DISCUSSION PAPER SERIES

DP11782

FIRST PRICE AUCTIONS WITH GENERAL INFORMATION STRUCTURES: IMPLICATIONS FOR BIDDING AND REVENUE

Dirk Bergemann, Benjamin A Brooks and Stephen Morris

INDUSTRIAL ORGANIZATION

ISSN 0265-8003

FIRST PRICE AUCTIONS WITH GENERAL INFORMATION STRUCTURES: IMPLICATIONS FOR BIDDING AND REVENUE

Dirk Bergemann, Benjamin A Brooks and Stephen Morris

Discussion Paper DP11782
Published 14 January 2017
Submitted 14 January 2017

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

FIRST PRICE AUCTIONS WITH GENERAL INFORMATION STRUCTURES: IMPLICATIONS FOR BIDDING AND REVENUE

Abstract

We explore the impact of private information in sealed-bid first-price auctions. For a given symmetric and arbitrarily correlated prior distribution over values, we characterize the lowest winning-bid distribution that can arise across all information structures and equilibria. The information and equilibrium attaining this minimum leave bidders indifferent between their equilibrium bids and all higher bids. Our results provide lower bounds for bids and revenue with asymmetric distributions over values. We also report further characterizations of revenue and bidder surplus including upper bounds on revenue. Our work has implications for the identification of value distributions from data on winning bids and for the informationally robust comparison of alternative bidding mechanisms.

JEL Classification: C72, D44, D82, D83

Keywords: First-price auction, information structure, Bayes correlated equilibrium, private values, interdependent values, common values, revenue, surplus, welfare bounds, reserve price.

Dirk Bergemann - dirk.bergemann@yale.edu
Yale University and CEPR

Benjamin A Brooks - babrooks@uchicago.edu
University of Chicago

Stephen Morris - smorris@princeton.edu
Princeton University

Acknowledgements
We acknowledge financial support through NSF Grant ICES 1215808 and SES 1459899. We thank the Editor, Joel Sobel, and five anonymous referees for helpful comments and suggestions. We also thank Ben Golub, Jason Hartline, Hanming Fang, Phil Reny, Vasilis Syrgkanis, Satoru Takahashi and Eva Tardos for productive conversations. Aron Tobias, Elia Sartori, and Kai Hao Yang provided valuable research assistance. We are also grateful to seminar participants at Arizona, Berkeley, Duke, Harvard, Hebrew University, MIT, Northwestern, Queen Mary, Stanford, and Tel Aviv University.

First-Price Auctions with General Information Structures: Implications for Bidding and Revenue*

Dirk Bergemann Benjamin Brooks Stephen Morris

October 13, 2016

Abstract

We explore the impact of private information in sealed-bid first-price auctions. For a given symmetric and arbitrarily correlated prior distribution over values, we characterize the lowest winning-bid distribution that can arise across all information structures and equilibria. The information and equilibrium attaining this minimum leave bidders indifferent between their equilibrium bids and all higher bids. Our results provide lower bounds for bids and revenue with asymmetric distributions over values. We also report further characterizations of revenue and bidder surplus including upper bounds on revenue. Our work has implications for the identification of value distributions from data on winning bids and for the informationally robust comparison of alternative auction mechanisms.

Keywords: First-price auction, information structure, Bayes correlated equilibrium, private values, interdependent values, common values, revenue, surplus, welfare bounds, reserve price.

JEL Classification: C72, D44, D82, D83.

*Bergemann: Department of Economics, Yale University, dirk.bergemann@yale.edu; Brooks: Department of Economics, University of Chicago, babrooks@uchicago.edu; Morris: Department of Economics, Princeton University, smorris@princeton.edu. We acknowledge financial support through NSF Grant ICES 1215808 and SES 1459899. We thank the anonymous referees for helpful comments and suggestions. We also thank Ben Golub, Jason Hartline, Hanming Fang, Phil Reny, Vasilis Syrgkanis, Satoru Takahashi and Eva Tardos for productive conversations. Aron Tobias, Elia Sartori, and Kai Hao Yang provided valuable research assistance. We are also grateful to seminar participants at Arizona, Berkeley, Chicago Booth, Duke, Harvard, Hebrew University, Johns Hopkins, LSE, MIT, Northwestern, École Polytechnique, Queen Mary, Stanford, and Tel Aviv University.

1 Introduction

The first-price auction has been the subject of extensive theoretical study for over fifty years. It is still fair to say, however, that its properties are well understood only in relatively special cases. Under complete information, the first-price auction reduces to classic Bertrand competition. Under incomplete information, most work focuses on the case of one-dimensional type spaces. For example, when bidders know their private values, it is typically assumed that each bidder *only* knows their private value and has no additional source of information. Beyond the private-values case, it is typically assumed that bidders have one-dimensional types that are jointly affiliated and that values are increasing in the profile of types (Milgrom and Weber, 1982). Thus, a strong relationship is assumed between each bidder's belief about his own value and his beliefs about others' information. But in first-price auctions, unlike in second-price auctions, bidders' beliefs about others' information are of central strategic importance, since what others know is informative about what they will bid in equilibrium. In many situations, it is unnatural to impose strong restrictions on the relationship between the conceptually distinct beliefs about one's own value and about others' information.

In this paper, we derive results about equilibrium behavior in the first-price auction that hold across *all* common-prior information structures. For a given prior distribution over value profiles, we study what can happen for all information structures specifying bidders' information about their own and others' values. Our setting thus incorporates all existing models of information. For any value distribution, we identify a lower bound on the distribution of winning bids in the sense of first-order stochastic dominance. In other words, no matter what the true information structure is, the equilibrium distribution of winning bids must first-order stochastically dominate the bound that we describe. In addition, when the prior distribution of values is symmetric, we construct an equilibrium and an information structure in which this lower bound is attained. This minimum winning-bid distribution therefore pins down the minimum amount of revenue that can be generated by the auction in expectation. Moreover, the minimum winning-bid distribution is attained in an efficient equilibrium. As a result, this equilibrium also attains an upper bound on the expected surplus of the bidders, which is equal to the maximum feasible surplus minus minimum revenue.[1]

Let us give a brief intuition for how our bounds are obtained. If the distribution of winning bids places too high of a probability on low bids, then some bidder would find that a modest increase in their bid would result in a relatively large increase in the probability

[1] Where no confusion results, we will write "revenue" for ex-ante expected revenue, "bidder surplus" for ex-ante expected surplus, etc.

of winning, so that such a deviation would be attractive. For example, it cannot be that all bidders tie with a bid of zero with probability one, for then some bidder could increase his bid a token positive amount and win the auction outright. This suggests that the relevant constraints for pinning down minimum bidding are those associated with deviating to higher bids. Indeed, we show that the minimum winning-bid distribution is characterized by bidders being indifferent to *all* upward deviations.

To characterize the minimum, it turns out to be sufficient to look at a relatively small class of such deviations: For some bid b, we say that a bidder *uniformly deviates up to* b if he switches to bidding b whenever he would have bid less than b in equilibrium. It is clearly necessary for equilibrium that the bidders should not want to uniformly deviate upward. Moreover, it turns out that the change in a bidder's surplus from a uniform upward deviation depends only on the distribution of winning bids, and not on the distribution of losing bids. This motivates a relaxed program in which we minimize the distribution of winning bids, subject only to the uniform upward incentive constraints. The solution to this relaxed program gives us a lower bound on the winning-bid distribution. With the additional assumption of symmetry of the prior, we are able to construct an information structure and equilibrium in which this lower bound is attained, thus verifying that it is indeed the minimum. We will further motivate and illustrate this proof strategy in Section 3 with an example in which there are two bidders and a uniformly distributed common value.

We report a number of further results about bidding, revenue, and bidder surplus. A straightforward upper bound on revenue, albeit a rather weak one, is the efficient surplus. We show that this bound is in general tight, and the bound is robust to standard equilibrium refinements, e.g., ruling out weakly dominated strategies. We also explore the whole welfare space of possible revenue and bidder-surplus outcomes, including those associated with inefficient equilibria. In the case of two bidders and independent values, we construct a maximally inefficient information structure and equilibrium strategy profile. In the resulting equilibrium, the bidders receive zero surplus, and the seller's revenue is the expectation of the lower of the two values.

Our primary focus in this paper is on developing insights about how general information structures can affect outcomes in the first-price auction and on the qualitative properties of the information structure that lead to different outcomes. The results we obtain can be used for a variety of applications, e.g., to partially identify the value distribution in settings where the information structure is unknown and to make informationally robust comparisons of selling mechanisms. We will discuss such applications in the concluding section.

Our work relates to a large literature on first-price auctions. The theoretical study of the first-price auction, going back to Vickrey (1961), has focused almost exclusively on the case where each bidder's information is one-dimensional. A distinctive feature of our analysis is that we allow the bidders to have multidimensional information about values. Fang and Morris (2006) and Azacis and Vida (2015) analyze a model with two bidders and two possible values, where each bidder knows his own value and observes a signal of the other bidder's value.[2] Another distinctive feature of our analysis is that we characterize bidding behavior in all equilibria for all information structures at once. Bergemann and Morris (2013, 2016) show that the range of such behavior can be described using a certain incomplete-information correlated equilibrium that they term *Bayes correlated equilibrium*. Although we do not explicitly use that solution concept in our analysis, a contribution of this paper is to characterize the Bayes correlated equilibria of the first-price auction. Others have also studied the first-price auction under solution concepts that are more permissive than Bayes Nash equilibrium. Motivated by collusion, Lopomo, Marx, and Sun (2011) study communication equilibria of the first-price auction when values are known and independent; they thus impose truth-telling constraints that do not arise in our setting. More recently, Feldman, Lucier, and Nisan (2016) analyze the correlated and coarse correlated equilibria in the first-price auction under complete information. Battigalli and Siniscalchi (2003) and Dekel and Wolinsky (2003) study rationalizable behavior in the first-price auction with a fixed affiliated information structure. Our analysis is more permissive than theirs, in that we consider all information structures, and it is also more restrictive, in that we use equilibrium rather than rationalizability as the solution concept.

The rest of this paper proceeds as follows. Section 2 presents our model of a first-price auction. Section 3 previews our results with a two-bidder example with a uniformly-distributed pure common value. Section 4 contains our main result, a characterization of minimum winning bids, minimum revenue, and maximum bidder surplus. Section 5 describes further results on welfare outlined above, and Section 6 concludes with a discussion of applications to identification and the robust comparison of selling mechanisms. Omitted proofs are contained in the Appendix.

2 Model

We consider the sale of a single unit of a good by a first-price auction. There are N individuals who bid for the good, indexed by $i \in \mathcal{N} = \{1, \ldots, N\}$, each of whom has a value which lies in the compact interval $V = [\underline{v}, \overline{v}] \subset \mathbb{R}_+ = [0, \infty)$. The bidders are assumed to be risk

[2] In an earlier version of this paper (Bergemann, Brooks, and Morris, 2013), we provide a complete analytic characterization of possible welfare outcomes in the setting of Fang and Morris (2006) and Azacis and Vida (2015).

neutral and to have quasilinear preferences over the allocation and payments. Values are jointly distributed according to a probability measure $\mu \in \Delta\left(V^N\right)$.[3]

We allow the common prior μ to be a general measure, not necessarily absolutely continuous, in order to encompass the pure-common-value case in the same analytic framework as non-common values. We will state our results for symmetric (i.e., exchangeable) common priors μ and we will give a formal definition of symmetry in Section 4 when the hypothesis is used (cf. the discussion before Lemma 3). We discuss what happens with asymmetric value distributions after stating the results for symmetric distributions. Some results extend as stated to asymmetric value distributions, most notably Proposition 3, and for the others, we will discuss weaker analogue results.

As we shall see, the distribution of the average of the $N-1$ lowest values will be essential to our analysis. For analytical convenience, we will assume that this distribution is non-atomic, so that its associated cumulative distribution function is continuous. Specifically, we assume that for each $x \in \mathbb{R}$, the event in which the $N-1$ lowest values sum to x has zero probability. In the pure-common-value case, this simply means that the distribution of the common value is non-atomic. This assumption can be dropped without any great conceptual difficulty, and the working paper Bergemann, Brooks, and Morris (2015) solves the present model when values are discrete.

Each individual $i \in \mathcal{N}$ submits a bid $b_i \in B = [0, \overline{v}]$, and the winner is selected uniformly from among the high bidders. For a profile of bids $b \in B^N$, let $W(b)$ be the set of high bidders,

$$W(b) = \{i \,|b_i \geq b_j, \, \forall j = 1, \dots, N\}.$$

Letting $\mathbb{I}_E$ denote the indicator function for an event E, the probability that bidder i receives the good if bids are $b \in B^N$ is[4]

$$q_i(b) = \frac{\mathbb{I}_{i \in W(b)}}{|W(b)|}.$$

Bidders may receive additional information about the profile of values, beyond knowing the prior distribution. This information comes in the form of signals that are correlated with the profile of values. An *information structure* is a collection $\mathcal{S} = \left(\{S_i\}_{i=1}^N, \pi\right)$, where the S_i are measurable spaces[5] and $\pi : V^N \to \Delta(S)$ is a measurable mapping from profiles of

[3]All sets considered in this paper are regarded as topological spaces with their standard topologies, wherever applicable, and endowed with the Borel σ-algebra. For a topological space X, $\Delta(X)$ denotes the set of Borel probability measures on X, endowed with the weak-* topology. For a measure $\mu \in \Delta(X)$ and a measurable function $f : X \to \mathbb{R}$, we denote the integral of f with respect to μ by $\int_{x \in X} f(x)\,\mu(dx)$.

[4]While we assume that ties are broken uniformly, none of our results depend on the particular choice of tie-breaking rule. See also the discussion following Lemma 1.

[5]Note that we do not impose any topological structure on the signal spaces.

values to probability measures over $S = \times_{i=1}^N S_i$. The interpretation is that S_i is the set of bidder i's signals and π describes the conditional joint distribution of signals given values.

For a fixed information structure $\mathcal{S}$, the first-price auction is a game of incomplete information, in which bidders' strategies are measurable mappings $\sigma_i : S_i \to \Delta(B)$ from signals to probability measures over bids. Let Σ_i denote the set of strategies for agent i. Fixing a profile of strategies $\sigma \in \Sigma = \times_{i=1}^N \Sigma_i$, bidder i's (ex-ante) surplus from the auction is

$$U_i(\mathcal{S}, \sigma) = \int_{v \in V^N} \int_{s \in S} \int_{b \in B^N} (v_i - b_i)\, q_i(b)\, \sigma(db|s)\, \pi(ds|v)\, \mu(dv),$$

where $\sigma : S \to \Delta\left(B^N\right)$ maps a profile of signals $s \in S$ into the product measure $\sigma_1(s_1) \times \cdots \times \sigma_N(s_N)$. The profile $\sigma \in \Sigma$ is a *Bayes Nash equilibrium*, or *equilibrium* for short, if and only if $U_i(\mathcal{S}, \sigma) \geq U_i(\mathcal{S}, \sigma'_i, \sigma_{-i})$ for all i and all $\sigma'_i \in \Sigma_i$. In the event that σ_i is a pure strategy, we will abuse notation slightly by writing $\sigma_i(s_i)$ for the bid that is made with probability one.

