

Optimal auction design with common values: An informationally-robust approach*

Benjamin Brooks Songzi Du

April 26, 2018

Abstract

A Seller can sell a single unit of a good to a group of bidders. The good is costly to produce, and the bidders have a pure common value that may be higher or lower than the production cost. The value is drawn from a prior distribution that is commonly known. The Seller does not know the bidders' beliefs about the value and values each auction mechanism according to the lowest expected profit across all Bayes Nash equilibria and across all common-prior type spaces that are consistent with the given known value distribution. We characterize and construct optimal auctions for such a Seller. The optimal auctions we construct have a simple structure, in which bidders make one-dimensional bids, the total allocation is a function of the aggregate bid, and individual allocations are proportional to bids. We report a number of further results on optimal auction design with common values and maxmin auction design.

KEYWORDS: Mechanism design, information design, optimal auctions, profit maximization, common value, interdependent value, type space, maxmin, Bayes correlated equilibrium, direct mechanism, revenue equivalence, information rent, duality, zero-sum game, gains from trade, competitive limit, central limit theorem, law of large numbers.

JEL CLASSIFICATION: C72, D44, D82, D83.

*Brooks: Department of Economics, University of Chicago, babrooks@uchicago.edu; Du: Department of Economics, Simon Fraser University, songzid@sfu.ca. We are extremely grateful for many helpful discussions with Dirk Bergemann and Stephen Morris, with whom one of the authors collaborated at an earlier phase of this project. We would also like to thank Gabriel Carroll, Piotr Dworzak, Elliot Lipnowski, Doron Ravid, and seminar audiences at Tel Aviv University, Washington University in St. Louis, Penn State, Pittsburgh, Boston University, Georgetown University, INFORMS, CIREQ, and the University of Toronto for helpful comments. We also thank Christian Baker for valuable research assistance. This research has been supported by NSF #1757222.

1 Introduction

1.1 Background

We study the design of profit-maximizing auctions when the bidders have a common value for the good being sold but partial and differential information about that value. Much of the existing work on optimal auction design has focused on the independent private value (IPV) model, in which the bidders have independently distributed values for the good being sold, and each bidder knows only their own value (Myerson, 1981). While the model has yielded valuable insights, there are many cases where it is unsatisfactory, e.g., auctions of financial assets, where to a first order the bidders have a common value for the asset’s future cash flows.

A number of conceptual issues arise in extending optimal auctions to common values. First, the optimal auction design depends heavily on what assumptions we make about bidders’ information. The IPV model has been broadly accepted as a useful benchmark when values are private, but there is no comparably canonical model when values are common. Moreover, even when values are private, the optimal auction can vary widely depending on the information model, from English auctions (Myerson, 1981; Bulow and Klemperer, 1996) to rather exotic full-surplus extracting mechanisms (Cr mer and McLean, 1985, 1988). The wide variation in the optimal form makes it hard to identify auctions that would work well even if the information model is misspecified. This concern is only amplified when we go to common values, where introspection yields even less guidance as to the right model of information.

To address these concerns, we model a Seller who knows the distribution of the common value, but faces ambiguity about the information that bidders have about the value, which is modeled as a common-prior *type space* (Harsanyi, 1967). The Seller is concerned about model misspecification, and evaluates each auction design according to its lowest expected profit across all type spaces and Bayes Nash equilibria, i.e., the auction’s *profit guarantee*. This is depicted in Figure 1, where each red curve corresponds to a different mechanism. These curves depict how profit varies as we vary bidders’ information but hold the mechanism fixed. The profit guarantee is the lowest point on the curve. The Seller’s Problem is to identify the auction that provides the largest such profit guarantee. Such a “maxmin” auction is represented in Figure 1 by the solid red curve.

1.2 Three examples

To further motivate this exercise, let us briefly consider how three “standard” mechanisms would perform under our worst-case criterion. The IPV model provides a profit-maximization foundation for first- and second-price auctions. The two mechanisms perform equally well in the IPV setting and if we consider the profit-maximizing equilibrium, but the second-price auction typically has other equilibria in which profit is very low. The second-price auction is therefore quickly dismissed since we consider minimum profit across all equilibria.

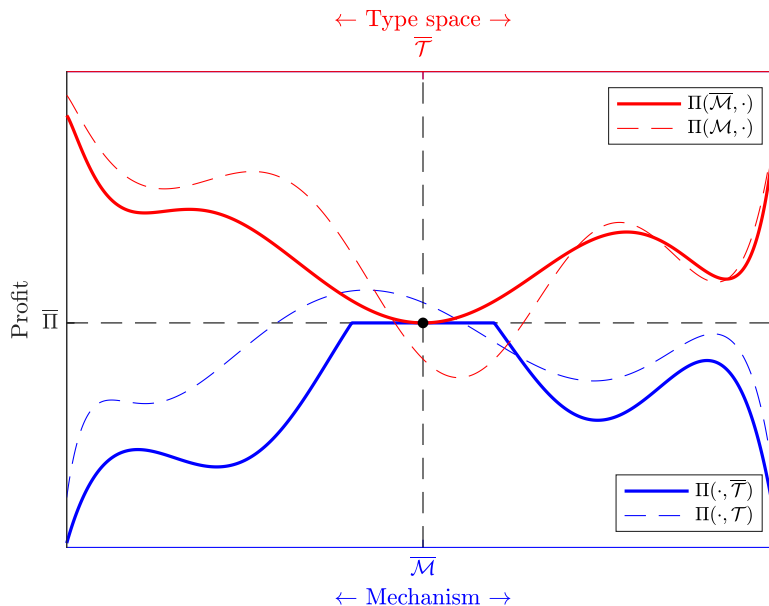


Figure 1: An artist’s impression of the solution to the Seller’s Problem.

The first-price auction, on the other hand, has a non-trivial guarantee. As Bergemann, Brooks, and Morris (2017) show, profit is minimized when the bidders have i.i.d. and one-dimensional signals, and the value is equal to the highest of the bidders’ signals. In this environment, there is very strong adverse inference about the value from winning the auction, which induces the bidders to shade aggressively to avoid a winner’s curse. This effect is so strong that profit would actually be higher if the Seller did not screen at all, and simply posted the highest price at which *all* types would be willing to buy. In fact, this “inclusive” posted price maximizes profit on this type space among auctions that always allocate the good (Bergemann, Brooks, and Morris, 2016b).

The posted price, however, has its own vulnerabilities. The probability of sale is minimized by information that is *public*, so that the events where each bidder is willing to purchase the good overlap as much as possible. The profit-minimizing public signal simply indicates whether the value is above or below a cutoff, and the cutoff is set so that the expected value conditional on being below the cutoff is equal to the price. Thus, learning that the value is below the cutoff makes bidders indifferent between purchasing and not purchasing, and profit is minimized when they break their indifference in favor of not buying the good. Ironically, if the bidders with public information were playing a first-price auction, they would compete the price up to the common expectation of the value and the Seller would obtain all of the surplus as profit.

The takeaway is that different auction formats have different strengths and vulnerabilities vis-à-vis the nature of bidders’ information. First-price auctions perform extremely well when information is public, but are vulnerable to environments with strong adverse selection. For posted prices, the reverse is true. An auction that maximizes the profit

guarantee must find a sweet spot in between these mechanisms and hedge against the environments just described, and all the others.

1.3 Main results

Our main result is to explicitly construct an auction mechanism $\overline{\mathcal{M}}$ that maximizes the profit guarantee and solves the Seller’s Problem. The optimal profit guarantee is denoted by $\overline{\Pi}$. When the number of bidders is large, $\overline{\Pi}$ is approximately the entire ex ante gains from trade, i.e., the ex ante expectation of the value minus the cost of production. While we do not formalize this idea, the guarantee seems to be a substantial share of surplus even when the number of bidders is small. For example, when there are two bidders and the value is standard uniform and there is zero cost of production, the auction we construct guarantees the Seller at least 56 percent of the total surplus as profit. To reiterate, this bound holds across *all* common-prior type spaces and equilibria.

The maxmin auction has a relatively simple structure, in which the bidders make one-dimensional bids, and the allocation and transfer rules are continuous functions of the bids. The allocation rule has the following form. The total probability of the good being sold is a function of the aggregate bid, i.e., the sum of the reports. Conditional on this total “supply,” the good is allocated to each bidder with a likelihood that is proportional to their bid. The fuzziness of the robust allocation rule is in some sense a compromise between the first-price auction, in which the allocation of the good is generically bang-bang,¹ and the posted price, where the good is randomly assigned. This hedges the ambiguity about which allocation generates the most profit, taking into account all effects that the allocation may have on incentives, e.g., inducing the bidders to shade to avoid a winners curse.

This allocation rule, and the accompanying transfer rule, solve a system of partial differential and integral equations that must be satisfied by any sufficiently smooth maxmin mechanism that can be “normalized” in a sense we describe shortly. These equations are a core part of our methodological contribution. The first of these equations pins down the divergence of the allocation probabilities, i.e., the sum of the partial derivatives of each bidder’s allocation with respect to their own message, which we refer to as the *total allocation sensitivity* (TAS). Under a suitable choice of units for messages, the total allocation sensitivity must only depend on the sum of the messages. This critical property is satisfied by the proportional rule.

There are two additional equations that motivate and are solved by the transfer rule we propose. One of these is a condition we refer to as *profit-incentive alignment* (PIA). This equation essentially links ex post profit to the bidders’ local incentive constraints in such a way that in *any* strategy profile at which bids are locally optimal, profit must be at least the optimal guarantee. The final equation is essentially a transversality condition that guarantees that transfers will remain bounded even when messages are large. Imposing boundedness on the transfers excludes pathological mechanisms that have no equilibria.

¹By bang-bang, we mean the good is allocated to a single bidder or is kept by the Seller with probability one.

While there may be many solutions to TAS, PIA, and transversality, the allocation and transfer rules we construct seem to be the simplest.

As we have said, our methodology is to treat the type space as a derived object according to a worst-case criterion. We identify a “minmax” type space that is the hardest for the Seller to optimize profit against, denoted $\overline{\mathcal{T}}$, which directly motivates the maxmin mechanism. $\overline{\mathcal{T}}$ has the following form: The bidders receive signals that are i.i.d. draws from a standard exponential distribution. The interim expectation of the value conditional on the realized signals depends only on the sum of the signals. While the exponential distribution by itself is a normalization, the combination of the exponential shape and the sum being the key statistic of the signal profile is a substantive restriction. We show that under this structure, the Seller is always indifferent as to which bidder should be allocated the good, when information rents due to local incentive constraints are taken into account. The critical minmax value function is the one that maximizes these information rents, subject to the constraint that allocating the good is always weakly better than withholding it. For this type space, optimal profit is equal to the optimal profit guarantee.

While the set of optimal auctions on the minmax type space is vast, this type space nonetheless contains the key to identifying the maxmin auction. When the type space is $\overline{\mathcal{T}}$, any equilibrium of any maxmin mechanism must result in the optimal profit guarantee $\overline{\Pi}$. The system of equations described above are derived from the hypothesis that the maxmin mechanism can be “reduced” by dropping all messages that are not sent in the minmax type space without affecting the profit guarantee. Note that there may well be maxmin mechanisms that are not reducible in this manner, and for which there are messages that are not sent at the minmax type space but would be sent on other type spaces and would lead to higher profit than the reduced mechanism. If a maxmin auction is reducible, we can simply normalize the messages so that at the minmax type space, there is an equilibrium in which bidders submit messages equal to their signals. TAS, PIA, and transversality are essentially necessary conditions on any normalized maxmin mechanism that is sufficiently smooth. We note that the existence of a normalized maxmin mechanism does not follow from the standard revelation arguments.

It is only verified by our construction. We refer to this structure as the *double revelation principle*: The maxmin mechanism is a direct mechanism on the minmax type space, and the minmax type space is a correlated equilibrium on the maxmin mechanism.

Figure 1 depicts the relationship between the minmax type space and the maxmin mechanism. For each type space, equilibrium profit as a function of the mechanism is depicted by a blue curve. The minmax type space is the one that minimizes the highest point on this curve. The solution $(\overline{\mathcal{M}}, \overline{\mathcal{T}})$ that we construct to the Seller’s Problem is a saddle point, in the sense that minimum profit for $\overline{\mathcal{M}}$ across type spaces and maximum profit for $\overline{\mathcal{T}}$ across mechanisms coincide at $\overline{\Pi}$. Note that the solid blue curve corresponding to the minmax type space has a flat at which optimal profit is attained, which reflects the fact that $\overline{\mathcal{M}}$ is one of many optimal mechanisms on this type space.

One can view this saddle point as being an equilibrium of a zero-sum game between the Seller, who chooses the mechanism to maximize profit, and adversarial Nature, who chooses the type space to minimize profit. A subtlety in modeling and solving this game is

that for a fixed mechanism and type space, there can be more than one equilibrium with potentially different levels of profit. Thus, one might be concerned that the solution to the Seller’s Problem would depend on the equilibrium selection rule. This turns out not to be the case. For the solution we identify, neither the Seller nor Nature can move profit in their preferred direction by changing the mechanism or type space, respectively, *even if the deviating party can also select the equilibrium*. The irrelevance of the equilibrium selection rule to payoffs is quite surprising and, in our view, a normatively desirable feature of the solution. Mechanism design has developed various tools for “killing off” bad equilibria, such as augmenting the mechanism with integer-game-like constructions from the full-implementation literature (Maskin, 1999) or perturbing the mechanism and using iterated deletion arguments (Abreu and Matsushima, 1992). These constructions introduce their own conceptual issues, related to unboundedness of the mechanism or the complexity of reasoning and sensitivity to small payoff differences. In contrast, our mechanism obtains the optimal profit guarantee exactly and in all equilibria, while using no more than the hypothesis that bidders’ strategies are locally maximizing a smooth objective function. Even though our description above involved a non-compact signal/message space, it is easy to compactify the mechanism by adding an infinite message, while maintaining continuity of the transfers, and introducing only a “mild” discontinuity in the allocation, so that the mechanism still satisfies permissive equilibrium existence theorems.

1.4 Additional results and related literature

This paper lies at the intersection of the literatures on mechanism design and information design. With respect to the former, we build on the seminal paper of Myerson (1981), and also subsequent work by Bulow and Klemperer (1996). We employ both the revelation principle and the revenue-equivalence formula to show that $\overline{\mathcal{M}}$ maximizes profit across all mechanisms and equilibria given the type space $\overline{\mathcal{T}}$. Similarly, we use tools from the information design literature to show that $\overline{\mathcal{T}}$ is a profit-minimizing information structure given the mechanism $\overline{\mathcal{M}}$. In particular, we use the incomplete information solution concept *Bayes correlated equilibrium* (BCE), which was first introduced by Bergemann and Morris (2013, 2016) and was applied to analyze the range of possible welfare outcomes in first-price auctions by Bergemann, Brooks, and Morris (2017). BCE is a counterpart to the revelation principle of Myerson (1981) for information design.

The most closely related papers are Du (2016) and Bergemann, Brooks, and Morris (2016a). Du (2016) solves our maxmin auction design problem in the limit when the number of bidders goes to infinity, and under the assumption that the Seller’s cost is zero. Specifically, the paper constructs a sequence of mechanisms, one for each number of bidders, and associated lower bounds on profit, such that in the limit as the number of bidders becomes large, the lower bounds converge to the expected surplus. The proof of the result uses duality arguments that are similar to the ones we employ. While the mechanisms from Du (2016) are optimal in the many-bidder limit, they do not achieve the optimal profit guarantee when the number of bidders is finite and more than one. In contrast, Bergemann, Brooks, and Morris (2016a) solves our maxmin auction design problem for the special case where there are exactly two bidders and two possible values,

one for which the gains from trade are positive and one for which the gains are zero. The proof strategy of Bergemann et al. shares some features with the one employed here, in that they construct a saddle point consisting of a maxmin mechanism and a minmax type space, and they also use duality arguments to bound profit for the maxmin mechanism. The minmax type space we construct generalizes the one they identified, but we construct different maxmin mechanisms that generalize more easily to other settings. Our contribution is to provide a flexible and general theory of maxmin auctions and also to give a clearer understanding of the essential properties that characterize maxmin auctions.

Optimal auction design for a common value good was previously studied by Crémer and McLean (1985, 1988), McAfee, McMillan, and Reny (1989), McAfee and Reny (1992), and Bulow and Klemperer (1996). The optimal auctions in those papers are tailored to particular type spaces. The first four of these papers consider the case where types are sufficiently correlated to permit the Seller to extract all of the surplus as profit. We are partly motivated by the concern that such mechanisms are fragile and overly sensitive to the specification of the environment. Bulow and Klemperer (1996) characterize optimal auctions in interdependent value settings that satisfy a generalized Myerson-regularity condition, which implies that the optimal auction is essentially an English auction, with a reserve price for the high bidder that depends on the $N - 1$ lowest reports.

Chung and Ely (2007), Chen and Li (2016), and Yamashita (2016) also study auction design when the Seller does not know the type space, but where values are private. The main result is that in some cases, the Seller can do no better than to use the optimal dominant-strategy mechanism, which is insensitive to bidders' beliefs when values are private. While the question is conceptually similar to ours, their optimal auctions are quite different and depend upon the assumptions that values are private and that the Seller's preferred equilibrium is always played. Our optimal auction generally does not have an equilibrium in weakly-dominant strategies or an ex post equilibrium. Other conceptually related studies of robust auction design are Neeman (2003), Brooks (2013), Yamashita (2015), and Carroll (2016).

We only consider ambiguity on the part of the Seller, and not on the part of the bidders. Bose, Ozdenoren, and Pape (2006) study optimal auction design in the IPV setting, but where both the Seller and the bidders have maxmin expected utility preferences and contemplate multiple priors over the values. Wolitzky (2016) studies bilateral trade where the agents have maxmin expected utility preferences.

After our main results, we extend the theory to the case where the Seller has to sell the good. All of our tools carry over with minor modifications. Also, the maxmin auction is not unique, and we present a generalized class of mechanisms which will also solve the Seller's Problem. In a sense the mechanism we construct is the simplest representative of this class. We will argue, however, that the optimal profit guarantee is unique, if one excludes exotic mechanisms and type spaces that manipulate the existence of equilibrium.

In Section 5, we explore a number of additional topics. In Section 5.1, we illustrate numerically how the optimal profit guarantee compares to those of other mechanisms considered in the literature, and as we vary the number of bidders. When there is one bidder, our optimal auction reduces to those in Carrasco et al. (2017) and Du (2016) when

values are binary and when there is a general value distribution, respectively. In Section 5.2 we show that the maxmin auction fully extracts the ex ante gains from trade as the number of bidders tends to infinity, regardless of how information changes as the number of bidders becomes large. This generalizes the result of Du (2016) that the optimal revenue guarantee is equal to the expected value in the competitive limit. We also characterize the optimal rate for this convergence, and we show that the result is true even if the Seller is restricted to efficient mechanisms that always sell the good.

Section 5.3 shows that in the competitive limit, the maxmin allocation simplifies as follows: there is a cutoff in the aggregate message, and the total supply is linearly increasing in the aggregate message below the cutoff and equal to the efficient supply above the cutoff. The transfer rule similarly simplifies, so that the Seller sets a price per-unit message that is equal to an estimate of the value of the good based on the aggregate message.

While we primarily consider a Seller who knows the value distribution but faces ambiguity about the type space, we consider in Section 5.4 what would happen if the value distribution is also misspecified. This connects our work to the literature on algorithmic mechanism design (e.g., Hartline and Roughgarden, 2009), which is concerned with performance guarantees that are independent of the value distribution but typically rely on more structure on information (e.g., IPV). We show that the maxmin mechanism provides a profit guarantee that is linear and weak-* continuous in the value distribution. Remarkably, the optimal profit guarantee converges to the ex ante gains from trade *even if the prior is misspecified*. Thus, whether the prior is correct is immaterial when the number of bidders is large. This is a special feature of the maxmin mechanism we construct, and does not extend to the aforementioned broader class of maxmin mechanisms.

Section 5.6 revisits the classic question of Bulow and Klemperer (1996) comparing the value of additional bidders versus the commitment power to withhold the good.² When the cost of production is zero, we replicate the earlier finding that the optimal can-keep profit guarantee is less than the optimal must-sell profit guarantee with even one additional bidder. If the cost is positive, however, the Seller could be better off with fewer bidders and the ability to commit to withhold the good.

In Section 5.7, we argue that the minmax type space can be microfounded as the information the bidders would *choose* if they could collude in advance on their information but were unable to commit to not compete against one another at the auction stage. This generalizes a result of Roesler and Szentes (2017), who studied this problem with a single bidder.

As a final topic, in Section 5.8 we consider what would happen to welfare on other type spaces that are not the minmax. Our constructions can be used to put bounds on the set of welfare outcomes that could obtain if the Seller uses a profit-maximizing mechanism for the true type space. We also compare the possible welfare outcomes from the maxmin mechanisms to other standard mechanisms, such as first-price auctions and posted prices.

²Bulow and Klemperer interpret their result as comparing simple mechanisms (English auctions) with more bidders to the more complicated optimal mechanism. We reinterpret their revenue ranking as comparing the ability to commit to withhold the good versus additional bidders.

The rest of the paper proceeds as follows. Section 2 describes our model and the Seller’s Problem. Section 3 presents an informal derivation of the solution. Section 4 presents our formal construction and characterization. Section 5 presents our extensions, and Section 6 briefly concludes.

2 Model

2.1 Primitives

For easy reference, all of our notation is compiled in Table 1 in Appendix D. A unit of a good can be allocated to one of N bidders. The bidders have a pure common value for the good v which is distributed according to the cumulative distribution function H on \mathbb{R}_+ . The non-negativity of the value is essentially an assumption that the good satisfies free disposal. We assume the support is bounded, with v_{\max} denoting the maximum of the support. We let $V = [0, v_{\max}]$.

The bidders have preferences over probabilities of receiving the good q_i and the amount they pay for it t_i . An equivalent interpretation is that there is a unit mass of the good and bidders’ valuations are linear in the amount q_i they are allocated. These preferences are represented by the state-dependent utility index

$$u_i(v, q_i, t_i) = vq_i - t_i.$$

There is a constant marginal cost of production $c \geq 0$.³ This cost can also represent the Seller’s value for the good. The Seller’s profit from the profile of allocations $q = (q_1, \dots, q_N)$ and transfers $t = (t_1, \dots, t_N)$ is

$$\Pi(q, t) = \sum_{i=1}^N (t_i - cq_i).$$

We make the non-degeneracy assumption that the support of H is not a singleton, and also that the expected value is at least c , so that the ex ante expected gains from trade are non-negative.⁴

2.2 Information

Recall that a real-valued random variable X is a *mean-preserving spread*, or garbling, of a random variable Y if X is equal to $Y + \epsilon$, where ϵ has zero conditional expectation given

³It is customary in private-value auction models to normalize the cost to zero. The reason we do not do so here is that we want to allow for the possibility that the gains from trade could be negative. As we shall see, allowing the gains from trade to be negative adds significant richness to the model. Screening the agents to determine whether it is efficient to trade is a non-trivial task, and whether it is even possible to do so depends on the type space, which we treat as an endogenous object.

