic Behavior*, 2003, *45*, 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*, 2527–2543.

__ **and Stephen Morris**, "Robust Mechanism Design," *Econometrica*, November 2005, *73* (6), 1771–1813.

__ **, Benjamin Brooks, and Stephen Morris**, "First-Price Auctions with General Information Structures: Implications for Bidding and Revenue," *Econometrica*, forthcoming 2016.

Binmore, Ken and Paul Klemperer, "The Biggest Auction Ever: The Sale Of The British 3G Telecom Licences," *The Economic Journal*, 2002, *112.*

Carroll, Gabriel, "Robustness and Linear Contracts," *American Economic Review*, 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.

de Clippel, Geoffroy, "Behavioral Implementation," *American Economic Review*, 2014, *104*, 2975–3002.

Dekel, Eddie and Asher Wolinsky, "Rationalizable outcomes of large private-value first-price discrete auctions," *Games and Economic Behavior*, 2003, *43*, 175–188.

Eliaz, Kfir, "Fault tolerant implementation," *Review of Economic Studies*, 2002, *69* (3), 589–610.

Engelbrecht-Wiggans, Richard, "The Effect of Regret on Optimal Bidding in Auctions," *Management Science*, 1989, *35* (6).

— **and Elena Katok**, "Regret and Feedback Information in First-Price Sealed-Bid Auctions," *Management Science*, 2008, *54* (4).

Filiz-Özbay, Emel and Erkut Y. Özbay, "Auctions with Anticipated Regret: Theory and Experiment," *American Economic Review*, September 2007, *97* (4), 1407–1418.

Gradwohl, Ronen and Omer Reingold, "Fault Tolerance in Large Games," *Games and Economic Behavior*, 2014, *86*, 438–457.

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.

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.

Levin, Dan and Emre Ozdenoren, "Auctions with uncertain numbers of bidders," *Journal of Economic Theory*, 2004, *118*, 229–251.

Linhard, P. B. and Roy Radner, "Minimax-Regret Strategies for Bargaining over Several Variables," *Journal of Economic Theory*, 1989, *48*, 152–178.

Lo, Kin Chung, "Sealed bid auctions with uncertainty averse bidders," *Economic Theory*, 1998, *12*, 1–20.

Renou, Ludovic and Karl Schlag, "Implementation in Minimax-Regret Equilibrium," *Games and Economic Behavior*, 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).

Savage, Leonard J., "The theory of Statistical Decision," *Journal of the American Statistical Association*, 1951, *46*, 55–67.

Schlag, Karl and Andriy Zapechelnyuk, "Robust sequential search," *Working paper*, 2016.