

Informationally Robust Optimal Auction Design*

Dirk Bergemann Benjamin Brooks Stephen Morris

April 12, 2017

Abstract

A single unit of a good is to be sold by auction to one of two buyers. The good has either a high value or a low value, with known prior probabilities. The designer of the auction knows the prior over values but is uncertain about the correct model of the buyers' beliefs. The designer evaluates a given auction design by the lowest expected revenue that would be generated across all models of buyers' information that are consistent with the common prior and across all Bayesian equilibria. An optimal auction for such a seller is constructed, as is a worst-case model of buyers' information. The theory generates upper bounds on the seller's optimal payoff for general many-player and common-value models.

KEYWORDS: Optimal auctions, common values, information structure, model uncertainty, ambiguity aversion, robustness, Bayes correlated equilibrium, revenue maximization, revenue equivalence, information rent.

JEL CLASSIFICATION: C72, D44, D82, D83.

*Bergemann: Department of Economics, Yale University, dirk.bergemann@yale.edu; Brooks: Department of Economics, University of Chicago, babrooks@uchicago.edu; Morris: Department of Economics, Princeton University, smorris@princeton.edu. This work has been presented at the University of Pennsylvania. We acknowledge financial support through NSF Grant ICES 1215808 and SES 1459899. We are grateful for discussions with Ilya Segal with whom we discussed related ideas for possible joint work in the past. We thank Elliot Lipnowski, Doron Ravid, and Phil Reny for helpful comments.

1 Introduction

In the standard model of optimal auction design, the seller and the potential buyers are assumed to have subjective expected utility preferences over the outcome of the auction, and their beliefs are assumed to be consistent with a common prior over the value of the goods being sold. In spite of having such a rich model of beliefs, the theory sometimes predicts that relatively simple auctions will be optimal. Such is the case in the celebrated work of Myerson (1981), in which a combination of independent signals, symmetry, and linearity of interim expected values leads to the conclusion that first- or second-price auctions with reserve prices will maximize expected revenue. But more generally, and especially when the bidders' information is correlated, the theory can lead to paradoxical and dizzyingly complex forms for the optimal auction. These mechanisms, which have highly desirable theoretical properties, are nonetheless impractical for reasons outside the model. Such is the case with the full-surplus extraction results of Cremer and McLean (1985, 1988).¹ One can give many reasons why such auctions would not be practically implementable; limited liability and risk aversion on the part of the buyers come to mind. But in our view, a critical failure is that in many settings, the designer of the auction faces uncertainty about the correct model of buyers' beliefs. He therefore might eschew an auction that is exactly optimal under a particular model of beliefs in favor of an auction that hedges performance across many different models.

In this paper, we fully embrace this aspect of the auction design problem. We consider a seller who has a single unit of a good for sale to one of two potential buyers. The good is known to have one of two values, high or low, and all of the agents (seller and buyers) agree on a common prior over that value. Moreover, the buyers have beliefs about the value of the good that are consistent with a common prior.² We assume, however, that the seller does not know the correct model of beliefs, and indeed, he considers all models to be possible that are consistent with the given prior over the value.³ For example, it might be that the buyers have no information about the value beyond the prior, and they both value the good at its

¹See also McAfee, McMillan, and Reny (1989) and McAfee and Reny (1992).

²The reader may rightly ask why we allow for model uncertainty on the part of the seller but not on the part of the buyers. We are entirely sympathetic to this view, and the possibility that the buyers face model uncertainty is eminently worthy of study. There are, however, cases in which the buyers may have been participating in a given market for a long time and therefore have a much better sense of the structure of information than does the seller, or they simply have insider knowledge of the process that generates information.

³The classical model where the designer is certain of the correct model of beliefs is extreme. A model in which the seller is completely agnostic about beliefs, as in our model, is also extreme, although it is extreme in a conservative direction that should lead to auctions that are more robust than necessary. Verily, a more reasonable model would fall somewhere in between. We consider the exploration of this middle ground to be a promising direction for future research.

ex-ante expected value. Or it may be that the buyers have noisy observations about the value, perhaps seeing a signal which is the value plus conditionally independent noise. The buyers' information can even be multidimensional, and come from signals that are correlated with one another and with the value in an arbitrary manner.

The seller seeks an auction format that will perform well regardless of which model of beliefs turns out to be the correct description of the world. In particular, for any candidate auction mechanism, the seller evaluates its performance by the lowest expected revenue across all models of information and across all Bayesian equilibria.⁴ The space of auctions that we allow the seller to choose from is vast: the seller can essentially choose any measurable sets of messages for the bidders to send to the mechanism, and arbitrary measurable mappings from messages to allocations and transfers. The only requirement is that the mechanism allow the buyers to "opt-out," in which case they receive the good with zero probability and make zero transfers. Our main result is to identify the best revenue guarantee that the seller can obtain across all auction formats, as well as to identify a particular auction that achieves this best lower bound. (Of course, this maxmin auction may perform even better than the guarantee if the model of beliefs is favorable.)

Let us describe our results, starting with the best revenue guarantee that the seller can obtain. This quantity can be simply understood through a particular worst-case model of information that the seller must guard against. The motivating idea for the worst-case model is the principle that many auctions should perform reasonably well, but none should stand out as better than the others. Put differently, if one auction stood out as superior to the others, then we might think it would be possible to construct a different model of information under which this auction is still the best, but has less of an advantage vis-a-vis the other auction formats than it did under the original type space.

Pursuing this idea, we construct a model of information in which there is an extreme amount of indifference across auction mechanisms. In the worst-case model, the buyers observe one-dimensional signals that are independently distributed, and as a normalization we can take them to be standard uniform random variables. Moreover, the interim expectation of the value given both signals is continuous and monotonically increasing in the signals. In such a setting, there is a generalization of the well-known revenue equivalence result of Myerson (1981), which is due to Bulow and Klemperer (1996). The revenue equivalence

⁴Thus, the seller exhibits an extreme form of ambiguity aversion with respect to the uncertainty about the correct model of beliefs, in the sense of Gilboa and Schmeidler (1989). We are implicitly taking the view that the seller has distinct attitudes towards the uncertainty about the correct model of beliefs, and the uncertainty about the realized beliefs conditional on a particular model. A complete justification of this view is beyond the scope of this paper. See (cf. Hansen and Sargent, 2001) for a more in-depth discussion of this perspective, in the context of monetary policy.

formula is based on the principles that (i) any Bayesian incentive compatible mechanism can be implemented as a direct mechanism in which bidders report their signals, and (ii) such a mechanism must deter “local” deviations, wherein a buyer changes their reported type by a marginal amount. Under these conditions, equilibrium transfers that buyers make to the seller, and hence expected revenue, are determined by the equilibrium allocation. In addition, expected revenue can be reformulated simply as the expectation of the “virtual value” of the buyer who is allocated the good, where “virtual value” is that buyer’s value for the good minus an information rent that comes from incentive constraints.

The value function in this “minmax” information structure has the property that *all bidders have the same virtual value*. As a result, the seller is always indifferent as to whether the mechanism allocates the good to one buyer or the other, conditional on it being allocated. Moreover, the signal space is divided into a low region and a high region. On the high region, the virtual value is constant and equal to the highest value, and on the low region, *the virtual value is identically zero*. Thus, on the low region, the seller is additionally indifferent between allocating the good and not allocating the good. Indeed, these type spaces that we construct are the only ones that satisfy all of these conditions, and they are indexed by two parameters, which are the interim expected value given the highest possible signals and the interim expected value given the lowest possible signals. The correct parameters are such that (i) the highest interim expectation is the highest ex-post value of the good, and (ii) the law of iterated expectations holds, i.e., the expectation of the interim-expected value is equal to the ex-ante expected value.

Given the conjecture for the worst-case model of beliefs, we then construct a mechanism that guarantees the seller this amount of revenue in all information structures.⁵ To do so, we built on the insights of a recent paper by Du (2016), which also studies a maxmin auction design problem with many bidders and general common value distributions (thus being more general than the model studied here). Two of us have previously shown (Bergemann and Morris, 2013, 2016) that given a fixed mechanism, the problem of calculating minimum revenue across all models of information and all Bayesian equilibria can be reformulated as a linear program, in which expected revenue is minimized across a class of incomplete-information correlated equilibria termed *Bayes correlated equilibria* (BCE). In this linear program, the optimization is over joint distributions of values and messages, and the constraints represent the incentive compatibility requirement that players would want to submit the message that they have “drawn” from the distribution, given the Bayesian inference that they can draw

⁵While one might have hoped that identifying the worst-case model of information would lead us to an optimal mechanism, this is not the case, since there are so many auctions that are optimal in that environment (such as first-price, second-price, or even posted-price mechanisms).

about others messages and the value of the good. The three of us subsequently applied this methodology in characterizing minimum revenue of the first-price auction (?).

Du (2016) innovated by using the dual of this linear program to provide lower bounds on revenue for particular mechanisms. In the dual program, the optimization is over multipliers on incentive constraints and on feasibility constraints for the correlated equilibrium distribution, and for a given mechanism, any choice of multipliers that are feasible for the dual problem generates a lower bound on minimum revenue. Du uses this principle to show that when the number of bidders is large, it is possible for the seller to extract virtually all of the efficient surplus. Specifically, he constructs a sequence of mechanisms and associated multipliers that are feasible for the dual program. While these mechanisms are not exactly optimal for any finite number of buyers, their associated revenue lower bounds converge to the ex-ante expected value of the good as the number of bidders goes to infinity.

