

Surveying and selling: Belief and surplus extraction in auctions

Job market paper

Benjamin A. Brooks*

December 18, 2013

Abstract

I consider the mechanism design problem of a seller who is uninformed about demand, while potential buyers are well-informed. The seller knows that (1) values are private, (2) buyers' beliefs are consistent with a common prior, and (3) the support of values is bounded relative to the expected efficient surplus. The seller's goal is to maximize the minimum ratio between expected revenue and the expected efficient surplus, which I term the extraction ratio.

I characterize simple mechanisms that maximize the minimum extraction ratio. In these mechanisms, the seller runs a second-price auction and simultaneously surveys the beliefs of buyers about others' values. Losing bidders' responses to the survey are used to set the reserve price for the winner. Such mechanisms guarantee the seller a substantial share of the efficient surplus, with the share parametrized by the bound on the highest value. If values can be up to ten times the expected efficient surplus, the seller is guaranteed a 20% revenue-share; up to 1,000 times, and the seller is guaranteed 10%.

Keywords: Auctions, beliefs, surplus extraction, maxmin preferences.

JEL classification: C72, D44, D81, D82, D83.

*Department of Economics, Princeton University, babrooks@princeton.edu. I would like to thank Dilip Abreu, Nemanja Antic, Dirk Bergemann, Sylvain Chassang, Rohit Lamba, and Stephen Morris for extremely helpful feedback and advice.

1 Introduction

1.1 Motivation

Consider a small municipality that is replacing a public school building. The replacement of schools is generally a rare event, with the average age of public schools in the U.S. being 42 years.¹ As such, it is reasonable to suppose that municipal officials do not have great expertise in assessing construction costs. On the other hand, the firms that bid for the contract are likely to have detailed knowledge of one another's costs and capabilities. Is it possible for the municipality to get the contractors to truthfully reveal what they know about one another's costs, even though the information they reveal will influence the award?

The elicitation of potential buyers' opinions by a seller is more than just a theoretical possibility. After the use of auctions for allocating radio spectrum was authorized by the U.S. Congress, the Federal Communications Commission (FCC) elicited feedback on its proposed rules from potential bidders and industry experts. The FCC received "written comments from 222 parties and reply comments from 169 parties" (FCC, 1997, p. 9). Such feedback was no doubt crucial to gauging the welfare effects of the new mechanism. The FCC does not specify when and how it incorporated this feedback into the auction design, but surely the responses of the interested parties were influenced by their strategic concerns vis-à-vis the ultimate allocation and costs of licenses.

In this paper, I will consider such situations, in which the seller of a good is *uninformed* about demand, whereas the potential buyers are *well-informed*. By well-informed, I mean that each agent knows their own private valuation for the good, and in addition, they have a belief about others' values which is derived from a common prior. The buyers' private valuations and beliefs can be thought of as being induced by informative signals, with the common prior corresponding to the ex-ante distribution over the signals. The set of signals together with the prior specify a *type space*, where a buyer's "type" is precisely the realized signal.² The seller could greatly benefit from knowing the type space: At the very least, such knowledge could facilitate the selection of a revenue enhancing reserve price, and in particular cases, the seller can even use variation in bidders'

¹According to the National Center for Education Statistics, as of 1999.

²Aside from requiring the common prior and that the set of types is finite (for tractability), I impose no additional restrictions on the type space. The set of possible environments is therefore quite general, and in fact includes type spaces for which there is no known characterization of the revenue maximizing Bayesian mechanism. See Farinha Luz (2013) for probably the most general characterization to date.

beliefs to extract all of the potential surplus as revenue. However, the seller in my model does not know the type space, and therefore he cannot build such detail-dependent features directly into the mechanism. Moreover the seller only interacts with the buyers after their types have been realized, and thus cannot easily incentivize them to reveal the prior. The remaining option, which the seller takes advantage of, is to use a mechanism that determines the allocation and transfers while simultaneously eliciting the buyers' interim beliefs, i.e., their beliefs after they learn their types but without knowing others' types. In this way, the outcome of the mechanism can be made to depend on the true type space. The buyers are of course aware of the effects of their reports, and will take advantage of any opportunity to misreport in order to favorably influence the outcome of the mechanism.

Given such large uncertainty about the type space, it is natural for the seller to use a worst-case criterion: The seller seeks a mechanism that will perform well irrespective of the true distribution of values and beliefs. Since it is possible for the buyers' valuations to be arbitrarily small, every mechanism has zero expected revenue in the worst-case. As a result, worst-case expected revenue is not a useful criterion to distinguish between mechanisms. Instead, I posit that the seller evaluates the performance of a mechanism by its expected revenue relative to the expected surplus that could be generated by allocating the good efficiently. I term this metric the *extraction ratio*: The ratio of expected revenue to expected efficient surplus. In addition, the seller makes no presumption that the buyers will behave according to his preferred equilibrium, so he evaluates a mechanism by its lowest extraction ratio over all type spaces and over all equilibria. Similar criteria have been considered in the literature, most notably by Neeman (2003) and by the computer science literature on mechanism design, surveyed in Hartline (2012). I will revisit the connections with these and other papers below and in some detail. By using a mechanism that maximizes the minimum extraction ratio, the seller will be guaranteed at least a minimum share of the expected efficient surplus, regardless of the true distribution of buyers' values and beliefs.

1.2 Overview of main results

My main result is that there is a simple class of mechanisms that the seller can use to achieve the max min extraction ratio. Moreover, this max min extraction ratio is economically substantial, as I will elaborate upon shortly. These mechanisms are essentially modified second-price sealed-bid

auctions, in which the buyers simultaneously submit bids as well as respond to a survey of their beliefs about the values of others. The high bidder will be “offered” the good at a price determined using others’ reports and ultimately receives the good if this price is less than the high bid. Because each buyer’s bid does not affect the price of the good, but only whether or not the good is received at an exogenously chosen price, truthful bidding is a weakly dominant strategy. A slight perturbation makes bidding one’s value strictly dominant. Also, given that others’ will bid truthfully, buyers can be incentivized to report their true beliefs about others’ values using a scoring rule. The use of scoring rules to elicit beliefs in mechanism design has also been considered by Azar et al. (2012). To calculate the price offered to the high bidder, the seller uses one of the losing bidders’ survey reports as a “consultation” about the conditional distribution of the highest value. This consult, together with the second-highest bid, is used to compute an optimal price to offer the winner. I give these mechanisms the descriptive moniker of *belief survey auctions* (BSA), since the seller uses a survey of losing bidders’ beliefs to set the winner’s price.

I derive the minimum extraction ratio for the BSA, and I show that no other mechanism could achieve a greater extraction ratio in the worst-case. The strict incentives to bid one’s value and report beliefs truthfully can be provided at arbitrarily small cost to the extraction ratio, so the BSA *virtually* achieves the max min. It turns out that if the support of valuations is unbounded, the max min extraction ratio is zero: It is possible to have arbitrarily large expected efficient surplus while holding the seller to finite expected revenue. However, these type spaces are extreme in that they have a lot of mass in the tail of the distribution of the highest value. A natural assumption is that the support of the highest value is bounded by a constant multiple of the expected efficient surplus, effectively limiting the dispersion of the highest value around its mean. I study how the max min extraction ratio changes with the bound on the dispersion of values. For any bound, the max min extraction ratio is strictly positive, and even for very generous bounds on values, the max min is economically substantial. As an example, if buyers’ values cannot be more than 10 times the expected efficient surplus, then the seller is guaranteed an extraction ratio of at least 20%. If values can be 1,000 times larger, the seller is still guaranteed a 10% extraction ratio.

It is particularly interesting that the seller is able to achieve these bounds with such simple mechanisms. The seller never recovers the prior distribution over values, but rather sets reserve prices using bidders’ interim beliefs. As I will argue below, this is actually a virtue of the mechanism;

bidders' interim beliefs are weakly more informative than the prior, and thus allow the seller to set better reserve prices at the interim stage. Moreover, the seller does not even need to elicit beliefs about the entire distribution of buyers' values; it is sufficient for the seller to ask bidders the conditional distribution of the top two valuations of other bidders and the number of bidders who tie for the highest value.

In addition to guaranteeing the seller a minimum extraction ratio, the BSA has desirable revenue properties away from the worst case. Aside from the small cost of providing strict incentives, the BSA guarantees the seller at least the revenue of a second-price auction with an optimal anonymous reserve price, i.e., a uniform reserve price for all bidders that maximizes expected revenue. As such, the BSA virtually maximizes expected revenue over all Bayesian mechanisms when the distribution of values is independent, symmetric, and regular, as in Myerson (1981).

A potential concern with the maxmin extraction ratio is that the seller seems to be indifferent between outcomes with very different expected revenues, as long as the expected efficient surplus varies proportionally. However, no such comparisons are necessary to justify the use of the BSA. An alternative way to model the seller's preference over mechanisms is the following conditional ordering: The seller prefers greater worst-case expected revenue conditional on the level of the expected efficient surplus, but he will not compare revenue outcomes between type spaces in which the social value of the good varies. As a result, one mechanism is preferred to another only if it has greater worst-case expected revenue conditional on *every* possible level of the expected efficient surplus. Observe that this ranking is only a partial order on the set of mechanisms, because the seller does not compare mechanisms whose worst-case revenue ranking switches depending on the surplus level. It turns out that the BSA is maximal with respect to this partial ordering: If the seller's wishes to maximize the minimum expected revenue conditional on a particular level of the expected efficient surplus, then he can select no better mechanism than the BSA.

1.3 The logic behind the BSA

Here I will give a brief summary of how my results are obtained. The BSA offers the good to the high bidder, so on average, the winner's valuation is drawn from the distribution of the highest value among the n bidders. If the seller knew the prior distribution over values, he could set the reserve price which maximizes revenue without knowing the identity of the winner, which is the

optimal anonymous reserve price. By assumption, the seller does not know the prior distribution over values, but the reports of the losing bidders allow the seller to set reserve prices conditional on more detailed information. Specifically, the seller learns the distribution of the winner's value conditional on (1) the winner not being the bidder who was consulted, (2) the realization of the second-highest bid, and (3) any extra information the consulted bidder has about the distribution, as encoded in his type. Conditional on (1)-(3), the seller can always set a reserve price that generates weakly more revenue than could be achieved with an optimally chosen anonymous reserve price. Thus, the BSA performs better on average than any second-price auction with an anonymous reserve, in spite of the fact that the seller never recovers the prior nor does he know the optimal anonymous reserve price.

There is an analogy to be made with third-degree price discrimination, in which a monopolist receives information that divides a market into segments. If the monopolist can set different prices in different segments, then this information must be weakly revenue increasing, since it is always feasible to set the optimal uniform price. A similar property holds in the auction setting. More informative reports by losing bidders allows the seller to set better reserve prices, and hence the worst-case environments for the BSA are ones in which (1)-(3) are minimally informative. These type spaces exhibit the property that bidders get no information beyond their private values, which minimizes the informativeness of (3). Moreover, the worst-case type spaces are lopsided, in the sense that at any time there is only one "serious" bidder who submits the high bid, and the other bidders know that they will not win, which minimizes the learning from (1) and (2). In a sense, this reduces the seller's problem to designing a mechanism for selling to a single serious buyer. Even so, multiple bidders will participate in the auction, and their reports are used by the seller to set an optimal reserve price for that single buyer. Finally, I derive the distribution for the serious bidder's value that minimizes revenue, subject to a given level of the efficient surplus.

With additional restrictions on the environment, the seller can achieve the same goals with mechanisms that are even simpler. Throughout the analysis, careful attention is paid to the possibility of multiple bidders having the same valuation, so that the winner is determined by a tie break. The tie break induces a selection effect: Conditional on winning the auction, the winner is less likely to have a valuation at which ties are likely to have occurred. For this reason, the seller must survey bidders' beliefs about the likelihood of ties. One might think of ties as being a non-

generic phenomenon, for example if values are drawn from a non-atomic distribution. If attention is restricted to type spaces in which ties do not occur with positive probability, the conditional distribution of the high bidder's value can be calculated much more simply. Also, as mentioned in the previous section, the seller elicits the buyers' beliefs about the top two valuations among other buyers. By leveraging the information about the highest value contained in the second-highest value, the BSA always perform better than a second-price auction with an optimal reserve price. However, if the seller is only concerned about worst-case extraction ratio, then the same bounds can be achieved with a mechanism that only elicits beliefs about the highest value of others.

1.4 Related literature

The results described above have a tight connection to the work of Neeman (2003), who studies the worst-case extraction ratio of the second-price auction. Neeman considers a seller who has three different levels of sophistication with regard to reservation prices. At the most basic level, the seller cannot set any reserve price. At the next level, the seller can use a fixed reserve price that is independent of the true type space. At the highest level of sophistication, the seller knows the distribution of values and is able to set the optimal anonymous reserve price for the true type space, although the seller is not sufficiently sophisticated to design and run the optimal auction. It is this last case that is the most relevant to the present paper. For this setting, Neeman derives bounds on the extraction ratio that are equal to my own, albeit with a slightly different parametrization of the set of type spaces. Indeed, since the BSA always generates as much revenue as a second-price auction with an anonymous reserve, and since the revenue of these two mechanisms coincides on the worst-case type spaces derived by Neeman, it is necessarily the case that both have the same minimum extraction ratio. However, I will give a direct proof of worst-case type spaces for the BSA, to better illuminate the connection with third-degree price discrimination described above.

Another paper which is closely related is that of Azar et al. (2012). They also consider a seller who is uninformed about the type space while the agents are well-informed, and they look for general mechanisms that achieve a favorable worst-case performance relative to the benchmark of maximum revenue in a dominant strategy ex-post individually rational mechanism. Similar to the present work, they extract buyers' beliefs using scoring rules. They consider a restricted class of environments, for which the gap between first-order beliefs and the prior distribution is relatively

small.³ By eliciting the buyers' first-order beliefs, the seller is able to recover a truncated view of the prior, and this is used as an input into a dominant strategy mechanism. In comparison, the present work is in much more general environments, in which buyers can have arbitrary conditional beliefs about the distribution of values. As such, very different arguments are required to arrive at my results. Also, I use a different benchmark which does not assume a restriction to a particular implementation concept. Nonetheless, to achieve the max min extraction ratio, it is sufficient for the seller to use simple mechanisms that only extract first-order beliefs about statistics of others' values.