⁴If either of these conditions fails, there are trivial solutions to the maxmin auction design problem. If the value is known then the Seller can extract all of the surplus with a first-price auction, and if the expected gains from trade are negative, then the optimal profit guarantee is zero, which is attained by any mechanism with non-negative transfers and with the bidders having no information about the value.

Y . An equivalent characterization, due to Blackwell and Girshick (1954) and Rothschild and Stiglitz (1970), is as follows. Let F_X denote the cumulative distribution function for X and F_Y that of Y . Then X is a mean-preserving spread of Y if and only if for all $x \in \mathbb{R}$,

$$\int_{z=-\infty}^x (F_X(z) - F_Y(z))dz \geq 0,$$

and the left-hand side is exactly equal to zero when $x = \infty$.

A *type space* \mathcal{T} consists of (i) a measurable set S_i of signals for each bidder i , (ii) a joint distribution $\pi \in \Delta(S)$ where $S = \times_{i=1}^N S_i$, and (iii) an interim value function $w : S \rightarrow \mathbb{R}$ such that v is a mean-preserving spread of $w(s)$. For a profile of signals s , $w(s)$ is the interim expectation of v conditional on s .⁵

2.3 Mechanisms

A *mechanism* \mathcal{M} consists of measurable sets of messages M_i , one for each bidder i , and measurable mappings

$$\begin{aligned} q_i &: M \rightarrow [0, 1] \\ t_i &: M \rightarrow \mathbb{R} \end{aligned}$$

for each i , where $M = \times_{i=1}^N M_i$ is the set of message profiles. A mechanism is *feasible* if for all m ,

$$\sum_i q_i(m) \leq 1.$$

All mechanisms are required to have an *opt-out message* $0 \in M_i$ such that

$$t_i(0, m_{-i}) = 0.$$

By sending this message, bidder i ensures that she will not make a transfer under any circumstances (although she may still be allocated the good, which we recall can be freely disposed).

2.4 The must-sell model

The model just described allows the Seller to keep the good if he wishes, for which reason we refer this as the *can-keep model*. We will also study a *must-sell model* which restricts the Seller to using mechanisms for which

$$\sum_{i=1}^N q_i(m) = 1$$

⁵This definition of the type space is equivalent to the perhaps more standard definition in which we specify the joint distribution of the signals and the value. The interim expectation of the value is a key object in our analysis, which is why we treat it as the primitive object.

for all $m \in M$. Importantly, this may result in the good being allocated even when it is ex post inefficient to do so, because $v < c$.

The must-sell model is useful for illustrating proof techniques. The must-sell constraint may also arise for economic reasons. For example, it may be that the Seller can commit to a mechanism for a day, but cannot credibly commit to not trying to sell the good in the future if it remains unsold. A Coasean logic suggests that if this competition between the Seller and his future self were very intense, then the only subgame perfect behavior may be to sell the good immediately. Such an argument has been made formally in the context of private-value auctions by Vartiainen (2013). Alternatively, it could be that participation is endogenous, and committing to sell is optimal because it induces bidders to enter the auction, as in McAfee and McMillan (1987) and McAfee (1993).

2.5 Equilibrium

A mechanism \mathcal{M} and a type space \mathcal{T} comprise a game of incomplete information. A (*behavioral*) *strategy* for bidder i is a transition kernel

$$\beta_i : S_i \rightarrow \Delta(M_i).$$

A profile of strategies $\beta = (\beta_1, \dots, \beta_N)$ is naturally identified with a transition kernel that associates to each $s \in S$ a distribution in $\Delta(M)$, where $\beta(s)$ is the product distribution $\beta_1(s_1) \times \dots \times \beta_N(s_N)$.

Given a strategy profile β , bidder i 's payoff is

$$U_i(\beta, \mathcal{M}, \mathcal{T}) = \int_{s \in S, m \in M} (w(s)q_i(m) - t_i(m))\beta(dm|s)\pi(ds).$$

A strategy profile β is a (*Bayes Nash*) *equilibrium* if for all i and strategies β'_i ,

$$U_i(\beta, \mathcal{M}, \mathcal{T}) \geq U_i(\beta'_i, \beta_{-i}, \mathcal{M}, \mathcal{T}).$$

Expected profit is

$$\Pi(\beta, \mathcal{M}, \mathcal{T}) = \int_{s \in S, m \in M} \sum_{i=1}^N (t_i(m) - cq_i(m))\beta(dm|s)\pi(ds).$$

2.6 The Seller's Problem

Informally, the Seller's Problem is to identify mechanisms that provide the best possible profit guarantee across all type spaces and equilibria. We can think of this problem as being a zero-sum game between the Seller, who chooses the mechanism to maximize equilibrium profit, and an adversarial Nature, who chooses the type space to minimize equilibrium profit.

While this formulation is natural, it raises conceptual issues: if there is more than one equilibrium for a particular type space and mechanism, which one will be played? What

is the profit guarantee if there are no equilibria (which may very well occur since we allow infinite mechanisms and type spaces)? Rather than taking a stand on these issues, we will instead impose a relatively strong solution concept. A *solution* of the Seller's Problem consists of a triple $(\mathcal{M}, \mathcal{T}, \beta)$ of a mechanism, type space, and strategy profile, with profit

$$\Pi = \Pi(\beta, \mathcal{M}, \mathcal{T}),$$

such that the following are satisfied:

1. For any type space \mathcal{T}' and any equilibrium β' of $(\mathcal{M}, \mathcal{T}')$,

$$\Pi \leq \Pi(\beta', \mathcal{M}, \mathcal{T}');$$

2. For any mechanism \mathcal{M}' and any equilibrium β' of $(\mathcal{M}', \mathcal{T})$,

$$\Pi \geq \Pi(\beta', \mathcal{M}', \mathcal{T});$$

3. β is an equilibrium of $(\mathcal{M}, \mathcal{T})$.

We refer to Π as the *value* of the solution.

Condition 1 says that the Seller is guaranteed at least Π in any equilibrium under any type space by choosing \mathcal{M} . Condition 2 says that the Seller cannot earn more than Π in any equilibrium of any mechanism if Nature chooses \mathcal{T} . Thus, neither the Seller nor Nature can improve their payoff by deviating, even if the deviator can select the equilibrium. Finally, Condition 3 says that these guarantees are not vacuous: there is some equilibrium in which Π is attained. In fact, the definition implies that for a solution $(\mathcal{M}, \mathcal{T}, \beta)$, all equilibria of $(\mathcal{M}, \mathcal{T})$ must generate profit equal to Π .

3 A roadmap to the solution

We will construct a general solution to the Seller's Problem in Section 4. This section gives an informal derivation and discussion of our solution. To be clear, the purpose of this section is to provide intuition, and formal arguments are in Section 4.

3.1 The structure of the solution

We will construct a solution $(\overline{\mathcal{M}}, \overline{\mathcal{T}}, \overline{\beta})$ to the Seller's Problem. Throughout, a bar accent will denote objects that are part of this solution. This solution has the following structure, which is illustrated in Figure 2. Both the signals for the type space and the messages for the mechanism lie in the same space, which is $M_i = S_i = \mathbb{R}_+$. Thus, there is a common language used for signals and messages. In addition, the equilibrium strategies specify that each bidder send a message that is equal to their signal: for all s_i ,

$$\overline{\beta}_i(s_i) = m_i,$$

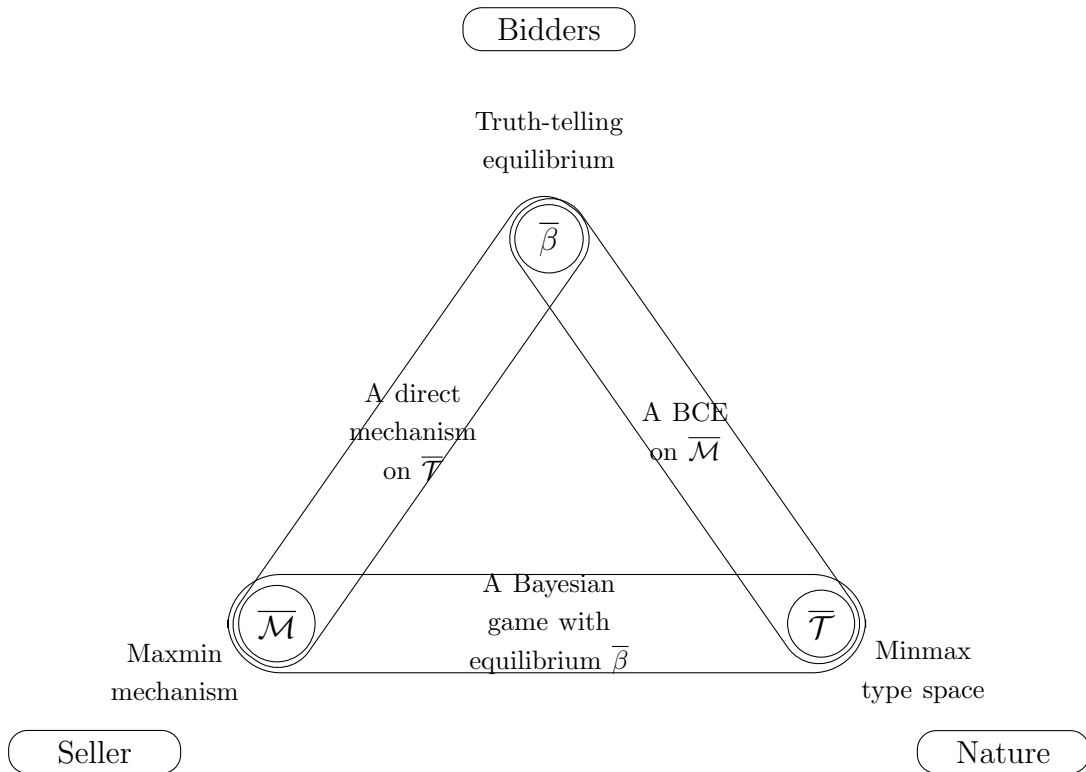


Figure 2: The structure of the solution.

i.e., bidders “report truthfully,” or, equivalently, bidders “obey their signals.”

The truthful equilibrium connects our work to two distinct methodologies in game theory: *direct mechanisms* and *correlated equilibria*. The revelation principle from mechanism design says that any outcome that can be induced with some mechanism and some equilibrium can also be induced with a *direct mechanism* and a *truthful equilibrium*, in which agents simply report their true types to the mechanism (Myerson, 1981). Similarly, epistemic results on Bayesian games say that for a fixed game form, any outcome that can arise under some type space and some equilibrium can also arise with a *Bayes correlated equilibrium* (BCE), in which agents receive signals that are actions in the game, and an *obedient equilibrium*, in which each agent plays an action equal to his signal (Bergemann and Morris, 2013, 2016). This is essentially a revelation principle for games.

Both of these revelation principles are at work in our solution: The maxmin mechanism is essentially a direct mechanism on the minmax type space, and the minmax type space is a BCE on the maxmin mechanism. This is the *double revelation principle* to which we referred in the introduction. Note that the existence of a solution of this form is not implied by the normalizations described in the preceding paragraph. In particular, for a fixed type space, any mechanism \mathcal{M} and equilibrium β has a corresponding direct mechanism \mathcal{M}' in which truth-telling is an equilibrium. But the mechanisms \mathcal{M} and \mathcal{M}' need not have the same set of equilibria, and our solution concept imposes strong

conditions on how profit varies across *all* equilibria. Similarly, replacing a given type space and equilibrium with the corresponding BCE may lead to a different set of equilibria.

3.2 The minmax type space

Having described the general structure and the equilibrium strategies, we turn our attention to the type space $\overline{\mathcal{T}}$. We will eventually show that for the $\overline{\mathcal{T}}$ we construct, the pair $(\overline{\mathcal{M}}, \overline{\beta})$ of the maxmin mechanism and truth-telling strategies maximize expected profit across all mechanisms and equilibria. To develop intuition, it is helpful to work backwards from the proof technique to understand why the type space we construct is the worst case from the Seller's perspective.

To characterize maximum profit on $\overline{\mathcal{T}}$, we will employ a version of the celebrated *revenue-equivalence formula* of Myerson (1981). The version that we use generalizes both Myerson's original result and the interdependent-value formulation derived by Bulow and Klemperer (1996). The latter applies to environments with independent and one-dimensional signals and values that are differentiable and increasing in the signals. The formula says that expected profit is the expectation of the *virtual value* of the bidder who receives the good, where the virtual value of bidder i when the signal profile is s is⁶

$$\psi_i(s) = w(s) - c - \frac{1 - F_i(s_i)}{f_i(s_i)} \frac{\partial w(s)}{\partial s_i}, \quad (1)$$

and where F_i is the absolutely continuous distribution of bidder i 's signal, with density f_i . Thus, the virtual value is equal to the gains from trade, minus the inverse hazard rate times the sensitivity of the value to bidder i 's signal. Note that this is a statement about what profit must be as a function of the allocation, not a characterization of which allocations can be implemented in equilibrium.

Our solution will turn out to have all of the structure that is required in order to apply the revenue-equivalence theorem. The signal space is \mathbb{R}_+ , the set of non-negative real numbers. The interim expectation of the common value is continuous and increasing in the signals. And the signals are independent draws from the same distribution, which we normalize to be exponential with a unitary arrival rate:

$$F_i(x) = 1 - \exp(-x).$$

Thus, the inverse hazard rate is constant and equal to one, and drops out of the virtual value formula.

Drawing on intuition from zero-sum games, we might suspect that the minmax type space $\overline{\mathcal{T}}$ creates lots of indifference on the part of the Seller as to which mechanism to choose. Such indifference would mean that $\overline{\mathcal{T}}$ is hard to best respond to, in the sense that while lots of mechanisms perform reasonably well, no mechanism stands out as exceptional. Indeed, at the solution we propose, there is tremendous indifference on the

⁶In the classic formulation of Myerson (1981), bidder i 's virtual value is their value minus the inverse hazard rate. We would obtain an analogous result if the bidders had different values, and $v_i(s) = s_i$, so that the partial derivative is identically one.

part of the Seller as to which allocation to implement. The minmax value function is always a continuous and increasing function of the *sum of the signals*, i.e.,⁷

$$w(s) = w(\Sigma s),$$

where

$$\Sigma s = s_1 + \dots + s_N.$$

(We will maintain this notational convention for the sum of a vector's elements throughout the paper.) As a result, the interim expected value is equally sensitive to all signals, all bidders have the same virtual value of $w(\Sigma s) - c - w'(\Sigma s)$, and the Seller is *always indifferent between allocating the good to different bidders*.

An important special case is the must-sell model described in Section 2.4 in which the Seller has to sell the good.⁸ In this case, the minmax value function is such that the distribution of $w(s)$ is equal to H , the prior distribution of the value. Let us assume for the remainder of this section that H is absolutely continuous (an assumption that we do not require for our general results). If we let G_N denote the distribution of the sum of the signals,⁹ with density g_N , then the *fully-revealing value function* \hat{w} is implicitly defined by

$$H(\hat{w}(x)) = G_N(x),$$

so that the percentile of the sum of the signals is equal to the percentile of the value.

The reason this is the worst case for the must-sell model is that when the good must be sold, minimizing profit is equivalent to maximizing bidder surplus. The bidders' information rents are in a broad sense proportional to the amount of private information they have about the value, so the more the signals reveal about the value, the greater is the bidders' share of surplus. Bidders' rents are therefore maximized when the signals perfectly reveal the value, as with \hat{w} . For example, when the distribution of the value is standard uniform, then $\hat{w}(x) = G_N(x)$. In Figure 3, we have plotted \hat{w} and the associated virtual value function for the uniform distribution in blue on the left- and right-hand panels, respectively.

Optimal profit on this type space can be easily computed as follows. Because the good must be sold and because the bidders all have the same virtual value, all mechanisms that allocate the good with probability one are optimal, as long as the lowest type of bidder gets zero surplus. So one optimal mechanism (although not a maxmin mechanism!) is to set a posted price equal to the expectation of the value conditional on having a signal of zero. The resulting profit would be

$$\hat{\Pi} = \int_{x=0}^{\infty} \hat{w}(x)g_{N-1}(x)dx - c, \tag{2}$$

⁷We hope we will not create confusion by using the same notation for the interim value as a function of the signal profile and for the interim value as a function of the sum of the signals.

⁸A formal characterization of the must-sell model is given in Theorem 2 below.

⁹Recall that the signals are i.i.d. draws from the standard exponential distribution. Both G_N and g_N have closed-form expressions, given as equations (10) and (11) below.




Figure 3: $N = 2$ and $v \sim U[0, 1]$. The minmax value functions in the must-sell and can-keep cases coincide above x^* . The fully-revealing value function (in blue) is equal to G_2 .

i.e., the expectation of the value as a function of the sum of $N - 1$ independent draws from the standard exponential distribution, minus the cost of production. All types would want to buy at this price, so that maximum profit is achieved.

For some distributions, \hat{w} will also be the minmax value function, even if the good does not have to be sold. In general, though, more work needs to be done. In the uniform example, $\hat{\Pi}$ is not optimal profit if the Seller can withhold the good. The reason is that the virtual value is strictly negative when the aggregate signal is low so that the Seller would be better off by withholding the good at such signal profiles. This can be seen in the right-hand panel of Figure 3. One way to make the Seller worse off would be to create even more indifference, between selling and not selling. For this to be the case, the virtual value must be exactly zero:

$$w(x) - c - w'(x) = 0,$$

i.e., the interim expected gains from trade, notated as $\gamma(x) = w(x) - c$, is of the form $k \exp(x)$ for some positive constant k . We refer to this $\gamma(x)$ as the *gains function*. This shape can be achieved by adding noise to the signal, so that we effectively pool realizations where the virtual value has different signs.

The gains function ends up being a key object for constructing the minmax type space. Note that the virtual value can be formulated in terms of the gains function as

$$\psi_i(s) = \gamma(s) - \frac{1 - F_i(s_i)}{f_i(s_i)} \frac{\partial \gamma(s)}{\partial s_i}, \quad (3)$$

Thus, we can equivalently conceptualize the bidders as getting private information about the gains from trade, from which they can back out the interim expected value $w(s) = \gamma(s) + c$.

In the uniform example, we can replace the fully-revealing gains function $\hat{\gamma}(x) = \hat{w}(x) - c$ when the value is low with an exponential shaped segment on an interval $[0, x^*]$.

This critical x^* is carefully chosen so that the mean-preserving spread constraints will be satisfied, and so that the exponential shape connects continuously with the fully-revealing gains function. We denote this new gains function by $\bar{\gamma}$, with an associated value function \bar{w} . In fact, the \bar{w} depicted in red in Figure 3 is the minmax gains function when the Seller can keep the good and when the cost is zero. The resulting virtual value is everywhere non-negative, so that any mechanism that always allocates the good maximizes profit. Hence, we can still use profit from a posted price mechanism to compute optimal profit on this type space:

$$\bar{\Pi} = \int_{x=0}^{\infty} \bar{\gamma}(x) g_{N-1}(x) dx.$$

More generally, the sign of the fully-revealing virtual value might switch back and forth (although it cannot be negative everywhere if the value has a finite expectation). As we shall see, it is always possible to transform the fully-revealing gains function $\hat{\gamma}$ into a modified gains function $\bar{\gamma}$ so that whenever the fully-revealing virtual value would have been negative, it will be zero under $\bar{\gamma}$, i.e., $\bar{\gamma}$ will grow exponentially on that region. We refer to this as *grading the gains function*, meaning we decrease the derivative of the gains function so that it does not grow faster than exponential. This may entail an exponential segment that is spliced in at the bottom, as in the uniform example, or there may be exponential “bridges” that start and end at strictly positive sums. Our general construction below identifies the correct places to put the exponential segments, so that the resulting gains function is continuous and increasing, satisfies the mean-preserving spread constraints, and results in a virtual value that is everywhere non-negative. Proposition 1 characterizes profit for this construction and gives a generalized upper bound $\bar{\Pi}$ for the optimal profit guarantee.

The grading procedure bears a high-level resemblance to the “ironing” of Myerson (1981) and also to the concavification methodologies that have been employed in the literature on Bayesian persuasion (Kamenica and Gentzkow, 2011). We will elaborate on these connections in Section 4.2.

3.3 Sufficient conditions for an optimal profit guarantee

The remaining piece of the solution is the maxmin mechanism. Before constructing this mechanism explicitly, we will first explain a sufficient condition for the mechanism to guarantee at least $\bar{\Pi}$ in all mechanisms and type spaces, which will then inspire the particular mechanism we construct. This condition is motivated by the twin conjectures that

- (a) profit is at least $\bar{\Pi}$ in all equilibria on all type spaces, and
- (b) truth-telling is an equilibrium under $\bar{\mathcal{T}}$.

Condition (a) just says that the mechanism is maxmin, and (b) conjectures that we can normalize the mechanism so that messages correspond to reported types at the minmax

type space. To these conjectures, we will add the hypothesis that the mechanism is sufficiently “smooth,” in the following sense. Recall that the message space of the mechanism is $M_i = \mathbb{R}_+$, and fix a putative maxmin allocation rule and a maxmin transfer rule. We will assume that the allocation and transfer rules are continuous and that for almost all message profiles $m \in \mathbb{R}_+^N$, $\partial q_i(m)/\partial m_i$ and $\partial t_i(m)/\partial m_i$ exist.

Let us define an *outcome* of this mechanism to be a joint distribution over message profiles and values $\sigma \in \Delta(V \times M)$. Every type space and equilibrium of this mechanism induces an outcome σ , where likelihood of m and v is simply the likelihood that v is the value drawn from H and m is the message profile induced by the bidders’ realized signals and strategies. As we said before, the set of outcomes that can be induced by some type space and equilibrium is equivalent to the set of *Bayes correlated equilibria* (BCE), where the BCE have to satisfy a family of linear *obedience constraints*, which is that for all i and m_i , m_i is a best response to the conditional distribution of (v, m_{-i}) induced by σ and conditioning on m_i , and also satisfy the *marginal constraints*, that the marginal distribution on V is the prior H .