We employ a similar technique to characterize the maxmin auction with two bidders and binary values. Motivated by our worst-case information structure, we can learn a great deal about the structure of a maxmin mechanism. The conjectured formula for minmax revenue predicts the correct multipliers on the total probability constraints. We know that any incentive compatible mechanism would effectively induce a direct mechanism on the worst-case type space, which has one-dimensional types, so we can guess that it is also sufficient to look at mechanisms with one-dimensional message spaces. Since local incentive constraints pin down the optimal revenue in the worst-case model of beliefs, we can guess that only local incentive constraints have non-zero multipliers. Following Du, we normalized the multipliers on local incentive constraints to be a constant. (This essentially corresponds to a particular choice of units for the messages.) The value of this multiplier is then essentially pinned down by the choice of the allocation rule. We then construct an allocation and transfer rule for which these multipliers are feasible for the dual problem. Thus, this mechanism must achieve the upper bound on maxmin revenue corresponding to our worst-case type space.

This mechanism turns out to have a relatively simple structure, and is distinct from that constructed by Du (2016). We can think of there being a unit quantity of the good which is to be divided between the buyers (with buyers' preferences being linear in the quantity of the good they receive). The bidders make one-dimensional "demands" for an amount of the good. The buyers' demands are then filled sequentially in a random order, so that if the sum of the demands exceeds the total supply of the good, then a buyer will only receive their full demand if they are served first. Otherwise, the buyer receives whatever is left over after the other buyer's demand has been filled. The maxmin transfer rule is somewhat more complicated. In a sense, the transfer is exponential in the buyer's own demand, and the

sum of the growth rates of the two buyers' transfers is constant, but how that growth rate is distributed across the buyers depends on their demands.

After our main theorem, we compare the maxmin mechanism to other well-known auction formats, such as the first- and second-price auctions, posted prices, and Du's mechanism. We also describe upper bounds on maxmin revenue for more general models, which are based on generalizations of the worst-case model of beliefs.

This paper relates to the literature on informationally robust mechanism design as well as informationally robust predictions in Bayesian games. The closest related papers are ? and Du (2016), described above. More recently, Bergemann, Brooks, and Morris (2016) studies optimal auction design on the information structure that is the worst-case for the first-price auction. That paper shows that any auction that induces a conditionally efficient allocation when values are independent, symmetric, and private, will be revenue equivalent to the first-price auction on the latter's worst-case. This shows that first-price auctions have greater minimum revenue than any other such auction, including second-price and all-pay auctions with or without reserve prices. Outside of this class, there are other auctions that improve on the first-price auction at its own worst-case, in particular posted-price mechanisms, thus showing that the first-price auction is *not* a maxmin mechanism. The present paper is the next logical step in this line of research, using the revenue equivalence principle to derive the minmax information structure space, and from there a maxmin mechanism.

In addition, there are a number of papers that treat similar questions with a single buyer. Bergemann and Schlag (2011, 2008) characterizes randomized posted price mechanisms that minimizes maximum regret in the face of uncertainty about the demand curve faced by a monopolist. Our maxmin mechanism reduces to theirs in the case of a single buyer. Condorelli and Szentes (2016) and Roesler and Szentes (2016) characterize optimal information structures for maximizing buyer surplus when the seller can make it a take-it-or-leave-it offer for the good. Under their interpretation, the buyer can choose the information structure before interacting with the monopolist, where the monopolist has all of the bargaining power when the two finally meet. Our worst-case type spaces reduce to theirs with a single buyer. Our model can also be interpreted in a similar manner: suppose the buyers can collude on designing their information, after which the seller will run a revenue maximizing auction. Then the optimal value for the buyers is given by the efficient surplus, less maxmin revenue.⁶

⁶Note that the good would not always be allocated under the maxmin mechanism that we construct, so that total surplus would be less than the efficient surplus. However, if the buyers chose the minmax information structure, then there are other seller-optimal mechanisms that the seller could choose, e.g., posted prices, under which the good is sold with probability one. Thus, an equilibrium of the sequential game is for the buyers to first choose the minmax information structure, followed by the seller choosing an optimal posted price.

Kos and Messner (2015) and Carrasco et al. (2015) construct a maxmin mechanism when there is a single buyer, which has the form of a randomized posted price. When there is only a mean constraint on the value, their models can be viewed as a special case of ours.

Chung and Ely (2007) pursue a conceptually similar exercise as ours. They also study surplus extraction where the auction designer is uncertain about the correct model of buyers’ beliefs. In contrast to our model, values are private and beliefs are not required to be consistent with a common prior. Moreover, they select for the seller’s most preferred equilibrium, if there is more than one. Their main conclusion is that there exists a maxmin mechanism in weakly dominant strategies. In contrast, we select for the worst equilibrium for the seller,⁷ and our mechanism does not generally have an equilibrium in weakly dominant strategies. Yamashita (2015) studies maxmin auction design in private value models within restricted classes of mechanisms and in the large market limit. Neeman (2003) and Brooks (2013b,a) also study robust surplus extraction in the private-value many-bidder model, but assume non-standard preferences on the part of the designer that are more in the spirit of the computer science literature on algorithmic mechanism design (e.g., Hartline and Roughgarden, 2009).

The rest of this paper proceeds as follows. In Section 2, we describe our model of optimal auction design with model uncertainty. Section 3 then constructs the worst-case model of information. Section 4 constructs the maxmin mechanism and presents our main theorem that characterizes maxmin revenue. Section 5 compares the maxmin auction with other mechanisms and discusses some extensions of the results. Section 6 concludes. Omitted proofs are contained in the Appendix.

2 Model

A single unit of a good is to be sold by auction to one of two bidders, indexed by $i \in \{1, 2\}$. The bidders assign the same value to the good $v \in V = \{0, 1\}$. Value v has ex-ante probability $p(v)$.

⁷It turns out that maxmin revenue does not depend on how one selects equilibrium, since any equilibrium on the worst-case type space would result in the same expected revenue. However, selecting the equilibrium that is preferred by the seller may expand the set of maxmin optimal mechanisms. For example, the seller could run a mechanism in which he asks the buyers to announce the type space. If there is a majority that announces the same type space, then the seller runs the optimal auction for that type space, whatever it may be, and otherwise runs some fixed mechanism. When there are at least three buyers, truthful reporting of the type space is an equilibrium, since no buyer is pivotal. There are many other equilibria, though, in which the buyers coordinate on untruthful reports. Our mechanism is much more compelling, in that *all* equilibria generate favorable revenue performance. We also give an explicit description of the mechanism, rather than relying on the black box of the “optimal mechanism” in type spaces for which the revenue maximizing auction is unknown.

An *auction mechanism* consists of (i) measurable sets of messages M_i for each bidder i , (ii) measurable allocation functions $q_i : M \rightarrow \mathbb{R}_+$ for $i = 1, 2$, where $M = M_1 \times M_2$, that satisfy the feasibility constraints

$$q_1(m) + q_2(m) \leq 1$$

for all $m \in M$, and (iii) measurable transfer functions $t_i : M \rightarrow \mathbb{R}$ for $i = 1, 2$. The allocation function $q_i(m)$ represents bidder i 's probability of being allocated the good if the message profile m is sent, and $t_i(m)$ is bidder i 's net transfer to the seller. We denote an auction by $\mathcal{A} = (\{M_i\}, \{q_i\}, \{t_i\})$. We require there exists a message $0_i \in M_i$ such that

$$q_i(0_i, m_j) = t_i(0_i, m_j) = 0.$$

Thus, sending a message of 0_i corresponds to opting out of the mechanism.

An *information structure* consists of (i) measurable sets of signals S_i for each bidder i and (ii) a probability transition kernel that maps values in V into probability measures over signal profiles in $S = S_1 \times S_2$:

$$\pi : V \rightarrow \Delta(S).$$

We denote an information structure by $\mathcal{S} = (\{S_i\}, \pi)$.

Given an auction mechanism \mathcal{A} and an information structure \mathcal{S} , we define a strategy of bidder i to be a measurable mappings from bidder i 's signals to distributions over messages. Thus, the set of bidder i 's strategies is

$$\Sigma_i(\mathcal{A}, \mathcal{S}) = \{\sigma_i : S_i \rightarrow \Delta(M_i)\}.$$

Given a profile of strategies

$$\sigma \in \Sigma(\mathcal{A}, \mathcal{S}) = \Sigma_1(\mathcal{A}, \mathcal{S}) \times \Sigma_2(\mathcal{A}, \mathcal{S}),$$

bidder i 's payoff is

$$U_i(\sigma, \mathcal{A}, \mathcal{S}) = \sum_{v \in V} \int_{s \in S} \int_{m \in M} (v q_i(m) - t_i(m)) \sigma(dm|s) \pi(s|v) p(v),$$

and revenue is

$$R(\sigma, \mathcal{A}, \mathcal{S}) = \sum_{v \in V} \int_{s \in S} \int_{m \in M} \sum_{i=1}^2 t_i(m) \sigma(dm|v) \pi(s|v) p(v).$$

We will suppress the dependence of these objects on $(\mathcal{A}, \mathcal{S})$ when it is clear from context. A *Bayes Nash equilibrium*, or *equilibrium* for short, is a profile of strategies $\sigma \in \Sigma(\mathcal{A}, \mathcal{S})$ such that

$$U_i(\sigma) \geq U_i(\sigma'_i, \sigma_j)$$

for all $\sigma'_i \in \Sigma_i(\mathcal{A}, \mathcal{S})$. We let $\Sigma^*(\mathcal{A}, \mathcal{S})$ denote the set of all equilibria of the Bayesian game $(\mathcal{A}, \mathcal{S})$.

The seller is ambiguity averse with respect to the information structure and the equilibrium. Thus, the seller evaluates a mechanism \mathcal{A} according to

$$R(\mathcal{A}) = \inf_{\mathcal{S}} \inf_{\sigma \in \Sigma^*(\mathcal{A}, \mathcal{S})} R(\sigma, \mathcal{A}, \mathcal{S}).$$

The seller's problem is to identify an auction mechanism that achieves

$$\sup_{\mathcal{A}} R(\mathcal{A}).$$

We let R^* denote the solution to the seller's problem. We will refer to R^* as *maxmin revenue*, and any mechanism that achieves the supremum will be referred to as a *maxmin mechanism*.