We denote the total surplus by

$$T(\mathcal{S}, \sigma) = \int_{v \in V^N} \int_{s \in S} \int_{b \in B^N} \sum\nolimits_i v_i q_i(b)\, \sigma(db|s)\, \pi(ds|v)\, \mu(dv).$$

This quantity is bounded above by the efficient total surplus:

$$\overline{T} = \int_{v \in V^N} \max\{v\}\, \mu(dv).$$

A strategy profile is *efficient* if $T(\mathcal{S}, \sigma) = \overline{T}$.

Throughout our analysis, we restrict attention to strategies in which each bidder never bids an amount which is sure to be strictly greater than his value. Formally, for any measurable set $X \subseteq B$, if in equilibrium $\Pr(b_i \in X) > 0$, then $\Pr(b_i \in X \text{ and } b_i \leq v_i) > 0$ as well. This requirement is similar to but slightly weaker than ruling out weakly dominated strategies, since we allow bidders to bid exactly their value with positive probability. Indeed, under complete information, the unique equilibrium involves winners bidding their values whenever there is a tie for highest value.

3 A Pure-Common-Value Example

Before giving our general results, we will illustrate where we are headed with a simple example. There are two bidders who share a common value for the good, which is uniformly distributed between 0 and 1. Since we assume there is no reservation price in the auction, the good is always allocated, regardless of the particular information structure and equilibrium, and as both bidders have the same value, all equilibria are socially efficient and result in a total surplus of 1/2. There may, however, be variation across information structures and equilibria in how this surplus is split between the bidders and the seller.

We allow bidders to observe arbitrary and possibly correlated signals about the common value. At one extreme the bidders' signals might be perfectly correlated, so that they have exactly the same information about the value. For example, the bidders might have no information beyond the prior, so that they both expect the good to be worth 1/2, or the bidders might both observe the true value of the good, so that they know the good's value exactly. In any such case, the bidders will compete the price up to the interim expected value of the good, which results in zero bidder surplus and expected revenue of 1/2. These examples illustrate our later general result that, unless we make additional assumptions about what the bidders know, a tight upper bound on revenue is the efficient surplus.

When the bidders have private information, the distribution of ex-ante surplus can be rather different. An important case has been studied by Engelbrecht-Wiggans, Milgrom, and Weber (1983): bidder 1 observes the true value while bidder 2 is uninformed and observes nothing. In the case of the uniform distribution, the resulting equilibrium involves the informed bidder 1 bidding $v/2$ and the uninformed bidder 2 randomizing uniformly over the interval $[0, 1/2]$. Let us briefly verify that this is an equilibrium. Any bid greater than 1/2 would obviously result in negative surplus for bidder 2. If bidder 2 bids $b \in [0, 1/2]$, she will win whenever bidder 1 bids less than b, which is when the true value is less than $2b$. The conditional distribution of v in the event that bidder 2 wins is therefore uniform on $[0, 2b]$, so that the expected value of v conditional on bidder 2 winning is exactly b. Thus, any bid of bidder 2 in $[0, 1/2]$ results in exactly zero surplus in expectation. On the other hand, conditional on the true value being v, bidder 1 wins with a bid of $b \in [0, 1/2]$ with probability $2b$, resulting in a surplus of $(v - b)\,2b$, which is maximized at $b = v/2$. This equilibrium results in a surplus of 1/6 for bidder 1, a surplus of 0 for bidder 2, and revenue of 1/3.

What can we say about the outcome of the auction more generally? A simple lower bound on revenue is zero, but this bound cannot be tight: revenue could be zero only if both bidders bid zero with probability one, in which case they must be tying, and each bidder obtains a surplus of 1/4. However, either bidder could deviate to bidding $\epsilon > 0$ all the time and obtain a surplus of $1/2 - \epsilon$. This suggests a general intuition that there cannot be too many bids too close to zero, lest the probability of winning increase too quickly as bidders deviate to higher bids. More generally, we might expect that the limit of how low bidding can go will be characterized by binding upward incentive constraints, i.e., bidders being indifferent to deviating to higher bids. Note that in the analysis of Engelbrecht-Wiggans et al., the informed bidder strictly prefers his equilibrium bid over any other bid. This suggests that it might be possible to construct other information structures in which revenue is even lower.

Here is one such construction: The two bidders receive signals $s \in [0, 1]$ that are independent draws from the cumulative distribution $F(s) = \sqrt{s}$, so that the distribution of the maximum signal is standard uniform, the same as the common value. Moreover, the signals and the value are correlated so that the maximum signal is exactly equal to the value:

$$v = \max\{s_1, s_2\}. \tag{1}$$

This information structure admits an equilibrium in which the bidders use the following monotonic pure strategy:

$$\sigma(s) = \frac{1}{\sqrt{s}} \int_{x=0}^{s} x \frac{dx}{2\sqrt{x}} = \frac{s}{3}.$$

Thus, the equilibrium bid given the signal s is the expectation of other bidder's signal, conditional on it being below s. We will presently verify that these strategies constitute an equilibrium, but let us first note the implied welfare outcomes: the winning bid will always be $\max_i s_i/3 = v/3$, so that revenue is $1/6$. Bidder surplus is therefore $1/3$, which is twice as much as the bidders obtained in the proprietary information model of Engelbrecht-Wiggans et al.[6]

Let us now verify that these strategies comprise an equilibrium. It is well known that these strategies are an equilibrium of a different but related model, in which the bidders receive the same signals drawn from the same distribution, but in which each bidder's signal is their private value. In other words, when there are independent private values (IPV), the equilibrium bid is the expectation of the other bidder's value (i.e., signal) conditional on it being less than one's own signal (Krishna, 2002).[7] Now, in our common-value model, a bidder with signal s who deviates by bidding $s'/3$ for some $s' < s$ will only win when their own signal was the highest signal, and therefore equal to the common value. Thus, a downward deviator's surplus looks exactly the same as in the as-if IPV setting, and we can immediately conclude that bidders do not want to deviate down. On the other hand, a bidder who deviates up to $s'/3$ with $s' > s$ continues to win on the event that s was the high

[6] In the online appendix we compute the ratio of the revenue in the equilibrium under the information structure (1) and under the proprietary information structure of Engelbrecht-Wiggans et al. (1983) for the class of power distribution functions, $P(v) = v^{\alpha}$. We show that this ratio increases from 1 to 4 as α increases from 0 to ∞, and thus the uniform case is the special result for $\alpha = 1$ with a revenue ratio of 2.

[7] Indeed, this is a necessary consequence of the revenue equivalence between first- and second-price auctions in the IPV setting (Myerson, 1981).

signal, and now wins on some events when it was the other bidder who had the high signal, which was the true value. The deviator's surplus is

$$\left(s-\frac{s'}{3}\right)\sqrt{s}+\int_{x=s}^{s'}\left(x-\frac{s'}{3}\right)\frac{1}{2\sqrt{x}}dx=\frac{2}{3}s\sqrt{s}$$

which is independent of s'. In other words, bidders are exactly indifferent to *all* upward deviations!

In fact, no matter how one structures the information or the equilibrium strategies, it is impossible for revenue to fall below the level attained in this example, i.e., 1/6 is a tight lower bound on revenue when there are two bidders and there is a pure common value that is standard uniform. Moreover, not only is it impossible for revenue to fall below the level of the example, but the distribution of winning bids in any equilibrium under any information structure must first-order stochastically dominate the winning-bid distribution in the equilibrium we just constructed.

The full proof of this result will be established in Theorem 1 below. To develop intuition, we will give a partial proof that revenue cannot fall below 1/6. First, notice that the equilibria in the no-information, complete-information, and independent-signal constructions all have the feature that the winning bid is a deterministic and weakly increasing function of the true value v. We denote the winning-bid function by $\beta : [0,1] \to [0,1]$ and emphasize that the winning bid $\beta(v)$ is an equilibrium outcome (as is the losing bid) and distinct from the bidding strategies of the individual bidders $\sigma_i(s)$. In the no-information case, the winning bid is always $\beta(v) = 1/2$; under complete information, the winning bid is $\beta(v) = v$; and in the independent-signal construction, $\beta(v) = v/3$. Let us explore more generally what can happen in symmetric equilibria of this form. In other words, we maintain the endogenous assumption that whatever the information structure and equilibrium strategies are, they induce an outcome in which (i) both players are equally likely to win at any given value v and (ii) the bidder who wins pays an amount that is a deterministic and increasing function $\beta(v)$ of the true value. For simplicity, we will also assume that ties occur with zero probability, and that β is strictly increasing. Notice that we are suppressing a great deal of information about the underlying information structure and equilibrium strategies that induce the winning bid. Nonetheless, the winning-bid function β provides sufficient information to calculate revenue:[8]

[8]To be clear, we are *not* making any particular assumption about the dimensionality of the signal space, nor are we assuming that bidders use pure strategies. The hypothesis that the winning bid only depends on the true value only requires that whichever bidder wins the auction is not randomizing given his information, but the bidders could be randomizing when they lose.

$$R = \int_{v=0}^{1} \beta(v)\, dv$$

and to calculate each bidder's expected surplus:

$$U_i = \frac{1}{2} \int_{v=0}^{1} (v - \beta(v))\, dv. \tag{2}$$

An equilibrium in which the winning bid is $\beta(v)$ must deter a large number of deviations, most of which we cannot assess without explicitly modeling the rest of the equilibrium. There is, however, one class of deviations which we can evaluate using only information about winning bids: for some $w \in [0, 1]$, bid $\beta(w)$ whenever the equilibrium bid would have been some $b \leq \beta(w)$. We refer to this as a *uniform deviation up to* $\beta(w)$.[9]

We can calculate the surplus of a bidder who uniformly deviates up to $\beta(w)$ as follows. First consider the event in which the true value is less than w, so that the equilibrium winning bid would have been some number less than $\beta(w)$. In this case, the upward deviator *always wins*: the deviator always bids $\beta(w)$, which is more than what either the winner or the loser would have bid in equilibrium. Now consider the event in which the true value v is strictly greater than w. In this case, if the deviator were going to win in equilibrium, then he would have bid $\beta(v) > \beta(w)$, so that behavior is unaffected by the deviation and he still wins. On the other hand, if the deviator were going to lose, then the other player would have bid $\beta(v) > \beta(w)$ and the deviator would have bid $b < \beta(v)$ in equilibrium, so that the deviator still loses even after the upward deviation. Thus, the upward deviator's surplus is

$$\int_{v=0}^{w} (v - \beta(w))\, dv + \frac{1}{2} \int_{v=w}^{1} (v - \beta(v))\, dv,$$

and the equilibrium payoff given by (2) deters uniform upward deviations only if

$$\frac{1}{2} \int_{v=0}^{w} (v - \beta(w))\, dv \leq \frac{1}{2} \int_{v=0}^{w} (\beta(w) - \beta(v))\, dv$$

for all w. This condition rearranges to

$$\beta(w) \geq \Lambda(\beta)(w), \tag{3}$$

[9] Feldman, Lucier, and Nisan (2016) also use the uniform upward deviation to show that winning bids can never be below the second-highest value in a correlated equilibrium of a complete-information first-price auction.

where

$$\Lambda(\beta)(w) = \frac{1}{2w}\int_{v=0}^{w}(v+\beta(v))\,dv$$

and $\Lambda(\beta)(0) = \beta(0)/2$. Thus, Λ is an operator that maps non-decreasing measurable functions into non-decreasing measurable functions.

A relaxation of the original problem of minimizing revenue over all information structures and equilibria (of this particular form) is to minimize revenue over all bidding functions that satisfy (3) and the condition that $\beta \geq 0$. The solution to this relaxed program has a simple form. Notice that Λ is monotonic in β. As a result, if β is feasible, then

$$\beta \geq \Lambda(\beta) \geq \Lambda(\Lambda(\beta)),$$

so that $\Lambda(\beta)$ is also feasible and is weakly lower than β. In addition, for any two candidate winning-bid functions β and $\widehat{\beta}$ that are weakly increasing,

$$\left|\Lambda(\beta)(w) - \Lambda\left(\widehat{\beta}\right)(w)\right| = \frac{1}{2w}\int_{v=0}^{w}\left|\beta(v)-\widehat{\beta}(v)\right|dv \leq \frac{1}{2}\left\|\beta-\widehat{\beta}\right\|,$$

where $\|\cdot\|$ denotes the L^1 norm, so that Λ is a contraction with modulus $1/2$. Thus, iteratively applying Λ to any feasible solution must generate a pointwise weakly-decreasing sequence of winning-bid functions that converges to the unique fixed point of Λ. The fixed point must itself be weakly lower than any other feasible solution. It is easily verified that

$$\underline{\beta}(w) = \frac{w}{3}$$

is the fixed point of Λ.[10]

Indeed, not only does this winning-bid function minimize expected revenue, it achieves a pointwise minimum across all feasible solutions to the relaxed program. As a result, it minimizes the distribution of winning bids in the sense of first-order stochastic dominance, within the class of equilibria we considered and subject only to the incentive constraints associated with uniform upward deviations. We will show in Section 4 that this bound continues to hold even if one allows asymmetric information structures, asymmetric equilibria, and equilibria in which the winning bid is stochastic conditional on v. This implies that revenue cannot fall below $1/6$ in any equilibrium under any information structure. Thus, the information structure and equilibrium that we constructed attain a global lower bound on the distribution of winning bids.

[10]This argument is the result of a conversation with Phil Reny, to whom we are extremely grateful.

4 Minimum Bidding, Minimum Revenue, and Maximum Bidder Surplus

4.1 Preview and Statement of Main Result

We shall see that the characterization of minimum winning bids can be generalized to any symmetric prior distribution over values and any number of bidders. For any such prior, there is a minimum winning-bid distribution that is first-order stochastically dominated by any distribution of equilibrium winning bids that can be induced by some information structure $\mathcal{S}$ and equilibrium σ. There also exists an $\mathcal{S}$ and σ which attain the generalized bound. The qualitative features of the solution, and the methods used to characterize it, closely resemble the arguments in the uniform example of the previous section: the relevant incentive constraints that pin down the minimum are those corresponding to uniform upward deviations, and at the lower bound, the winner's bid will turn out to be a deterministic function of the profile of values.

