

Inference on Auctions with Weak Assumptions on Information*

Vasilis Syrgkanis[†]

Elie Tamer[‡]

Juba Ziani[§]

October 12, 2017

Abstract

Given a sample of bids from independent auctions, this paper examines the question of inference on auction fundamentals (e.g. valuation distributions, welfare measures) under weak assumptions on information structure. The question is important as it allows us to learn about the valuation distribution in a robust way, i.e., without assuming that a particular information structure holds across observations. We leverage recent contributions in the robust mechanism design literature that exploit the link between Bayesian Correlated Equilibria and Bayesian Nash Equilibria in incomplete information games to construct an econometrics framework for learning about auction fundamentals using observed data on bids. We showcase our construction of identified sets in private value and common value auctions. Our approach for constructing these sets inherits the computational simplicity of solving for correlated equilibria: checking whether a particular valuation distribution belongs to the identified set is as simple as determining whether a *linear* program is feasible. A similar linear program can be used to construct the identified set on various welfare measures and counterfactual objects. For inference and to summarize statistical uncertainty, we propose novel finite sample methods using tail inequalities that are used to construct confidence regions on sets. We also highlight methods based on Bayesian bootstrap and subsampling. A set of Monte Carlo experiments show adequate finite sample properties of our inference procedures. We also illustrate our methods using data from OCS auctions.

*We thank participants at the Econometric Society China Meeting in Wuhan - 2017 and at the R. Porter Festschrift at Northwestern University in May 2017 for comments. Code for monte-carlo simulations and ocs data analysis using our methods can be found at https://github.com/vsyrgkanis/information_robust_econometrics_auctions

[†]Microsoft Research, New England; vasy@microsoft.com

[‡]Department of Economics, Harvard University; elietamer@fas.harvard.edu

[§]Department of Computing and Mathematical Sciences, California Institute of Technology; jziani@caltech.edu

1 Introduction

A recent literature in robust mechanism design studies the following question: given a game (e.g. an auction), what are the possible outcomes - such as welfare or revenue - that arise under different information structures. This literature is motivated by *robustness*, i.e., characterizing outcomes that can occur in a given game under weak assumptions on information (See Bergemann and Morris [2013]). For example, in auction models, in addition to specifying the details of the game in terms of bidder utility function, and bidding rules, one needs to specify the information structure (what players know about the state of the world and the information possessed by other players) to derive the Nash equilibrium.

For instance, in the Independent Private Values (IPV) setting, players know their value for the item, which is assumed independent from other player values, and receive no further information about their opponents' values. The latter typically yields a unique equilibrium outcome. However, different assumptions on what signals players have about opponent values prior to bidding in the auction, lead to different equilibrium outcomes. Auction data rarely contain information on what bidders knew and what their information sets include, and given that this information leads to different outcomes, it would be interesting to analyze what can happen when we relax the independence assumption in such auctions by allowing bidders to know some information about their opponents' valuations. Bergemann, Brooks, and Morris [2017] (BBM) examine exactly this question in an auction game and provide achievable bounds on various outcomes, such as the revenue of the auction as a function of auction fundamentals like the distribution of the common value.

In this paper, we address the following econometrics question, which is the reverse of the one posed by BBM above: given an i.i.d. sample of auction data (independent copies of bids from a set of auctions), what can we learn about auction fundamentals, such as the distribution of values, when we make weak assumptions on the information structure? We maintain throughout that players use Bayesian Nash equilibrium (BNE) but allow these bidders to have different information structures in different auctions, i.e. receive different types of signals prior to bidding. In particular, we use these observations (bids and other observables) from a set of independent auctions to construct *sets* of valuation distributions that are consistent with both the observed distribution of bids and the known auction rules maintaining that players are Bayesian. We exploit the robust predictions in a given auction à la BBM to conduct econometrically robust predictions of auction fundamentals

given a set of data; i.e., robust economic prediction leads to robust inference.

Key to our approach is the characterization of sharp sets of valuation distributions (and other functionals of interest) via computationally attractive procedures. This is a result of the equivalence between a particular class of *Bayes Correlated Equilibria* (or BCE) and Bayes Nash Equilibria (or BNE) for a similar game with an arbitrary information structure. It is well known that BCEs can be computed efficiently since they are solutions to linear programs (as opposed to BNE which are hard to compute). Moreover, exploiting a result of Bergemann and Morris [2016] (see also Aumann [1987]), we show that there is an equivalence between the set of fundamentals that obey the BCE restrictions and the fundamentals that obey the BNE constraints under *some* information structure. This equivalence is the key to our econometrics approach. The formal statistics program that ensues is one where the sharp set satisfies a set of linear equality and inequality constraints. If we knew the true distribution of bids, then this becomes a simple computational problem of solving a linear program. We do not observe the true bid distribution, but this distribution can be estimated consistently from the observed data. So, we use the estimated bid distribution to solve for an estimate of the identified set. We are also able to characterize sampling uncertainty to obtain various notions of confidence regions covering the identified with a prespecified probability.

In addition to learning about auction primitives, we show how our approach using data on auctions can be used to construct identified sets for auction welfare measures, seller surplus and other objects. Information on auction primitives along with these measures obtained using our procedures that combine data with the theory can be used to guide future market designers to better study particular auction setups (or use our results in other markets).

Importantly, we address the problem of counter-factual estimation: what would the revenue or surplus have been had we changed the auction rules. We formulate notions of informationally robust counter-factuals and we show that such counter-factual questions can also be phrased as solutions to a single linear program, simultaneously capturing equilibrium constraints in the current auction as well as the new target auction. We show that even without recovering the information structure from the data, an analyst can perform robust counterfactuals: under an arbitrary information structure in the current auction which produced the data at hand, what is the best and worst value of the given quality measure (e.g. welfare, revenue) in the new target auction under an arbitrary information structure also? Thus we can get estimates of the upper and lower bounds of a given quantity in the new auction design, in a way that is robust to information structure and without the

need of recovering it. So, this approach to inference requires minimal assumptions on information in the DGP that generated the data, but also minimal assumptions in the counterfactual auction that a market designer is contemplating to run.

The closest work to our paper is Bergemann, Brooks, and Morris 2017, where the authors provide worst-case bounds on the revenue of a common value first price auction as a function of the distribution of values. Their approach does not use the bid distribution as input, unlike our approach which obtains an estimate of the bid distribution from the data. The main approach in BBM is to claim that the revenue cannot be too small since at BCE, no player wants to deviate to any other action and so players do not want to deviate to a specific style of a deviation which is: conditional on your bid, deviate uniformly at random above your bid (upwards deviation). So, the bound on the mean they provide uses a *subset* of the set of best response deviations that are allowed so as to bound the bid of a player as a function of his value. Hence, in drawing a connection between the equilibrium bid and the value, this bound is by definition loose (and can be very loose - bound twice as large as identified set - as we show in an example in the appendix). Given data, we are able to learn the bid distribution and hence we are not constrained to look at only these bid-distribution-oblivious upwards deviations. We can instead compute an optimal best-response bid for this given bid distribution and use the constraint that the player does not want to deviate to this tailored best-response. This allows us to bound the unobserved value of the player as a function of the observed bid leading to a *sharp characterization* of auction fundamentals using data.

Another paper that uses a similar insight of studying the econometrics of games with weak information is the recent work of Magnolfi and Roncoroni [2016] on inference in entry game models. The approach used there, though similar in motivation, does not transport easily to studying general auction mechanisms.

The paper is organized as follows. Section 2 introduces the problem and provides formal definitions of the objects of interest. We then state our identification results given an iid set of data on bids. This identification is constructed via a linear program where we show how various constraints (such as symmetry, parametric restrictions, etc) can be used. We also show how sharp sets for the expected value of moments of the fundamentals amounts to solving two linear programs and how robust counter-factuals to changes in the auction, with respect to some metric function can also be easily handled. We then provide two example applications of the general setup: one for

common value auctions (Section 3) and another for private value auctions (Section 4). Section 5 provides our estimation approach for constructing confidence intervals on the estimated quantities from sampled datasets using sub-sampling methods and finite sample concentration inequality approaches. Section 6 examines the finite sample performance of the large linear program using a set of Monte Carlo simulations. These show adequate performance in IPV and CV setups. Section 7 illustrates our inference approach using auction data from OCS wildcat oil auctions and show how the statistical algorithm can be used to derive bounds on valuation distributions. Finally, the Appendix contains results on the sharpness of the BBM bounds, and bounds on the mean of the valuation in common value auctions with different smoothness assumptions.

2 Bayes-Correlated Equilibria and Information Structure Uncertainty

We consider a game of incomplete information among n players. There is an unknown payoff-relevant state of the world $\theta \in \Theta$. This state of the world enters directly in each player's utility. Each player i can pick from among a set of actions A_i and receives utility which is a function of the payoff-relevant state of the world θ and the action profile of the players $\mathbf{a} \in A \equiv A_1 \times \dots \times A_n$: $u_i(\mathbf{a}; \theta)$. This along with a prior on θ (defined below) will represent the game structure that we denote by G , as separate from the information structure which we will define next.

Conditional on the state of the world each player receives some minimal signal $t_i \in T_i$. The state of the world $\theta \in \Theta$ and the vector of signals $\mathbf{t} \in T \equiv T_1 \times \dots \times T_n$ are drawn from some joint measure¹ $\pi \in \Pi \subseteq \Delta(\Theta \times T)$. The signals t_i 's can be arbitrarily correlated with the state of the world. We denote such signal structure with S . This defines the game (G, S) .

We consider a setting where prior to picking an action each player receives some additional information in the form of an extra signal $t'_i \in T'_i$. The signal vector $\mathbf{t}' = (t'_1, \dots, t'_n)$ can be arbitrarily correlated with the true state of the world and with the original signal vector $\mathbf{t} = (t_1, \dots, t_n)$. We denote such augmenting signal structure with S' and the set of all possible such augmenting signal structures with \mathcal{S}' . This will define a game (G, S') . Subsequent to observing the signals t_i and t'_i , the player picks an action a_i . A Bayes-Nash equilibrium or BNE in this game (G, S') is a mapping $\sigma_i : T_i \times T'_i \rightarrow \Delta(A_i)$ from the pair of signals t_i, t'_i to a distribution over actions, for each player i , such that each player maximizes his expected utility conditional on the signals s/he received.

¹The setup is general in that the set T is unrestricted.

A fundamental result in the literature on robust predictions (see Bergemann and Morris [2013, 2016]), is that the set of joint distributions of outcomes $a \in A$, unknown states θ and signals t , that can arise as a BNE of incomplete information under an arbitrary additional information structure in (G, S') , is equivalent to the set of Bayes-Correlated Equilibria, or BCE in (G, S) . So, every BCE in (G, S) is a BNE in $(G, S;)$ for some augmenting information structure S' . We give the formal definition of BCE next.

Definition 1 (Bayes-Correlated Equilibrium). *A joint distribution $\psi \in \Delta(\Theta \times T \times A)$ is a Bayes-correlated equilibrium of (G, S) if for each player i , signal t_i , action a_i and deviating action a'_i :*

$$\mathbb{E}_{(\theta, t, a) \sim \psi} [u_i(a; \theta) \mid a_i, t_i] \geq \mathbb{E}_{(\theta, t, a) \sim \psi} [u_i(a'_i, a_{-i}; \theta) \mid a_i, t_i] \quad (1)$$

and such that the marginals with respect to signals and payoff states are preserved, i.e.:

$$\forall \theta \in \Theta, t \in T : \sum_{a \in A} \psi(\theta, t, a) = \pi(\theta, t) \quad (2)$$

An equivalent and simpler way of phrasing the Bayes-correlated equilibrium conditions is that:

$$\forall t_i, a_i, a'_i : \sum_{\theta, t_{-i}, a_{-i}} \psi(\theta, t, a) \cdot (u_i(a; \theta) - u_i(a'_i, a_{-i}; \theta)) \geq 0 \quad (3)$$

We state the main result in Bergemann and Morris [2016] next.

Theorem 1 (Bergemann and Morris [2016]). *A distribution $\psi \in \Delta(\Theta \times T \times A)$ can arise as the outcome of a Bayes-Nash equilibrium under some augmenting information structure $S' \in \mathcal{S}'$, if and only if it is a Bayes-correlated equilibrium in (G, S) .*

The robustness property of this result is as follows. The set of BNEs for (G, S') (think of an auction with unknown information) is the same as the set of BCEs for (G, S) where S' is an augmented information structure derived from S . So, we will not need to know what is in S' , but rather we could compute the set of BCEs for (G, S) and the Theorem shows that for each BCE, there exists a corresponding information structure S' in \mathcal{S}' and a BNE of the game (G, S') that implements the same outcome.

2.1 The Econometrics Inference Question

We consider the question of inference on auction fundamentals using data under weak assumptions on information. In particular, assume we are given sample of observations of action profiles

a^1, \dots, a^N from an incomplete information game G . Assume that we do not know the exact augmenting signal structure S^1, \dots, S^N that occurred in each of these samples where it is implicitly maintained that S^i can be different from S^j for $i \neq j$, i.e., the signal structure in the population is drawn from some unknown mixture. Also, maintaining that players play Nash, or that a^t was the outcome of some Bayes-Nash equilibrium or BNE under signal structure S^t , We study the question of inference on the distribution π of the fundamentals of the game. Under the maintained assumption, we characterize the sharp set of possible distributions of fundamentals π that could have generated the data. This allows for policy analysis within the model without making strong restrictions on information.

A similar question was recently analyzed in the context of entry games by Magnolfi and Roncoroni [2016], where the goal was the identification of the single parameter of interaction when both players choose to enter a market. In this work we ask this question in an auction setting and attempt to identify the distribution of the unknown valuations non-parametrically.