More broadly, my work is part of the large literature on robust mechanism design (Bergemann and Morris, 2012b, provide an overview). At least since the critique of Wilson (1987), the mechanism design literature has held as a desideratum that mechanisms should be *detail-free*, in the sense that the rules of the game should not vary with fine details of the environment. This is in contrast to classical auction design, e.g., Myerson (1981) and Crémer and McLean (1988), in which the mechanism can be tailored to specific and highly structured type spaces. A more recent contribution of Farinha Luz (2013) considers very general type spaces but still allows the mechanism to depend on the type space.

The robust mechanism design literature has explored various ways to operationalize the Wilson critique. Much of the literature focuses on more stringent implementation concepts. For example, Bergemann and Morris (2009, 2011) require that a particular social choice function be implemented regardless of the beliefs of the agents. I consider auction formats that are compatible with with a slightly different interpretation of the detail-free criterion: The mechanisms that the uninformed seller can use are detail-free in that the distribution of values and beliefs of the agents cannot be hard-wired into the mechanism. However, the outcome of the mechanism can depend on details of the environment through equilibrium behavior, if these details are known to the agents.

Other authors have considered criteria akin to max min extraction ratio. As discussed above, the closest such related work is that of Neeman (2003). Bergemann and Schlag (2011) consider a monopolist facing unknown demand from a single buyer, and characterize the pricing rule that

³Specifically, they consider environments in which bidders' beliefs are derived from a common prior in the following manner: Each bidder is associated with a partition of others' values, and bidders learn their own value and the cell of the partition containing other bidders' values. The first-order beliefs of different types of the same bidder have disjoint supports, and beliefs are always proportional to the prior distribution on their support.

achieves min max regret, which is the absolute difference between expected revenue and expected efficient surplus. Chassang (2011) studies dynamic incentive contracts, and solves for contracts that achieve a target that is analogous to max min extraction ratio. Carroll (2012) also considers max min preferences over contracts in a static setting. Chung and Ely (2007) give a foundation for dominant strategy mechanisms by positing a seller with worst-case preferences and who knows the distribution of private values but not the beliefs of the agents, which may be inconsistent with a common prior.

The criterion of max min extraction ratio is similar to the competitive ratios studied by computer scientists (see Hartline (2012) for a comprehensive survey). This literature looks at worst-case revenue ratios, with a variety of benchmarks in the denominator. The benchmark is often tailored to a specific solution concept, such as maximum revenue over all dominant strategy mechanisms. The efficient surplus is in a sense a more demanding benchmark, as it does not presume a restriction to a particular class of mechanisms. An assumption throughout some of the literature is that mechanisms can only elicit one-dimensional bids, which precludes the belief extraction approach of the present model. Chawla et al. (2007) and Hartline and Roughgarden (2009) study worst-case competitive ratios for the second-price auction with optimal reserve prices, which presumes that the seller knows the prior. Goldberg et al. (2004) and Goldberg and Hartline (2003) look at mechanisms which do not depend on the prior, with a benchmark which is the revenue the seller could generate selling $k \geq 2$ units of the good at the k th highest price. Such a benchmark could be zero in cases where the efficient surplus is positive.

Others have considered how a seller can learn about demand. Baliga and Vohra (2003) and Segal (2003) consider a seller who forecasts the distribution of values using past realizations. In contrast, I will look at a situation where the seller asks agents for their beliefs, rather than dynamic learning based on reported values. Caillaud and Robert (2005) construct detail-free mechanisms that use agents' beliefs to partially implement the optimal auction of Myerson (1981). Choi and Kim (1999) consider belief extraction in the context a public goods problem, but assume the existence of an ex-ante stage at which the seller can extract prior beliefs, before the realization of agents' private information.

Finally, this work is part of my broader investigation into mechanisms that harness the agents' beliefs about the environment, to make up for a lack of knowledge on the part of the designer.

I see the present model as a midpoint in the tradeoff between the simplicity of the mechanism and the strength of the optimality criterion. In Brooks (2013a), I investigate the limits of how much the seller could learn about the environment. This relates to a classic “folk argument” in the mechanism design literature, that if a common prior were known to the agents and not to the designer, then the designer could recover the prior for free (Bergemann and Morris, 2012a), in the sense that the need to recover the prior does not restrict the social choice functions that the seller can implement once the prior is known. I show that the designer can indeed extract the prior, without compromising on how the prior will be used, by using a mechanism which elicits bidders’ infinite hierarchy of beliefs. While complexity is not explicitly modeled, it is safe to say that this mechanism would be much more challenging to implement than the BSA. At the other end of the spectrum, Brooks (2013b) looks at mechanisms in which the seller runs a second price auction and simply asks each bidder to suggest a reserve price for the other bidders. The seller incentivizes truth-telling by sharing revenue generated through a bidder’s suggestion. In more structured type spaces, this mechanism has a natural equilibrium in which bids are close to values, and bidders suggest reserve prices that are approximately optimal. I revisit the broader agenda in Section 5.

The rest of this paper is organized as follows. In Section 2, I describe the model and the seller’s mechanism design problem. In Section 3, I present a simple example that illustrates some of the main ideas of the paper. Section 4 presents the main results. Section 5 is a discussion, and Section 6 concludes. Omitted proofs are in the Appendix.

2 Model

There are n potential buyers for a single unit of a private good, indexed by $i \in N = \{1, \dots, n\}$. I adopt the usual convention that $-i = \{j \in N | j \neq i\}$, and vectors x_S denotes the sub-vector of x containing indices in S , e.g., $t = (t_i, t_{-i}) \in T$ is a profile of types. For real vectors x , $x^{(1)}$ denotes the highest value in x , $x^{(2)}$ denotes the second-highest value, and $x^{(1,2)}$ is the ordered pair of the highest and second-highest values. If x only has a single coordinate, then $x^{(2)} = -\infty$.

The values and beliefs of the bidders are modeled with the language of type spaces. In particular, there is a finite set of types $T = \times_{i \in N} T_i$ and a joint distribution $\pi \in \Delta(T)$. The notation $\Delta(X)$ denotes the set of probability measures on X with finite support. Each type $t_i \in T_i$ is associated

with a private value

$$\phi_i(t_i) \in \mathbb{R}.$$

Together, a *type space* is the triple $\mathcal{T} = (T, \pi, \phi)$. I will write $\pi(t_{-i}|t_i)$ for the conditional distribution of types given t_i , and $\pi_i(t_i)$ for the marginal distribution on T_i . For each type space \mathcal{T} , the expected surplus generated if the good were allocated efficiently is

$$S(\mathcal{T}) = \sum_t \phi^{(1)}(t) \pi(t). \quad (1)$$

I write \underline{v} and \bar{v} for smallest and largest values in the support of the measure over values induced by π under the mapping ϕ . For easy reference, all notation is compiled in Table 1, which appears at the end of the paper.

A type space is *symmetric* if (T_i, ϕ_i) is the same for all bidders, and π is exchangeable in the types, i.e., $\pi(t_1, \dots, t_n) = \pi(t_{\psi(1)}, \dots, t_{\psi(n)})$ where ψ is a permutation of N . A *payoff* type spaces has the property that $|\phi_i^{-1}(v_i)| \leq 1$ for all $v_i \in \mathbb{R}_+$. In other words, each valuation is associated with at most one type, so all bidders with a given valuation have the same conditional beliefs about other bidders' types as well as other bidders' valuations. I will say that a type space is *lopsided* if with probability one, at most one bidder has a valuation above the minimum of the support. These type spaces are lopsided in the sense that the winner's valuation tends to be much larger than the second-highest value.

The seller must design an auction for the sale of the good. A mechanism consists of a measurable space of messages M_i for each player, $M = \times_{i \in N} M_i$, and mappings $q : M \rightarrow \mathbb{R}^n$ and $p : M \rightarrow \mathbb{R}^n$. These are respectively the allocation rule and net transfer to the seller: $q_i(m)$ is the probability that agent i is allocated the good and $p_i(m)$ are agent i 's net transfer when the message profile m is sent. Naturally, $q_i(m)$ is required to be non-negative and $\sum_{i \in N} q_i(m) \leq 1$. A *mechanism* is the triple $\mathcal{M} = (M, q, p)$.

A mechanism and a type space together define a Bayesian game, in which each player's strategy set is $\Sigma_i(\mathcal{M}, \mathcal{T}) = \{\sigma_i : T_i \rightarrow \Delta(M_i)\}$. I write $\sigma_i(dm_i; t_i)$ for the probability measure over bidder i 's messages m_i given type t_i . For a strategy profile $\sigma \in \Sigma(\mathcal{M}, \mathcal{T}) = \times_{i \in N} \Sigma_i(\mathcal{M}, \mathcal{T})$ and type $t_i \in T_i$,

bidder i 's payoff is

$$u_i(\sigma, t_i) = \sum_{t_{-i} \in T_{-i}} \pi(t_{-i}|t_i) \int_{m \in M} [\phi_i(t_i)q_i(m) - p_i(m)] \sigma(dm; t).$$

where $\sigma(dm; t) = \times_{i \in n} \sigma_i(dm_i; t_i)$ is the product measure on M . A profile σ is a *Bayesian Nash equilibrium* if

$$\sigma_i \in \arg \max_{\sigma'_i \in \Sigma_i(\mathcal{M}, \mathcal{T})} \sum_{t_i \in T_i} \pi_i(t_i) u_i((\sigma'_i, \sigma_{-i}), t_i).$$

I denote by $\text{BNE}(\mathcal{M}, \mathcal{T})$ the set of all Bayesian Nash equilibria. Note that this set may be empty for particular choices of $(\mathcal{M}, \mathcal{T})$. The revenue of \mathcal{M} under a particular type space \mathcal{T} and strategy profile σ is

$$R(\mathcal{M}, \mathcal{T}, \sigma) = \sum_{t \in T} \pi(t) \int_{m \in M} \sum_{i \in N} p_i(m) \sigma(dm; t). \quad (2)$$

The corresponding *extraction ratio* is

$$E(\mathcal{M}, \mathcal{T}, \sigma) = \frac{R(\mathcal{M}, \mathcal{T}, \sigma)}{S(\mathcal{T})}, \quad (3)$$

with the convention that when $S = 0$, $E = 1$.

The seller's goal is to select find mechanisms that solve

$$\sup_{\mathcal{M}} \inf_{\mathcal{T}} \inf_{\sigma \in \text{BNE}(\mathcal{M}, \mathcal{T})} E(\mathcal{M}, \mathcal{T}, \sigma). \quad (4)$$

The interpretation of this problem is: The seller must select a mechanism \mathcal{M} , following which Nature⁴ will select both the type space \mathcal{T} and the equilibrium $\sigma \in \text{BNE}(\mathcal{M}, \mathcal{T})$. If $\text{BNE}(\mathcal{M}, \mathcal{T})$ is empty, then our convention is that the infimum is zero. I will refer to the value (4) as the max min extraction ratio.

The formulation of (4) is quite demanding: If the seller chooses a mechanism for which there are multiple equilibria, Nature will select the one with the lowest extraction ratio. The exposition

⁴Throughout, I use Nature as the personification of all minimization operations beyond the control of the designer.

is simplified by initially allowing the seller to choose the equilibrium, which permits us to achieve a “partial” max min extraction ratio. The results will later be strengthened to “full” max min by allowing Nature to select the equilibrium with the lowest ratio. This terminology is modeled after the partial and full implementation concepts in mechanism design, although to be clear, full max min does not require equilibrium uniqueness, but just that the extraction ratio be maximized in the worst type space and worst equilibrium. When it is clear which equilibrium is used, I will simply write $E(\mathcal{M}, \mathcal{T})$ for the extraction ratio.

3 Example

Let us start by considering a simple example that will illustrate some of the main ideas. There are two potential buyers $i = 1, 2$, and each buyer could be of type L or type H . Type L thinks the good is worth v , an type H thinks the good is worth $2v$. The distribution of types is given by the table below.

	L	H
L	$1 - 2\rho - \psi$	ρ
H	ρ	ψ

Relative to the general model, I have assumed that the type space is a symmetric and payoff type space, and that the support of values is of the form $\{v, 2v\}$. These features are common knowledge among the buyers and the seller. However, the parameters ρ , ψ , and v are known to the buyers and unknown to the seller. These assumptions are made for the simplicity of the example, and will be relaxed for the main results.

The maximum surplus that can be generated by allocating the good efficiently is:

$$\begin{aligned}
 S &= v(2(\psi + 2\rho) + 1 - 2\rho - \psi) \\
 &= v(1 + 2\rho + \psi).
 \end{aligned}$$

The seller has to select a mechanism to sell the good which is independent of v , ρ , and ψ , and therefore independent of S as well. As discussed in the introduction, the seller is highly uncertain about S , and lacks beliefs about which parameters are likely to obtain. As a result, the seller

compares mechanisms by a worst-case performance metric which is scale free in S : the minimum extraction ratio. Our seller will first pick a mechanism, and then Nature will choose the parameters to minimize the extraction ratio.

For starters, let us consider what would happen if the seller were to use a second-price auction. If there is a positive reserve price $r > 0$, then Nature could always select v such that $2v < r$. S would be positive, but $R = 0$ (since the reserve is greater than the highest value). This is an important observation: introducing a positive reserve price that is totally unresponsive to the parameters of the model leads to extremely unfavorable outcomes in the worst-case. With a reserve price of zero, revenue is

$$\begin{aligned} R &= v(2\psi + 1 - \psi) \\ &= v(1 + \psi). \end{aligned}$$

Hence, the extraction ratio would be

$$E = \frac{1 + \psi}{1 + 2\rho + \psi}.$$

Clearly, to minimize the ratio, Nature should make ψ as small as possible and ρ as large as possible, so $\psi = 0$ and $\rho = \frac{1}{2}$. The resulting extraction ratio is $E = \frac{1}{2}$. Note well that it does not actually matter what level v Nature chooses: The model is “scale-free” in v .