As a segue to stating our main result, we briefly describe how the results exposited in Section 3 generalize to progressively broader classes of models. First, consider the case where there is a pure common value that is drawn from an arbitrary continuous distribution P on the interval $[\underline{v}, \overline{v}]$. In this case, the minimum distribution over winning bids is characterized by a deterministic winning bid that is given by the following function of the value:

$$\underline{\beta}(w) = \frac{1}{\sqrt{P(w)}} \int_{x=\underline{v}}^{w} x \frac{P(dx)}{2\sqrt{P(x)}}$$

for $w > \underline{v}$ and $\underline{\beta}(\underline{v}) = \underline{v}$. Thus, the winning bid when the true value is w is equal to the expectation of a random variable drawn from the cumulative distribution $\sqrt{P(x)/P(w)}$ on the range $[\underline{v}, w]$, which is necessarily strictly increasing on the support of P. This induces a minimum winning-bid distribution $\underline{H}(b)$ defined by:

$$\underline{H}(b) = P\left(\underline{\beta}^{-1}(b)\right).$$

where $\underline{\beta}^{-1}(b) = \sup\left\{w | \underline{\beta}(w) \leq b\right\}$. These formulae reduce to $\underline{\beta}(v) = v/3$ and $\underline{H}(b) = 3b$ when P is standard uniform.

Next, consider the case where there are two bidders whose values are drawn from a joint distribution $\mu(dv_1, dv_2)$. In this case, there continues to be a minimum winning-bid function that pins down $\underline{H}$, but it is now a function of the lower of the two values. Thus, if we write $Q(w) = \mu(\{v | \min\{v_1, v_2\} \leq w\})$ for the cumulative distribution of the lowest value, then the minimum winning-bid function is

$$\underline{\beta}(w) = \frac{1}{\sqrt{Q(w)}} \int_{x=\underline{v}}^{w} x \frac{Q(dx)}{2\sqrt{Q(x)}}, \tag{4}$$

and

$$\underline{H}(b) = Q\left(\underline{\beta}^{-1}(b)\right).$$

The intuition for this structure is as follows. The relevant constraints that pin down $\underline{H}$ continue to be those associated with uniform upward deviations, and when a player uniformly deviates upward, they continue to win whenever they would have won in equilibrium. This means that the change in surplus from such a deviation only depends on the deviator's values when he loses in equilibrium. As a result, the minimum winning bid only depends on the distribution of the value of the player who loses. All else equal, lower losing values result in a weaker incentive to deviate up, so that the distribution of winning bids is minimized at an efficient allocation in which the losing value is the lowest value.

Finally, consider the general case with N bidders whose values are jointly distributed according to the probability measure $\mu(dv_1, \ldots, dv_N)$. $\underline{H}$ continues to be characterized by a winning bid that is a deterministic function of the losing bidders' values. Moreover, for the same reasons as with two bidders, $\underline{H}$ is attained with an efficient allocation. The remaining question is how the winning bid depends on the $N-1$ lowest values that lose the auction in equilibrium.

To develop intuition, let us reason by analogy with the benchmark of complete information, in which all bidders see the entire profile of values. The equilibria in this information structure involve the bidder with the highest value winning the good, and paying a price which is equal to the maximum of the $N-1$ lowest values. Now suppose that instead of knowing the entire profile of values, the bidders only learn (i) whether or not they have the highest value and (ii) the realized $N-1$ lowest values. Importantly, the low-value bidders do not learn which value they have, but only the empirical distribution of values among the $N-1$ lowest. Due to symmetry of μ, low-value bidders believe that they are equally likely to fall anywhere in that distribution, so they expect their value to be the *average* of the $N-1$ lowest:

$$\alpha(v) = \frac{1}{N-1}\left(\sum_{i=1}^{N} v_i - \max v\right). \tag{5}$$

The equilibria in this information structure still involve the high-value bidder winning the auction. However, the high-value bidder now faces less intense competition from the low-value bidders: whereas under complete information the most optimistic losing bidder would compete the price up to the second-highest value, now losing bidders are only willing to compete the price up to the average of the $N-1$ lowest values given by $\alpha(v)$. Revenue will therefore be substantially lower than under complete information.

Based on this intuition, we can guess that the generalized minimum winning bid will depend only on the average of the losing buyers' values. Let Q denote the distribution of $\alpha(v)$:

$$Q(w) = \mu\left(\left\{v \in V^N | \alpha(v) \leq w\right\}\right),$$

and write $\underline{w}$ and $\overline{w}$ for the minimum and maximum of the support of Q, respectively. (Recall that we assumed in Section 2 that the events $\alpha^{-1}(x)$ have zero probability, so that $Q(x)$ is continuous.) Let

$$\underline{\beta}(w) = \frac{1}{Q^{\frac{N-1}{N}}(w)} \int_{x=\underline{w}}^{w} x \frac{N-1}{N} \frac{Q(dx)}{Q^{\frac{1}{N}}(x)}, \tag{6}$$

and let

$$\underline{H}(b) = Q\left(\underline{\beta}^{-1}(b)\right). \tag{7}$$

For a given information structure $\mathcal{S}$ and equilibrium σ, the winning-bid distribution is defined by:

$$H(b; \mathcal{S}, \sigma) = \int_{v \in V^N} \int_{s \in S} \sigma\left([0,b]^N \middle| s\right) \pi(ds|v)\, \mu(dv), \tag{8}$$

where $\sigma\left([0,b]^N \middle| s\right)$ is the conditional probability that all bids are less than $b \in B$ given signal profile s. Our main result is the following:

Theorem 1 (Minimum Winning Bids)**.**

(i) For any information structure $\mathcal{S}$ and equilibrium σ, the distribution of winning bids $H(\mathcal{S}, \sigma)$ first-order stochastically dominates $\underline{H}$, i.e., $H(b; \mathcal{S}, \sigma) \leq \underline{H}(b)$ for all b;

(ii) There exists an information structure $\mathcal{S}$ and an efficient equilibrium σ such that the distribution of winning bids $H(\mathcal{S}, \sigma)$ is equal to $\underline{H}$.

Immediate implications of Theorem 1 are characterizations of minimum revenue and maximum bidder surplus. Let $\underline{R}$ be defined by

$$\underline{R} = \int_{x=\underline{w}}^{\overline{w}} \underline{\beta}(x)\; Q(dx),$$

and similarly define the revenue from a given information structure $\mathcal{S}$ and equilibrium σ as

$$R(\mathcal{S},\sigma)=\int_{b\in B} b\,H(db;\mathcal{S},\sigma).$$

Corollary 1 (Minimum Revenue).
Any equilibrium σ under any information structure $\mathcal{S}$ must result in revenue $R(\mathcal{S},\sigma)\geq \underline{R}$. Moreover, there exists an equilibrium σ for some information structure $\mathcal{S}$ in which $R(\mathcal{S},\sigma)=\underline{R}$.

Next, recall that $\overline{T}$ is the total surplus that is generated by the efficient allocation, and let the maximal bidder surplus be:

$$\overline{U}=\overline{T}-\underline{R}.$$

Corollary 2 (Maximum Bidder Surplus).
Any equilibrium σ under any information structure $\mathcal{S}$ must result in bidder surplus $U(\mathcal{S},\sigma)\geq \overline{U}$. Moreover, there exists an equilibrium σ for some information structure $\mathcal{S}$ in which $U(\mathcal{S},\sigma)=\overline{U}$.

The rest of this section will be devoted to the proof of Theorem 1. The proof consists of two main pieces. We will first argue that (7) is a lower bound on the distribution of winning bids that can arise in equilibrium. In particular, we will show that (7) is the solution to a relaxed program in which we minimize the distribution of winning bids subject to only the uniform upward incentive constraints. We will then construct an information structure and an equilibrium that exactly attain the solution to the relaxed program. This information structure and equilibrium turn out to be remarkably simple: the bidders receive one-dimensional signals and use a symmetric monotonic pure strategy which is equal to the minimum winning-bid function.

4.2 The Relaxed Program

For an information structure $\mathcal{S}$ and equilibrium strategies σ, let

$$H_i(b|v;\mathcal{S},\sigma)=\int_{s\in S}\int_{x\in[0,b]^N} q_i(x)\,\sigma(dx|s)\,\pi(ds|v) \tag{9}$$

denote the probability that bidder i wins and the winning bid is less than or equal to b when the profile of values is v. Thus,

$$H(b|v;\mathcal{S},\sigma)=\sum_{i=1}^{N} H_i(b|v;\mathcal{S},\sigma)$$

denotes the total probability that the winning bid is less than or equal to b when the realized profile of values is v. The aggregate distribution of winning bids can be written as

$$H(b; \mathcal{S}, \sigma) = \int_{v \in V^N} H(b|v; \mathcal{S}, \sigma)\, \mu(dv).$$

We will hereafter suppress the dependence of these distributions on a particular $(\mathcal{S}, \sigma)$, and we simply write $H_i(b|v)$, $H(b|v)$, and $H(b)$.

Just as β was sufficient to calculate revenue and bidder surplus in the example of Section 3, so too are the winning-bid distributions $\{H_i(\cdot|\cdot)\}$ sufficient to calculate these welfare outcomes for the general model. Revenue is simply

$$R = \int_{x \in B} x H(dx)$$

and bidder i's surplus is

$$U_i = \int_{v \in V^N} \int_{x \in B} (v_i - x)\, H_i(dx|v)\, \mu(dv).$$

Thus, bounds on the distribution of winning bids, in addition to the correlation structure between winning bids and values, immediately imply bounds on revenue and bidder surplus.

The $\{H_i(\cdot|\cdot)\}$ are a family of marginal distributions that are induced by the equilibrium, and they contain insufficient information to evaluate the merits of all deviations. We can, however, use them to evaluate the previously introduced class of uniform upward deviations: for some cutoff $b \in B$, bid $\max\{x, b\}$, where x is the equilibrium bid. Let us heuristically derive the incentive constraint corresponding to the uniform upward deviation to b, assuming now that $H(\cdot)$ does not have a mass point at b (so that we can ignore ties at this particular bid). Note that a bidder who uniformly deviates up to b will win at a bid of b whenever the equilibrium winning bid would have been a number strictly less than b. When the winning bid would have been strictly greater than b, however, the upward deviator wins if and only if he would have won in equilibrium, and at the bid at which he would have won in equilibrium. For if the uniform upward deviation changes the bid to b, then that bid is still too low to change the outcome of the auction. Otherwise, the deviator's bid is unchanged, as is the outcome of the auction. The surplus from uniformly deviating up to b is therefore

$$\int_{v \in V^N} \left((v_i - b)\, H(b|v) + \int_{x=b}^{\overline{v}} (v_i - x)\, H_i(dx|v) \right) \mu(dv). \tag{10}$$

Thus, the uniform deviation up to b is unattractive if and only if

$$\int_{v \in V^N} (v_i - b) H(b|v) \mu(dv) \leq \int_{v \in V^N} \int_{x=0}^{b} (v_i - x) H_i(dx|v) \mu(dv). \tag{11}$$

We summarize the preceding discussion with the following result, whose rigorous proof appears in the Appendix.

Lemma 1 (Uniform Upward Incentive Constraints). *Any equilibrium σ under any information structure $\mathcal{S}$ must induce winning-bid distributions $\{H_i(\mathcal{S}, \sigma)\}$ that satisfy (11) for all $i \in \mathcal{N}$ and $b \in B$.*

Note that this constraint must still be satisfied even if there is a mass point at b, which represents the possibility that the uniform deviation up to b would induce ties with positive probability. For bidders must not want to uniformly deviate up to any $b' > b$, and since H can have only countably many mass points, the infimum surplus over all such $b' > b$ is precisely (10). In fact, the uniform upward incentive constraints would continue to hold as stated even if we used a different and non-uniform tie breaking rule. In the sequel, we will find it useful to have an alternate form of (11) obtained by rearranging and integrating by parts:

$$\int_{v \in V^N} (v_i - b)(H(b|v) - H_i(b|v)) \mu(dv) \leq \int_{v \in V^N} \int_{x=0}^{b} H_i(x|v) \, dx \, \mu(dv). \tag{12}$$

To characterize minimum bidding, we will consider a family of relaxed programs, indexed by $\hat{b} \in B$, in which the objective is to maximize

$$\int_{v \in V^N} \sum_{i=1}^{N} H_i\left(\widehat{b} \middle| v\right) \mu(dv) \tag{13}$$

over all mappings $H : V^N \times \mathcal{N} \times \mathbb{R}_+ \to [0,1]$, written $H_i(b|v)$, such that: $H_i(b|v)$ is measurable with respect to v for every (i, b); $H_i(b|v)$ is weakly increasing and right-continuous in b for every (i, v); $\sum_{i=1}^{N} H_i(\overline{v}|v) = 1$ for every v; and the incentive constraint (12) is satisfied for all i and b. It will turn out that there exists a feasible solution that simultaneously maximizes (13) for all $\widehat{b}$. This solution is therefore a lower bound on the distribution of winning bids satisfying (12) in the sense of first-order stochastic dominance, and by Lemma 1, it is also a lower bound on the distribution of winning bids that can arise in equilibrium.

The rest of this subsection will be devoted to solving the relaxed program. The following result, whose proof appears in the Appendix, verifies that a solution exists:

Lemma 2 (Existence)**.**
A solution to the relaxed program exists.

Our next four results show that a solution can be found within a relatively small family of candidate optima that are (i) symmetric, (ii) are associated with efficient allocations, (iii) only depend on the average of the losing values, and (iv) correspond to a winning bid that is a deterministic and increasing function of that average. Once these results have been established, we will use essentially the same argument as in Section 3 to characterize the optimal winning-bid function, with the average of the $N-1$ lowest values playing the role of the common value.

Let us first show that it is without loss of generality to look at solutions that are symmetric. Let Ξ denote the set of permutations of the bidders' identities, i.e., bijective mappings from $\mathcal{N}$ into itself. We associate each $\xi \in \Xi$ with a mapping from V^N into itself, where $v' = \xi(v)$ if $v'_{\xi(i)} = v_i$ for all i. We have assumed that the distribution μ is symmetric, by which we formally mean that $\mu(X) = (\mu \circ \xi)(X) = \mu(\xi(X))$ for all measurable sets $X \subseteq V^N$. In addition, we say that a solution $\{H_i(\cdot|v)\}$ is *symmetric* if for all $\xi \in \Xi$, $v \in V^N$, and $b \in B$,

$$H_{\xi(i)}(b|\xi(v)) = H_i(b|v).$$

In other words, the probability of a bidder winning with bid less than b only depends on (i) the bidder's own value and (ii) the distribution of the values of the other bidders, but it does not depend on how others' values are matched to bidders' individual identities.