The key question for our approach is to allow for observations on different auctions to use different (and unobserved to the econometrician) information structures, and, given the information structures, that different markets or observations on auctions, to use a different BNE. Given this equivalence of the set of Bayes-Nash equilibria under some information structure and the set of Bayes-correlated equilibria, and given that the set of BCEs is convex, allowing for this kind of heterogeneity is possible. Heuristically, given a distribution over action profiles $\phi \in \Delta(A)$ (which is constructed using the data), there exists a mixture of information structures and equilibria under which ϕ was the outcome. The process by which we arrived at ϕ is by first picking an information structure from this mixture and then selecting one of the Bayes-Nash equilibria for this information structure. This is possible if and only if there exists a distribution $\psi(\theta, t, a)$, that is a Bayes-correlated equilibrium and such that $\sum_{\theta, t} \psi(\theta, t, a) = \phi(a)$ for all $a \in A$. Again, we start with the elementary information structure S and maintain that observations in the data are expansions of this information structure. So, for a given market i , any BNE using information structure S^i is a BCE under S . The data distribution of action is a mixture of such BNEs over various information structures and hence it would map into a mixture of BCEs under the same S . Since the set of BCEs under S is convex, any mixtures of elements in the set is also a BCE. So then intuitively, the set of primitives that are consistent with the model and the data is the set of BCEs, $\psi(\theta, t, a)$ such that $\sum_{\theta, t} \psi(\theta, t, a) = \phi(a)$ for all $a \in A$.

To conclude, the convexity of the set of BCEs allows us to relate a distribution of bids from an iid

sample to a mixture of BCEs. This is possible since the distribution of bids uses a mixture of signal structures, which essentially coincides with a mixture of BCEs, itself another BCE by convexity. We summarize this discussion with a formal result.

Lemma 2. *Consider a model which conditional on the unobservables, yields a convex set of possible predictions on the observables. Then the sharp identified set of the unobservables under the assumption that exactly one of these predictions is selected in our dataset, is identical to the sharp identified set under the assumption that our dataset is a mixture of selections from these predictions.*

Proof. Suppose that what we observe in the data is a convex combination of feasible predictions of our model. Then by convexity of the prediction set, this convex combination is yet another feasible prediction of our model. Hence, this is equivalent to the assumption that this single prediction is selected in our dataset. ■

Finally, we provide next the main engine that allows for construction of the observationally equivalent set of primitives that obey model assumptions and result in a distribution on the observables that match that with the data. We state this as a Result.

Result 3. *Let there be a distribution ϕ defined on the space of action profiles. Given the setup and results above, the set of feasible joint distributions of signals and types π that are consistent with ϕ are the set of distributions for which the following linear program is feasible:*

$$LP(\phi, \pi) \quad \forall t_i, a_i, a'_i : \quad \sum_{\theta, t_{-i}, a_{-i}} \phi(a) \cdot x(\theta, t|a) \cdot (u_i(a; \theta) - u_i(a'_i, a_{-i}; \theta)) \geq 0 \quad (4)$$

$$\forall(\theta, t) : \quad \pi(\theta, t) = \sum_{a \in A} \phi(a) \cdot x(\theta, t|a) \quad (5)$$

$$x(\cdot, \cdot | \alpha) \in \Delta(\Theta \times T) \quad (6)$$

where $\sum_{\theta, t} \psi(\theta, t, \alpha) = \phi(\alpha)$ for all $\alpha \in A$, and $\psi(\theta, t, \alpha) = \phi(\alpha)x(\theta, t|\alpha)$.

Equivalently, the sharp set for the distribution $\pi \in \Delta(\Theta \times T)$:

$$\Pi_I(\phi) = \{ \pi \in \Pi : LP(\phi, \pi) \text{ is feasible} \}$$

The above result is generic, in that it handles general games with generic states of the world θ . In particular, it nests both standard private and common value auction models and provides a mapping between the distribution of bids ϕ and the set of feasible distributions over signals and θ . An iid assumption on bids along with a large sample assumption allow us to learn the function $\phi(\cdot)$ (asymptotically). So, given the data, we can consistently estimate ϕ . This is a maintained assumption that we require throughout. Given ϕ , the above result tells us how to map the estimate of ϕ to the set of BNEs that are consistent with the data and are robust to any information structure that is an expansion of a minimal information structure S . Suppose we assume that both t and θ take finitely many values (an assumption we maintain throughout), then a joint distribution on (t, θ) , $\pi(t, \theta)$ is consistent with the model and the data if and only if it solves the above *linear program*. The LP formulation is general, but in particular applications, it is possible to use parametric distributions for the π . In addition, it is possible for the above LP to allow for observed heterogeneity by using covariate information whereby this LP can be solved accordingly (see an example such adaptation in the common value Section 3). Though the above LP holds in general (and covers both common and private values for example), we specialize in the next Sections the above LP to standard cases studied in the auction literature, mainly common values and private values models.

2.2 Identified Sets of Moments of Fundamentals

In the case of non-parametric inference where we put no constraint on the distribution of fundamentals, i.e. $\Pi = \Delta(\Theta \times T)$, observe that the sharp set is linear in the density function of the fundamentals $\pi(\cdot, \cdot)$. Therefore, maximizing or minimizing any linear function of this density can be performed via solving a single linear program. This implies that we can evaluate the upper and lower bounds of the expected value of any function $f(\theta, t)$ of these fundamentals, in expectation over the true underlying distribution. The latter holds, since the expectation of any function with respect to the underlying distribution is a linear function of the density. We state this as a corollary next.

Corollary 4. *Let Result (3) hold. Also, let $f(\cdot, \cdot)$ be any function of the state of the world and the profile of minimal signals (eg, $f(\theta, t) = \theta$ or $f(\theta, t) = \theta^2$, ...). The sharp identified set for the expected value of $f(\cdot, \cdot)$ w.r.t. the true distribution of fundamentals, i.e. $E_{(\theta, t) \sim \pi}[f(\theta, t)]$ is an*

interval $[L, U]$ such that:

$$L = \min_{\pi \in \Pi_I(\phi)} \sum_{\theta \in \Theta, t \in T} f(\theta, t) \cdot \pi(\theta, t) \quad (7)$$

$$U = \max_{\pi \in \Pi_I(\phi)} \sum_{\theta \in \Theta, t \in T} f(\theta, t) \cdot \pi(\theta, t) \quad (8)$$

Without further assumptions on Π , i.e., when $\Pi = \Delta(\Theta, T)$, these are two linear programming problems.

2.3 Robust Counter-factual Analysis

Suppose that we wanted to understand the performance of some other auction when deployed in the same market, with respect to some objective or metric: $F : \Theta \times T \times A \rightarrow \mathbb{R}$ that is a function of the unknown fundamentals and the action vector. Examples of such metrics in single-item auctions could be social welfare: $F(\theta, t, a) = \sum_{i=1}^n \theta_i x_i(a)$ or revenue, i.e. $F(\theta, t, a) = \sum_{i=1}^n p_i(a)$, where θ_i is the value of player i for the item at sale, $x_i(a)$ is the probability of allocating to player i under action profile a and $p(a)$ is the expected payment of player i under action profile a .

We are interested in computing an upper and lower bound on this metric under this new auction which has different utilities $\tilde{u}_i(a; \theta)$ and under any Bayes-correlated equilibrium which would map into a BNE with an augmenting information structures. This is important since it allows us to obtain bounds on welfare or other metrics in a new environment . The welfare bounds computed in this manner will inherit the robustness property in that they will be valid under all information structures.

This is straightforward in our setup since computing a sharp identified set for any such counter-factual can be done in a computationally efficient manner, in both the common value setting and in the correlated private value setting. The upper bound of the counter-factual can be obtained using the following linear program, that takes as input the observed distribution of bids in our current auction, the metric F , and the primitive utility form $\tilde{u} = (\tilde{u}_1, \dots, \tilde{u}_n)$ under the alternative auction. We state this result in the next Theorem.

Theorem 5. *Given a metric function $F : \Theta \times B \rightarrow \mathbb{R}$, a distribution over action profiles $\phi()$ and vector of alternative auction utilities \tilde{u} , the sharp upper and lower bounds on the expected metric*

under the new auction can be computed using the following LP:

$$\begin{aligned}
LP(\phi, F, \tilde{u}) \quad & \min_{\tilde{\psi}} / \max_{\tilde{\psi}} \sum_{\theta, t, a} F(\theta, t, a) \cdot \tilde{\psi}(\theta, t, a) \\
& \forall t_i, a_i, a'_i : \sum_{\theta, t_{-i}, a_{-i}} \tilde{\psi}(\theta, t, a) \cdot (\tilde{u}_i(a; \theta) - \tilde{u}_i(a'_i, a_{-i}; \theta)) \geq 0 \\
& \forall (\theta, t) \in \Theta \times T : \pi(\theta, t) = \sum_{a \in A} \tilde{\psi}(\theta, t, a) \\
& \tilde{\psi} \in \Delta(\Theta \times T \times A) \text{ and } \pi \in \Pi_I(\phi)
\end{aligned}$$

In the non-parametric case, where $\Pi = \Delta(\Theta \times T)$, the latter is a linear program.

Note that getting sharp bounds on welfare measures for example using the above Theorem does not require one to infer in a prior step the distribution over the primitives. Rather, the above procedure provides sharp bounds on this welfare measure using a linear program.

3 Common Value Auctions

We specialize the above results to important classes of auctions. We begin with the common value model in which the game of incomplete information G is a single item common value auction. In this case the unknown state of the world is the unknown common value of the object v , which we assume to take values in some finite set V . Moreover, we initially assume that the minimal information structure is degenerate, i.e., players receive no minimal signal about this unknown common value². The signal set T becomes a singleton and is irrelevant. Thus we will denote with $\pi \in \Delta(V)$ the distribution of the unknown common value, which is the parameter that we wish to identify. This is a particularly simple model to illustrate the structure of the LP approach and showcase the flexibility of our methods. So, in this particular model, we want to learn the distribution of the state of the world which is the common valuation distribution.

Prior to bidding in the auction, the players receive some signal which is drawn from some distribution; this signal can be correlated with the unknown common value and with the signals of his opponents. We wish to be ignorant about which information structure realized in each auction sample and want to identify the sharp identified set for π . Moreover, we will assume that the players' bids take values in some discrete set B and players play a BNE. The characterization of the identified set for π in this model is stated in the Theorem below.

²Other constraints on the initial signals are allowed and here we take the degenerate signal for simplicity.

Theorem 6. *Let the common value model above hold with n bidders. Given a distribution of bid profiles ϕ supported on a set $S \subseteq B^n$, the set of distributions of bids $\pi \in \Delta(V)$ that are consistent with ϕ are ones where the following program is feasible:*

$$\begin{aligned}
LP(\phi, \pi) \quad & \forall b_i^*, b'_i \in B : & \sum_{v \in V, \mathbf{b} \in S: b_i = b_i^*} \phi(\mathbf{b}) \cdot x(v|\mathbf{b}) \cdot (u_i(\mathbf{b}; v) - u_i(b'_i, \mathbf{b}_{-i}; v)) & \geq 0 \\
& \forall v \in V : & \pi(v) = \sum_{\mathbf{b} \in S} \phi(\mathbf{b}) \cdot x(v|\mathbf{b}) \\
& \forall \mathbf{b} \in S : & x(\cdot|\mathbf{b}) \in \Delta(V)
\end{aligned}$$

Observe, that in this setting, the latter linear program is also linear in π . Thus we get that the sharp identified set is a convex set and is defined as the set of solutions to the above linear program, where π is also a variable.

This observation also allows us to easily infer upper and lower bounds on any linear function of the unknown distribution π . This is stated next as a Corollary to the above Theorem.

Corollary 7. *Let Result (6) hold. Also, let $f(\cdot)$ be any function of valuations (such as $f(v) = v$ or $f(v) = v^2, \dots$). Then, using the LP above, we can get upper and lower bounds on such moments as follows:*

$$\max_{\pi \in \Pi_I(\phi)} \sum_{v \in V} f(v) \cdot \pi(v) \tag{9}$$

$$\min_{\pi \in \Pi_I(\phi)} \sum_{v \in V} f(v) \cdot \pi(v) \tag{10}$$

Observe that the latter linear expressions are simply: $\mathbb{E}_{v \sim \pi}[f(v)]$ and so the above shows that we can compute in polynomial time upper and lower bounds of any moment of the unknown distribution of the common value.

Also, note that the Corollary above shows that to do set inference with respect to any moment of the unknown distribution, we do not need to discretize the space of probability distributions and enumerate over all probability vectors, checking whether they are inside the sharp identified set. Rather we can just solve the above LP.

Essentially this observation says that we can easily compute the support function $h(z; \Pi_I(\phi))$ of the identified set $\Pi_I(\phi)$ at any direction z , by simply solving a linear program.

In Section (A.2), we provide an upper bound on the mean valuation distribution when the latter is continuous. This upper bound is derived in terms of the observed bids distribution.

Remark 1 (Winning bid). *The above procedure can be easily modified if indeed as it may be the case, only winning bids are observed. In particular, given that I observe the CDF $F(\cdot)$ of the winning bid, I know that an equilibrium is consistent with said CDF if and only if for every possible bid in B ,*

$$F(x) = \sum_{v \in V} \sum_{\mathbf{b} / b_i \leq x \forall i} \psi(\mathbf{v}, \mathbf{b})$$

This has up to B linear constraints of B^{n+1} variables, but if we assume the bid vector distribution has small support of size K , we only need $K \cdot B$ variables.