The seller's problem is essentially a zero-sum game between the seller, who chooses a mechanism, and nature, who chooses the information structure and the equilibrium strategies adversarially in order to make the expected revenue as low as possible. In the above formulation, the seller first selects the mechanism and nature then selects the type space and equilibrium. Alternatively, one could have reversed the order of moves with nature first choosing \mathcal{S} and then the seller choosing \mathcal{A} . In a zero-sum game with finite actions, the minimax theorem would imply that the two problems are payoff-equivalent, and that there exists an equilibrium saddle point $(\mathcal{A}^*, \mathcal{S}^*)$ such that each player's action is optimal regardless of the order of moves. In this setting, it is far from obvious that a minimax theorem should hold, due to the fact that the equilibrium correspondence is not lower-hemicontinuous with respect to the information and game form. However, it remains true that R^* is weakly less than the *minmax revenue* that the seller would obtain if Nature had to choose the information structure first. An optimal such information structure would be a *minmax information structure*. As Theorem 1 below affirms, the minimax property *will* hold in our setting, and we will prove our main result by constructing a saddle point for the seller's problem.

Finally, there are two technical points that need to be made. First, the seller's problem involves a sup over auction mechanisms and an inf over information structures, even though we have not properly defined the available sets of mechanisms and information structures. In order to make the "set of all mechanisms (information structures)" well-defined, we can

fix some ambient and sufficiently rich measurable spaces \overline{M} (\overline{S}) and impose that M_i is a measurable subset of \overline{M} (and ditto for S_i and \overline{S}). We shall see that it will be sufficient if these ambient spaces contain subsets that are isomorphic to the unit interval $[0, 1]$ with the standard Borel σ -algebra. Second, the definition of $R(\mathcal{A})$ involves an infimum over equilibria of the Bayesian game $(\mathcal{A}, \mathcal{S})$, even though we have not provided sufficient conditions for the existence of *any* equilibrium. We could, without altering our results, restrict the seller and nature to choosing only mechanisms and information structures for which an equilibrium exists, e.g., those satisfying the hypotheses of Milgrom and Weber (1985), since as we shall see, the maxmin mechanism and minmax information structure both fall into this class. For our purposes, it is sufficient to require that (i) the seller must pick a mechanism for which an equilibrium exists under *some* information structure, and (ii) given the mechanism, Nature must pick an information structure for which an equilibrium exists.

3 An upper bound on R^*

We will proceed by first establishing an upper bound on R^* , and then constructing a mechanism that achieves the bound. The bound is generated by a particular choice of information structure for which we can easily pin down revenue. On this type space, $S_i = [0, 1]$, and the bidders' signals are independent draws from the standard uniform distribution.⁸ Instead of specifying the conditional distribution of the signals given the value, we shall describe the information structure via the interim expectation of the value given the pair of signals:

$$v(s_1, s_2) = \min \left\{ \frac{a}{(1-s_1)(1-s_2)}, 1 \right\}$$

for some $a \in [0, 1]$. The parameter a is determined as follows. Given our distributional assumptions and the functional form of v , one can compute the expected value as

$$\int_{s_1=0}^1 \int_{s_2=0}^1 v(s_1, s_2) ds_1 ds_2 = a \left(1 - \log(a) + \frac{1}{2} \log^2(a) \right).$$

⁸We emphasize that the fact that the worst-case information structure has independent and one-dimensional signals is a *conclusion*, rather than an assumption. Independence is a fairly intuitive property for the worst-case, since any correlation between signals creates opportunities for the seller to separate types and extract additional surplus using side bets in the style of Cremer and McLean (1985, 1988). The unidimensionality of signals and monotonicity of interim expected values is far less obvious. In our analysis of the first-price auction ?, we also concluded that the worst-case model of information had independent and unidimensional signals and monotonic interim expected values (in the common value case). This corresponds to a strong ordering on belief hierarchies, so that types that have higher expectations of the value also have higher expectations of others expectations, etc. It is curious that for both of these problems, the most effective way to distribute information for the purposes of depressing revenue is to order the types in a consistent manner at all belief levels.

In order for the law of iterated expectations to hold, it must be that the expectation of $v(s_1, s_2)$ is equal to the ex-ante expected value, so that

$$a \left(1 - \log(a) + \frac{1}{2} \log^2(a) \right) = p(1). \quad (1)$$

One can easily verify that the left-hand side is strictly increasing in a and is equal to 0 and 1 when a is 0 or 1, respectively, so that there is a unique solution to this equation by the intermediate value theorem. We denote this information structure by \mathcal{S}^* .

\mathcal{S}^* has the features that (i) the buyers have one-dimensional and independent signals and (ii) the interim expected value is a weakly increasing function of the buyers' signals. In addition, the value is less than one on the low-signal region in which $(1 - s_1)(1 - s_2) < a$, and it is identically equal to one when this inequality does not hold.

It turns out that revenue on this type space can be computed relatively easily, using the generalization of the revenue equivalence formula described in Bulow and Klemperer (1996). We can heuristically derive this formula as follows. A direct mechanism in which the allocation and transfer rule are given by measurable functions $q_i : [0, 1]^2 \rightarrow [0, 1]$ and $t_i : [0, 1] \rightarrow \mathbb{R}$ respectively. The equilibrium utility of a type s_i is

$$U_i(s_i) = \max_{s'_i} \int_{s_j=0}^1 (v(s_i, s_j) q_i(s'_i, s_j) - t_i(s'_i, s_j)) ds_j.$$

The envelope formula implies that

$$U'_i(s_i) = \int_{s_j=0}^1 \frac{\partial v}{\partial s_i}(s_i, s_j) q_i(s_i, s_j) ds_j.$$

The ex-ante surplus of bidder i is therefore

$$\begin{aligned} U_i &= \int_{s_i=0}^1 \int_{x=0}^{s_i} \int_{s_j=0}^1 \frac{\partial v}{\partial s_i}(x, s_j) q_i(x, s_j) ds_j dx ds_i \\ &= \int_{s_i=0}^1 \int_{s_j=0}^1 \frac{\partial v}{\partial s_i}(s_i, s_j) q_i(s_i, s_j) (1 - s_i) ds_j ds_i \end{aligned}$$

where we have applied Fubini's theorem. The seller's revenue on this type space when the allocation is $q_i(s)$ must be the difference between total surplus and the bidders' rents, which is

$$R = \sum_{i=1}^2 \int_{s \in [0,1]^2} \left(v(s) - (1 - s_i) \frac{\partial v}{\partial s_i}(s) \right) q_i(s) ds.$$

As an aside, we note that this formula tells us what revenue must be as a function of the allocation that is implemented in equilibrium, though it does not tell us which allocations can be implemented in a globally incentive compatible mechanism. In the case studied in Myerson (1981), each buyer’s value is linearly increasing in their own signal, and in that case monotonicity of the interim expected allocation is a necessary and sufficient condition for an allocation to be globally incentive compatible. In this non-linear setting, we know of no analogous characterization. However, this turns out not to be central to our analysis.

What the model does have in common with the linear model of Myerson (1981) is that revenue is simply the expectation of the “virtual value” of the buyer who is allocated the good, where the generalized virtual value is

$$\psi_i(s) = v(s) - (1 - s_i) \frac{\partial v}{\partial s_i}(s).$$

If we now plug in the particular choice of v into this expression, we obtain

$$\psi_i(s) = \begin{cases} 0 & \text{if } v(s) < 1; \\ 1 & \text{if } v(s) = 1. \end{cases}$$

This type space therefore has two remarkable features. First, the two bidders always have the same virtual value. Thus, conditional on an allocation being implementable, the seller is indifferent between allocating the good to buyer 1 or to buyer 2. Second, on the low region where the value is less than one, the virtual value is zero, so that the seller is additionally indifferent between allocating the good or not allocating the good. In effect, this type space creates a tremendous amount of indifference on the part of the seller between mechanisms. It is for this reason that it seems to be a good candidate for a worst-case type space from the seller’s perspective.

Indeed, the only requirement for a mechanism to maximize revenue on this type space is that the good is allocated whenever the value is exactly 1. As a result, optimal revenue is simply the probability that the value is one, which is

$$\int_s \mathbb{I}_{v(s)=1} ds = a(1 - \log(a)) = \bar{R}. \tag{2}$$

For any mechanism that the seller can propose, it is always possible that the information structure is S^* , in which case equilibrium revenue can be no more than (2). We therefore have the following proposition:

Proposition 1 (Revenue Upper Bound).

Any mechanism can generate at most \bar{R} in revenue when the information structure is S^* . Thus, R^* must be weakly less than \bar{R} .

Incidentally, there are simple incentive compatible mechanisms that would implement a pointwise optimal allocation, so that \bar{R} is indeed optimal revenue on the type space S^* . In particular, the seller could offer the good at a posted price of $a(1 - \log(a))$, and then randomly allocate the good among the subset of bidders who are willing to purchase it at that price. It is easily verified that all types would want to buy the good at this price (which is the interim expectation of the good conditional on having a signal of $s_i = 0$).

