

Extremal Information Structures in the First Price Auction*

Dirk Bergemann[†] Benjamin Brooks[‡] Stephen Morris[§]

First Version: May 2013

Current Version: December 29, 2013

Abstract

We study how the outcomes of a private-value first price auction can vary with bidders' information, for a fixed distribution of private values. In a two bidder, two value, setting, we characterize all combinations of bidder surplus and revenue that can arise, and identify the information structure that minimizes revenue. The extremal information structure that minimizes revenue entails each bidder observing a noisy and correlated signal about the other bidder's value.

In the general environment with many bidders and many values, we characterize the minimum bidder surplus of each bidder and maximum revenue across all information structures. The extremal information structure that simultaneously attains these bounds entails an efficient allocation, bidders knowing whether they will win or lose, losers bidding their true value and winners being induced to bid high by partial information about the highest losing bid. Our analysis uses a linear algebraic characterization of equilibria across all information structures, and we report simulations of properties of the set of all equilibria.

KEYWORDS: First Price Auction, Mechanism Design, Robust Predictions, Private Information, Bayes Correlated Equilibrium.

JEL CLASSIFICATION: C72, D44, D82, D83.

*We gratefully acknowledge financial support from NSF SES 0851200 and ICES 1215808. We would like to thank Satoru Takahashi for many conversations about the present work. We also acknowledge support by the high performance computer center at Princeton University, TIGRESS.

[†]Department of Economics, Yale University, New Haven, U.S.A., dirk.bergemann@yale.edu.

[‡]Department of Economics, Princeton University, Princeton, U.S.A., babrooks@princeton.edu.

[§]Department of Economics, Princeton University, Princeton, U.S.A., smorris@princeton.edu.

1 Introduction

It has been a little over fifty years since Vickrey (1962) presented a general analysis of the first-price sealed-bid auction when values are private, symmetric, and independent, i.e., bidders know their own value for sure and know only the prior distribution about others' preferences. Since that time, the analysis has been extended to some special cases involving complete information or positive correlation, and we review this literature below. However, these analyses have generally been carried out under the maintained assumption that bidders' types are one-dimensional, so that bidders know nothing more than their own value and the prior distribution. A comprehensive analysis of what can happen in a first-price auction when bidders receive more information than their private value has thus far been lacking.

In this paper, we seek to fill the gap by characterizing what can happen in all equilibria for all information structures of a first price auction, holding fixed the distribution of private values of the bidders. Solving for all possible information structures sounds challenging. But we will exploit an argument in Bergemann and Morris (2013a), showing that the set of all joint distributions of values and bids that could arise from a Bayes Nash equilibrium for some information structure is equivalent to the set of a version of incomplete information correlated equilibrium dubbed Bayes correlated equilibrium (BCE). In the case of private values considered in this paper, this reduces to the Bayesian solution of Forges (1993). In the context of a first price auction, a Bayes correlated equilibrium of a first price auction is a joint distribution of bids and values with the property that, conditional on any given value and bid of a bidder, that bidder does not have an incentive to substitute an alternative bid. Once we can identify extremal Bayes correlated equilibria, such as those maximizing revenue of the seller, we can mechanically identify the information structure that would give rise to that distribution as a Bayes Nash equilibrium.

Our results in this preliminary version comprise three parts. First, we consider a single-unit first-price auction with two bidders who have only two possible values but can make any real bid.¹ In this auction, the low valuation bidder is essentially in Bertrand competition with the other low valuation bidder, and always bids his value. The high valuation bidder never has an incentive to bid below the low value, and consequently the outcome is always efficient. Nonetheless, by varying the beliefs of the high valuation bidder, one can achieve substantial variation in revenue and bidders' surplus.

¹To ensure existence, we use an efficient tie breaking rule, and we also rule out the use of dominated strategies in which a player bids above his valuation.

For symmetric distributions over values, we provide a complete characterization of the set of possible bidders' surpluses and revenue that can arise. We construct equilibria along the frontier, and we show that these equilibria achieve bounds on bidders' surplus. For asymmetric distributions we are able to construct equilibria that trace out a frontier which is consistent with evidence from simulations. In the binary valuation example, the worst outcome for bidders is achieved with the complete information type space in which bidders are perfectly informed about all bidders' valuations. This induces Bertrand competition between the bidders, and consequently the winner bids the second highest value. However, the best outcome for bidders (and lowest revenue for the seller) is attained when each bidder observes a noisy and correlated signal of the other bidder's value. The noise and correlation in bidders' signals can be used to maximize the probability that high valuation bidders assign to facing an opponent making a low bid, thus increasing bidder surplus and reducing revenue.

We also analyze the many bidder many value case. In this case, it is possible for the bidders to be strictly worse off than under complete information Bertrand competition. In fact, there is a straightforward lower bound on each bidder's surplus.² Since we assume that dominated strategies are not played, each bidder knows that his opponents will never bid above their respective values. Thus, a worst-case bound on the distribution of opponents' highest bid is the distribution of their highest value. A bidder could always ignore any additional information and best respond to this distribution, and guarantee himself a minimum payoff. Of course, in equilibrium, his opponents cannot in fact always be bidding their values. We show that it is always possible to construct a more complicated joint distribution of values and bids such that each bidder is held down to this lower bound. Moreover, each bidder can be held to this bound for a range of payoffs of the other bidders. Since the BCE that achieves this bound allocates the good efficiently, we therefore also obtain an upper bound on the revenue of the seller. The extremal information structure attaining these bounds entails bidders knowing whether they will win or lose, losing bidders bidding their values and winning bidders receiving partial information about the highest losing bid in a way that maximizes their average bid.

Our third set of results is a computational exploration of the set of Bayes correlated equilibria. We consider discretized models in which players values and bids are drawn from a grid. One of the attractive features of BCE is that, like complete information correlated equilibria, the object of interest is a joint distribution satisfying a series of linear incentive constraints. When combined with a linear objective such as bidder surplus or revenue, this has the structure of a linear program. Thus, we are able to use large-scale linear programming software to compute the set of bidders surpluses and

²We are grateful for Satoru Takahashi for pointing this out to us.

revenue for fairly rich specifications of the model. Using this computational approach, we show what happens to revenue and the shape of the bidder surplus set as the number of valuations grows.

Our focus in the paper is identifying what can happen in equilibrium if bidders know at least their private values but have additional information about other bidders' values. But suppose that we knew that, in addition to knowing their private values, bidders had also observed at least some information about others' values. Thus we had a higher lower bound on the amount of information held by bidders. The general formulation of Bayes correlated equilibrium in Bergemann and Morris (2013a) allows for arbitrary lower bounds on the amount of information held by bidders.³ As the lower bound increases, the set of Bayes correlated equilibria must shrink towards the complete information Nash equilibrium. We illustrate this phenomenon by revisiting the binary-value conditionally-independent model of Fang and Morris (2006). We compute the set of BCE for a range of informativeness of the conditionally independent signals to illustrate the shrinking.

Finally, it is well known that reserve prices and entry fees can boost revenue for a fixed information structure. We consider the effect of these variations in the context of robust predictions. These exercises illustrate the possible use of the methodology of this paper in evaluating how mechanisms perform over a wide range of assumptions about information structures. The results indicate that both reserve prices and entry fees can boost minimum revenue, with reserve prices being substantially more beneficial in this regard. A positive entry fee can also raise the maximum revenue, though reserve prices unequivocally depress the upper bound on revenue. Thus, if the designer is concerned with worst-case performance of the mechanism, then both devices are beneficial and we provide a robust recommendation of which reserve price or fee to use.

A small number of papers have solved for equilibria of private value first price auctions where bidders know their own value but have partial information about other bidders' values. Kim and Che (2004) consider the case where bidders are partitioned into groups, and there is complete information of valuations within elements of the partition, but no information about the valuations of bidders not in the same element of the partition. Equilibria in this setting are inefficient and thus reduce seller revenue. The implications of specific information structures in auctions, and their implication for online advertising market design, are analyzed in recent work by Abraham, Athey, Babaioff, and Grubb (2012) and Celis, Lewis, Mobius, and Nazerzadeh (2012). Both papers are motivated by asymmetries

³Bergemann and Morris (2013b) analyze the Bayes correlated equilibrium in an environment with linear quadratic payoffs and normally distributed uncertainty. There, the lower bound on the information is described by the variance of the noise terms in the signals received by the players.

in bidders' ability to access additional information about the object for sale. Consequently, they examine the role of the distributions of valuations resulting from the private acquisition of data by a single bidder. In particular, Abraham, Athey, Babaioff, and Grubb (2012) focus on second price auctions in a common value environment, while Celis, Lewis, Mobius, and Nazerzadeh (2012) propose an approximately optimal mechanism in a private values model. In a closely related contribution to these two papers, Kempe, Syrganis, and Tardos (2012) study the first-price, common-value auction with asymmetrically informed bidders.

By contrast, we focus on extremal information structures which give rise to extremal values of bidder surplus and revenue and this involves efficient allocations at least in the symmetric cases which we can solve. Fang and Morris (2006) and Azacis and Vida (2013) consider the two bidder, two value, two signal case. We review their results in detail in the next Section. We end up solving for "Bayes correlated equilibria" for first price auctions (because these characterize what can happen in equilibrium for different information structures). Other papers have examined outcome in private value first price auctions under solution concepts weaker than Bayes Nash equilibrium (for a fixed information structure). Battigalli and Siniscalchi (2003) and Dekel and Wolinsky (2003) examine rationalizable outcomes. The set of rationalizable outcomes they consider are neither a subset nor a superset of the BCE outcomes we consider: more restrictive than us, they maintain that all bidders' interim beliefs about opponents' values are the same as the prior distribution; however, our solution concept maintains the common prior assumption. Lopomo, Marx, and Sun (2011) examine bidder collusion in first price auctions. They model bidder collusion as a mechanism design problem, and so the set of attainable equilibria corresponds to the set of communication equilibria in the sense of, e.g., Forges (1993), and they give analytic and computational results showing the impossibility of collusion. Communication equilibrium is another version of incomplete information correlated equilibrium which imposes "truth-telling" constraints (players must have an incentive to truthfully report their types to a mediator) that do not arise in Bayes correlated equilibria.

The rest of the paper proceeds as follows: In Section 2, we give a parameterized two value example to illustrate our result. In Section 3, we describe the general auction model which is used throughout the rest of the paper. In Section 4, we specialize to the case of two bidders and binary valuations. Section 5 gives the general analysis of the bidder surplus lower bounds and revenue upper bound. Section 6 reports the computational results.

2 A Binary Example

We will begin with a simple example to illustrate the role that information can play in a first price auction. There are two bidders, each of whom has either a low (0) or high (1) valuation. Valuations are independently distributed, and are low with probability $\frac{1}{3}$ and high with probability $\frac{2}{3}$. Thus, the probability distributions over value profiles is given by

value	0	1
0	$\frac{1}{9}$	$\frac{2}{9}$
1	$\frac{2}{9}$	$\frac{4}{9}$

where rows correspond to the valuation of bidder 1 and columns correspond to the valuation of bidder 2. With probability $\frac{8}{9}$, at least one bidder has positive value of 1, and thus the ex-ante efficient surplus is $\frac{8}{9}$. The results reported in this Section are a special case of the more general analysis of the two type case in Section 4.

This example gives rise to a surprisingly rich set of outcomes, depending on what information bidders receive about one another's values. The example has been considered by other authors, with various assumptions on the structure of information. For the remainder of this Section, we will summarize some key contributions of this literature to give context and motivation to our own results. For each structure and equilibrium considered, we will characterize the surplus generated for the bidders, with the results being collected in Figure 1.

Complete Information and Zero Information: A Let us begin with two benchmark information structures: complete information and zero information (beyond the common prior). If there is complete information, then conditional on a realization of the profile of valuations, the bidders are essentially in Bertrand competition with one another. The unique undominated equilibrium has each agent always bid 0 except when both have valuations 1, when both bid 1. In this equilibrium, each bidder gets positive surplus of 1 when his valuation is high and the other bidder's valuation is low, which occurs with probability $\frac{2}{9}$. Thus, each bidder's expected surplus is $\frac{2}{9} \approx 0.22$. The complete information point is marked as A in Figure 1. The allocation is efficient in equilibrium, and thus the seller gets expected revenue of $\frac{4}{9}$.

The case where bidders have no information about the other's value has been studied by Maskin and Riley (2003). In this case, each bidder always thinks his opponent has a high value with probability $\frac{2}{3}$ and a low value with probability $\frac{1}{3}$. There is a unique Bayes Nash equilibrium where low valuation




Figure 1: Outcomes of the binary valuation example under different informational assumptions.

bidders always bid 0 and high valuation bidders randomize over bids in the interval $[0, \frac{2}{3}]$ according to a distribution function:

$$F^*(b) = \frac{1}{2} \left(\frac{b}{1-b} \right).$$

A high valuation bidder who expects his opponent to bid 0 with probability $\frac{1}{3}$ and to bid according to distribution F^* with probability $\frac{2}{3}$, is indifferent between all bids in the interval $[0, \frac{2}{3}]$, since bidding any $b \in [0, \frac{2}{3}]$ would give expected payoff:

$$(1-b) \left(\frac{1}{3} + \frac{2}{3} \frac{1}{2} \left(\frac{b}{1-b} \right) \right) = \frac{1}{3} (1-b) + \frac{1}{3} b = \frac{1}{3}.$$

Since a bidder is indifferent to always bidding 0, expected surplus of each bidder is again $\frac{2}{9}$ and expected revenue is again $\frac{4}{9}$. Thus revenue equivalence holds, and this also corresponds to point A in Figure 1.