	L	H
L	0	$\frac{1}{2}$
H	$\frac{1}{2}$	0

The second-price auction with no reserve is of course just one mechanism. Let us consider a simple modification. In addition to accepting bids b_i , the seller canvasses the bidders for what they believe about the distribution of others’ bids. Each bidder’s response to the survey will be used to set the reserve price only if that bidder does not win. Specifically, bidders submit a quantity w_i , which is the bidder’s report for the value of v , and a quantity μ_i , which is the bidder’s reported probability that $b_j = 2v$. If bidder i has the high bid, or if bids are equal and i wins a uniform tie break, then bidder i will be “offered” the good at a reserve price r_i that only depends on (b_j, w_j, μ_j) .

Moreover, this price is always at least b_j . Thus, bidder i is facing an exogenous distribution of prices, depending on the other bidder's reports, and he will receive the good and pay the price as long as $b_i \geq r_i$. As in the second-price auction, truthful bidding is a weakly dominant strategy: $b_i = v_i$. In the following, I impose that bidders follow this strategy.

Bidders will receive a small side reward for their survey response (w_i, μ_i) . In particular, bidder i is paid according to the scoring rule

$$\epsilon b_j \left(\mu_i \mathbb{I}_{b_j=2w_i} + (1 - \mu_i) \mathbb{I}_{b_j=w_i} - \frac{(\mu_i)^2 + (1 - \mu_i)^2}{2} \right). \quad (5)$$

Given a report $w_i = v$, then there is a strict incentive to report $\mu_i = Pr(b_j = 2v|v_i)$, and in fact the bidder receives a positive net payment from the seller. Moreover, any report other than $w_i = v$ induces a smaller payoff: If $w_i \neq 2v$ as well, then the bidder always makes a net transfer to the seller, and if $w_i = 2v$, bidder i only gets paid when $b_j = 2v$, and not when $b_j = v$. Thus, in any equilibrium in which the buyers bid their values, they must also truthfully report $w_i = v$ and their conditional belief that $b_j = 2v$. Note that the payment is scaled by $b_j = v_j > 0$, so that in expectation, the transfer to bidder i from (5) is no more than $\epsilon \mathbb{E}[v_j] \leq \epsilon S$.

The seller offers the winner the good at the revenue maximizing price, conditional on the winner being the high bidder and winning any tie breaks, and also conditional on the loser's reported beliefs. Suppose bidder i has the high bid and wins a tie break. Clearly, if the second-highest bid is $b_j = 2w_j = 2v$, then the seller should set i 's price at $2v$. If $b_j = v$, then the probability of $v_i = 2v$ is

$$\frac{2\mu_j}{3} = \frac{2}{3} \frac{\rho}{1 - \rho - \psi},$$

and the conditional probability of $v_i = v$ is

$$\frac{1 - \mu_j}{3} = \frac{1}{3} \frac{1 - 2\rho - \psi}{1 - \rho - \psi}.$$

These formulae exhibit the selection effect of the tie: The "raw" probability of both players having valuation v is $1 - \mu_j$, and one having $2v$ is μ_j . But if both have a low value, bidder i only wins half the time, thus leading to the formulae above. The optimal price in this case is $2v$ if $\rho \geq \frac{1-\psi}{4}$ and

v if $\rho \leq \frac{1-\psi}{4}$.

Thus, Nature has two options. If $\rho \leq \frac{1-\psi}{4}$, then revenue is $R_1 = v(1 + \psi)$, and the extraction ratio is

$$E_1 \geq \frac{R_1}{S} = \frac{1 + \psi}{1 + 2\rho + \psi},$$

which is minimized by making ρ as large and ψ as small as possible. Hence, E_1 is minimized at $\rho = \frac{1}{4}$ and $\psi = 0$. At these values, $E_1 = \frac{2}{3}$.

If $\rho \geq \frac{1-\psi}{4}$, then the price is always $2v$, and revenue is $R_2 = 2v(2\rho + \psi)$ and the extraction ratio (not counting transfers associated with (5)) is

$$E_2 \geq \frac{R_2}{S} = \frac{2(2\rho + \psi)}{1 + 2\rho + \psi}.$$

The ratio is decreasing in ρ and ψ . Substituting in $\rho = \frac{1-\psi}{4}$, the ratio is decreasing in ψ , so again the optimal values are $\rho = \frac{1}{4}$ and $\psi = 0$. Hence, $E_2 = \frac{2}{3}$. Taking into account at most ϵS in lost revenue for each bidder due to (5), the extraction ratio for this mechanism is at least $\frac{2}{3} - 2\epsilon$.

	L	H
L	$\frac{1}{2}$	$\frac{1}{4}$
H	$\frac{1}{4}$	0

Note that this type space is lopsided, in the sense introduced in Section 2: With probability one, at most one bidder has a valuation greater than v , which is the bottom of the support.

The bottom line: A simple modification of the second-price auction yields a substantial improvement in the worst-case extraction ratio from $\frac{1}{2}$ to $\frac{2}{3}$. This mechanism accepts bids and also surveys bidders' beliefs about the distribution of others' bids, with the truthful revelation of this information being incentivized with a scoring rule. Each bidder's survey response is used to set the reserve price when the other bidder wins, which protects the seller from low revenue when there is a large gap between the highest and second-highest values.

It is worth noting that this mechanism, while a significant improvement over the second-price auction with no reserve, does *not* maximize the minimum possible extraction ratio. It is easy to see that in the worst-case distribution, bidders' beliefs determine their preferences in the sense of

Neeman (2004), since H puts zero probability on the other bidder being of type H . With a more complicated mechanism in which the seller elicits second-order beliefs and introduces side-bets, this property could be exploited to extract all of the efficient surplus. The only type spaces which do not have this property are those in which types are drawn independently, for which the extraction ratio is minimized when $Pr(v_i = 2v) = \sqrt{2} - 1$, and the max min extraction ratio is approximately 0.7071.

This wedge is entirely due to the assumption that the support of values is restricted to being of the form $\{v, 2v\}$. In the rest of the paper, I will pursue a similar analysis but in the more general setting of Section 2, without restrictions on the number of bidders, on the support of valuations, or on the kinds of information that bidders might learn about the distribution. It will turn out that the worst-case type spaces approach a continuous distribution of values, in contrast to this inherently discrete example. A straightforward generalization of the mechanism described above achieves the max min extraction ratio for this more general problem.

4 Characterizing the max min extraction ratio

4.1 Preamble

I now proceed to characterize the max min extraction ratio and present simple mechanisms that achieve the max min. I begin by defining a particular auction, the *belief survey auction* (BSA). This mechanism is a modified second-price auction in which the seller accepts bids and also elicits reports of first-order beliefs. A bidder's reported belief is used to set the reserve price when one of the other bidders wins the auction. I show that truthful reporting of values and beliefs is incentive compatible, and in this truthful equilibrium, the seller is guaranteed a tight lower bound on the extraction ratio. In particular, Lemma 1 shows that there is a small subset of type spaces, namely symmetric and lopsided payoff type spaces, within which the extraction ratio for the BSA can be minimized. The argument proceeds by taking a given type space as input, and producing a new symmetric and lopsided payoff type space with the same efficient surplus and weakly lower revenue, and hence a lower extraction ratio. These type spaces have a simple interpretation: The reports of losing bidders allow the seller to set reserve prices, and the more informative the losers' reports are, the better reserves the seller is able to set. In these worst-case type spaces, bidders' reports are

minimally informative: Since they are symmetric and payoff type spaces, bidders beliefs contain no information beyond the value, and all bidders' reports are equally informative. And since all losing bidders have the same value equal the bottom of the support, there is just one belief that is used to set the reserve price.

The extraction ratio of the BSA on such type spaces is completely determined by the distribution of the highest value. I show that for a given level of revenue, the efficient surplus is maximized by drawing the highest value from a particular Pareto distribution. If the support of the distribution is unbounded, the highest value has infinite expected value, and the resulting extraction ratio is zero. As a result, I consider type spaces in which the support of the highest value is bounded as a constant multiple γ of the efficient surplus. This constant parametrizes the set of type spaces, and for each value of γ I characterize the max min extraction ratio.

In addition, the BSA turns out to be an optimal auction on symmetric and lopsided payoff type spaces. As a result, the bound on the extraction ratio is tight: No mechanism can have a higher extraction ratio in the worst case. This establishes the partial max min extraction ratio result of Theorem 1. Finally, I show that if the seller rewards bidders for their reported beliefs using a scoring rule, truthful reporting can be made the unique strategy profile that survives iterated deletion of dominated strategies. Hence, the partial result is strengthened to full max min in Theorem 2.

I note that there are in fact many mechanisms which approach the solution to (4). The mechanisms I consider are notable for their simplicity, but in Section 5 I will discuss some alternatives.

4.2 The belief survey auction

Our foundation for constructing the BSA is the second-price auction. This auction has an important property: each bidder is facing a random price at which he could purchase the good, where the distribution of the price is completely determined by the strategies of other bidders. In the second-price auction, this price is the highest bid made by other bidders, $b_{-i}^{(1)}$. The own bid b_i is the cutoff such that bidder i would like to purchase the good if the realized price is less than b_i . Since it is optimal to buy the good at any price below the bidder's value, $b_i = v_i$ is a weakly dominant strategy.

The BSA will retain this property: Each buyer submits a bid b_i which is the cutoff at which they accept a price which is a function of other buyers' reports. Our point of departure is that this

price is not $b_{-i}^{(1)}$, but rather incorporates more information that is elicited from the other bidders. Specifically, in addition to a bid, each bidder will submit a report of their beliefs about the joint distribution of (1) the highest bid of others $b_{-i}^{(1)}$, (2) the second-highest bid of others $b_{-i}^{(2)}$, and (3) the number k of high bidders amongst the other players, i.e., the number of players j such that $b_j = b_{-i}^{(1)}$. Naturally, if $b_{-i}^{(1)} > b_{-i}^{(2)}$, then under a truthful report, $k = 1$ with probability 1. Assuming the report is truthful, bidder j 's reported beliefs allow the seller to determine an optimal price to charge the winning bidder conditional on the winner not being bidder j , and conditional on the second-highest bid of others. The report of the number of tied bidders allows the seller to control for the selection effect induced by tie breaking.

This mechanism strikes a balance between the amount of information about the environment that the seller elicits from bidders and the range of type spaces in which the mechanism maximizes revenue. In particular, by canvassing beliefs about $b_{-i}^{(2)}$ in addition to $b_{-i}^{(1)}$, the seller is able to set reserve prices that are better on average than the optimal anonymous reserve price in the second-price auction (Proposition 2). If the seller collected less information, namely beliefs about $b_{-i}^{(1)}$ and the number of high bidders, he could still achieve max min extraction ratio (Theorems 1 and 2) but would no longer be guaranteed to do as well as the second-price auction.

More formally, I define a mechanism \mathcal{M}^{BSA} as follows. Each message m_i consists of a bid b_i and a distribution μ_i in $\Delta(\mathbb{R}_+^2 \times \mathbb{N})$, where $\mathbb{R}_+ = [0, \infty)$ and \mathbb{N} is the set of positive integers. For a vector x , let

$$W(x) = \left\{ i \mid x_i = x^{(1)} \right\} \quad (6)$$

denote the set of maximal indices in x . The interpretation is that μ_i is bidder i 's reported beliefs about the distribution of $b_{-i}^{(1,2)}$ and $|W(b_{-i})|$. Thus, $M_i = \mathbb{R}_+ \times \Delta(\mathbb{R}_+^2 \times \mathbb{N})$. A typical message will be written $m_i = (b_i, \mu_i)$. The allocation rule is specified as follows. Suppose bidder i submits the highest bid, $b_i = b^{(1)}$. If there are ties, the mechanism selects i uniformly from the set of high bidders $W(b)$. We will then pick a bidder $j \neq i$ uniformly to calibrate the price $r_j(m_j, b_{-j}^{(2)}) \geq b^{(2)} = b_{-i}^{(1)}$ for bidder i . If $b_i \geq r_j(m_j, b_{-j}^{(2)})$, bidder i wins the good and pays $r_j(m_j, b_{-j}^{(2)})$. Otherwise, the

good remains unallocated. Hence,

$$q_i^{BSA}(m) = \frac{1}{|W(b)|} \frac{1}{n-1} \sum_{j \neq i} \mathbb{I}_{b_i \geq r_j(m_j, b_{-j}^{(2)})},$$

$$p_i^{BSA}(m) = \frac{1}{|W(b)|} \frac{1}{n-1} \sum_{j \neq i} \mathbb{I}_{b_i \geq r_j(m_j, b_{-j}^{(2)})} r_j(m_j, b_{-j}^{(2)}),$$

where \mathbb{I}_C is the indicator function, equal to one if condition C is met and zero otherwise. The price $r_j(m_j, b_{-j}^{(2)})$ in fact does not depend on m_i when i is allocated the good, since $b_i \geq b^{(2)}$. Thus, bidding one's value is a weakly dominant strategy and for now I impose that this occurs in equilibrium.

Also, note that bidder i 's report of μ_i has no effect on any price $r_j(m_j, b_{-j}^{(2)})$ when bidder $j \neq i$ is consulted, nor does it affect whether or not i is offered the good at any price. Hence, any report of μ_i is incentive compatible. I consider the “truth-telling” equilibrium in which bidders report

$$\mu_i(v_{-i}^{(1,2)}, k) = \sum_{\left\{ t_{-i} \mid \begin{array}{l} \phi_{-i}^{(1,2)}(t_{-i}) = v_{-i}^{(1,2)}, \\ |W(\phi_{-i}(t_{-i}))| = k \end{array} \right\}} \psi(t_{-i} | t_i).$$

In plain language, bidders report the conditional joint distribution of the first two order statistics of others' bids, and the number of high bidders among the other players. We will subsequently see that this strategy can be made the unique equilibrium, for any \mathcal{T} , at an arbitrarily small cost to the extraction ratio.