4 A maxmin mechanism

4.1 Reformulating the revenue minimization problem

We shall argue that there exists a mechanism that is guaranteed to generate at least \bar{R} in revenue. To do so, we will employ two simplifications. First, for a fixed auction mechanism $\mathcal{A} = (\{M_i\}, \{q_i\}, \{t_i\})$, Bergemann and Morris (2013, 2016) have shown that the set of outcomes that can arise under some information structure and equilibrium is equivalent to the set of *Bayes correlated equilibria* (BCE). This solution concept generalizes correlated equilibrium to games of incomplete information by allowing the players' actions to be correlated with payoff relevant states of Nature. In our setting, we can define a Bayes correlated equilibrium μ to be an element of $\mathcal{M}(V \times M^2)$, which is the set of Borel (and not necessarily probability) measures on $V \times M^2$, that satisfies the following conditions. First, for every $v \in V$, μ must satisfy the probability constraint

$$\mu(\{v\} \times M^2) = p(v). \quad (3)$$

Second, μ must satisfy the following obedience constraints: for every measurable *deviation mapping* $\eta_i : M_i \rightarrow M_i$, it must be that

$$\begin{aligned} U_i(\mu, \mathcal{A}) &= \int_{(v,m) \in V \times M} (v q_i(m) - t_i(m)) \mu(dv, dm) \\ &\geq \int_{(v,m) \in V \times M} (v q_i(\eta_i(m_i), m_j) - t_i(\eta_i(m_i), m_j)) \mu(dv, dm). \end{aligned} \quad (4)$$

One interpretation of the BCE μ is that there is an omniscient mediator who knows the true value of the good and can make private and correlated recommendations to the bidders

of what message they should send the auction. Equation (4) says that players must prefer to follow the recommendation m_i they receive from the mediator rather than deviate as a function of their recommendation to $\eta_i(m_i)$. We denote by $BCE(\mathcal{A})$ the set of BCE for a given auction \mathcal{A} .

Any equilibrium σ of \mathcal{A} under an information structure \mathcal{S} must induce a joint distribution over the value and messages, which is obtained by integrating out signals. For any measurable set $X \subseteq M$,

$$\mu(v, X) = \int_{s \in \mathcal{S}} \int_{m \in X} \sigma(X|v) \pi(s|v) p(v).$$

It is a fact that this distribution must also be a BCE. Similarly, given a BCE, one can construct an information structure \mathcal{S} and an equilibrium σ that induce the same distribution over values and actions. See Bergemann et al. (2015a) for a formal proof of the epistemic equivalence between the two solution concepts.

The upshot is that we can reformulate the seller's problem by taking the infimum revenue over all BCE, rather than over all information structures and equilibria. In particular,

$$\inf_{\mu \in \mathcal{M}(V \times M)} \int_{(v,m) \in V \times M} \sum_{i=1}^2 t_i(m) \mu(dv, dm) \text{ subject to (3) and (4)} \quad (\text{P})$$

must be equal to $R(\mathcal{A})$. This optimization problem is essentially an infinite dimensional linear program, since both the objective, and the constraints (3) and (4) that characterize BCE are all linear in the measure μ .

We will study this optimization problem for a special class of mechanisms in which $M_i = [0, 1]$ and the functions q_i and t_i are both assumed to be differentiable. Any auction satisfying these assumptions will be referred to as *regular*. Following Du (2016), we will study a dual of a relaxed version of the primal problem (P). At the first step, we will drop all of the constraints except for those corresponding to “local upward deviations.” Specifically, consider the class of deviation mappings of the form

$$\eta_{i,k,X}(m_i) = \mathbb{I}_{m_i \in X} \min \left\{ m_i + \frac{1}{k}, 1 \right\}$$

for some positive integer k and some measurable set $X \subseteq M_i$. For each such deviation, the incentive constraint (4) becomes

$$k \int_{(v,m) \in V \times X} \left[v \left(q_i \left(\min \left\{ m_i + \frac{1}{k}, 1 \right\}, m_j \right) - q_i(m) \right) - \left(t_i \left(\min \left\{ m_i + \frac{1}{k}, 1 \right\}, m_j \right) - t_i(m) \right) \right] \mu(dv, dm) \leq 0.$$

The limit as $k \rightarrow \infty$ is simply

$$\int_{(v,m) \in V \times X} \left[v \frac{\partial q_i}{\partial m_i}(m) - \frac{\partial t_i}{\partial m_i}(m) \right] \mathbb{I}_{m_i < 1} \mu(dv, dm) \leq 0. \quad (5)$$

Note that wherever q_i and t_i are not differentiable, the partial derivative is implicitly defined by right-continuity. Thus, the obedience constraints (4) imply that the local upward obedience constraint (5) must be satisfied for all measurable sets $X \subseteq M_i$. We can therefore relax the primal problem (P) by dropping all other constraints except these local ones, to obtain the following optimization problem:

$$\inf_{\mu \in \mathcal{M}(V \times M)} \int_{(v,m) \in V \times M} \sum_{i=1}^2 t_i(m) \mu(dv, dm) \text{ subject to (3) and (5)}. \quad (P')$$

Let us denote by $\tilde{R}(\mathcal{A})$ the solution to this optimization problem. Since (P') is a relaxation of (P), and since (P) is equivalent to the seller's problem, we can conclude that $\tilde{R}(\mathcal{A}) \leq R(\mathcal{A})$ for any regular mechanism.

Again, following Du (2016), we will study a different optimization problem that is dual to (P'). The choice variables in the dual will be ‘‘multipliers’’ on the constraints in the primal problem. Specifically, let Γ denote the set of real-valued functions of values, with a typical element being

$$\gamma : V \rightarrow \mathbb{R},$$

and let A denote the set of measurable non-negative functions on $[0, 1]$, with a typical element being

$$\alpha : [0, 1] \rightarrow \mathbb{R}_+.$$

These functions represent multipliers on probability constraints (3) and multipliers on the local incentive constraints (5), respectively. Also let

$$\Phi(\gamma, \{\alpha_i\}) = \inf_{\mu \in \mathcal{M}(V \times M)} \left\{ \sum_{v \in V} \gamma(v) p(v) + \hat{\Phi}(\gamma, \{\alpha_i\}, \mu) \right\},$$

where

$$\begin{aligned} \hat{\Phi}(\gamma, \{\alpha_i\}, \mu) = \int_{(v,m) \in V \times M} & \left[\sum_{i=1}^2 t_i(m) - \gamma(v) \right. \\ & \left. + \sum_{i=1}^2 \alpha(m_i) \left(v \frac{\partial q_i}{\partial m_i}(m) - \frac{\partial t_i}{\partial m_i}(m) \right) \mathbb{I}_{m_i < 1} \right] \mu(dv, dm). \end{aligned}$$

Then the dual problem is

$$\sup_{(\gamma, \{\alpha_i\}) \in \Gamma \times A^2} \hat{\Phi}(\gamma, \{\alpha_i\}). \quad (D)$$

Let us denote by $\widehat{R}(\mathcal{A})$ the optimal value for the dual problem.

Lemma 1 (Weak Duality).

For any regular mechanism \mathcal{A} , the optimal value for the dual problem (D) is weakly less than the optimal value for the relaxed primal problem (P'), and hence weakly less than the optimal value of the seller's problem, i.e., $\widehat{R}(\mathcal{A}) \leq \widetilde{R}(\mathcal{A}) \leq R^$.*

This result, which is standard for finite dimensional linear programs, is only slightly less trivial in this infinite dimensional setting. For the sake of completeness, we have included a proof in the Appendix.

Note that it is always possible to achieve a finite value for $\Phi(\gamma, \{\alpha_i\})$ by choosing the multipliers to be identically zero. Indeed, a necessary condition for Φ to be non-negative is that

$$\sum_{i=1}^2 \left(t_i(m) + \alpha(m_i) \left(v \frac{\partial q_i}{\partial m_i}(m) - \frac{\partial t_i}{\partial m_i}(m) \right) \mathbb{I}_{m_i < 1} \right) \geq \gamma(v). \quad (6)$$

Otherwise, it is possible to choose a μ that places infinite mass on this particular (v, m) , which would yield a value for Φ of $-\infty$. Moreover, any choice of μ is obviously weakly dominated by a μ that puts zero mass on those (v, m) for which (6) holds as a strict inequality. Thus, it is without loss of generality to rewrite the dual problem as

$$\sup_{(\gamma, \{\alpha_i\}) \in \Gamma \times A^2} \sum_{v \in V} p(v) \gamma(v) \quad \text{subject to (6)}. \quad (D')$$

This is the final form of the optimization problem, from which we will derive a maxmin mechanism.

4.2 A maxmin mechanism

We have asserted that the upper bound on R^* constructed in Section 3 is tight. To demonstrate this fact, we will construct a mechanism \mathcal{A}^* for which $\widehat{R}(\mathcal{A}^*)$, the optimal value for the dual problem (D'), is exactly equal to \overline{R} . In turn, this will be demonstrated by constructing a particular pair of multiplier functions such that the dual objective value is \overline{R} .

Under the assumption that \overline{R} is the solution to the seller's problem, we already know a great deal about the solution to the dual problem for a maxmin mechanism. In particular, $\gamma(v)$ can be interpreted as the shadow value of relaxing the probability constraints in terms of minimum revenue. But given the formula for minimum revenue defined in (1) and (2), we can compute this shadow value directly. Specifically,

$$a \left(1 - \log(a) + \frac{1}{2} \log^2(a) \right) = \frac{p(1)}{p(0) + p(1)}.$$

If we take the derivative of both sides with respect to the probabilities, we conclude that

$$\begin{aligned} \frac{\partial a}{\partial p(0)} &= -\frac{2}{\log^2(a)} \frac{p(1)}{(p(0) + p(1))^2} = -\frac{2p(1)}{\log^2(a)}; \\ \frac{\partial a}{\partial p(1)} &= \frac{2}{\log^2(a)} \frac{p(0)}{(p(0) + p(1))^2} = \frac{2p(0)}{\log^2(a)}, \end{aligned}$$

since $p(0) + p(1) = 1$. Moreover, total revenue is

$$R = (p(0) + p(1)) a (1 - \log(a)),$$

so that the derivatives of revenue with respect to the probabilities are

$$\begin{aligned} \frac{dR}{dp(0)} &= a(1 - \log(a)) + \frac{2}{\log(a)} p(1) \equiv \gamma^*(0); \\ \frac{dR}{dp(1)} &= a(1 - \log(a)) - \frac{2}{\log(a)} p(0) \equiv \gamma^*(1). \end{aligned}$$

We will construct a feasible solution to (D') with γ^* as multipliers on probability constraints.