Let us denote by $\bar{\sigma}$ the outcome induced by $(\bar{\mathcal{T}}, \bar{\beta})$. Our conditions (a) and (b) above are equivalent to saying that $\bar{\sigma}$ is the profit-minimizing BCE of the mechanism and that profit in this BCE is $\bar{\Pi}$. It will turn out that the only obedience constraints that are relevant for our problem are those associated with *local* obedience constraints, i.e., that for all i and m_i ,

$$\int_{(v, m_{-i}) \in V \times M_{-i}} \left(v \frac{\partial q_i(m_i, m_{-i})}{\partial m_i} - \frac{\partial t_i(m_i, m_{-i})}{\partial m_i} \right) \sigma(dv, dm_{-i} | m_i) = 0.$$

The BCE $\bar{\sigma}$ is therefore the solution to an infinite dimensional linear program, for which the associated Lagrangian is¹⁰

$$\begin{aligned} \mathcal{L}(\sigma, \{\alpha_i\}, \lambda) = & \sum_{i=1}^N \int_{(v, m) \in V \times M} (t_i(m) - cq_i(m)) \sigma(dv, dm) \\ & + \sum_{i=1}^N \int_{(v, m) \in V \times M} \alpha_i(m_i) \left(v \frac{\partial q_i(m)}{\partial m_i} - \frac{\partial t_i(m)}{\partial m_i} \right) \sigma(dv, dm) \quad (4) \\ & + \int_{(v, m) \in V \times M} \lambda(v) (H(dv) - \sigma(dv, dm)). \end{aligned}$$

This Lagrangian has three terms: profit induced by the BCE, the sum of local obedience constraints times their corresponding multipliers (the functions $\alpha_i(m_i)$), and the sum of marginal constraints times their corresponding multipliers (the function $\lambda(v)$). A necessary first-order condition for $\bar{\sigma}$ to be the profit-minimizing BCE is that for all (v, m) ,

$$\sum_{i=1}^N \left[t_i(m) - cq_i(m) + \alpha_i(m_i) \left(v \frac{\partial q_i(m)}{\partial m_i} - \frac{\partial t_i(m)}{\partial m_i} \right) \right] - \lambda(v) \geq 0, \quad (5)$$

¹⁰Since we are just providing informal motivation, we will not rigorously argue that this Lagrangian is equivalent to the profit-minimization program. The fact that this is the correct Lagrangian is a consequence of Theorem 1 below.

with the constraint holding as an equality for all $(v, m) \in \text{supp } \bar{\sigma}$. In other words, for the optimal multipliers (which maximize the minimum value of the Lagrangian across all σ), making any pair (v, m) more likely must weakly increase the Lagrangian, and making pairs in the support of the profit-minimizing BCE more likely must have no effect (otherwise we could decrease profit by making them less likely).

This equation contains a tremendous amount of information about what a maxmin mechanism must look like. It imposes a tight relationship between the mechanism and the Lagrange multipliers on the profit-minimization program. Moreover, there are very natural guesses for what those multipliers should be. By the envelope theorem, the Lagrange multiplier $\lambda(v)$ must be the derivative of minimum profit in the maxmin mechanism with respect to the prior probability of v , which we denote by $\bar{\lambda}(v)$.¹¹ Let us now make explicit the dependence of the upper bound on the profit guarantee $\bar{\Pi}(H)$ on the distribution of v . Under the retained hypothesis that $\bar{\Pi}(H)$ is the optimal profit guarantee, we know exactly what $\bar{\lambda}(v)$ has to be: it is the local linear approximation of $\bar{\Pi}(\cdot)$ around H , i.e., for H' close to H ,

$$\bar{\Pi}(H') \approx \bar{\Pi}(H) + \int_{v \in V} \bar{\lambda}(v)(H'(dv) - H(dv)).$$

Indeed, if $\bar{\lambda}$ were *not* the gradient of $\bar{\Pi}(\cdot)$, then we could find a direction in which minimum profit from the maxmin mechanism for H would be increasing faster than the profit guarantee. We will give an explicit formula for $\bar{\lambda}$ in Section 4. The key insight is that $\bar{\lambda}$ is determined by the function $\bar{\Pi}$ (which was independently motivated by our conjecture for the minmax type space) and is not a feature of any particular maxmin mechanism.

As for the multipliers on local obedience, there is an even simpler answer: the α_i are all constant and equal to one. Why should this be the case? The Lagrangian (4) is very similar to the Lagrangian for the linear program for *maximizing* profit given the *fixed* type space $\bar{\mathcal{T}}$, where we would fix $\sigma = \bar{\sigma}$ and treat the direct mechanism (q_i, t_i) as the choice variables. The local obedience constraints would then be reinterpreted as local truth-telling constraints. As is well known, the optimal auction is associated with all of the local downward incentive constraints binding, and in fact the multiplier on local incentive compatibility at a type m_i is equal to the inverse hazard rate of the distribution. This hazard rate has been normalized to be equal to one by the assumption that signals are standard exponentials.¹²

Now, the dual constraints (5) bind on the support of $\bar{\sigma}$, which consists of any pair (v, m) such that (i) $v = \hat{w}(\Sigma m)$ when the gains from trade are fully revealed, or (ii) $(v, \Sigma m) \in [\hat{w}(a), \hat{w}(b)] \times [a, b]$, where the gains function has been graded on the interval

¹¹We could obtain $\bar{\lambda}(v)$ directly by taking the directional derivative of $\bar{\Pi}(H)$ in the direction of the Dirac measure on v , and indeed, we did this ourselves a long time ago. But this is now unnecessary: we will give an explicit formula in Section 4, and the fact that this is the correct multiplier is a consequence of Theorem 1.

¹²We have been unable to find a published work that makes the specific point that the multiplier on local incentive compatibility in the standard auction design problem is equal to the inverse hazard rate. The dual of the mechanism design problem has been studied in numerous places, including lecture notes of Roger Myerson, Myerson (1983), and Vohra (2011).

$[a, b]$. Let us first consider a message m where the gains function is not graded at Σm . Let

$$Q(m) = \sum_{i=1}^N q_i(m)$$

denote the aggregate supply. We can rearrange the terms in the first-order condition to

$$\sum_{i=1}^N \left(\frac{\partial t_i(m)}{\partial m_i} - t_i(m) \right) \leq v \sum_{i=1}^N \frac{\partial q_i(m)}{\partial m_i} - \bar{\lambda}(v) - cQ(m) \quad (6)$$

for all $v \in V$, and the equation must hold as an equality when $v = \hat{w}(\Sigma m)$. This means that the right-hand side is minimized at this value. As we shall see, the $\bar{\lambda}$ we construct¹³ is almost everywhere differentiable, so the following first-order condition must be satisfied:

$$\sum_{i=1}^N \frac{\partial q_i(m)}{\partial m_i} = \bar{\lambda}'(\hat{w}(\Sigma m)). \quad (\text{TAS})$$

We refer to the left-hand side as the *total allocation sensitivity*, which we denote by μ .

In economic terms, this equation says the following. With the Lagrangian approach, we have essentially dropped the constraints in Nature's problem and replaced them with costs for the local gains from deviating for the bidders and costs for the realized values. Now, consider perturbing the minmax solution by increasing the value by ϵ when the signal profile is m . If $\bar{\sigma}$ truly minimizes profit, then this perturbation should not change the total cost. There are two effects: the first is to change the gains from local deviations, in proportion to the total allocation sensitivity at m . The second is to "save" the cost of $\hat{w}(\Sigma m)$, which is $\bar{\lambda}(\hat{w}(\Sigma m))$, but also to incur a cost of $\bar{\lambda}(\hat{w}(\Sigma m) + \epsilon)$ from making $\hat{w}(\Sigma m) + \epsilon$ more likely. (TAS) says these two effects must exactly balance one another.

In fact, by plugging in the explicit formula for $\bar{\lambda}$, we can compute the exact optimal total allocation sensitivity. Whenever the value function is fully-revealing at the aggregate signal $x = \Sigma m$, the total allocation sensitivity of a maxmin mechanism must be

$$\bar{\mu}(x) = \frac{N-1}{x}. \quad (7)$$

On regions where the value is graded, $\bar{\lambda} \circ \hat{w}$ is linear in the aggregate message, so that $\bar{\mu}$ is constant. Thus, the right-hand side of (6) will be minimized at all value and message profile pairs in the graded region. We will give the explicit formula in Section 4. The main takeaway is that if the maxmin auction is smooth and satisfies (a) and (b), then the total allocation sensitivity is completely pinned down. Note, however, that there are additional conditions that the allocation rule must satisfy in order for truth-telling to maximize profit

¹³ $\bar{\lambda}$ is uniquely defined only for values in the support of H . We construct a particular $\bar{\lambda}$ that is continuous and differentiable for all v . As we shall see, this has desirable normative properties, but it does restrict the class of mechanisms we consider. See the discussion in Section 4.6.

on $\overline{\mathcal{T}}$. In particular, it must be that the good is always allocated at aggregate messages for which the virtual value is positive, which is when the value function is fully revealing.

What does (6) tell us about the transfers? We can view the binding constraints as a partial differential equation that relates the transfers to the total allocation sensitivity. In particular, let

$$\xi_i(m) = \frac{\partial t_i(m)}{\partial m_i} - t_i(m)$$

denote the *excess growth* of the transfer in bidder i 's message relative to exponential growth. It is easy to show that given an excess growth function $\xi_i(m)$, the implied transfer function is

$$t_i(m) = \exp(m_i) \int_{x=0}^{m_i} \exp(-x) \xi_i(x, m_{-i}) dx. \quad (8)$$

Thus, the binding first-order condition (6) reduces to

$$\sum_{i=1}^N \xi_i(m) = \widehat{w}(\Sigma m) \bar{\mu}(\Sigma m) - \bar{\lambda}(\widehat{w}(\Sigma m)) - cQ(m). \quad (\text{PIA})$$

We refer to this as the *profit-incentive alignment* (PIA) equation. We let $\Xi(m)$ denote the right-hand side of (PIA), which is the *total excess growth*. In words, equation (PIA) says that the total excess growth of the transfers must equal $\Xi(m)$, and the total excess growth function itself must be given by the right-hand side of (PIA), where the only remaining endogenous object is the total supply function $Q(m)$. Moreover, all of the terms except possibly the total supply are only a function of the aggregate message. At a higher level, (PIA) imposes a linkage between ex post profit, $\Sigma(t - cq)$, and the sum of the bidders' local incentive constraints, $v(\nabla \cdot t) - (\nabla \cdot q)$. This ensures that as long as local incentive constraints are satisfied, profit cannot fall too low.

We have been using (5) to derive conditions that a maxmin auction must satisfy. But as Proposition 2 shows, these conditions are actually *sufficient* for a mechanism to guarantee profit of at least $\overline{\Pi}(H)$. Specifically, if we have a smooth mechanism with the features that (i) the total allocation sensitivity coincides with the critical value $\bar{\mu}$ given by the right-hand side of (TAS) and (ii) the total excess growth of the transfers is at least Ξ , then profit must be at least $\overline{\Pi}(H)$ in all type spaces and all equilibria. The proof is essentially an application of the weak duality theorem of linear programming. Importantly, the total allocation sensitivity and total excess growth are determined from $\overline{\Pi}(H)$ and its gradient $\bar{\lambda}$ and do not depend directly on other assumptions about the maxmin mechanism other than the facts that it is smooth and that messages are normalized so that truth-telling is an equilibrium on the type space $\overline{\mathcal{T}}$.

3.4 Construction of a maxmin mechanism

The last step is to explicitly construct a mechanism that completes the solution. The strategy is to make sure this mechanism satisfies the sufficient conditions of Proposition

2 and also that truth-telling is an equilibrium at $\bar{\mathcal{T}}$. Note that the Proposition 2 does not imply that truth-telling is an equilibrium, and this is something we will have to separately verify for the particular mechanism we construct.

Let us first construct the allocation rule that has the required total allocation sensitivity. The form for the allocation rule can be motivated by considering the case of two bidders and zero cost. Fix an aggregate-message at which the value function is fully revealing, so that profit maximization demands that the good is always allocated, and hence $q_2(m_1, m_2) = 1 - q_1(m_1, m_2)$. Thus, the total allocation sensitivity is

$$\frac{\partial q_1(m_1, m_2)}{\partial m_1} - \frac{\partial q_1(m_1, m_2)}{\partial m_2} = \frac{1}{m_1 + m_2}.$$

Now consider a level curve of the aggregate message where $m_1 + m_2 = x$. Then we can view the left-hand side as the total derivative with respect to m_1 along the parametric curve $m_2(m_1) = C(x) - m_1$, so that integrating both sides, we obtain

$$q_1(m_1, m_2) = \frac{m_1}{x} + C(x).$$

But in order to have $q_1 \in [0, 1]$, the constant that depends on the aggregate message must be zero. Thus, the implied allocation probability is simply the bidder's share of the aggregate message.

More generally, we use an allocation that has the following functional form, which we refer to as the *proportional allocation rule*: when $\Sigma m > 0$, we set

$$q_i(m) = \frac{m_i}{\Sigma m} Q(\Sigma m). \quad (9)$$

When $\Sigma m = 0$, allocations are $q_i(0) = Q(0)/N$. Thus, total supply only depends on the aggregate message, and conditional on the total supply, allocations are proportional to messages.

Given this functional form for the allocation, the optimal total supply function \bar{Q} is completely pinned down from the known total allocation sensitivity $\bar{\mu}$ via the following differential equation:

$$\bar{\mu}(x) = \frac{N-1}{x} \bar{Q}(x) + \bar{Q}'(x) = \bar{\lambda}'(\hat{w}(x)).$$

In particular, $\bar{Q}(x) = 1$, so that the probability that bidder i gets the good is simply their share of the aggregate message. This is also consistent with profit maximization, which requires the good to be allocated with probability one on such regions. When the value function is graded, the good must be rationed to keep the total allocation sensitivity constant and equal to $\bar{\lambda}'$. The optimal supply function \bar{Q} also pins down the allocation rule and the total excess growth, $\bar{\Xi}$, which is also a function of the aggregate message.

This leaves the transfers. At first glance, Proposition 2 seems to leave tremendous flexibility in how to specify the transfers. Indeed, equation (6) has to hold as an equality if our conjecture about the profit-minimizing BCE is correct. But if it held as an inequality,

our dual argument would still imply that $\bar{\Pi}$ is a lower bound on profit for all type spaces and equilibria. In fact, we could even do better: what is the harm in picking even higher excess growth functions? That would allow us to accommodate even higher λ 's, which would result in an even higher profit guarantee! The fallacy lurking here is that the mechanisms with higher excess growth would correspond to transfer rules that blow up as messages become large, so that no equilibria exist on *any* type space, let alone $\bar{\mathcal{T}}$.

Intuitively, this is related to well-known issues of non-existence of equilibrium in games with discontinuous payoffs and/or non-compact action spaces. This suggests that we should impose a kind of “transversality” condition that would make the transfers continuous and bounded at infinity, so as to restore equilibrium existence. This brings us to the second critical condition for the maxmin mechanism. Inspecting the formula for the transfers (8), it is clear that a necessary condition for the transfers to not explode as messages become large is that

$$\int_{m_i=0}^{\infty} \xi_i(m_i, m_{-i}) \exp(-m_i) dm_i = 0. \quad (\text{ST})$$

In other words, it must be that each bidder has zero excess growth in expectation across their own report, holding fixed the other bidders’ messages. We refer to this condition as *strong transversality*. If, in addition, the excess growth functions are themselves continuous and bounded, then L’Hôpital’s rule applied to the limit of (8) implies that

$$\lim_{m_i \rightarrow \infty} t_i(m_i, m_{-i}) = \lim_{m_i \rightarrow \infty} -\xi_i(m_i, m_{-i}),$$

so that the transfers will be well-behaved, and we could apply standard equilibrium existence arguments.

Because it implies that the transfers are well-behaved when messages are large, the zero excess growth condition is normatively appealing. As Lemma 8 shows, this condition is also necessary and sufficient for truth-telling to be an equilibrium at the minmax type space, when we use the proportional allocation rule with two bidders. When there are three or more bidders, however, Lemma 8 shows there is a slightly weaker necessary and sufficient condition for transfer rules to be incentive compatible for the proportional allocation rule, which we refer to as *weak transversality*:

$$\int_{m_{-i} \in \mathbb{R}_+^{N-1}} \xi_i(m_i, m_{-i}) \exp(-\Sigma m_{-i}) dm_{-i} = \int_{x=0}^{\infty} \Xi(m_i + x) g_{N-1}(x) dx. \quad (\text{WT})$$

In other words, conditional on their own message, bidder i expects to get all of the excess growth on average across other bidders’ messages. Remarkably, when we use the proportional allocation rule, weak transversality is *equivalent* to local optimality of truth-telling, and it is also sufficient for global incentive compatibility.

The remaining question is whether there *exists* a transfer function that solves the excess growth equations. At first glance, (ST) and (WT) seem to contradict (PIA). If every bidder expects to get no excess growth on average as per (ST), then how can they split the growth between themselves as per (PIA)? The way out of the “paradox,” of

course, is that the ex ante expected total excess growth must be zero. This is indeed the case for the total allocation sensitivity that we derived and the total supply that we constructed. In fact, using the fact that expected total excess growth is zero, we are able to explicitly construct excess growth functions $\bar{\xi}_i$ that satisfy (PIA) and strong transversality (ST). From these, we compute the transfer rules \bar{t}_i , to complete the specification of our solution.

The constructed excess growth functions have a simple form, which can be understood as follows. We arrange the bidders in a random order. The first bidder i_1 receives the total excess growth for the realized message profile. The second bidder i_2 then makes a “transfer” of excess growth to bidder i_1 , so that in expectation across m_{i_1} , bidder i_1 ’s excess growth is zero. Explicitly, the transfer is the negative of the expected total excess growth across m_{i_1} , given m_{-i_1} . We then have the third bidder i_3 make a transfer to bidder i_2 to cancel out bidder i_2 ’s growth transfer to i_1 in expectation across m_{i_2} , and have bidder i_4 ’s transfer cancel out bidder i_3 ’s transfer to i_2 , etc. The k th of these transfers is the expectation of the total excess growth across $(m_{i_1}, \dots, m_{i_k})$, and conditional on $(m_{i_{k+1}}, \dots, m_{i_N})$. Thus, the last transfer is the ex ante expected total excess growth, which is zero.

Theorem 1 shows that this construction completes the specification of a maxmin mechanism, and therefore completes the solution to the Seller’s Problem. The optimal allocation and transfer rules are plotted for the two-bidder/uniform/zero-cost example in Figure 4 for message profiles in $[0, 5]^2$. For comparison, we have also plotted the corresponding objects for the must-sell model. Essentially all of our arguments carry over to the must-sell case with minor modifications. We will explain this in more detail in Section 4.5.2.

4 A general solution to the Seller’s Problem

4.1 Guide to the formal arguments

We now present a general and rigorous solution to the Seller’s Problem. The order roughly corresponds to the informal exposition that we just concluded. For easy reference, Figure 5 summarizes the structure of and logical relationships between our formal results.

4.2 The type space

4.2.1 Construction of $\bar{\mathcal{T}}$

For a general distribution of values H , $\bar{\mathcal{T}}$ is defined as follows. The N bidders have signal spaces $S_i = \mathbb{R}_+$, with the standard measurable structure. The distribution of the signals will be

$$\bar{\pi}(ds) = \exp(-\Sigma s) ds.$$

In other words, the signals are independent draws from the exponential distribution with arrival rate 1.

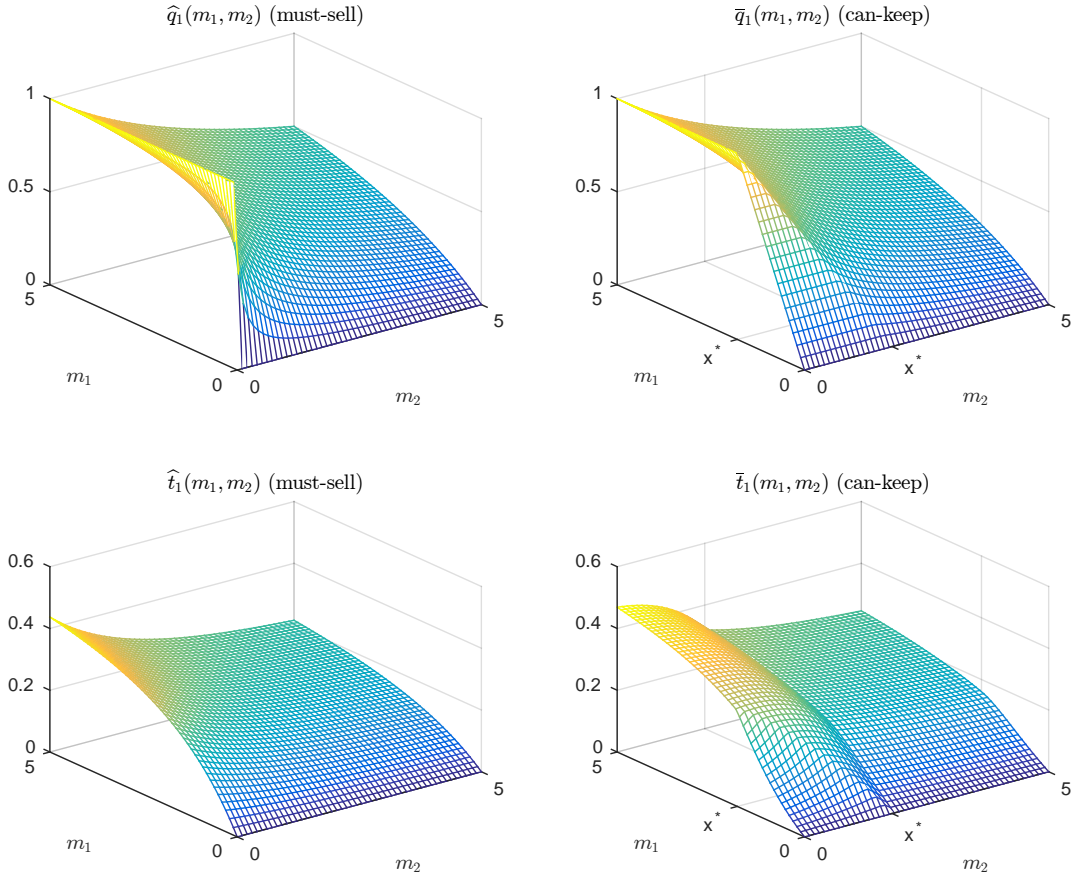


Figure 4: The maxmin mechanism and transfer rule in the uniform case with $N = 2$.