I still have to specify the prices that bidders are offered. What I would like to implement is a “monopoly” price with respect to the conditional distribution of the winner's value when bidder j is consulted. Each bidder reports their beliefs μ_j conditional on t_j , but the seller only consults bidder j in particular situations, namely when j is not a high bidder or when j is a high bidder but loses a tie break. Hence, μ_j is *not* the distribution of $(v_{-j}^{(1,2)}, |W(v_{-j})|)$ conditional on j being consulted. Rather, the mechanism takes into account the fact that $v_{-j}^{(1)} \geq b_j$ and that j must have lost any and all tie breaks. Finally, since the price can depend on any information from losing bidders, the mechanism additionally uses the fact that $v_{-j}^{(1)} \geq b_{-j}^{(2)}$. As previously discussed, by conditioning on the second highest bid of $-j$, the seller makes sure that the auction generates weakly greater revenue than a second-price auction with the optimal reserve price. These computations result in

an upper cumulative conditional distribution of the winner's value when j is consulted, which is

$$G_j \left(r; m_j, b_{-j}^{(2)} \right) = \sum_{\left\{ v_{-j}^{(1,2)}, k \mid \begin{array}{l} v_{-j}^{(1)} \geq \max\{r, b_{-j}^{(2)}\}, \\ v_{-j}^{(2)} = b_{-j}^{(2)} \end{array} \right\}} \frac{k}{\mathbb{I}_{b_j = v_{-j}^{(1)}} + k} \mu_j \left(v_{-j}^{(1,2)}, k \right). \quad (7)$$

This is the probability that the winner's value is at least r , conditional on bidder j being consulted and on $b_{-j}^{(2)}$. The price induced by bidder j 's report is a monopoly price with respect to this distribution:

$$r_j \left(m_j, b_{-j}^{(2)} \right) \in \arg \max_r r G_j \left(r; m_j, b_{-j}^{(2)} \right). \quad (8)$$

Note that $r_j \left(m_j, b_{-j}^{(2)} \right)$ is always at least $b_{-j}^{(2)}$, since $G_j \left(r; m_j, b_{-j}^{(2)} \right)$ is constant for $r \leq b_{-j}^{(2)} = \max \left\{ b_j, b_{-j}^{(2)} \right\}$.

4.3 Worst-case extraction ratio for the BSA

In this section, I characterize the minimum extraction ratio and the minimizing type spaces for the BSA under the truth-telling equilibrium. I will give an informal argument, with a rigorous proof in the Appendix.

The reports of the losing bidders contain information that the seller uses to optimally set the winner's price. In particular, the seller conditions on

- (1) Bidder j is not being offered the good, i.e., $v^{(1)} \geq v_j$ and j loses any tie breaks;
- (2) The winner's value is greater than the realized values $b_j = v_j$ and $b_{-j}^{(2)} = v_{-j}^{(2)}$;
- (3) The consulted bidder j 's realized type t_j .

These three pieces of information are incorporated into G_j and the optimized price r_j .

Note that the distribution of $v_{-j}^{(1)}$ conditioning on (1)-(3) will on average be the distribution of $v_{-j}^{(1)}$, given that j is not being offered the good, which is just (1). The fact that j is not offered the good means that $v_{-j}^{(1)} \geq v_j$ and j loses any tie breaks. This average distribution does not condition on the fact that v_j and $v_{-j}^{(2)}$ have particular realized values b_j and $b_{-j}^{(2)}$, respectively, and it does not incorporate the additional information contained in t_j .

An analogy with third-degree price discrimination highlights the impact of (1)-(3). The aggregate market is the distribution of $v_{-j}^{(1)}$ conditional on (1). Instead of having to price with respect to this aggregate market, the seller sees demand broken up into pieces conditional on (2) and (3). Such price discrimination is always beneficial to the seller, since the seller could ignore the extra information and set uniform prices. Hence, it is weakly worse for the seller to have less information, which is when b_j , $b_{-j}^{(2)}$, and t_j are less informative about $v^{(1)}$.

Given a particular type space \mathcal{T} with $S(\mathcal{T}) = S$, the seller generally learns more from (1)-(3) than from just (1). However, it is possible to find another type space \mathcal{T}' in which the distribution of $v^{(1)}$ conditional on (1) is the same, but in which the seller learns nothing from (2) and (3). This alternative type space \mathcal{T}' has the same efficient surplus, but revenue must be weakly lower $R(\mathcal{M}^{BSA}, \mathcal{T}') \leq R(\mathcal{M}^{BSA}, \mathcal{T})$, since the seller does not benefit by setting discriminatory reserve prices based on (2) and (3). Hence, the extraction ratio is lower as well.

How is this type space obtained? First, if bidders have more than one type t_j associated with a particular realization v_j , then \mathcal{T}' can be defined so that these types are effectively merged. In other words, bidder j 's types $t_j \in \phi_j^{-1}(v_j)$ are replaced with a single type, so that bidder j only learns that he has one of the types such that $\phi_j(t_j) = v_j$, i.e., \mathcal{T}' is a payoff type space with the same marginal distribution over values as \mathcal{T} . This makes bidder j 's report μ_j weakly less informative. Second, the realized values b_j and $b_{-j}^{(2)}$ are also informative. We can modify the type space so that every bidder except the winner has the minimum valuation in the support \underline{v} , meaning the type space is lopsided. Thus, the losing bidders' values are completely uninformative as lower bounds on the winner's valuation. Finally, there may be asymmetries wherein one bidder's losing report is on average more informative than others'. We can "symmetrize" the distribution so that all bidders' losing reports are equally informative. This is formalized in the following:

Lemma 1 (Worst-case type spaces). *For any $\epsilon > 0$, there exists a symmetric and lopsided payoff type space \mathcal{T} such that*

$$E(\mathcal{M}^{BSA}, \mathcal{T}) < \inf_{\mathcal{T}'} E(\mathcal{M}^{BSA}, \mathcal{T}') + \epsilon,$$

Any type space within the class described in Lemma 1 is of the following form: Pick one bidder $i \in N$ uniformly, and set $v_j = \underline{v}$ for $j \neq i$. Bidder i 's value is drawn from the distribution

$F^{(1)} \in \Delta(\mathbb{R}_+)$, which is the unconditional distribution of $v^{(1)}$. Losing bidders j always report the conditional belief that $v_{-j}^{(1)} \sim F^{(1)}$, and $b_j = b_{-j}^{(2)} = \underline{v}$. The reserve price r^* is a solution to

$$\max_{r \geq 0} r G^{(1)}(r),$$

where $G^{(1)}(r) = 1 - \lim_{v \uparrow r} F^{(1)}(v)$ is the probability of the offered price of r being “accepted” by the high bidder, when the bidder buys whenever indifferent. Revenue $R(\mathcal{T})$ is simply this maximum.

The spirit of Lemma 1 is that holding fixed the expected efficient surplus, there is a certain class of type spaces within which revenue can be minimized. It is now instructive to reverse the question: Suppose we wanted to maintain $R(\mathcal{T}) \leq R$. Which distributions $F^{(1)}$ will maximize $S(\mathcal{T})$ subject to this revenue constraint? It must be that for ever $r \geq 0$, $r G^{(1)}(r) \leq R$, so $F^{(1)} \geq 1 - \frac{R}{r}$. On the other hand, pushing down the cumulative distribution of $v^{(1)}$ always increases $\mathbb{E}[v^{(1)}] = S(\mathcal{T})$. Thus, the supremum of $E(\mathcal{M}^{BSA}, \mathcal{T})$ will be attained when $F^{(1)}$ is precisely

$$F^{(1)}(v) = \begin{cases} 0 & \text{if } v < R \\ 1 - \frac{R}{v} & \text{if } R \leq v < \bar{v} \\ 1 & \text{if } v \geq \bar{v} \end{cases}, \quad (9)$$

where \bar{v} is the largest valuation in the support, which is a truncated Pareto distribution with scale R and shape of 1. An example of such a distribution is given in Figure 1.

If the distribution of $v^{(1)}$ is given by (9), then the efficient surplus is given by the Riemann-Stieltjes integral

$$\begin{aligned} S &= \int_{v=R}^{\bar{v}} v \frac{R}{v^2} dv + \bar{v} \frac{R}{\bar{v}} \\ &= R(1 + \log(\bar{v}) - \log(R)). \end{aligned} \quad (10)$$

Note that (9) implies a mass point on \bar{v} of size $\frac{R}{\bar{v}}$. It is evident that S is increasing without bound in \bar{v} . If the distribution of values can be unbounded, then a fixed level of revenue is consistent with arbitrarily large efficient surplus. However, this requires putting a lot of mass on extremely large valuations, far from the efficient surplus. It is natural to ask how the extraction ratio behaves when there are limits to how dispersed values can be. One way to accomplish this goal is to require that values not be too much larger than S . This is formalized in Assumption 1:




Figure 1: An example of the lower bound on $F^{(1)}$ for $R = 1$ and $\bar{v} = 2$. The efficient surplus in this case is approximately 2.68, and the extraction ratio is approximately 0.37.

Assumption 1 (Bounded support). The support of values is contained in $[0, \gamma S(\mathcal{T})]^n$ for some $\gamma \geq 1$.

Let

$$\mathbf{T}(\gamma) = \{\mathcal{T} \mid \text{supp}(\phi_*\pi) \subset [0, \gamma S(\mathcal{T})]^n\},$$

where $\phi_*\pi$ is the pushforward measure on values, i.e., the distribution on values induced by the distribution π and the mapping ϕ . Under this assumption, γS and E can be substituted for \bar{v} and $\frac{R}{S}$, so that (10) becomes

$$E(1 + \log(\gamma) - \log(E)) = 1. \tag{11}$$

This equation represents an accounting identity. When the highest value is drawn from (9), the expected efficient surplus S must be equal to the Riemann-Stieltjes integral with respect to (9), where γS is the upper limit of the integral. As shown in the proof of the following proposition, this equation has a unique solution, denoted $E^*(\gamma)$. We have the following:

Proposition 1 (min extraction ratio for BSA). *For any $\gamma > 0$, the worst-case extraction ratio for the BSA under the truth-telling equilibrium when type spaces are restricted to $\mathbf{T}(\gamma)$ is the unique $E^*(\gamma)$ which solves (11). This extraction ratio is attained by type spaces of the form described in Lemma 1, with the distribution of the highest value approaching (9).*

This concludes the characterization of the worst-case extraction ratio for the BSA.

4.4 max min extraction ratio

In fact, for the worst-case type spaces, in which at most one bidder has a positive value, \mathcal{M}^{BSA} is an optimal auction. On these type spaces, the seller's problem is formally equivalent to the selling of a single unit to a single buyer. It is well known that the optimal mechanism is a posted price, and the reports of the losers allow the seller to set the optimal price (cf. Riley and Zeckhauser, 1983).

This implies that the lower bound of $E^*(\gamma)$ is tight. In general,

$$\sup_{\mathcal{M}} \inf_{\mathcal{T} \in \mathbf{T}(\gamma)} E(\mathcal{M}, \mathcal{T}) \leq \inf_{\mathcal{T} \in \mathbf{T}(\gamma)} \sup_{\mathcal{M}} E(\mathcal{M}, \mathcal{T}),$$

since any mechanism that the seller chooses when forced to move first could also be chosen when moving second, and therefore guarantee at least as large of a payoff. For many problems, it turns out that the inequality is in fact an equality, as in the minimax theorems of zero sum games. This is not automatically the case here since the setup does not satisfy the regularity conditions of the minimax theorems known to the author.⁵

However, a solution to the LHS is given by the seller using \mathcal{M}^{BSA} and Nature choosing a type space satisfying the conditions of Lemma 1. Moreover, for the RHS, Nature could always use the same type spaces, and the seller can do no better than with \mathcal{M}^{BSA} . Hence, the two sides are in fact equal. This observation, combined with Proposition 1 gives us the following:

Theorem 1 (Partial max min extraction ratio). *The solution to (4) restricted to $\mathcal{T} \in \mathbf{T}(\gamma)$ is no greater than $E^*(\gamma)$. Hence, the BSA under the truth-telling equilibrium partially solves (4).*

⁵Specifically, von Neumann's minimax theorem only applies to finite domains, and Sion's minimax theorem requires the domains to be linear topological spaces. This structure is lacking on type spaces and mechanisms.

4.5 Characterizing performance

The number $E^*(\gamma)$ gives the maxmin extraction ratio that the seller is guaranteed by using the BSA. But is this lower bound economically meaningful? It would be useful to know that the lower bound guarantees the seller a substantial revenue-share of the efficient surplus. On the right panel of Figure 2, $E^*(\gamma)$ is plotted for values ranging from 1 to 50. The left hand panel gives E^* for six values of γ . We see that $E^*(\gamma)$ is monotonically decreasing, quickly for small γ , with the rate of decrease falling rapidly. For example, going from $\gamma = 1.1$ to $\gamma = 2$ entails a decreasing from $E^* = 0.67$ to $E^* = 0.37$, whereas the difference between $\gamma = 100$ and $\gamma = 1,000$ is only 3 percentage points. In the latter case, the seller is guaranteed at least a 10% revenue-share of the efficient surplus. Even for $\gamma = 10,000,000$, the seller is guaranteed approximately a 5% revenue-share.