Remark 2 (Covariate Heterogeneity). *Suppose we have covariates and want to allow for observed heterogeneity where we maintain the assumption that the vector of covariates x takes finitely many values. Then, one nonparametric approach is to repeat and solve the above LP for $\pi(v|x = x_0)$ for every value x_0 that x takes. In addition, in cases where we have $E[v|x] = x'\beta$, then we can solve directly for the identified set for β by solving the following LP:*

$$\begin{aligned} LP(\phi, \beta) \quad & \forall b_i^*, b'_i \in B, x_0 \in X : \quad \sum_{v \in V, \mathbf{b} \in S: b_i = b_i^*} \phi(\mathbf{b}|x_0) \cdot x(v|\mathbf{b}, x_0) \cdot (u_i(\mathbf{b}; v) - u_i(b'_i, \mathbf{b}_{-i}; v)) \geq 0 \\ & \forall x_0 \in X : \quad x'_0 \beta = \sum_{\mathbf{b} \in S} \sum_v v \phi(\mathbf{b}|x_0) \cdot x(v|\mathbf{b}, x_0) \\ & \forall \mathbf{b} \in S, x_0 \in X : \quad x(\cdot|\mathbf{b}, x_0) \in \Delta(V) \end{aligned}$$

Here, the $n \times k$ vector $x'_0 \beta$ allows for different players to have different β 's (and the x 's in this case would be auction specific heterogeneity where the different β 's would allow the mean valuation of different players to depend differently on auction characteristics).

The latter is attractive as it allows us to couple the identified set of the mean of the unknown common value across multiple co-variate realizations, in a computationally tractable manner. Otherwise we would have to generate the identified sets of the mean for each co-variate realization and then solve a second stage problem which would try to find the set of joint solutions in each of these identified sets (via some form of joint grid search) that are consistent with a model of how the conditional mean $E[v|x]$ varies as a function of x . However, joint grid search would grow exponentially with the number of realizations of the co-variates. The latter remark, shows that when the model of $E[v|x]$ is linear, then we can save this exponential blow-up in the computation whilst leveraging the statistical power of coupling data from separate co-variate realizations.

4 Private Value Auctions

We now consider the case of a private value single item auction. In this case the (unknown) state of the world is the a vector of private values $\mathbf{v} = (v_1, \dots, v_n) \in V^n$. We assume that these private values come from some unknown joint distribution $\pi \in \Delta(V^n)$. Moreover, we initially assume that players know at least their own private value. Thus the (minimal) signal set T_i is equal to V and moreover, we have that conditional on a value vector v , $t_i = v_i$, deterministically. Since the signal is a deterministic function of the unknown state of the world, we will again denote with $\pi \in \Delta(V^n)$ the distribution of the unknown valuation vector, which is the parameter that we wish to identify. Here, each player first draws a valuation (as an element of the state of the world), and then each player's own valuation is revealed to the player through a signal. After that, a signal is further revealed before players play a BNE given this signal.

In this setting the sharp identified set is again slightly simplified. The result is stated in the next Theorem.

Theorem 8. *Let the above private values auction model hold. Given a distribution of bid profiles ϕ supported on a set $S \subseteq B^n$, the set of distributions of bids $\pi \in \Delta(V^n)$ that are consistent with ϕ are the ones where the following linear program is feasible:*

$$\begin{aligned}
 LP(\phi, \pi) \quad & \forall v_i^* \in V, b_i^*, b'_i \in B : & \sum_{\mathbf{v}: v_i=v_i^*, \mathbf{b} \in S: b_i=b_i^*} \phi(\mathbf{b}) \cdot x(\mathbf{v}|\mathbf{b}) \cdot (u_i(\mathbf{b}; v_i) - u_i(b'_i, \mathbf{b}_{-i}; v_i)) &\geq 0 \\
 & \forall \mathbf{v} \in V^n : & \pi(\mathbf{v}) = \sum_{\mathbf{b} \in S} \phi(\mathbf{b}) \cdot x(\mathbf{v}|\mathbf{b}) \\
 & \forall \mathbf{b} \in S : & x(\cdot|\mathbf{b}) \in \Delta(V^n)
 \end{aligned}$$

Observe, that in this setting, the latter linear program is also linear in π . Thus we get that the sharp identified set is a convex set and is defined as the set of solutions to the above linear program, where π is also a variable. Note also here that no assumption is made on the correlation between player valuation. This result allows for the recovery of valuation distribution with arbitrary correlation (and general signaling structures). As a special case of the above, we study next the IPV model of auctions.

Independent Private Values. The situation becomes more complex if we also want to impose an extra assumption that the distribution of private values is independent. In that case, we have

the extra condition that π must be a product distribution which is a non-convex constraint. For instance, if we want to assume that the value of each player is independently drawn from the same distribution ρ , and hence ρ is what we wish to identify, then we also have the extra constraint that:

$$\forall \mathbf{v} \in V^n : \pi(\mathbf{v}) = \rho(v_1) \cdot \dots \cdot \rho(v_n) \quad (11)$$

Adding this constraint into the above LP, makes the LP non-convex with respect to the variables $\rho(v)$ (even though checking whether a given ρ is in the identified set, is still an LP). Thus in this case we cannot compute in polynomial time upper and lower bounds on the moments of the distribution ρ using the above LP.

However, we make the following observation which simplifies the constraints of the LP: we note that conditional on a player's valuation and on a bid profile b , the effect of a deviation $u_i(\mathbf{b}; v_i) - u_i(b'_i, \mathbf{b}_{-i}; v_i)$ is independent of the values of opponents, in a private value setting. Thus we can re-write the best response constraint as:

$$\sum_{\mathbf{b} \in S: b_i = b_i^*} \phi(\mathbf{b}) \cdot x_i(v_i^* | \mathbf{b}) \cdot (u_i(\mathbf{b}; v_i^*) - u_i(b'_i, \mathbf{b}_{-i}; v_i^*)) \geq 0 \quad (12)$$

where $x_i(v_i | b) = \Pr[V_i = v_i | B = b]$, where V_i is the random variable representing player i 's value and B is the random variable representing the bid profile at a BCE.

Then we can formulate the consistency constraints, by simply imposing a constraint per player, i.e. if ρ_i is the distribution of player i 's value, then it must be that:

$$\rho_i(v_i) = \sum_{\mathbf{b} \in S} \phi(\mathbf{b}) \cdot x_i(v_i | \mathbf{b}) \quad (13)$$

These are constraints that are still linear in $\rho_i(v_i)$. We state the LP as a corollary next.

Corollary 9. *Assume that the above IPV model hold. The following LP characterizes the sharp identification under the independent private values model:*

$$\begin{aligned} LP(\phi, \rho) \quad & \forall i \in [n], v_i \in V, b_i^*, b'_i \in B : & \sum_{\mathbf{b} \in S: b_i = b_i^*} \phi(\mathbf{b}) \cdot x_i(v_i | \mathbf{b}) \cdot (u_i(\mathbf{b}; v_i) - u_i(b'_i, \mathbf{b}_{-i}; v_i)) & \geq 0 \\ & \forall i \in [n], v_i \in V : & \rho_i(v_i) = \sum_{\mathbf{b} \in S} \phi(\mathbf{b}) \cdot x_i(v_i | \mathbf{b}) \\ & \forall \mathbf{b} \in S : & x_i(\cdot | \mathbf{b}) \in \Delta(V) \end{aligned}$$

In particular if we assume that player's are symmetric, i.e. $\rho_i(v) = \rho(v)$, then we can compute

upper and lower bounds on any moment $E[f(v)]$ of the common value distribution:

$$\begin{aligned}
& \max_{x_i(\cdot|\mathbf{b})} \sum_{v \in V} f(v) \cdot \rho(v) \\
& \forall i \in [n], v_i \in V, b_i^*, b'_i \in B : \quad \sum_{\mathbf{b} \in S: b_i = b_i^*} \phi(\mathbf{b}) \cdot x_i(v_i|\mathbf{b}) \cdot (u_i(\mathbf{b}; v_i) - u_i(b'_i, \mathbf{b}_{-i}; v_i)) \geq 0 \\
& \forall i \in [n], v \in V : \quad \rho(v) = \sum_{\mathbf{b} \in S} \phi(\mathbf{b}) \cdot x_i(v|\mathbf{b}) \\
& \forall i \in [n], \mathbf{b} \in S : \quad x_i(\cdot|\mathbf{b}) \in \Delta(V)
\end{aligned}$$

A by-product of this analysis is that the linear program allows us to test for symmetry in the independent private values model. In particular it is not clear that when we assume that all marginals are the same, then the LP is feasible. Thus by checking feasibility of the LP we can refute the assumption of symmetric independent private values.

Note here that given that we allow the augmenting signals to be arbitrary correlated, we are not able to infer any information about the joint distribution of valuation, such as correlation, given a bid profile.

4.1 Point Identification in First Price Auctions with Continuous Bids

Generally in the independent private values model, if we allow the bids and valuations to be continuous, we show here that the bidder specific valuation distributions are all identified even allowing for general information structures. This is important since the information here is allowed to be correlated but yet the valuation distributions ρ_i of each player i (under independence) are point identified.

Denote with $G_{-i}(\cdot|b_i)$, the CDF of the maximum other bid at a BCE, conditional on a bid b_i of player i , and let $g_{-i}(\cdot|b_i)$ be the density. These are observable in the data. Now the utility of a player from submitting a bid b'_i conditional on being recommended by BCE to player b_i and observing a value of v_i is:

$$U_i(b'_i; v_i, b_i) = (v_i - b'_i) \cdot G_{-i}(b'_i|b_i) \quad (14)$$

Since, b_i is maximizing the above quantity, by the best-response constraints, then we get that the derivative of the utility with respect to b'_i has to be equal to 0 at $b'_i = b_i$. The latter implies:

$$(v_i - b_i) \cdot g_{-i}(b_i|b_i) - G_{-i}(b_i|b_i) = 0 \quad (15)$$

We summarize our result in the following Theorem.

Theorem 10. *Let the IPV model above hold and further assume the bid distribution is continuous and admits a density. Thus we can write the value of a player as a function of his bid and of the conditional maximum other bid distribution:*

$$v_i = b_i + \frac{G_{-i}(b_i|b_i)}{g_{-i}(b_i|b_i)} \quad (16)$$

Hence, the valuation distribution for each player is point identified.

Hence, if we know the population bid distribution and if we assume that it is continuous and admits a density, then given the bid of a player we can invert and uniquely identify his value. Thus we can write the CDF of the value distribution of each player as a function of the observables.

A similar result was shown for the first price IPV model (without signals) in Guerre, Perrigne, and Vuong [2000] and later also generalized to the affiliated private values setting by Li, Perrigne, and Vuong [2002]. Interestingly, the inversion formula that we arrived to in the independent private values setting, but under robustness to the information structure, is the same as the inversion formula of Li, Perrigne, and Vuong [2002] under the correlated private value setting. It is interesting to note that under the assumption that the bid of a player is strictly increasing in his value at any Bayes-Correlated equilibrium, then we can re-do the analysis in this section in the more general correlated private values setting and show that we can invert the value of a player from the observed correlated bid distribution. The inversion formula is identical to Equation (16) and identical to the inversion derived in Li, Perrigne, and Vuong [2002]. This essentially shows that in terms of inference, robustness to information structure is in some sense equivalent to robustness to correlation in values.

Finally, as in the common values case, it is possible to adjust the LP to account for observing only the winning bids in the Private Values case.

5 Statistical Inference on Sharp Sets in CV Auctions

In general, we do not have access to the distribution $\phi(\mathbf{b})$. Instead, we have access to N i.i.d. observations $\omega_1, \dots, \omega_N$ from said distribution ; let $\omega_t(\mathbf{b}) = 1\{\omega_t = \mathbf{b}\}$. The sampled distribution is then given by

$$\phi_N(\mathbf{b}) = \frac{1}{N} \sum_{t=1}^N \omega_t(\mathbf{b}) \quad (17)$$

In this section, we develop techniques for estimating the sharp identified set using samples from $\phi(\mathbf{b})$. We showcase our techniques for the CV setup for simplicity. Also, we focus on inference

on the identified set for π which is the object of inference here. Often times, we parametrize the distribution of the common values, and so inference will be on the identified set for the vector of parameters³. We start with finite sample approaches to constructing confidence regions for sets using concentration inequalities. These seem to be the first application of such results on using such inequalities to settings with partial identification⁴. We also show how existing set inference methods can also be used to construct confidence regions.

5.1 Finite Sample Inference Using Concentration Inequalities

We explore first the question of set inference using finite sample concentration inequalities. This allows us to obtain a confidence set for the identified set where the coverage property holds for every sample size. In addition, we highlight tools from the concentration of measure literature applied to inference on sets in partially identified models.

As a reminder, given a probability distribution over bids $\phi(\mathbf{b})$, the sharp set of compatible equilibria $\phi(v, \mathbf{b}) = \phi(\mathbf{b})x(v|\mathbf{b})$ is the set Θ_I of joint probability distributions satisfying:

$$\forall b_i^*, b'_i \in B : \sum_{v \in V, \mathbf{b} \in S: b_i = b_i^*} \phi(\mathbf{b}) \cdot x(v|\mathbf{b}) \cdot (u_i(\mathbf{b}; v) - u_i(b'_i, \mathbf{b}_{-i}; v)) \geq 0 \quad (18)$$

for the proper $x(\cdot|\cdot)$. See the statement of Theorem (6) above. Let $\pi(\cdot)$ be the vector characterizing the distribution of the common value. Let $\{F_j(x; \pi, \phi) : j \in M\}$, denote the negative of the best-response and density constraints, associated with the BCE LP in Theorem (6). Then in the population π is feasible iff:

$$\min_x \max_{j \in M} F_j(x; \pi, \phi) \leq 0 \quad (19)$$

Then we have that the identified set is defined as:

$$\Pi_I = \{\pi : \min_x \max_{j \in M} F_j(x; \pi, \phi) \leq 0\} \quad (20)$$

Now we consider a finite sample analogue. Observe that $F_j(x; \pi, \phi) = \mathbb{E}[f_j(x; \pi, \omega)]$, where expectation is over the random vector ω and the function $f_j(\cdot; \pi, \omega)$ takes the form:

$$f_j(x; \pi, \omega) = \sum_{v \in V, \mathbf{b} \in S: b_i = b_i^*} \omega(\mathbf{b}) \cdot x(v|\mathbf{b}) \cdot (u_i(b'_i, \mathbf{b}_{-i}; v) - u_i(\mathbf{b}; v)) \quad (21)$$

³It is possible to construct the CI for the unknown parameters -rather than the identified set by inverting test statistics.