Lemma 3 (Symmetry)**.**
For any feasible solution to the relaxed program, there exists a symmetric feasible solution with the same aggregate distribution of winning bids.

The idea behind the proof is that if we had a feasible solution that was asymmetric, it is possible to "symmetrize" the solution by creating new winning-bid distributions that are the average of the winning-bid distributions over all permutations of the bidders' identities:

$$\widetilde{H}_i(b|v) = \frac{1}{N!} \sum_{\xi \in \Xi} H_{\xi(i)}(b|\xi(v)).$$

This new solution $\left\{\widetilde{H}_i(\cdot|\cdot)\right\}$ is symmetric, and since the objective and constraints for the relaxed program are all linear in the H_i, the constraints that were previously satisfied will still be satisfied at the symmetrized solution. In light of Lemma 3, we will henceforth assume that the solution we are working with is symmetric.

Next, we say that a solution to the relaxed program is *efficient* if $H_i(b|v) = 0$ for all b whenever $v_i < \max v$. The second property that we can assume without loss of generality is that the solution is efficient.

Lemma 4 (Efficiency)**.**
For any feasible solution to the relaxed program, there exists an efficient feasible solution with the same aggregate distribution of winning bids.

This result is proven by taking a candidate solution $H_i(b|v)$, and constructing a new efficient solution according to

$$\widetilde{H}_i(b|v) = \frac{\mathbb{I}_{i \in W(b)}}{|W(b)|} H(b|v).$$

Since only losing values appear in (12) (on the left-hand side), and since this new allocation has weakly lower losing values on average across all bidders, one can show that this new solution relaxes the critical uniform upward incentive constraints, while maintaining the same aggregate distribution of bids.

The observation that only losing values appear in (12) allows us to collapse the relaxed program into a somewhat more compact form. In particular, the fact that the H_i are symmetric implies that the distribution of winning bids will be the same for all permutations of the losing buyers' values. Thus, a bidder who uniformly deviates up to b will win with probability $H(b|v)$ when the profile of values is v, but he will also win with the same probability when the profile of values is $\xi(v)$ for any permutation $\xi \in \Xi$. In short, this bidder believes that he will win with the same probability for all elements of the equivalence class $[v] = \{\xi(v) | \xi \in \Xi\}$. Symmetry of μ implies that the deviator is equally likely to have any of the values in the profile v. Since the good is always won by a bidder with the highest value, the expected value conditional on losing in equilibrium and conditional on $[v]$ must be $\alpha(v)$ as defined by (5), which is the average of the $N-1$ lowest realized values. This must also be the expected surplus that the bidder gains by winning when he would have lost in equilibrium on the event $[v]$.

Since upward deviators only care about $\alpha(v)$, it is without loss of generality to consider solutions to the relaxed program in which the distribution of winning bids only depends on the average of the $N-1$ lowest values. Specifically, we can restrict attention to solutions for which

$$\alpha(v) = \alpha(v') \implies H(\cdot|v) = H(\cdot|v').$$

For such a solution, we can then define the conditional distribution of winning bids given the average losing value as $H_\alpha : [\underline{w}, \overline{w}] \to \Delta(B)$ by $H_\alpha(\cdot|w) = H(\cdot|v)$ for any $v \in \alpha^{-1}(w)$.

Recalling that Q is the marginal distribution of $\alpha(v)$, (12) can be rewritten as

$$\frac{N-1}{N}\int_{w=\underline{w}}^{\overline{w}}(w-b)\,H_\alpha(b|w)\,Q(dw)\leq\frac{1}{N}\int_{w=\underline{w}}^{\overline{w}}\int_{x=0}^{b}H_\alpha(x|w)\,dx\,Q(dw)\,. \tag{14}$$

The following lemma verifies that this incentive constraint is equivalent to (12).

Lemma 5 (Average Losing Values)**.**
For any feasible solution to the relaxed program, there exists a feasible solution H with the same aggregate distribution of winning bids such that $H(\cdot|v)=H(\cdot|v')$ whenever $\alpha(v)=\alpha(v')$. Moreover, for such solutions, (14) is equivalent to (12).

In light of Lemma (5), we will now treat $H_\alpha(\cdot|\cdot)$ as the choice variable for the relaxed program, although we retain the notation $H(\cdot)$ for the aggregate distribution of winning bids, and we note that

$$H(b)=\int_{w=\underline{w}}^{\overline{w}}H_\alpha(b|w)\,Q(dw)\,.$$

Next, we say that the solution H_α is *monotonic* if $H_\alpha(b|w)<1$ implies that $H_\alpha(b|w')=0$ for all $w'>w$. In other words, the supports of the winning-bid distributions are monotonically increasing in the average losing value. Our next result will show that it is without loss of generality to restrict attention to solutions that are monotonic. The reason is the following. Note that the incentive constraint (14) can be rewritten as

$$\frac{N-1}{N}\int_{w=\underline{w}}^{\overline{w}}w\,H_\alpha(b|w)\,Q(dw)\leq\frac{N-1}{N}b\,H(b)+\frac{1}{N}\int_{x=0}^{b}H(x)\,dx. \tag{15}$$

Let us consider how this constraint is affected by varying the solution $H_\alpha(b|w)$ while maintaining a fixed aggregate distribution of winning bids $H(b)$. The only piece which depends on the correlation between winning bids b and values w is the left-hand side, which is proportional to the expectation of the average losing value conditional on the winning bid being less than b. On the whole, decreasing this expectation relaxes the constraint, as lower losing values result in smaller gains from a uniform upward deviation. Monotonicity essentially says that the lowest losing values should be associated with the lowest winning bids. In fact, this structure minimizes the expectation of $\alpha(v)$ conditional on the winning bid being less than b, pointwise for every b, and thereby relaxes the constraints as much as possible while maintaining the given $H(b)$.

As a result, it is without loss of generality to consider solutions to the relaxed program that correspond to a deterministic winning bid $\beta(w)$ as a function of the average losing value

$w = \alpha(v)$, which is defined by

$$\beta(w) = \min\{b | H(b) \geq Q(w)\}. \tag{16}$$

Since H is right-continuous and Q is continuous, the set of b such that $H(b) \geq Q(w)$ must be closed. The minimum therefore always exists, is non-decreasing in w, and is continuous from the left. However, β will fail to be right-continuous when H has mass points, and it will have flats on regions that the measure Q assigns zero probability. We can therefore rewrite (12) one last time as

$$\frac{N-1}{N}\int_{x=\underline{w}}^{w}(x-\beta(w))\,Q(dx) \leq \frac{1}{N}\int_{x=\underline{w}}^{w}(\beta(w)-\beta(x))\,Q(dx). \tag{17}$$

If we define

$$\beta^{-1}(b) = \max\{w | \beta(w) \leq b\},$$

then maximizing $H\left(\hat{b}\right)$ is equivalent to minimizing $\beta^{-1}\left(\hat{b}\right)$.

Lemma 6 (Monotonicity).
For any feasible solution to the relaxed program, there exists a monotonic feasible solution with the same aggregate distribution of winning bids. Moreover, for such solutions, (17) is equivalent to (12).

It is now apparent that our situation is quite close to what we assumed in the example of Section 3. Our object of choice in the relaxed program is a deterministic winning bid as a function of a one-dimensional statistic. In the case of pure common values, this statistic was the true value of the good, and in the general model it is the average of the $N-1$ lowest values.

Lemma 7 (Minimum winning-bid function).
The lowest winning-bid function that satisfies (17) is given by (6).

The proof of this result is as follows. As in the uniform example of Section 3, we can rewrite (17) as

$$\beta \geq \Lambda(\beta),$$

where the operator Λ maps the set of non-decreasing measurable functions from $[\underline{w}, \overline{w}]$ into B into itself according to

$$\Lambda(\beta)(w) = \frac{1}{Q(w)}\int_{x=\underline{w}}^{w}\left(\frac{N-1}{N}x + \frac{1}{N}\beta(x)\right)Q(dx) \tag{18}$$

for $w > \underline{w}$, and

$$\Lambda(\beta)(\underline{w}) = \frac{N-1}{N}\underline{w} + \frac{1}{N}\beta(\underline{w}).$$

It is easily verified that $\underline{\beta}$ given by (6) satisfies (17) everywhere, and is therefore a fixed point of Λ. Moreover, Λ is a contraction mapping of modulus $1/N$ in the L^1 norm, so that $\underline{\beta}$ is the unique fixed point. Λ is also monotonic in β, so that if β is feasible for (17), then $\Lambda(\beta)$ is feasible as well. Thus, starting from any feasible β^0, the sequence of functions β^k defined by $\beta^k = \Lambda\left(\beta^{k-1}\right)$ is feasible, decreasing, and converges pointwise to the unique fixed point. This proves that (6) is lower than any other feasible winning-bid function.

The solution to the relaxed program is summarized as follows:

Proposition 1 (Solution to the Relaxed Program)**.**
$\underline{H}\left(\widehat{b}\right)$ given by (7) is the solution to the relaxed program for every $\widehat{b}$. Thus, $H(b|\mathcal{S},\sigma) \leq \underline{H}(b)$ for all bids b, information structures $\mathcal{S}$, and equilibria σ.

4.3 A Minimum-Bidding Information Structure and Equilibrium

At this point, we have proven the first part of Theorem 1, which is that the minimum winning-bid distribution $\underline{H}$ must be first-order stochastically dominated by any equilibrium winning-bid distribution. To show that the bounds are tight, we will construct an efficient equilibrium in which the distribution of winning bids is precisely $\underline{H}$. This construction will generalize the independent-signal information structure of Section 3.

In the winning-bid-minimizing information structure, bidders receive signals that are independent draws from the distribution $F(s) = (Q(s))^{1/N}$ on the support $S = [\underline{w}, \overline{w}]$. This distribution is chosen so that the highest signal is distributed according to Q. Indeed, signals will be correlated with values so that:

(i) the highest signal is equal to the realized average losing value, i.e., $\max s = \alpha(v)$;[11]

(ii) the bidder with the highest value receives the highest signal.

We note that we could have alternatively specified the same information structure by first drawing a profile of values v from μ, then giving the highest-value bidder a signal $w = \alpha(v)$

[11] In the pure-common-value model, the winning-bid-minimizing information structure coincides with an environment that Bulow and Klemperer (2002) refer to as "the maximum game." In this case, the average losing value is simply the common value which coincides with the maximum of the independent signals. Bulow and Klemperer show that in the maximum game, the bidder with the highest signal has the lowest marginal revenue (as described by the virtual utility of the bidder), thus hinting at the low revenue properties of this information structure. In Bergemann, Brooks, and Morris (2016), we characterize revenue maximizing auctions for this information structure.

(breaking ties uniformly), and then giving the other bidders signals which are independent draws from the conditional distribution $F(s)/F(w)$.

In equilibrium, a bidder with signal s bids $\underline{\beta}(s)$ with probability one. These strategies are an equilibrium under a different but related model in which bidders' signals are distributed as above, but each bidder's signal is their value. Specifically, in the independent private-value (IPV) model in which the bidders' values are independent draws from F, there is an equilibrium in monotonic pure strategies in which a bidder with value s_i bids the expected highest signal of others, $\max s_{-i}$, conditional on $\max s_{-i}$ being below s_i, which is precisely $\underline{\beta}(s_i)$ (Krishna, 2002).

We can use this connection to prove that $\underline{\beta}$ is an equilibrium under the interdependent values information structure constructed above. Consider a bidder who receives a signal s and bids $\underline{\beta}(w)$ with $w < s$. Such a deviator would only win when the other bidders' signals are less than w, in which case the deviator must have the highest signal and expects his value to be the maximum value conditional on $\alpha(v) = s$, which we can denote by $\tilde{v}(s)$. Note that $\tilde{v}(s)$ must be weakly greater than s. As a result, if we write $G(w) = (F(w))^{N-1} = (Q(w))^{(N-1)/N}$ for the probability that $N-1$ of the independent signals are below w, then the deviator's surplus is

$$\left(\tilde{v}(s) - \underline{\beta}(w)\right) G(w) = (\widetilde{v}(s) - s)\, G(w) + \left(s - \underline{\beta}(w)\right) G(w).$$

The second piece is exactly the surplus under the IPV interpretation, which must be less than $\left(s - \underline{\beta}(s)\right) G(s)$. In addition, $(\widetilde{v}(s) - s)\, G(w) \leq (\widetilde{v}(s) - s)\, G(s)$, so that downward deviations are not attractive.

Now consider an upward deviation to $\underline{\beta}(w)$ with $w > s$. Such a deviator believes that he has the highest signal with probability $G(s)$, in which case his expected value is $\widetilde{v}(s)$. In addition, the deviator will also win when the highest signal of others is in $[s, w]$, and the expected value conditional on winning when $\max s_{-i} = x$ is precisely x, since the highest signal of others must equal the average losing value.[12] The surplus from the upward deviation is therefore

$$\tilde{v}(s)\, G(s) + \int_{x=s}^{w} x G(dx) - \underline{\beta}(w)\, G(w) = \int_{x=\underline{w}}^{s} (\widetilde{v}(s) - x)\, G(dx), \tag{19}$$

[12]Note that for non-common-value models, $\tilde{v}(s)$ is generally strictly greater than s. This means that a bidder's expectation of his own value is non-monotonic in others' signals: when $\max_{j \neq i} s_j < s_i$, bidder i expects his value to be $\tilde{v}(s_i) > s_i$, but when $\max_{j \neq i} s_j > s_i$, bidder i expects his value to be $\max_{j \neq i} s_j$. This non-monotonicity puts our information structure outside the affiliated-values model of Milgrom and Weber (1982), though this violation is not so severe as to preclude the existence of an equilibrium in monotonic pure strategies.

due to the fact that

$$\underline{\beta}(w) = \frac{1}{G(w)} \int_{x=\underline{w}}^{w} x\, G(dx).$$

Thus, (19) is independent of w, and we conclude that bidders are indifferent to upward deviations. We have proved the following:

Proposition 2 (Minimum Bidding with Symmetric Distributions)**.**
There exists an information structure and efficient equilibrium in which the winning bid is a deterministic function of the average losing value and is given by $\underline{\beta}(w)$.

Proposition 2 completes the proof of Theorem 1. We have constructed an information structure and equilibrium in which the distribution of winning bids is precisely the solution to the relaxed program, so that minimum revenue is attained. Moreover, it is always a bidder with a high value who receives the good, so that the equilibrium allocation is efficient. We remark that the similarity in strategies between our construction and the "as-if" IPV model leads to the following interpretation of the minimum winning-bid distribution: it is the distribution of winning bids that would arise in an independent private-values model in which the distribution of the highest value is equal to $Q(v)$, the true distribution of the average of the $N-1$ lowest values.