In addition, we will choose the multipliers on local upward incentive constraints to be constant and equal to

$$\alpha^* = -\frac{1}{\log(a)}.$$

This deserves some explanation. As with Du (2016), we are assuming that the shadow cost of relaxing the local incentive constraint is the same for both players and the same for all messages. In addition, our relaxation is effectively assuming that only local incentive constraints bind for a revenue minimizing BCE on a maxmin mechanism. Why should this be the case? The sufficiency of these constraints is motivated by the minmax type space, in which optimal revenue is completely determined by local incentive constraints. As to why the multipliers should be constant, that is in a sense a normalization that is without loss of generality. For if the local incentive constraints did not have constant multipliers, then we could relabel the messages by moving some closer together and others further apart (but preserving the same ordering), which would effectively change the “size” of a local deviation so as to make all of the multipliers the same. The fact that the correct constant value is $-1/\log(a)$ is more subtle, though this choice shall be vindicated by our subsequent analysis.

Our maxmin mechanism \mathcal{A}^* consists of the message space $M_i^* = [0, 1]$, the allocation rule:

$$q_i^*(m_i, m_j) = \begin{cases} m_i & \text{if } m_1 + m_2 < 1; \\ \frac{1}{2}(1 + m_i - m_j) & \text{if } m_1 + m_2 \geq 1, \end{cases}$$

and the transfer rule:

$$t_i^*(m_i, m_j) = \log(a) a^{-m_i} \int_{x=0}^{m_i} a^x \xi(x, m_j) dx,$$

where

$$\xi(x, y) = \begin{cases} \frac{\gamma^*(0)}{2} & \text{if } m_1 + m_2 < 1; \\ \frac{1}{2}(\gamma^*(1) - \alpha^* + \log(a) ((x - \beta)^2 - (y - \beta)^2)) & \text{if } m_1 + m_2 \geq 1, \end{cases} \quad (7)$$

and

$$\beta = \frac{1}{2} \left(1 - \frac{1}{\log(a)} \right).$$

This completes the specification of the mechanism. Figure 1 depicts the allocation and transfer rules when $p(0) = p(1) = 1/2$.

The allocation rule q^* has multiple interpretations. Aumann and Maschler (1985) point out that this allocation is associated with various standard solutions of an associated cooperative game, in which each player has a “claim” on a good with fixed total supply. This claim represents the maximum amount of the good that the player can use, and below this amount utility is proportional to the quantity received. The characteristic function for this game assigns to each coalition the total supply of the good less the claims of players who are not in the coalition, or zero in the event that others’ claims exceed the supply. In the two




Figure 1: The mechanism \mathcal{A}^* when $p(0) = p(1) = 1/2$. Clockwise from top left: The allocation rule $q_i^*(m_i, m_j)$, the transfer rule $t_i^*(m_i, m_j)$, $\xi(m_i, m_j)$, and the total probability that some bidder is allocated the good.

player case, q^* is the allocation of the good that is identified by both the Shapley value as well as the nucleolus. The former has the interpretation that the buyers arrive in a random order, and that the claims are satisfied in the order that players arrive. Thus, each buyer is equally likely to arrive first or second. The first buyer's claim is completely filled, and the second buyer gets the maximum of his claim and the amount left over from the satisfying first buyer's claim.

The transfer rule we find harder to interpret. Clearly, there is an exponential shape to the transfer as a function of a buyer's own report. Note that t_i is the solution to the following differential equation:

$$\frac{\partial t_i^*}{\partial m_i}(m_i, m_j) = \frac{1}{\alpha^*} [t_i^*(m_i, m_j) - \xi(m_i, m_j)].$$

Thus, $-\xi$ can be interpreted as an "excess" growth rate of the transfer rule. (We say excess, though ξ can be positive or negative depending on the message profile.) The dual incentive constraints when $m_i < 1$ for $i = 1, 2$ can therefore be interpreted as a lower bound on the sum of the excess growth rates,

$$\sum_{i=1}^2 \xi(m_i, m_j) \geq \gamma^*(v) - \alpha^* v \sum_{i=1}^2 \frac{\partial q_i(m)}{\partial m_i}.$$

However, these constraints allow for the growth in transfers to be distributed asymmetrically across the buyers. According to (7), the growth rate is symmetric when the demands sum to less than one, but are asymmetric when the demands are incompatible. This is necessary in order to satisfy the dual constraints at the boundary when $m_i = 1$ for at least one i .

Let us now verify that (γ^*, α^*) are feasible for the dual problem (D') for mechanism \mathcal{A}^* . Note that for $m_i \in [0, 1)$,

$$\xi(m_i, m_j) + \xi(m_j, m_i) = \begin{cases} \gamma^*(0) (= \gamma^*(1) - 2\alpha^*) & \text{if } m_1 + m_2 < 1; \\ \gamma^*(1) - \alpha^* & \text{if } m_1 + m_2 \geq 1 \end{cases}$$

and

$$\sum_{i=1}^2 \frac{\partial q_i^*}{\partial m_i}(m) = \begin{cases} 2 & \text{if } m_1 + m_2 < 1; \\ 1 & \text{if } m_1 + m_2 \geq 1. \end{cases}$$

Thus, as long as both messages are strictly less than one, the left-hand side of (6) with $v = 0$ reduces to $\gamma^*(0)$ if $m_1 + m_2 \leq 1$ and $a(1 - \log(a)) \geq \gamma^*(0)$ if $m_1 + m_2 > 1$. When $v = 1$,

the left-hand side of (6) reduces to $\gamma^*(0) + 2\alpha^* = \gamma(1)$ when $m_1 + m_2 < 1$, and it reduces to $a(1 - \log(a)) + \alpha^* = \gamma^*(1) \geq \gamma^*(0)$ when $m_1 + m_2 \geq 1$.

Thus, it only remains to check that the dual constraints are satisfied on the boundary region where $m_i = 1$ for some i . The proof is not technically difficult, but it is lengthy, so we have relegated it to the Appendix. This completes the proof of the following result:

Lemma 2 (Dual Feasibility).

The multipliers (γ^, α^*) are feasible for the dual problem for the mechanism \mathcal{A}^* . As a result, $\widehat{R}(\mathcal{A}^*) \geq \overline{R}$.*

In turn, we have shown that:

Proposition 2 (Lower Bound on Revenue).

Revenue under auction \mathcal{A}^ is at least \overline{R} in every information structure and equilibrium. Thus, R^* is at least \overline{R} .*

Proof of Proposition 2. For the multipliers $(\gamma^*, \{\alpha_i^*\})$, the dual objective is simply

$$p(0)\gamma^*(0) + p(1)\gamma^*(1) = a(1 - \log(a)) = \overline{R}.$$

Thus, the solution to the primal problem with local constraints (P') must be weakly greater than \overline{R} , as must the solution to the primal problem (P). \square

As a corollary of Propositions 1 and 2, we obtain our main result:

Theorem 1 (Solution to the Seller's Problem).

Maxmin revenue is $\overline{R} = a(1 - \log(a))$, where a solves (1). Moreover, \mathcal{A}^ is a maxmin mechanism, and \mathcal{S}^* is a minmax information structure.*

We remark that both \mathcal{A}^* and \mathcal{S}^* satisfy the hypotheses of Milgrom and Weber (1985) for there to exist an equilibrium in distributional strategies, which implies the existence of an equilibrium in behavioral strategies (i.e., strategies as we have defined them). Proposition 1 then implies that revenue in this equilibrium must be precisely \overline{R} . In fact, it turns out that there is an equilibrium in the monotonic, symmetric, and pure strategies

$$\sigma_i^*(s_i) = \begin{cases} \frac{\log(1-s_i)}{\log(a)} & \text{if } m_i < 1 - a; \\ 1 & \text{if } m_i \geq 1 - a. \end{cases}$$

In this equilibrium, the good is allocated with probability one if and only if the interim expected value is exactly one. We formalize this as the following result, whose proof is in the Appendix:

Proposition 3 (Saddle-Point Equilibrium).

The strategies σ^ constitute an equilibrium of the maxmin mechanism on the minmax type space.*

This concludes our characterization of maxmin revenue and a maxmin auction. At this point, it is useful to revisit our original motivation, and ask where we stand relative to the goal of identifying simple, robust, detail-free mechanisms. We have demanded that our maxmin mechanism provide favorable revenue guarantees even under arbitrarily rich models of information, in which signals can be high-dimensional and correlated in an arbitrary manner, thus generating very rich forms of interdependence. Out of this universe of models, we have concluded that the critical worst-case model has one-dimensional signals, signals that are statistically independent, and an interim expected value that is a continuous and increasing function of the signals.

We have also allowed the seller to use mechanisms that employ arbitrarily rich message spaces and we have allowed allocation and transfer rules that can depend on messages in an essentially arbitrary manner. All we have required is that the seller allow the buyers to walk away from the transaction if they so desire. We have concluded that there exists a maxmin mechanism with one-dimensional messages, allocation and transfer rules that are continuous and monotonically increasing in the messages, and have a simple interpretation in terms of bidders making demands, which are settled according to a simple and egalitarian procedure. Thus, there is already a strong case that our methodology has dramatically reduced the complexity of the auction design problem and it has led us to a tractable and well-behaved auction design.