Let us denote by $G_K(x)$ the distribution of the sum of K i.i.d. draws from an exponential distribution, with density $g_K(x)$. This is also known as an Erlang distribution. It is straightforward to show that

$$g_K(x) = \frac{x^{K-1}}{(K-1)!} \exp(-x) \quad (10)$$

and that

$$G_K(x) = 1 - \sum_{k=1}^K g_k(x). \quad (11)$$

The interim expected value will also be a function of the sum of the signals: $w(s) \equiv \bar{w}(\Sigma s)$. This \bar{w} will be an increasing function, and is defined according to the following *grading procedure*. First define the *fully-revealing value function* $\hat{w}(x)$ according to

$$\hat{w}(x) = \min \{v \mid H(v) \geq G_N(x)\}.$$

We define $\hat{\gamma}(x) = \hat{w}(x) - c$ to be the *fully-revealing gains function*.

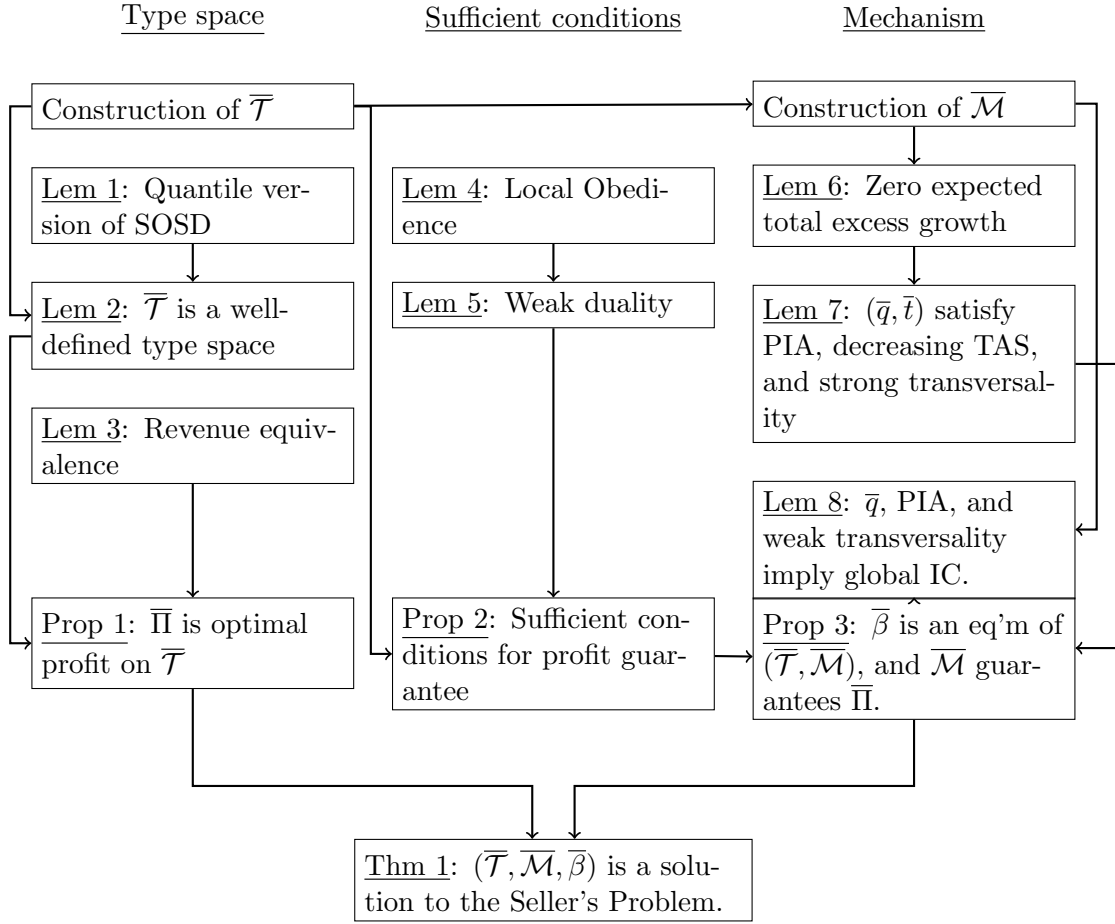


Figure 5: The formal arguments and logical dependencies.

Next, let

$$\hat{\Gamma}(x) = \int_{y=0}^x \hat{\gamma}(y) g_N(y) dy.$$

Also let

$$E(x) = \int_{y=0}^x \exp(y) g_N(y) dx,$$

which is strictly increasing, and hence it has a well-defined inverse E^{-1} . Let $\text{cav}(\hat{\Gamma} \circ E^{-1})$ denote the smallest concave function that is everywhere above $\hat{\Gamma} \circ E^{-1}$. We then set $\bar{\Gamma} = \text{cav}(\hat{\Gamma} \circ E^{-1}) \circ E$, and define

$$\begin{aligned} \bar{\gamma}(x) &= \frac{1}{g_N(x)} \frac{d}{dx} \bar{\Gamma}(x), \\ \bar{w}(x) &= \bar{\gamma}(x) + c \end{aligned}$$

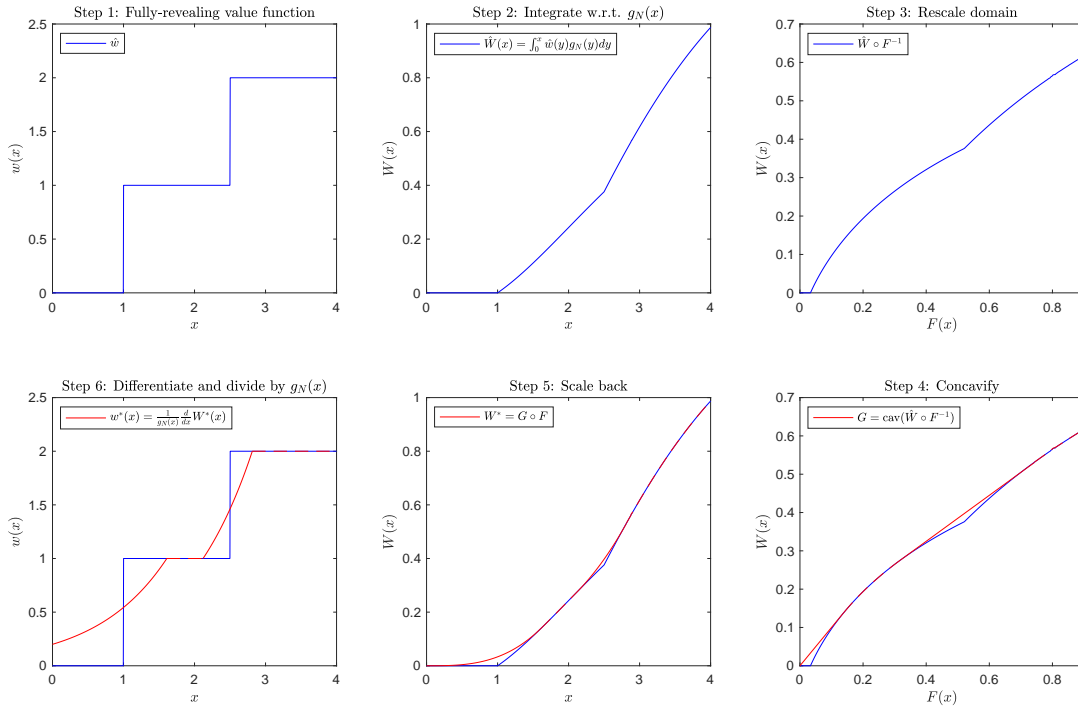


Figure 6: Grading the value function when $N = 3$ and $\hat{\gamma}(x) = \mathbb{I}_{x>1} + \mathbb{I}_{x>2} - 0.5$.

where the derivative is taken from the right. We refer to $\bar{\gamma}$ and \bar{w} as the *graded gains function* and the *graded value function*, respectively.

The grading procedure is illustrated in Figure 6 for an example with three bidders, and where $c = 0.5$ and $\hat{\gamma}(x) = \hat{w}(x) - c = \mathbb{I}_{x>1} + \mathbb{I}_{x>2} - 0.5$. In other words, the support of H is $\{0, 1, 2\}$, with $H(0) = G_3(1)$, $H(1) = G_3(2)$, and $H(2) = 1$. Essentially, we want to even out $\hat{\gamma}$ so that its growth rate is always sub-exponential. This can be reformulated as a constraint on the rate of growth in the derivative of $\bar{\gamma}$, i.e., that we do not want $\bar{\gamma}$ to be “too convex.” The maximum amount of convexity is precisely that of E , which is what would obtain if the value function grew at exactly an exponential rate. By renormalizing the domain so that E grows linearly, then our constraint is that the renormalized $\bar{\Gamma}$ grows slower than linear, i.e., it is concave. At the same time, we need $\bar{\Gamma}$ to coincide with $\hat{\Gamma}$ whenever the growth of the derivative is sub-exponential, and we need $\bar{\Gamma}$ and $\hat{\Gamma}$ to meet smoothly, so that we define $\bar{\Gamma} \circ E^{-1}$ to be the smallest concave function that is above $\hat{\Gamma} \circ E^{-1}$. As we shall see, defining $\bar{\gamma}$ and \bar{w} in this manner ensures that v is a mean-preserving spread of $\bar{w}(x)$, so that it is part of a well-defined type space.

There is a strong resemblance between grading and the “ironing” of Myerson (1981), and also the concavification arguments used in Bayesian persuasion, e.g., Kamenica and Gentzkow (2011). In our view, the relationship is as follows. In optimal auctions with independent private values, the ironing procedure is used to solve the *Seller’s* auction design problem, by transforming a non-monotonic virtual value function into a monotonic “pseudo”-virtual value, for which the pointwise optimal allocation will respect global

incentive constraints. In contrast, our procedure is not about the Seller’s design problem, but rather it concerns the amount of *information* that the bidders should have about the value when they combine their various signals. Nature designs that information so that the Seller is willing to always allocate the good. This requires that the bidders’ signals should never be too informative about the value, in the sense that the value function should not be increasing too quickly at any given point.

Because it is inherently about how to design the bidders’ information, grading seems to have more in common with the Bayesian persuasion setting, where Nature is the Sender and the bidders are collectively the Receiver. In Bayesian persuasion, the concavity of the Sender’s value function represents the law of iterated expectations, that the expectation of the posterior belief must be equal to the prior. For us, concavity represents a limit to how quickly the interim expectation of the value can grow in the sum of the signals, in order to induce the Seller to sell the good. Our “ironed” value function is, however, consistent with the law of iterated expectations, as represented by the stochastic dominance constraints. Thus, our procedure is related on a technical level but is conceptually distinct from those earlier applications of concavification.

4.2.2 Properties of $\bar{\mathcal{T}}$

This concludes the definition of the minmax type space. We now show that this is in fact a well-defined type space. To show that v is a mean-preserving spread of $\bar{w}(x)$, we will use the following characterization of mean-preserving spreads, which is of some independent interest.

Given a cumulative distribution $F(x)$ (i.e., a non-decreasing and right-continuous function into $[0, 1]$), let us define its *inverse* to be

$$F^{-1}(z) = \min\{x \mid F(x) \geq z\}.$$

This is the *quantile function* corresponding to F . Because F is increasing and continuous from the right, there is a closed set of values that have a higher cumulative probability than z , so this minimum is well-defined. Moreover, $F^{-1}(x)$ is an increasing function and is continuous from the right, and it has discontinuities when there are gaps in the support of F .

Lemma 1 (Mean-preserving spreads). *Fix real-valued random variables X and Y with cumulative distributions F_X and F_Y . Then X is a mean-preserving spread of Y if and only if for all $\alpha \in [0, 1]$,*

$$\int_{y=0}^{\alpha} (F_Y^{-1}(y) - F_X^{-1}(y)) dy \geq 0. \tag{12}$$

The proof is in the Appendix, but the idea is demonstrated visually in Figure 7.¹⁴ The left-hand panel depicts two cumulative distribution functions, where the blue one is a mean-preserving spread of the red one. The center panel depicts the Blackwell definition

¹⁴The proof of this result benefited greatly from a discussion with Piotr Dworzak and Doron Ravid.




Figure 7: Cumulative versus quantile definitions of mean-preserving spreads.

of second-order stochastic dominance: When we integrate the area under the CDFs, the area under the blue CDF must be greater than the area under the red. Equivalently, the area under blue and above red must be non-negative (with the difference being negative when red is above blue). If we integrate up to the given \hat{x} , then there is a blue shaded area where the blue CDF is higher, so that the signed area is positive, and a red shaded area where red is above blue, so that the signed area is negative.

The right-hand panel depicts our definition (12), where we integrate the quantile functions up to a given cutoff probability. Here, we are integrating the area to the left of the red quantile function and to the right of the blue quantile function. It is not hard to see that when the cutoff is a point where the two curves cross, e.g., \underline{x} or \bar{x} , the signed area between the curves is the same whether we integrate in cumulative or in quantile direction. Thus, the signed area between the curves is the same at these crossings, and the two definitions are equivalent at these points.

Moreover, in between the crossing points, the direction of change in the signed area is the same. For example, the signed area is decreasing in \hat{x} or $\hat{\alpha}$ when the cutoff is in $[\underline{x}, \bar{x}]$ or in $[\underline{\alpha}, \bar{\alpha}]$, respectively. Thus, the extremes of the signed area must be attained at the crossing points. If it has the correct sign at these crossing points, then it has the correct sign in between as well, and we are done.

This characterization essentially reframes stochastic dominance in terms of ordering of conditional expectations. Indeed, the integrals in (12) are simply expectations among the lowest α realizations. Thus, X is a mean-preserving spread of Y if for every $\alpha \in [0, 1]$, the average of the α lowest realizations of X is less than the average of the α lowest realizations of Y .

Let us define a *graded interval* to be an interval $[a, b]$ such that $\bar{\Gamma}(x) = \hat{\Gamma}(x)$ for $x \in \{a, b\}$ and $\bar{\Gamma}(x) > \hat{\Gamma}(x)$ for $x \in (a, b)$. We are now ready to state the following:

Lemma 2 (Properties of $\bar{\mathcal{T}}$). *The following properties hold for $\bar{\mathcal{T}}$: (a) $\bar{\mathcal{T}}$ is a type space for $v \sim H$; (b) if x is not in a graded interval, then $\bar{\gamma}(x) = \hat{\gamma}(x)$, and if $[a, b]$ is a graded interval, then for $x \in [a, b]$, $\bar{\gamma}(x) = \bar{\gamma}(a) \exp(x - a)$; (c) $\bar{\gamma}$ is continuous and increasing; and (d) for all x and $y \geq x$,*

$$\bar{\gamma}(y) \leq \bar{\gamma}(x) \exp(y - x). \quad (13)$$

Let us sketch the argument. The functions $\widehat{\Gamma}$ and $\overline{\Gamma}$ are in fact integrated quantile functions for the variables $\widehat{\gamma}(\Sigma s)$ and $\overline{\gamma}(\Sigma s)$, where the former has the distribution of $v - c$, $v \sim H$. Thus, v is a mean-preserving spread of $\overline{w}(\Sigma s)$ iff $\overline{\Gamma}$ is everywhere larger than $\widehat{\Gamma}$, and if they coincide at the extremes. But this follows from the fact that $\overline{\Gamma} \circ E^{-1}$ is everywhere above $\widehat{\Gamma} \circ E^{-1}$, since the former is the concavification of the latter. Lemma 1 then implies the result.

For (b), $\overline{\gamma}$ and $\widehat{\gamma}$ coincide on any interval where $\overline{\Gamma} = \widehat{\Gamma}$, since the derivatives of these functions must then coincide on that interval as well. On the regions where we concavified, $\overline{\Gamma} \circ E^{-1}$ is linear, which implies that $\overline{\gamma}$ grows exponentially, as

$$\begin{aligned} \frac{d}{dx} \overline{\Gamma}(x) &= \frac{d}{dz} (\overline{\Gamma}(E^{-1}(z)))|_{z=E(x)} E'(x) \\ &= C \exp(x) g_N(x). \end{aligned}$$

Part (c) follows partly from the fact that whenever $\overline{\Gamma}$ coincides with $\widehat{\Gamma}$, it must be that $\widehat{\gamma}$ is continuous, since the function $\widehat{\Gamma} \circ E^{-1}$ will be strictly convex around a point of discontinuity. Moreover, $\widehat{\gamma}$ is always increasing, so that $\overline{\gamma}$ is continuous and increasing wherever it coincides with $\widehat{\gamma}$. Moreover, $\overline{\gamma}$ is also continuous and increasing on any concavified interval where it has an exponential shape. At the points where these intervals meet, it follows from the fact that the slope of $\overline{\Gamma} \circ E^{-1}$ must coincide with the slope of $\widehat{\Gamma} \circ E^{-1}$ at the endpoints of the concavified interval, so that $\overline{\Gamma}$ and $\widehat{\Gamma}$ also have the same slope at these points as well. (Note that $\widehat{\Gamma}$ cannot have a concave kink, since this would require $\widehat{\gamma}$ to jump down, but $\widehat{\gamma}$ is monotonically increasing.) Finally, we prove (d) by showing that at any point where $\widehat{\gamma}$ grows superexponentially, $\widehat{\Gamma} \circ E^{-1}$ will be strictly convex, so that we have to concavify around this point to obtain $\overline{\gamma}$.

4.2.3 Upper bound on profit

Next, we say that a type space \mathcal{T} is *standard* if S_i is a closed and connected subset of \mathbb{R}_+ , w is monotonic in s and right-continuous in s_i for all i , and if the s_i are independent draws from absolutely continuous distributions F_i . Write f_i for the density of s_i , and let $f(s)$ denote the density of the profile of signals. Note that $\overline{\mathcal{T}}$ is a standard type space.

Lemma 3 (Revenue equivalence). *Suppose an incentive compatible mechanism implements the direct allocation $q(s)$ on a standard type space \mathcal{T} and that the lowest type of bidder i has interim utility \underline{U}_i . Then expected profit is bounded above by*

$$\sum_i \left[\int_s q_i(s) f(s) \psi_i(ds) - \underline{U}_i \right], \quad (14)$$

where

$$\psi_i(ds) = (w(s) - c) ds - \frac{1 - F_i(s_i)}{f_i(s_i)} \omega_i(ds_i; s_{-i}) ds_{-i}, \quad (15)$$

and where $\omega_i(ds_i; s_{-i})$ is the outer measure on \mathbb{R}_+ induced by

$$\omega_i([a, b]; s_{-i}) = w(b, s_{-i}) - w(a, s_{-i}). \quad (16)$$

The measure $\psi_i(ds)$ represents a generalized virtual value, or marginal profit, from allocating to bidder i . An analogous result can be found in Bulow and Klemperer (1996) for the case where w is differentiable in s_i . We have given this more general formulation so as to encompass the case where w is not differentiable, i.e., when the distribution of $w(s)$ has gaps in the support or is not absolutely continuous. A formal proof appears in the Appendix. The proof overlaps significantly with that of the analogous result in Myerson (1981). Briefly, expected profit in an incentive compatible direct mechanism can always be accounted as the difference between total surplus (a function only of the allocation) and the sum of the bidders' expected surpluses. The latter can be computed by taking the expectation of the interim utility as a function of the type, denoted $U_i(s_i)$. An increase in the type s_i entails two changes: an increase in the true type and an increase in the reported type. In any incentive compatible direct mechanism, truth-telling must be a local optimum, so that at the margin, the change in U_i from a change in the reported type is zero. This envelope-theorem argument implies that the derivative $U'_i(s_i)$ is just the expected sensitivity of the value to s_i , times the equilibrium allocation. By integrating over s and applying the Fubini theorem, one obtains the formula (14). Note that gaps in the support of the value introduce some slack, so that the profit formula (14) does not have to hold exactly in all incentive compatible direct mechanisms, but it is always a valid upper bound.

We will use Lemma 3 to prove the following:

Proposition 1 (Upper bound on maxmin profit). *Maximum profit across all mechanisms and equilibria when the type space is $\overline{\mathcal{T}}$ is*

$$\overline{\Pi} = \int_{x=0}^{\infty} \overline{w}(x)g_{N-1}(x)dx - c. \quad (17)$$

Proof of Proposition 1. First, the Seller can always attain profit $\overline{\Pi}$ by setting a posted price of $\overline{\Pi} + c$. In other words, run a mechanism in which the bidders have binary actions, *Buy* or *Not Buy*. The good is randomly allocated to one of the bidders who say *Buy*, and a bidder who wins the good has to pay $\overline{\Pi} + c$. Since $\overline{\Pi} + c$ is the expectation of the value conditional on a signal of $s_i = 0$, and since \overline{w} is increasing in the sum of the messages, all bidders would want to purchase at that price, so that the good would be sold at a price of $\overline{\Pi} + c$ with probability one.

Second, under $\overline{\mathcal{T}}$, the virtual value reduces to

$$\overline{\psi}_i(ds) = \overline{\gamma}(\Sigma s)ds - \overline{\gamma}(ds_i + \Sigma s_{-i})ds_{-i}.$$

This is because we assumed that the distribution of the signals was exponential so the hazard rate is constant and equal to one, and also the value is only a function of the sum, so the partial derivative is the same for all bidders. Moreover, Lemma 2 (d) shows that $\overline{\gamma}$ grows sub-exponentially, which implies that this measure is non-negative. In particular,

for any interval $[a, b]$, the bound (13) implies that

$$\begin{aligned}
& \int_{s_i=a}^b (\bar{\gamma}(\Sigma s) ds_i - \bar{\gamma}(ds_i + \Sigma s_{-i})) \\
& \geq \int_{s_i=a}^b \bar{\gamma}(b + \Sigma s_{-i}) \exp(s_i - b) ds_i - \bar{\gamma}(b + \Sigma s_{-i}) + \bar{\gamma}(a + \Sigma s_{-i}) \\
& = \bar{\gamma}(b + \Sigma s_{-i})(1 - \exp(a - b)) - \bar{\gamma}(b + \Sigma s_{-i}) + \bar{\gamma}(a + \Sigma s_{-i}) \\
& = \bar{\gamma}(a + \Sigma s_{-i}) - \bar{\gamma}(b + \Sigma s_{-i}) \exp(a - b) \geq 0,
\end{aligned}$$

where the last inequality again follows from (13).

As a result of the virtual-value measure being non-negative, the upper bound on profit is always increasing in q_i , so that it is optimal to always allocate the good. Moreover, ψ_i is always the same for all bidders, in the sense that the measure is exchangeable, so that all allocations that always allocate the good generate precisely the same profit.

Thus, an upper bound on profit is

$$\begin{aligned}
\int_{x=0}^{\infty} g_N(x) (\bar{\gamma}(x) dx - d\bar{\gamma}(x)) &= \int_{x=0}^{\infty} \bar{\gamma}(x) (g_N(x) - g'_N(x)) dx \\
&= \int_{x=0}^{\infty} \bar{\gamma}(x) g_{N-1}(x) dx,
\end{aligned}$$

where we have used the integration-by-parts formula for the Lebesgue-Stieltjes integral and then used the facts that for $N \geq 2$, $g_N(0) = g_N(\infty) = 0$ (and $\bar{\gamma}$ is bounded) and that

$$g'_N(x) = g_{N-1}(x) - g_N(x).$$

□

4.3 Sufficient conditions for profit guarantee

We say a mechanism \mathcal{M} is *standard* if the message space is $M_i = \mathbb{R}_+$, and the allocation q_i and transfer t_i are differentiable over $m_i \in M_i$.