Is there some deeper reason why this information structure is able to attain the bounds? Put differently, what are the essential properties of this construction that make it work? While there are some degrees of freedom in how one constructs the winning-bid-minimizing information structure and equilibrium, there are certain necessary features that could be *derived* from the solution to the relaxed program. First, the bidder with the highest value must end up winning the good. Second, the bidders must be indifferent to all uniform upward deviations. This second property actually implies that bidders must be indifferent to *all* upward deviations, uniform or otherwise. For if a bidder strictly preferred not to deviate up to some b for a set of equilibrium bids $X \subset [0, b]$ that arise with positive probability in equilibrium, then indifference to the uniform upward deviation implies that they strictly prefer to deviate up to b when the equilibrium bid is in $[0, b] \setminus X$. Now, the value of a "pointwise" upward deviation from an equilibrium bid x up to $b > x$ depends on the likelihood of losing with a bid of x to a winning bid less than b for each possible value, and it turns out that indifference to all upward deviations exactly pins down those likelihoods, and hence the marginal distribution of each bidder's losing bids. Moreover, it is always possible to normalize signals so that bidders bid b after receiving a signal $\underline{\beta}^{-1}(b)$, in which case the marginal distribution of each bidder's losing signal is exactly as in our construction. In the case of $N = 2$ and a pure common value, this describes an information structure that is essentially unique. For $N > 2$, there is some flexibility as to how losing bidders' signals are correlated with one another and with the highest value, neither of which affects incentive constraints. It is always possible, however, to make losing signals independent of one another and independent of $\max v$, conditional on $\alpha(v)$, as we have done.

As to a deeper reason why this information structure minimizes the distribution of winning bids, the best explanation we can give is that because the highest signal is $\alpha(v)$ and bidders use monotonic pure strategies, as a bidder deviates up, their value on the marginal event that they win increases in lockstep with their bid. Moreover, independence of the signals means that all bidders face the same tradeoff as they deviate up between additional wins and a higher price for all wins. This makes it possible to have all types of bidders indifferent to deviating up, regardless of the signal they received, thereby filling up all of the upward incentive constraints and pushing bids down as much as possible. It is nonetheless remarkable that all of these properties can be achieved by a single information structure and equilibrium.

4.4 Asymmetric Value Distributions

Finally, let us reconsider the maintained assumption that the distribution of values is symmetric. This assumption was used in Lemma 3, in which we proved that it was without loss of generality to restrict attention to symmetric solutions to the relaxed program, and also in the construction of the information structure and equilibrium that attain the bounds. We do not have a tight characterization of minimum bidding for general asymmetric distributions of values. We can, however, use our methods to obtain a lower bound on the equilibrium winning-bid distribution. Specifically, for an arbitrary distribution μ, we continue to define Q to be the distribution of the average $N-1$ lowest values, and similarly define $\underline{\beta}$ and $\underline{H}$ as in equations (6) and (7). The following proposition is proven in the Online Appendix:

Proposition 3 (Minimum Bidding with Asymmetric Distributions)**.** *For any information structure $\mathcal{S}$ and equilibrium σ, the induced distribution of winning bids $H(\mathcal{S},\sigma)$ must first-order stochastically dominate $\underline{H}$.*

Let us briefly describe the logic behind this result. Notice that Lemmas 1 and 2 did not use the hypothesis that μ is symmetric. Thus, the relaxed program is well-defined for asymmetric μ, it has a solution, and for each b, the solution gives an upper bound on the probability that the winning bid is less than b. Now, given a feasible solution for the relaxed program for μ, it turns out that we can construct a feasible solution for the relaxed program for a "symmetrized" prior $\widetilde{\mu}$, which is generated from μ by randomly permuting the identities of the bidders, and this solution for $\widetilde{\mu}$ induces the same marginal distribution over winning bids. Moreover, the distribution of the average of the $N-1$ lowest values is the same for

μ and for $\widetilde{\mu}$. Since we know that the solution to the relaxed program for $\widetilde{\mu}$ is $\underline{H}$, we can conclude that the latter is a lower bound on the distributions of winning bids that can arise under μ.

5 Further Results on Revenue and Bidder Surplus

5.1 An Independent Value Example

We will now broaden our view and consider the whole set of revenue and bidder surplus pairs that could arise in the first-price auction, in addition to those corresponding to minimum revenue and maximum bidder surplus. Our leading example for this section is one with two bidders whose values are independent draws from the standard uniform distribution. Note the contrast with the example of Section 3, in which the bidders have the same value so that every allocation is efficient. In the present case, the surplus that is generated by the auction depends on the allocation. The average of the $N-1$ lowest values is simply the lower of the two uniform draws, which has distribution $Q(v) = 1 - (1-v)^2$. Thus, the revenue-minimizing bidding function (4) reduces to

$$\underline{\beta}(v) = \frac{1}{\sqrt{1-(1-v)^2}} \int_{x=0}^{v} \frac{x(1-x)}{\sqrt{1-(1-x)^2}} dx.$$

Minimum revenue does not have a simple analytical expression, but it numerically integrates to $\underline{R} \approx 0.096$. In addition, the efficient surplus is $\overline{T} = 2/3$, so that maximum bidder surplus is $\overline{U} \approx 0.571$.

For comparison, in the independent private-values model—when each bidder only knows his own value but maintains the common prior regarding the other bidder's value—the bidders' surplus is 1/3 and revenue is 1/3. This is the same outcome as would obtain in the complete information model where both bidders observe both values. Maximum bidder surplus is therefore approximately 1.7 times larger than that predicted by either of those information structures. By contrast, in the no-information environment in which each bidder knows nothing about the values except the prior distribution, the bidders compete the price up to their expected values of 1/2. As a result, the allocation will be ex-post inefficient, revenue is 1/2, and bidder surplus is 0.

Figure 1 illustrates results for this example. Possible combinations of bidder surplus (on the x-axis) and revenue (on the y-axis) are plotted. As the maximum total surplus is 2/3, the efficient allocations correspond to the -45 degree line on the right of the picture. The worst case for efficiency would be that the object is always sold to the bidder with the lowest

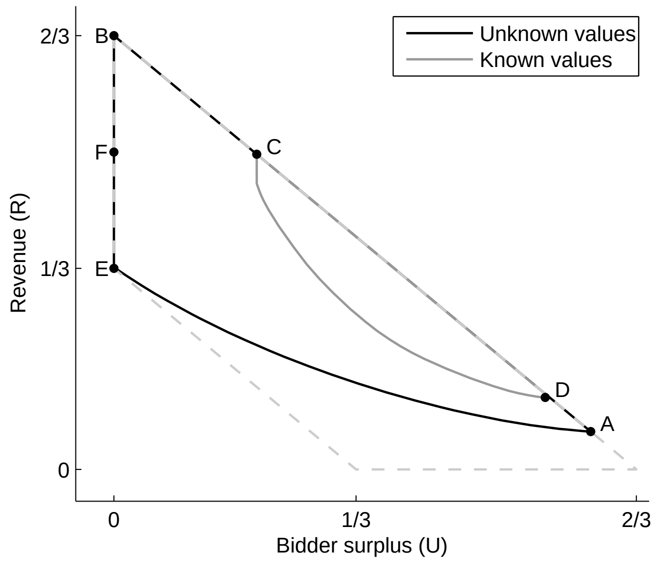

Figure 1: The set of revenue-bidder surplus pairs that can arise in equilibrium for some information structure.

value, which would generate a total surplus of 1/3. Thus, the dashed trapezoid represents the surplus pairs that satisfy this range restriction on total surplus and also give non-negative surplus to both the bidders and the seller.

We will now consider the whole range of revenue and bidder surplus across all possible information structures and equilibria. This includes, in particular, maximum revenue, minimum bidder surplus, and minimum total surplus. We start by considering all information structures, including those in which the bidders' signals do not reveal their own values exactly. We refer to this model as one of *unknown values*, to distinguish it from the model we consider next. The set of surplus pairs that can arise in the unknown-values model is the area enclosed by the dark grey curve in Figure 1. The revenue-minimizing outcome identified in the previous section is attained at point A. Beyond minimum revenue, we can see that there is wide a range of possible welfare outcomes. There are two extreme points which stand out: At point B, the allocation is efficient but bidder surplus is zero, which necessarily attains maximum revenue. At point C bidder surplus is again zero, but the outcome is minimally efficient and the revenue is held down to the minimum feasible surplus. We show that an analogous but somewhat weaker inefficiency result holds for many bidders and independent values.

A natural question to ask is what would happen to our prediction if we assume that the bidders have more information. In the *known-values* model, we assume that the bidders at least know their own values, and their signals may contain additional information about others' signals. The set of welfare outcomes that can arise under known-values information structures and equilibria is enclosed by the light grey curve in Figure 1.[13] In contrast to unknown values, known values implies that each bidder can guarantee himself a strictly positive surplus. We will describe a lower bound on bidder surplus and a corresponding upper bound on revenue which turn out to be tight, and in Figure 1 they are attained at point D. The working paper Bergemann, Brooks, and Morris (2015) identifies such bounds in general settings and proves that they are tight.

Bergemann and Morris (2016) characterize the set of joint distributions of bids and values that can be induced by some information structure and equilibrium as a particular class of incomplete-information correlated equilibria that they term *Bayes correlated equilibria* (BCE). Loosely speaking, a distribution $\phi \in \left(V^N \times B^N\right)$ is an unknown-values BCE if b_i is a best response to the conditional distribution $\phi\left(v, b_{-i}|b_i\right)$ of players' values and others' bids given that bidder i's bid is b_i. Similarly, ϕ is a known-values BCE if b_i is a best response to $\phi\left(v_{-i}, b_{-i}|b_i, v_i\right)$ for every (v_i, b_i). These incentive constraints are linear, so that for a discretized version of the auction with finitely many values and bids, the set of BCE is a convex polytope, and the problem of maximizing a linear objective, e.g., revenue, over all BCE is a linear program. We used this methodology to compute the range of welfare outcomes for cases in which we do not have an analytical characterization. In particular, while points A–D are derived analytically, as described above, other points are computed numerically for an independent uniform distribution with grids of 10 values and 50 bids between 0 and 1. The axes have been re-scaled to match moments with the continuum limit; for the discretized example, the efficient surplus and minimum surplus are respectively 41/60 and 19/60, as opposed to their limit values of 2/3 and 1/3.

5.2 Maximum Revenue

We will now give a general characterization of maximum revenue with unknown values. To maximize revenue and minimize bidder surplus, we would like to generate a highly competitive environment where the highest bid is equal to the highest value. To develop intuition for how this can be done, we will reason by analogy with a slightly different model in which the tie-breaking rule endogenously depends on the realized values. Consider the information structure in which each bidder receives a signal that is equal to the highest value, independent

[13]The fact that the known-values set is contained within the unknown-values set is a reflection of the general observation that providing the bidders with more information decreases the set of outcomes that can be rationalized as an equilibrium with even more information. In other words, the set of Bayes correlated equilibria is decreasing in the minimum information of the players. Bergemann and Morris (2016) formalize the notion of "more information" and give a precise statement of this result.

of who has the highest value:

$$s_i = \max v, \text{ for all } i.$$

Let us also suppose that ties are broken in favor of the bidder with the highest value. Now consider the strategy profile in which each bidder bids his signal. Under this strategy profile, each bidder is indifferent between all bids less than or equal to his signal s_i: bidding less than the signal results in losing the auction for sure, and by bidding the signal, efficient tie breaking means that the winner has the highest value, which is equal to the signal and the bid. Moreover, no one would want to deviate to a higher bid, since this would result in winning for sure and paying a price which is greater than the highest value. Thus, these strategies are an equilibrium, and the resulting winning bid is the maximum value, so that revenue is the efficient surplus.

This argument relied on the endogenous tie-breaking rule. However, it is possible to achieve approximately the same outcome with the uniform (and thus potentially inefficient) tie-breaking rule by suitably perturbing the information structure:

Theorem 2 (Maximal Revenue and Minimum Bidder Surplus).
For all $\varepsilon > 0$, there exists an information structure and equilibrium such that revenue is at least $\overline{T} - \epsilon$ and bidder surplus is less than ϵ.

To establish the result under the uniform tie-breaking rule, consider the following information structure. First, if there is a tie for the highest value, then the bidders receive signals that tell them that this event has occurred and in fact reveal the entire profile of values. The bidders with the highest value will then compete the price up to their value, so that the seller obtains all of the surplus. Otherwise, let us assume that the highest value $v^{(1)}$ is strictly larger than the second-highest value $v^{(2)}$. In this case, the bidder with the highest value observes a signal that is a convex combination of the highest and second-highest values, $xv^{(1)} + (1-x)\,v^{(2)}$, where x has a distribution $F(x)$ that is non-atomic and full support on $[0,1]$. The losing bidders then observe conditionally independent signals $yv^{(1)} + (1-y)v^{(2)}$, where the weight y is drawn independently across bidders on the interval $[0,x]$ from the distribution

$$y \sim F(y|x) = \left(\frac{y}{1-y}\frac{1-x}{x}\right)^{\frac{1}{N-1}}.$$

In equilibrium, bidders follow the pure strategy of bidding their signal. As we show in the Online Appendix, these strategies form an equilibrium. Moreover, we can choose the distribution of x so that it is arbitrarily close to a Dirac measure on $\{1\}$ in the weak-* topology, e.g., $F(x) = x^k$ where k is arbitrarily large. In the limit as F converges to a point mass on one, revenue approaches the efficient surplus, but the bidders surplus is arbitrarily close to zero.

Thus, in an environment with unknown values, the private information of each bidder might be sufficiently confounding to induce very aggressive bidding behavior. The bidders are willing to bid a large amount because they think that the bid is less than their value conditional on *winning*, although their value might be quite a bit lower than their bid conditional on *losing*. As a result, the strategy of bidding one's signal is weakly undominated.

We note that the construction of the bid distribution described above exploits symmetry among the bidders, but the argument could be extended to asymmetric distributions of values, assuming only that the support is symmetric.

5.3 Minimum Efficiency

Thus far, our analysis has led us to equilibria in which the allocation of the good was efficient, so that the welfare outcome lay on the northeast frontier of Figure 1. As the figure plainly shows, there are also a large number of possible outcomes in which the allocation is inefficient. That some inefficiency might arise is obvious, for example when the bidders have no information about values except the prior. What is more striking is the extent of this potential inefficiency. In particular, point C attains a maximally inefficient outcome in which the good is always allocated to the buyer with the *lowest* value, all while giving the bidders zero rents.