Informed Bidder vs. Uninformed Bidder: B Now consider an asymmetric information structure. Suppose that bidder 1 knows bidder 2's value, but bidder 2 has no information about bidder 1's value. What happens now? A technical problem arises in solving for equilibria. If a high valuation bidder 1 knows that bidder 2 has a low value and thus will bid 0, then there is an openness problem: bidder 1 has an incentive to bid strictly greater than 0, to guarantee that he will win, but wants the bid to be as small as possible. We will use a standard technical trick—in this example and in the

remainder of the paper—to get around this problem: we assume an efficient tie-breaking rule, so that if both bids are equal, the good is always sold to the highest valuation bidder. Under this assumption, there is a unique equilibrium.

Bidder 1 will bid 0 either if he has a low valuation, or if he has a high valuation, but knows that bidder 2 has a low valuation. If bidder 1 has a high valuation and knows that bidder 2 has a high valuation, then he will bid according to the distribution F^* described above. The low valuation bidder 2 will bid 0; the high valuation bidder 2 will bid 0 with probability $\frac{1}{3}$ and bid according to F^* with probability $\frac{2}{3}$. This strategy profile is optimal, because a high valuation bidder 1 who knows he is facing a high valuation bidder 2, and a high valuation bidder 2 who does not know what type of bidder 1 he is facing, assigns probability $\frac{1}{3}$ to an opponent bidding 0 and $\frac{2}{3}$ to an opponent bidding according to F^* . In this equilibrium, bidder 2's expected surplus remains $\frac{2}{9}$, but bidder 1's surplus increases to $\frac{10}{27}$, since he is indifferent to always bidding 0 when he has a high valuation and would win with probability $\frac{2}{9} + \frac{1}{3}(\frac{4}{9})$ if he did so. This point and the corresponding case where the bidders' roles are reversed are marked as point B in Figure 1. Thus we observe that a bidder does get rent from extra information as long as the other bidder does not have the analogous information.

Conditionally Independent Signals: C, D We have now covered the four cases where each bidder either knows or does not know the other bidder's value. But what about intermediate cases where bidders have some information about the other bidder's value? A natural special case to consider then, in addition to knowing his own value, is a bidder who observes a *conditionally independent* noisy signal of his opponent's value. This problem was studied by Fang and Morris (2006) who suppose, in particular, that each agent observes a binary signal, 0 or 1, and that the "accuracy" of the signal is $q \in [\frac{1}{2}, 1]$, so that with (conditionally independent) probability q , an agent's signal is equal to the other agent's value, while with probability $1 - q$ it is not. In this case, there is also a unique Bayes Nash equilibrium. Low valuation agents always bid 0. High valuation bidders who observe a low signal bid on the interval $[0, \underline{b}]$, where

$$\underline{b} = \frac{2(1-q)^2}{2(1-q)^2 + q},$$

according to c.d.f.

$$F_0(b) = \frac{q}{2(1-q)^2} \frac{b}{1-b}.$$

High valuation bidders who observe a high signal bid on the interval $[\underline{b}, \bar{b}]$, where

$$\bar{b} = \frac{2q^2 + (1+2q)(1-q)\underline{b}}{2q^2 + (1+2q)(1-q)},$$

according to c.d.f.

$$F_1(b) = \frac{(1+2q)(1-q)}{q^2} \frac{b-b}{1-b}.$$

In this case, one can show that the expected surplus of each agent is given by

$$S(q) = \frac{2}{9} \left(\frac{2q(1-q) + q}{2(1-q)^2 + q} \right).$$

Consistent with our earlier calculations, if $q = \frac{1}{2}$, we have the independent private values case and surplus $S(\frac{1}{2})$ equals $\frac{2}{9}$, and if $q = 1$, we have complete information and $S(1)$ equals $\frac{2}{9}$. However, surplus is strictly greater than $\frac{2}{9}$ for all $\frac{1}{2} < q < 1$ and surplus is maximized when $q = \frac{3}{4}$ and expected bidder surplus is $\frac{2}{7} \approx 0.29$. This point is marked as point C in Figure 1.

Azadis and Vida (2013) have generalized the analysis of Fang and Morris (2006), examining what happens when each bidder observes n conditionally independent signals of the other bidder's values, and solving numerically the highest expected surplus that bidders can obtain in this setting in a symmetric n conditionally independent signal model. The surplus is increasing in n and appears to converge to around 0.31 as n increases. These points are marked as D in Figure 1.

Correlated Signals: E,F But the restriction to conditionally independent information structures is also restrictive. Azadis and Vida (2013) give an example showing that if the bidders' signals about the other bidder's value are correlated (conditional on the realized value), then it is possible to construct equilibria which give higher symmetric surplus to the bidders than any conditionally independent signal structure. This is point E in Figure 1. We conclude this section by giving an example of a correlated information structure which, it will turn out, maximizes the sum of the surplus of the two bidders and minimizes revenue.

We maintain the assumption of binary signals. Now suppose that when both bidders have high valuations, the probability that both bidders observe a "correct" signal, i.e., equal to the value of the other bidder, was $2 - \sqrt{3} \approx 0.27$. Suppose that the probability that one bidder observes a correct signal while the other bidder observes an incorrect signal is $\frac{1}{2}(\sqrt{3} - 1) \approx 0.37$. The probability that both observe incorrect signals is 0. Finally, when one bidder has a high valuation and the other's is low, the high valuation bidder observes the correct signal. Now there is an equilibrium where low valuation bidders, and high valuation bidders who observe a low signal, always bid 0. High valuation bidders with a high signal bid on the interval $\left[0, \frac{4-2\sqrt{3}}{3-\sqrt{3}}\right]$ according to the distribution:

$$\hat{F}(b) = \frac{\sqrt{3} - 1}{2(2 - \sqrt{3})} \frac{b}{1-b}.$$

The key feature of the information structure supporting this equilibrium is that a high valuation bidder faces the same distribution of bids independent of the signal he observes in equilibrium. With probability

$$\frac{\frac{2}{9}}{\frac{2}{9} + \frac{2}{9}(\sqrt{3} - 1)} = \frac{\frac{2}{9}(\sqrt{3} - 1)}{\frac{2}{9}(\sqrt{3} - 1) + \frac{4}{9}(2 - \sqrt{3})},$$

he expects his opponent to bid 0, while with complementary probability he expects his opponent to make a strictly positive bid distributed according to \hat{F} . Independent of his signal, he is indifferent between all bids in the interval $\left(0, \frac{4-2\sqrt{3}}{3-\sqrt{3}}\right]$. In this equilibrium, each bidder's surplus is given by $\frac{2}{2\sqrt{3}} \approx 0.58$, which is marked as point F in Figure 1.

In this Section, we have described equilibria for particular information structures and illustrated in Figure 1 how they translate into bidder surplus. In this example, all equilibria will be efficient for all information structures and so Figure 1 also maps out all revenue outcomes. In the next Section, we will describe the methodology for characterizing what happens in all information structures at once. With this methodology, we will be able to use results in Section 4 to show that the set of bidder surplus pairs that could arise in any information structure is given by the shaded area in Figure 1.

3 Model

There are I agents who are bidding for a single unit of a good. In this Section, we will describe the general model for the case of finite values and finite bids. In later Sections, we will sometimes work with exactly this model (i.e., in the computational model) but sometimes work with the analogous definitions with continuum values and/or bids.

We let $V \subseteq \mathbb{R}_+$ be a discrete set of possible values. The valuation of the good to agent i is v_i , where $v_i \in V$. The prior distribution on values is given by $\psi \in \Delta(V^I)$. We will write ψ_i for the marginal distribution over bidder i 's value, and

$$\psi(v_{-i}|v_i) = \frac{\psi(v_i, v_{-i})}{\sum_{\tilde{v}_{-i} \in V^{I-1}} \psi(v_i, \tilde{v}_{-i})}$$

to be the conditional distribution of v_{-i} given v_i .

The agents compete in a first-price auction. Each agent i selects a bid $b_i \in B$, where $B \subseteq \mathbb{R}_+$ is a discrete set of possible bids. For example, B might be the set equal to V , as it is for some of the simulations in Section 6. Note that for economy of notation only, we have chosen the sets of possible values and bids to be the same for all agents. The efficient surplus that could be generated

by allocating the good to the bidder who values it the most is

$$\bar{W} = \sum_{v \in V^I} \psi(v) \max_i \{v_i\}.$$

We will study a standard first price auction. However, we will maintain the efficient tie-breaking rule introduced in the previous Section, so that, in the case of ties, the agent with the higher valuation wins. If multiple bidders with the same valuation also make the same bid, the winner is chosen uniformly among this group. Given the realized profile of bids $b \in B^I$ and values $v \in V^I$, we denote by $H(b, v)$ the set of high bidders under the tie-breaking rule

$$H(b, v) = \{i | b_i \geq b_j \text{ and } v_i \geq v_j \text{ if } b_i = b_j \text{ for all } j \neq i\}.$$

The payoff to bidder i is then given by

$$u_i(b, v) \triangleq \begin{cases} \frac{v_i - b_i}{\#H(b, v)} & \text{if } i \in H(b, v); \\ 0 & \text{if otherwise.} \end{cases}$$

An *information structure* is given by a set of signals for each bidder, T_i , and a probability distribution mapping profiles of values to profiles of signals:

$$\pi : V^I \rightarrow \Delta(T).$$

The information structure (T, π) , combined with the utility functions defined above, parametrizes a game of incomplete information. A strategy for agent i in the incomplete information game (T, π) is a mapping $\beta_i : V \times T_i \rightarrow \Delta(B)$. A strategy is *undominated* if $\beta_i(b_i | v_i, t_i) > 0$ implies $b_i \leq v_i$. A strategy profile $\beta = (\beta_i)_{i=1}^I$ is a *Bayes Nash equilibrium* (BNE) if it is undominated and $\beta_i(b_i | v_i, t_i) > 0$ implies

$$b_i \in \arg \max_{b'_i \in B} \sum_{v_{-i}, b_{-i}, t_{-i}} \psi(v_i, v_{-i}) \left(\prod_{j \neq i} \beta_j(b_j | v_j, t_j) \right) \pi((t_i, t_{-i}) | (v_i, v_{-i})) u_i((b'_i, b_{-i}), (v_i, v_{-i})).$$

The number of games induced by different information structures is large, and for each information structure the set of BNE is a complicated object. We will rely on a simpler set of objects to achieve our characterization. A *decision rule* $\sigma : V^I \rightarrow \Delta(B^I)$ specifies a joint distribution over bids as a function of the profile of valuations. A decision rule is *undominated* if no agent bids above his value, so that

$$\sigma(b | v) > 0 \Rightarrow b_i \leq v_i \text{ for all } i.$$

A decision rule is *obedient* if

$$\begin{aligned} & \sum_{v_{-i}, b_{-i}} \psi(v_i, v_{-i}) \sigma(b_i, b_{-i} | v_i, v_{-i}) u_i((b_i, b_{-i}), (v_i, v_{-i})) \\ & \geq \sum_{v_{-i}, b_{-i}} \psi(v_i, v_{-i}) \sigma(b_i, b_{-i} | v_i, v_{-i}) u_i((b'_i, b_{-i}), (v_i, v_{-i})) \end{aligned} \quad (1)$$

for all i , v_i , b_i and b'_i . A decision rule is a *Bayes correlated equilibrium* (BCE) if it is undominated and obedient.

Our motive for studying BCE is that the set of BCE characterize all behavior that could arise in an BNE for any information structure. A strategy profile β induces a decision rule

$$\sigma(b|v) = \sum_{t \in T} \left(\prod_{j=i}^I \beta_j(b_j | v_j, t_j) \right) \pi(t|v)$$

This motive is summarized in the following result:

Theorem 1 (Equivalence)

Decision rule σ is a undominated Bayes correlated equilibrium if and only if there exists an information structure (T, π) and an undominated Bayes Nash equilibrium β of the auction with information structure (T, π) such that β induces σ .

This is a special case of Theorem 1 of Bergemann and Morris (2013a), with the proviso the added restriction that only undominated actions are played. For completeness, we will summarize some of the main ideas of the proof. If σ is an obedient decision rule, then we can use it to construct an information structure in which bidders are told what they would have played under σ , which we can think of as a "recommended" bid. In other words, $T_i = B$ and $\pi(b|v) = \sigma(b|v)$. We define a strategy profile $\beta : B \rightarrow B$ which is just the identity from recommendations to bids. If bidders $-i$ bid their recommendations, then obedience tells us that bidder i has a weak incentive to bid the recommendation as well.

In the other direction, if β is a BNE for information structure (T, π) , then the decision rule induced by (T, π) and β must be obedient. Conditional on b_i being the random draw from the decision rule, it is as if bidder i learns that one of the t_i must have been realized such that b_i is in the support of $\beta(\cdot | t_i)$. Since b_i is a best response conditional on the realization of any of these signals, it must be a best response to learning that just one was realized.