Under a type space $\mathcal{T} = (S, \pi, w)$, an equilibrium on a mechanism \mathcal{M} consists of strategies $\beta_i : S_i \rightarrow \Delta M_i$ for each bidder i . Because the prior H is a mean-preserving spread of the interim values $w(s)$ where s is distributed according to $\pi \in \Delta(S)$, there exists a transition kernel $\tilde{w}(s) \in \Delta(V)$ that

1. has a mean of $w(s)$, i.e., $w(s) = \int_v v \tilde{w}(s)(dv)$ for every s ,
2. induces H when integrated with respect to π , i.e., $H(B) = \int_s \tilde{w}(s)(B) \pi(ds)$ for every measurable $B \subseteq V$.

An outcome $\sigma \in \Delta(V \times M)$ is then obtained by integrating the transition kernel $\tilde{w}(s) \times \beta_1(s_1) \times \cdots \times \beta_N(s_N)$ with respect to π .

Lemma 4 (Local obedience). *At any outcome $\sigma \in \Delta(V \times M)$ that is induced by an equilibrium on any type space, we have*

$$\int_{(v,m)} \left(v \frac{\partial q_i}{\partial m_i}(m) - \frac{\partial t_i}{\partial m_i}(m) \right) \sigma(dv, dm) \leq 0 \quad (18)$$

for every bidder i .

This is a local version of the epistemic characterization of BCE due to Bergemann and Morris (2013, 2016). In a sense, the outcomes that satisfy (18) are a kind of “local upward” Bayes correlated equilibrium, in which bidders do not want to make local upward deviations. Specifically, in any type space and equilibrium, it must be that for all i and m_i , bidder i has no incentive to deviate to $m_i + \epsilon$, where $\epsilon > 0$ is an arbitrarily small amount of deviation. These constraints turn out to be sufficient to pin down minimum profit of the maxmin mechanism.

The following is a standard weak duality result, which says that any feasible solution to the dual constraints in (19) provides a lower bound on equilibrium profit for any smooth mechanism.

Lemma 5 (Weak duality). *Fix a function $\lambda : V \rightarrow \mathbb{R}$. Suppose a standard mechanism (q_i, t_i) satisfies*

$$\sum_{i=1}^N t_i(m) - cQ(m) + \sum_{i=1}^N \left(v \frac{\partial q_i}{\partial m_i}(m) - \frac{\partial t_i}{\partial m_i}(m) \right) \geq \lambda(v), \quad (19)$$

for all $(v, m) \in V \times M$. Then under any equilibrium of any type space, the expected profit of (q_i, t_i) is at least $\int_v \lambda(v) H(dv)$.

Proof. Fix an outcome $\sigma \in \Delta(V \times M)$ induced by a type space and an equilibrium. We have

$$\begin{aligned} & \int_{(v,m)} \left(\sum_i t_i(m) - cQ(m) \right) \sigma(dv, dm) \\ & \geq \int_{(v,m)} \left(\sum_i t_i(m) - cQ(m) + \sum_i \left(v \frac{\partial q_i}{\partial m_i}(m) - \frac{\partial t_i}{\partial m_i}(m) \right) \right) \sigma(dv, dm) \\ & \geq \int_{(v,m)} \lambda(v) \sigma(dv, dm) = \int_v \lambda(v) H(dv), \end{aligned}$$

where in the second line we apply the inequality in (18). \square

Proposition 2. *Suppose in a standard mechanism, the total allocation sensitivity $\mu(m)$ depends only on Σm (i.e., $\mu(\Sigma m) \equiv \mu(m)$) and is decreasing over Σm . Define*

$$\lambda(v) = \Pi + \int_{x=0}^{\infty} \mu(x) G_N(x) d\hat{w}(x) - \int_{\nu=v}^{v_{max}} \mu(G_N^{-1}(H(\nu))) d\nu, \quad (20)$$

and suppose the profit-incentive alignment is satisfied. Then under any equilibrium in any type space, the expected profit is at least Π .

Proof. We verify that the dual constraints in Lemma 5 are satisfied. By construction, the dual constraint is satisfied with an equality for any m and $v = \widehat{w}(\Sigma m)$.

Let $C \equiv \Pi + \int_{x=0}^{\infty} \mu(x) G_N(x) d\widehat{w}(x)$. Suppose $v \neq \tilde{v} = \widehat{w}(\Sigma m)$, then we have

$$\begin{aligned} \lambda(v) &= C - \int_{\nu=v}^{\nu_{\max}} \mu(G_N^{-1}(H(\nu))) d\nu \\ &\leq (v - \tilde{v})\mu(\Sigma m) + C - \int_{\nu=\tilde{v}}^{\nu_{\max}} \mu(G_N^{-1}(H(\nu))) d\nu \\ &= (v - \tilde{v})\mu(\Sigma m) + \tilde{v}\mu(\Sigma m) - \Xi(m) - cQ(m) \\ &= v\mu(\Sigma m) - \Xi(m) - cQ(m), \end{aligned}$$

where the second line follows from the assumption that $\mu(x)$ is decreasing in x , and the third line follows from the profit-incentive alignment for m and $\tilde{v} = \widehat{w}(\Sigma m)$. Therefore, the dual constraint is satisfied for every v and m .

Since $\widehat{w}(x) = H^{-1}(G_N(x))$ and \widehat{w} is right-continuous, the substitution rule for Riemann-Stieltjes integrals implies that

$$\int_{x=0}^{\infty} \mu(x) G_N(x) d\widehat{w}(x) = \int_{v=0}^{\nu_{\max}} \int_{\nu=v}^{\nu_{\max}} \mu(G_N^{-1}(H(\nu))) d\nu H(dv).$$

(For a helpful discussion, see Falkner and Teschl (2012).) That is, $\int_v \lambda(v) H(dv) = \Pi$. Therefore, Lemma 5 implies that the profit guarantee is at least Π . \square

Note that in the proof of Proposition 2 we have not used any property of Π , so the profit guarantee of Π might as well have been any arbitrarily large number. As we discussed in Section 3, a potential fallacy here is that the profit guarantee of the mechanism in Proposition 2 may be vacuous, for the mechanism may not have any equilibrium under any type space. Thus, in our subsequent application of Proposition 2 we must impose additional assumptions on the mechanism, and in particular on the total allocation sensitivity, to ensure equilibrium existence and hence obtain a non-vacuous profit guarantee.

4.4 Construction of a maxmin mechanism

We now construct a maxmin mechanism $\overline{\mathcal{M}}$; we denote its allocation and transfer functions as \overline{q}_i and \overline{t}_i . It is a standard mechanism with a message space $M_i = \mathbb{R}_+$ coinciding with the signal space of $\overline{\mathcal{T}}$.

We construct a *proportional allocation rule* as described in (9), and where the total supply function is

$$\overline{Q}(x) = \begin{cases} 1 & \text{if } \widehat{\Gamma}(x) = \overline{\Gamma}(x), \\ \frac{x C(a,b)}{N} + \frac{D(a,b)}{x^{N-1}} & \text{if } x \in [a, b], \text{ where } [a, b] \text{ is a graded interval,} \end{cases} \quad (21)$$

and

$$C(a, b) = N \frac{b^{N-1} - a^{N-1}}{b^N - a^N},$$

$$D(a, b) = \frac{b - a}{b^N - a^N} a^{N-1} b^{N-1}.$$

Our transfer rule is:

$$\bar{t}_i(m) = \frac{1}{N} \sum_{n=1}^N \frac{1}{\binom{N-1}{n-1}} \sum_{\substack{I \subseteq \mathcal{N} \setminus \{i\} \\ |I|=n-1}} \int_{x=0}^{\infty} (\bar{\Xi}(\Sigma m_I + x) - \bar{\Xi}(m_i + \Sigma m_I + x)) g_{N-n+1}(x) dx, \quad (22)$$

where $\mathcal{N} = \{1, \dots, N\}$, and $\bar{\Xi}$ is the total excess growth given by the profit-incentive alignment:

$$\bar{\Xi}(x) = \bar{\mu}(x) \hat{w}(x) - \bar{\lambda}(\hat{w}(x)) - c \bar{Q}(x), \quad (23)$$

and

$$\bar{\lambda}(v) = \bar{\Pi} + \int_{x=0}^{\infty} \bar{\mu}(x) G_N(x) d\hat{w}(x) - \int_{\nu=v}^{\nu_{\max}} \bar{\mu}(G_N^{-1}(H(\nu))) d\nu. \quad (24)$$

In the equations above $\bar{\mu}$ is the total sensitivity from the allocation in (9) and is spelled out in (25).

Finally, let $\bar{\beta}_i$ be the truthful strategy in the mechanism $\bar{\mathcal{M}}$ under the type space $\bar{\mathcal{T}}$: $\bar{\beta}_i(s_i) = m_i$ for every $s_i = m_i \in S_i = M_i$.

Theorem 1. *The tuple $(\bar{\mathcal{M}}, \bar{\mathcal{T}}, \bar{\beta})$ is a solution to the Seller's Problem. The solution has a value of $\bar{\Pi}$ defined by (17).*

Before proceeding to the proof, we note some immediate properties of the allocation rule in $\bar{\mathcal{M}}$. First, outside the graded region the total supply $\bar{Q}(\Sigma m)$ is always equal to one, and the allocation is simply proportional to the message: $\bar{q}_i(m) = m_i / \Sigma m$. On the other hand, inside a graded interval the total supply $\bar{Q}(\Sigma m)$ is always strictly less than one.¹⁵ Interestingly, the total supply $\bar{Q}(\Sigma m)$ can decrease with the aggregate message Σm ; that is, the Seller might want to withhold more of the good when the bidders have more optimistic signals about the value.

Second, the total allocation sensitivity depends only on the aggregate message $x = \Sigma m$ and has the following form:

$$\bar{\mu}(x) = \begin{cases} \frac{N-1}{x} & \text{if } \hat{\Gamma}(x) = \bar{\Gamma}(x), \\ C(a, b) & \text{if } x \in [a, b], \text{ where } [a, b] \text{ is a graded interval.} \end{cases} \quad (25)$$

¹⁵On a graded interval (a, b) , $Q(x) < 1$ is equivalent to:

$$D(a, b) > \frac{1/x - 1/a}{1/x^N - 1/a^N} = \frac{1}{\sum_{n=0}^{N-1} 1/x^n \cdot 1/a^{N-1-n}}.$$

Clearly, $\frac{1/x - 1/a}{1/x^N - 1/a^N}$ increases with x , and equals $D(a, b)$ at $x = b$. Therefore, the above inequality holds for $x \in (a, b)$.

In particular, solving the differential equation $\bar{\mu}(x) = \frac{N-1}{x}\bar{Q}(x) + \bar{Q}'(x) = C(a, b)$ on a graded interval $[a, b]$ gives the functional form (21). Moreover, we choose the constant $D(a, b)$ so that \bar{Q} is continuous and equal to 1 at the endpoints. Thus, given this choice of $D(a, b)$, the total supply \bar{Q} and the allocation rules \bar{q}_i are continuous functions, even though the total allocation sensitivity $\bar{\mu}$ is discontinuous.

We present the proof of Theorem 1 in two steps. In Step 1 we show that the mechanism $\bar{\mathcal{M}}$ has the optimal profit guarantee and satisfies some additional properties. In Step 2 we show that these additional properties imply the truth-telling equilibrium.

4.4.1 Step 1: Characterization of $\bar{\mathcal{M}}$

As alluded after Proposition 2, we design the total allocation sensitivity $\bar{\mu}$ so that the following condition holds for the the total excess growth $\bar{\Xi}$, which will facilitate the existence of equilibrium:

Lemma 6. *The total excess growth $\bar{\Xi}$ has zero expectation under $\bar{\mathcal{T}}$:*

$$\int_{m \in M} \bar{\Xi}(\Sigma m) \exp(-\Sigma m) dm = 0. \quad (26)$$

Proof. Using the formula for $\bar{\Xi}$ in equation (23), we see that equation (26) holds if and only if:

$$\int_{x=0}^{\infty} (\bar{\mu}(x)\hat{w}(x) - \bar{\Pi} - c\bar{Q}(x)) g_N(x) dx = 0. \quad (27)$$

To see why (27) is true, suppose first that $c = 0$. Writing out the expression for $\bar{\Pi}$ and using $\hat{w}(x) = \hat{\gamma}(x)$, we see that (27) is equivalent to

$$\bar{\Pi} = \int_{x=0}^{\infty} \bar{\gamma}(x)g_{N-1}(x)dx = \int_{x=0}^{\infty} \hat{\gamma}(x)\bar{\mu}(x)g_N(x)dx.$$

When $\bar{\Gamma}(x) = \hat{\Gamma}(x)$, we have $\bar{\gamma}(x) = \hat{\gamma}(x)$ and $\bar{\mu}(x) = (N-1)/x$, so $\bar{\mu}(x)g_N(x) = g_{N-1}(x)$ and the two integrands above are exactly equal. On the other hand, over a graded interval (a, b) , we construct $\bar{\mu}(x)$ to be the constant so that the above two integrals over (a, b) are equal:

$$C(a, b) = \frac{\int_{x=a}^b \bar{\gamma}(x)g_{N-1}(x)dx}{\int_{x=a}^b \hat{\gamma}(x)g_N(x)dx} = \frac{\int_{x=a}^b \bar{\gamma}(a) \exp(x-a)g_{N-1}(x)dx}{\int_{x=a}^b \bar{\gamma}(a) \exp(x-a)g_N(x)dx},$$

since $\bar{\gamma}(x) = \bar{\gamma}(a) \exp(x-a)$ and $\hat{\Gamma}(b) - \hat{\Gamma}(a) = \bar{\Gamma}(b) - \bar{\Gamma}(a)$ (cf. Lemma 2).

For the case of $c > 0$, we note that $\bar{\mu}(x) = \bar{Q}'(x) + \frac{N-1}{x}\bar{Q}(x)$, and

$$c \int_{x=0}^{\infty} \bar{Q}(x)g_N(x) dx = c \int_{x=0}^{\infty} (\bar{Q}'(x) + (N-1)\bar{Q}(x)/x)g_N(x) dx$$

by an integration by parts, so we can replace $-c\bar{Q}(x)$ by $-c\bar{\mu}(x)$ in equation (27), and the above reasoning again implies that equation (27) holds, since $\hat{\gamma}(x) = \hat{w}(x) - c$. \square

We now summarize the main properties of the mechanism $\overline{\mathcal{M}}$.

Lemma 7. *The total allocation sensitivity $\overline{\mu}(\Sigma m)$ is decreasing in Σm . The excess growth $\overline{\xi}_i(m) = \frac{\partial \overline{t}_i}{\partial m_i}(m) - \overline{t}_i(m)$ satisfies profit-incentive alignment and strong transversality.*

Proof. The statements on decreasing total allocation sensitivity and profit-incentive alignment follow immediately from our construction.

We first explicitly construct excess growth $\overline{\xi}_i(m)$ that satisfies strong transversality, and then show that they imply the transfers in (22). We use the following analogy: It is as if the bidders share a common account that, as a function of the bidders' messages, has a balance $\overline{\Xi}(\Sigma m)$. We want to allocate funds to the bidders, as a function of m , so that (i) the sum of the bidders' individual accounts is equal to $\overline{\Xi}(\Sigma m)$ and (ii) conditional on m_{-i} , bidder i 's account has an average balance of zero with respect to the exponential density $\exp(-m_i)$. Note that the individual accounts can either be positive or negative, but they must always sum to $\overline{\Xi}$ (which can also be positive or negative).

How can such a division be accomplished? Consider the following procedure. After realizing the type profile m , we will arrange the bidders in a random order, with all orders being equally likely. We denote a realized order by a bijective mapping $\zeta : \{1, \dots, N\} \rightarrow \{1, \dots, N\}$, with $\zeta(1)$ denoting the identity of the first bidder, $\zeta(2)$ denoting the second bidder, etc. Equivalently, $\zeta^{-1}(i)$ is the position of bidder i in the line. Let Z denote the set of all such permutations.

Given a randomly drawn ζ , we start by giving bidder $\zeta(1)$ all of the account $\overline{\Xi}(\Sigma m)$. Now, if each bidder i receives all of the account when she is first in line and there were no further adjustments, then we would generally violate (ii) above, since the expectation of $\overline{\xi}_i(m_i, m_{-i})$ with respect to m_i would be the expectation of $\overline{\Xi}(m_i + \Sigma m_{-i})/N$ across m_i , which need not be zero. So, we will have the other bidders in line make transfers¹⁶ that on average will cancel out bidder i 's balance.

In particular, we have the second bidder $\zeta(2)$ make a net transfer to bidder $\zeta(1)$ in the amount of

$$\begin{aligned} \tau_2(\zeta, m) &= -\mathbb{E}_{m'_{\zeta(1)}} \left[\overline{\Xi}(m'_{\zeta(1)} + \Sigma m_{\zeta(\geq 2)}) \right] \\ &= -\int_{x=0}^{\infty} \overline{\Xi}(x + \Sigma m_{\zeta(\geq 2)}) g_1(x) dx, \end{aligned}$$

where $m_{\zeta(\geq n)}$ denotes the subvector of m with indices $\zeta(k)$ for $k \geq n$ (i.e., the messages of the bidders who are in the n th place or later), and we define $m_{\zeta(\leq n)}$ analogously. Thus, in expectation across $m_{\zeta(1)}$, bidder $\zeta(1)$'s expected account is zero. Of course, now bidder $\zeta(2)$ may be left with a non-zero balance across $m_{\zeta(2)}$, which we do not want. The solution is then to have bidder $\zeta(3)$ pay bidder $\zeta(2)$ an amount equal to

$$\begin{aligned} \tau_3(\zeta, m) &= -\mathbb{E}_{m'_{\zeta(\leq 2)}} \left[\tau_2(\zeta, m'_{\zeta(\leq 2)}, m_{\zeta(\geq 3)}) \right] \\ &= -\int_{x=0}^{\infty} \overline{\Xi}(x + \Sigma m_{\zeta(\geq 3)}) g_2(x) dx. \end{aligned}$$

¹⁶Here we do not use the word transfer to mean the transfer between the bidders and the Seller, but rather it is a figurative transfer of growth rate between bidders in order to balance out the excess growth.

This cancels out bidder $\zeta(2)$'s transfer, on average across $m_{\zeta(2)}$, but now creates a non-zero expected balance for bidder $\zeta(3)$.

Continuing in this manner, the n th bidder pays an amount

$$\begin{aligned}\tau_n(\zeta, m) &= -\mathbb{E}_{m'_{\zeta(\leq n-1)}} [\tau_{n-1}(\zeta, m'_{\zeta(\leq n-1)}, m_{\zeta(\geq n)})] \\ &= -\int_{x=0}^{\infty} \bar{\Xi}(x + \Sigma m_{\zeta(\geq n)}) g_{n-1}(x) dx.\end{aligned}$$

to bidder $\zeta(n-1)$. Our convention is that $\tau_1(\zeta, m) = -\bar{\Xi}(\Sigma m)$ (meaning the first bidder takes all of $\bar{\Xi}$). Finally, the last bidder $\zeta(N)$ makes a transfer to bidder $\zeta(N-1)$, which on average across $m_{\zeta(N)}$ is equal to

$$-\int_{m_{\zeta(N)}=0}^{\infty} \int_{x=0}^{\infty} \bar{\Xi}(x + m_{\zeta(N)}) g_{N-1}(x) dx \exp(-m_{\zeta(N)}) dm_{\zeta(N)} = 0,$$

since this is just the unconditional expectation of $\bar{\Xi}(\Sigma m)$. As a result, when we integrate each bidder i 's message, all of the ‘‘adjustments’’ cancel out, and

$$\mathbb{E}_{m_{\zeta(n)}} [\tau_{n+1}(\zeta, m) - \tau_n(\zeta, m)] = 0,$$

where $\tau_{N+1} = 0$. Averaging across all orders, bidder i 's expected excess growth is

$$\bar{\xi}_i(m) = \frac{1}{N!} \sum_{\zeta \in Z} (\tau_{\zeta^{-1}(i)+1}(\zeta, m) - \tau_{\zeta^{-1}(i)}(\zeta, m)). \quad (28)$$

We can give a more explicit formula for \bar{t}_i as follows. Note that because the expectation of $\bar{\xi}_i$ is zero given m_{-i} , we can rewrite (8) as

$$\begin{aligned}\bar{t}_i(m) &= \exp(m_i) \int_{x=0}^{m_i} \bar{\xi}_i(x, m_{-i}) \exp(-x) dx \\ &= -\exp(m_i) \int_{x=m_i}^{\infty} \bar{\xi}_i(x, m_{-i}) \exp(-x) dx \\ &= -\int_{x=0}^{\infty} \bar{\xi}_i(m_i + x, m_{-i}) \exp(-x) dx.\end{aligned}$$

As a result,

$$\bar{t}_i(m) = \frac{1}{N!} \sum_{\zeta \in Z} \int_{x=0}^{\infty} (\bar{\Xi}(x + m_{\zeta(\geq \zeta^{-1}(i))}) - \bar{\Xi}(x + m_i + m_{\zeta(\geq \zeta^{-1}(i))})) g_{N-\zeta^{-1}(i)}(x) dx$$

The second term inside the sum is simply the expectation of $\bar{\Xi}$ when we condition on the messages of the bidders who come after bidder i , according to the order ζ . The first term is the same expectation, but with the sum of the messages increased by m_i . The expected post transfer is equal to the difference between these two expectations, averaged across all orders. Finally, it is easy to see that the transfer in the above equation is exactly equal to that in (22). \square

4.4.2 Step 2: Truth-telling equilibrium

We now show that truth-telling is an equilibrium.

Lemma 8. *Given the allocation \bar{q}_i in equations (9) and (21), if the transfer rule satisfies profit-incentive alignment (PIA) and weak transversality (WT), then the strategy profile $\bar{\beta}$ is an equilibrium under \bar{T} . That is:*

$$\begin{aligned} & \int_{m_{-i} \in M_{-i}} (\bar{w}(\Sigma m) \bar{q}_i(m) - \bar{t}_i(m)) \exp(-\Sigma m_{-i}) dm_{-i} \\ & \geq \int_{m_{-i} \in M_{-i}} (\bar{w}(\Sigma m) \bar{q}_i(m'_i, m_{-i}) - \bar{t}_i(m'_i, m_{-i})) \exp(-\Sigma m_{-i}) dm_{-i}, \end{aligned} \quad (29)$$

for every bidder i and every pair $(m_i, m'_i) \in M_i \times M_i$.