⁴Finite sample inference results are particularly attractive in models with partial identification since standard asymptotic approximations are not typically uniformly valid especially in models that are close to/or are point identified.

for some triplet (i, b_i^*, b'_i) for the case of best-response constraints and similarly for density consistency constraints. Then we consider the sample analogue:

$$F_j^N(x; \pi) = \frac{1}{N} \sum_{t=1}^N f_j(x; \pi, \omega_t) \quad (22)$$

Observe that due to the linearity of F_j with respect to ϕ , we can re-write: $F_j^N(x; \pi) = F_j(x; \pi, \phi_N)$. One can then define the estimated identified set by analogy, i.e., replacing $\phi(b)$ with $\phi_N(b)$:

$$\Pi_I^N = \{\pi : \min_x \max_{j \in M} F_j(x; \pi, \phi_N) \leq \sigma^N\} \quad (23)$$

for some decaying tolerance constant σ^N (which can be set to zero).

Because the pdf of a non-parametric distribution is a very high-dimensional object, inference on it will require many samples. Hence, we will instead focus on two more structured inference problems. In the first one we are interested in inferring the identified set for *a moment* of the distribution and in the second we make assume that the said pdf is known up to a finite dimensional parameter and infer the identified set of these lower dimensional parameters. One can in principle recover non-parametric inference by simply making the parameters be the values of the pdf at the discrete support points, albeit at a cost in the sample complexity.

5.1.1 Inference on Identified Set for Moments.

We begin with the non-parametric setting and show how to construct inference on the identified set of any moment function $m : V \rightarrow [-H, H]$ of the common value distribution that is valid in finite samples. Observe that the identified set for any moment $m : V \rightarrow [-H, H]$ is an interval $[L, U]$ defined by:

$$\begin{aligned} L &= \min_{x: \max_{j \in M} F_j(x; \phi) \leq 0} \sum_{v \in V} m(v) \cdot \sum_{b \in S} \phi(\mathbf{b}) \cdot x(v|\mathbf{b}) \\ U &= \max_{x: \max_{j \in M} F_j(x; \phi) \leq 0} \sum_{v \in V} m(v) \cdot \sum_{b \in S} \phi(\mathbf{b}) \cdot x(v|\mathbf{b}) \end{aligned} \quad (24)$$

where x in both optimizations is ranging over the convex set of conditional distributions, i.e. $x(\cdot|\mathbf{b}) \in \Delta(V)$, which we omit for simplicity of notation. We can then define their finite sample analogues as:

$$\begin{aligned} L^N(\sigma^N) &= \min_{x: \max_{j \in M} F_j^N(x) \leq \sigma^N} \sum_{v \in V} m(v) \cdot \sum_{b \in S} \phi_N(\mathbf{b}) \cdot x(v|\mathbf{b}) \\ U^N(\sigma^N) &= \max_{x: \max_{j \in M} F_j^N(x) \leq \sigma^N} \sum_{v \in V} m(v) \cdot \sum_{b \in S} \phi_N(\mathbf{b}) \cdot x(v|\mathbf{b}) \end{aligned} \quad (25)$$

Where $F_j^N(x)$ and ϕ_N are given in Equations (22) and (17) respectively.

The following result gives finite sample high probability bounds on the coverage of the interval $[L^N(\sigma^N) - \epsilon^N, U^N(\sigma^N) + \epsilon^N]$. As a reminder, we use n to designate the number of bidders, N is sample size, and $|B|$ is the number of support points for bids.

Theorem 11. *Suppose that $u_i(\mathbf{b}; v) \in [-H, H]$ for all $i \in [n]$, $\mathbf{b} \in B^n$ and $v \in V$ and let $\sigma^N = 2H\sqrt{\frac{\log(4n|B|^2/\delta)}{N}}$ and $\epsilon^N = 2H\sqrt{\frac{\log(4/\delta)}{N}}$. Then:*

$$\Pr [[L, U] \subseteq [L^N(\sigma^N) - \epsilon^N, U^N(\sigma^N) + \epsilon^N]] \geq 1 - \delta \quad (26)$$

where L, U are defined by Equation (24) and $L^N(\sigma^N), U^N(\sigma^N)$ are defined by Equation (25).

Proof. We will show that with probability at least $1 - \delta/2$: $L^N(\sigma^N) - \epsilon^N \leq L$. The theorem follows by also showing that $U^N(\sigma^N) + \epsilon^N \geq U$ with probability at least $1 - \delta/2$ and then using a union bound of the bad events. The second inequality, follows along identical lines by simply replacing $m(v)$ with $-m(v)$. So it suffices to show the first inequality.

We remind the reader the definitions of the two quantities:

$$L = \min_{x: \max_{j \in M} F_j(x; \phi) \leq 0} \sum_{v \in V} m(v) \cdot \sum_{b \in S} \phi(\mathbf{b}) \cdot x(v|\mathbf{b}) \quad (27)$$

$$L^N(\sigma^N) = \min_{x: \max_{j \in M} F_j^N(x) \leq \sigma^N} \sum_{v \in V} m(v) \cdot \sum_{b \in S} \phi_N(\mathbf{b}) \cdot x(v|\mathbf{b}) \quad (28)$$

Let x^* be the optimal solution to the population LP of Equation (27). We will argue that w.h.p. for the choice of σ^N it remains a solution of the finite sample LP of Equation (28) and the value of the finite sample LP under x^* is at least the value of the population LP plus ϵ^N .

Observe that $F_j^N(x^*)$ is the sum of N i.i.d. random variables bounded in $[-H, H]$ with mean $F_j(x^*; \phi)$. By Hoeffding's inequality with probability $1 - \kappa$:

$$F_j^N(x^*) \leq F_j(x^*; \phi) + 2H\sqrt{\frac{\log(1/\kappa)}{N}} \leq 2H\sqrt{\frac{\log(1/\kappa)}{N}} \quad (29)$$

where the second inequality follows by feasibility of x^* for the LP of Equation (27). By a union bound over all $j \in M$ and since $|M| = n|B|^2$ we get that with probability at least $1 - \kappa n|B|^2$:

$$\max_{j \in M} F_j(x^*; \phi) \leq 2H\sqrt{\frac{\log(1/\kappa)}{N}} \quad (30)$$

For $\kappa = \frac{\delta}{4n|B|^2}$ and $\sigma^N = 2H\sqrt{\frac{\log(4n|B|^2/\delta)}{N}}$ we get that with probability $1 - \frac{\delta}{4}$: $\max_{j \in M} F_j^N(x^*) \leq \sigma^N$ and thereby x^* is feasible for the finite sample LP.

Finally, the value of the finite sample solution can be written as $Q^N(x) = \frac{1}{N} \sum_{t=1}^N q(x; \omega_t)$ with:

$$q(x; \omega_t) = \sum_{v \in V} m(v) \sum_{\mathbf{b} \in S} \omega_t(\mathbf{b}) \cdot x(v|\mathbf{b}) \quad (31)$$

Thus it is also the sum of N i.i.d. random variables bounded in $[-H, H]$,⁵ and with mean $Q(x; \phi) = \sum_{v \in V} m(v) \cdot \sum_{\mathbf{b} \in S} \phi(\mathbf{b}) \cdot x(v|\mathbf{b})$. Hence, by Hoeffding's inequality, with probability at least $1 - \delta/4$:

$$Q_j^N(x^*) \leq Q_j(x^*; \phi) + 2H \sqrt{\frac{\log(4/\delta)}{N}} = L + \epsilon^N \quad (32)$$

By a union bound we get that with probability at least $1 - \delta/2$, x^* is feasible for the finite sample LP and achieves value $Q_j^N(x^*) \leq L + \epsilon^N$. In that case, by the definition of $L^N(\sigma^N)$, we get: $L^N(\sigma^N) \leq Q_j^N(x^*) \leq L + \epsilon^N$ and the theorem follows.⁶ ■

5.1.2 Inference on Identified Set for Parametric Distributions.

For the parametric case, we assume that the distribution $\pi(v)$ is parametric of the form $\pi(v, \theta)$ for some finite parameter set Θ . In the parametric setting we need to augment the constraint set M , apart from containing the best-response constraints, to contain the parametric form consistency constraints:

$$\forall v \in V : \pi(v, \theta) = \sum_{\mathbf{b} \in S} \phi(\mathbf{b}) \cdot x(v|\mathbf{b}) \quad (33)$$

These can also be written of the form $F_j(x; \phi) \leq 0$ for some function $F_j(x; \phi) = \mathbb{E}[f_j(x; \omega)]$.⁷ We overload notation and let M be this augmented set of constraints. Then the parameter of interest is θ and the identified set for θ takes the form:

$$\Theta_I = \{\theta \in \Theta : \min_x \max_{j \in M} F_j(x; \theta, \phi) \leq 0\} \quad (34)$$

and its sample equivalent

$$\Theta_I^N(\sigma^N) = \{\theta \in \Theta : \min_x \max_{j \in M} F_j(x; \theta, \phi_N) \leq \sigma^N\} \quad (35)$$

for some decaying tolerance constant σ^N (which can be set to zero).

The next Theorem provides finite sample high probability coverage bounds for Θ_I .

⁵Since $m(v) \in [-H, H]$ and $x(\cdot|\mathbf{b}) \in \Delta(V)$.

⁶We could have used an empirical Bernstein inequality instead of Hoeffding's inequality and replace the H in the quantities σ^N by $\sqrt{\max_j \text{Var}_N(F_j^N(x^*))}$, where Var_N is the empirical variance, at the expense of adding a lower order term of $O(\frac{H}{N})$ (see e.g. Maurer and Pontil [2009], Peel et al. [2010]). Similarly, for ϵ^N . However, the latter requires knowledge of x^* and taking a supremum over x^* in the latter seems to be as conservative as a Hoeffding bound.

⁷Simply add one constraint of the form $\pi(v, \theta) - \sum_{\mathbf{b} \in S} \phi(\mathbf{b}) \cdot x(v|\mathbf{b}) \leq 0$ and one of the form $\sum_{v \in V} \phi(\mathbf{b}) \cdot x(v|\mathbf{b}) - \pi(v, \theta) \leq 0$.

Theorem 12. Consider Θ_I and $\Theta_I^N(\sigma^N)$ as defined in Equations (34) and (48). Suppose that $u_i(\mathbf{b}; v) \in [-H, H]$ for all $i \in [n]$, $\mathbf{b} \in B^n$ and $v \in V$. If $\sigma^N = 2H\sqrt{\frac{\log(|\Theta| \cdot (n|B|^2 + |V|)/\delta)}{N}}$, then:

$$\Pr [\Theta_I \subseteq \Theta_I^N(\sigma^N)] \geq 1 - \delta \quad (36)$$

Proof. To show the statement we need to show that with probability $1 - \delta$, for all $\theta \in \Theta_I$ it must be that $\theta \in \Theta_I^N(\sigma^N)$. Equivalently, if $\min_x \max_{j \in M} F_j(x; \theta, \phi) \leq 0$ then $\min_x \max_{j \in M} F_j^N(x; \theta) \leq \sigma^N$. The latter follows along similar lines as in the proof of Theorem 11. Let $\theta \in \Theta_I$ and let x^* be the solution to the population problem. Then by Hoeffding's inequality and a union bound over M ,⁸ for the chosen σ^N we have that with probability at least $1 - \frac{\delta}{|\Theta|}$: $\max_{j \in M} F_j^N(x^*; \theta) \leq \sigma^N$. Hence, the latter inequality also holds for the minimum over x and thereby $\theta \in \Theta_I^N(\sigma^N)$. By a union bound the latter holds uniformly over θ with probability at least $1 - \delta$. ■

Sample Variance Based Bounds. We now show how to improve upon the prior analysis and getting tolerance variables σ^N whose leading $1/\sqrt{N}$ term does not depend on H , but rather on the sample variance of the constraints. We will modify the finite sample identified set $\Theta_I^N(\sigma^N)$ defined in Equation (48), as follows: instead of using a uniform tolerance σ^N for all the constraints, we will define tolerance for each constraint differently based on its sample variance:

$$\hat{\Theta}_I^N(\lambda, \sigma^N) = \left\{ \theta : \min_x \max_{j \in M} \left(F_j^N(x; \theta) - \lambda \sqrt{\frac{\text{Var}_N(f_j(x; \theta, \omega))}{N}} \right) \leq \sigma^N \right\} \quad (37)$$

Where $\text{Var}_N(X)$ for a random variable X , denotes the sample variance: $\frac{1}{N(N-1)} \sum_{1 \leq t < t' \leq n} (X_t - X_{t'})^2$. The extra variance modification, which is reminiscent of *sample variance penalization* in empirical risk minimization [Maurer and Pontil, 2009] and *optimism* in bandit algorithms [Audibert et al., 2009], will allow us to set σ^N to be of order $O(H/N)$ rather than $O(H/\sqrt{N})$.