For the model underlying Figure 1, with two bidders and independent uniform valuations, there is an extremely simple information structure and equilibrium which attains this outcome: each bidder observes the *other* bidder's valuation, and bids half of what they observe. To see that this is an equilibrium, consider a bidder i who has observed a signal s_i. Because of independence, this signal contains no information about the bidder's own value v_i, which has a posterior distribution that is uniform. Now, conditional on bidding some b_i, bidder i wins whenever b_j is less than b_i, which is when v_i is less than $2b_i$. Thus, the expectation of v_i conditional on winning with a bid of b_i is just the expectation of a uniform random variable conditional it being below $2b_i$, which is precisely b_i. No matter what bidder i bids in the range $[0, 1/2]$, the expected value conditional on winning equals the bid. In the end, bidders receive zero rents in equilibrium, and since bids are monotonic in the other bidder's value, it is the bidder with the lowest value who wins the auction.

While we have not explored minimum efficiency in its full generality, we report in the Online Appendix a class of information structures and equilibria which generalize this example to the case of many bidders and symmetric independent values drawn from a cumulative distribution $P(v)$. Specifically, each bidder observes the maximum of others' values $s_i = \max v_{-i}$ and bids the expectation of a value drawn from the prior conditional on it being below s_i. In the resulting equilibrium, each of the $N-1$ low-value bidders wins the object with equal probability. When there are two bidders, this construction attains the maximally inefficient outcome, and more generally, we conjecture that this equilibrium minimizes total surplus subject to the constraint that bidder surplus is zero.

To sum up, we have characterized three "corners" of the unknown-values set (for the two-bidder case), and by convexity, we can generate both the western and northeastern flats of the blue region in Figure 1. The remaining feature of the unknown values set in Figure 1, hitherto unexplained, is the apparently strictly convex southwestern frontier that runs from the maximally inefficient equilibrium to the efficient revenue-minimizing equilibrium. In the working paper, Bergemann, Brooks, and Morris (2015) we give a complete description of the class of equilibria that generate this southwestern frontier. They are members of a class of "conditionally-revenue-minimizing" equilibria, which minimize revenue conditional on a fixed allocation of the good. As the allocation ranges from efficient to maximally inefficient, we move between points A and C, and there is a particular class of allocations that generate the frontier.

5.4 Minimum and Maximum Revenue with Known Values

In the environment with arbitrary interdependence in values, we have a complete characterization of minimum and maximum revenue and bidder surplus. One might ask how our results would change if we imposed additional restrictions on how much bidders can learn about their values from the outcome of the auction. An extreme assumption, but one that is ubiquitous in the literature, is that each bidder knows his own value for sure. This is what we call the *known-values* case.

The assumption that bidders know their own values substantially affects the set of possible outcomes. It is no longer the case that bidder surplus can be driven all the way down to zero. In the working paper Bergemann, Brooks, and Morris (2015), we derive an elementary lower bound on each bidder's surplus, which we describe here in the context of the two-bidder independent uniform example. As each bidder knows his value for the object, any weakly undominated strategy profile requires that the bidders never bid above their values. Thus, each bidder knows that if they bid b, then they will necessarily win whenever the other bidder's value is less than b. In the independent uniform example, the surplus from bidding b when the value is v is at least $(v-b)\,b$, and maximizing this quantity over all b implies that the bidder is guaranteed at least $v^2/4$ in surplus. In ex-ante terms, bidder i must receive at least

$$\underline{U}_i = \frac{1}{4}\int_{v=0}^{1} v^2 dv = \frac{1}{12}.$$

In Bergemann, Brooks, and Morris (2015) we establish in Theorem 3 that this lower bound is in fact tight: there is an information structure and equilibrium in which bidder i receives *exactly* $1/12$ in surplus. Moreover, bidders can be held to this lower bound while maintaining an efficient allocation, so that this equilibrium simultaneously minimizes bidder surplus and maximizes the revenue of the seller at $2/3 - 2\,(1/12) = 1/2$. This result can be generalized well beyond the two-bidder independent uniform example, to any symmetric or asymmetric joint distribution of values μ.

With regard to minimum revenue, the model with unknown values provides a lower bound on revenue for the model with known values. In Bergemann, Brooks, and Morris (2015), we provide a complete characterization of minimum revenue with known values for the special case where there are only two possible values for the bidders, i.e., $|V| = 2$. This characterization employs a similar methodology of (i) formulating and solving a relaxed program for revenue, and (ii) extending the solution to the relaxed program to an information structure and equilibrium. We also discuss the scope for extending this program beyond binary values, though at the time of writing we do not have a general analytical characterization of minimum revenue in known-values models. This remains an interesting avenue for future research.

The known-values surplus set is significantly smaller than the unknown-values surplus set, and—except for the known-values maximum-revenue result (point D)—is derived from computations. It is notable that the inefficiencies that can arise in the known-values model are relatively small compared to what can happen with unknown values, as visually expressed by the slimness of the red lens that describes the set of all possible equilibrium surplus realizations. This observation is in line with the results of Syrgkanis and Tardos (2013) and Syrgkanis (2014) who show that the efficiency loss in the independent private-value auction expressed in terms of the ratio between realized surplus and efficient surplus in the first-price auction is bounded below by $1 - 1/e$.

5.5 The Many-Bidder Limit

We conclude this section by relating our model to the classic question of the performance of auctions when the number of bidders becomes large. Consider a sequence of economies indexed by N, each of which is associated with a joint distribution of the potential bidders' values. Our preceding analysis tells us that the features of the economy that matter for minimum revenue and maximum bidder surplus are (i) the distribution of the average losing value, which determines the minimum winning-bid distribution, and (ii) the total surplus that is generated by an efficient allocation. Thus, we can characterize behavior in the many-bidder limit by characterizing the behavior of these two objects: if the distribution of the average losing value converges to a limit $Q(w)$, then minimum winning-bid function converges to the limit as $N \to \infty$ of (6), i.e.,

$$\underline{\beta}(w) = \frac{1}{Q(w)} \int_{x=\underline{w}}^{w} x\, Q(dx), \tag{20}$$

and minimum revenue converges to

$$\underline{R} = \int_{w=\underline{w}}^{\overline{w}} \underline{\beta}(w)\, Q(dw).$$

In addition, if the efficient surplus converges to $\overline{T}$, then maximum bidder surplus converges to $\overline{U} = \overline{T} - \underline{R}$.

In the pure-common-value case, the distribution of the average losing value is just the distribution of the common value. If we hold this distribution fixed as N grows large, then revenue and bidder surplus converge to the expressions above where Q is the distribution of the common value. Thus, revenue is bounded away from the total surplus and bidder surplus is bounded away from zero. This conclusion, while perhaps surprising, is not novel. For example, the same result is obtained by Engelbrecht-Wiggans, Milgrom, and Weber (1983) when one bidder is informed and $N-1$ bidders are uninformed. In that case, both the informed bidder's strategy and surplus are independent of N, uninformed bidders obtain zero rents, and the uninformed bidders strategies adjust so as to support the informed bidder's behavior. In general, though, maximum bidder surplus is strictly greater than that obtained in the model of Engelbrecht-Wiggans et al. In fact, the limiting winning-bid function (20) is exactly equal to the strategy of the informed bidder, though the informed bidder loses the auction with non-vanishing probability. For example, in the case of a uniform distribution, the informed bidder's surplus is 1/6, whereas maximum bidder surplus converges to 1/4 as N grows large.

Another leading case is the one where bidders' values are independent draws from fixed prior $P \in \Delta([0,1])$. When the number of bidders is large, the distribution of the average of the $N-1$ lowest values converges to a Dirac measure on the mean of P:

$$\hat{v} = \int_{v=\underline{v}}^{\overline{v}} v\, P(dv).$$

In the limit, the winning bid converges almost surely to $\hat{v}$. The allocation is also efficient so total surplus converges to $\overline{v}$ and bidder surplus converges to $\overline{v} - \hat{v}$. Thus, revenue is what would obtain if the bidders had no information beyond the prior, but the allocation is asymptotically efficient and all of the additional surplus goes to the bidders.

The conclusion that minimum revenue in the many-bidder limit is bounded away from the efficient surplus stands in stark contrast with the literature on information aggregation in large markets (Wilson, 1977; Milgrom, 1979; Bali and Jackson, 2002). This literature considers two distinct but related issues: first, is information aggregated in large markets, in the sense that the winning bid reveals the true maximal value of the good? And second, does this revelation of information induce the bidders to compete away their rents? The positive results in this literature rely upon assumptions about information which our constructions violate. Both Wilson (1977) and Bali and Jackson (2002) assume that there is a uniform lower bound on the quality of bidders' information even as the number of bidders becomes large. In contrast, in the revenue-minimizing information structure, the marginal distribution of each losing bidder's signal converges to a point mass on the lowest possible average losing value, so that in the limit the signals are individually uninformative about the maximum value. Thus, the sale price is fully revealing about the value without competition forcing the price up to the true value. With non-common-values, we have both a failure of information aggregation and a failure of full surplus extraction in the limit, although enough information is aggregated to ensure an efficient allocation.

6 Discussion

This paper has provided new and general characterizations of equilibrium bidding in the first-price auction. More broadly, this paper contributes to the growing literature on information-free predictions in Bayesian games. We believe that this perspective has value for both theoretical and applied work. There are many real-world scenarios involving Bayesian games in which practitioners cannot reasonably be expected to know the nature of information held by strategic agents. In such a setting, it would be sensible for researchers and policy makers to base their inference or institutional design on the admittedly weaker but also much safer information-free methodology that we have employed in the present paper. In that spirit, we will conclude by highlighting some implications of our results.

First, there is a large and active literature on auction econometrics that aims to back out payoff-relevant fundamentals from observed bidding behavior, and a sizable portion of this literature focuses on the first-price auction (e.g., Laffont, Ossard, and Vuong, 1995, Athey and Haile, 2007, and Somaini, 2015). This work interprets the received model of the first-price auction quite literally, and assumes that bids are generated by an equilibrium under a classical—i.e., affiliated—information structure. Alternatively, one could use our theory to derive bounds on the value distribution that rationalizes an observed distribution of winning bids H. Let us suppose that H was generated by a Bayesian equilibrium under some common prior information structure about a pure common value distributed according to $P(v)$. Then it must be that the minimum winning-bid distribution $\underline{H}$ associated with P is first-order stochastically dominated by H. Loosely speaking, P cannot put too much mass on high values, or else H will not dominate the minimum winning-bid distribution for P. Since we cannot analytically invert $\underline{\beta}$ to evaluate (7), it is more convenient to express this constraint in terms of percentiles as

$$\underline{\beta}\left(P^{-1}(z)\right) = z^{-\frac{N-1}{N}} \left(z^{\frac{N-1}{N}} P^{-1}(z) - \int_{x=0}^{P^{-1}(z)} (P(x))^{\frac{N-1}{N}} dx \right) \leq H^{-1}(z) \qquad (21)$$

for all probabilities $z \in [0,1]$, where the expression for $\underline{\beta}(P^{-1}(z))$ follows from integrating (6) by parts. We can rewrite this constraint as

$$\Gamma(z;P) \leq H^{-1}(z)\, z^{\frac{N-1}{N}},$$

where

$$\Gamma(z;P) = z^{\frac{N-1}{N}} P^{-1}(z) - \int_{x=0}^{P^{-1}(z)} (P(x))^{\frac{N-1}{N}} dx$$

denotes the area below the line $y = z^{\frac{N-1}{N}}$ and above the curve $(P(w))^{\frac{N-1}{N}}$, which is depicted in Figure 2(a).[14] In sum, we think a promising direction for future research is to further investigate how information-free predictions can be used for partial identification of value distributions from bidding data.

Second, our approach can be used to compare the set of possible welfare outcomes across mechanisms. Myerson (1981) showed that introducing a positive reserve price can increase revenue for a fixed information structure, and we can also improve our lower bound on revenue using reserve prices. In the Online Appendix, we extend our results on minimum and maximum revenue to the first-price auction with a reserve price for pure-common-value environments. Figure 2(b) reports these revenue bounds for the standard-uniform pure-common-value example. With a zero reserve price, the mechanism is exactly the first-price

[14]Note that Γ is monotonic in P in the first-order stochastic dominance ordering. Thus, if a distribution P satisfies (21) for all $z \in [0,1]$ and if $\widetilde{P}$ is first-order stochastically dominated by P, then $\widetilde{P}$ also satisfies (21) for all z. This does not imply, however, that the set of distributions satisfying (21) has a maximal element in the first-order stochastic dominance ordering, nor is it true that every distribution of winning bids H can be rationalized as the minimum winning-bid distribution for some distribution of a common value.

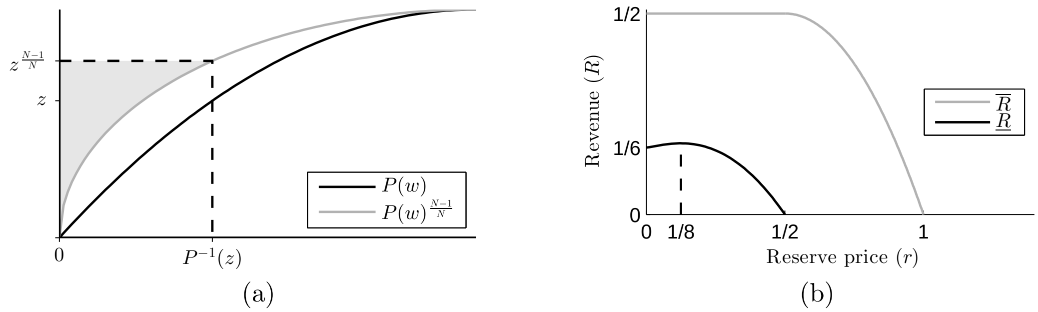

Figure 2: (a) Drawing inference about the distribution of values: for all $z \in [0,1]$, the green shaded area must be less than $z^{\frac{N-1}{N}} H^{-1}(z)$. (b) Maximum revenue $\overline{R}(r)$ and minimum revenue $\underline{R}(r)$ as functions of the reserve price r.

auction studied in Section 3, and maximum revenue $\overline{R}$ and minimum revenue $\underline{R}$ being 1/2 and 1/6, respectively. For $r > 0$, $\overline{R}(r)$ and $\underline{R}(r)$ are concave functions, with maximum $\overline{R}(r)$ being attained at any $r < 1/2$ and maximum $\underline{R}(r)$ being attained at $r^* = 1/8$. The latter has the following interpretation: Consider a seller who knows that there is a standard-uniform common value, is ambiguity averse with respect to both the information structure and selection of Bayesian equilibrium, and has to sell using a first-price auction with reserve price. Then the reserve price for such a seller is precisely 1/8.