The complete proof of Lemma 8 is in the Appendix. Here is a sketch of the argument. As a first step, we prove the local incentive compatibility:

$$\int_{m_{-i}} \left(\bar{w}(m_i + \Sigma m_{-i}) \frac{\partial \bar{q}_i(m)}{\partial m_i} - \frac{\partial \bar{t}_i(m)}{\partial m_i} \right) \exp(-\Sigma m_{-i}) dm_{-i} = 0 \quad (30)$$

for all i and m_i . When there is no grading,

$$\begin{aligned} & \int_{m_{-i}} \bar{w}(m_i + \Sigma m_{-i}) \frac{\partial \bar{q}_i(m)}{\partial m_i} \exp(-\Sigma m_{-i}) dm_{-i} \\ & = - \int_{x=0}^{\infty} \bar{w}(m_i + x) g_N(x) d\bar{\mu}(m_i + x). \end{aligned} \quad (31)$$

This is fairly easy to see, since in this case the allocation is $\bar{q}_i(m) = m_i / \Sigma m$, which implies $\bar{\mu}(x) = (N-1)/x$ and $d\bar{\mu}(x) = \bar{\mu}'(x)dx = -((N-1)/x^2)dx$. Thus, we have

$$\frac{\partial \bar{q}_i(m)}{\partial m_i} = \frac{\Sigma m_{-i}}{(\Sigma m)^2} = -\bar{\mu}'(\Sigma m) \frac{\Sigma m_{-i}}{N-1}$$

so that

$$\frac{\partial \bar{q}_i(m)}{\partial m_i} g_{N-1}(\Sigma m_{-i}) = -\bar{\mu}'(m_i + \Sigma m_{-i}) g_N(\Sigma m_{-i}).$$

Integrating with respect to $x = \Sigma m_{-i}$, which has density g_{N-1} , we have the result. Our proof in the Appendix shows that (31) can be suitably generalized when there are graded intervals.

We then show that (31) is exactly equal to the expectation of $\partial \bar{t}_i / \partial m_i$ over m_{-i} . Because of weak transversality, we have

$$\begin{aligned} & \int_{m_{-i}} \frac{\partial \bar{t}_i(m)}{\partial m_i} \exp(\Sigma m_{-i}) dm_{-i} \\ & = \int_{m_{-i}} (\bar{t}_i(m) + \bar{\xi}_i(m_i, m_{-i})) \exp(\Sigma m_{-i}) dm_{-i} \\ & = \int_{y=0}^{m_i} \int_{x=0}^{\infty} \bar{\Xi}(y+x) g_{N-1}(x) dx \exp(m_i - y) dy + \int_{x=0}^{\infty} \bar{\Xi}(m_i + x) g_{N-1}(x) dx, \end{aligned}$$

where we use the fact that $\bar{t}_i(m) = \int_{y=0}^{m_i} \exp(m_i - y) \bar{\xi}_i(y, m_{-i}) dy$ before applying the weak transversality condition in the first part of the third line.

But remembering that the unconditional expectation of $\bar{\Xi}$ is zero, it must be that

$$\begin{aligned}
& \int_{y=0}^{m_i} \int_{x=0}^{\infty} \bar{\Xi}(y+x) g_{N-1}(x) dx \exp(m_i - y) dy \\
&= - \int_{y=m_i}^{\infty} \int_{x=0}^{\infty} \bar{\Xi}(y+x) g_{N-1}(x) dx \exp(m_i - y) dy \\
&= - \int_{y=0}^{\infty} \int_{x=0}^{\infty} \bar{\Xi}(m_i + y + x) g_{N-1}(x) dx \exp(-y) dy \\
&= - \int_{x=0}^{\infty} \bar{\Xi}(m_i + x) g_N(x) dx.
\end{aligned}$$

Thus,

$$\begin{aligned}
\int_{m_{-i}} \frac{\partial \bar{t}_i(m)}{\partial m_i} \exp(\Sigma m_{-i}) dm_{-i} &= \int_{x=0}^{\infty} \bar{\Xi}(m_i + x) (g_{N-1}(x) - g_N(x)) dx \\
&= - \int_{x=0}^{\infty} g_N(x) d\bar{\Xi}(m_i + x),
\end{aligned}$$

where we used the integration-by-parts formula for the Riemann-Stieltjes integral, and the fact that $g'_N = g_{N-1} - g_N$. Finally, the profit-incentive alignment in (23) implies that $d\Xi = \hat{w}d\mu$ when there is no grading, as the total supply \bar{Q} is always 1, thus proving the local incentive compatibility (30) in this case since $\hat{w} = \bar{w}$. In the Appendix, we generalize this argument when there is grading.

Now, by itself this only guarantees that truthful reporting is a local optimum. In the proof in the Appendix, we show that truth-telling is globally optimal. Once again, this is relatively easy to see when there is no grading, since in this case the allocation is $\bar{q}_i(m) = m_i/\Sigma m$, which is always increasing in one's own message. It is well known that when both the value function and allocation are increasing in one's own message, local incentive compatibility implies global incentive compatibility. The argument for the general case is slightly more complicated, since the bidders' utility from misreporting need not be single-peaked in the report.

Combining Proposition 2 with Lemmas 7 and 8, we conclude

Proposition 3. *In the mechanism $\overline{\mathcal{M}}$, the expected profit under any type space and any equilibrium is at least $\bar{\Pi}$. Moreover, the truth-telling strategy profile $\bar{\beta}$ is an equilibrium of $\overline{\mathcal{M}}$ under the type space $\overline{\mathcal{T}}$.*

Theorem 1 clearly follows from Propositions 1 and 3.

4.4.3 Graphical illustration

Figure 8 depicts the objects used in the construction of $\overline{\mathcal{M}}$ for an example in which $N = 3$ and the distribution of v is uniform on $[0, 0.8]$ with probability 0.8 and is uniform

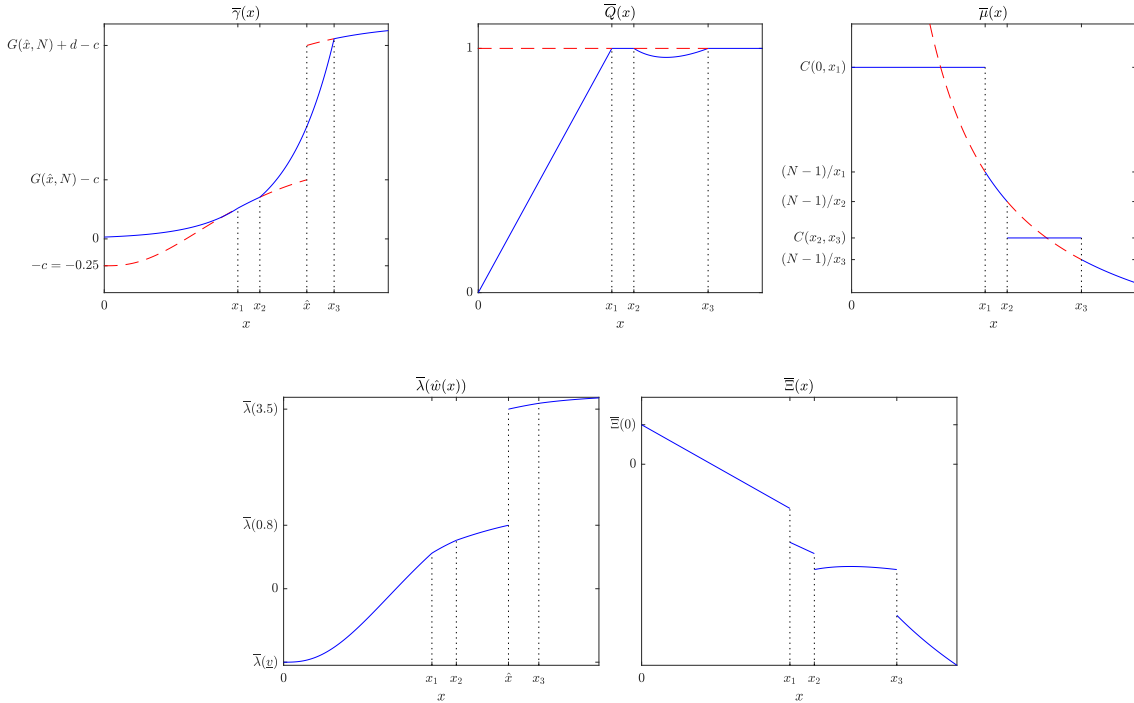


Figure 8: Objects used in the construction of $\bar{\mathcal{M}}$, for a modified uniform example in which $N = 3$ and the distribution of v is uniform on $[0, 0.8]$ with probability 0.8 and is uniform on $[3.5, 3.7]$ with probability 0.2. The red dashed lines show the corresponding objects in the must-sell case. There are two regions where the gains function is graded to obtain $\bar{\gamma}$: around $x = 0$, and around the discontinuity at $x = \hat{x}$.

on $[3.5, 3.7]$ with probability 0.2, and $c = 0.25$. For this distribution, the fully-revealing value function is $\hat{w} = G_N(x) + 2.7\mathbb{1}_{x > \hat{x}}$, where $G_N(\hat{x}) = 0.8$, which is approximately $\hat{x} = 4.279$. There are two graded intervals, which we denote by $[0, x_1]$ and $[x_2, x_3]$, where $0 < x_1 < x_2 < x_3$ and \hat{x} is contained in the latter interval. Interestingly, while \hat{w} and $\bar{\lambda} \circ \hat{w}$ have discontinuities around $x = \hat{x}$, these discontinuities cancel one another out so $\bar{\Xi}$ is continuous at \hat{x} .

4.5 Two special cases

4.5.1 Single-crossing distributions

There are two special cases in which our solution takes on a particularly simple form. First, there is a class of distributions for which there is at most one graded region of the form $[0, x^*]$, so that the only rationing takes the form of a linear supply rule $Q(x) = \min\{x/x^*, 1\}$. This kind of rationing seems to be quite important, and as we show in Section 5.2, this is the only kind of grading/rationing that survives when the number of bidders is large.

We say a distribution H is *single-crossing* if the virtual value associated with $\hat{\gamma}(x)$ crosses zero only once: there exists a cutoff $\bar{x} \geq 0$ such that $\hat{\gamma}(x) - d\hat{\gamma}(x) \leq 0$ for $x \in [0, \bar{x})$ and $\hat{\gamma}(x) - d\hat{\gamma}(x) \geq 0$ for $x \in [\bar{x}, \infty)$. Both the standard uniform and the binary values case studied by Bergemann, Brooks, and Morris (2016a) are single crossing for all N .

Suppose we have a single-crossing distribution with cutoff \bar{x} . Given the functions $\hat{\Gamma}$ and E from the grading procedure, we first note that the function $\hat{\Gamma} \circ E^{-1}$ is convex on $[0, E(\bar{x}))$ and concave on $(E(\bar{x}), \infty)$. Recall that $\bar{\Gamma} \circ E^{-1}$ is the concavification of $\hat{\Gamma} \circ E^{-1}$. It must therefore be linear on $[0, z^*)$ and coincides with $\hat{\Gamma} \circ E^{-1}$ on (z^*, ∞) , for some $z^* > E(\bar{x})$. Setting $x^* = E^{-1}(z^*)$, Lemma 2 then implies that the minmax type space has the following gains function:

$$\bar{\gamma}(x) = \begin{cases} \bar{\gamma}(x^*) \exp(x - x^*) & x < x^* \\ \hat{\gamma}(x) & x \geq x^* \end{cases}.$$

Next we turn to the maxmin allocation rule. On the graded interval $[0, x^*]$, we have $\bar{Q}(x) = xC(0, x^*)/N = x/x^*$, since $D(0, x^*) = 0$. Therefore, the maxmin allocation is simply:

$$\bar{q}_i(m_i, m_{-i}) = \frac{m_i}{\max\{x^*, \Sigma m\}}.$$

We can interpret m_i as bidder i 's nominal demand for the good and x^* as the unit of the demand, so m_i/x^* is the demand in units of probability of being allocated the good. With this interpretation, our allocation rule simply states that the bidders get their demands if the total demand is feasible (i.e., less than 1); if the total demand is not feasible, then the good is rationed in proportion to the demands.

This structure is similar to the allocation rule studied by Bergemann, Brooks, and Morris (2016a). Indeed, the binary value distribution they looked at, which has support $\{0, 1\}$, always satisfies the single-crossing property regardless of the number of bidders. For when the value is zero, a strictly increasing value must be growing super-exponentially. As we discuss below in Section 4.6, our proportional allocation rule is slightly different than the one they used, and it generalizes more readily.

4.5.2 The must-sell case

An important variant of our model is when the Seller has to sell the good, i.e.,

$$\sum_{i=1}^N q_i(m) = 1 \tag{32}$$

for every message profile m . All of our existing tools carry over to the must-sell setting and almost immediately give us the solution. Consider the type space $\hat{\mathcal{T}}$ where the signals are i.i.d. and exponential draws on \mathbb{R}_+ , and the value function is \hat{w} , i.e., the fully-revealing value function. Also consider the mechanism $\hat{\mathcal{M}}$ corresponding to the efficient

proportional rule:

$$\hat{q}_i(m_i, m_{-i}) = \frac{m_i}{\Sigma m},$$

and the allocations are equal to $1/N$ if $\Sigma m = 0$. We define the transfers \hat{t}_i according to the same formulae for λ , Ξ , and ξ_i , so that they satisfy PIA and strong transversality. Finally, let $\hat{\beta}$ denote the same truthful/obedient strategies as were part of our general solution.

Theorem 2 (Must-sell solution). *The triple $(\widehat{\mathcal{M}}, \widehat{\mathcal{T}}, \widehat{\beta})$ is a solution to the must-sell Seller's Problem. The solution has a value of $\widehat{\Pi}$ given by (2).*

Proof of Theorem 2. The proofs of Propositions 2 and 3 remain valid with $\hat{\gamma}$ in place of $\bar{\gamma}$; in particular, the proof sketch of Lemma 8 is a complete proof of the lemma that the truthful strategy profile is an equilibrium of $\widehat{\mathcal{M}}$ under $\widehat{\mathcal{T}}$, since by definition there is no grading in $\widehat{\mathcal{T}}$. The only minor complication is the discontinuity in the allocation rule at zero. But since the messages are all zero with zero probability under the truthful strategies, the bidders' payoffs are continuous in their messages when they deviate, and the rest of our proof goes through. Thus, the mechanism $\widehat{\mathcal{M}}$ guarantees the Seller at least $\widehat{\Pi}$ in any equilibrium, and there is an equilibrium on the type space $\widehat{\mathcal{T}}$.

Indeed, the only place where our argument changes is in the proof of Proposition 1. If the Seller could keep the good, then $\widehat{\Gamma}$ is not an upper bound on optimal profit, since in general the Seller can do better by sometimes keeping the good. In the must-sell case, however, such mechanisms are not available. Moreover, it is still the case that all bidders have the same virtual value (due to the value being a function of the sum of the signals, and the i.i.d. exponential distribution of signals). In particular, the virtual value reduces to

$$w(s)ds - w(ds_i, s_{-i})ds_{-i} - c = \hat{\gamma}(\Sigma s)ds - \hat{\gamma}(d\Sigma s).$$

This means that *all* must-sell allocations generate exactly the same profit on $\widehat{\mathcal{T}}$, which is $\widehat{\Pi}$. Thus, it is impossible for the Seller to deviate to a must-sell mechanism and equilibrium with higher profit. \square

4.6 Other solutions to the Seller's Problem

We have constructed a particular solution to the Seller's Problem. We know for a fact that there are others. While a complete characterization of the set of solutions is beyond our present capabilities, we are able to describe a class of solutions which generalizes the one described in Theorem 1. All of these solutions will use the same minmax type space $\bar{\mathcal{T}}$ ¹⁷ but different maxmin mechanisms, and they will all result in precisely the same profit guarantee. The trade-off is that the mechanisms will be implicitly described as the

¹⁷It is an open question whether there are other minmax type spaces. We strongly suspect that $\bar{\mathcal{T}}$ is unique if we assume independent signals and increasing values. However, we do not even have a tool to characterize maximum profit if we go outside this class. Intuitively, correlation in the signals would only make it easier for the Seller to separate types with Crémer-McLean type constructions.

solution to a system of equations, rather than obtained via an explicit construction as $\overline{\mathcal{M}}$. In effect, these are the maxmin mechanisms that can be reduced to direct mechanisms on $\overline{\mathcal{T}}$.

A critical juncture at which we used the particular structure of our mechanism is in proving that truth-telling is an equilibrium. To do so, we used the specific functional form of \overline{q} , together with (PIA) and (WT) to show that all global incentive constraints are satisfied. An alternative approach, which we now adopt, is to construct a mechanism \mathcal{M} that has a profit guarantee of $\overline{\Pi}$ and that is nice enough that we know there exists *some* equilibrium β on the minmax type space $\overline{\mathcal{T}}$. A fortiori, profit in this equilibrium will have to be $\overline{\Pi}$, so that $(\mathcal{M}, \overline{\mathcal{T}}, \beta)$ is a solution with value $\overline{\Pi}$.

Key concerns for equilibrium existence are (i) unbounded action spaces and (ii) discontinuous payoffs. Thus, to ensure equilibrium existence, we extend and compactify our message space by adding an infinite message, so that the message space is $\overline{M}_i = [0, \infty]$. We will also extend the allocation and transfer rules to this infinite message in such a way that any discontinuities are “mild” enough that standard equilibrium existence arguments still apply.

Let \mathbf{Q} be the set of allocation functions satisfying the following conditions. First, for all i , q_i should be continuous on $M \setminus \{0\}$, upper semi-continuous in m_i and lower semi-continuous in m_{-i} on \overline{M} ; moreover, the functions

$$\mu(m) = \sum_{i=1}^N \frac{\partial q_i(m)}{\partial m_i} \mathbb{I}_{m_i < \infty} \quad \text{and} \quad Q(m) = \sum_{i=1}^N q_i(m) \quad (33)$$

are continuous on $\overline{M} \setminus M$. Note that we now let the total allocation sensitivity and the total supply depend on the entire message profile, rather than just being a function of the aggregate message.

Second, it must satisfy

$$\begin{aligned} \Xi(m; q) &= \widehat{w}(\Sigma m) \sum_{i=1}^N \frac{\partial q_i(m)}{\partial m_i} - \overline{\lambda}(\widehat{w}(\Sigma m)) - cQ(m) \\ &\geq v \sum_{i=1}^N \frac{\partial q_i(m)}{\partial m_i} - \overline{\lambda}(v) - cQ(m) \end{aligned} \quad (34)$$

for all $m \in M$ and for all $v \in \text{supp } H$, where $\Xi(\cdot; q)$ is the total excess growth induced by the allocation rule. These conditions guarantee that $\overline{\mathcal{T}}$ is still a profit-minimizing type space. Third, the allocation rule must induce zero expected total excess growth:

$$\int_{m \in \mathbb{R}_+^N} \Xi(m; q) \exp(-\Sigma m) dm = 0.$$

Once again, the zero expected total excess growth condition ensures that there will exist bounded transfers that also satisfy (PIA).

Next, for each $q \in \mathbf{Q}$, let $\mathbf{T}(q)$ be the set of transfers such that for every $m \in \overline{M}$,

$$t_i(m) = \exp(m_i) \int_{x=0}^{m_i} \exp(-x) \xi_i(x, m_{-i}) dx, \quad (35)$$

where the excess growth ξ_i is bounded and continuous on \overline{M} and satisfies (PIA) and strong transversality (ST). Since Ξ has zero expectation, we know it is possible to solve this system, e.g., using the construction in Lemma 7.¹⁸ We let \mathbf{M} be the set of pairs (q, t) of allocation and transfer rules such that $q \in \mathbf{Q}$ and $t \in \mathbf{T}(q)$.

It is easy to extend our construction from Section 4.4 so that it is in \mathbf{M} . By defining

$$\bar{q}_i(m) = \begin{cases} \frac{\bar{Q}(0)}{N} & \text{if } \Sigma m = 0; \\ \frac{m_i}{\Sigma m} \bar{Q}(\Sigma m) & \text{if } 0 < \Sigma m < \infty; \\ \frac{1}{|\{j|m_j=\infty\}|} & \text{if } \Sigma m = \infty, \end{cases}$$

and it is easy to see that the extended \bar{q} belongs to \mathbf{Q} . Our transfer function \bar{t} from Section 4.4 is also defined by an excess growth ξ_i that is bounded and satisfies (PIA) and (ST). Thus, after extending \bar{t} from M to \overline{M} , the extended \bar{t} belongs to $\mathbf{T}(\bar{q})$.

The following proposition shows that any transfer rule for a mechanism in \mathbf{M} is nicely behaved:

Proposition 4. *Suppose $\xi_i(m)$ is bounded and continuous over $m \in [0, \infty]^N$. The transfer $\bar{t}_i(m)$ defined by equation (35) is continuous and bounded over $m \in [0, \infty]^N$ if and only if the excess growth satisfies strong transversality. If strong transversality holds, then $\bar{t}_i(\infty, m_{-i}) = -\xi_i(\infty, m_{-i})$ for every m_{-i} .*

Proof. Since $|\xi_i(m)|$ is bounded above by a constant, the continuity of $\int_{y=0}^{m_i} \exp(-y)\xi_i(y, m_{-i}) dy$ over $m \in [0, \infty]^N$ follows from the dominated convergence theorem. Thus, t_i is continuous over $m \in [0, \infty]^N$. Clearly, t_i is bounded as $m_i \rightarrow \infty$ if and only if strong transversality is satisfied and ξ_i is itself bounded, in which case L'Hôpital's rule implies $\lim_{m_i \rightarrow \infty} t_i(m) = -\xi_i(\infty, m_{-i})$. \square

Extending the message space to \overline{M}_i does not change the profit guarantee, as we now see:

Lemma 9. *For any mechanism $(\bar{q}, \bar{t}) \in \mathbf{M}$, under any equilibrium of any type space, the expected profit of (\bar{q}_i, \bar{t}_i) is at least $\overline{\Pi}$.*

Proof. It is easy to see that Lemma 5 continues to hold with the following dual constraints: for all $v \in V$ and $m \in \overline{M}$,

$$\sum_{i=1}^N \left(t_i(m) + \left(v \frac{\partial q_i}{\partial m_i}(m) - \frac{\partial t_i}{\partial m_i}(m) \right) \cdot \mathbb{I}_{m_i < \infty} \right) - cQ(m) \geq \lambda(v).$$

Let $\lambda(v) = \bar{\lambda}(v)$ for v in the support of H , and we let $\lambda = -\infty$ outside the support (this does not affect the objective and essentially drops all of the dual constraints for values outside of the support). When $m \in M$, the dual constraint is the same as before and

¹⁸Presumably there are many other solutions to the system. Bergemann, Brooks, and Morris (2016a) construct excess growth functions that do not coincide with our linear sharing rule for a special case.

holds for all v by our assumption in (34). For a message profile $m \in \overline{M} \setminus M$, the dual constraint is

$$\sum_{i=1}^N \left(v \frac{\partial q_i(m)}{\partial m_i} \mathbb{I}_{m_i < \infty} - \xi_i(m) \right) - cQ(m) \geq \lambda(v),$$

since $-\xi_i(m) = t_i(m)$ for any i such that $m_i = \infty$ by Proposition 4. The above dual constraint is satisfied because we can take a sequence $\{m^n\} \subset M$ that converges to m . As $n \rightarrow \infty$, the left-hand side of the dual constraint for m^n converges to the left-hand side of the dual constraint for m by condition (33) and the assumption that ξ_i is continuous. \square

We say that the type space \mathcal{T} is *product continuous* if the joint distribution $\pi \in \Delta S$ is absolutely continuous with respect to the independent product measure $\times_{i=1}^N \pi_i$, where π_i is the marginal distribution of π over S_i (cf. Milgrom and Weber (1985)). Clearly, our minmax type space $\overline{\mathcal{T}}$ is product continuous, since its signals are independently distributed.