This mechanism is unlike any other than we know of, and we do not know of any auctions that are in regular use that fit its description. One reason may be that the mechanism requires the seller to randomize over the allocation, conditional on the reports. Such randomization may be difficult for a seller to commit to in a credible way, though it seems to be a necessary feature of the optimal mechanism as we have defined the seller's problem. It is an interesting question for future research how the form of the optimal mechanism might change if we imposed determinacy as an additional constraint on the mechanism. As a final remark, we have conducted extensive simulations, which were helpful in identifying the form of the maxmin auction constructed above. We have not proven, nor do we believe it to be the case, that the minmax type space and the maxmin auction are unique. The ones that we have found seem extremely "regular," in the sense that the message spaces are compact and allocation and transfer rules are continuous. We suspect that there may exist other maxmin auctions which satisfy the same, and might be even easier to interpret and implement. We regard this as a promising avenue for further exploration.

5 Discussion

5.1 Comparison with other mechanisms

To put the maxmin mechanism in context, let us compare it to other mechanisms that one might consider using, in the example where $p(0) = p(1) = 1/2$. First, let us observe that for this example, the parameter that solves (1) is $a \approx 0.6897$, which yields maxmin revenue of $\bar{R} \approx 0.2534$. This is approximately 50.7% of the total surplus of $1/2$.

How would other well-known mechanisms perform? Under certain solution concepts and in certain settings, the second-price auction is thought to be quite robust. In particular, in a private value model, where other buyer's signals contain no information about one's own value, the second-price auction has a unique equilibrium in weakly-undominated strategies. In the present interdependent value model, however, the second-price auction performs spectacularly poorly: there is always an equilibrium in which one bidder bids zero and the other bids one, which would result in revenue of zero. We select for the worst equilibrium for the seller, but even if we did not, it is not clear how to select for an equilibrium that is better for the seller. The reason is that weak dominance has very little bite when there can be arbitrary interdependence in values. Bidding a very small amount is weakly undominated as long as there is a non-zero chance that the value conditional on winning with a higher bid would be small.

What about posted-price mechanisms? The seller offers the good at a price of $\rho \in [0, 1]$, and bidders indicate whether or not they want to buy the good at that price. The good is then randomly allocated between the bidders who want to buy, and the winner of the good pays the posted price. It turns out that the worst model of information for this mechanism is symmetric, with the bidders never purchasing the good when the value is zero and not purchasing when the value is high with probability x such that

$$\frac{x}{1+x} = \rho \implies x = \frac{\rho}{1-\rho}.$$

Thus, conditional on not purchasing, the expected value is exactly equal to the price, so that buyers are indifferent. The optimal price is therefore $\rho = 1 - 1/\sqrt{2}$, which yields a maximum revenue of $3 - 2\sqrt{2} = 0.1716$, which yields about 34.3% of the total surplus and 67.7% of maxmin revenue.

One can do one step better by using a randomized posted price, rather than a deterministic posted price. In particular, Kos and Messner (2015) and Carrasco et al. (2015) have shown that a randomized posted price can attain maxmin revenue with a single buyer, and indeed the maxmin revenue is given by the parameter b that solves $b(1 - \log(b)) = 1/2$

(cf. the discussion of generalized minmax type spaces in the next section). The solution is $b \approx 0.1867$, which is approximately 73.7% of maxmin surplus.

Next, the working paper Bergemann et al. (2015b) characterizes minimum revenue for first-price auctions with discrete and common values. The revenue minimizing distribution of bids has the feature that bidders are indifferent to all *uniform upward deviations*, where a uniform upward deviation up to a bid of b means that bidding the maximum of b and whatever bid x would have been made in equilibrium. If $H(b|v)$ is the cumulative distribution of bids in equilibrium conditional on the value, then the uniform upward incentive constraint is that

$$\sum_{v \in V} (v - b) H(b|v) p(v) \leq \sum_{v \in V} p(v) \int_{x=0}^b H(x|v) dx.$$

One can show that it is without loss of generality to take $H(0|0) = 1$, so that this integral equation becomes

$$(1 - b) H(b|1) - b \leq b + \int_{x=0}^b H(x|1) dx.$$

The solution to this inequality as an equality, with the initial condition $H(0|1) = 0$, is given by

$$H(b|1) = \frac{b(2 - b)}{(1 - b)^2},$$

which hits 1 at

$$\bar{b} = 1 - \frac{1}{\sqrt{2}}.$$

Thus, each bidder is indifferent between their equilibrium strategy and winning with probability one at a bid of \bar{b} . This tells us that each bidder's equilibrium surplus must be

$$U_i^{FPA} = \frac{1}{2} - \bar{b} = \frac{1}{\sqrt{2}} - \frac{1}{2},$$

so that revenue is

$$R^{FPA} = \frac{1}{2} - 2U_i^{FPA} = \frac{3 - 2\sqrt{2}}{2} \approx 0.0858,$$

which is half of what the seller could obtain with the optimal deterministic posted price and about 33.9% of optimal maxmin revenue. Naturally, this bound could be improved if one also used deterministic or randomized reserve price.

Finally, Du (2016) has constructed a class of mechanisms in which $M_i = [0, 1]$,

$$q_i(m_i, m_j) = m_i - \min\{m_1, m_2\} / 2$$

and

$$t_i(m_i, m_j) = k(1 - c^{-m_i})$$

for some constants c and k . Maxmin revenue within this class of mechanisms (maximizing across (k, c) and minimizing across information structures and equilibria) is approximately 0.2403, which is about 94.8% of maxmin revenue.

Thus, the maxmin mechanism provides a significantly better guarantee than many other well-known mechanisms. At the same time, one advantage of having characterized maxmin revenue is that it allows us to put into context the revenue performance of various other mechanisms. How much revenue might be lost by using a mechanism that is non-maxmin optimal? Put differently, what cost would we have to assign to the complexity of the maxmin mechanism to be willing to use one of these alternative auction formats? We can see that some of these simpler auctions do fairly well relative to the best one could hope for. The first-price auction, for example, is completely detail free, in that its specification does not even depend on the prior distribution of values, and it still obtains upwards of one-third of the best revenue guarantee. The mechanism of Du (2016), which has a significantly simpler transfer rule than the maxmin mechanism, obtains nearly 95% of the best revenue guarantee and asymptotically extracts all of the surplus.

5.2 Beyond two players and binary values

Our construction of the maxmin mechanism is (for now) special to two players and the case in which the ex-post value has only two points in its support. We are optimistic that the model can be substantially generalized, to models with many bidders and many values. Indeed, there is a natural generalization of the minmax information structure \mathcal{S}^* to these cases, which generates an upper bound on maxmin revenue and a conjecture for the maxmin itself.

Consider a model with $n \geq 1$ bidders and in which there is a value that is drawn from some cumulative distribution function $P(v)$ with support contained in $[0, 1]$. In the generalized minmax type space, the bidders once again have signals that are independent uniform draws, and the interim expectation of the value conditional on the entire profile of signals is

$$v(s_1, \dots, s_n) = \min \left\{ \frac{a}{(1 - s_1) \cdots (1 - s_n)}, b \right\},$$

where $0 \leq a \leq b \leq 1$ are parameters of the information structure. Let us denote this information structure by $\mathcal{S}_{n,a,b}^*$. This information structure again admits a revenue equivalence

formula, such that if the equilibrium allocation is described by $q_i(s)$, then revenue is

$$\int_{s \in [0,1]^n} \sum_{i=1}^n \psi_i(s) q_i(s) ds,$$

where

$$\psi_i(s) = v(s) - (1 - s_i) \frac{\partial v}{\partial s_i}(s) = \begin{cases} 0 & \text{if } v(s) < 1; \\ 1 & \text{if } v(s) = b. \end{cases}$$

Thus, this type space preserves the tremendous indifference on the part of the seller: any mechanism is optimal as long as it always allocates the good when the value is one, and the seller is indifferent as to whom the good is allocated and whether or not the good is allocated when $v(s) < b$.

Let us denote by $F_{n,a,b}(v)$ the cumulative distribution of the interim expected value $v(s)$ on this type space. Then necessary and sufficient conditions for $\mathcal{S}_{n,a,b}^*$ to be consistent with the ex-ante distribution of the value is that P is a mean preserving spread of $F_{n,a,b}$, or that

$$\int_{x=0}^v F_{n,a,b}(x) dx \leq \int_{x=0}^v P(x) dx$$

for every $x \in [0, 1]$ and

$$\int_{x=0}^1 F_{n,a,b}(x) dx = \int_{x=0}^1 P(x) dx.$$

We conjecture that for a general distribution P , the minmax type space is of the form $S_{n,a,b}^*$ for some (a, b) .

In fact, it is fairly easy to compute optimal revenue on this type space. Let $\bar{R}_{n,a,b}$ denote optimal revenue on $\mathcal{S}_{n,a,b}^*$, and let $E_{n,a,b}$ denote the expected value. Then

$$\bar{R}_{n,a,b} = E_{n-1,a,b}$$

and

$$E_{n,a,b} = \int_{x=0}^{1-\frac{a}{b}} E_{n-1,a/(1-x),b} dx + a.$$

The reason is that the expected value, conditional on one bidder having a signal $s_i \in [0, 1 - a/b]$, is exactly $E_{n,a/(1-s_i),b}$, since the signal effectively increases the constant a in the definition of the interim expected value. On the other hand, if $s_i \in [a/b, 1]$, then the value is necessarily equal to b , and this event occurs with probability a/b . Moreover, an optimal mechanism for this type space is setting a posted price equal to the expected value given a signal of $s_i = 0$, which is simply the expectation of the interim expected value with

$n - 1$ buyers. Straightforward calculation reveals that

$$E_{n,a,b} = \bar{R}_{n,a,b} = a \sum_{k=0}^n \frac{1}{k!} \left(-\log \left(\frac{a}{b} \right) \right)^k.$$

To see this, simply observe that

$$\int_{x=0}^1 \frac{1}{1-x} \left(\log \left(\frac{a}{b} \right) - \log(1-x) \right)^k dx = -\frac{1}{k+1} \left(\log \left(\frac{a}{b} \right) \right)^k.$$

In the case where the value is either zero or one (i.e., there are only mean and support restrictions), revenue is minimized by setting $b = 1$ and setting a such that $E_{n,a,1}$ is equal to the ex-ante expected value. More generally, if there are more values in the support of P , then the optimal b for minimizing revenue may be less than 1, so as to satisfy the second-order stochastic dominance constraints. For example, for the uniform distribution with two bidders, the optimal parameters for minimizing revenue are approximately $a = 0.0891$ and $b = 0.8023$, which yield an upper bound $\bar{R}_{2,a,b} \approx 0.2848$.