The characterization in Theorem 1 allows us to reduce our original goal, characterizing all BNE for all information structures, to characterizing the set of obedient decision rules. This is a collection

of families of joint distributions satisfying the linear obedience constraints of (1). Under the decision rule σ , we can define the welfare outcomes:

$$\begin{aligned} U_i(\sigma) &= \sum_{v,b} \psi(v) \sigma(b|v) u_i(v,b) \\ R(\sigma) &= \sum_{v,b} \psi(v) \sigma(b|v) \max_i \{b_i\} \\ TS(\sigma) &= R(\sigma) + \sum_{i \in I} U_i(\sigma), \end{aligned}$$

which are respectively bidder i 's surplus, revenue, and the total surplus. These quantities are also linear functions of σ . Our main results will be a characterization of the set of possible (U_1, \dots, U_I, R) that can arise under an obedient decision rule.

4 The Two Bidder, Two Type Case

4.1 Overview

In Section 2, we gave several examples of BCE in the specialized symmetric model with two bidders and two values. These examples illustrated the richness of how information can impact the outcome of the auction, and also give a sense of the understanding achieved thus far in the literature. We also gave examples of the new equilibria we have discovered. In this Section, we will provide a more systematic study of the BCE in the binary value model. Specifically, we will construct a class of parametrized BCE that attain the bidder surplus frontier depicted in Figure 1. Moreover, we will argue that these equilibria trace out the entire set of surpluses that can be achieved in BCE, i.e., the bounds in Figure 1 are tight. Aside from being interesting in its own right, the two bidder two type example will allow us to illustrate the Bayes correlated equilibrium methodology. At the end, we will also construct a frontier of surpluses for the asymmetric version of this model, although we will not prove tightness.

Relative to the general model of Section 3, we specialize to $I = 2$ and $V = \{\underline{v}, \bar{v}\}$, where $0 \leq \underline{v} < \bar{v}$. Let the symmetric probability distribution ψ over values be given by the following table:

value distribution	\underline{v}	\bar{v}
\underline{v}	$1 - 2p - r$	p
\bar{v}	p	r

(2)

where $p \in [0, \frac{1}{2}]$ is the probability that any one bidder has a high value and the other has a low value

and $r \in [0, 1 - 2p]$. The expected efficient surplus in this example is

$$\bar{W} = \underline{v} + (2p + r)(\bar{v} - \underline{v}).$$

The example in Section 2 corresponds to the special case where $\underline{v} = 0$, $\bar{v} = 1$, $p = \frac{2}{9}$ and $r = \frac{4}{9}$.

4.2 A Class of BCE

We first construct a family of Bayes correlated equilibria. We will illustrate Theorem 1 by showing that these BCE correspond to BNE of particular information structures. In Theorem 2, we will establish that this class of equilibria attain all of the welfare outcomes attainable in any BCE. These equilibria can be thought of as being based on the following information structure: Low valuation bidders get no additional information beyond their value, but high types can receive one of two signals, low L and high H . After receiving a signal L , a bidder knows for sure that the other player could not have value \bar{v} and have received L . Rather, either the other player has a low value, or has a high value and received a high signal H . However, after receiving a high signal, any valuation and signal combination is possible for the other player. This information structure is summarized in (3) below. Since low valuation bidders are in competition with the other low type and play an undominated strategy, a player with value \underline{v} always bids \underline{v} . We will construct an equilibrium in which high bidders who receive low signals always bid \underline{v} , and high bidders who receive high signals always randomize on an interval (\underline{v}, \bar{b}) . Thus, high bidders who bid \underline{v} in equilibrium put zero probability on their high opponent bidding \underline{v} .

Specifically, this class of BCE is parameterized by probabilities x_1, x_2 and c . We will focus on decision rules where a low valuation bidder always bids \underline{v} and, if a high valuation bidder does not bid \underline{v} , then he always independently selects a strictly positive bid from the interval

$$\left(\underline{v}, \frac{r - x_1 - x_2}{c + r - x_1 - x_2} (\bar{v} - \underline{v}) \right]$$

according to c.d.f.

$$F(b) = \frac{c}{r - x_1 - x_2} \frac{b - \underline{v}}{\bar{v} - b}.$$

Thus, for each bidder, there are three cases to consider: (i) he has valuation \underline{v} and bids \underline{v} ; (ii) he has valuation \bar{v} and bids \underline{v} ; and (iii) he has valuation \bar{v} and selects a strictly positive bid according to F . In the following table, we describe the joint distribution over these three (value,bid) pairs for the two

bidders under the parametrized decision rules:

	(\underline{v})	(\bar{v}, L)	(\bar{v}, H)
(\underline{v})	$1 - 2p - r$	$p + x_2 - c$	$c - x_2$
(\bar{v}, L)	$p + x_1 - c$	0	x_2
(\bar{v}, H)	$c - x_1$	x_1	$r - x_1 - x_2$

(3)

Feasibility requires that

$$\begin{aligned}
 x_1 + x_2 &\leq r, \\
 0 \leq c - x_1 &\leq p, \\
 0 \leq c - x_2 &\leq p.
 \end{aligned}
 \tag{4}$$

By construction, a bidder of type (\bar{v}, H) is indifferent between all bids. We must ensure that the obedience constraint of a bidder of type (\bar{v}, L) is satisfied. This requires that the probability that this type assigns to facing a bid of \underline{v} is at least as high as that for type (\bar{v}, H) . This requires that

$$\frac{p + x_i - c}{x_j} \geq \frac{c}{r - x_i - x_j}.
 \tag{5}$$

If these constraints are satisfied, then the surplus of bidder i is

$$U_i = (p + x_i)(\bar{v} - \underline{v})$$

and revenue is

$$\underline{v} + (r - x_1 - x_2)(\bar{v} - \underline{v})$$

Thus the ex ante surplus profile (U_1, U_2) of the two bidders is attainable within this class of BCE if and only if there exist (x_1, x_2, c) satisfying (4) and obedience (5) such that

$$U_1 = (p + x_1)(\bar{v} - \underline{v}) \text{ and } U_2 = (p + x_2)(\bar{v} - \underline{v}).
 \tag{6}$$

For this reason, the quantity x_i is each bidder's excess surplus over p , which a bidder could always obtain by bidding \underline{v} and only winning when the other bidder has valuation \underline{v} .

Figure 1 illustrates a number of equilibria from this class. If $x_1 = x_2 = 0$ and $c = p$, then the model reduces to the no information equilibrium that hits point A. If $x_i = 0$ and $c = x_j = \frac{pr}{p+r}$ (to maintain (5)), then we obtain points B. And finally, if $x_1 = x_2 = c = \sqrt{p(p+r)} - p$, we attain point F, which maximizes the joint surplus of the bidders. We have not previously described the frontier equilibria between points B and F, but they are in this class as well. In fact, they are attained if we set $c = x_i$ and $x_j = \frac{(r-x_i)(p+x_i)}{p+r}$, then we parametrize the portion of the boundary where $x_j \geq x_i$, as x_i ranges from 0 to $\sqrt{p(p+r)} - p$.

4.3 Bounds Attained

Our main result for this section is that the set of surplus pairs attainable in the parametric BCE traces out the entire set of bidder payoffs that can arise in obedient decision rules. We will prove this result in a generalization of the model of Section 3 that allows for a continuum of bids. Specifically, we redefine a decision rule to be a mapping from profiles of valuations to joint cumulative distributions $F(b_1, b_2 | v_1, v_2)$ on V^2 , where $V = [\underline{v}, \bar{v}]$. For the ease of exposition, we will make the following assumptions:

1. The low valuation bidder always bids \underline{v} .
2. The high valuation bidders never both bid \underline{v} .
3. $F(b_1, b_2 | v_1, v_2)$ is differentiable on $[\underline{v}, \bar{v}]^2$ and has a continuous density $f(b_1, b_2 | v_1, v_2)$ on $(\underline{v}, \bar{v})^2$.

Note that (3) does not preclude there being mass along the boundary of $[\underline{v}, \bar{v}]^2$. We simply require that the derivative exists and is unique when taking limits from within the box. In light of these assumptions, we can economize on notation a bit by writing $F(b_1, b_2) = F(b_1, b_2 | \bar{v}, \bar{v})$, $G_1(b_1) = F(b_1, \underline{v} | \bar{v}, \underline{v})$, and $G_2(b_2) = F(\underline{v}, b_2 | \underline{v}, \bar{v})$. We can also define the marginal distribution $F_1(b_1) = F(b_1, \bar{v})$. Since F and G are absolutely continuous, the marginal distributions have densities f_i and g_i such that

$$\begin{aligned} F_1(b) &= F_1(\underline{v}) + \int_{x=\underline{v}}^b f_1(x) dx, \\ G_1(b) &= G_1(\underline{v}) + \int_{x=\underline{v}}^b g_1(x) dx. \end{aligned}$$

Similarly, $F(b_1, b_2)$ can be written as the integral of a F_1 -almost everywhere defined function $\frac{\partial F}{\partial b_1}(b_1, b_2)$, and we define $F_1(b_2 | b_1) = \frac{\frac{\partial F}{\partial b_1}(b_1, b_2)}{f(b_1)}$ where this function exists. All of these objects are defined analogously for bidder 2. Our obedience constraint for the high type can now be stated as

$$b_i \in \arg \max_{b'} (\bar{v} - b') (r F_i(b' | b_i) dF_i(b_i) + p dG_i(b_i)) \text{ for } b_i > \bar{v},$$

where this obedience constraint is required to hold $(r F_i + p G_i)$ -almost everywhere. For $b_i > 0$, $dG_i(b_i) = \frac{\partial G_i(b_i)}{\partial b_i} = g_i(b)$, $dG_i(0) = G_i(0) \equiv \gamma_i$, and analogously for dF_i . Note that we have introduced special notation for $G_i(0)$. For the remainder of this section, when we refer to BCE, the aforementioned structure and (1)-(3) are assumed. We are now ready to state our result.

Theorem 2 (Surplus Boundary of BCE)

The set of bidder surpluses attainable in BCE equals the set of bidder surpluses attainable in BCE in the parameterized class.

The proof is rather involved and will proceed via a series of Lemmas. Here is an overview of the argument. We will consider the problem of maximizing the weighted sum of bidders' surpluses $S(\lambda) = S_1 + \lambda S_2$ over all BCE, where $\lambda \in [-1, 1]$. The reason for restricting to $\lambda \geq -1$ will be seen over the course of the argument; for $\lambda < -1$, we will always have a corner solution, and so this is without loss of generality. The case when $\lambda > 1$ is symmetric. The parametric equilibria of the previous Subsection attain a particular function $S(\lambda)$, which will be characterized in Lemma 1.

While BCE have a relatively tractable linear structure, there are a lot of obedience constraints, and making sure that all of them are satisfied is a challenging task. Our shortcuts are threefold: (i) we will define one-dimensional marginal distributions of the BCE that are sufficient to pin down bidder surplus; (ii) we will simply drop a bunch of the obedience constraints, in particular we drop all of bidder 1's obedience constraints, and any obedience constraint of bidder 2 that involves deviating to less than the recommendation; and (iii) we will aggregate up the remaining constraints into a restriction on the lower dimensional objects specified in (i). Step (i) will be accomplished in Lemma 2 and discussion preceding it. Steps (ii) and (iii) will be the subject of Lemma 3. These three tasks leave us with a relaxed version of our original weighted surplus maximization problem, in which we maximize over the low dimensional objects. We will then solve the relaxed problem for the maximum surplus in the direction $(1, \lambda)$, and verify that one of our parametrized BCE achieves the same level in this direction.

To start, let us define the level $S(\lambda)$ that we will show is tight.