Proposition 5. *Fix a mechanism $\mathcal{M} = (q, t) \in \mathbf{M}$. If a type space \mathcal{T} is product continuous, then there exists a Bayes Nash equilibrium of the game $(\mathcal{M}, \mathcal{T})$. Moreover, expected profit in this equilibrium is at least $\overline{\Pi}$.*

The proof of Proposition 5 is in the Appendix and uses equilibrium existence results from Carbonell-Nicolau and McLean (2017), who extend and refine Reny’s 1999 notion of “payoff security” to an incomplete information setting. We obtain the following theorem:

Theorem 3. *For any mechanism $\overline{\mathcal{M}} = (\overline{q}, \overline{t}) \in \mathbf{M}$, let β be an equilibrium of $\overline{\mathcal{M}}$ under the minmax type space $\overline{\mathcal{T}}$ from Section 4.2. The tuple $(\overline{\mathcal{M}}, \overline{\mathcal{T}}, \beta)$ is a solution to the Seller’s Problem, and the value of the solution is $\overline{\Pi}$.*

A few remarks are in order. Note that the linkage between the allocation and transfer rule in a maxmin mechanism is only through the total excess growth functions. Thus, if two maxmin mechanisms (q, t) and (q', t') have the same total excess growth $\Xi(q) = \Xi(q')$, then (q, t') and (q', t) are both maxmin mechanisms as well. In this limited sense, allocation and transfer rules are interchangeable.

Next, recall that in our sufficient conditions in Section 4.3, we highlighted the fact that the allocation sensitivity should be decreasing, and the constructed solution involved a particular total allocation sensitivity $\mu(m) = \overline{\mu}(\Sigma m)$. The dual conditions (34) are seemingly more permissive and in principle allow the total allocation sensitivity to vary non-monotonically, and even within message profiles that have the same aggregate message. When the support of H is convex, (34) implies that $\mu(\Sigma m) = \overline{\lambda}'(\widehat{w}(\Sigma m))$, but when the support may have gaps, there is additional slack introduced that allows for different solutions. For example, Bergemann, Brooks, and Morris (2016a) studied the present problem when there are only two bidders, zero cost, and the value has support $\{0, 1\}$.¹⁹ They constructed a solution with the same type space (parametrized differently), but a

¹⁹The high value is a normalization, but a low value of zero is a substantive assumption, as it means that there are no gains from trade when the value is low.

different mechanism. Bidders make “demands” in $[0, 1]$ which are then “filled” in a random order. We refer to this as the *Shapley rule*, since it is analogous to the manner in which surplus is allocated within the grand coalition by the Shapley value in cooperative games with transferable utility. In Appendix B, we show that the Shapley rule achieves maxmin profit even if there are many bidders, but still when the value is in $\{0, 1\}$.²⁰

The Shapley rule induces a μ that takes on many different values for message profiles at which $\widehat{w}(\Sigma m) = 1$. This multiplicity makes the multipliers $\bar{\lambda}(v)$ infeasible for $v \in (0, 1)$. But since the support of H is only $\{0, 1\}$, we can change these multipliers without affecting the profit guarantee, and in fact we can make these multipliers minus infinity, so that (34) is slack. A critical feature that the Shapley rule shares with the proportional rule is that both induce a total excess growth that has zero expectation. This allows us to satisfy strong transversality and construct bounded transfers that also satisfy (PIA). However, the fact that the Shapley rule has non-monotonic total allocation sensitivity means that it *cannot* be part of a maxmin mechanism when the support is convex, the reason being that it would be impossible to satisfy all of the dual constraints (34). Moreover, the proportional rule has the normatively desirable property that even if the support of H is non-convex, we can still “fill in” the multipliers $\lambda(v)$ so that they are continuous. This is the key property used in Proposition 7 below that show that the profit guarantee is robust to misspecification of the prior. We also subsequently show that $\overline{\mathcal{M}}$ will extract all of the ex ante gains from trade in the competitive limit when the number of bidders is large, *even if the prior is misspecified*. This proof relies on particular choice of multipliers off of the support of H , and only goes through when the total allocation sensitivity is $\bar{\mu}$.

While Theorem 3 covers a wide range of maxmin mechanisms, there may very well be other maxmin mechanisms that are outside \mathbf{M} . The dual constraints (34), (PIA), and transversality only make sense when we consider standard mechanisms with real messages and smooth allocations and transfers. Thus, we have excluded mechanisms with more complex messaging protocols, such as when the Seller has the bidders report the type space itself. As Börgers (2017) has pointed out, the Seller could also embellish the mechanism with additional features and messages that are not used at the minmax type space, but could generate higher profit in other type spaces. Our subjective view is that the maxmin mechanisms we consider are simpler than these alternatives, in that they strip away everything but those essential features for attaining the optimal profit guarantee, and they do so while maintaining a natural “bidding” interface with the auction participants.

The fact that there could be other solutions not characterized by our theorem raises the concern that these other solutions could have different values as well. While we do not know of any solution with a value that is not $\bar{\Pi}$, any such solution would have to be quite exotic. This follows from the fact that the mechanism and type space we constructed

²⁰With two players, this rule is also referred to as the “contested garment rule,” and is discussed in Aumann and Maschler (1985), and it is implied by a number of cooperative solution concepts that coincide for two players but diverge for $N > 2$ (Thomson, 1994). Bergemann et al. (2016a) correctly conjectured, based on numerical evidence, that the Shapley rule would generalize to many bidders, but it seems that other generalizations, such as the Talmudic rule of Aumann and Maschler, would not work more generally.

satisfy fairly permissive existence theorems. In particular, consider an alternative solution $(\mathcal{M}, \mathcal{T}, \beta)$, that has a value Π . Proposition 2 implies that any equilibrium β' of the game $(\mathcal{M}, \overline{\mathcal{T}})$ must have expected profit less than $\overline{\Pi}$. Thus, as long as an equilibrium of $(\mathcal{M}, \overline{\mathcal{T}})$ exists, Π must be weakly less than $\overline{\Pi}$, else Nature could deviate to $(\overline{\mathcal{T}}, \beta')$ and thereby decrease profit. Similarly, Proposition 1 implies that as long as $(\overline{\mathcal{M}}, \mathcal{T})$ has an equilibrium β'' , then Π must be weakly greater than $\overline{\Pi}$, else the Seller could deviate to $(\overline{\mathcal{M}}, \beta'')$ and thereby increase profit. As Proposition 5 shows, an equilibrium will exist under fairly permissive conditions.

We can formalize this as follows. Let us say that the mechanism \mathcal{M} is *regular* if the game $(\mathcal{M}, \overline{\mathcal{T}})$ has an equilibrium. Similarly, the type space \mathcal{T} is *regular* if $(\overline{\mathcal{M}}, \mathcal{T})$ has an equilibrium. A solution $(\mathcal{M}, \mathcal{T}, \beta)$ is *regular* if both \mathcal{M} and \mathcal{T} are regular. The following theorem then follows immediately from the preceding argument and Theorem 1:

Theorem 4 (Uniqueness). *Every regular solution has a value $\overline{\Pi}$.*

Two corollaries of this theorem assert that $\overline{\mathcal{M}}$ and $\overline{\mathcal{T}}$ can legitimately be considered a maxmin mechanism and a minmax type space, respectively.

Corollary 1 (Maxmin mechanism). *$\overline{\mathcal{M}}$ maximizes minimum profit, where the maximum is taken across regular mechanisms, and the minimum is taken across regular type spaces and Bayes Nash equilibria.*

Corollary 2 (Minmax type space). *$\overline{\mathcal{T}}$ minimizes maximum profit, where the minimum is taken across regular type spaces, and the maximum is taken across regular mechanisms and Bayes Nash equilibria.*

Proposition 5 shows that regularity is quite a permissive condition on type spaces, and indeed, a type space will be regular whenever it is product continuous. We believe it is similarly quite permissive with regard to mechanisms. The type space $\overline{\mathcal{T}}$ has i.i.d. signals that are one-dimensional, and a value that is a monotonic, symmetric, and differentiable function of the signals. While the signal space is not compact, it is easy to compactify the signal space by adding an infinite signal that has zero probability, so that S_i is the extended non-negative reals. Thus, $\overline{\mathcal{T}}$ is about as well-behaved as one could possibly hope for in an uncountably infinite and interdependent value setting, and it satisfies fairly permissive existence theorems. For example, Milgrom and Weber (1985) have shown that an equilibrium will exist as long as the allocation rule and transfer rule are continuous and have compact domains. For then, the family of payoff functions indexed by the signal profile will be equicontinuous (since the value is bounded). This type space is “affiliated” in the sense of Milgrom and Weber (1982), so an equilibrium will exist for a wide range of mechanisms with discontinuous allocations and transfers, e.g., first-price, second-price, and English auctions, with reserve prices or entry fees. Also, Carbonell-Nicolau and McLean (2017) have shown that all-pay auctions and wars of attrition will have (possibly mixed) equilibria.

5 Maxmin auctions in the large and other results

5.1 Profit comparison

In this section, we will further explore the properties of the maxmin auction and the optimal profit guarantee. Let us begin by comparing the maxmin mechanism to various other mechanisms for the uniform distribution example discussed in Section 3 with $c = 0$ and as we vary the number of bidders. It turns out that the uniform distribution is single crossing for all N , so that the maxmin value function \bar{w} has an exponential shaped $k^* \exp(x)$ on $[0, x^*]$, and is fully revealing above x^* . The fully-revealing value is simply $\hat{w}(x) = G_N(x)$. Thus, for these two pieces to meet smoothly, it must be that $k^* = \exp(-x^*)G_N(x^*)$. Moreover, for the mean-preserving spread constraint to bind at x^* , it must be that

$$\begin{aligned} \int_{v=0}^{\bar{w}(x^*)} H(v)dv &= \frac{(G_N(x^*))^2}{2} \\ &= \exp(-x^*)G_N(x^*) \int_{x=0}^{x^*} \exp(x)g_N(x)dx \\ &= G_N(x^*)g_{N+1}(x^*). \end{aligned}$$

Thus, the cutoff x^* is the (unique) solution to

$$G_N(x^*) = 2g_{N+1}(x^*).$$

Maxmin profit in the can-keep case is therefore

$$\begin{aligned} \bar{\Pi} &= \int_{x=0}^{x^*} k^* \exp(x)g_{N-1}(x)dx + \int_{x=x^*}^{\infty} G_N(x)g_{N-1}(x)dx \\ &= G_N(x^*)g_N(x^*) + \int_{x=x^*}^{\infty} G_N(x)g_{N-1}(x)dx, \end{aligned}$$

while in the must-sell case, it is only

$$\hat{\Pi} = \int_{x=0}^{\infty} G_N(x)g_{N-1}(x)dx.$$

In Figure 9, we have plotted these optimal guarantees for N ranging from 1 up to 30. The can-keep guarantee is depicted in blue stars, while the must-sell guarantee is in red circles.

For comparison, we have also plotted profit guarantees associated with three other mechanisms. First, the profit guarantee of Du's (2016) mechanism with parameters tuned to the uniform distribution is plotted in magenta crosses. As Du showed, these mechanisms are maxmin optimal when $N = 1$, thus coinciding with the profit guarantee of our family of mechanisms, although they diverge for $N > 1$. Second is the profit guarantee of the first-price auction, as computed by Bergemann, Brooks, and Morris (2017), which turns out to be $(N-1)/(4N-2)$ for the standard uniform distribution. Finally, we have plotted

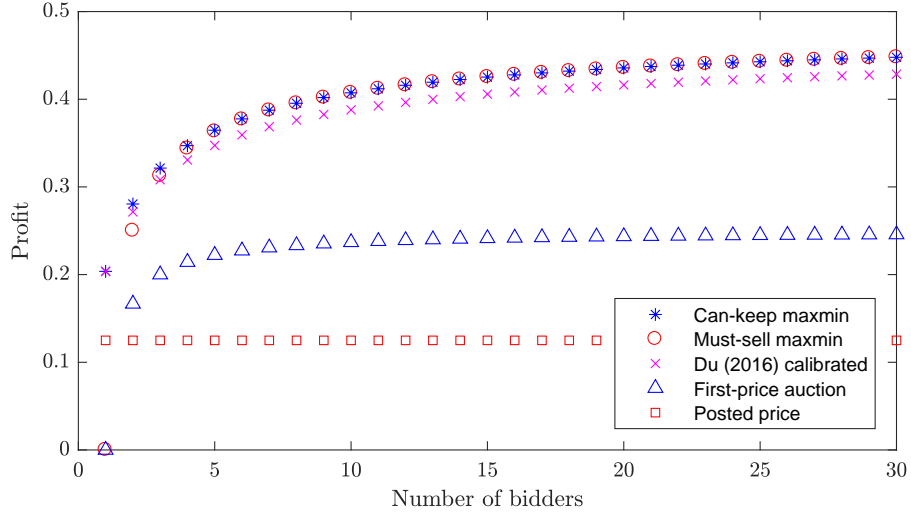


Figure 9: Comparing the maxmin mechanism to other auctions.

the best guarantee from a posted price mechanism, which is obtained by setting a price of $1/4$, which guarantees profit at least $1/8$.²¹

There are a number of striking features of this picture that we explore in the rest of this section. First, note that the profit guarantee increases in N and appears to be converging towards 0.5 . The latter is the ex ante expected value, which is obviously an upper bound on profit in any mechanism. In fact, as N goes to infinity, the profit guarantee converges to the expected surplus. This remarkable fact is actually implied by the earlier result of Du (2016), who constructed a particular sequence of mechanisms for which the profit guarantee converges to total surplus, which we can see as the magenta crosses, corresponding to Du’s mechanisms, also converge to full surplus. We extend this result by characterizing the optimal rate at which the bound can be obtained, and also by characterizing the optimal limit guarantee when the Seller has a positive cost. We will argue that the optimal profit guarantee can be attained even if the Seller misspecifies the prior distribution of values. Another striking feature of this picture is how close together are the optimal can-keep and must-sell guarantees. We will argue that the optimal must-sell profit guarantee is always sandwiched between the optimal can-keep guarantees for N and $N+1$ bidders, and in fact they eventually coincide under a wide range of environments. As we shall see, the optimal informationally-robust auctions simplify greatly in the many bidder limit, and have an intuitive structure whereby allocations are determined by the simple proportional rules corresponding to single-crossing distributions, and bidders are charged a simple price per unit bid that depends on simple statistics from the realized message profile.

²¹The worst-case information is when bidders learn publicly whether the value is below $1/2$. If it is below $1/2$, none of them buy, and if it is above, they strictly prefer to purchase.

5.2 Information and welfare in the competitive limit

Let us proceed to characterize welfare outcomes in the competitive limit as the number of bidders becomes large. There is a substantial literature on this subject, going back to Wilson (1977) and Milgrom (1979). Those papers showed that in the mineral rights model (pure common values with conditionally independent signals), the equilibrium high bid in a first-price auction will converge to the value, so that information is aggregated in the price and bidders compete away all of their rents. This result was generalized by Bali and Jackson (2002), who characterized a class of auctions that similarly extract all the surplus in the many-bidder limit, subject to smoothness restrictions on the bidders' information. It is well known that there are type spaces violating these conditions under which limiting behavior in first-price auctions is not competitive, e.g., the proprietary information model of Engelbrecht-Wiggans, Milgrom, and Weber (1983). Bergemann, Brooks, and Morris (2017) characterized exactly the minimum limit profit of the first-price auction, where the minimum is taken across type spaces and equilibria, and showed that it is generally strictly less than full surplus.

These papers left open the question of whether or not there are mechanisms that always result in a competitive outcome when the number of bidders becomes large, regardless of the form of the information. This question was answered in the affirmative by Du (2016), who constructed a sequence of mechanisms and associated profit lower bounds that converge to the expected value as the number of bidders becomes large.

We have several new results to contribute to this literature. First, we characterize the limit worst-case type space. We then show that the optimal profit guarantee converges to the *ex ante gains from trade*:

$$\int_{v=0}^{v_{\max}} vH(dv) - c. \quad (36)$$

This generalizes Du's result, which implicitly assumed common knowledge of positive gains from trade. Note that the limit profit guarantee is less than that obtained with first-price auctions under the mineral rights model. In that setting, the high bid converges to the value, so that trade would be ex post efficient in the many-bidder limit, and surplus would converge to

$$\int_{v=0}^{v_{\max}} \max\{v - c, 0\}H(dv).$$

In the informationally-robust model, however, neither the Seller nor the bidders is guaranteed to learn whether trade is ex post efficient. For example, one possible type space is no information, in which case the Seller would obtain at most (36) independent of N . Thus, (36) is always an upper bound on $\bar{\Pi}_N(H)$.

In the process of characterizing this guarantee, we show that the optimal profit guarantee converges at a rate of $1/\sqrt{N}$, which is strictly faster than the profit guarantees in Du (2016).²² Moreover, we show that the optimal guarantee and rate is the same for both

²²An important subtlety is that while we know our guarantees are tight, it is possible that Du's mechanisms have profit that is significantly higher than the guarantees that he proves, so that the true gap between the mechanisms is smaller.

the can-keep and must-sell models. Thus, the ability to ration the good does not benefit the Seller in the competitive limit.

We now proceed more formally. To study welfare under the limiting minmax type spaces, it is necessary to normalize the bidders' signals so that the distribution of the aggregate signal converges. Since the signals are i.i.d. exponential random variables, we use the Central Limit Theorem (CLT) to normalize them so that the distribution of the aggregate message converges to a standard Normal. The normalized signals are:

$$s_i^C = \frac{s_i - 1}{\sqrt{N}}.$$

Thus, the sum is

$$x^C = \frac{1}{\sqrt{N}} \left(\sum_{i=1}^N s_i - N \right).$$

(Throughout, a C superscript will denote the CLT normalized signals, although we will drop the C for arguments of functions to prevent notation from getting out of hand.) The distribution of the normalized sum is

$$G_N^C(x) = G_N(\sqrt{N}x + N),$$

with density

$$g_N^C(x) = \sqrt{N}g_N(\sqrt{N}x + N).$$

Lemma 10 (Central Limit Theorem). *G_N^C converges pointwise to Φ , the cumulative distribution of a Normal random variable with zero mean and unit variance. Moreover, g_N^C converges pointwise to ϕ , the density of the standard Normal.*

Proof of Lemma 10. The convergence in distribution of x^C to a standard Normal is a consequence of the Lindberg-Lévy Theorem. This theorem does not imply that the densities converge. However, the normalized Erlang distributions g_N^C satisfy the sufficient conditions of Boos (1985) for convergence in distribution to imply pointwise convergence of densities, namely equicontinuity and boundedness of the densities. Indeed, the density of the Erlang is $g_N(x) = x^{N-1} \exp(-x) / ((N-1)!)$, which is maximized at $x = N-1$, and the maximum value converges to zero. Also, equicontinuity follows from the fact that the density goes to zero pointwise. \square

Let us denote by $\hat{\gamma}_N$ the fully-revealing gains function when the signals are i.i.d. exponential, and let

$$\hat{\gamma}_N^C(x) = \hat{\gamma}_N(\sqrt{N}x + N)$$

denote the normalized fully-revealing gains function. It is straightforward to verify that as N goes to infinity, $\hat{\gamma}_N^C$ converges pointwise to

$$\hat{\gamma}_\infty^C(x) = \min\{v - c \mid H(v) \geq \Phi(x)\},$$

i.e., the function that matches quantiles in the gains-from-trade distribution with their corresponding quantiles in the standard Normal.

Similarly, we can denote by

$$\bar{\gamma}_N^C(x) = \bar{\gamma}_N(\sqrt{N}x + N)$$

the normalized graded gains function. The limit of this object is more subtle, but as the following proposition shows, it must converge to the following function. Let

$$x^* = \min \left\{ x \mid \int_{x'=-\infty}^x \hat{\gamma}_\infty^C(x') \phi(x') dx' \geq 0 \right\}.$$

This is the smallest limit aggregate message x such that the expected gains from trade conditional on an aggregate message being less than x is non-negative. Let

$$\bar{\gamma}_\infty^C(x) = \begin{cases} 0 & \text{if } x < x^*; \\ \hat{\gamma}_\infty^C(x) & \text{if } x \geq x^*. \end{cases}$$

Lemma 11 (Limiting information). *As N tends to infinity, the must-sell normalized min-max gains function $\hat{\gamma}_N^C$ converges pointwise to $\hat{\gamma}_\infty^C$, and the can-keep normalized minmax gains function $\bar{\gamma}_N^C$ converges pointwise to $\bar{\gamma}_\infty^C$.*

The proof of the convergence for the must-sell model is almost immediate, but the proof for the can-keep model is non-trivial. We can, however, give a straightforward intuition from the profit equivalence formula. When we renormalize messages, we change the constant hazard rate of signals to $1/\sqrt{N}$, so that the fully-revealing virtual-value measure becomes

$$\hat{\gamma}_N^C(x) dx - \frac{1}{\sqrt{N}} \hat{\gamma}_N^C(dx).$$

Now, the grading procedure is designed to modify this virtual value so that it is everywhere non-negative. But as N grows large, the normalized fully-revealing gains function converges to the bounded limit $\hat{\gamma}_\infty^C$, but the inverse hazard rate shrinks to zero. This suggests that the fully-revealing virtual value will eventually be positive wherever $\hat{\gamma}_N^C(x) > 0$, and thus the gains function will not need to be graded at all. This intuition is sometimes complete, but it does not apply when there are jumps in the gains function, or points where the slope becomes infinite, both of which are allowed in our model. If there are jumps, then even when N is large, the gains function will be graded around the discontinuity, but the graded region shrinks as N grows and asymptotically disappears.