Theorem 13. Consider Θ_I and $\hat{\Theta}_I^N(\lambda, \sigma^N)$ as defined in Equations (47) and (37). Suppose that $u_i(\mathbf{b}; v) \in [-H, H]$ for all $i \in [n]$, $\mathbf{b} \in B^n$ and $v \in V$. If $\lambda = \sqrt{2 \log(2|\Theta||M|/\delta)}$ and $\sigma^N = \frac{14H \log(2|\Theta||M|/\delta)}{3(N-1)}$ (with $|M| = n|B|^2 + |V|$), then:

$$\Pr [\Theta_I \subseteq \hat{\Theta}_I^N(\lambda, \sigma^N)] \geq 1 - \delta \quad (38)$$

Proof. To show the statement we need to show that with probability $1 - \delta$, for all $\theta \in \Theta_I$ it must be that $\theta \in \hat{\Theta}_I^N(\lambda, \sigma^N)$. Let $\theta \in \Theta_I$ and let x^* be the solution to the population problem, i.e. $x^* \in \arg \min_x \max_{i \in M} F_i(x; \theta, \phi)$.

⁸Now M also contains the density consistency constraints and thereby $|M| = n|B|^2 + |V|$.

By the empirical Bernstein bound (Theorem 4 of Maurer and Pontil [2009]), we have that with probability at least $1 - \kappa$:

$$F_j^N(x^*; \theta) \leq F_j(x^*; \theta, \phi) + \sqrt{\frac{2\text{Var}_N(f_j(x^*; \theta, \omega)) \log(2/\kappa)}{N}} + \frac{14H \log(2/\kappa)}{3(N-1)} \quad (39)$$

Since x^* is feasible $F_j(x^*; \theta, \phi) \leq 0$ and thereby with probability $1 - \kappa$:

$$F_j^N(x^*; \theta) - \sqrt{\frac{2\text{Var}_N(f_j(x^*; \theta, \omega)) \log(2/\kappa)}{N}} \leq \frac{14H \log(2/\kappa)}{3(N-1)} \quad (40)$$

Hence, also:

$$F_j^N(x^*; \theta) - \sqrt{\frac{2\text{Var}_N(f_j(x^*; \theta, \omega)) \log(2/\kappa)}{N}} \leq \frac{14H \log(2/\kappa)}{3(N-1)} \quad (41)$$

By the union bound the latter happens for all $j \in M$ with probability $1 - \kappa|M|$. Thus setting $\kappa = \frac{\delta}{|\Theta||M|}$, we have that with probability $1 - \frac{\delta}{|\Theta|}$:

$$\max_{j \in M} F_j^N(x^*; \theta) - \sqrt{\frac{2\text{Var}_N(f_j(x^*; \theta, \omega)) \log(2|\Theta||M|/\delta)}{N}} \leq \frac{14H \log(2|\Theta||M|/\delta)}{3(N-1)} \quad (42)$$

The latter then also holds for the minimum over x of the LHS:

$$\min_x \max_{j \in M} F_j^N(x; \theta) - \sqrt{\frac{2\text{Var}_N(f_j(x; \theta, \omega)) \log(2|\Theta||M|/\delta)}{N}} \leq \frac{14H \log(2|\Theta||M|/\delta)}{3(N-1)} \quad (43)$$

Hence, $\theta \in \hat{\Theta}_I^N(\lambda, \sigma^N)$ for the λ and σ^N stated in the Theorem. The Theorem then follows by a union bound over Θ . ■

The latter theorem is asymptotically an improvement over Theorem 12, when the variances of the constraints are not as large as their worst case bound of $4H^2$, which is essentially what is assumed in Theorem 12. However, one drawback of this theorem is that the new estimated set $\hat{\Theta}_I^N(\lambda, \sigma^N)$ requires solving more than linear programs. In particular for every parameter $\theta \in \Theta$ we need to solve an optimization problem of the form:

$$\min_x \max_j \left(\langle \alpha_j, x \rangle - \kappa_j \sqrt{\sum_{b, b' \in S} \gamma_{b, b'} \left(\sum_v (x_{vb} \beta_{jvb} - x_{vb'} \beta_{jvb'}) \right)^2} \right) \quad (44)$$

The negative part is a concave function of x , as it can be thought of as a norm of a vector that is a linear function of x . Thus this is a non-convex minimization problem. Thus even though it offers a statistical improvement, it is at the expense of computational efficiency.

5.2 Other set inference approaches

Here, we use existing set inference approaches and adapt them to the problem at hand. The first approach exploits the mapping between the reduced form parameter (the distribution of bids) and the identified set via the linear program, and the second uses subsampling approaches to set inference.

5.2.1 Inference via the Bayesian Bootstrap

In this section, we describe approaches to inference on the identified set via a Bayesian Bootstrap for both parametric and nonparametric models. We start with the nonparametric case. As a reminder, the identified set in this nonparametric case is

$$\Pi_I = \{\pi : \min_x \max_{j \in M} F_j(x; \pi, \phi) \leq 0\} \quad (45)$$

This is a case of a separable problem where knowing $\phi^* = \{\phi^*(b_1), \dots, \phi^*(b_B)\}$ we can solve for the identified set via the mapping:

$$\Pi(\phi) = \{\pi : \min_x \max_{i \in M} F_i(x; \pi, \phi^*) \leq \sigma^N\} \quad (46)$$

for some decaying tolerance constant σ^N (which can be set to zero). Inference on parameter sets in separable models is analyzed in (Kline and Tamer [2016]) where (posterior) probability statements related to the identified set⁹ can be computed using the mapping between the reduced form parameters, $\phi(\mathbf{b})$ in this case, and the structural parameters π . Intuitively, for every draw s from the posterior for ϕ (this is a multinomial distribution and so obtaining draws from the multinomial is simple using the Bayesian bootstrap¹⁰) we can solve via (46) for a “copy” of the implied identified set Π_N^s by solving the LP above. Using this procedure, we can get a sequence $\{\Pi_N^s\}_{s=1}^S$ that we can use to answer probability statements about the identified set Π_I . The computational constraint in this nonparametric problem is that the parameter of interest π is a vector of probabilities with $|V|$ support points. So, the bigger $|V|$ is the larger the number of parameters the more difficult it is to solve for the identified set via (23) above. Hence, with the finite support condition on the

⁹Examples of such statements are: the probability that the identified set belongs to a particular set (e.g, when Θ_I is an interval, the probability that this interval lies in $[a, b]$) or the probability that a particular vector belongs to the identified set, etc.

¹⁰For example, under a limiting uninformative Dirichlet prior for $\phi(\mathbf{b})$, the posterior for this $\phi(\mathbf{b})$ approaches the Dirichlet posterior $Dir(n_1, \dots, n_B)$ where $n_j = \sum_{i=1}^N 1[b_i = b_j]$. Therefore, using results that connect the Gamma and Dirichlet distributions, a draw from given posterior can be approximated by a weighted Bootstrap with Gamma weights. See also Chamberlain and Imbens [2003].

bids standard (Bayesian) inference on $\phi(\mathbf{b})$ can be easily mapped into inference on the set Π_I . The exact same bootstrap procedure can be used to provide (posterior) probability statement on identified intervals $[L, U]$ that are defined through the mapping in (25) above. This is relevant if one is interested in linear functional of the distribution. Inference on such objects reduces the computational burden substantively.

In the parametric case, we assume that the distribution function of v belongs to a parametric class that is known up to a finite dimensional parameter θ and so the LP in Theorem 6 is modified by setting $\pi(v) = \pi(v, \theta)$. Then the parameter of interest is θ and the identified set for θ takes the form:

$$\Theta_I = \{\theta : \min_x \max_{j \in M} F_j(x; \theta, \phi) \leq 0\} \quad (47)$$

and so the separable mapping between ϕ^* and θ takes the form again of

$$\Theta_I(\phi^*; \sigma^N) = \{\theta : \min_x \max_{j \in M} F_j(x; \theta, \phi^*) \leq \sigma^N\} \quad (48)$$

for some decaying tolerance constant σ^N (which can be set to zero).

Here again, we can use the approach in (Kline and Tamer [2016]) to answer probability statements on Θ_I . The computational step here is much easier than the nonparametric case as now solving for the identified set for every draw from the posterior for ϕ is a lower dimensional problem. The Bayesian bootstrap is used to draw vectors of bid probabilities from the multinomial distribution for bids.

In addition to this Bayesian bootstrap approach, we also provide methods based on subsampling next.

5.2.2 Parametric Case via Subsampling

Subsampling methods can be used to conduct inference on the identified set in a large sample framework. In particular, let θ be the parameter of interest. Again, our problem is isomorphic to one where the objective function $Q(\theta)$ is as follows:

$$Q(\theta) = \min_x \max_{j \in M} F_j(x; \theta) \quad (49)$$

with a the corresponding sample analogue

$$Q^N(\theta) = \min_x \max_{j \in M} F_j^N(x; \theta) \quad (50)$$

We first re-state the above random variable as a minimax problem over convex and compact sets. Let $p \in \Delta_M$ lie on the simplex on M constraints. Then $Q^N(\theta)$ is equivalently defined as:

$$Q^N(\theta) = \min_x \sup_p \sum_{j \in M} p_j F_j^N(x; \theta) \quad (51)$$

and

$$Q^*(\theta) = \min_x \sup_p \sum_{j \in M} p_j F_j(x; \theta) \quad (52)$$

where Q^* is zero for all feasible θ 's. In addition, following the approach in Chernozhukov, Hong, and Tamer [2007], we get a confidence region for the identified set $\Theta_I = \{\theta : Q^*(\theta) \leq 0\}$ by studying the asymptotic distribution¹¹ of

$$\mathcal{C}_N = \sup_{\theta \in \Theta_I} \min_x \sup_p \sum_{j \in M} p_j F_j^N(x; \theta)$$

Let $\tau_N(1 - \alpha)$ be the $(1 - \alpha)$ -quantile of \mathcal{C} where \mathcal{C} is the nondegenerate limit of $\sqrt{N}\mathcal{C}_N$. Then, define the $C_N(1 - \alpha)$ as follows

$$C_N(1 - \alpha) = \{\theta : Q^N(\theta) \leq \tau_N^+(1 - \alpha)\}$$

where $\tau_N^+(1 - \alpha) = \max(\tau_N(1 - \alpha), 0)$. Notice here that the event $\Theta_I \subseteq C_N(1 - \alpha)$ is equivalent to the event $\mathcal{C}_N \leq \tau_N^+(1 - \alpha)$. The next Theorem states the result.

Theorem 14. *Assume that N increases to infinity, and $\sqrt{N}(\mathcal{C}_N)$ converges in distribution to \mathcal{C} , a nondegenerate random variable. Then the set $C_N(1 - \alpha)$ defined above has the following coverage property:*

$$\lim_{N \rightarrow \infty} \Pr[\Theta_I \subseteq C_N(1 - \alpha)] = 1 - \alpha.$$

The asymptotic distribution \mathcal{C} above can be characterized using for example results from Shapiro [2009] and sufficient conditions for nondegeneracy are given there which requires second moments to hold (these hold in our case trivially since we maintain the assumptions that bids take finitely many values).

Finally, to be able to feasibly implement the above approach, one needs to get a value for the cutoff $\tau_N(1 - \alpha)$. It is clear here that the standard nonparametric bootstrap may not work. Even

¹¹Alternatively, we can define the objective function $\tilde{Q}^N(\theta) = [Q^N(\theta)]_+$ which is the positive part of $Q^N(\theta)$. The identified set now is the minimizer of $\tilde{Q}^N(\theta)$ and then similar approaches to Chernozhukov et al. [2007] can be used to construct a CI based on subsampling.

if the asymptotic distribution is normal¹², we may not be able to estimate its variance because it depends on the number of binding constraints. Given the nondegeneracy of the limit, one approach is to use subsampling to compute the cutoff. This can be accomplished by getting m subsamples of the data such that $m/N \rightarrow 0$ as $N \rightarrow \infty$. For every subsample, we compute C_N^m using a preliminary estimate $\hat{\Theta}_I$ and using the sequence $\{\mathcal{C}_N^m\}_{m=1}^M$ to get its upper $(1 - \alpha)$ quantile. This will result in an estimate of $C_N(1 - \alpha)$. We can then use that as our new $\hat{\Theta}_I$ and iterate one or two times. This approach was implemented in the Monte Carlo section below and in the empirical application and it provided adequate results.

6 Monte-Carlo simulations for common value auctions

In this section, we examine the techniques of Section 5 to obtain an estimated set on simulated data. We assume in all the simulations that the values and bids are taken from discrete sets, respectively V and B . In the whole section, we fix the following parameters:

- The maximum value/bid H is given by $H = 20$.
- The set of possible bids is given by $B = \{0, \dots, H\}$. The set of possible common values is $V = B = \{0, \dots, H\}$.

We generate equilibrium observations as follows: First, we generate a density $f(v)$ of valuations with support V . We consider the four following densities:

- $f_n(\cdot)$ is the density of a normal random variable with mean parameters $\mu = 4$ and standard deviation parameter $\sigma = 1$, discretized and truncated to have support V , i.e.

$$f_n(v) = \frac{\exp(-(v - \mu)^2/\sigma^2)}{\sum_{w \in V} \exp(-(w - \mu)^2/\sigma^2)} \quad (53)$$

- $f_p(\cdot)$ is the density of a Poisson distribution with parameters $\lambda = 4$, truncated inside V , i.e.:

$$f_p(v) = \frac{\lambda^v \exp(-\lambda)/v!}{\sum_{w \in V} \lambda^w \exp(-\lambda)/w!} \quad (54)$$

- $f_b(\cdot)$ is the density of a binomial random variable with probability $p = 0.2$ and number of draws $n = H = 20$, i.e.:

$$f_b(v) = \binom{N}{v} p^v (1 - p)^{H-v} \quad (55)$$

¹²The distribution of \mathcal{C} is likely to be the supremum of a gaussian process since the optimal value of the LP is generally not unique.