The slow rate of decay of $E^*(\gamma)$ can be formalized as follows. Asymptotically,

$$E^*(\gamma) = O\left(\frac{1}{\log(\gamma)}\right),$$

This follows from (11), since

$$\begin{aligned} 1 &= \lim_{\gamma \rightarrow \infty} \frac{\frac{1}{E^*(\gamma)} + \log(E^*(\gamma))}{1 + \log(\gamma)} \\ &= \lim_{\gamma \rightarrow \infty} \frac{\frac{1}{E^*(\gamma)} + \log(E^*(\gamma))}{\log(\gamma)} \end{aligned}$$

since the expression inside the limit is equal to 1 for all γ . It is easy to see that $\lim_{\gamma \rightarrow \infty} E^*(\gamma) = 0$, since if it were bounded away from 0, the left hand side of (11) would blow up. Hence, by L'Hôpital's rule,

$$\lim_{\gamma \rightarrow \infty} \frac{\frac{1}{E^*(\gamma)}}{\frac{1}{E^*(\gamma)} + \log(E^*(\gamma))} = \lim_{\gamma \rightarrow \infty} \frac{1}{1 + E^*(\gamma)} = 1$$

where the derivative exists because of the implicit function theorem. Since both limits exist, the

γ	$E^*(\gamma)$
1.1	0.6656
2	0.3734
10	0.2045
100	0.1309
1,000	0.0977
10^7	0.0497

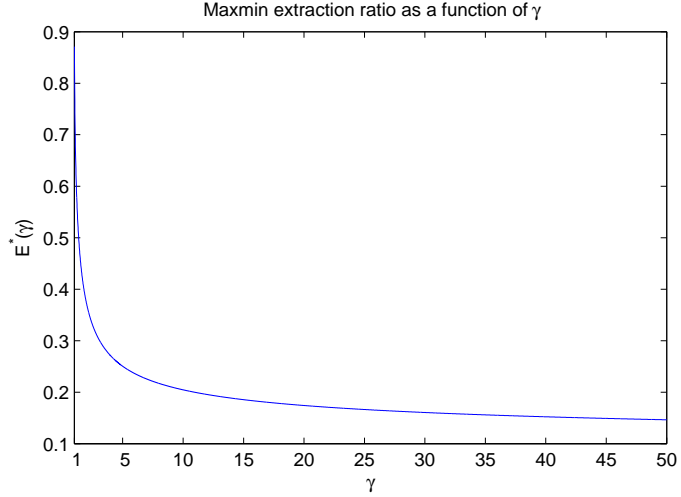


Figure 2: $E^*(\gamma)$ for a range of values.

limit of the product is equal to the product of the limits, and

$$\begin{aligned}
 1 &= \lim_{\gamma \rightarrow \infty} \frac{\frac{1}{E^*(\gamma)} + \log(E^*(\gamma))}{\log(\gamma)} \frac{\frac{1}{E^*(\gamma)}}{\frac{1}{E^*(\gamma)} + \log(E^*(\gamma))} \\
 &= \lim_{\gamma \rightarrow \infty} \frac{\frac{1}{\log(\gamma)}}{E^*(\gamma)}
 \end{aligned}$$

which proves the result. In sum, $E^*(\gamma)$ goes to zero exponentially slower than the rate of growth of γ , so even for very generous bounds on the dispersion in values, the seller will still be guaranteed a substantial share of the efficient surplus.

Another potential concern is that the type spaces used to achieve the lower bound are highly stylized. Symmetry of the type space is not necessary, but what is necessary is the consequence of the lopsided property, that there is a large gap between the average highest and second-highest values. In some situations, this property could be quite natural. In the school construction example from the introduction, it is possible that there is one dominant contractor that tends to have the lowest cost, and this asymmetry is common knowledge among the firms. Nonetheless, in many situations one would not expect to find such a large gap. Is the extraction ratio in the worst-case radically different from extraction ratios in the kinds of environments that are more frequently modeled?

Figure 3 gives examples of four different distributions over values in which the gap between highest and second-highest values is modest. Each type space has a different efficient surplus and

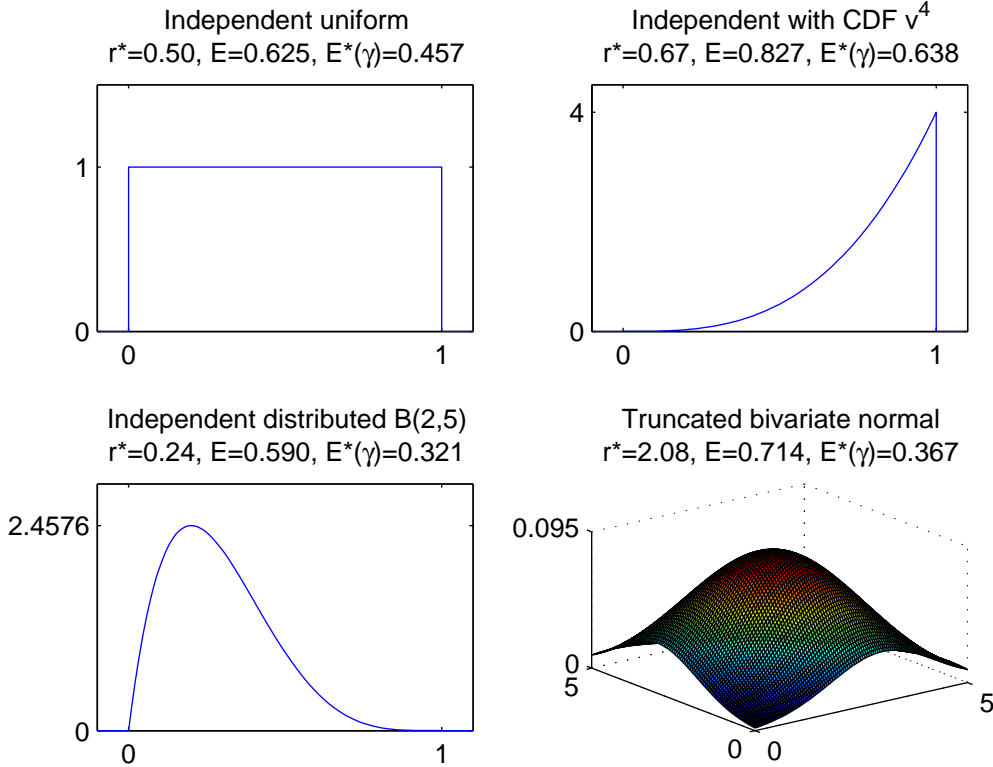


Figure 3: Examples of the extraction ratio on particular type spaces in $\mathbf{T}(\gamma)$ versus $E^*(\gamma)$. For the first three frames, values are i.i.d. and drawn from the depicted PDF. In the fourth frame, values are correlated and drawn from the depicted joint PDF. In each case, E^* is on the same order of magnitude as the actual extraction ratio that would obtain with the second-price auction with an optimal anonymous reserve.

consequently different $\gamma = \frac{\bar{v}}{S}$. For example, in the first panel values are independent and uniformly distributed on $[0, 1]$. With two bidders, the efficient surplus is 0.67, so $\gamma = 1.5$. For each of these type spaces, the extraction ratio the seller obtains with a second-price auction and the optimally chosen anonymous reserve price is compared to the lower bound guaranteed by the BSA for the same γ . In the uniform example, the seller could set the optimal reserve of 0.5 and obtain the optimal extraction ratio of 0.625. In contrast, $E^*(1.5) = 0.4569$. The point of these examples is that although the worst-case type spaces are stylized, the lower bound is not orders of magnitude different from the extraction ratio on “typical” examples with similar γ .

In fact, there are classes of environments, namely symmetric payoff type spaces with regular and independent distributions, in which the BSA will implement the optimal auction. The reason is simply that each bidder will report the independent distribution from which other bidders’ values

are drawn, and the seller will set the winner's price equal to the maximum of the second-highest bid and the optimal reserve price, which is where the virtual valuation is zero. More generally, in any type space the BSA has to generate at least as much revenue as a second-price auction with an optimal uniform reserve price. The seller's pricing problem in the BSA is broken up into a bunch of conditional pricing problems. For each of these problems, it is always feasible for the seller to set the price equal to $\max\{r^*, b^{(2)}\}$, where r^* is an anonymous reserve price. Hence, the optimal pricing rule for each of these problems must generate weakly more revenue than the fixed reserve rule. We have the following:

Proposition 2 (Comparison with second-price auction). *Expected revenue in the BSA when the seller does not know the prior is weakly greater than expected revenue of the second-price auction when the seller knows the prior and sets the optimal anonymous reserve price. If the distribution of values is independent, symmetric, and regular, then the BSA is an optimal auction.*

The bottom line is that the BSA guarantees the seller a relevant lower bound on the worst-case extraction ratio, and also does not greatly disadvantage the seller away from the worst-case. Revenue is always weakly better than in the second-price auction with anonymous reserve, which is probably the most widely used auction format in the world and is known to be an optimal auction in benchmark environments.

4.6 Equilibrium uniqueness

In this section, I extend the partial max min result of Theorem 1 to full max min. This is facilitated by simple perturbations of \mathcal{M}^{BSA} that make the truth-telling equilibrium unique. Specifically, I will construct a mechanism \mathcal{M}^ϵ for every $\epsilon > 0$ with the message space M^{BSA} . This mechanism implements the same allocation and transfers as \mathcal{M}^{BSA} with probability $1 - \epsilon$, but is perturbed in such a way that truth-telling is the unique strategy profile that survives iterated deletion of strictly dominated strategies. In particular, the message space is $M^\epsilon = M^{BSA}$, and the allocation and payoff rules are

$$q^\epsilon(m) = (1 - \epsilon)q^{BSA}(m) + \epsilon q^1(m),$$

$$p^\epsilon(m) = (1 - \epsilon)p^{BSA}(m) + \epsilon p^1(m) + \epsilon p^2(m),$$

where q^1 , p^1 , and p^2 will be defined presently.

Truthful bidding is a weakly dominant strategy of \mathcal{M}^{BSA} . The functions q^1 and p^1 make it strictly dominant, by adding a small probability event that each bidder is selected to be offered the good at a price drawn from a distribution $G(r)$ with positive density $g(r)$ and support equal to \mathbb{R}_+ . Specifically, define

$$q_i^1(m) = \frac{1}{n} \int_{r=0}^{b_i} g(r) dr,$$

$$p_i^1(m) = \frac{1}{n} \int_{r=0}^{b_i} r g(r) dr.$$

Since p^2 will not depend on b_i at all, $b_i = v_i$ is uniquely optimal. This trick is similar to that used by Bergemann and Morris (2012a).

The second new component of the transfer p^2 is a modified scoring rule that rewards bidders for correctly guessing the distribution of $(b_{-i}^{(1,2)}, |W(b_{-i})|)$. I say a modified scoring rule, as the transfer is weighted it so that the seller never has to pay too much in expectation to incentivize bidders to report their beliefs. In particular,

$$p_i^2(m) = \frac{1}{n} \left[\sum_{v_{-i} \in \mathbb{R}_+^{n-1}} v_{-i}^{(1)} \frac{\left(\mu \left(v_{-i}^{(1,2)}, |W(v_{-i})| \right) \right)^2}{2} - b_{-i}^{(1)} \mu_i \left(b_{-i}^{(1,2)}, |W(b_{-i})| \right) \right].$$

Since bidders report $b_{-i} = v_{-i}$ in equilibrium, by reporting μ_i , bidder i 's expected payoff is

$$\mathbb{E}[p_i^2(m)|t_i] = \frac{1}{n} \sum_{v_{-i} \in \mathbb{R}_+^{n-1}} v_{-i}^{(1)} \left[\frac{\left(\mu \left(v_{-i}^{(1,2)}, |W(v_{-i})| \right) \right)^2}{2} - \pi \left(v_{-i}^{(1,2)}, |W(v_{-i})| \mid t_i \right) \mu_i \left(v_{-i}^{(1,2)}, |W(v_{-i})| \right) \right],$$

so the first-order condition implies (as long as $v_{-i}^{(1)} > 0$) that the type t_i reports

$$\mu_i \left(v_{-i}^{(1,2)}, k \right) = \pi \left(v_{-i}^{(1,2)}, t_i \right),$$

for all $v_{-i}^{(1,2)}$ and k . There is a unique v_{-i} such that $v_{-i}^{(1)} = 0$. Since bidders must report a distribution, they report the probability of this event accurately as well. Thus, \mathcal{M}^ϵ has a truth-telling as the unique equilibrium for all \mathcal{T} .

Finally, observe that in equilibrium, it must be that

$$\begin{aligned}
\mathbb{E}[p_i^2(m)] &= -\frac{1}{n} \sum_{t_i} \pi_i(t_i) \sum_{v_{-i} \in \mathbb{R}_+^{n-1}} v_{-i}^{(1)} \frac{\left(\pi \left(v_{-i}^{(1,2)}, |W(v_{-i})| \mid t_i \right) \right)^2}{2} \\
&\leq -\frac{1}{n} \sum_t \phi_{-i}^{(1)}(t) \pi(t) \\
&\leq -\frac{1}{n} \sum_t \phi^{(1)}(t) \pi(t) \\
&= -\frac{1}{n} S(\mathcal{T}).
\end{aligned}$$

On average, the seller makes such a transfer for each of the n bidders, at a total cost of at most $-S(\mathcal{T})$. Since $\mathbb{E}[p^1(m)] \geq 0$, it will be true that

$$E(\mathcal{M}^\epsilon, \mathcal{T}) \geq E(\mathcal{M}^{BSA}, \mathcal{T}) - 2\epsilon.$$

The results of Proposition 1 and Theorem 1, together with the fact that \mathcal{M}^ϵ has a unique equilibrium that always has an extraction ratio within 2ϵ of the truth-telling equilibrium of \mathcal{M}^{BSA} imply the following theorem:

Theorem 2 (Full max min extraction ratio). *$E^*(\gamma)$ is the solution to (4), and the mechanisms \mathcal{M}^ϵ guarantee the seller an extraction ratio that is at least $E^*(\gamma) - 2\epsilon$.*

4.7 Conditional preferences

The max min extraction ratio criterion implicitly assumes that the seller compares expected revenues for different type spaces not one-for-one, but relative to the expected efficient surplus. I regard this as reasonable if the seller uses the expected efficient surplus as the target for revenue. However, the manner in which the seller compares revenue on type spaces with different surpluses is not essential to my results.