More generally, we have shown that the first-price auction (even without a reserve price) is guaranteed to generate positive revenue, regardless of the information structure and equilibrium. This is not true of the second-price auction. Even when buyers know their own values, the second-price auction admits weakly-dominated equilibria in which revenue is zero, and there are information structures that add small uncertainty about one's values in which revenue is essentially the same but equilibrium strategies are no longer weakly dominated. In our view, an exciting direction for future research is to look for mechanisms that provide even more favorable revenue guarantees than first-price auctions with reserve prices, or more broadly, to look for mechanisms that are guaranteed to perform well regardless of the structure of information and equilibrium.

References

Athey, S. and P. Haile (2007): "Nonparametric Approaches to Auctions," in *Handbook of Econometrics*, ed. by J. Heckman and E. Leamer, Amsterdam: Elsevier, vol. 6A, 3847–3965.

Azacis, H. and P. Vida (2015): "Collusive Communication Schemes in a First-Price Auction," *Economic Theory*, 58, 125–160.

Bali, V. and M. Jackson (2002): "Asymptotic Revenue Equivalence in Auctions," *Journal of Economic Theory*, 106, 161–176.

Battigalli, P. and M. Siniscalchi (2003): "Rationalizable Bidding in First-Price Auctions," *Games and Economic Behavior*, 45, 38–72.

Bergemann, D., B. Brooks, and S. Morris (2013): "Extremal Information Structures in First Price Auctions," Cowles Foundation for Research Discussion Paper 1926, Yale University.

——— (2015): "First Price Auctions with General Information Structures: Implications for Bidding and Revenue," Tech. Rep. 2018, Cowles Foundation for Research in Economics, Yale University.

——— (2016): "Optimal Auction Design with Common Values," Tech. rep., Yale University, University of Chicago, and Princeton University.

Bergemann, D. and S. Morris (2013): "Robust Predictions in Games with Incomplete Information," *Econometrica*, 81, 1251–1308.

——— (2016): "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, 11, 487–522.

Bulow, J. and P. Klemperer (2002): "Prices and the Winner's Curse," *RAND Journal of Economics*, 33, 1–21.

Çinlar, E. (2011): *Probability and Stochastics*, New York: Springer.

Dekel, E. and A. Wolinsky (2003): "Rationalizable Outcomes of Large Private-Value First-Price Discrete Auctions," *Games and Economic Behavior*, 43, 175–188.

Engelbrecht-Wiggans, R., P. Milgrom, and R. Weber (1983): "Competitive Bidding and Proprietary Information," *Journal of Mathematical Economics*, 11, 161–169.

Fang, H. and S. Morris (2006): "Multidimensional Private Value Auctions," *Journal of Economic Theory*, 126, 1–30.

Feldman, M., B. Lucier, and N. Nisan (2016): "Correlated and Coarse Equilibria of Single-Item Auctions," Tech. rep., Tel Aviv University, Microsoft Research, and Hebrew University.

Krishna, V. (2002): *Auction Theory*, San Diego: Academic Press.

Laffont, J.-J., H. Ossard, and Q. Vuong (1995): "Econometrics of First-Price Auctions," *Econometrica*, 953–980.

Lopomo, G., L. Marx, and P. Sun (2011): "Bidder Collusion at First-Price Auctions," *Review of Economic Design*, 15, 177–211.

Milgrom, P. (1979): "A Convergence Theorem for Competitive Bidding with Differential Information," *Econometrica*, 47, 679–88.

Milgrom, P. and R. Weber (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089–1122.

Myerson, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

Somaini, P. (2015): "Identification in Auction Models with Interdependent Costs," Tech. rep., Stanford University.

Syrgkanis, V. (2014): "Efficiency of Mechanisms in Complex Markets," Ph.D. thesis, Cornell University.

Syrgkanis, V. and É. Tardos (2013): "Composable and Efficient Mechanisms," in *STOC '13 Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, 211–220.

Vickrey, W. (1961): "Counterspeculation, Auctions and Competitive Sealed Tenders," *Journal of Finance*, 16, 8–37.

Wilson, R. (1977): "A Bidding Model of Perfect Competition," *Review of Economic Studies*, 44, 511–518.

A Appendix

Proof of Lemma 1. Fix an information structure $\mathcal{S}$ and strategies σ, and define the uniform deviation up to b by its distribution function $\sigma_i^b\left([0,x]\,|s\right) = \mathbb{I}_{x\geq b}\sigma_i\left([0,x]\,|s\right)$. Now, fix $b \in B$. Observe that the marginal distribution of each bidder's bids

$$\int_{v\in V^N} \sigma_i\left(db_i|s_i\right)\pi\left(ds|v\right)\mu\left(dv\right)$$

can have only countably many mass points. Thus, it must be possible to find ϵ arbitrarily small such that

$$\int_{v\in V^N} \sigma_i\left(\{b+\epsilon\}\,|s_i\right)\pi\left(ds|v\right)\mu\left(dv\right) = 0$$

for all i. Fix such an ϵ. Then for σ to be an equilibrium, it must be that $U_i\left(\mathcal{S},\sigma\right) \geq U_i\left(\mathcal{S},\sigma_i^{b+\epsilon},\sigma_{-i}\right)$. The latter utility is

$$U_i\left(\mathcal{S},\sigma_i^{b+\epsilon},\sigma_{-i}\right) = \int_{v\in V^N}\int_{s\in S}\int_{x\in B^N}\left(v_i - x_i\right)q_i\left(x\right)\left(\sigma_i^{b+\epsilon},\sigma_{-i}\right)\left(dx|s\right)\pi\left(ds|v\right)\mu\left(dv\right).$$

Note that for each i, the joint distribution of bids is unchanged outside of $\left\{x\in B^N|x_i\leq b+\epsilon\right\}$, and bidder i deterministically bids $b+\epsilon$ on this event, so we can rewrite this surplus as

$$\begin{aligned}
&\int_{v\in V^N}\int_{s\in S}\int_{x_{-i}\in B^{N-1}}\left(v_i-b-\epsilon\right)q_i\left(b+\epsilon,x_{-i}\right)\sigma_{-i}\left(dx_{-i}|s_{-i}\right)\sigma_i\left([0,b+\epsilon)\,|s_i\right)\pi\left(ds|v\right)\mu\left(dv\right)\\
&\quad+\int_{v\in V^N}\int_{s\in S}\int_{x\in B^N}\left(v_i-x_i\right)\mathbb{I}_{x_i\geq b+\epsilon}q_i\left(x\right)\sigma\left(dx|s\right)\pi\left(ds|v\right)\mu\left(dv\right)\\
&=\int_{v\in V^N}\int_{s\in S}\left(v_i-b-\epsilon\right)\sigma\left([0,b+\epsilon)^N\,|s\right)\pi\left(ds|v\right)\mu\left(dv\right)\\
&\quad+\int_{v\in V^N}\int_{s\in S}\int_{x\in B^N}\left(v_i-x_i\right)\mathbb{I}_{x_i\geq b+\epsilon}q_i\left(x\right)\sigma\left(dx|s\right)\pi\left(ds|v\right)\mu\left(dv\right).
\end{aligned}$$

where we have used the facts that each bidder bids $b+\epsilon$ with (ex-ante) probability zero and that bidder i wins with probability one when bids are in $[0,b+\epsilon)^N$. Thus, the incentive constraint is equivalent to

$$\begin{aligned}
&\int_{v\in V^N}\left(v_i-b-\epsilon\right)\int_{s\in S}\sigma\left([0,b+\epsilon)^N\,|s\right)\pi\left(ds|v\right)\mu\left(dv\right)\\
&\leq\int_{v\in V^N}\int_{s\in S}\int_{x\in B^N}\left(v_i-x_i\right)\mathbb{I}_{x_i<b+\epsilon}q_i\left(x\right)\sigma\left(dx|s\right)\pi\left(ds|v\right)\mu\left(dv\right)\\
&=\int_{v\in V^N}\int_{s\in S}\int_{x\in[0,b+\epsilon)^N}\left(v_i-x_i\right)q_i\left(x\right)\sigma\left(dx|s\right)\pi\left(ds|v\right)\mu\left(dv\right)
\end{aligned}$$

where the second equality is coming from the fact that $\mathbb{I}_{x_i<b+\epsilon} q_i(x) > 0$ only if $x \in [0, b+\epsilon)^N$. Now, as $\epsilon \to 0$,

$$\int_{s\in S} \sigma\left([0,b+\epsilon)^N |s\right) \pi(ds|v) \to \int_{s\in S} \sigma\left([0,b]^N |s\right) \pi(ds|v) = H(b|v;\mathcal{S},\sigma),$$

which follows from summing (9) over $i \in \mathcal{N}$ and the fact that $\sum_{i\in\mathcal{N}} q_i(x) = 1$, and invoking countable additivity. Similarly,

$$\int_{s\in S}\int_{x\in[0,b+\epsilon)^N} q_i(x)\,\sigma(dx|s)\,\pi(ds|v) \to \int_{s\in S}\int_{x\in[0,b]^N} q_i(x)\,\sigma(dx|s)\,\pi(ds|v) = H_i(b|v;\mathcal{S},\sigma).$$

Plugging in these expressions yields (11). □

Proof of Lemma 2. We can identify the domain of the relaxed program with the subset D of $\Delta\left(B \times \mathcal{N} \times V^N\right)$ that have $\mu(dv)$ as the marginal over V^N, and satisfy

$$\int_{v\in V^N}\int_{x\in B}\sum_{j\in\mathcal{N}\setminus\{i\}} \mathbb{I}_{x\le b}(v_i - b)\,\phi(dx,j,dv) \le \int_{x\in[0,b]}\left(\int_{v\in V^N}\int_{y\in[0,x]}\phi(dy,i,dv)\right)dx. \quad (22)$$

for every b. For any such distribution ϕ disintegrates into the marginal μ and a probability transition kernel $K : V^N \to \Delta(B \times \mathcal{N})$, for which $H_i(b|v) = K([0,b]\times\{i\}|v)$ (Çınlar, 2011, Theorem IV.2.18). Under this transformation, (22) reduces to (12). The objective of the relaxed program is simply to maximize

$$\int_{v\in V^N}\int_{x\in B}\sum_{i\in\mathcal{N}} \mathbb{I}_{x\le\hat{b}}\phi(dx,i,dv),$$

which is upper semi-continuous in ϕ. Thus, we can verify existence of a solution by proving that D is weak-* compact. This is not trivial, however, since for each individual constraint (22), the set of distributions satisfying the constraint is not closed. We will show, however, that the set D as a whole is closed. Note that $\Delta\left(B\times\mathcal{N}\times V^N\right)$ is weak-* compact, so that compactness of D will follow from the fact that it is closed.

Let $\left\{\phi^k\right\}$ be a sequence D that converges to some ϕ. It is trivial to show that ϕ has μ as a marginal on V^N, and we only need to verify that ϕ satisfies (22). We know that

$$\begin{aligned}\int_{v\in V^N}\int_{x\in B}\sum_{j\in\mathcal{N}\setminus\{i\}} \mathbb{I}_{x\le b}(v_i-b)\,\phi^k(dx,j,dv) &\le \int_{x\in[0,b]}\left(\int_{v\in V^N}\int_{y\in[0,x]}\phi^k(dy,i,dv)\right)dx \\ &= \int_{x\in[0,b]} H_i^k(x)\,dx\end{aligned}$$

for all k, where $H_i^k(x) = \phi^k\left([0,x] \times \{i\} \times V^N\right)$. Because the $H_i^k(x)$ are distribution functions that are weakly converging to $H_i(x) = \phi\left([0,x] \times \{i\} \times V^N\right)$, it must be that $H_i^k(x) \to H_i(x)$ whenever H_i is continuous at x, which must be true for all but countably many values of x. Thus, the integral on the right-hand side must converge to $\int_{x\in[0,b]} H_i(x)\,dx$ as $k \to \infty$. If we additionally assume that $\phi\left(\{b\} \times \mathcal{N} \times V^N\right) = 0$, then

$$\lim_{k\to\infty} \int_{v\in V^N} \int_{x\in B} \sum_{j\in\mathcal{N}\setminus\{i\}} \mathbb{I}_{x\le b}(v_i - b)\,\phi^k(dx, j, dv) = \int_{v\in V^N} \int_{x\in B} \sum_{j\in\mathcal{N}\setminus\{i\}} \mathbb{I}_{x\le b}(v_i - b)\,\phi(dx, j, dv),$$

since $[0,b] \times \mathcal{N} \times V^N$ is a continuity set of ϕ.

Moreover, just as in the proof of Lemma 1, if (22) holds for all bids b at which there is no atom, then it must hold everywhere. Explicitly, since $H(b) = \phi\left([0,b] \times \mathcal{N} \times V^N\right)$ can have at most countably many mass points, then for every b and for every $\epsilon > 0$, we can find a $b' \in (b, b+\epsilon)$ such that $\phi\left(\{b'\} \times \mathcal{N} \times V^N\right) = 0$ and such that $\phi\left((b,b'] \times \mathcal{N} \times V^N\right) < \epsilon/\overline{v}$. Fixing such an ϵ and b', we know that

$$\begin{aligned}
\int_{v\in V^N} \int_{x\in B} \sum_{j\in\mathcal{N}\setminus\{i\}} \mathbb{I}_{x\le b}(v_i - b)\,\phi(dx, j, dv) &\le \int_{v\in V^N} \int_{x\in B} \sum_{j\in\mathcal{N}\setminus\{i\}} \mathbb{I}_{x\le b'}(v_i - b)\,\phi(dx, j, dv) + \epsilon \\
&\le \int_{v\in V^N} \int_{x\in B} \sum_{j\in\mathcal{N}\setminus\{i\}} \mathbb{I}_{x\le b'}(v_i - b')\,\phi(dx, j, dv) + 2\epsilon \\
&\le \int_{x\in[0,b']} \left(\int_{v\in V^N} \int_{y\in[0,x]} \phi(dy, i, dv) \right) dx + 2\epsilon \\
&\le \int_{x\in[0,b]} \left(\int_{v\in V^N} \int_{y\in[0,x]} \phi(dy, i, dv) \right) dx + 3\epsilon.
\end{aligned}$$

Since this must be true for every $\epsilon > 0$, we conclude that (22) holds at mass points as well, thus completing the proof that D is closed. □

Proof of Lemma 3. Given a feasible solution $\{H_i(\cdot|v)\}$, we can explicitly define a symmetrized solution by

$$\widetilde{H}_i(b|v) = \frac{1}{N!} \sum_{\xi\in\Xi} H_{\xi(i)}(b|\xi(v))$$

and $\widetilde{H}(b|v) = \sum_{i=1}^N \widetilde{H}_i(b|v)$. It is clear that this solution is symmetric, since $\Xi = \{\xi \circ \xi' | \xi \in \Xi\}$ for each $\xi' \in \Xi$, so that

$$\widetilde{H}_{\xi'(i)}(b|\xi'(v)) = \frac{1}{N!} \sum_{\xi\in\Xi} H_{\xi\circ\xi'(i)}(b|\xi\circ\xi'(v)) = \frac{1}{N!} \sum_{\xi\in\Xi} H_{\xi(i)}(b|\xi(v)) = \widetilde{H}_i(b|v).$$

Moreover, $\left\{\widetilde{H}_i\right\}$ will clearly still be increasing and satisfy the probability bounds, since they are just obtained by averaging the H_i. In addition, since μ is symmetric,

$$\begin{aligned}\widetilde{H}(b) &= \int_{v\in V^N} \widetilde{H}(b|v)\,\mu(dv) = \frac{1}{N!}\sum_{\xi\in\Xi}\int_{v\in V^N} H(b|\xi(v))\,\mu(dv)\\ &= \frac{1}{N!}\sum_{\xi\in\Xi}\int_{v\in V^N} H(b|v)\left(\mu\circ\xi^{-1}\right)(dv) = \frac{1}{N!}\sum_{\xi\in\Xi}\int_{v\in V^N} H(b|v)\,\mu(dv)\\ &= \int_{v\in V^N} H(b|v)\,\mu(dv) = H(b),\end{aligned}$$

where $\mu\circ\xi^{-1} = \mu$ by the assumption of exchangeability. Thus, symmetrizing the solution does not change the aggregate distribution of winning bids.