- $f_g(\cdot)$ is the density of a geometric random variable with probability $p = 0.2$ truncated to have support V , i.e.:

$$f_g(v) = \frac{p(1-p)^v}{\sum_{w \in V} p(1-p)^w} \quad (56)$$

For each of those densities, we then generate one distribution of equilibrium bids ϕ , through solving the BCE linear program for the given distribution of values and with variables $\phi(\mathbf{b})$ rather than $\pi(v)$. We then generate N samples of bid vectors from ϕ , to generate the observed empirical bid distribution ϕ_N .

In the non-parametric setting, since we cannot directly formulate the estimated set of the variance (as it cannot be written as $E[m(v)]$ for some m), we obtain a superset for the identified set as follows: First obtain upper and lower bounds $E_{min}, E_{max}, E_{min}^{2nd}, E_{max}^{2nd}$ for the first and second moments respectively. Then set the bounds on the standard deviation by the conservative ones: $E_{min}^{2nd} - E_{max}^2 \leq Var \leq E_{max}^{2nd} - E_{min}^2$. A computationally more tedious procedure of estimating the identified set for the variance is to first get the identified set for the (discrete) distribution of V and then using that we can “solve” for a bound on the variance¹³.

In the parametric case, bounds on the variance can be obtained directly from recovering bounds on the possible parameters for the distribution we consider, and tight bounds on the parameters imply tight bounds on the second and higher order moments. All simulations in the parametric case are therefore presented in terms of identified and estimated sets on the parameters of the distribution of values.

6.1 Comparing the parametric and non-parametric identified sets

In this section, we compare the identified sets when we use and do not use parametric knowledge of the distribution of v . Figure 1 uses the true distribution ϕ and shows:

- The set of (mean, standard deviation) pairs that belong to the identified set when solving the linear program with parametric constraints, in brown.
- The set of (mean, standard deviation) pairs that belong to the identified set when solving the original LP without parametric constraints, in green.

We do so for the four different distributions of the common value (gaussian, poisson, binomial and geometric) mentioned above. We remark that the non-parametric linear programs seems to recover

¹³For example, for every “draw” from the identified set for the distribution of V , we can obtain a variance. We can repeat the process to build the identified set for the variances.

the mean of the common value accurately in all cases. However, the bounds obtained on the second moment/standard deviation of the distribution of common values are far from being tight.

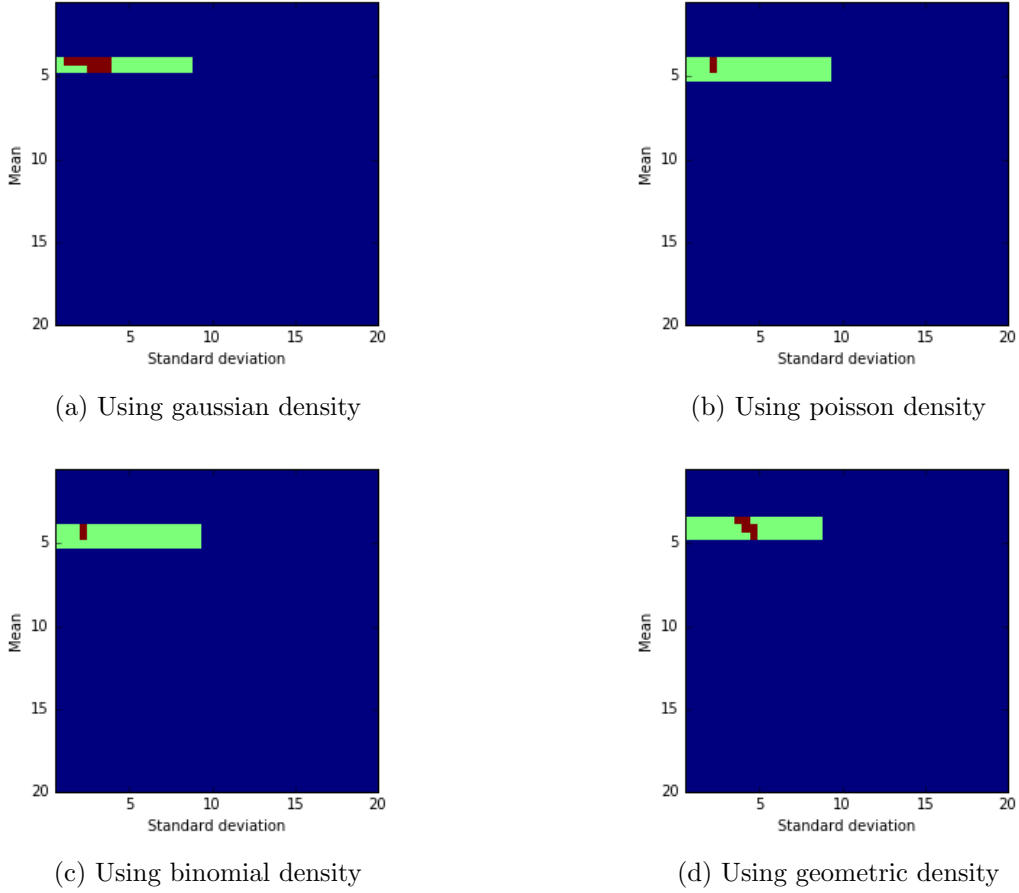


Figure 1: Plots of the non-parametric and parametric identified sets in the (mean, standard deviation) space

Figure 2 compares the identified set for the true distribution to the estimated set using Hoeffding with $\delta = 0.10$, as described in Section 5.1 for the Gaussian density function, for a number of samples $N \in \{10^3, 10^4, 10^5, 10^6\}$. We see that as N grows larger, the estimated set grows smaller and smaller and closer to the true identified set.

6.2 Parametric identified sets

In this section, we characterize the identified and estimated sets in the parametric case for the four distributions described above. For the gaussian distribution, we provide figures of the identified parameters in the (μ, σ) space. For the other, 1-dimensionally parameterized distributions, we provide intervals for the parameter of the chosen distribution. We consider two different techniques




Figure 2: Plots of the parametric identified sets and the non-parametric estimated set through Hoeffding with $\delta = 0.10$ in the (mean, standard deviation) space

to determine the estimated set:

- Using tolerances determined through Hoeffding's inequality, as discussed in Section 5.1.2
- Using quantiles of the tolerance via subsampling, as discussed in Section 5.2.2

In all figures, the brown region is the true identified set while the union of the brown and the green region is the estimated set. Figure 3 plots the identified and estimated set when using a Gaussian distribution for the common value and tolerances determined through Hoeffding with 90 percent confidence ($\delta = 0.10$), for the number of samples $N \in \{10^3, 10^4, 10^5, 10^6\}$. We remark that the number of samples needs be large (of the order of at least 10^5) for Hoeffding to perform well.

Figure 4 plots the identified and estimated set when using a Gaussian distribution for the common value and tolerances determined through subsampling and quantile estimation for the 90, 95 and 99 percent quantiles, as seen in Section 5.2.2; we only use $N = 100$ samples for the bid




Figure 3: Plots of the parametric identified sets and the parametric estimated set through Hoeffding with $\delta = 0.10$ in the (μ, σ) space

distribution in the three figures and $k = 50$ subsamples of size $s = N/4 = 25$. We note that the quantile estimation technique covers the true identified set fairly sharply even though N is only equal to 100 and hence should be preferred to Hoeffding when small amounts of data are available.

Table 6.2 gives the true identified interval for the parameters of the poisson, binomial and geometric distributions and the estimated interval using subsampling for quantile estimation, for the 90, 95 and 99 percent quantiles – see 5.2.2 – for a bid distribution sampled with $N = 100$. We see that for the binomial distribution, the estimated set is very close to the true identified set. While the recovered sets for the poisson distribution are not as sharp as the recovered sets for the binomial distribution, they still restrict the space of possible parameters in a reasonable way: the recovered interval is $[1.5, 6.5]$ while the space of possible parameters is $[0, 20]$. However, the recovered sets for the geometric distribution contains more than half of the possible parameters,




Figure 4: Plots of the parametric identified sets and the parametric estimated set through the quantile method in the (μ, σ) space for $N = 100$

which is unsatisfying. A reason for this comes from the fact that the geometric distribution has a much higher variance than a poisson or binomial with comparable means, hence there is a lot of variability accross different subsamples that can lead to large tolerances.

Distribution	True identified set	0.90 quantile set	0.95 quantile set	0.99 quantile set
Binomial	[0.19,0.22]	[0.17,0.25]	[0.16,0.25]	[0.15,0.27]
Poisson	[4.0,4.5]	[1.5,6.5]	[1.5,6.5]	[1.5,6.5]
Geometric	[0.17,0.22]	[0.06,0.56]	[0.04,0.66]	[0.02,0.70]

Using $N = 500$ samples and $k = 50$ subsamples of size $s = N/4 = 125$, we obtain Table 6.2. We can see that the estimated sets for the binomial and geometric distributions are now significantly tighter.

Distribution	True identified set	0.90 quantile set	0.95 quantile set	0.99 quantile set
Binomial	[0.19,0.22]	[0.18,0.24]	[0.17,0.25]	[0.17,0.25]
Poisson	[4.0,4.5]	[3.5,5.0]	[3.0,5.5]	[2.5,5.5]
Geometric	[0.17,0.22]	[0.16,0.25]	[0.15,0.28]	[0.14,0.29]

Finally, figure 5 plots the identified and estimated set when using a Gaussian distribution for the common value and tolerances determined through subsampling and quantile estimation for the 90, 95 and 99 percent quantiles, as Figure 5. However, we now use $N = 500$ samples for the bid distribution in the three figures and $k = 50$ subsamples of size $s = N/4 = 125$. We note that the estimated set now almost coincides with the true identified set.

7 Empirical Illustration: OCS Auctions

In this section, we illustrate our framework for common value auctions on real data. We use the Outer Continental Shelf (OCS) Auction Dataset that was used in the seminar work of Hendricks and Porter (see Hendricks and Porter [1988]). The dataset contains bidding information on 3036




Figure 5: Plots of the parametric identified sets and the parametric estimated set through the quantile method in the (μ, σ) space for $N = 500$

tract auctions in Louisiana and Texas. In particular, for each auction, our dataset contains:

- The acreage of the tract
- The total bid of each participant in the auction

We assume that the bidders participate in a first price common value auction, where the value is defined per acre; our goal is i) to show that indeed, the bidders' behavior in the data can be explained by a common value auction (via the testable restriction of whether the estimated identified set is empty), and ii) to recover the first and second order moments of the distribution of said value. This is under weak assumptions on information in that the framework allows bidders in different auctions to know more information about the environment.

Pre-processing of data: The dataset contains 3036 auction with varying number of players. We consider 2-player common value auctions, hence we only keep the entries in the dataset that contain exactly 2 bidders; there are 584 such auctions. We model the two bidders as being the same over the 584 auctions, and assign bidders' identities to be 1 or 2 uniformly at random in each auction. Many of the bids we have are zero and Figure 6 plots the distribution of bids; it has mean \$991.48 and standard deviation \$1825.43.




Figure 6: Distribution of bids

We assume the distribution of the common value per acre has bounded finite support $V = \{0, \dots, H\}$. We renormalize the bids per acre to be in $[0, \lceil \frac{H}{2} \rceil]$ – we pick $H/2$ because the common value could have a distribution whose support goes beyond the observed bids –, and discretize the set of bids to be $\{0, \dots, \lceil \frac{H}{2} \rceil\}$; we do so by rounding each renormalized bid in each auction to the closest integer. We remark that the dataset contains a few outliers whose bid per acre is significantly higher than in all other auctions; we therefore delete the auctions that contain bids over threshold $t = \$20000$. We further assume that the distribution of the common value is given by a truncated normal distribution that takes discrete values in $\{0, \dots, H\}$, exactly as described in the simulations of Section 6, and parameterize all optimizations problems we solve accordingly.

Results: In all figures, the value of (μ, σ) are given as the values in dollars instead of the corresponding discretized and renormalized value, for the sake of comparison with the \$20,000 threshold and the corresponding maximum value of \$40,000. Figure 7 plots a heatmap of the estimated set as a function of the chosen tolerance in the (μ, σ) space for two different values of H . Each color on the heatmap corresponds to a tolerance level, and the mapping from tolerance levels to colors is given by the colorbar on the right of each figure. The color that is assigned to any given (μ, σ) pair corresponds to the minimum level of tolerance that the analyst needs to add to the equilibrium constraints for (μ, σ) to belong to the estimated set; therefore, the heatmap shows how the estimated set grows as the tolerance picked by the analyst increases.

Figure 8 plots the estimated set when using the tolerance determined by the subsampling




Figure 7: Heatmap of minimum tolerance needed for (μ, σ) pairs to belong to the estimated set

approach of Section 5.2.2, in green using 95% level, and compares it to the estimated set using the minimum tolerance for which the estimated set is non-empty, in brown, for $H = 400$. No matter what method is used for picking the tolerance and determining the corresponding confidence intervals, the region that is obtained cannot be possibly smaller than the minimum tolerance set unless it is empty; in this sense, the minimum tolerance set is the best possible set that one could hope to obtain through *any* method for determining tolerances. We remark that i) the estimated set for (μ, σ) remains small relatively to the upper bound of \$20,000 on the bids and of \$40,000 on the maximum common value, and ii) the estimated set is not much bigger compared to the best estimated set we could hope to obtain, indicating that on top of covering the true identified set with high probability, our techniques cannot possibly overestimate the size of the identified by too much. We can see from the confidence regions that the mean of the common values varies from zero 0 to around \$4000 while the standard deviation varies from close to zero to 6000. It is possible to estimate these means as functions of covariates and hence to allow for observed heterogeneity.