In this Section, I consider a much weaker preference in which the seller does not compare outcomes across different levels of the efficient surplus. Define

$$\mathbf{T}(S, \gamma) = \left\{ \mathcal{T} \mid \begin{array}{l} S(\mathcal{T})=S \\ \text{supp}(\phi_*\pi) \subseteq [0, \gamma S]^n \end{array} \right\}$$

Consider the incomplete preference over mechanisms, where mechanism \mathcal{M} is weakly preferred to mechanism \mathcal{M}' (denoted $\mathcal{M} \succeq \mathcal{M}'$) if for every (S, γ) ,

$$\inf_{\mathcal{T} \in \mathbf{T}(S, \gamma)} \inf_{\sigma \in \text{BNE}(\mathcal{M}, \mathcal{T})} R(\mathcal{M}, \mathcal{T}, \sigma) \geq \inf_{\mathcal{T} \in \mathbf{T}(S, \gamma)} \inf_{\sigma \in \text{BNE}(\mathcal{M}', \mathcal{T})} R(\mathcal{M}', \mathcal{T}, \sigma) \quad (12)$$

This preference is a partial ordering, since two mechanisms \mathcal{M} and \mathcal{M}' are incomparable if (12) holds for one (S, γ) , but not for (S', γ') . However, $\mathcal{M} \succeq \mathcal{M}'$ indicates a strong notion of dominance in that for every efficient surplus, \mathcal{M} performs better in terms of worst-case revenue. A mechanism is maximal in the ordering \succeq if it solves

$$\sup_{\mathcal{M}} \inf_{\mathcal{T} \in \mathbf{T}(S, \gamma)} \inf_{\sigma \in \text{BNE}(\mathcal{M}, \mathcal{T})} R(\mathcal{M}, \mathcal{T}, \sigma) \quad (13)$$

for every (S, γ) . Denote the solution to this problem by $R^*(S, \gamma)$. My next result is that the mechanisms \mathcal{M}^ϵ are virtually maximal according to \succeq :

Proposition 3. *The solution to (13) is $R^*(S, \gamma) = E^*(\gamma)S$. For every (S, γ) ,*

$$\inf_{\mathcal{T} \in \mathbf{T}(S, \gamma)} \inf_{\sigma \in \text{BNE}(\mathcal{M}^\epsilon, \mathcal{T})} R(\mathcal{M}^\epsilon, \mathcal{T}, \sigma) \geq (E^*(\gamma) - 2\epsilon)S.$$

The proof is quite straightforward.

Proof. For the first part, clearly it cannot be that $R^*(S, \gamma) < E^*(\gamma)S$, since this implies that

$$\sup_{\mathcal{M}} \inf_{\mathcal{T}} \inf_{\sigma \in \text{BNE}(\mathcal{M}, \mathcal{T})} \frac{R(\mathcal{M}, \mathcal{T}, \sigma)}{S} \leq \sup_{\mathcal{M}} \inf_{\mathcal{T} \in \mathbf{T}(S, \gamma)} \inf_{\sigma \in \text{BNE}(\mathcal{M}, \mathcal{T})} \frac{R(\mathcal{M}, \mathcal{T}, \sigma)}{S} < E^*(\gamma).$$

Moreover, if $R^*(S, \gamma) > E^*(\gamma)S$ for some (S, γ) , then it would have to be strictly larger for every (S, γ) , since every type space $\mathcal{T} = (T, \phi, \pi) \in \mathbf{T}(\gamma)$ with strictly positive $S(\mathcal{T})$ can be mapped to some $\mathcal{T}' = (T', \phi', \pi') \in \mathbf{T}(S, \gamma)$ by defining $T' = T$, $\pi' = \pi$, and $\phi'(t) = \phi(t) \frac{S}{S(\mathcal{T})}$. Hence, if $R^*(S, \gamma) > (E^*(\gamma) + \epsilon)S$ for some S , then $\frac{R^*(S, \gamma)}{S} > E^*(\gamma)$ for all S , a contradiction.

The second part follows almost directly, since the extraction ratio is invariant to scaling of valuations as in the previous paragraph. As \mathcal{M}^ϵ achieves $E^*(\gamma) - 2\epsilon$ for some value of S , it must achieve the same extraction ratio for all S , and therefore revenue is at least $(E^*(\gamma) - 2\epsilon)S$. \square

In sum, it does not matter how the seller compares revenue across environments in which the expected social value of the good is different. As long as the seller has max min preferences over revenue for a fixed expected efficient surplus, the BSA is an optimal auction.

5 Discussion

5.1 Belief extraction

This paper has been focused on the selection of a mechanism by a seller who evaluates mechanisms by their worst-case extraction ratio. Under such preferences, the seller is tolerant of suboptimal extraction ratios on particular type spaces, as long as this ratio is greater than the worst-case. The BSA can result in such suboptimal extraction ratios, since it collects rather limited information about the environment and therefore will not maximize revenue on most type spaces. In particular, the seller only asks each bidder to estimate the distribution of the top two order statistics of other bidders' values, and the number of ties. In principle, the seller could have collected information more ambitiously. Is there a limit to how much the seller could learn, by asking the bidders more complicated questions? Could the seller, for example, collect enough information and in such a way that a revenue maximizing auction is always implemented?

This question is related to a folk argument that has existed in the mechanism design literature: If a common prior is known to the agents, but not to the designer, then the prior could be extracted by the designer for free (Bergemann and Morris, 2012a). By “free”, I mean that having to incentivize truthful revelation of the prior does not restrict the class of social choice functions that can be implemented, according to any solution concept. There is an obvious partial implementation solution to this problem: Ask all of the agents to simultaneously announce the prior, and if they disagree, punish all of the agents severely. Of course, this is not entirely satisfactory because this mechanism would also enforce coordinated misreporting of the prior, which is counter to the full implementation philosophy of the present work.

However, in a related paper, Brooks (2013a) provides a stronger resolution of the folk argument. In general, it is possible for the seller to extract agents' beliefs in such a way that the common prior is revealed to the seller in *every* equilibrium. A caveat is that the general mechanism accomplishing this goal is quite complicated, and requires the seller to elicit each agent's infinite hierarchy of

beliefs. Infinite iterated deletion of strictly dominated strategies forces the agents to report their hierarchy truthfully in any equilibrium. Nonetheless, if such complex mechanisms are permitted, then it is possible for the seller to implement a revenue-maximizing mechanism for every realized type space, as long as the common prior is known to the agents.

The folk argument is also related to the works of Neeman (2003) and Azar et al. (2012). If the seller can extract the prior for free, then it is possible to implement any mechanism, including the second-price auction with an optimally chosen anonymous reserve price or the optimal dominant strategy and ex-post individually rational mechanism. The distinct contribution of this paper is to show that for a particular criterion, i.e., max min extraction ratio, the seller need not use such an elaborate mechanism. It is sufficient for the seller to extract simple statistics, and use these statistics to guard against the downside risk associated with a large gap between the highest and second-highest values, as in lopsided type spaces.

5.2 Simpler mechanisms

The previous section discussed more complicated mechanisms that the seller could use to achieve optimal performance in a wider range of environments. But what about the other direction: Are there classes of environments in which the BSA can be further simplified, without greatly compromising performance?

There are at least two dimensions along which the BSA can be easily simplified. First, because of the selection effect induced by ties, the BSA needs to extract bidders beliefs about the number of bidders who will make high bids. Ties would not occur with positive probability if valuations were drawn from a non-atomic distribution, or if the finite supports of values were non-overlapping. Furthermore, ties do not occur with positive probability in the worst-case type spaces for extraction ratios. If attention is restricted to type spaces in which ties occur with zero probability, then clearly the seller can get away with just extracting bidders beliefs about $b_{-i}^{(1,2)}$. Second, extracting beliefs about $b_{-i}^{(2)}$ is necessary to make sure that the BSA does as well as the second-price auction. If the seller is purely concerned with max min extraction ratio, then the seller could extract beliefs about the just highest bid of others, and set an optimal reserve price conditional on $b^{(1)} = b_{-j}^{(1)} \geq b_j$. The same analogy with third-degree price discrimination applies, and the seller must obtain at least as much revenue as if the seller sold the good to the high bidder at the optimal monopoly price with

respect to $F^{(1)}$.

Perhaps the most natural method of aligning the incentives of the seller and buyers would be to give buyers a direct stake in the revenue generated by their reports. The seller could for example share a small portion of revenue with bidder j whenever a sale is made with a price based on bidder j 's report. However, this creates complicated incentives to influence the allocation of the good: For example, in the BSA, the marginal event affected by bidding $b_i = v_i$ is when the bidder is allocated the good at a price equal to v_i . In this case, the marginal surplus from being allocated the good is zero. The bidder may have an incentive to “throw” the auction at such marginal events, so as to instead obtain a positive share of revenue from selling to others at price v_i . Brooks (2013b) studies auctions of this form, and shows that in reasonably structured environments, an equilibrium exists in which the bidding strategy equates the marginal surplus from the allocation and the marginal surplus from sharing in revenue.

Finally, the BSA sets a price using the interim beliefs of a single losing bidder, combined with the relatively sparse information of the second-highest bid amongst all other bidders. It is natural to ask if there is a straightforward way to aggregate all of the losers' information, so that the seller sets an optimal price for the winning bidder i conditional on t_{-i} . This could easily be accomplished by having bidders report their entire hierarchy of beliefs as described above. Unfortunately, there is not an obvious simpler solution. One possibility would be to allow the losing bidders to “converse” about the optimal price, by iteratively reporting their conditional beliefs about the winner's value. Arguments in the vein of Geanakoplos and Polemarchakis (1982) would show that if such communication was allowed over multiple rounds, the bidders would eventually agree on a posterior. However, this posterior need not coincide with the true posterior conditional on t_{-i} . Moreover, if only losing bidders can have this conversation and receive the rewards that incentivize truth-telling, then these rents could create an incentive for bidders to throw the auction. Even so, such mechanisms are a promising direction for future research.

5.3 Common values

The assumption of private values is more appropriate in some settings than others. In the motivating example of a school construction project, it is reasonable to suppose there are private value components to firms' costs, such as prior commitments, worker abilities, etc. However, the firms

might also have a common value in the idiosyncrasies of the project, such as the suitability of the land on which the school is to be built. Auction design with interdependent values can be challenging due to the buyers updating their preferences upon winning the auction. I briefly sketch the scope for generalizing my results to this broader setting. I have concluded that given reasonable assumptions on the interdependence, the same max min extraction ratio obtains even if Nature is allowed to choose type spaces with interdependent preferences, although a much more complicated mechanism is required than the BSA.

As in much of the literature, I distinguish between “information” types t_i and “payoff” types $\theta_i \in \Theta_i$. A buyer’s valuation is a function $\phi_i(\theta)$ of the profile of payoff types but does not depend on t . Thus, the definition of a type space is expanded to $\mathcal{T} = (\Theta, T, \pi, \phi)$. With interdependent values, the seller needs to elicit not just bidders’ beliefs about θ , but also the form of the interdependence, i.e., $\phi(\theta)$.

Let us suppose for the moment that the seller knows Θ and ϕ . Many of the positive results in the literature require the assumption that Θ_i is one-dimensional, and that $\phi_i(\theta)$ is monotonically increasing.⁶ This assumption, combined with a single-crossing property on ϕ , is sufficient for the existence of an efficient equilibrium of the English auction (Dasgupta and Maskin, 2000; Maskin, 1992; Krishna, 2003; Birulin and Izmalkov, 2011). Starting from this efficient equilibrium, the seller can partially implement a mechanism similar to the BSA, where bidders report θ_i and beliefs about θ_{-i} . These reported beliefs can be used to find the reserve price for when bidder i wins that maximizes

$$r \Pr \left(\left\{ \tilde{\theta}_i \mid \phi_i \left(\tilde{\theta}_i, \theta_{-i} \right) \geq r \right\} \mid \theta_{-i}, t_{-i} \right)$$

The seller only sells the good to bidder i at this price if the realized value conditional on θ is at least r . Since the seller optimally accepts this price on behalf of the winning bidder, and the price does not depend on bidder i ’s report, truthful reporting is still an equilibrium. I will call this the interdependent belief survey auction (IBSA).

Under this mechanism, the efficient surplus and revenue with an interdependent value type space $\mathcal{T} = (\Theta, T, v, \pi)$ are the same as under the BSA with a particular private value type space

⁶For multi-dimensional Θ_i , existence of mechanisms with efficient equilibria becomes problematic. See Jehiel and Moldovanu (2001).

$\mathcal{T}' = (T', \phi', \pi')$. Intuitively, I would like to find a private value type space which has the same distribution of the winner's value conditional on losers' information. The bidders cannot simply be told their ex-post values are, since they will reveal this to the seller, who might then be able to identify the winner's value from losers' reported values. However, a buyer can be told his value $\phi_i(\theta)$ on events where he wins, while losing bidders observe θ_i but receive private values values of zero. Under such a private value type space, losers' beliefs about the winner's value are the same as under \mathcal{T} , and the distribution of the highest value is the same. This discussion implies that the IBSA achieves the same extraction ratio with \mathcal{T} as the BSA achieves on the private value type space \mathcal{T}' . Hence, when minimizing the extraction ratio for this richer mechanism, it is sufficient to look at private value type spaces, and therefore the same lower bound $E^*(\gamma)$ obtains.

Finally, I return to the issue of extracting the form of the interdependence. The seller needs a general “detail-free” language in which to have the buyers communicate what they know. The preference hierarchies of Bergemann et al. (2011) are just such a language, in which bidders report a sequence of state-dependent preferences over Anscombe-Aumann acts. At the first level, the preference is over a state space with a single element corresponding to (θ_i, t_i) . This preference is the player's willingness to pay for the good unconditional on other buyers' information. The second-order preference is over acts that depend on the first-order preferences of other buyers, and so on. At each level, types are separated by their preferences conditional on what others have revealed about their types, and these separated types can then be used to separate more types, in a manner analogous to Abreu and Matsushima (1992). It turns out that bidders can be given strict incentives to truthfully reveal their interdependent preferences in this language, using the techniques of Subsection 4.6, and these reports can be used as an input to construct the IBSA. It is important to note that the seller can only provide strict incentives to recover a coarsened state space which corresponds to the distinguishable types of Bergemann et al. (2011, 2012).

5.4 The role of the common prior

I have assumed throughout that the buyers' beliefs are derived from a common prior. A natural question to ask is whether or not my results can be extended to environments in which beliefs do not satisfy this restriction. The meaning of the common prior has been debated and critiqued in the literature (see Aumann, 1987; Morris, 1995; Gul, 1998). There are two possible interpretations:

One is that there is some ex-ante stage before private information is realized, at which point there is common knowledge of the distribution over future private information. With such a temporal structure, the consistency of interim beliefs with a common prior is a consequence of ex-ante common knowledge. In the other interpretation, there is no ex-ante stage, but rather the common prior is a restriction on agents' higher-order beliefs (see Samet (1998) for a characterization of the common prior in terms of interim beliefs).