Lemma 1 (Parametric frontier)

In the direction $(1, \lambda)$ in bidder surplus space, the parametric BCE attain at least the level

$$S(\lambda) = (\bar{v} - \underline{v}) \left[p + \frac{(1 + \lambda^2)(p + r)}{4} \right]. \quad (7)$$

Proof. We will look at the subset of the parametrized equilibria for which $c = x_1$ and $x_1 \geq x_2$. As such, the second obedience constraint in (1) implies the first, as

$$\frac{p}{x_2} \geq \frac{p}{x_1} \geq \frac{p + x_2 - x_1}{x_1},$$

so we will impose that the second constraint binds, so that

$$x_1^2 = (p + x_2 - x_1)(r - x_1 - x_2) \implies x_1 = \frac{(r - x_2)(p + x_2)}{p + r}. \quad (8)$$

Given that x_i gives bidder i 's excess surplus, we simply wish to maximize

$$x_1 + \lambda x_2,$$

over non-negative x_1 and x_2 that satisfy (8) and $x_1 - x_2 \leq p$. Substituting in (8) and taking a first-order condition, we obtain

$$\begin{aligned} x_1(\lambda) &= \frac{(1 - \lambda^2)(p + r)}{4}, \\ x_2(\lambda) &= \frac{\lambda(p + r) + r - p}{2}, \end{aligned}$$

and hence

$$\begin{aligned} S_1(\lambda) &= \left(p + \frac{(1 - \lambda^2)(p + r)}{4} \right) (\bar{v} - \underline{v}), \\ S_2(\lambda) &= \frac{(1 + \lambda)(p + r)}{2} (\bar{v} - \underline{v}), \end{aligned}$$

and the sum is

$$\begin{aligned} S(\lambda) &= S_1(\lambda) + \lambda S_2(\lambda) \\ &= (\bar{v} - \underline{v}) \left[p + \frac{(1 - \lambda^2)(p + r)}{4} + \lambda \frac{(1 + \lambda)(p + r)}{2} \right] \\ &= (\bar{v} - \underline{v}) \left[p + \frac{(1 + \lambda^2)(p + r)}{4} \right]. \end{aligned}$$

Since we showed above that these equilibria trace out a closed convex set, and since all of these boundary equilibria solve the relaxed problem, we have shown that these equilibria in fact attain the entire boundary. ■

We will now maximize the weighted sum

$$\int_{b_1=\underline{v}}^{\bar{v}} (\bar{v} - b_1) (r F_1(b_1|b_1) dF_1(b_1) + p dG_1(b_1)) + \lambda \int_{b_2=\underline{v}}^{\bar{v}} (\bar{v} - b_2) (r F_2(b_2|b_2) dF_2(b_2) + p dG_2(b_2)),$$

subject to the obedience constraint, and show that this maximum value coincides with $S(\lambda)$. For the next step, we define the lower dimensional objects we will be working with. This is step (i) of our aforementioned program. Let us define

$$h_i(b) = r F_i(b|b) f_i(b) + p g_i(b)$$

to be the density of bidder i winning the auction with a bid of $b > \underline{v}$. We also introduce the shorthand notation

$$\gamma_i = G_i(\underline{v}).$$

We can therefore simplify the objective function to

$$\begin{aligned} & \int_{b=\underline{v}}^{\bar{v}} (\bar{v} - b) (h_1(b) + \lambda h_2(b)) db + (\bar{v} - \underline{v}) p (\gamma_1 + \lambda \gamma_2) \\ = & \int_{b=\underline{v}}^{\bar{v}} (H_1(b) + \lambda H_2(b)) db + (\bar{v} - \underline{v}) p (\gamma_1 + \lambda \gamma_2). \end{aligned}$$

The equality follows from integration by parts, where

$$H_i(b) = \int_{x=\underline{v}}^b h_i(x) dx,$$

is the cumulative probability of bidder i winning with a bid less than b and greater than \underline{v} . Thus, our problem has been reduced from characterizing a two-dimensional distribution and two one-dimensional distributions to just two one-dimensional distributions and the real valued parameters γ_1 and γ_2 . Note that H_i is not a probability distribution, in the sense that $H_1(\bar{v}) + H_2(\bar{v}) = r + p(2 - \gamma_1 - \gamma_2)$, which is the total probability of the high valuation bidders winning with bids in the half open interval $(\underline{v}, \bar{v}]$. The following Lemma uses the differentiable structure of the decision rule.

Lemma 2 (Shared surplus from $b_i > \underline{v}$)

In any BCE, we must have $H_1(b) = H_2(b) = H(b)$. Hence, the bidders derive equal surplus from positive bids.

Proof. Our obedience constraint implies that for $b > 0$, we must have the following first-order condition hold at $b' = b$:

$$\begin{aligned} 0 &= (\bar{v} - b) r \left. \frac{\partial F_i(x|b)}{\partial x} \right|_{x=b} - f_i(b) - r F_i(b|b) f_i(b) - p g_i(b) \\ &= (\bar{v} - b) r f(b, b) - h_i(b) \\ \implies h_i(b) &= (\bar{v} - b) r f(b, b), \end{aligned}$$

which implies that $h_1(b) = h_2(b) = h(b)$. In other words, both bidders must be equally likely to win with a bid $b > \underline{v}$. As a result, we also have

$$H_1(b) = H_2(b) = H(b)$$

Since the surplus a bidder derives from positive bids is precisely $\int_{b=\underline{v}}^{\bar{v}} H(b) db$, we have the result. ■

Consequently, we can simplify our objective to

$$(1 + \lambda) \int_{b=\underline{v}}^{\bar{v}} H(b) db + (\bar{v} - \underline{v}) p (\gamma_1 + \lambda \gamma_2).$$

With this result, we are down to a single one-dimensional distribution and the two real parameters. Our next result combines steps (ii) and (iii). We will use a subset of the obedience constraints to define a law of motion on $H(b)$. This law of motion must hold for any obedient decision rule, and it is the only consequence of obedience that we will retain in our relaxed problem.

Lemma 3 (Aggregated obedience constraint)

In any BCE, $H(b)$ must obey the law of motion

$$H(b) \leq \frac{1}{\bar{v} - b} \int_{x=\underline{v}}^b H(x) dx + p \left(\frac{b - \underline{v}}{\bar{v} - b} \gamma_2 + G_1(b) - \gamma_1 \right) \quad (9)$$

Proof. The non-local obedience constraint tells us that

$$\begin{aligned} (\bar{v} - b') (r F_i(b'|b) f_i(b) + p g_i(b)) &\leq (\bar{v} - b) (r F_i(b|b) f_i(b) + p g_i(b)) \\ &= (\bar{v} - b) h(b) \end{aligned}$$

which implies that

$$r (F_i(b'|b) - F_i(b|b)) f_i(b) \leq \frac{b' - b}{\bar{v} - b'} h(b). \quad (10)$$

For $b = \underline{v}$, we have the similar constraint that

$$(\bar{v} - b') (r F_i(b'|\underline{v}) F_i(\underline{v}) + p \gamma_i) \leq (\bar{v} - \underline{v}) p \gamma_i,$$

so

$$r F_i(b'|\underline{v}) F_i(\underline{v}) \leq \frac{b' - \underline{v}}{\bar{v} - b'} p \gamma_i.$$

Moreover,

$$\begin{aligned} H(b) &= \int_{x=\underline{v}}^b (r F_1(x|x) f_1(x) + p g_1(x)) dx \\ &= r \int_{x=\underline{v}}^b \int_{y=\underline{v}}^x dF_1(y|x) f_1(x) dx + p (G_1(b) - \gamma_1) \\ &= r \int_{x=\underline{v}}^b \int_{y=\underline{v}}^x f(y, x) dy dx + r \int_{x=\underline{v}}^b dF_1(\underline{v}|x) f_1(x) dx + p (G_1(b) - \gamma_1), \end{aligned} \quad (11)$$

but

$$\begin{aligned}
r \int_{x=\underline{v}}^b \int_{y=\underline{v}}^x f(y, x) dx &= r \int_{y=\underline{v}}^b \int_{x=y}^b f(y, x) dy dx \\
&= r \int_{x=\underline{v}}^b (F_2(b|x) - F_2(x|x)) f_2(x) dx \\
&\leq \int_{x=\underline{v}}^b \frac{b-x}{\bar{v}-b} h(x) dx \\
&= \frac{1}{\bar{v}-b} \int_{x=\underline{v}}^b H(x) dx.
\end{aligned} \tag{12}$$

The first line is Fubini's Theorem, the second follows from the definition of the conditional distributions, the third is substituting in (10), and the last line is integration by parts. Similarly,

$$\begin{aligned}
r \int_{x=\underline{v}}^b dF_1(\underline{v}|x) f_1(x) dx &= r F_2(b|\underline{v}) F_2(\underline{v}) \\
&\leq \frac{b-\underline{v}}{\bar{v}-b} p \gamma_2.
\end{aligned} \tag{13}$$

Substituting (12) and (13) into (11) and simplifying gives the desired result. ■

We will now consider a relaxation of our original problem:

$$\max_{H, G_1, G_2} (1 + \lambda) \int_{b=\underline{v}}^{\bar{v}} H(b) db + (\bar{v} - \underline{v}) p (\gamma_1 + \lambda \gamma_2)$$

subject to $G_i \leq 1$, $H \leq \frac{r+p(2-\gamma_1-\gamma_2)}{2}$ and (9). The remainder of the proof will characterize the solution, and verify that this solution is attained by one of the parametric equilibria.

Proof of Theorem 2. Note that $G_1(b)$ only enters our problem through (9), and by making $G_1(b)$ as large as possible, we relax the constraint. Hence, it is without loss of generality to set $G_1(b) = 1$ for all $b > \underline{v}$. Next, since $\lambda \geq -1$, the objective is weakly increasing in $H(b)$, so we should make $H(b)$ as large as possible. As long as $H(b) < \frac{r+p(2-\gamma_1-\gamma_2)}{2}$, it must be that (9) is binding. Otherwise we could increase $H(b)$ by setting it equal to the RHS of (9), which weakly raises $H(b)$ for all values of b , and consequently relaxes further (9).

To summarize thus far, we know that there is a solution to the relaxed problem in which $G_1(b) = 1$ and H solves (9) as an equality until it hits $r+(2-\gamma_1-\gamma_2)$. We are therefore left with a two-parameter objective, which is a function of γ_1 and γ_2 . H solves the following differential equation

$$(\bar{v} - b) dH = 2H + p(\gamma_2 + \gamma_1 - 1)$$

subject to the initial condition $H(\underline{v}) = 0$. Note that we must impose $\gamma_2 + \gamma_1 \geq 1$, since otherwise H is decreasing at 0, and there is no way to satisfy the obedience constraints. This differential equation has the general solution:

$$H(b) = \frac{(b - 2\bar{v})bp(\gamma_1 + \gamma_2 - 1) + C}{2(\bar{v} - b)^2}.$$

Evaluating at $b = \underline{v}$, we determine that

$$C = -p(\gamma_1 + \gamma_2 - 1)\underline{v}(2\bar{v} - \underline{v}),$$

and therefore

$$H(b) = \frac{p(\gamma_2 + \gamma_1 - 1)}{2(\bar{v} - b)^2} [b(2\bar{v} - b) - \underline{v}(2\bar{v} - \underline{v})].$$

This function hits $\frac{r+p(2-\gamma_1-\gamma_2)}{2}$ at precisely

$$\bar{b} = \bar{v} - (\bar{v} - \underline{v}) \sqrt{\frac{p(\gamma_2 + \gamma_1 - 1)}{p + r}}.$$

Therefore,

$$\begin{aligned} \int_{b=\underline{v}}^{\bar{v}} H(b) db &= (\bar{v} - \bar{b})H(\bar{b}) - p[(\bar{b} - \underline{v})\gamma_2 + (\bar{v} - \bar{b})(1 - \gamma_1)] + (\bar{v} - \bar{b})H(\bar{b}) \\ &= (\bar{v} - \bar{b})[r + p(2 - \gamma_1 - \gamma_2)] - p[(\bar{b} - \underline{v})\gamma_2 + (\bar{v} - \bar{b})(1 - \gamma_1)] \\ &= (\bar{v} - \bar{b})(p + r) - p(\bar{v} - \underline{v})\gamma_2, \end{aligned}$$

where we have used (9) to calculate $\int_{x=\underline{v}}^{\bar{b}} H(b) db$, and the fact that $H(\bar{b}) = \frac{r+p(2-\gamma_1-\gamma_2)}{2}$. Plugging back into our objective, we find

$$(1 + \lambda) [(\bar{v} - \bar{b})(p + r) - p(\bar{v} - \underline{v})\gamma_2] + p(\bar{v} - \underline{v})(\gamma_1 + \lambda\gamma_2).$$

Hence, the derivative with respect to γ_1 is

$$-(1 + \lambda)(p + r) \frac{\partial \bar{b}}{\partial \gamma_1} + p(\bar{v} - \underline{v}).$$

It is clear that \bar{b} is in fact decreasing in γ_1 , so it is always optimal to set γ_1 as large as possible, namely $\gamma_1 = 1$. Substituting in the definition of \bar{b} , our relaxed problem has simplified to

$$\max_{\gamma_2} (\bar{v} - \underline{v}) \left[(1 + \lambda) \sqrt{\gamma_2 p(p + r)} + p(1 - \gamma_2) \right].$$

The first order condition for γ_2 is

$$(1 + \lambda) \sqrt{p(p + r)} \frac{1}{2\sqrt{\gamma_2}} = p,$$

which has the unique solution

$$\gamma_2^*(\lambda) = \frac{(1 + \lambda)^2 (p + r)}{4p}.$$

Note that we are ignoring extra constraints on $\gamma_2^*(\lambda)$, like that it be in $[0, 1]$. However, ignoring these constraints simply relaxes the problem further. Hence, the optimal objective of the relaxed problem is

$$S^*(\lambda) = (\bar{v} - \underline{v}) \left(p + \frac{(1 + \lambda)^2 (p + r)}{4} \right).$$

Since this coincides with the $S(\lambda)$ we derived for the parametric equilibria, we are done. ■

To be clear, we have shown that our parametric class of BCE attain bounds on how far one can go in a given direction $(1, \lambda)$. The fact that the entire surplus set is attained is a consequence of the convexity of the set of obedient decision rules; we could take weighted sums of our parametric equilibria to construct any point within the surplus set.

The argument we have presented is notable as a demonstration of one way in which the BCE solution concept can be used to generate a robust prediction for a game. We will use the same method in Section 5, simultaneously constructing equilibria and showing that these equilibria attain bounds from a relaxed problem, to give a generalized lower bound on bidder surplus in the first price auction.