Finally, we verify that (12) will still be satisfied. Note that if $\xi(i) = j$, then

$$\int_{v\in V^N} H_{\xi(i)}(b|\xi(v))\,\mu(dv) = \int_{v\in V^N} H_{\xi(i)}(b|v)\left(\mu\circ\xi^{-1}\right)(dv) = \int_{v\in V^N} H_j(b|v)\,\mu(dv).$$

Thus,

$$\begin{aligned}&\frac{1}{N!}\sum_{\xi\in\Xi}\int_{v\in V^N}(v_i - b)\left(H(b|\xi(v)) - H_{\xi(i)}(b|\xi(v))\right)\mu(dv)\\ &\quad= \frac{1}{N!}\sum_{\xi\in\Xi}\int_{v\in V^N}\left(\xi(v)_{\xi(i)} - b\right)\left(H(b|\xi(v)) - H_{\xi(i)}(b|\xi(v))\right)\mu(dv)\\ &\quad= \frac{1}{N!}\sum_{\xi\in\Xi}\int_{v\in V^N}\left(v_{\xi(i)} - b\right)\left(H(b|v) - H_{\xi(i)}(b|v)\right)\left(\mu\circ\xi^{-1}\right)(dv)\\ &\quad= \frac{1}{N}\sum_{j=1}^{N}\int_{v\in V^N}(v_j - b)\left(H(b|v) - H_j(b|v)\right)\mu(dv),\end{aligned}$$

so that the left-hand side of (12) is simply averaged across bidders. By a similar argument,

$$\frac{1}{N!}\sum_{\xi\in\Xi}\int_{v\in V^N} H_{\xi(i)}(x|\xi(v))\,\mu(dv) = \frac{1}{N}\sum_{j=1}^{N}\int_{v\in V^N} H_j(x|v)\,\mu(dv),$$

so that

$$\int_{v\in V^N}(v_i-b)\left(\widetilde{H}(b|v)-\widetilde{H}_i(b|v)\right)\mu(dv)=\frac{1}{N}\sum_{j=1}^{N}\int_{v\in V^N}(v_j-b)\left(H(b|v)-H_j(b|v)\right)\mu(dv)$$
$$\leq\frac{1}{N}\sum_{j=1}^{N}\int_{v\in V^N}\int_{x=0}^{b}H_j(x|v)\,dx\,\mu(dv)=\int_{v\in V^N}\int_{x=0}^{b}\widetilde{H}_i(x|v)\,dx\,\mu(dv),$$

as desired. □

Proof of Lemma 4. Suppose that we have an inefficient solution. We can then define an alternative solution
$$\widetilde{H}_i(b|v)=\frac{\mathbb{I}_{i\in\arg\max v}}{|\arg\max v|}H(b|v).$$

It is clear that $\widetilde{H}(b|v)=H(b|v)$, so that the aggregate distribution of winning bids is unchanged. We therefore only have to check that $\widetilde{H}$ satisfies (12). By Lemma 3, we can take H_i to be symmetric, so that

$$\int_{v\in V^N}\widetilde{H}_i(b|v)\mu(dv)=\frac{1}{N}\sum_{j=1}^{N}\int_{v\in V^N}\widetilde{H}_j(b|v)\mu(dv)=\frac{1}{N}\int_{v\in V^N}H(b|v)\mu(dv),$$

and the right-hand side of (12) is unchanged. Symmetry implies that

$$\begin{aligned}\int_{v\in V^N}v_i\widetilde{H}_i(b|v)\mu(dv)&=\frac{1}{N}\sum_{j=1}^{N}\int_{v\in V^N}v_j\widetilde{H}_j(b|v)\mu(dv)\\=\frac{1}{N}\int_{v\in V^N}(\max v)\ H(b|v)\mu(dv)&=\int_{v\in V^N}(\max v)\ H_i(b|v)\mu(dv).\end{aligned}$$

The left-hand side of (12) can be rewritten as

$$\int_{v\in V^N}(v_i-b)\left(H(b|v)-H_i(b|v)\right)\mu(dv)+\int_{v\in V^N}(v_i-b)\left(H_i(b|v)-\widetilde{H}_i(b|v)\right)\mu(dv).$$

The second piece reduces to

$$\int_{v\in V^N}(v_i-\max v)\,H_i(b|v)\,\mu(dv)\leq 0,$$

so that (12) is satisfied. □

Proof of Lemma 5. By Lemmas 3 and 4, we can assume that the $\{H(\cdot|v)\}$ are symmetric, efficient, and satisfy (12). Define $\phi \in \Delta\left(V^N \times B\right)$ by

$$\phi(X) = \int_{(v,b)\in X} H(db|v)\,\mu(dv)$$

for measurable sets $X \subseteq V^N \times B$. Since α is measurable, so is the mapping $(\alpha \times \mathbf{I})$: $V^N \times B \to [\underline{w}, \overline{w}] \times B$ defined by $(\alpha \times \mathbf{I})(v,b) = (\alpha(v), b)$. Thus, we can define $\psi \in \Delta([\underline{w}, \overline{w}] \times B)$ by $\psi = \phi \circ (\alpha \times \mathbf{I})^{-1}$. Note that the marginal distribution of ψ on B must be H, i.e., $\psi([\underline{w}, \overline{w}] \times [0,b]) = \phi\left(V^N \times [0,b]\right) = H(b)$ for all $b \in B$, and the marginal distribution of ψ on $[\underline{w}, \overline{w}]$ must be Q, i.e., $\psi([\underline{w}, w] \times B) = \phi(\alpha^{-1}([\underline{w}, w]) \times B) = Q(w)$ for all $\underline{w} \leq w \leq \overline{w}$. Again, the Disintegration Theorem implies that there exists a probability transition kernel $K : [\underline{w}, \overline{w}] \to \Delta(B)$ such that $\psi(dw, db) = K(db|w)\,Q(dw)$, and we can define $H_\alpha(b|w) = K([0,b]\,|w)$.

Now let us argue that $H(\cdot|v)$ satisfies (12) if and only if $H_\alpha(\cdot|w)$ satisfies (14). Given that we can take H to be efficient and symmetric, we can conclude that

$$\int_{v\in V^N} v_i\, H_i(b|v)\,\mu(dv) = \frac{1}{N}\int_{v\in V^N} (\max v)\ H(b|v)\,\mu(dv).$$

Moreover, symmetry implies that

$$\int_{v\in V^N} v_i\, H(b|v)\,\mu(dv) = \int_{v\in V^N} \frac{1}{N}\sum_{j=1}^{N} v_j\, H(b|v)\,\mu(dv).$$

Letting $\gamma(w,b) = w$, the left-hand side of (14) is simply

$$\begin{aligned}
\int_{v\in V^N} \frac{1}{N}\left(\sum_{i=1}^{N} v_i - \max v\right) H(b|v)\,\mu(dv) &= \frac{N-1}{N}\int_{v\in V^N} \alpha(v)\, H(b|v)\,\mu(dv) \\
&= \frac{N-1}{N}\int_{(v,x)\in V^N\times[0,b]} \gamma\left((\alpha\times\mathbf{I})^{-1}(v,x)\right)\phi(dv,dx) \\
&= \frac{N-1}{N}\int_{(w,x)\in[\underline{w},\overline{w}]\times[0,b]} \gamma(w,x)\,\psi(dw,dx) \\
&= \frac{N-1}{N}\int_{w=\underline{w}}^{\overline{w}} w\, H_\alpha(b|w)\,Q(dw),
\end{aligned}$$

where we have used the change of variables formula (cf. Çınlar, 2011, Theorem I.5.2). Thus, we conclude that

$$\begin{aligned}\int_{v\in V^N}(v_i-b)\left(H(b|v)-H_i(b|v)\right)\mu(dv) &= \frac{N-1}{N}\left(\int_{w=\underline{w}}^{\overline{w}} w\,H_\alpha(b|w)\,Q(dw)-b\,H(b)\right)\\ &= \frac{N-1}{N}\int_{w=\underline{w}}^{\overline{w}}(w-b)\,H_\alpha(b|w)\,Q(dw)\end{aligned}$$

and similarly

$$\int_{v\in V^N}\int_{x=0}^{b} H_i(x|v)\,dx\,\mu(dv)=\frac{1}{N}\int_{x=0}^{b}H(x)\,dx=\frac{1}{N}\int_{w=\underline{w}}^{\overline{w}}\int_{x=0}^{b}H_\alpha(x|w)\,dx\,Q(dw),$$

so that (12) and (14) are equivalent.

Finally, we can always define a new solution

$$\widetilde{H}(b|v)=H_\alpha(b|\alpha(v))$$

that induces the same joint distribution of winning bids and $\alpha(v)$, so that it is feasible for the relaxed program and only depends on $\alpha(v)$. □

Proof of Lemma 6. Notice that the only piece of the incentive constraint (15) that depends on how b is correlated with $\alpha(v)$ is through the left-hand side, and all things equal, making $\int_{w=\underline{w}}^{\overline{w}} w\,H_\alpha(b|w)\,Q(dw)$ smaller relaxes the incentive constraint and makes uniform upward deviations less attractive. Let β be defined as in (16). Then $\int_{w=\underline{w}}^{\overline{w}} w\,\widetilde{H}_\alpha(b|w)\,Q(dw)$ is minimized pointwise and for all $b\in B$ by setting $\widetilde{H}_\alpha(b|w)=\mathbb{I}_{b\geq\beta(w)}$. To see this, consider the functions

$$G(w)=\int_{x=\underline{w}}^{w}H_\alpha(b|x)\,Q(dx)$$

and

$$\widetilde{G}(w)=\int_{x=\underline{w}}^{w}\widetilde{H}_\alpha(b|x)\,Q(dx),$$

where b is in the support of $H(db)$. Claim: $G(w)\leq\widetilde{G}(w)$ for all w, i.e., G first-order stochastically dominates $\widetilde{G}$. For if we let $\hat{w}=\beta^{-1}(b)$, then for $w\leq\hat{w}$, we have $H_\alpha(b|w)\leq 1=\widetilde{H}_\alpha(b|w)$, and for $w>\hat{w}$, we have that $\widetilde{G}(w)=H(b)$, and $G(w)\leq H(b)$. Thus,

$$\int_{w=\underline{w}}^{\overline{w}} w\,H_\alpha(b|w)\,Q(dw)=\int_{w=\underline{w}}^{\overline{w}} w\,G(dw)\geq\int_{w=\underline{w}}^{\overline{w}} w\,\widetilde{G}(dw)=\int_{w=\underline{w}}^{\overline{w}} w\,\widetilde{H}_\alpha(b|w)\,Q(dw).$$

We now prove that (17) is equivalent to (12) and (14). Let us first verify that β defined by (16) satisfies (17) for all w. Assuming that $H_\alpha(b|w)$ is monotonic, we must have

$$\begin{aligned}\int_{x=\underline{w}}^{\overline{w}}(x-\beta(w))\,H_\alpha(\beta(w)\,|x)\,Q(dx) &= \int_{x=\underline{w}}^{\overline{w}}(x-\beta(w))\,\mathbb{I}_{\beta(w)\geq\beta(x)}Q(dx)\\ &= \int_{x=\underline{w}}^{w}(x-\beta(w))\,Q(dx)\,.\end{aligned}$$

Note that while $\mathbb{I}_{\beta(w)\geq\beta(x)}$ may be positive over a larger range than $[\underline{w},w]$, β can only be constant on regions where the measure Q places zero probability, so that this equality continues to hold. Moreover,

$$\begin{aligned}\int_{x=\underline{w}}^{\overline{w}}\int_{b=0}^{\beta(w)}H_\alpha(b|x)\,db\,Q(dx) &= \int_{x=\underline{w}}^{\overline{w}}\int_{b=0}^{\beta(w)}\mathbb{I}_{b\geq\beta(x)}db\,Q(dx)\\ &= \int_{x=\underline{w}}^{\overline{w}}\mathbb{I}_{\beta(w)\geq\beta(x)}\,(\beta(w)-\beta(x))\,Q(dx)\\ &= \int_{x=\underline{w}}^{w}(\beta(w)-\beta(x))\,Q(dx)\,.\end{aligned}$$

Thus, (14) implies (17).

Going in the other direction, if we have a monotonic and left-continuous function $\beta : [\underline{w},\overline{w}]\to B$, we can simply define $H_\alpha(b|x)=\mathbb{I}_{b\geq\beta(x)}$. Thus, if we let w be the maximum of $\beta^{-1}([0,b])$, i.e., the largest value such that $\beta(w)\leq b$, which is well defined by left continuity of β, then

$$\begin{aligned}\int_{x=\underline{w}}^{w}(x-\beta(w))\,Q(dx) &= \int_{x=\underline{w}}^{w}(x-b)\,Q(dx)+(b-\beta(w))\,Q(w)\\ &\geq \int_{x=\underline{w}}^{\overline{w}}(x-b)\,H_\alpha(b|x)\,Q(dx)\end{aligned}$$

and similarly

$$\begin{aligned}\int_{x=\underline{w}}^{w}(\beta(w)-\beta(x))\,Q(dx) &= \int_{x=\underline{w}}^{w}(b-\beta(x))\,Q(dx)-(b-\beta(w))\,Q(w)\\ &\leq \int_{x=\underline{w}}^{\overline{w}}\int_{y=0}^{b}H_\alpha(y|x)\,dy\,Q(dx)\,.\end{aligned}$$

Combining these inequalities with (17) yields (14). $\square$