In my model, even if there is an ex-ante stage, the seller only interacts with the potential buyers *after* private information is realized. Hence, anything the seller learns about the prior must be obtained through the buyers' interim beliefs. The most natural interpretation is that there is some physical process which generates signals that the buyers see, and this process is known to the buyers but not the seller. As such, all of the bidders' interim beliefs about the profile of valuations are distributed around the average belief generated by this signal structure.

The prior distribution provides a neutral perspective from which the seller can calculate expected revenue and efficient surplus. Also, since the buyers' reported beliefs will average to the prior, their interim reports give the seller access to a "segmentation" of the prior distribution of the highest valuation. Without a common prior, it would still be possible to elicit interim beliefs from the agents, but these reports would require a more complex interpretation to be useful to the seller. Thus, there is no immediate generalization of my result to non-common prior type spaces. However, as discussed by Azar et al. (2012), the result does not require each agent's entire hierarchy of beliefs to be consistent with a common prior. A similar result would obtain as long as the agents' first-order beliefs average to the same prior over values, and the seller uses this prior to calculate expected revenue and the expected efficient surplus.

6 Conclusion

This paper has considered the mechanism design problem faced by an uninformed seller, who believes that agents are well-informed. The seller uses mechanisms that survey bidders' beliefs in addition to their private values. This information is used to optimize the prices offered to winning bidders. Such mechanisms achieve an optimal lower bound on the share of the efficient surplus that the seller can extract as revenue, and also perform well away from the worst-case.

In practice, auctions are much more complicated than the stylized models studied by economists. For example, in the case of auctions for government contracts or natural resources, the auction designers conduct extensive research into demand, and surely engage in informal discussion with potential buyers about the pros and cons of different formats. This paper has characterized particular ways in which the seller can elicit useful information from the buyers without distorting incentives to bid truthfully. But there is a more general message to be gleaned: If the seller has sufficient commitment power with regard to how information will be used, then it is indeed possible to have these informal discussions without allowing for adverse manipulation of the auction format.

References

- Abreu, Dilip and Hitoshi Matsushima (1992), “Virtual implementation in iteratively undominated strategies: Incomplete information.”
- Aumann, Robert J (1987), “Correlated equilibrium as an expression of bayesian rationality.” *Econometrica: Journal of the Econometric Society*, 1–18.
- Azar, Pablo, Jing Chen, and Silvio Micali (2012), “Crowdsourced bayesian auctions.” In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, 236–248, ACM.
- Baliga, Sandeep and Rakesh Vohra (2003), “Market research and market design.” *Advances in Theoretical Economics*, 3.
- Bergemann, Dirk and Stephen Morris (2009), “Robust implementation in direct mechanisms.” *The Review of Economic Studies*, 76, 1175–1204.
- Bergemann, Dirk and Stephen Morris (2011), “Robust implementation in general mechanisms.” *Games and Economic Behavior*, 71, 261–281.
- Bergemann, Dirk and Stephen Morris (2012a), “Robust mechanism design: An introduction.” In *Robust Mechanism Design: The Role of Private Information and Higher Order Beliefs* (Dirk Bergemann and Stephen Morris, eds.), volume 2, World Scientific Publishing Company.
- Bergemann, Dirk and Stephen Morris (2012b), *Robust Mechanism Design: The Role of Private Information and Higher Order Beliefs*, volume 2. World Scientific Publishing Company.
- Bergemann, Dirk, Stephen Morris, and Satoru Takahashi (2011), “Interdependent preferences and strategic distinguishability.”
- Bergemann, Dirk, Stephen Morris, and Satoru Takahashi (2012), “Efficient auctions and interdependent types.” *The American Economic Review*, 102, 319–324.
- Bergemann, Dirk and Karl Schlag (2011), “Robust monopoly pricing.” *Journal of Economic Theory*, 146, 2527–2543.

- Birulin, Oleksii and Sergei Izmalkov (2011), “On efficiency of the english auction.” *Journal of Economic Theory*, 146, 1398–1417.
- Börger, Tilman (2013), “An introduction to the theory of mechanism design.” *Teaching Notes, University of Michigan (authors personal website)*.
- Brooks, Benjamin A. (2013a), “Extracting common priors: Formalizing a folk argument.” *mimeo*.
- Brooks, Benjamin A. (2013b), “Revenue sharing in second-price auctions.” *mimeo*.
- Caillaud, Bernard and Jacques Robert (2005), “Implementation of the revenue-maximizing auction by an ignorant seller.” *Review of Economic Design*, 9, 127–143.
- Carroll, Gabriel (2012), “Robustness and linear contracts.” Technical report, Stanford University Working Paper.
- Chassang, Sylvain (2011), “Calibrated incentive contracts.” *Princeton University Economic Theory Center Working Paper*.
- Chawla, Shuchi, Jason D Hartline, and Robert Kleinberg (2007), “Algorithmic pricing via virtual valuations.” In *Proceedings of the 8th ACM conference on Electronic commerce*, 243–251, ACM.
- Choi, Jaewon and Taesung Kim (1999), “A nonparametric, efficient public good decision mechanism: Undominated bayesian implementation.” *Games and Economic Behavior*, 27, 64–85.
- Chung, Kim-Sau and Jeffrey C Ely (2007), “Foundations of dominant-strategy mechanisms.” *The Review of Economic Studies*, 74, 447–476.
- Crémer, Jacques and Richard P McLean (1988), “Full extraction of the surplus in bayesian and dominant strategy auctions.” *Econometrica: Journal of the Econometric Society*, 1247–1257.
- Dasgupta, Partha and Eric Maskin (2000), “Efficient auctions.” *The Quarterly Journal of Economics*, 115, 341–388.
- Farinha Luz, Vitor (2013), “Surplus extraction with rich type spaces.” *Journal of Economic Theory*.
- Federal Communications Commission (1997), “The FCC report to Congress on spectrum auctions.” *WT Docket*.

- Geanakoplos, John D and Heraklis M Polemarchakis (1982), “We can’t disagree forever.” *Journal of Economic Theory*, 28, 192–200.
- Goldberg, Andrew V and Jason D Hartline (2003), “Envy-free auctions for digital goods.” In *Proceedings of the 4th ACM conference on Electronic commerce*, 29–35, ACM.
- Goldberg, Andrew V, Jason D Hartline, Anna R Karlin, and Michael Saks (2004), “A lower bound on the competitive ratio of truthful auctions.” In *STACS 2004*, 644–655, Springer.
- Gul, Faruk (1998), “A comment on aumann’s bayesian view.” *Econometrica*, 66, 923–927.
- Hartline, Jason D (2012), “Bayesian mechanism design.” *Foundations and Trends® in Theoretical Computer Science*, 8, 143–263.
- Hartline, Jason D and Tim Roughgarden (2009), “Simple versus optimal mechanisms.” In *Proceedings of the 10th ACM conference on Electronic commerce*, 225–234, ACM.
- Jehiel, Philippe and Benny Moldovanu (2001), “Efficient design with interdependent valuations.” *Econometrica*, 69, 1237–1259.
- Krishna, Vijay (2003), “Asymmetric english auctions.” *Journal of Economic Theory*, 112, 261–288.
- Maskin, Eric (1992), “Auctions and privatization.” In *Privatization* (Horst Siebert, ed.), p. 115–136, J.C.B. Mohr.
- Morris, Stephen (1995), “The common prior assumption in economic theory.” *Economics and philosophy*, 11, 227–227.
- Myerson, Roger B (1981), “Optimal auction design.” *Mathematics of operations research*, 6, 58–73.
- Neeman, Zvika (2003), “The effectiveness of english auctions.” *Games and Economic Behavior*, 43, 214–238.
- Neeman, Zvika (2004), “The relevance of private information in mechanism design.” *Journal of Economic Theory*, 117, 55–77.
- Riley, John and Richard Zeckhauser (1983), “Optimal selling strategies: When to haggle, when to hold firm.” *The Quarterly Journal of Economics*, 98, 267–289.

Samet, Dov (1998), “Iterated expectations and common priors.” *Games and economic Behavior*, 24, 131–141.

Segal, Ilya (2003), “Optimal pricing mechanisms with unknown demand.” *The American economic review*, 93, 509–529.

Wilson, Robert (1987), “Game-theoretic analyses of trading processes.” In *Advances in Economic Theory: Fifth World Congress* (Truman Bewley, ed.), chapter 2, 33–70, Cambridge University Press.

A Proofs

Proof of Lemma 1. Fix a type space \mathcal{T} . I will show that there exists a payoff type space $\widehat{\mathcal{T}}$ such that $S(\mathcal{T}) = S(\widehat{\mathcal{T}})$ and $R(\mathcal{T}) \geq R(\widehat{\mathcal{T}})$. To that end, define $R_j(t_j)$ to be the revenue generated when bidder j is consulted. Since bidders report truthfully, this is

$$R_j(t_j) = \frac{1}{n-1} \sum_{v_{-j}^{(2)}} \max_r r \sum_{v_{-j}^{(1)} \geq r, k} \frac{k}{\mathbb{I}_{\phi_j(t_j)=v_{-j}^{(1)}} + k} \pi \left(v_{-j}^{(1,2)}, k \mid t_j \right),$$

where

$$\pi \left(v_{-j}^{(1,2)}, k \mid t_j \right) = \sum_{\left\{ t_{-j} \in T_{-j} \mid \begin{array}{l} \phi_{-j}^{(1,2)}(t_{-j}) = v_{-j}^{(1,2)}, \\ |W(v_{-j})| = k \end{array} \right\}} \pi(t_{-j} \mid t_j).$$

Consider the payoff type space in which $\widehat{T} = \phi(T)$, $\widehat{\phi}_i(t_i) = t_i$, and $\widehat{\pi}(t) = \sum_{t' \in \phi^{-1}(t)} \pi(t')$. The type space $\widehat{\mathcal{T}} = (\widehat{T}, \widehat{\pi}, \widehat{\phi})$. Since the marginal distribution over values is the same between the two type spaces, clearly the same efficient surplus obtains. Also, the revenue generated by consulting bidder j when $\phi_j(t_j) = v_j$ is

$$\begin{aligned} \sum_{t_j \in \phi_j^{-1}(v_j)} \pi_j(t_j) R_j(t_j) &= \sum_{t_j \in \phi_j^{-1}(v_j)} \frac{\pi_j(t_j)}{n-1} \sum_{v_{-j}^{(2)}} \max_r r \sum_{v_{-j}^{(1)}, k} \frac{k}{\mathbb{I}_{\phi_j(t_j)=v_{-j}^{(1)}} + k} \pi \left(v_{-j}^{(1,2)}, k \mid t_j \right) \\ &\geq \frac{1}{n-1} \sum_{v_{-j}^{(2)}} \max_r r \sum_{v_{-j}^{(1)}, k} \frac{k}{\mathbb{I}_{\phi_j(t_j)=v_{-j}^{(1)}} + k} \sum_{t_j \in \phi_j^{-1}(v_j)} \pi_j(t_j) \pi \left(v_{-j}^{(1,2)}, k \mid t_j \right) \\ &= \frac{1}{n-1} \sum_{v_{-j}^{(2)}} \max_r r \sum_{v_{-j}^{(1)}, k} \frac{k}{\mathbb{I}_{\phi_j(t_j)=v_{-j}^{(1)}} + k} \widehat{\pi}(v_{-j} \mid v_j) \\ &= \widehat{R}_j(v_j). \end{aligned}$$

Thus, it is without loss of generality to consider payoff type spaces.

Symmetry follows. For a payoff type space \mathcal{T} can always be made symmetric by uniformly randomizing over the $n!$ permutations $\xi : N \rightarrow N$ of the players' identities. Let the set of such

permutations be denoted Ξ . Players types are $\widehat{T}_i = \Xi \times \cup_{i \in N} T_i$, $\widehat{\phi}_i(\xi_i, v_i) = v_i$, and

$$\widehat{\pi}(\xi, v) = \begin{cases} 0 & \text{if } \xi_i \neq \xi_j \text{ for some } i \text{ and } j \\ \frac{1}{n!} \pi(v^\xi) & \text{otherwise} \end{cases},$$

where $v_{\xi(i)}^\xi = v_i$. Clearly, revenue and the efficient surplus are the same under the type space $\widehat{\mathcal{T}}$, but if the new types ξ_i are integrated out, S stays the same and R weakly decreases, and we are left with the symmetric and payoff type space $\widetilde{\mathcal{T}}$ which has the distribution over values

$$\widetilde{\pi}(v) = \frac{1}{n!} \sum_{\xi \in \Xi} \pi(v^\xi),$$

types $\widetilde{T} = \cup_{i \in N} T_i$ and $\widetilde{\phi}_i(v_i) = v_i$.

Finally, a similar argument shows that revenue can always be lowered by giving every bidder except the winner have a value of \underline{v} . Starting with a symmetric and payoff type space \mathcal{T} , define $\widehat{\mathcal{T}}$ that has the same support for values T_i , but distribution

$$\widehat{\pi}(v) = \begin{cases} f(x) & \text{if } V_i = x, v_j = \underline{v} \forall j \neq i \\ \pi(v) & \text{if } v_i = \underline{v} \forall i \\ 0 & \text{otherwise} \end{cases},$$

where

$$f(x) = \sum_{\{\tilde{v} | v_i = \tilde{v}^{(1)} = x\}} \frac{\pi(v)}{|W(v)|},$$

which is independent of i by symmetry.