4.4 Beyond the Symmetric Common Prior

We can generalize the distribution over values used this far, see (2), to an asymmetric prior:

value distribution	\underline{v}	\bar{v}
\underline{v}	$1 - p_1 - p_2 - r$	p_2
\bar{v}	p_1	r

(14)

where $p_1 + p_2 + r \leq 1$, and p_i gives the probability that bidder i has a high value and bidder $j \neq i$ has a low value. For concreteness, suppose that $p_1 > p_2$. We can still construct the parametric equilibria introduced before, however the feasibility and obedience constraints are now

$$\begin{aligned} x_1 + x_2 &\leq r \\ 0 \leq c - x_1 &\leq p_1 \\ 0 \leq c - x_2 &\leq p_2 \end{aligned}$$

By construction, a bidder of type (\bar{v}, H) is indifferent between all bids. We must ensure that the obedience constraint of a bidder of type (\bar{v}, L) is satisfied. This requires that the probability that this

type assigns to facing a bid of \underline{v} is at least as high as that for type (\bar{v}, H) . This requires that

$$\frac{p_i + x_i - c}{x_j} \geq \frac{c}{r - x_i - x_j}.$$

In the symmetric model, we were able to construct frontier equilibria that trace out the frontier of a convex set of bidder surpluses, as depicted in Figure 1. However, with $p_1 > p_2$, the set of payoffs traced out by the parametric class is no longer convex. Let us consider the Pareto frontier of this class. Since c is an upper bound on x_i , the excess surpluses, we would like to make c as large as possible that is still consistent with feasibility. The smallest such value is $c = \max\{x_1, x_2\}$. Set $x_2 = x_1 = c = x = \sqrt{p_2(p_2 + r)} - p_2$, so that the second obedience constraint binds. At this point, the first obedience constraint is necessarily slack, since

$$\frac{p_1}{x} > \frac{p_2}{x} = \frac{x}{r - 2x}.$$

Now consider moving along the Pareto frontier increasing x_2 . In this case, $c = \max\{x_1, x_2\} = x_2$, the first obedience constraint is slack, and we solve

$$\frac{p_2}{x_1} = \frac{x_2}{r - x_1 - x_2} \implies x_1 = p_2 \left(\frac{p_2 + r}{x_2 + p_2} - 1 \right),$$

so that x_1 is a convex function of x_2 . Hence, the Pareto frontier is convex at this point and the set of surpluses attained by this class is non-convex.

Of course, we can always achieve any point in the convex hull of this set by taking linear combinations of the three-type equilibria. But this discussion suggests that for asymmetric distributions, there are other classes of equilibria which might further push out the frontier of bidder surplus. Indeed, this is the case, and we will now construct such equilibria, which require the high value bidders to receive *three* signals:

	\underline{v}	(\bar{v}, L)	(\bar{v}, M)	(\bar{v}, H)
\underline{v}	$1 - p_1 - p_2 - r$	p_2	0	0
(\bar{v}, L)	f	0	a_2	b_2
(\bar{v}, M)	$p_1 - f$	a_1	d	$c - b_2$
(\bar{v}, H)	0	b_1	$c - b_1$	e

As before, the bidder with valuation \underline{v} always bids \underline{v} , as does the bidder of type (\bar{v}, L) . In addition, the equilibrium has two bid cutoffs $\underline{v} \leq b_M \leq b_H \leq \bar{v}$. We will construct the equilibrium so that (\bar{v}, M) is indifferent to bidding anywhere in (\underline{v}, b_M) and (\bar{v}, H) is indifferent to bidding in (b_M, b_H) . This pins

down the shape of the bid distribution for types that bid on these regions, in a manner analogous to the three type construction. As such, the bidders' payoffs will be:

$$S_i = (\bar{v} - \underline{v})(p_i + a_i) + (\bar{v} - b_M)c.$$

We will impose the following incentive constraints, as indicated in this constraint incidence chart:

		Indifference		
Type		\underline{v}	b_M	b_H
Player 1	(\bar{v}, L)	x	x	-
	(\bar{v}, M)	x	x	-
	(\bar{v}, H)	-	x	x
Player 2	(\bar{v}, L)	x	x	x
	(\bar{v}, M)	x	x	x
	(\bar{v}, H)	x	x	x

The interpretation is that if there is an "x" in a box, that means that the player with the given type is indifferent to that bid. There are in fact eight incentive constraints, since the two players' type (\bar{v}, H) indifference between b_M and b_H are redundant. In addition, there is an adding up constraint

$$a_1 + a_2 + b_1 + b_2 + 2c + d + e = r.$$

Since there are ten variables (two bid cutoffs and eight probability variables) the system is underidentified by one parameter, which we will leave as f . We spare the reader the derivation, and jump straight to the results. It turns out that without using the adding up constraint, we can cleanly eliminate all of the variables except for f and c :

$$\begin{aligned} a_1 &= \frac{p_2(p_1-f)}{f-p_2}, & a_2 &= \frac{f(p_1-f)}{f-p_2}, \\ b_1 &= \frac{cp_2(f-p_2)}{f(p_1+p_2)-f^2-p_2^2}, & b_2 &= \frac{c(f-p_2)}{p_1-p_2}, \\ d &= \frac{f(p_1-f)^2}{(f-p_2)^2}, & e &= \frac{c^2(f-p_2)^2}{(p_1-p_2)(f(p_1+p_2)-f^2-p_2^2)}, \\ b_M &= \underline{v} + (\bar{v} - \underline{v}) \frac{p_1-f}{p_1-p_2}, & b_H &= \underline{v} + (\bar{v} - \underline{v}) \frac{c(f-p_2)^2 + (p_1-f)(f(p_1+p_2)-f^2-p_2^2)}{c(f-p_2)^2 + (p_1-p_2)(f(p_1+p_2)-f^2-p_2^2)}. \end{aligned}$$

The variable c can then be obtained from the adding up equation. In the interest of saving trees, we will not print the closed form expression of c as a function of f , although it can be provided upon request. In Figure 2, we show the three-type equilibria that have a non-convex frontier, and also these four-type equilibria that convexify the frontier. These equilibria are meant to demonstrate that it is




Figure 2: The set of bidder surpluses for an asymmetric example. In solid lines are the three type equilibria constructed previously. Note the non-convexity of the northeast frontier. The four type equilibria in dashed lines push out the set.

possible to extend our analysis to more general models, although the complexity of such constructions grows quite quickly, even for the modest generalization of going from three types to four. In the next section, we will construct a rich class of equilibria of the many player, many valuation model, which are analytically tractable in spite of the generality.

5 Many Values: Lower Bounds of Bidder Surplus and Upper Bounds of Revenue

We now return to the general model of Section 3 with many valuations and many types. The goal of this section is to give a tight characterization of the lower limit of bidder surplus and the upper limit of revenue over all BCE. This characterization consists of a theoretical bound and the construction of equilibria that attain the bound. For most of this section, we will study a model where the set of valuations is equal to the set of bids, i.e., $B = V$. As the analysis of Section 4 has illustrated, our

techniques can handle models with a continuum of bids and we will give a continuum bid example. However, our bounds and constructed equilibria will only involve bids that are in the support of bidders' values. Thus, it is expositionally easier and without loss of generality for us to stay within this simpler framework.

The following analysis is closely connected to ?, where we analyze the limits of third-degree price discrimination induced with respect to the private information. The problem of choosing a bid when facing a fixed distribution of opponents' bids is formally equivalent to the pricing problem of a monopolist when facing a fixed distribution of buyers' valuations. In the auction setting, the bid b represents a cutoff bid of others below which the buyer receives the good with surplus $v - b$, whereas in the monopoly setting, the price p is a cutoff valuation above which the monopolist makes a sale and earns profit $p - c$, where c is the cost of production. ? study the possible effects of information on the monopolist's pricing problem, which is closely related to how information can influence bidding behavior when facing a fixed distribution of opponents' bids. The key difference which greatly complicates the auction problem is that unlike an exogenous distribution of consumer valuations, the distribution of opponents' bids is endogenous, and must itself be generated by an obedient decision rule. Nonetheless, arguments similar to those used in characterizing price discrimination will be employed to construct equilibria attaining our lower bound on bidder surplus.

5.1 Bounds on Bidder Surplus and Revenue

Throughout our analysis, we have assumed that bidders do not use dominated strategies in which they bid above their own values. Hence, in any BCE, bidders must believe that whatever their opponents' bidding strategy, it is bounded above by the conditional distribution of their opponents' values. Since the surplus a bidder can achieve by best responding is decreasing (in the sense of first-order stochastic dominance) in the distribution of opponents' bids, the equilibrium surplus must be weakly better than what a bidder could have attained if opponents bid their values. This property can be exploited to give a straightforward bound on the surplus a bidder can guarantee himself in equilibrium. In particular, the epistemic result of Theorem 1 shows that a perfectly legitimate interpretation of a BCE is that bidders receive information and best respond to their opponents' behavior conditional on this information. However, a bidder could always *ignore* this extra information, and simply best respond to this worst-case conjecture about their opponent's bids, and be guaranteed a minimum surplus.

Formally, the equilibrium bid distribution faced by a bidder with valuation v_i is bounded above by $\psi(v_{-i}|v_i)$. We define $\underline{U}_i(v_i)$ to be the maximum surplus a bidder could obtain when opponents bid

their values, and $\underline{b}_i(v_i)$ is a bid that attains it,

$$\begin{aligned} \underline{U}_i(v_i) &= \max_{b_i \in B} (v_i - b_i) \sum_{\left\{v_{-i} \in V^{I-1} \mid b_i > \max_{j \neq i} v_j\right\}} \psi(v_{-i} | v_i). \\ \underline{b}_i(v_i) &\in \arg \max_{b_i \in B} (v_i - b_i) \sum_{\left\{v_{-i} \in V^{I-1} \mid b_i > \max_{j \neq i} v_j\right\}} \psi(v_{-i} | v_i). \end{aligned}$$

This gives an ex ante lower bound on bidder surplus for bidder i of

$$\underline{U}_i = \sum_{v_i \in V} \psi_i(v_i) \underline{U}_i(v_i). \quad (15)$$

Recall that the efficient surplus is given by:

$$\bar{W} = \sum_{v \in V^I} \psi(v) \max_i \{v_i\}.$$

Since bidders have to receive at least \underline{U}_i in equilibrium, the maximum revenue the seller could receive is the total feasible surplus \bar{W} minus the sum of these bounds for each player. Hence, an upper bound on revenue \bar{R} is the efficient surplus minus the surplus that each agent can guarantee himself,

$$\bar{R} = \bar{W} - \sum_{i=1}^I \underline{U}_i. \quad (16)$$

We summarize these results in the following proposition:

Proposition 1 (Surplus and revenue bounds)

In any BCE, bidders must receive a surplus weakly greater than \underline{U}_i , and revenue can be no more than \bar{R} .

5.2 Equilibria that Attain the Bounds

Our next result constructs a class of BCE which attain the bounds from the previous section. Before launching into the details, we will give some intuition for the construction. Recall, we wish to construct an efficient equilibrium in which bidders are held down to \underline{U}_i . This BCE implicitly has an information structure, which we can think of as sending bidders two kinds of signals. The first signal L is sent to bidders who do not have a strictly highest valuation, and in equilibrium, bidders with valuation v_i who receive L as their signal *know* that someone else is bidding $b_j \geq v_i$. Hence, it is a best reply to bid $b_i = v_i$, and this is what happens.

If bidder i does not receive the L signal, he receives the signal H and a further instruction to bid a particular value b . H means that other bidders' values are strictly less than v_i . As a consequence, other bidders received signal L and are all bidding their values. Now, in order for the equilibrium to hang together, it must be that the instructed bid b is always greater than $\max_{j \neq i} v_j$. We show that in fact there is a way to "suggest" bids to the H bidder so that this constraint is satisfied. Moreover, we can structure recommendations so that the bidder is always indifferent between following the recommendation b and bidding $\underline{b}_i(v_i)$. Since the latter strategy would simply result in the payoff \underline{U}_i , this shows that the bidders are held down to their lower bound surplus.

Theorem 3 (Tightness of bounds)

There is an efficient undominated BCE where the bidder surplus lower bound \underline{U}_i and revenue upper bound \bar{R} are simultaneously attained.