The above plots the mean and variance of a normal density that is then truncated. Since the truncated normal has a different mean and variance than the underlying normal, we plot in Figure 9 the minimum tolerance and the estimated sets as a function of this overall mean and standard deviation of the distribution of common values that we identify, instead of the parameters μ, σ of the truncated normal distribution. We obtain these plots by computing a mapping from (μ, σ) pairs to (mean, standard deviation) pairs, and plot the image of Figure 8 by said mapping in the (mean, standard deviation) space. We note that the estimated set is fairly small, indicating that




Figure 8: Estimated sets using the 90, 95 and 99 percent quantiles, for $H = 400$.

our approach identifies the first two moments of the true, underlying distribution of the common value in an accurate fashion, despite only having access to a limited number of samples.




Figure 9: Estimated sets using the 90, 95 and 99 percent quantiles, for $H = 400$.

8 Conclusion

We provide a framework for inference on auction fundamentals without making strong restrictions on information. Using data, we use the recent results in theory to characterize the identified set for these primitives by exploiting the linear programming structure of the set of Bayesian correlated equilibria. We have several applications of the approach mainly to common value and private value auctions and other scenarios. Our results can also be used by mechanism designers in that the data allows us to restrict the domain of signal/state of the world distribution which would lead to sharper mechanisms. The paper also provides various approaches to inference, and proposes finite sample approaches to building confidence regions for sets in partially identified models.

References

- J.-Y. Audibert, R. Munos, and C. Szepesvri. Explorationexploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19):1876 – 1902, 2009. ISSN 0304-3975. doi: <https://doi.org/10.1016/j.tcs.2009.01.016>. URL <http://www.sciencedirect.com/science/article/pii/S030439750900067X>. Algorithmic Learning Theory.
- R. J. Aumann. Correlated equilibrium as an expression of bayesian rationality. *Econometrica: Journal of the Econometric Society*, pages 1–18, 1987.
- D. Bergemann and S. Morris. Robust predictions in games with incomplete information. *Econometrica*, 81(4):1251–1308, 2013. ISSN 1468-0262. doi: 10.3982/ECTA11105. URL <http://dx.doi.org/10.3982/ECTA11105>.
- D. Bergemann and S. Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016. URL <https://econtheory.org/ojs/index.php/te/article/view/20160487/0>.
- D. Bergemann, B. Brooks, and S. Morris. First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica*, 85(1):107–143, 2017. ISSN 1468-0262. doi: 10.3982/ECTA13958. URL <http://dx.doi.org/10.3982/ECTA13958>.
- G. Chamberlain and G. W. Imbens. Nonparametric applications of bayesian inference. *Journal of Business & Economic Statistics*, 21(1):12–18, 2003.
- V. Chernozhukov, H. Hong, and E. Tamer. Estimation and confidence regions for parameter sets in econometric models. *Econometrica*, 75(5):1243–1284, 2007.
- E. Guerre, I. Perrigne, and Q. Vuong. Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3):525–574, 2000.
- K. Hendricks and R. H. Porter. An empirical study of an auction with asymmetric information. *The American Economic Review*, pages 865–883, 1988.
- D. Hoy, D. Nekipelov, and V. Syrgkanis. Robust data-driven efficiency guarantees in auctions. *CoRR*, abs/1505.00437, 2015. URL <http://arxiv.org/abs/1505.00437>.
- B. Kline and E. Tamer. Bayesian inference in a class of partially identified models. *Quantitative Economics*, 7(2):329–366, 2016.

- T. Li, I. Perrigne, and Q. Vuong. Structural estimation of the affiliated private value auction model. *The RAND Journal of Economics*, 33(2):171–193, 2002. ISSN 07416261. URL <http://www.jstor.org/stable/3087429>.
- L. Magnolfi and C. Roncoroni. Estimation of discrete games with weak assumptions on information. Technical report, mimeo, 2016.
- A. Maurer and M. Pontil. Empirical Bernstein Bounds and Sample Variance Penalization. *ArXiv e-prints*, July 2009.
- T. Peel, S. Anthoine, and L. Ralaivola. Empirical bernstein inequalities for u-statistics. In *NIPS*, 2010.
- A. Shapiro. 5. *Statistical Inference*, pages 155–252. 2009. doi: 10.1137/1.9780898718751.ch5. URL <http://epubs.siam.org/doi/abs/10.1137/1.9780898718751.ch5>.
- V. Syrgkanis and E. Tardos. Composable and efficient mechanisms. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 211–220, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-2029-0. doi: 10.1145/2488608.2488635. URL <http://doi.acm.org/10.1145/2488608.2488635>.

A Non-sharpness of the BBM bound

We examine the upper and lower bounds on the mean from BMM. In a common value auction an easy worst-case bound on the mean of the common value distribution comes from the following fact: at any BCE players are getting non-zero expected utility. Thus an obvious lower bound on mean value is the observed revenue. We can try to show that mean value is also not much more than the revenue or some function of the observed bids that is close to the revenue. For instance, the BBM bound shows that:

$$v(q) - \frac{1}{q^{(n-1)/n}} \int_0^{v(q)} (F(x))^{(n-1)/n} dx \leq B(q) \quad (57)$$

where $q \in [0, 1]$ is a quantile, $v(q) = F^{-1}(q)$, $F(x)$ is the cdf of the value distribution and $B(q)$ is the quantile function of the maximum bid distribution. This is equivalent to:

$$\frac{1}{q^\alpha} \int_0^q \alpha y^{\alpha-1} v(y) dy \leq B(q) \quad (58)$$

From this we can try to upper bound the mean of F (i.e. $\int_0^1 v(q) dq$) as a function of the revenue, which is simply the mean of the maximum value distribution (i.e. $\int_0^1 B(q) dq$). If the CDF is convex, $F(0) = 0$ and $F(H) = 1$, then $F(x) \leq x/H$ hence the BBM bound gives:

$$v(q) \left(1 - \frac{n}{2n-1} \frac{1}{H^{(n-1)/n}} \left(\frac{v(q)}{q} \right)^{(n-1)/n} \right) \leq B(q) \quad (59)$$

By integrating and if we let μ be the mean of the common value and R be the observed revenue of the auction:

$$\mu \leq R + \frac{n}{2n-1} \frac{1}{H^{(n-1)/n}} \int_0^1 \frac{v(q)^{(2n-1)/n}}{q^{(n-1)/n}} dq$$

This shows that there is a reasonable upper bound. For instance for the uniform distribution, i.e. $v(q) = q \cdot H$, this gives $\mu \leq R + \frac{n}{2n-1} H \int_0^1 q dq = R + H \frac{n}{2n-1} \frac{1}{2} \leq R + H/3$. Observe that if the uniform distribution is allowed in our set of possible value distributions then we cannot hope to prove any better upper bound.

For the case of two bidders we can now show that the BBM bound implies that the mean of the auction must be at most:

$$\mu \leq R \cdot \left(2\sqrt{\frac{H}{R}} - 1 \right) \quad (60)$$

where R is the observed revenue of the auction and H is the assumed upper bound on the distribution. Thus if the revenue is at least H/α for some α , then we can conclude that the mean must

leave in the range: $[R, R(2\sqrt{\alpha} - 1)]$ which is a non-trivial sharp identified set. In other words the ratio of the upper to lower bound of the sharp identified set is at most $2\sqrt{\alpha} - 1$.

For n bidders we show (see Section A.1 below) that the BBM bound implies that the mean of the distribution cannot be more than: $\sqrt{\frac{2n}{n-1}} \cdot R \cdot H$ or alternatively $\sqrt{\frac{2n}{n-1}} \cdot R \cdot (\frac{1}{n}R + L)$ if the density of values is bounded from below by $1/L$. Thus if the revenue is at least H/α , then we have that μ lies in the range: $[R, R \cdot \sqrt{\frac{2n}{n-1}}\alpha]$. In fact the above bound is tight if one only uses the deviation of BBM, i.e. if all we know is the inequality implied by these deviations, then there exists a distribution that satisfies that inequality and which achieves the above bound.

Non-sharpness of BBM: The upper bound of BBM can be very far from the identified set. Before we give a concrete example, we first argue why the bound from Bergemann et al. [2017] cannot possibly be sharp: The main point of the work of BBM was getting worst-case bounds on the revenue of a common value first price auction as a function of the distribution of values. For that reason in the analysis, the bid distribution is not part of the input. Their main approach is to claim that the revenue cannot be too small. The idea behind the analysis is as follows: since something is a BCE, no player wants to deviate to any other action. Hence, they do not want to deviate to a specific style of a deviation which is: conditional on your bid deviate uniformly at random above your bid (upwards deviation). Such deviation arguments were also for previously used in Hoy, Nekipelov, and Syrgkanis [2015], Syrgkanis and Tardos [2013] to give bounds on the welfare of non-truthful auctions such as the first price auction.

However, this means that the final bound given in (57) is the product of a subset of the best-response deviation constraints. If we actually knew what the bid distribution is in a BCE (which is the case in the econometrics task), then we wouldn't look at only these bid-oblivious upwards deviations. We would instead compute an optimal best-response bid for this given bid distribution and right the constraint that the player does not want to deviate to this tailored best-response.

We now give a concrete “extreme” example where the upper bound provided by (57) can be as large as twice the sharp bound. In fact, in this example the mean can be point identified, but the bound (57) gives a large interval.

Example. Consider the case of two bidders, where the observed maximum bid distribution is a singleton i.e. $\{b^*\}$. First we argue that in this case, we can conclude that the mean μ of the common value is equal to b^* . First, we know that the total expected utility of the bidders must be

non-negative, hence the mean of the common value is at least the expected revenue of the auction, which is equal to b^* , i.e. $\mu \geq b^*$.

Now suppose that the mean was strictly larger than b^* , i.e. $b^* + \epsilon$. We will show that this yields a contradiction.

First consider the case where in the support of the BCE, there exist bid vectors of the form $(b, 5)$ or $(5, b)$ for $b < 5$. Then a player when seeing a bid of $b < 5$, he wants to deviate to bidding $b^* + \zeta$ for some $\zeta \in (0, \epsilon)$. The reason is that with a bid of b he knows he is not winning and hence he is getting zero utility, while with a bid of $b^* + \zeta$ he knows he is deterministically winning, getting a value of μ and paying $b^* + \zeta < \mu$, yielding strictly positive utility. Thus it must be that the BCE contains only one bid vector, i.e. (b^*, b^*) . But under this BCE, it is obvious that the expected value conditional on winning is μ and it is also clear that one of the two players is winning with probability less than 1, i.e. $1 - \delta$. Thus if this player deviates to $b^* + \zeta$, he wins deterministically and pays $b^* + \zeta$. Thus the net effect of this deviation is $\mu - b^* - \zeta - (1 - \delta)(\mu - b^*) = \delta \cdot (\mu - b^*) - \zeta = \delta\epsilon - \zeta$. Taking $\zeta \rightarrow 0$, yields a deviation with positive net effect.

Thus we get a contradiction, and we conclude that the mean is point identified and $\mu = b^*$.

Now consider the bound that is derived from Equation (57) for this case of a bid distribution. First we note that for a singleton bid distribution we have that $B(q) = b^*$ for all $q \in [0, 1]$. Thus the constraint on the distribution of values implied by Equation (58) is simplified to:

$$\forall q \in [0, 1] : \frac{1}{\sqrt{q}} \int_0^q \frac{v(y)}{2\sqrt{y}} dy \leq b^* \quad (61)$$

We now give a distribution of values, i.e. a function $v(\cdot)$, that satisfies the above constraint and whose mean μ^* is at least $2 \cdot b^* - (b^*)^2/H$, where H is an externally assumed upper bound on the distribution of values. If H is large, then this gives an interval such that the ratio of the upper to the lower bound is of size 2, which is far from point identification.

Consider a value function $v(\cdot)$, as follows (two point-mass distribution):

$$v(q) = \begin{cases} H & \text{if } q \geq \theta \\ 0 & \text{o.w.} \end{cases} \quad (62)$$

Then the constraint simplifies to:

$$\begin{aligned} b^* &\geq \frac{1}{\sqrt{q}} \int_{\min\{\theta, q\}}^q \frac{H}{2\sqrt{y}} dy \\ &= \frac{1}{\sqrt{q}} H [\sqrt{y}]_{\min\{\theta, q\}}^q = H \left(1 - \sqrt{\frac{\min\{\theta, q\}}{q}} \right) \end{aligned}$$

If $\theta \leq q$, then this constraint is definitely satisfied, since $b^* \geq 0$. Thus we need to only check the constraints for $q > \theta$. The tightest of these constraints with respect to θ , is when $q = 1$, leading to:

$$\theta \geq \left(\frac{H - b^*}{H} \right)^2 \quad (63)$$

By setting θ to be equal to the above lower bound, we get a feasible value function and this value function has mean:

$$\mu^* = H(1 - \theta) = H - \frac{(H - b^*)^2}{H} = \frac{H^2 - H^2 - (b^*)^2 + 2b^*H}{H} = 2b^* - \frac{(b^*)^2}{H}$$

which concludes the point of the example. ■

A.1 Upper bound on mean of common value

Theorem 15. *Let $R = \mathbb{E}_{b \sim D} [\max_i b_i]$ is the expected revenue of a first price auction when bids are drawn from some BCE and let $\mu = \mathbb{E}[v]$ be the expected common value. If we assume that $v \in [0, H]$ and that the inverse of the CDF of the distribution of values is continuously differentiable in $[0, H]$, then it holds that:*

$$\mu \leq \sqrt{\frac{2n}{n-1} H \cdot R} \quad (64)$$

If we also assume that the inverse of the CDF of values is L -lipschitz (equivalently the density of values is bounded below by $1/L$), then we don't need to assume an upper bound on the distribution and we alternatively get:

$$\mu \leq \sqrt{\frac{2n}{n-1} \left(\frac{1}{n} R + L \right) \cdot R} = R \sqrt{\frac{2}{n-1}} + \sqrt{\frac{2n}{n-1} L \cdot R} \quad (65)$$

Proof. By the main theorem of Bergemann et al. [2017] we have that if F is the CDF of common values and H is the CDF of the maximum bid distribution at a BCE, while $v(\cdot) = F^{-1}(\cdot)$ and $B(\cdot) = H^{-1}(\cdot)$ are the quantile functions of these distributions, then for any $q \in [0, 1]$:

$$B(q) \geq v(q) - \frac{1}{q^\alpha} \int_0^{v(q)} (F(x))^\alpha dx \quad (66)$$

with $\alpha = \frac{n-1}{n}$. Let $I(q)$ denote the right hand side.