We conjecture that this generalized upper bound is tight. We also hope that the proof technique used in the two-buyer case can be generalized to many-buyer and many-value models. In particular, under the hypothesis that the upper bound is tight, we can derive the correct multipliers on probability constraints from the formula for minmax revenue. We can also conjecture that the generalized maxmin mechanism continues to have one-dimensional signals that can be interpreted as “demands,” and that all of the demands are filled when they sum to less than one. The constant multiplier on the local incentive constraints would then be pinned down by the requirement that all of the dual constraints bind when the demands sum to less than one. Indeed, in the many-player binary-value case, this leads to generalized multipliers of

$$\begin{aligned} \gamma^*(0) &= E_{n-1,a,1} + \frac{np(1)}{\log(a)}; \\ \gamma^*(1) &= E_{n-1,a,1} - \frac{np(0)}{\log(a)}; \\ \alpha^* &= -\frac{1}{\log(a)}. \end{aligned}$$

Moreover, we conjecture that at least in the binary-value model, there is a maxmin mechanism in which the allocation has the form that the buyers’ demands are filled sequentially, in a random order, as in the two-player case. If we let Z denote the set of permutations of

$\{1, \dots, n\}$, then this allocation is

$$q_i^*(m) = \frac{1}{n!} \sum_{\zeta \in Z} \min \left\{ m_i, \max \left\{ 0, 1 - \sum_{\{j | \zeta(j) < \zeta(i)\}} m_j \right\} \right\}.$$

This allocation corresponds to the Shapley value of the n -player cooperative bargaining game described in Section 4. These conjectures are supported by numerical simulations that we have conducted with three buyers. The only step remaining to go from conjecture to theorem is to identify a maxmin transfer rule such that these multipliers and the random sequential service allocation are feasible for the dual constraints. We are actively pursuing such a generalized transfer rule.

6 Conclusion

The purpose of this paper has been to characterize auctions which provide the best possible revenue guarantees for a seller who faces large model uncertainty. We have shown that the worst model for the seller has a simple description in which bidders' signals are independent and the value function has a particular form that makes the seller indifferent over a wide array of equilibrium allocations. We have shown that there exists a maxmin auction in which the bidders make demands, and demands are filled sequentially in a random order. When paired with carefully chosen transfer rule, this allocation is part of a maxmin optimal auction. The auction is simple and well-behaved: allocations and transfers are continuous and increasing functions of the reports, and moreover, buyers can always opt out by sending the lowest demand, in which case they pay no transfer and are not allocated the good.

Many questions remain. Are there other mechanisms that are also maxmin optimal, and that would be superior to the one we constructed for reasons outside the model? What would maxmin auctions look like with more than two buyers, or with more general distributions of values? What would happen if we impose reasonable restrictions on the models of beliefs that would exclude the minmax type space constructed above? In the future, we hope to extend our knowledge of optimal informationally-robust auctions in all of these directions.

References

- AUMANN, R. J. AND M. MASCHLER (1985): “Game theoretic analysis of a bankruptcy problem from the Talmud,” *Journal of economic theory*, 36, 195–213.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015a): “Bayes Correlated Equilibrium in Infinite Games,” Tech. rep., Yale University and University of Chicago and Princeton University.
- (2015b): “First Price Auctions with General Information Structures: Implications for Bidding and Revenue,” Tech. Rep. 2018, Cowles Foundation for Research in Economics, Yale University.
- (2016): “Optimal Auction Design in a Common Value Model,” Tech. rep., Yale University, University of Chicago and Princeton University.
- BERGEMANN, D. AND S. MORRIS (2013): “Robust Predictions in Games with Incomplete Information,” *Econometrica*, 81, 1251–1308.
- (2016): “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games,” *Theoretical Economics*, 11, 487–522.
- BERGEMANN, D. AND K. SCHLAG (2008): “Pricing Without Priors,” *Journal of the European Economic Association Papers and Proceedings*, 6, 560–569.
- (2011): “Robust Monopoly Pricing,” *Journal of Economic Theory*, 146, 2527–2543.
- BROOKS, B. A. (2013a): “Revenue Sharing in Second-Price Auctions,” Tech. rep., Mimeo, University of Chicago.
- (2013b): “Surveying and selling: Belief and surplus extraction in auctions,” Tech. rep., Mimeo, University of Chicago.
- BULOW, J. AND P. KLEMPERER (1996): “Auctions vs Negotiations,” *American Economic Review*, 86, 180–194.
- CARRASCO, V., V. F. LUZ, P. MONTEIRO, AND H. MOREIRA (2015): “Robust Selling Mechanisms,” *Mimeo, University of British Columbia*.
- CHUNG, K.-S. AND J. ELY (2007): “Foundations of Dominant Strategy Mechanisms,” *Review of Economic Studies*, 74, 447–476.

- CONDORELLI, D. AND B. SZENTES (2016): “Buyer-Optimal Demand and Monopoly Pricing,” Tech. rep., Mimeo, London School of Economics and University of Essex.
- CREMER, J. AND R. MCLEAN (1985): “Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands are Interdependent,” *Econometrica*, 53, 345–361.
- (1988): “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, 56, 1247–1258.
- DU, S. (2016): “Robust Mechanisms under Common Valuation,” Tech. rep., Mimeo, Simon Fraser University.
- GILBOA, I. AND D. SCHMEIDLER (1989): “Maxmin Expected Utility with Non-Unique Prior,” *Journal of Mathematical Economics*, 18, 141–153.
- HANSEN, L. P. AND T. J. SARGENT (2001): “Acknowledging misspecification in macroeconomic theory,” *Review of Economic Dynamics*, 4, 519–535.
- HARTLINE, J. D. AND T. ROUGHGARDEN (2009): “Simple versus optimal mechanisms,” in *Proceedings of the 10th ACM conference on Electronic commerce*, ACM, 225–234.
- KOS, N. AND M. MESSNER (2015): “Selling to the Mean,” Tech. rep.
- MCAFEE, P. AND P. RENY (1992): “Correlated Information and Mechanism Design,” *Econometrica*, 60, 395–421.
- MCAFEE, R. P., J. MCMILLAN, AND P. J. RENY (1989): “Extracting the surplus in the common-value auction,” *Econometrica*, 1451–1459.
- MILGROM, P. R. AND R. J. WEBER (1985): “Distributional strategies for games with incomplete information,” *Mathematics of operations research*, 10, 619–632.
- MYERSON, R. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- NEEMAN, Z. (2003): “The Effectiveness of English Auctions,” *Games and Economic Behavior*, 43, 214–238.
- ROESLER, A.-K. AND B. SZENTES (2016): “Buyer-Optimal Learning and Monopoly Pricing,” Tech. rep., Mimeo, London School of Economics and University of Michigan.
- YAMASHITA, T. (2015): “Revenue guarantee in auction with common prior,” Tech. rep., Mimeo, Toulouse School of Economics.

A Appendix

Proof of Lemma 1. Let μ be any BCE that is feasible for the primal problem (P'). Then the probability and incentive constraints (3) and (5) are satisfied everywhere, so that for any multipliers $\gamma \in \Gamma$, we have

$$\sum_{v \in V} \gamma(v) (p(v) - \mu(\{v\} \times [0, 1]^2)) = 0$$

and for any non-negative measurable function $\alpha_i : [0, 1] \rightarrow \mathbb{R}_+$, we have

$$\int_{(v,m) \in V \times M} \alpha_i(m_i) \left(v \frac{\partial q_i}{\partial m_i}(m) - \frac{\partial t_i}{\partial m_i}(m) \right) \mathbb{I}_{m_i < 1} \mu(dv, dm) \leq 0.$$

Essentially, any given $\mu \in \mathcal{M}(\mathcal{A})$ defines a (possibly signed) measure $\nu \in \mathcal{M}(\mathcal{A})$ according to

$$\nu(dv, dm) = \left(v \frac{\partial q_i}{\partial m_i}(m) - \frac{\partial t_i}{\partial m_i}(m) \right) \mathbb{I}_{m_i < 1} \mu(dv, dm)$$

that is absolutely continuous with respect to μ . If μ is feasible for the primal problem, then μ only assigns positive measure to those sets on which the average local gains from deviating up are non-positive, so that ν is also a positive measure. Hence, the expectation of any non-negative function under ν is also non-negative.