Proof. We will construct an obedient decision rule that attains the bounds. We can divide the set of value realizations into subsets based on who is the winner. Let

$$X_i(v_i) = \{\tilde{v} \in V^I \mid \tilde{v}_i = v_i > \tilde{v}_j \ \forall j \neq i\}$$

and

$$Y = \left\{ \tilde{v} \in V^I \mid \# \left(\arg \max_i \tilde{v}_i \right) \geq 2 \right\}.$$

In plain words, X_i is the set of profiles of valuations on which bidder i has the strictly highest valuation, and Y is the set of profiles on which at least two bidders tie for highest valuation. The decision rule for $v \in Y$ will be

$$\sigma(b|v) = \begin{cases} 1 & \text{if } b = v; \\ 0 & \text{if otherwise.} \end{cases}$$

Now consider $v \in X_i(v_i)$. Let $\underline{b}_i(v_i)$ be defined as in the previous Subsection, which clearly is not more than v_i . We will define the decision rule so that for $v \in X_i(v_i)$, $\sigma(b|v) > 0$ only if $b_{-i} = v_{-i}$, i.e., bidders other than i always bid their values. If bidder i were just told that the valuation profile is in $X_i(v_i)$, then it must be that $\underline{b}_i(v_i)$ maximizes bidder i 's conditional surplus, since $\underline{b}_i(v_i)$ is obviously superior to any bid $b' \geq v_i$, which generates non-positive surplus, whereas virtue of the fact that $X_i(v_i) \neq \emptyset$, we know that it is possible for bidder i to achieve strictly positive surplus. Moreover, the relative probabilities of events of the form

$$E_i(b_i) = \left\{ v_{-i} \in V^{I-1} \mid b_i > \max_{j \neq i} v_j \right\},$$

for $b_i < v_i$ are the same conditional on $X_i(v_i)$ as they are unconditional on $X_i(v_i)$.

We will construct the decision rule on $X_i(v_i)$ so that (1) bidder i is always indifferent to bidding $\underline{b}_i(v_i)$, and (2) $b_i \geq v_j$ for all $j \neq i$. Together with the efficient tie-breaking rule, this implies that bidder i always wins the auction on event $X_i(v_i)$. We can denumerate the values that arise as $w = \max_{j \neq i} v_j$ for $v \in X_i(v_i)$ as

$$w_0 > \cdots > w_K,$$

and further divide the set $X_i(v_i)$ into

$$X_i(v_i, w) = \left\{ v \in X_i(v_i) \mid w = \max_{j \neq i} v_j \right\}.$$

We can define

$$\xi_0(w) = \sum_{v \in X_i(v_i, w)} \psi(v)$$

to be the distribution over $\max_{j \neq i} v_j$ restricted to $X_i(v_i)$. As such,

$$\underline{b}_i(v_i) \in \arg \max_b (v_i - b) \sum_{w \leq b} \xi_0(w).$$

We will inductively define α_k to be the solution to

$$(v_i - \underline{b}_i(v_i)) \sum_{w \leq \underline{b}_i(v_i)} \alpha_k \xi_k(w) = (v_i - w_k) \left(\sum_{w < w_k} \alpha_k \xi_k(w) + \xi_k(w_k) \right)$$

for $k \geq 0$ and set

$$\mu_k(w) = \begin{cases} \xi_0(w) \prod_{l < k} (1 - \alpha_l), & \text{if } w = w_k; \\ \xi_0(w) \alpha_k \prod_{l < k} (1 - \alpha_l), & \text{if } w < w_k; \\ 0, & \text{otherwise;} \end{cases}$$

and

$$\xi_{k+1}(w) = \begin{cases} \xi_0(w) \prod_{l \leq k} (1 - \alpha_l) = (1 - \alpha_k) \xi_k(w), & \text{if } w < w_k; \\ 0, & \text{if } w \geq w_k. \end{cases}$$

Finally, for $v \in X_i(v_i, w_l)$, we define

$$\sigma(b|v) = \begin{cases} \prod_{l=0}^{k-1} (1 - \alpha_l), & \text{if } b_i = w_k \text{ and } b_{-i} = v_{-i}; \\ \alpha_k \prod_{l=0}^{k-1} (1 - \alpha_l), & \text{if } b_i = w_l \text{ with } l < k \text{ and } b_{-i} = v_{-i}; \\ 0, & \text{if } b_i = w_l \text{ with } l > k \text{ or } b_{-i} \neq v_{-i}. \end{cases}$$

To be clear, this defines $\sigma(b|v)$ for all $v \in V^I$, because each v is in some $X_i(v_i, w_k)$ for some v_i and w_k . Also, bidder i always wins when $v \in X_i(v_i, w_k)$, since the probability of recommendation w_l is zero when $w_l < w_k$.

We will show that this is a well-defined and obedient decision rule. Let $P[k]$ be the statement that

$$\underline{b}_i(v_i) \in \arg \max_b (v_i - b) \sum_{w \leq b} \xi_k(w).$$

By assumption, $P[0]$ is true. We will argue that $P[k] \implies P[k+1]$. Observe that if $\xi_k(w_k) > 0$, we have

$$\alpha_k = \frac{(v_i - w_k) \xi_k(w_k)}{(v_i - \underline{b}_i(v_i)) \sum_{w \leq \underline{b}_i(v_i)} \xi_k(w) - (v_i - w_k) \sum_{w < w_k} \xi_k(w)}$$

which is less than 1, since

$$(v_i - \underline{b}_i(v_i)) \sum_{w \leq \underline{b}_i(v_i)} \xi_k(w) \geq (v_i - w_k) \sum_{w \leq w_k} \xi_k(w),$$

and if $\xi_k(w_k) = 0$, $\alpha_k = 0$. Hence, $\alpha_k \in [0, 1]$, and ξ_{k+1} is well defined and proportional to ξ_0 below w_k . Clearly,

$$\begin{aligned} (v_i - \underline{b}_i(v_i)) \sum_{w \leq \underline{b}_i(v_i)} \xi_{k+1}(w) &= (v_i - \underline{b}_i(v_i)) \sum_{w \leq \underline{b}_i(v_i)} \xi_0(w) \prod_{l \leq k} (1 - \alpha_l) \\ &\geq (v_i - b) \sum_{w \leq b} \xi_0(w) \prod_{l \leq k} (1 - \alpha_l) \\ &\geq (v_i - b) \sum_{w \leq b} \xi_{k+1}(w), \end{aligned}$$

which proves $P[k+1]$. Finally, if $\underline{b}_i(v_i) = w_{k^*}$ and $\xi_{k^*}(w_{k^*}) > 0$, then we must have $\alpha_{k^*} = 1$, so the algorithm has to terminate at some $\hat{k} \leq k^*$. If $v \in X_i(v_i, w_k)$ with $k \geq \hat{k}$, then

$$\begin{aligned} \sum_{l=0}^k \sigma(w_l, v_{-i} | v) &= \prod_{m=0}^{k-1} (1 - \alpha_m) + \sum_{l=0}^{k-1} \alpha_l \prod_{m=0}^{l-1} (1 - \alpha_m) \\ &= (1 - \alpha_{k-1} + \alpha_{k-1}) \prod_{m=0}^{k-2} (1 - \alpha_m) + \sum_{l=0}^{k-2} \alpha_l \prod_{m=0}^{l-1} (1 - \alpha_m) \\ &= \prod_{m=0}^{k-2} (1 - \alpha_m) + \sum_{l=0}^{k-2} \alpha_l \prod_{m=0}^{l-1} (1 - \alpha_m) \\ &\quad \vdots \\ &= 1 \end{aligned}$$

and hence, the decision rule is well-defined.

Next we will show obedience. Observe that if $b_i = v_i$, then it must be either (1) $v \in Y$ or $v \in X_j(v_j)$ for some j . Either way, if $\sigma(b_i, b_{-i} | v) > 0$, it must be that there exists $j \neq i$ such that $b_j \geq b_i = v_i$.

Hence, the bidder's conditional surplus must be zero, and any deviation at which he would win with positive probability requires $b' > v_i$, which would lead to non-positive surplus. Hence, the decision rule is obedient whenever $b_i = v_i$.

If $b_i = w_k < v_i$, then it must be that $v \in X_i(v_i, w_l)$ for $l \leq k$. As such,

$$\sum_{v_{-i}, b_{-i}} \psi(v_i, v_{-i}) \sigma(b_i, b_{-i} | v_i, v_{-i}) u_i((b'_i, b_{-i}), (v_i, v_{-i})) = (v_i - b'_i) \sum_{w \leq b'_i} \mu_k(w),$$

since the probability of getting recommendation w_k when $v \in X_i(v_i, w_l)$ is precisely $\mu_k(w_l)$. Obedience will follow from our final claim, which is that

$$w_k \in \arg \max_b (v_i - b) \sum_{w \leq b} \mu_k(w)$$

which is a consequence of

$$\begin{aligned} (v_i - w_k) \sum_{w \leq w_k} \mu_k(w) &= (v_i - \underline{b}_i(v_i)) \sum_{w \leq \underline{b}_i(v_i)} \mu_k(w) \\ &= (v_i - \underline{b}_i(v_i)) \sum_{w \leq \underline{b}_i(v_i)} (1 - \alpha_k) \xi_k(w) \\ &\geq (v_i - b) \sum_{w \leq b} (1 - \alpha_k) \xi_k(w) \text{ for all } b < w_k \\ &= (v_i - b) \sum_{w \leq b} \mu_k(w) . \end{aligned}$$

This also shows that $\underline{b}_i(v_i)$ is always a weak best reply, and hence

$$\begin{aligned} &\sum_{v_i, b_i, v_{-i}, b_{-i}} \psi(v_i, v_{-i}) \sigma(b_i, b_{-i} | v_i, v_{-i}) u_i((b_i, b_{-i}), (v_i, v_{-i})) \\ &= \sum_{v_i} \sum_{\underline{b}_i(v_i), v_{-i}, b_{-i}} \psi(v_i, v_{-i}) \sigma(b_i, b_{-i} | v_i, v_{-i}) u_i((\underline{b}_i(v_i), b_{-i}), (v_i, v_{-i})) \\ &= \sum_{v_i} \psi_i(v_i) \underline{U}_i(v_i) \\ &= \underline{U}_i. \end{aligned}$$

And finally, the decision rule is efficient, so $R = \bar{R}$. ■

Note that even though we have restricted bidders to using bids in the support of valuations, they have no incentive to bid outside this set. The only situation in which a bidder ties at $b < v_i$ is when $v_j < v_i$, and hence the efficient tie-breaking negates any incentives to bid more.

The decision rule constructed in the proof of Theorem 3 has a special feature that the rule is "compartmentalized" to the sets $X_i(v_i)$. In particular, we could modify the decision rule on any one

such set, without affecting the obedience constraints for the rest of the equilibrium, as long as for $v \in X_i(v_i)$, $\sigma(b_i, v_{-i} | v_i, v_{-i}) > 0$ implies that $b_i \geq \max_{j \neq i} v_j$. For example, on such decision rule would define for $v \in X_i(v_i)$,

$$\sigma(b|v) = \begin{cases} 1, & \text{if } b_i = \max_{j \neq i} v_j, b_{-i} = v_{-i}; \\ 0, & \text{otherwise.} \end{cases}$$

The signal structure has a simple interpretation: For bidder i who receives the signal H , he also learns the profile of other bidders' valuations v_{-i} , and consequently is able to bid the second-highest value. As a result, bidder i must obtain the same payoff as he would get in the complete information equilibrium. However, bidders $j \neq i$ are still receiving \underline{U}_j . As such, we have the following Corollary:

Corollary 1

There exist undominated BCE that hold player i to \underline{U}_i , while giving other bidders a range of surpluses which extends at least up to the bidder surplus attained under complete information.

5.3 An Example with a Continuum of Values

The bounds readily generalize to models with a continuum of values. As an example, let us consider a setting with 2 agents whose values are drawn from the interval $[0, 1]$ according to the cumulative distribution function $F(v) = v^\alpha$. In this case, bidders conjecture that in the worst case their opponents bids are also distributed with the cumulative distribution b^α . The symmetric lower bound on surplus for a bidder with valuation v is given by

$$\underline{W}(v) = \max_{b \in [0,1]} (v - b) b^\alpha.$$

The maximum is attained by setting

$$\underline{b}_i(v) = \frac{\alpha}{1 + \alpha} v,$$

and thus

$$\begin{aligned} \underline{W}(v) &= \left(v - \frac{\alpha}{1 + \alpha} v \right) \left(\frac{\alpha}{1 + \alpha} v \right)^\alpha \\ &= \frac{1}{1 + \alpha} \left(\frac{\alpha}{1 + \alpha} \right)^\alpha v^{1+\alpha}. \end{aligned}$$

The ex ante lower bound on surplus is then

$$\begin{aligned} \underline{W} &= \int_{v=0}^1 \frac{1}{1+\alpha} \left(\frac{\alpha}{1+\alpha} \right)^\alpha v^{1+\alpha} \alpha v^{\alpha-1} dv \\ &= \int_{v=0}^1 \left(\frac{\alpha}{1+\alpha} \right)^{\alpha+1} v^{2\alpha} dv \\ &= \frac{1}{1+2\alpha} \left(\frac{\alpha}{1+\alpha} \right)^{\alpha+1}. \end{aligned}$$

One can show that expected surplus in this example is $\overline{W} = \frac{2\alpha}{1+2\alpha}$ and thus the upper bound on revenue is

$$\overline{R} = \left(1 - \left(\frac{1}{1+\alpha} \right) \left(\frac{\alpha}{1+\alpha} \right)^\alpha \right) \frac{2\alpha}{1+2\alpha}.$$

The example with zero information beyond the common prior has a unique Bayes Nash equilibrium, in which the revenue is $\left(\frac{\alpha}{1+\alpha} \right) \frac{2\alpha}{1+2\alpha}$. In the special case of $\alpha = 1$ (a uniform distribution), we have that $\underline{U}(v) = \frac{1}{4}v^2$, and thus the minimum ex ante bidder surplus is $\underline{U} = \frac{1}{12}$, and the total surplus is $TS = \overline{W} = \frac{2}{3}$. The upper bound on revenue is $\overline{R} = \frac{1}{2}$, and by contrast the revenue in the BNE is $R = \frac{1}{3}$. As $\alpha \rightarrow 0$, the upper bound on revenue converges to 0, but the ratio of the upper bound on revenue to the BNE revenue converges to $+\infty$:

$$\lim_{\alpha \rightarrow 0} \frac{\left(1 - \left(\frac{1}{1+\alpha} \right) \left(\frac{\alpha}{1+\alpha} \right)^\alpha \right) \frac{2\alpha}{1+2\alpha}}{\left(\frac{\alpha}{1+\alpha} \right) \frac{2\alpha}{1+2\alpha}} = \lim_{\alpha \rightarrow 0} \frac{1 + \alpha - \left(\frac{\alpha}{1+\alpha} \right)^\alpha}{\alpha} = +\infty.$$

Note that the conjectured behavior that generates the bounds is far from equilibrium: Each bidder best responds to the belief that others will bid their values. But we will generally have the best response $\underline{b}_i(v)$ be strictly less than v , so this conjecture must turn out to be false in equilibrium. However, we will see in the next Subsection that there is a BCE in which bidders are held down to the bound, and moreover this equilibrium is efficient.