We first do a change of variables in the integral: let $x = v(y)$. Then $dx = v'(y)dy$ and since the integral ranges from $x \in [0, v(q)]$, in the new variable the integral ranges from $y \in [0, q]$. Hence:

$$I(q) = v(q) - \frac{1}{q^\alpha} \int_0^q (F(v(y)))^\alpha v'(y) dy = v(q) - \frac{1}{z^\alpha} \int_0^q y^\alpha v'(y) dy \quad (67)$$

Applying integration by parts on the integral we have:

$$\begin{aligned}
I(q) &= v(q) - \frac{1}{q^\alpha} \left([y^\alpha v(y)]_0^q - \int_0^q \alpha y^{\alpha-1} v(y) dy \right) \\
&= v(q) - v(q) + \frac{1}{q^\alpha} \int_0^q \alpha y^{\alpha-1} v(y) dy \\
&= \frac{1}{q^\alpha} \int_0^q \alpha y^{\alpha-1} v(y) dy
\end{aligned}$$

Now observe that, $\mu = \int_0^1 v(q) dq$. This can be easily verified pictorially, but also algebraically as follows:

$$\mu = \int_0^H 1 - F(x) dx = \int_0^1 (1 - F(v(q))) v'(q) dq = [(1 - q)v(q)]_0^1 + \int_0^1 v(q) dq = \int_0^1 v(q) dq \quad (68)$$

Similarly, $R = \int_0^1 B(q) dq$.

Thus integrating the inequality $B(q) \geq I(q)$ over q we get:

$$R \geq \int_0^1 \frac{1}{q^\alpha} \int_0^q \alpha y^{\alpha-1} v(y) dy dq \quad (69)$$

Exchanging the integration order, we get:

$$R \geq \int_0^1 \int_y^1 \frac{1}{q^\alpha} \alpha y^{\alpha-1} v(y) dq dy = \int_0^1 \alpha y^{\alpha-1} v(y) \int_y^1 \frac{1}{q^\alpha} dq dy \quad (70)$$

$$= \int_0^1 \alpha y^{\alpha-1} v(y) \left[\frac{q^{1-\alpha}}{1-\alpha} \right]_y^1 dy \quad (71)$$

$$= \int_0^1 \frac{\alpha}{1-\alpha} y^{\alpha-1} v(y) (1 - y^{1-\alpha}) dy \quad (72)$$

$$= \int_0^1 \frac{\alpha}{1-\alpha} v(y) (y^{\alpha-1} - 1) dy \quad (73)$$

Thus by re-arranging we have that:

$$\int_0^1 v(y) (y^{\alpha-1} - 1) dy \leq \frac{1-\alpha}{\alpha} R \quad (74)$$

Therefore if we are given as fixed the revenue of the auction R , then the mean of the common value distribution can be at most the solution to the following optimization program over the set of all possible quantile functions $v(\cdot)$:

$$\max_{v(\cdot) \in [0, H]} \int_0^1 v(y) dy \quad (75)$$

$$s.t. \int_0^1 v(y) (y^{\alpha-1} - 1) dy \leq \frac{1-\alpha}{\alpha} R \quad (76)$$

Now consider the function $w(y) = y^{\alpha-1} - 1$. This is a monotone decreasing and non-negative function of y , starting from ∞ at $y = 0$ and ending at 0 at $y = 1$.

The goal of the above linear program is to push as much possible value in $v(y)$ as possible, while satisfying the constraint. Observe that a quantity $v(y)$ is multiplied by a smaller value in the above constraint than any $v(y')$ for $y' \leq y$. Thus it is easy to see that the optimal solution to the above linear program puts as much value on the high y 's and sets the remainder y 's to zero (another way of arguing this is considering arbitrarily small discretizations of the y space and then arguing that the solution to the linear program with respect to $v(y)$, has the above form). Thus the optimal $v(\cdot)$ takes the form:

$$v(y) = \begin{cases} 0 & \text{if } y \leq \theta \\ H & \text{o.w.} \end{cases} \quad (77)$$

for some threshold $\theta \in [0, 1]$.

Thus the optimal value of the above program simplifies to:

$$\max_{\theta \in [0, 1]} H(1 - \theta) \quad (78)$$

$$s.t. \quad H \int_{\theta}^1 (y^{\alpha-1} - 1) dy \leq \frac{1 - \alpha}{\alpha} R \quad (79)$$

Thus we want to minimize θ , such that:

$$\frac{1 - \alpha}{\alpha} \frac{R}{H} \geq \int_{\theta}^1 (y^{\alpha-1} - 1) dy = \frac{1}{\alpha} - \frac{\theta^{\alpha}}{\alpha} - 1 + \theta \quad (80)$$

Case of $n = 2$ For $n = 2$, i.e. $\alpha = 1/2$, we can exactly solve the latter, since it corresponds to a quadratic inequality, i.e.:

$$\frac{R}{H} \geq 1 - 2\sqrt{\theta} + \theta = (1 - \sqrt{\theta})^2 \quad (81)$$

which yields that: $\sqrt{\theta} \geq 1 - \sqrt{\frac{R}{H}}$ or equivalently, $\theta \geq 1 + \frac{R}{H} - 2\sqrt{\frac{R}{H}}$. Therefore, the highest possible mean is at most: $H(1 - \theta) = H \left(2\sqrt{\frac{R}{H}} - \frac{R}{H} \right) = 2\sqrt{R \cdot H} - R$.

General n . For arbitrary $\alpha \in [0, 1]$, we can lower bound the right hand side with a second order Taylor expansion around $\theta = 1$: i.e. for any $\theta \in [0, 1]$, for some $t \in [\theta, 1]$:

$$\theta^{\alpha} = 1 + (\theta - 1)\alpha + \frac{\alpha(\alpha - 1)}{2} t^{\alpha-2} (\theta - 1)^2 \leq 1 + (\theta - 1)\alpha + \frac{\alpha(\alpha - 1)}{2} (\theta - 1)^2 \quad (82)$$

(since for $\alpha \in [0, 1]$, $t \in [0, 1]$: $t^{\alpha-2} \geq 1$ and $\alpha \leq 1$) Replacing the above approximation in the constraint, gives:

$$\frac{1 - \alpha}{\alpha} \frac{R}{H} \geq \frac{1}{\alpha} - \frac{1}{\alpha} - \theta + 1 - \frac{(\alpha - 1)}{2} (\theta - 1)^2 - 1 + \theta = \frac{1 - \alpha}{2} (\theta - 1)^2 \quad (83)$$

Simplifying further, we want to minimize θ such that: $(\theta - 1)^2 \leq \frac{2}{\alpha} \frac{R}{H}$. The optimal value sets:

$$\theta = 1 - \sqrt{\frac{2}{\alpha} \frac{R}{H}} \quad (84)$$

Leading to a mean of:

$$\mu^* = H(1 - \theta) \leq \sqrt{\frac{2}{\alpha} \cdot R \cdot H} \quad (85)$$

For the second part of the theorem, we simply observe that if the function $v(\cdot)$ is L -Lipschitz, then it must be that $v(1) \leq (1 - \alpha)R + L$, thereby, we can do the exact same analysis as above, but with $H = (1 - \alpha)R + L$. The upper bound on $v(1)$, comes from the fact, that if $v(\cdot)$ is lipschitz, then $v(y) \geq v(1) - L(1 - y) \geq v(1) - L$. Thus we have:

$$\frac{1 - \alpha}{\alpha} R \geq \int_0^1 v(y) (y^{\alpha-1} - 1) dy \geq (v(1) - L) \int_0^1 (y^{\alpha-1} - 1) dy = \frac{v(1) - L}{\alpha} \quad (86)$$

Re-arranging gives the property and concludes the proof of the theorem. ■

A.2 Upper bound on mean of common value via first order conditions

We show an upper bound on the mean of the common value from the observed bid distribution, in the case where bids and their distributions are continuous and admit a density.

Consider a single player i . We will look at the auction from the perspective of that player and so we will be dropping the index i . Consider one bid b for player i that is in the support of the BCE distribution. It must be that this player is maximizing his utility by submitting b rather than any other bid, when he is recommended to bid b . Let B denote the random variable that corresponds to the maximum bid of all players and let B_i be the random variable corresponding to the recommended bid for player i . Moreover, let $G_B(\cdot|b)$ denote the conditional CDF of the maximum bid, conditional on player i being recommended a bid b , and $g_B(\cdot|b)$ be the PDF of this distribution. Note that even if the bid distribution of the bidders is continuous, both G_B and g_b are in general discontinuous at b : for example, $g_B(b'|b) = 0$ for $b' < b$ but $g_B(b|b) = P(B_{-i} \leq b|B = b)$ (where B_{-i} is the maximum bid of the other players), which can be non-zero. However for $b' > b$, we have that

$$G_B(b'|b) = P[B_{-i} \leq b'|B_i = b] = G_{B_{-i}}(b'|b) \quad (87)$$

$$g_B(b'|b) = g_{B_{-i}}(b'|b) \quad (88)$$

where $g_{B_{-i}}(\cdot|b)$ is the density of B_{-i} conditional on $B_i = b$ and $G_{B_{-i}}(b'|b)$ the corresponding cdf. This in particular implies that $G_B(b'|b)$ and $g_B(b'|b)$ are continuous for $b' > b$, and that

$$\lim_{b' \rightarrow b^+} G_B(b'|b) = G_{B_{-i}}(b'|b) \quad (89)$$

$$\lim_{b' \rightarrow b^+} g_B(b'|b) = g_{B_{-i}}(b|b) \quad (90)$$

as G_B and g_B are continuous for all b' .

Let's consider the expected utility of a player from deviating to a bid b' when being recommended a bid b such that $b' > b$ (this does not apply for $b' < b$):

$$U(b'|b) = (\mathbb{E}[v|B \leq b', B_i = b] - b') G_B(b|b) \quad (91)$$

We get that because b maximizes $U(b'|b)$ for $b' > b$, it must be the case that

$$\lim_{b' \rightarrow b^+} \frac{\partial U(b'|b)}{\partial b'} \leq 0 \quad (92)$$

First we consider the derivative of the quantity $E[v|B \leq b', B_i = b]G_B(b'|b)$ for any $b' > b$:

$$E[v|B \leq b', B_i = b]G_B(b'|b) = E[v \cdot 1\{B \leq b'\} | B_i = b] = \int_0^{b'} \left(\int_0^\infty v \cdot f(v|z, b) dv \right) g_B(z|b) dz$$

where $f(v|z, b)$ is the density of the conditional distribution of the common value, conditional on $B = z$ and $B_i = b$. Thus the derivative of the latter with respect to b' is simply:

$$\frac{\partial E[v|B \leq b', B_i = b]G_B(b'|b)}{\partial b'} = \left(\int_0^\infty v \cdot f(v|b', b) dv \right) g_B(b'|b) = E[v|B = b', B_i = b]g_B(b'|b)$$

Combining with the derivative of the term $b'G_B(b'|b)$ we get:

$$\lim_{b' \rightarrow b^+} (E[v|B = b', B_i = b]g_B(b'|b) - b'g_B(b'|b) - G_B(b'|b)) \leq 0 \quad (93)$$

Re-arranging yields:

$$\lim_{b' \rightarrow b^+} E[v|B = b', B_i = b] \leq \lim_{b' \rightarrow b^+} b' + \frac{\lim_{b' \rightarrow b^+} G_B(b'|b)}{\lim_{b' \rightarrow b^+} g_B(b'|b)} \quad (94)$$

Using equations (87), (88), (89) and (90), we obtain:

$$\lim_{b' \rightarrow b^+} E[v|B = b', B_i = b] \leq b + \frac{G_{B_{-i}}(b|b)}{g_{B_{-i}}(b|b)} \quad (95)$$

Note that assuming continuity of $f(v|b', b)$ (i.e. the joint distribution of realized value and bids is continuous), we have continuity of $E[v|B = b', B_i = b] = E[v|B_{-i} \leq b', B_i = b]$ for all b' and we can write:

$$E[v|B = b, B_i = b] \leq b + \frac{G_{B_{-i}}(b|b)}{g_{B_{-i}}(b|b)} \quad (96)$$

Assuming any two players never play the same bid at the same time with non-zero probability (this happens as long as $g_{B_{-i}}(\cdot|b)$ is not degenerate for all bids b and all players i), we have that for all b , the events $(B = b, B_i = b)$ are disjoint accross possible values of i and b (only one can happen at a time). Additionally, there always exist a player i such that $B = b, B_i = b$ for some bid b , hence at least one of the events must happen. So exactly one of $(B = b, B_i = b)$ happens at a time for all possible bids b and bidders i , and we can condition on said events:

$$E[v] = \int_b \sum_i E[v|B = b, B_i = b] P[B = b, B_i = b] db \quad (97)$$

Plugging in Equation (96), we get:

$$\begin{aligned} E[v] &\leq \sum_i \int_b \left(b + \frac{G_{B_{-i}}(b|b)}{g_{B_{-i}}(b|b)} \right) \cdot G_B^i(b|b) g_{B_i}(b) db \\ &= R + \sum_i \int_b \frac{G_B^i(b|b)^2}{g_{B_{-i}}(b|b)} \cdot g_{B_i}(b) db \end{aligned}$$