As a result,

$$\sum_{v \in V} \gamma(v) p(v) + \hat{\Phi}(\gamma, \{\alpha_i\}, \mu) \leq \int_{(v,m) \in V \times M} \sum_{i=1}^2 t_i(m) \mu(dv, dm).$$

Since $\hat{\Phi}(\gamma, \{\alpha_i\})$ is the infimum of the left-hand side over all measures $\mu \in \mathcal{M}(\mathcal{A})$ (even those that are not feasible for the primal problem), it must be the case that

$$\hat{\Phi}(\gamma, \{\alpha_i\}) \leq \int_{(v,m) \in V \times M} \sum_{i=1}^2 t_i(m) \mu(dv, dm)$$

for all μ that are feasible for problem (P'). Since this is true for all choices of multipliers, we conclude that $\hat{\Phi}^* \leq R^{**}$. \square

Proof of Lemma 2. It remains to verify that the dual constraints (6) are satisfied for the boundary cases where $m_i = 1$ for at least one player. Let us first consider the case in which

exactly one player sends a message of 1. We will show that

$$t^*(1, x) + \xi(x, 1) = \gamma^*(1) - \frac{\alpha^*}{2} \geq \gamma^*(0) - \frac{\alpha^*}{2},$$

so that the dual constraints are satisfied for both $v = 0$ and $v = 1$. Let us first verify that the equality holds at $x = 0$. In this case,

$$t^*(1, 0) = \frac{\gamma^*(0)}{2} \left(1 - \frac{1}{a}\right)$$

and

$$\begin{aligned} \xi(0, 1) &= \frac{1}{2} (\gamma^*(1) - \alpha^* + \log(a)(1 - \beta)^2 - \log(a)\beta^2) \\ &= \frac{1}{2} (\gamma^*(1) - \alpha^* - 1). \end{aligned}$$

Thus, using the fact that $\gamma^*(0) = \gamma^*(1) - 2\alpha^*$,

$$\begin{aligned} t^*(1, 0) + \xi(0, 1) &= \frac{1}{2} (\gamma^*(0) + \gamma^*(1) - \alpha^* - 1) - \frac{\gamma^*(0)}{2a} \\ &= \gamma^*(1) - \frac{\alpha^*}{2} - \alpha^* - \frac{1}{2} - \frac{\gamma^*(0)}{2a}. \end{aligned}$$

Finally,

$$\begin{aligned} -\alpha^* - \frac{1}{2} - \frac{\gamma^*(0)}{2a} &= \frac{1}{\log(a)} - \frac{1}{2} - \frac{a(1 - \log(a)) + \frac{2p(1)}{\log(a)}}{2a} \\ &= \frac{1}{\log(a)} - \frac{1}{2} - \frac{(1 - \log(a))}{2} + \frac{(1 - \log(a) + \frac{1}{2}\log(a)^2)}{\log(a)} \\ &= 0 \end{aligned}$$

as desired.

Next, we will show that

$$\frac{d}{dx} [t^*(1, x) + \xi(x, 1)] = 0. \tag{8}$$

Observe that

$$\begin{aligned}
\frac{d}{dx}t^*(1, x) &= -\frac{\log(a)}{a} \int_{y=1-x}^1 a^y \log(a) (x - \beta) dy \\
&\quad - \log(a) a^{-x} \left(\frac{\gamma(1) - \frac{1}{\log(a)} - \gamma(0)}{2} - \frac{1}{2} \log(a) (1 - x - \beta)^2 + \frac{1}{2} \log(a) (x - \beta)^2 \right) \\
&= -\log(a) (1 - a^{-x}) (x - \beta) - a^{-x} \left(\frac{1}{2} + \frac{\log^2(a)}{2} (2\beta - 1) (2x - 1) \right) \\
&= \frac{1}{2} \left[- (1 - a^{-x}) (1 + \log(a) (2x - 1)) - a^{-x} \left(1 + \frac{\log(a)}{a} (2x - 1) \right) \right] \\
&= -\frac{1}{2} (1 + \log(a) (2x - 1))
\end{aligned}$$

In addition,

$$\begin{aligned}
\frac{d}{dx}\xi(x, 1) &= \log(a) (x - \beta) \\
&= \frac{1}{2} (1 + \log(a) (2x - 1))
\end{aligned}$$

so that (8) holds.

Finally, we will show that

$$2t^*(1, 1) = \gamma^*(1) \geq \gamma^*(0).$$

Note that

$$\begin{aligned}
t^*(1, 1) &= \frac{\log(a)}{a} \int_{x=0}^1 a^x \xi(x, 1) dx \\
&= \frac{\log(a)}{a} \int_{x=0}^1 a^x (\gamma^*(1) - \alpha^* - \log(a) (1 - \beta)^2) dx \\
&\quad + \frac{\log^2(a)}{a} \int_{x=0}^1 a^x (x - \beta)^2 dx.
\end{aligned}$$

It is easily verified that

$$\int a^x (x - \beta)^2 dx = \frac{a^x}{\log^3(a)} (\log^2(a) (x - \beta)^2 - 2 \log(a) (x - \beta) + 2) + C.$$

Thus,

$$\begin{aligned}
\int_{x=0}^1 a^x (x - \beta)^2 dx &= \frac{a}{\log^3(a)} (\log^2(a) (1 - \beta)^2 - 2 \log(a) (1 - \beta) + 2) \\
&\quad - \frac{1}{\log^3(a)} (\log^2(a) \beta^2 + 2 \log(a) \beta + 2) \\
&= \frac{a}{\log^3(a)} \left(\log^2(a) \frac{1}{4} \left(1 + \frac{1}{\log(a)}\right)^2 - \log(a) \left(1 + \frac{1}{\log(a)}\right) + 2 \right) \\
&\quad - \frac{1}{\log^3(a)} \left(\log^2(a) \frac{1}{4} \left(1 - \frac{1}{\log(a)}\right)^2 + \log(a) \left(1 - \frac{1}{\log(a)}\right) + 2 \right) \\
&= \frac{1}{\log^3(a)} \left(\frac{\log^2(a) + 5}{4} (a - 1) - \frac{\log(a)}{2} (1 + a) \right).
\end{aligned}$$

In addition,

$$\begin{aligned}
&\frac{\log(a)}{a} \int_{x=0}^1 a^x (\gamma^*(1) - \alpha^* - \log(a) (1 - \beta)^2) dx \\
&= (\gamma^*(1) - \alpha^* - \log(a) (1 - \beta)^2) \left(1 - \frac{1}{a}\right) \\
&= \gamma^*(1) + \frac{1}{\log(a)} - \log(a) \frac{1}{4} \left(1 + \frac{2}{\log(a)} + \frac{1}{\log^2(a)}\right) \\
&\quad - \frac{1}{a} \left(a(1 - \log(a)) + \frac{2p(1)}{\log(a)} - \frac{1}{\log(a)} - \log(a) \frac{1}{4} \left(1 + \frac{2}{\log(a)} + \frac{1}{\log^2(a)}\right) \right) \\
&= \gamma^*(1) + \frac{1}{\log(a)} - \log(a) \frac{1}{4} \left(1 + \frac{2}{\log(a)} + \frac{1}{\log^2(a)}\right) \\
&\quad - (1 - \log(a)) + \frac{2(1 - \log(a) + \frac{1}{2} \log^2(a))}{\log(a)} - \frac{1}{a \log(a)} - \frac{\log(a)}{a} \frac{1}{4} \left(1 + \frac{2}{\log(a)} + \frac{1}{\log^2(a)}\right) \\
&= \gamma^*(1) - \frac{1}{\log(a)} \left(\frac{\log^2(a) + 5}{4} \left(1 - \frac{1}{a}\right) - \frac{\log(a)}{2} \left(1 + \frac{1}{a}\right) \right).
\end{aligned}$$

Combining these expressions into the above formula for $t^*(1, 1)$ yields the desired result. \square

Proof of Proposition 3. We can reparametrize the type space so that types are identified with their equilibrium bids, and so that the maxmin mechanism is in fact a direct mechanism. Under this parametrization, the type s_i is associated with $\min\{1, \log(1 - s_i) / \log(a)\}$, so that a message $m_i \in [0, 1)$ is associated with $s_i = 1 - a^{m_i}$ and a message $m_i = 1$ is associated with the types $[1 - a, 1]$. Thus, the interim expected value given messages m is

$$w(m) = \min\{1, a^{1-m_1-m_2}\},$$

and the distribution of messages is independent and given by the following cumulative distribution function:

$$F(m) = \begin{cases} 1 - a^m & \text{if } m < 1; \\ 1 & \text{otherwise.} \end{cases}$$

This distribution has a density of $-\log(a) a^m$ on the range $[0, 1)$, and it has a mass point of size a on a message of 1.

One can show that the equilibrium utility of a type m_i that reports m'_i is

$$\begin{aligned} U(m'_i|m_i) &= \int_{m_j=0}^1 [w(m_i, m_j) q_i(m'_i, m_j) - t_i(m'_i, m_j)] dF(m_j) \\ &= \frac{am'_i}{2} - a^{1-m'_i} \frac{3a^{m'_i} - 3 - 2\log(a)(a^{m'_i} + m'_i - 1)}{2\log(a)} \\ &\quad + \begin{cases} \frac{a(a^{-m'_i} - 1 - m'_i \log(a) + 2a^{-m_i} m'_i \log(a)(1 - \log(a)(1 - m_i)))}{2\log(a)} & \text{if } m_i > m'_i; \\ \frac{a^{1-m_i} (2 - 2a^{m_i} + 2\log(a)(m_i + m'_i - a^{m_i} m'_i) + \log^2(a)((m_i + m'_i)^2 - 4m'_i))}{4\log(a)} & \text{if } m_i \leq m'_i. \end{cases} \end{aligned}$$

The derivative with respect to m'_i is therefore

$$U'(m'_i|m_i) = \begin{cases} g(m_i) - g(m'_i) & \text{if } m_i > m'_i; \\ \frac{1}{2} (g(m_i) - g(m'_i) + (a^{1-m_i} - a^{1-m'_i})(1 - \log(a)(1 - m'_i))) & \text{if } m_i \leq m'_i, \end{cases}$$

where

$$g(x) = a^{1-x} (1 - \log(a)(1 - x)).$$

Note that the function g is increasing in x . We conclude that $U'(m'_i|m_i) > 0$ when $m'_i < m_i$. When $m'_i \geq m_i$, we can additionally observe that $a^{1-m_i} < a^{1-m'_i}$ (since $a < 1$) so that $U'(m'_i|m_i) \leq 0$. Thus, the function $U(m'_i|m_i)$ is single peaked with a maximum at $m'_i = m_i$, so that the direct mechanism is incentive compatible, as desired. \square