We briefly illustrate what the constructed BCE looks like in the case where there are two bidders whose values are uniformly distributed on $[0, 1]$; this corresponds to the case of $\alpha = 1$. The construction uses an example ?, adapted to the auction setting. We first draw two values uniformly from $[0, 1]$, where v_h and v_l denote the highest and lowest values. In addition, we also draw a "tentative recommendation" $r \in \left[\frac{v_h}{2}, v_h \right]$ for each bidder according to the cumulative distribution

$$H(r) = \frac{v_h}{2r - v_h} e^{1 - \frac{v_h}{2r - v_h}},$$

and write

$$h(r) = \frac{4v_h(v_h - b)}{(2r - v_h)^3} e^{1 - \frac{v_h}{2r - v_h}}$$

for the corresponding density. Note the identity:

$$H(r) = h(r) \frac{(2r - v_h)^2}{4(v_h - r)}.$$

After drawing the values, bidders are informed whether they are the higher value ("winning") bidder or the lower value ("losing") bidder. In addition, the winning bidder will observe a "final recommendation" set equal to $b = \max\{r, v_l\}$.

Under this information structure, there is an equilibrium where the bidder with valuation bids v_l and the high valuation player bids the recommendation b . Let us verify that this is incentive compatible. Conditional on observing b with valuation v_h , it could have been that either (1) $b = r \geq v_l$ and v_l is uniformly distributed on $[0, r]$, or (2) $b = v_l > r$, so there is a mass point on b . Hence, the cumulative distribution of v_l conditional on observing the signal b is

$$F(x|b, v_h) \propto \begin{cases} h(b) \frac{x}{v_h}, & \text{if } x < b; \\ h(b) \frac{b}{v_h} + H(b) \frac{1}{v_h}, & \text{otherwise.} \end{cases}$$

Since $F(x|b, v_h)$ is constant for all $x \geq b$, $x = b$ dominates all bids greater than b . The payoff from a bid of $x = b$ is

$$\begin{aligned} (v_h - b) F(b|b, v_h) &\propto \frac{h(b)}{v_h} (v_h - b) \left(b + \frac{(2b - v_h)^2}{4(v_h - b)} \right) \\ &= \frac{h(b)}{v_h} (v_h - b) \frac{4b(v_h - b) + (4b^2 - 4bv_h + v_h^2)}{4(v_h - b)} \\ &= \frac{h(b)}{v_h} \frac{v_h^2}{4}. \end{aligned}$$

We have to verify that there are no bids $x < b$ that are better. For $x < b$, we have

$$(v_h - x) F(x|t, v_h) \propto \frac{h(b)}{v_h} (v_h - x) x.$$

This function is concave and has a maximum on $[0, v_h)$ at $x = \frac{v_h}{2}$, at which point the payoff is equal to the payoff from $x = b$.

Hence, we conclude that following the recommended bid is incentive compatible, and moreover that no matter the recommendation, bidders are indifferent to a bid of $\frac{v_h}{2}$, which is the bid that guarantees them the lower bound payoff. We conclude that in this BCE, bidders are held down to the lower bound.

6 Computational Results

In this Section, we push beyond our analytical results to obtain a more complete picture of how information influences the outcome of a first-price auction. Availing ourselves of the linear structure of BCE and the associated linear programs, we are able to solve for extremal BCE of discretized examples in which values and bids are confined to a grid. We use these simulated equilibria to investigate the limiting characteristics of BCE as the support of values and/or bids converges to the continuum, thus approximating a continuous distribution of values. The simulations tell us a great deal about how information can influence the auction for general distributions of values. In particular, the simulations deliver numerical bounds on revenue and surplus, they show us the shape of the bidder surplus frontier and the revenue-total surplus frontier, and the simulations allow us to assess the impact of entry fees and reserve prices. We also will look more closely at the binary valuation model, and calculate comparative statics as we vary bidders' information.

Though we will mainly report summary statistics about the computed BCE, we in fact solve for the entire joint distribution of bids and values. In our experience, this distribution has a detailed and complex structure. This is at least partly a consequence of the multiplicity of BCE attaining a particular objective (e.g. maximizing revenue). As a result, we tend to see a mashup equilibrium which is a convex combination of the various optimal distributions. However, we can distill some general features of the BCE, such as which constraints *must* be binding to obtain a particular objective, and which value/bid profiles must be in the support of the BCE.

Let us briefly comment on our methods. As stated above, we solve for extremal BCE for fine discretizations of the bid and value space. Even modest discretizations will result in very large linear programs, with hundreds of thousands or even millions of variables. For example, with 2 players, 35 bids and 35 values per player (the largest model we will report below), we already have approximately 1.5 million possible combinations of bids and values (35^4). The number of constraints is smaller, but still substantial: Approximately 40,000 incentive constraints per player, and 1225 probability constraints. In fact, these raw numbers are exponential in the number of players and polynomial in the numbers of values and bids. As a result, we will restrict our numerical investigations to models with two bidders. We can also economize substantially by ruling out dominated behavior, such as bidding above one's value, or by taking advantage of symmetry. But even with these restrictions, the complexity is daunting.

To tackle such large problems, we have made use of the state-of-the-art linear programming package

CPLEX. We have written programs in C++ that construct and solve discretized models. In fact, these programs can be used to compute the BCE of any game, not just first-price auctions, as long as the game is suitably specified within the object model of our program. For analysis, we have also developed graphical tools in MATLAB to explore and understand the computed BCE. These tools allow us to see the conditional joint distribution of actions and the state from a particular player’s perspective, and we can tell at a glance which incentive constraints are binding for which types. We plan to make further use of these tools in subsequent research.

Many of our simulation results are reported for uniform distributions over values. In fact, we have run similar simulations using various distributions for values, both independent and correlated. The stylized facts that we highlight seem quite robust to alternate specifications of symmetric distributions. Indeed, we know that the bounds on minimum bidder surplus and maximum revenue hold true for any number of bidders and distribution of values, even for asymmetric distributions.

6.1 Limit of Revenue and Bidder Surplus

We start our numerical analysis with a description of the ranges of revenue and bidder surplus that can be achieved in some BCE. Our analytic results show that the lower bound of bidder surplus and the upper bound of revenue are determined by Satoru’s bound. In the other direction, how large can bidder surplus be, and how low can revenue fall? We computed minimum revenue and maximum bidder surplus for a range of examples with two bidders in which there are 35 bids and the number of valuations varies between 2 and 35, with bids and values evenly spaced in $[0, 1]$. The joint distribution values is uniform. Figure 3 gives the results, along with Satoru’s bound. To be clear, by bidder surplus we mean the sum of the individual surpluses of the two bidders.

As a preliminary observation, we know that one information structure which is always feasible is complete information, in which the realization of values is publicly known. Bertrand competition results in both players bidding the second-highest valuation, and the high-valuation player wins the tie break. The outcome is efficient, so total surplus is the expected highest value, and revenue is the expected second-highest value. So as a rough estimate, we know that the range of feasible revenues will contain the expected lowest value, and the range of feasible total bidder surpluses will contain the expected difference between the highest and lowest values.

Indeed, this is precisely what we see, and moreover it appears that the revenue and surplus bounds track their complete information quantities as we increase the number of values. Bounds on bidder surplus are given by red lines and bounds on revenue are given by blue lines. In dashed lines are the

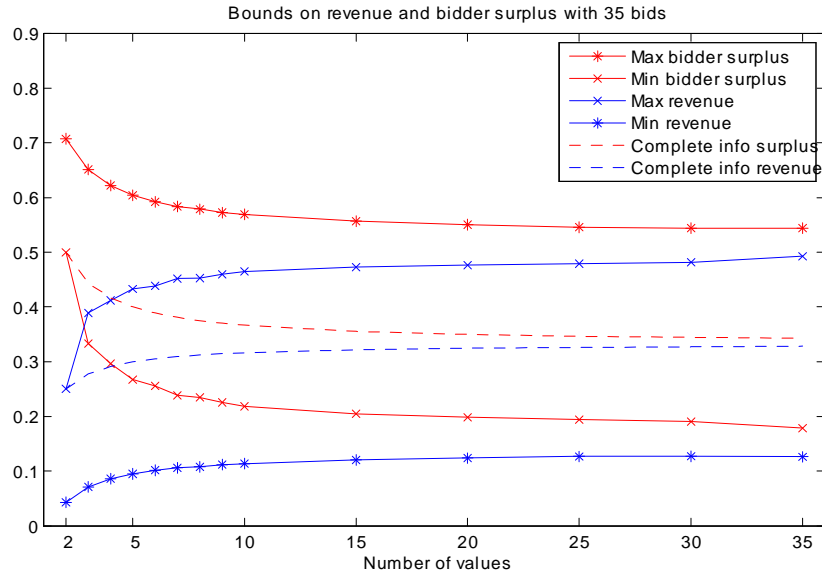


Figure 3: Bounds on revenue and bidder surplus. Maximum revenue and minimum bidder surplus are achieved in a common and efficient equilibrium.

expected difference between highest and lowest values in red, and expected lowest value in blue. With two values, we observe that the expected highest value is 0.75 and the expected lowest is 0.25, giving us the minimum bidder surplus and maximum revenue of 0.25 and 0.5 respectively. Also, as we know from our earlier analysis, maximum bidder surplus and minimum revenue are substantially higher and lower, respectively.

For more than two values, we see the bounds on revenue and surplus roughly tracking the complete information benchmarks. For 35 values, the largest model computed, revenue lies in $[0.13, 0.49]$ and bidder surplus lies in $[0.18, 0.54]$, and appears to be converging to values nearby. At the very least, these numbers serve as a sanity check: revenue and surplus seem to converge in a reasonable manner, and track the relevant benchmarks in the complete information case. In addition, the simulations tell us that revenue and bidder surplus cut wide swaths around their complete information benchmarks, indicating that information retains a powerful affect on the outcome even when there is a large number of values. For minimum bidder surplus and maximum revenue, this is expected given that Satoru's bound can be attained in efficient equilibria. In fact, this implies that the maximum surplus and minimum revenue can be achieved with the same equilibrium.

Maximum bidder surplus and minimum revenue are more mysterious, and we are actively engaged in understanding the extremal information structures that achieve these goals. We conjecture that

the BCE that minimize revenue are efficient, so that maximum bidder surplus and minimum revenue are achieved by a common equilibrium. In fact, if one adds to this picture the sum of maximum bidder surplus and minimum revenue, and also the sum of minimum bidder surplus and maximum revenue, the two sums virtually coincide for each number of values, and therefore also coincide with the expected highest value. We have omitted this additional information for the sake of clarity. In the next subsection, we will present further evidence for this conjecture.

6.2 Feasible Surplus Pairs

Maximizing and minimizing joint bidder surplus and revenue represent just a handful of directions in which we could look for extremal BCE. What are the extremal equilibria in other directions? For example, what is the set of all feasible bidder surplus pairs, or the set of all revenue-total surplus pairs? In the binary case, we gave a complete description of these sets with a continuum of bids. In Figure 4, we show the theoretical frontier and the simulated frontier of bidder surplus for a discretized binary value example with 50 bids. Bidders' values are independent, and each has a valuation of zero with probability $\frac{1}{3}$. As anticipated, the two sets are quite close. Though the theoretical frontier has a kink at the northeast corner, the computed example is somewhat rounded over. By always bidding 0, a bidder is able to guarantee themselves a minimum payoff of $\frac{2}{9}$, when one has a high value and the other has a low value, which gives us the southwest corner of the set.

For models with many values, we do not have a concise characterization of the entire frontier. However, there are certain features that we know should be present. First, it is possible to hold each bidder to Satoru's bound while inducing a range of surpluses for the other bidder. Thus, there should be a right angle at the southwest corner of the set of bidder surpluses. The flats emanating from the corner extend at least to the complete information payoff, where the high valuation bidder wins and pays the lowest value. But beyond these features, there is much we do not know about the set. Do the flats extend beyond the complete information case, as in the binary values example? Also, what does the rest of the frontier look like? In the binary example, every equilibrium is efficient. But which frontier equilibria are efficient with many values? Is there a kink at the point that maximizes total bidder surplus?