I verify that \mathcal{T} and $\widehat{\mathcal{T}}$ have the same efficient surplus by checking that the probability that

$v^{(1)} = x$ is the same for both type spaces:

$$\begin{aligned}
\sum_{\{v|v^{(1)}=x\}} \pi(v) &= \sum_{\{v|v^{(1)}=x\}} \sum_i \mathbb{I}_{v_i=x} \frac{\pi(v)}{|W(v)|} \\
&= \sum_i \sum_{\{v|v_i=v^{(1)}=x\}} \frac{\pi(v)}{|W(v)|} \\
&= \sum_i \sum_{\{v|v_i=x, v_j=\underline{v} \ \forall j \neq i\}} \hat{\pi}(v) \\
&= \sum_{\{v|v^{(1)}=x\}} \hat{\pi}(v).
\end{aligned}$$

Revenue revenue is lower, as

$$\begin{aligned}
\sum_{v_j} \pi_j(v_j) R_j(v_j) &= \sum_{v_j} \sum_y \max_{r \geq 0} r \sum_{x \geq \max\{v_j, r\}} \sum_{\{\tilde{v} | \tilde{v}_j = v_j, \tilde{v}_{-j}^{(1,2)} = (x, y)\}} \frac{|W(\tilde{v})| - \mathbb{I}_{\tilde{v}_j=x}}{|W(\tilde{v})|} \pi(\tilde{v}) \\
&\geq \max_{r \geq 0} r \sum_{v_j} \sum_y \sum_{x \geq \max\{v_j, r\}} \sum_{\{\tilde{v} | \tilde{v}_j = v_j, \tilde{v}_{-j}^{(1,2)} = (x, y)\}} \frac{|W(\tilde{v})| - \mathbb{I}_{v_j=x}}{|W(\tilde{v})|} \pi(\tilde{v}) \\
&= \max_{r \geq 0} r \sum_{x \geq \max r} \sum_{\{\tilde{v} | \tilde{v}^{(1)}=x\}} \frac{|W(\tilde{v})| - \mathbb{I}_{\tilde{v}_j=x}}{|W(\tilde{v})|} \pi(\tilde{v}) \\
&= \max_{r \geq 0} r \sum_{x \geq r} \frac{|W(\tilde{v})| - \mathbb{I}_{\tilde{v}_j=x}}{|W(\tilde{v})|} f(x) \\
&= \max_{r \geq 0} r \sum_{\{\tilde{v} | \tilde{v}_j = \underline{v}, v^{(1)} \geq r\}} \hat{\pi}(\tilde{v}) = \hat{R}_j.
\end{aligned}$$

□

Proof of Proposition 1. It is clear that symmetric payoff type spaces in which at most one bidder has a positive value are defined by the distribution of the highest value $F^{(1)}$. Moreover, for a given R , it must be that (9) is a lower bound on the distribution. Thus, it must be that

$$\begin{aligned}
S(\mathcal{T}) &= \int_{v=0}^{\gamma S(\mathcal{T})} v dF^{(1)}(v) \\
&\geq R(1 + \log(\gamma) + \log(S(\mathcal{T})) - \log(R)),
\end{aligned}$$

and therefore

$$1 \geq E(\mathcal{M}^{BSA}, \mathcal{T})(1 + \log(\gamma) - \log(E(\mathcal{M}^{BSA}, \mathcal{T})))$$

for any \mathcal{T} . Now let us consider the quantity

$$h(x) = 1 - x(1 + \log(\gamma) - \log(x)).$$

It is straightforward to derive

$$h'(x) = -(1 + \log(\gamma) - \log(x)) + 1 = \log(x) - \log(\gamma) < 0,$$

since $x \in [0, 1]$ and $\gamma > 1$, and h is strictly decreasing. Also,

$$\begin{aligned} \lim_{x \rightarrow 0} h(x) &= 1 - \lim_{x \rightarrow 0} \frac{1 + \log(\gamma) - \log(x)}{x^{-1}} \\ &= 1 - \lim_{x \rightarrow 0} \frac{-x^{-1}}{-x^{-2}} = 1 \end{aligned}$$

via L'Hôpital's rule and $h(1) = 1 - (1 + \log(\gamma)) = -\log(\gamma) < 0$. Thus, there exists a unique point x^* at which $h(x^*) = 0$, and $h(x) > 0$ iff $x < x^*$. This implies that $E(\mathcal{M}^{BSA}, \mathcal{T}) \geq E^*(\gamma)$ for every $\gamma > 1$.

All that remains to be seen is that there is a sequence of type spaces \mathcal{T}^k such that $E(\mathcal{M}^{BSA}, \mathcal{T}^k) \rightarrow E^*(\gamma)$. We are careful to make sure each \mathcal{T}^k has support in $[0, \gamma S(\mathcal{T}^k)]$. Take S^k any sequence converging to $\frac{R}{E^*(\gamma)}$, and $V^k = \{v_0, \dots, v_{m_k}\}$ is the support of \mathcal{T}^k with $v_l = \frac{l}{m_k} \gamma S^k$. We use the CDF of the highest value $F_k^{(1)}$ defined by $F_k^{(1)}(v) = F^{(1)}(v_{l+1})$ for all $v \in [v_l, v_{l+1})$ with $F^{(1)}$ as in (9). For each k , as $m_k \rightarrow \infty$,

$$\int_{v=0}^{\gamma S^k} v dF_k^{(1)}(v) \xrightarrow{m_k \rightarrow \infty} R(1 + \log(\gamma) + \log(S^k) - \log(R)) > S^k,$$

since $h\left(\frac{R}{S^k}\right) < 0$. Thus, m_k can be taken large enough so that $S(\mathcal{T}^k) > S^k$, and therefore satisfies Assumption 1. Moreover, $S(\mathcal{T}^k) \leq \frac{R}{E^*(\gamma)}$ by the argument of the previous paragraph, so $S(\mathcal{T}^k) \rightarrow \frac{R}{E^*(\gamma)}$ by the squeeze theorem. \square

Proof of Theorem 1. I show that for the sequence \mathcal{T}^k constructed in the proof of Proposition 1, the supremum of $E(\mathcal{M}, \mathcal{T}^k, \sigma)$ over all \mathcal{M} and σ converges to $E^*(\gamma)$. Since the supremum over σ is weakly greater than the infimum over σ , this will prove the lemma. The rest of the proof is standard, and follows Myerson (1981) or Börgers (2013).

Since the designer is allowed to pick σ , it is sufficient to look at direct revelation mechanisms where $M = T$. We can divide profiles of valuations into those in which bidder i has a positive value, and all other bidders have a zero valuation. Let $q_i(v)$ and $p_i(v)$ be bidder i 's allocation and transfer if i has a positive value v , and other bidders have zero values. Let $\underline{q}_j^i(v)$ and $\underline{p}_j^i(v)$ be bidder j 's allocation and transfer when j has a zero value, and bidder i has a positive value v . If $v > v'$, then it must be

$$\begin{aligned} v q_i(v) - p_i(v) &\geq v q_i(v') - p_i(v') \\ v' q_i(v') - p_i(v') &\geq v' q_i(v) - p_i(v) \\ \implies v(q_i(v) - q_i(v')) &\geq p_i(v) - p_i(v') \geq v'(q_i(v) - q_i(v')), \end{aligned}$$

so $q_i(v) \geq q_i(v')$, and the allocation must be weakly increasing. Moreover, it must be that

$$\begin{aligned} u_i(v) - u_i(v') &= v q_i(v) - p_i(v) - v' q_i(v') + p_i(v') \\ &\geq v q_i(v') - p_i(v') - v' q_i(v') + p_i(v') \\ &= (v - v') q_i(v'), \end{aligned}$$

and thus, if the support of values is indexed by $\{v^0 = 0, \dots, v^L = \bar{v}\}$, then

$$u_i(v^l) \geq u_i(0) + \sum_{m=0}^{l-1} (v^{m+1} - v^m) q_i(v^m),$$

and thus

$$\begin{aligned} p_i(v^l) &= v^l q_i(v^l) - u_i(v^l) \\ &\leq v^l q_i(v^l) - u_i(0) - \sum_{m=0}^{l-1} (v^{m+1} - v^m) q_i(v^m). \end{aligned}$$

Finally, individual rationality tells us that $u_i(0) \geq 0$.

Also, since the utility of the low type is always $0q_j^i(v) - p_i^j(v) = -p_i^j(v)$, which must be non-negative to satisfy individual rationality, it follows that $p_i^j(v) \leq 0$. It is never beneficial for the seller to allocate the good to a bidder with valuation zero, since they will never pay a positive amount. The seller might as well leave the good unallocated.

Hence, an upper bound on the seller's revenue is

$$\sum_{i \in N} \sum_{l=0}^L p_i(v^l) \pi_i(v^l),$$

which is a linear function of the $q_i(v^l)$, as shown above. The set of weakly increasing $q_i(v^l)$ is a convex set and its extreme points are those functions for which $q_i(v^l) \in \{0, 1\}$. These are precisely the allocations that are implemented by posted price rules, where there are bidder specific reservation prices r_i^* . But since each bidder's valuation has distribution proportional to $F_k^{(1)}$ when positive, an optimal reserve price is by construction $r_i^* = R$. At this price, revenue is exactly R , since the bidder with the high value always buys. This proves the result. \square

Proof of Proposition 2. The seller would achieve the same revenue as in a second-price auction with anonymous reserve r^* if the price set for winning bidders is $\max\{r^*, b^{(2)}\}$. Observe,

$$\begin{aligned} R_j(t_j) &= \frac{1}{n-1} \sum_{v_{-j}^{(2)}} \max_r r \sum_{v_{-j}^{(1)} \geq r, k} \frac{k}{\mathbb{I}_{\phi_j(t_j)=v_{-j}^{(1)}} + k} \pi(v_{-j}^{(1,2)}, k | t_j) \\ &\geq \frac{1}{n-1} \sum_{v_{-j}^{(2)}} \max\{r^*, v_j, v_{-j}^{(2)}\} \sum_{v_{-j}^{(1)} \geq \max\{r^*, v_j, v_{-j}^{(2)}\}, k} \frac{k}{\mathbb{I}_{\phi_j(t_j)=v_{-j}^{(1)}} + k} \pi(v_{-j}^{(1,2)}, k | t_j) \\ &= \frac{1}{n-1} \sum_{\{t_{-j} | \phi_{-j}^{(1)}(t_{-j}) \geq \max\{\phi_j(t_j), r^*\}\}} \max\{r^*, \phi^{(2)}(t)\} \frac{|W(\phi(t))| - \mathbb{I}_{\phi_j(t_j)=\phi^{(1)}(t)}}{|W(\phi(t))|} \pi(t_{-j} | t_j). \end{aligned}$$

Hence, the total revenue from the auction satisfies

$$\begin{aligned}
& \sum_j \sum_{t_j \in T_j} \pi_j(t_j) R_j(t_j) \\
& \geq \sum_{j, t_j \in T_j} \frac{\pi_j(t_j)}{n-1} \sum_{\{t_{-j} | \phi_{-j}^{(1)}(t_{-j}) \geq \max\{r^*, \phi_j(t_j)\}\}} \max\{r^*, \phi^{(2)}(t)\} \frac{|W(\phi(t))| - \mathbb{I}_{\phi_j(t_j)=\phi^{(1)}(t)} \pi(t_{-j}|t_j)}{|W(\phi(t))|} \pi(t_{-j}|t_j) \\
& = \frac{1}{n-1} \sum_j \sum_{\{t \in T | \phi^{(1)}(t) \geq r^*\}} \max\{r^*, \phi^{(2)}(t)\} \frac{|W(\phi(t))| - \mathbb{I}_{\phi_j(t_j)=\phi^{(1)}(t)} \pi(t)}{|W(\phi(t))|} \pi(t) \\
& = \frac{1}{n-1} \sum_{\{t \in T | \phi^{(1)}(t) \geq r^*\}} \max\{r^*, \phi^{(2)}(t)\} \pi(t) \sum_j \frac{|W(\phi(t))| - \mathbb{I}_{\phi_j(t_j)=\phi^{(1)}(t)} \pi(t)}{|W(\phi(t))|} \\
& = \sum_{\{t \in T | \phi^{(1)}(t) \geq r^*\}} \max\{r^*, \phi^{(2)}(t)\} \pi(t),
\end{aligned}$$

since

$$\frac{|W(\phi(t))| - \mathbb{I}_{\phi_j(t_j)=\phi^{(1)}(t)} \pi(t_{-j}|t_j)}{|W(\phi(t))|}$$

is zero if $\phi_j(t_j) > \phi_{-j}^{(1)}(t_{-j})$ and also

$$\sum_j \frac{|W(\phi(t))| - \mathbb{I}_{\phi_j(t_j)=\phi^{(1)}(t)} \pi(t)}{|W(\phi(t))|} = n - 1.$$

Since $\sum_{\{t \in T | \phi^{(1)}(t) \geq r^*\}} \max\{r^*, \phi^{(2)}(t)\} \pi(t)$ is revenue under a second-price auction with the anonymous reserve r^* , this proves the result. \square

Table 1: Notation.

Symbol	Page	Meaning
$\Delta(X)$	10	Set of distributions on a measurable space X with finite support.
$x^{(1)}, x^{(2)}, x^{(1,2)}$	10	The highest, second-highest, and vector of highest and second-highest values of the vector x .
$W(x)$	19	The set of indices i such that $x_i = x^{(1)}$.
T_i, T	11	Set of types.
$\phi_i : T_i \rightarrow \mathbb{R}$	11	Mapping from type to private value.
$\pi \in \Delta(T)$	11	Prior distribution over types.
$\mathcal{T} = (T, \pi, \phi)$	11	Type space.
\underline{v}, \bar{v}	11	Smallest and largest values in the support.
M_i, M	11	Set of messages
$q : M \rightarrow \mathbb{R}^n, p : M \rightarrow \mathbb{R}^n$	11	Allocation and transfer rules (net to seller).
$\mathcal{M} = (M, q, p)$	11	Mechanism.
$\text{BNE}(\mathcal{M}, \mathcal{T})$	12	Bayes Nash equilibria with mechanism \mathcal{M} and type space \mathcal{T} .
$S(\mathcal{T})$	11	Efficient surplus.
$R(\mathcal{M}, \mathcal{T}, \sigma)$	12	Revenue.
$E(\mathcal{M}, \mathcal{T}, \sigma)$	12	Extraction ratio.
\mathcal{M}^{BSA}	19	The belief survey auction (BSA).
$r_j \left(m_j, b_{-j}^{(2)} \right)$	21	Reserve price when bidder j is consulted.
$G_j \left(\cdot; m_j, b_{-j}^{(2)} \right)$	21	Upper cumulative distribution of winner's value when bidder j is consulted.
γ	23	Bound on the ratio between values and the efficient surplus.
$E^*(\gamma)$	25	max min extraction ratio for the parameter γ .
\mathcal{M}^ϵ	29	The perturbed BSA.