In Figure 5, we show the frontier of the set of bidder surpluses for a model with 20 valuations and 20 bids. A couple of features immediately stand out. As expected, we see the corner on the southwest frontier of the set. The complete information payoffs of approximately 0.175 are in the interior of the set, though the "flats" extend well beyond this point, to about 0.235. We are intrigued by this

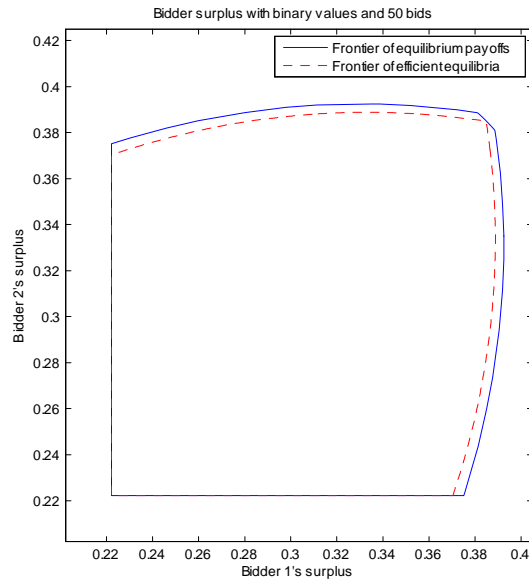


Figure 4: The set of bidder surpluses arising in Bayes correlated equilibria when $\underline{v} = 0$, $\bar{v} = 1$, $p = \frac{2}{9}$, and $r = \frac{4}{9}$. In blue is the simulated boundary, and in red is the theoretical frontier.

feature: It is possible to give each player payoffs well above complete information while holding the other player to their minimum payoff.

In addition to the unconstrained frontier, we also show the frontier of bidder surpluses that can be attained in an efficient equilibrium. For these equilibria, the good is always allocated to the bidder with the highest valuation. It appears that the maximum possible total bidder surplus can be achieved in an efficient equilibrium. Though not clear from this picture, further computational results indicate that efficiency is necessary to maximize total bidder surplus. Thus, the efficient point on the 45 degree line also represents a revenue minimizing equilibrium.

In the simulation, it appears that there is a region of the northeast frontier of bidder surpluses that is also efficient. However, we have observed that the region of overlap tends to shrink as we make the bid grid finer. Our conjecture is that this is the same rounding over effect that we see in Figure 4, and that in a model with a continuum of bids, the only efficient point on the frontier of bidder surpluses is the symmetric point.

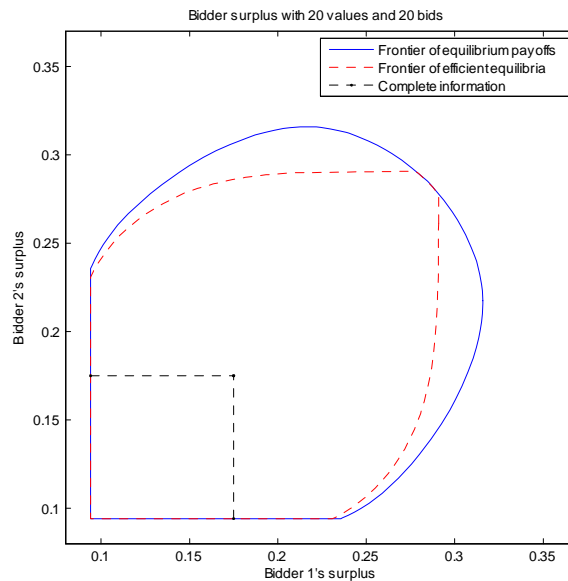


Figure 5: The set of bidder surpluses in Bayes correlated equilibria with many values. Each bidder can be held to lower bound surplus for a range of surpluses of the other bidder.

6.3 Reserve Prices and Entry Fees

Throughout the paper, we have been exploring a single auction format: The first-price auction with an efficient tie-breaking rule. Much of the auction design literature is concerned with finding the optimal mechanism for a particular environment. While beyond the scope of the present work, we are interested in characterizing robust optimal auctions that have favorable revenue performance even when there is large uncertainty about players' beliefs. As a preliminary step in this direction, we explore the impact of two simple extensions to the first-price auction format: reserve prices and entry fees.

Suppose the seller can either set a minimum admissible bid r or charge bidders a flat fee e for submitting a positive bid. It is well known that these devices can enhance revenue in first-price auctions, for a fixed information structure (see Milgrom and Weber (1982), and references therein). On some occasions there is a preference for one device or another. Second-price auctions with reserve prices are in fact optimal auctions in regular symmetric case, see Myerson (1981).

In Figure 6, we report the maximum and minimum revenue for a model with 25 values and 25 bids, as the reserve price and entry fee range from 0 to 1. The results are striking. Positive reserve prices and entry fees both raise the minimum possible revenue. Positive entry fees raise the maximum revenue, whereas maximum revenue is monotonically decreasing in the reserve. For a reserve price of

1, only the high type participates, and bids his value. On the other hand, with an entry fee of 1, no type submits positive bids and revenue is zero.

For both pictures, we have included a “complete information” analogue to calibrate our expectations of what should happen. For reserve prices, the dashed curve is the exact outcome of a model with the given reserve price and 25 values and 25 bids, where the profile of valuations is common knowledge. We see that it roughly tracks the center of the revenue range, and follows the hump shape of the minimum revenue curve.

For entry fees, we do not know the equilibrium for complete information. Instead, we have plotted revenue from a second-price auction with an entry fee, for which it is well known how to construct the equilibrium: There is a cutoff type $\bar{v}(e)$ who participates, solving $\bar{v}(e)F(\bar{v}(e)) = c$, and values above $\bar{v}(e)$ bid their values. For the discretized model, $\bar{v}(e)$ jumps discontinuously and the revenue curve has a saw tooth pattern that lies within the revenue bounds. For clarity, we have plotted an idealized version with 1,000 valuations. Again, the second-price auction curve roughly tracks the center of the range, and the hump shape of the bounding curves.

Let us summarize. Positive entry fees robustly raise revenue, in the sense that a positive entry fee raises both the minimum and maximum revenue over all BCE. Positive reserve prices also raise the minimum possible revenue, but lower the maximum. If a designer had worst-case preferences, and sought to maximize the minimum revenue over all BCE, then a large reserve price would be advisable: A reserve price of 0.58 raises the minimum revenue to 0.40, from 0.13. For entry fees, a fee of 0.33 raises the maximum revenue from 0.49 to its maximum of 0.57, and a fee of 0.25 raises minimum revenue to its maximum of 0.25.

6.4 The Role of Additional Information

For our last numerical exercise, we study the effect of giving the bidders extra information beyond their private value. The set of BCE is a “robust prediction” that encompasses all possible Bayesian equilibria that could result when players have access to additional signals beyond their private values. Some of these equilibria require bidders to be well-informed about the other bidders’ values, and some equilibria require less-informed bidders. As an example, in the asymmetric extremal BCE of the binary valuation model, one bidder receives more information from the “suggested” bid than the other bidder. If we impose lower bounds on players’ information, by forcing them to observe additional signals, we will invariably rule out equilibria in which players’ signals need to be uninformative.

Specifically, let us consider the model of Fang and Morris (2006), in which valuations are in $\{0, 1\}$.





Figure 6: Bounds on revenue as we introduce reserve prices and entry fees for a model with 25 values and 25 bids, where values are uniformly distributed.

In addition to learning v_i , each player receives a noisy signal s_i about the other bidder's valuation. Conditional upon v_j , $Pr(s_i = v_j | v_j) = \sigma$, where $\sigma \geq 0.5$. Thus, σ represents the informativeness of the signal. We look for BCE when players observe at least (v_i, s_i) .

When $\sigma = 0.5$, this is our baseline model in which each player only knows his private value. When $\sigma = 1$, players learn (v_i, v_j) , and we have a unique BCE corresponding to the complete information equilibrium. Moreover, we know that if $\sigma > \sigma'$, then the signals in the model σ are more informative than the signals in the model σ' in the sense of Blackwell (1951). Hence, the set of BCE with information structure σ is contained in that of the information structure σ' , since any BCE under σ could be achieved by giving players additional information beyond what they receive under σ' .

In Figure 6.4, we depict the surplus set of the bidders for a model in which values are independent, and $Pr(v_i = 0) = \frac{2}{3}$. Players bids are confined to a grid of 35 values, equally spaced between 0 and 1. For values of σ close to 0.5, the set of BCE coincides with the baseline model, except that we shave off the corners of the set. This is intuitive: Our equilibria at the northwest and southeast corners require



that one bidder not learn from his bid about the other bidder's value. As soon as we give this bidder additional information through s_i , it is no longer possible that his first-order beliefs always coincide with the prior distribution.

On the other hand, the symmetric bidder surplus maximizing equilibrium requires both players to learn a modest amount from their bids: Beliefs are different after recommendations of $b_i = 0$ and $b_i = 1$. For low values of σ , the equilibrium bid distribution is more informative than s_i . However, for larger values of σ , s_i becomes more informative, and it is no longer possible to achieve the total bidder surplus maximizing equilibrium. As σ approaches 1, the surplus set converges to the complete information payoff of $\frac{2}{9}$. It is notable that even for $\sigma \approx 0.98$, there still exist BCE which substantially raise the surplus of the bidders above the complete information level.

7 Conclusion

The literature on auctions has documented the fact that the outcome of the first-price auction is sensitive to bidders' information. We have applied the recently developed methodology of Bayes correlated equilibria to systematically study exactly how much of an impact information can have on the welfare outcomes of the auction. Our key finding is that information can have a surprisingly large

effect, which we have attempted to quantify both theoretically and computationally. So far, we have learned a great deal about both the complexity of the problem and the subtle structure that sometimes emerges. While information can have a large effect, it does have its limits, and with regard to the binary valuation example and the lower bound on bidders surplus, we have given these limits a tight characterization. As our final results on reserve prices and entry fees show, with well chosen auction formats, it is possible to control how good, and how bad, this informational effect can be.

This a progress report rather than our final word on the subject. The computational results of Section 6 leave us with many unanswered questions and many conjectures about the role of information, particularly with regard to maximum bidder surplus and minimum revenue. We continue to be actively engaged in investigating these questions.

References

- ABRAHAM, I., S. ATHEY, M. BABAIOFF, AND M. GRUBB (2012): “Peaches, Lemons, and Cookies: Designing Auction Markets with Dispersed Information,” Discussion paper, Microsoft Research, Stanford, and MIT.
- AZACIS, H., AND P. VIDA (2013): “Collusive Communication Schemes in a First-Price Auction,” *Economic Theory*, forthcoming.
- BATTIGALLI, P., AND M. SINISCALCHI (2003): “Rationalizable Bidding in First Price Auctions,” *Games and Economic Behavior*, 45, 38–72.
- BERGEMANN, D., AND S. MORRIS (2013a): “The Comparison of Information Structures in Games: Bayes Correlated Equilibrium and Individual Sufficiency,” Discussion Paper 1909, Cowles Foundation for Research in Economics, Yale University.
- (2013b): “Robust Predictions in Games with Incomplete Information,” *Econometrica*, 81, 1251–1308.
- BLACKWELL, D. (1951): “Comparison of Experiments,” in *Proc. Second Berkeley Symp. Math. Statist. Probab.*, pp. 93–102. University of California Press, Berkeley.
- CELIS, L., G. LEWIS, M. MOBIUS, AND H. NAZERZADEH (2012): “Buy-it-now or Take-a-chance: A simple sequential screening mechanism,” Discussion paper, Harvard University.
- DEKEL, E., AND A. WOLINSKY (2003): “Rationalizable Outcomes of Large Independent Private-Value First Price Discrete Auctions,” *Games and Economic Behavior*, 43, 175–188.
- FANG, H., AND S. MORRIS (2006): “Multidimensional Private Value Auctions,” *Journal of Economic Theory*, 126, 1–30.
- FORGES, F. (1993): “Five Legitimate Definitions of Correlated Equilibrium in Games with Incomplete Information,” *Theory and Decision*, 35, 277–310.
- KEMPE, D., V. SYRGANIS, AND E. TARDOS (2012): “Information Asymmetries in first-price common-value auctions, or, How to Sell Third-Party Cookies,” Discussion paper, USC.
- KIM, J., AND Y.-K. CHE (2004): “Asymmetric information about rivals’ types in standard auctions,” *Games and Economic Behavior*, 46(2), 383–397.

- LOPOMO, G., L. M. MARX, AND P. SUN (2011): “Bidder collusion at first-price auctions,” *Review of Economic Design*, 15(3), 177–211.
- MASKIN, E., AND J. RILEY (2003): “Uniqueness of Equilibrium in Sealed High-Bid Auctions,” *Games and Economic Behavior*, 45, 395–409.
- 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.
- VICKREY, W. (1962): “Auctions and bidding games,” in *Recent advances in game theory*, ed. by O. Morgenstern, and A. W. Tucker, vol. 29, pp. 15–27. Princeton Univ. Press.