As an almost immediate consequence, we obtain a characterization of the optimal profit guarantees when the number of bidders becomes large:

Proposition 6 (Limiting optimal profit guarantee). *In the limit as N goes to infinity, both the can-keep and must-sell optimal profit guarantees converge to the ex ante gains from trade at a rate of $1/\sqrt{N}$.*

Proof of Proposition 6. We can rewrite the optimal must-sell profit guarantee in terms of the normalized signals as

$$\begin{aligned} & \int_{x=-\sqrt{N}}^{\infty} g_N^C(x) \hat{\gamma}_N^C(x) dx - \frac{1}{\sqrt{N}} \int_{x=-\sqrt{N}}^{\infty} g_N^C(x) \hat{\gamma}_N^C(dx) \\ &= \int_{v=0}^{v_{\max}} v H(dv) - c - \frac{1}{\sqrt{N}} \int_{x=-\sqrt{N}}^{\infty} g_N^C(x) \hat{\gamma}_N^C(dx). \end{aligned}$$

From the Dominated Convergence Theorem, the second integral converges to

$$\int_{x=-\infty}^{\infty} \phi(x) \hat{\gamma}_{\infty}^C(dx),$$

which is a positive finite quantity, since both the standard Normal density and the outer measure induced by the gains function are bounded. Thus, the profit lost to information rents goes to zero at a rate of $1/\sqrt{N}$.

The proof for the can-keep case is identical, where we replace the fully-revealing gains function with the graded gains function $\bar{\gamma}_N^C$, which according to Lemma 11 converges to a bounded monotonic function $\bar{\gamma}_{\infty}^C$, so that again the information rent with large N is approximately

$$\frac{1}{\sqrt{N}} \int_{x=-\infty}^{\infty} \phi(x) \bar{\gamma}_N^C(dx),$$

where the integral is a finite but strictly positive quantity. \square

5.3 The limit mechanisms

The previous subsection characterizes information and welfare in the competitive limit. Arguably even more important is to understand what happens to the mechanisms themselves. What happens to allocations and transfers as we approach the competitive limit? And what is it about the structure of these limit mechanisms that guarantees the Seller profit equal to the ex ante gains? We now explore this issue, using a heuristic and informal approximation of the limit mechanisms.

We can start to answer these questions using the central limit theorem (CLT) normalization used in the previous subsection to examine limit welfare. Under this normalization, we have already shown that the limiting gains function has a simple form, which is graded below some cutoff $x^* \leq 0$, and fully-revealing above. This means that the total supply has the same shape as that described for single-crossing distributions, namely linearly increasing below x^* and constant and equal to 1 above x^* . Conditional on this total supply, the allocation is proportional to messages. A subtlety is that even though there might be rationing in the limit, under this normalization, the rationing appears to go to zero. The reason is that when there are N bidders, the CLT-normalized aggregate message has support $[-\sqrt{N}, \infty]$, so that when N is large, the supply must linearly interpolate between 0 and 1 on $[-\sqrt{N}, x^*]$, meaning that the formula for supply as a function of the aggregate message is close to $(x + \sqrt{N})/(x^* + \sqrt{N})$, so that supply converges to 1 pointwise.

We can also study the limiting transfers using the CLT normalization. Some care needs to be taken in translating the optimal auction into CLT-normalized units. This is done rigorously in Appendix C. One way to implement bidder i 's transfer would be to first pick a random subset I of the bidders other than i , and then charge bidder i the quantity

$$\bar{t}_i^C(m, I) = \frac{1}{\sqrt{N}} \int_{x=0}^{\infty} \left(\bar{\Xi}_N^C(x + \Sigma m_I) - \bar{\Xi}_N^C(x + \Sigma m_I + m_i) \right) g_{N-|I|}^C(x) dx,$$

where g^C and $\bar{\Xi}_N^C$ are the CLT-normalized counterparts of g and $\bar{\Xi}$. Note that the message profile in the argument of \bar{t}_i^C is also in CLT-normalized units. The quantity x can be interpreted as a random imputation for the aggregate message among bidders not in I .

Let us study how this transfer behaves as the number of bidders becomes large. In particular, consider a sequence of CLT-normalized message profiles and a sequence of subsets of bidders I with $|I| = \delta/N$ such that the normalized message Σm_I converges to some α . We will also assume that $m_i \rightarrow 0$, so that bidder i 's message is a negligible fraction of the aggregate message. (This is equivalent to assuming that the unnormalized message is bounded.) We will compute the limiting object

$$p^C(\alpha, \delta) = \lim_{N \rightarrow \infty} \frac{N \bar{t}_i^C(m, I)}{m_i}.$$

This is the transfer per unit message, and scaled up by N so that it is in per capita terms. By dividing by m_i , the terms inside the integral become the derivative of the total excess growth, which reduces to:

$$\frac{d}{dx} \bar{\Xi}_N^C(x) = \frac{d}{dx} \bar{\mu}_N^C(x) \widehat{w}_N^C(x) - c \frac{d}{dx} \bar{Q}_N^C(x).$$

When the good is fully supplied, only the first term is non-zero and μ' converges to 1, so that when N is large, the derivative reduces to $-\widehat{w}_\infty(x)$. If the good is linearly rationed on the low region, the total allocation sensitivity is constant, and only the second term is non-zero. But the derivative of total supply converges to 1, and so $\sqrt{N} \frac{d}{dx} \bar{\Xi}_N^C(x) \approx -c$. Thus, in both cases, we find that $-\sqrt{N} \frac{d}{dx} \bar{\Xi}_N^C(x)$ converges to $\bar{w}_\infty^C(x)$, the graded value function in the limit, and hence

$$p^C(\alpha, \delta) = \int_{x=0}^{\infty} \bar{w}_\infty^C(\delta x + \alpha) \phi(x) dx,$$

where ϕ is still the density of the normal distribution (the limit of g_n^C). In words, when the number of bidders is large, the bidders essentially pay a price per unit message. This price is an estimated value of the good, based on the aggregate message in a subsample plus randomly imputed bids for the rest of the population.

The CLT normalization blows up signals so that the uncertainty about the aggregate message settles down to a manageable distribution. This is like looking through a "microscope" at what happens to the maxmin mechanism right around the average message

under the minmax type space, which is where much of the delicate structure of the mechanism lies. At the same time, why should we expect messages to be in this range on other type spaces? To understand what might happen at other message profiles far from the center, we can take a “wide angle” perspective on the mechanism by instead using the law of large numbers (LLN) normalization:

$$m_i^L = \frac{m_i}{N},$$

(A superscript L will denote the LLN normalization. As before, we drop the superscript for arguments.) The normalized support remains \mathbb{R}_+ , and the average message at the saddle point converges in distribution to 1.

Let us reconsider the supply. Under the CLT normalization, we found that total supply converges to 1 pointwise, regardless of whether there was common knowledge of gains from trade. This is still true when there are gains from trade, in which case the limit gains function is fully revealing and the LLN-normalized supply rule converges to $\bar{Q}_\infty^L(x) = 1$. When the gains from trade might be negative, however, there will be linear rationing when the CLT-normalized message is below x^* , where $\Phi(x^*) \in (0, 1/2)$. In LLN units, this cutoff is converging to 1 (since the probability of Σm^L being less than any $x < 1$ is going to zero). When the gains from trade might be negative, the limiting supply rule is therefore

$$\bar{Q}_\infty^L(x) = \min\{x, 1\}.$$

Thus, the good is still rationed when the aggregate message is significantly below one.

The transfers similarly take on more extreme values than what were possible under the CLT normalization. As before, we can consider a sequence of message profiles where Σm_I^L converges to α and where $|I| = \delta N$. (Note that this is *not* the same limit as before. Under these assumptions, the CLT-normalized aggregate message in I must converge to $-\infty$ or ∞ , unless $\alpha = 1$.) Again assume that m_i^L goes to zero.

We can similarly define the transfer per unit message and per capita $p^L(\alpha, \delta)$. This price will be an average of expected graded values, where we take the expectation with respect to the realized subsample and imputed values for the other bidders. The difference is that now the imputed values will be extremely close to 1, by the law of large numbers. If the report in the subsample is also close to 1, then the CLT approximation is very good. But if it is significantly different from 1, then the expectation of the value will be heavily skewed towards one tail or the other. For example, if in the large N limit the average value in the realized subsample $\Sigma m_I^L/|I|$ is significantly above 1, then in CLT-normalized units, this would be the same as if the message were blowing up, in which case the expected value (and hence the price per unit message) converges to

$$p^L(\alpha, \delta) = \bar{w}_\infty^C(\infty) = v_{\max}.$$

On the other hand, if $\Sigma m_I^L/|I|$ is significantly less than 1, the price per unit is converging to

$$p^L(\alpha, \delta) = \bar{w}_\infty^C(-\infty) = c.$$

These transfers essentially force the LLN-normalized aggregate message to be close to 1 with high probability when the number of bidders is large. For if not, the price bidders would be paying for the good would either be very low, in which case they would want to bid more to increase their allocation probability and obtain a good deal, or the price would be very high, in which case they would want to shade their bids to avoid paying too high of a price.

5.4 Robustness to the prior

We have heretofore assumed that the Seller faces large ambiguity about the type space but knows precisely the common prior over the value. This is quite an extreme assumption. It turns out, however, that the profit guarantee is robust to misspecification of the prior, in that the loss from misspecification is a continuous and linear function of the true distribution. Moreover, the result that the optimal profit guarantee converges to the ex ante gains from trade as the number of bidders becomes large remains true *even if the prior is misspecified*.

Let us suppose that the Seller runs the mechanism $\overline{\mathcal{M}}_N(H)$ that provides the optimal profit guarantee when the prior is H . Let $\overline{\Pi}_N(H)$ denote the associated profit guarantee when the prior is H . Let $\overline{\lambda}_N(v; H)$ denote the associated optimal dual multipliers given by (20).

Now suppose that that prior is not H , but rather the value is distributed according to the distribution H' . Note that these multipliers will give us a lower bound on profit via the argument in Lemma 5 for any prior. Thus, even if the prior is H' , the multipliers $\overline{\lambda}_N(v; H)$ are still feasible for the dual program (together with unitary multipliers on local incentive constraints). Indeed, all that changes is the lower bound itself, which is the expectation of $\overline{\lambda}_N(v; H)$ under H' . Continuity of $\overline{\lambda}_N(v; H)$ immediately yields the following result:

Proposition 7 (Profit guarantee for misspecified prior). *In any equilibrium of $\overline{\mathcal{M}}(H)$ for any type space where the marginal distribution of the value is H' , expected profit must be at least*

$$\overline{\Pi}_N(H, H') = \int_{v=0}^{v_{max}} \overline{\lambda}_N(v; H) H'(dv). \quad (37)$$

Thus, expected profit from the mechanism $\overline{\mathcal{M}}(H)$ is bounded below by a continuous and linear function of the true prior that coincides with the optimal guarantee when the prior is H .

Not only is the maxmin mechanism robust to misspecification of the prior at fixed N , but in fact the optimal profit guarantee is always obtained even if the prior is misspecified. This will follow from the following lemma:

Lemma 12 (Limiting multipliers). *In the limit as N goes to infinity, $\overline{\lambda}_N(v; H)$ converges to $v - c$ for v in the convex hull of the support of H .*

Proof of Lemma 12. For all N ,

$$\bar{\lambda}'_N(v; H) = \bar{\mu}_N(G_N^{-1}(H(v))). \quad (38)$$

We claim that this quantity converges to 1 as $N \rightarrow \infty$. There are basically two cases to consider. If $\Phi^{-1}(H(v)) > x^*$, then asymptotically the gains function will be fully-revealing when the gains from trade are $v - c$, and the total allocation sensitivity is $\mu(x) = (N - 1)/x$. Thus, (38) is the inverse of $G_N^{-1}(\alpha)/(N - 1)$, which is the α th quantile of the distribution of x/N . The law of large numbers implies that as N goes to infinity, x/N converges almost surely to 1, meaning that $G_N^{-1}(\alpha)/N$ converges to 1 as well. On the other hand, if $\Phi^{-1}(H(v)) < x^*$, then asymptotically the gains function is graded and $\mu(x) = N/(\sqrt{N}x^* + N)$, which converges to 1 as N goes to infinity. Thus, we conclude that $\bar{\lambda}_\infty(v; H) = v + C$ for some constant C . But since $\bar{\Pi}_N(H)$ converges to the ex ante gains from trade, it must be that the constant is $-c$ and $\bar{\lambda}_N(v; H)$ converges to $v - c$. \square

As a result, Lebesgue's dominated convergence theorem implies that the profit guarantee for $\bar{\mathcal{M}}(H)$ converges to the ex ante gains from trade for all distributions H' . We thus obtain the following proposition:

Proposition 8 (Prior-independent limiting profit guarantee). *Suppose the support of H' lies in the convex hull of the support of H . In the limit as N goes to infinity, $\bar{\Pi}_N(H, H')$ converges to the ex ante gains from trade under H' .*

We note that analogues of Propositions 7 and 8 also hold for the must-sell profit guarantee $\hat{\Pi}_N(H)$. The only difference in the arguments is that we should use the optimal must-sell multipliers $\hat{\lambda}_N$ for Proposition 7 and the optimal must-sell total allocation sensitivity $\hat{\mu}_N(x) = (N - 1)/x$ for Proposition 8.

Interestingly, while the profit guarantee converges to the ex ante gains from trade under the true prior H' , that guarantee need not be positive. Thus, $\bar{\mathcal{M}}(H)$ may not be a maxmin mechanism in the competitive limit, in the event that the optimal profit guarantee is zero and it is better to shut down production entirely.

We conclude this section with a redux of Figure 9, where we now show the profit guarantee associated with the must-sell maxmin mechanisms for a misspecified prior, namely the exponential distribution $H(v) = 1 - \exp(-v)$, whereas the true distribution is standard uniform. Suppose $c = 0$. The fully-revealing gains function is

$$\hat{\gamma}_N(x) = -\log(1 - G_N(x)).$$

Thus, the linear approximation for the must-sell maxmin auction is

$$\hat{\lambda}_N(\hat{\gamma}_N(x)) = \int_{y=0}^{\infty} \left(\frac{G_N(y)}{1 - G_N(y)} - \log(1 - G_N(y)) \right) g_{N-1}(y) dy - \int_{y=x}^{\infty} \frac{g_{N-1}(y)}{1 - G_N(y)} dy.$$

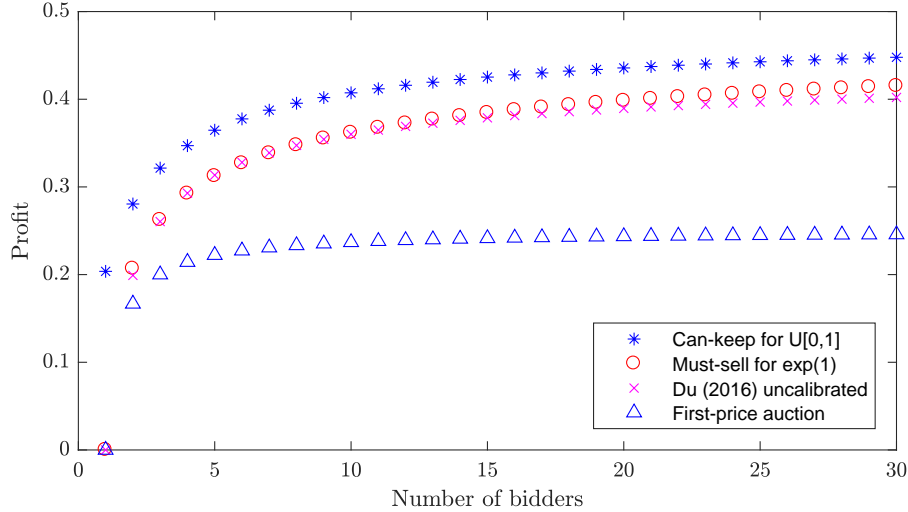


Figure 10: Comparison of “uncalibrated” auctions.

The lower bound for profit when the true distribution is standard uniform is therefore

$$\begin{aligned}
\int_{x=0}^{\hat{\gamma}_N^{-1}(1)} \hat{\lambda}_N(\hat{\gamma}_N(x)) \hat{\gamma}'_N(x) dx &= \int_{x=0}^{\infty} \left(\frac{G_N(x)}{1 - G_N(x)} - \log(1 - G_N(x)) \right) g_{N-1}(x) dx \\
&\quad - \int_{x=\hat{\gamma}_N^{-1}(1)}^{\infty} \frac{g_{N-1}(x)}{1 - G_N(x)} dx \\
&\quad - \int_{x=0}^{\hat{\gamma}_N^{-1}(1)} \frac{g_{N-1}(x)}{1 - G_N(x)} (-\log(1 - G_N(x))) dx
\end{aligned}$$

We have plotted this lower bound in Figure 10, which also retains the optimal can-keep profit guarantee for the standard uniform, as well as a prior-independent mechanism proposed by Du (2016) and the first-price auction. As we can see, the profit guarantee of the must-sell maxmin auction for the exponential distribution is over 90% that of the optimal can-keep profit guarantee when $N \geq 15$.

5.5 Simple mechanisms

In this section we illustrate with two examples that our mechanism can be significantly simplified for some prior H . Our objective in working out these simple mechanisms is twofold: First, we want to illustrate our construction in Section 4 with some concrete examples. Second, these simple mechanisms give good profit guarantees when the prior is close to the assumed prior, and in our second example, the profit guarantee converges to the ex ante gains from trade as $N \rightarrow \infty$ at the optimal rate for *all* priors with the assumed upper bound on value.

5.5.1 Tullock contest

In our first example, we take the must-sell maxmin mechanism from Section 4.5.2. Without loss of generality, suppose $c = 0$. Suppose the prior H_N , which varies with N , satisfies:

$$\widehat{w}_N(x) = H_N^{-1}(G_N(x)) = \frac{x^2}{N(N+1)}. \quad (39)$$

We choose the normalization factor $\frac{1}{N(N+1)}$ so that the expected value is one:

$$\int v H_N(dv) = \int_{x=0}^{\infty} \frac{x^2}{N(N+1)} g_N(x) dx = \int_{x=0}^{\infty} g_{N+2}(x) dx = 1.$$

The allocation rule is

$$\widehat{q}_i(m) = \frac{m_i}{\Sigma m},$$

so the total allocation sensitivity is $\widehat{\mu}_N(x) = \frac{N-1}{x}$. The linear approximation to the maxmin profit is:

$$\begin{aligned} \widehat{\lambda}_N(\widehat{w}_N(x)) &= \int_{y=0}^{\infty} \widehat{w}_N(y) g_{N-1}(y) dy + \int_{y=0}^{\infty} \widehat{\mu}_N(y) G_N(y) \widehat{w}'_N(y) dy - \int_{y=x}^{\infty} \widehat{\mu}_N(y) \widehat{w}'_N(y) dy \\ &= \int_{y=0}^{\infty} \widehat{w}_N(y) g_{N-1}(y) dy - \int_{y=0}^{\infty} \widehat{\mu}_N(y) (1 - G_N(y)) \widehat{w}'_N(y) dy + \int_{y=0}^x \widehat{\mu}_N(y) \widehat{w}'_N(y) dy \\ &= C_N + \frac{2(N-1)}{N(N+1)} x, \end{aligned}$$

where C_N is a constant independent of x .

Therefore, the total excess growth is

$$\widehat{\Xi}_N(x) = \widehat{w}_N(x) \widehat{\mu}_N(x) - \widehat{\lambda}_N(\widehat{w}_N(x)) = -\frac{(N-1)}{N(N+1)} x - C_N,$$

which implies the simple transfer rule:

$$\begin{aligned} \widehat{t}_i(m) &= \frac{1}{N} \sum_{n=1}^N \frac{1}{\binom{N-1}{n-1}} \sum_{\substack{I \subseteq N \setminus \{i\} \\ |I|=n-1}} \int_{x=0}^{\infty} \left(\widehat{\Xi}(\Sigma m_I + x) - \widehat{\Xi}(m_i + \Sigma m_I + x) \right) g_{N-n+1}(x) dx \\ &= \frac{(N-1)}{N(N+1)} m_i. \end{aligned}$$

The mechanism $(\widehat{q}, \widehat{t})$ is known as the Tullock contest (Tullock, 1980). Our conclusion is thus that the Tullock contest is a must-sell maxmin mechanism for the prior distribution in (39).

By solving for x in $\widehat{w}_N(x) = v$, we see that

$$\widehat{\lambda}_N(v) = C_N + 2(N-1) \sqrt{\frac{v}{N(N+1)}}.$$

Thus, unlike in the previous section, we cannot conclude that the profit guarantee of the Tullock contest under an alternative prior H' converges to the ex ante gains from trade under H' as $N \rightarrow \infty$. The difference is due to the fact that the prior H_N that gives Tullock contest varies with N ; in fact, as N increases H_N becomes more and more concentrated on $v = 1$.

5.5.2 A simple, asymptotically optimal mechanism

In our second example, suppose the upper bound on the value $v_{\max} < \infty$, and take

$$q_i(m) = \frac{m_i}{\max\{\Sigma m, N\}} \quad (40)$$

as the allocation rule. The total allocation sensitivity and the total supply is then

$$\mu_N(x) = \begin{cases} 1 & x \leq N \\ \frac{N-1}{x} & x > N \end{cases}, \quad Q_N(x) = \begin{cases} \frac{x}{N} & x \leq N \\ 1 & x > N \end{cases}.$$

Suppose H places probability 1 on v_{\max} . Consider the linear approximation

$$\lambda_N(v) = \int_{x=0}^{\infty} (v_{\max} \mu_N(x) - c Q_N(x)) g_N(x) dx - (v_{\max} - v), \quad v \in [0, v_{\max}].$$

This linear approximation comes from equation (20) in Proposition 2, where we take $\Pi = \int_{x=0}^{\infty} (v_{\max} \mu_N(x) - c Q_N(x)) dx$, and use the facts that $\widehat{w}_N(x) = v_{\max}$ and $\mu_N(G_N^{-1}(H(v))) = 1$ for all $v < v_{\max}$. Note that Π is not the optimal profit guarantee under H , as q_i is not the maxmin allocation rule for H ; we choose this Π so that the total excess growth from the profit-incentive alignment

$$\Xi_N(x) = v_{\max} \mu_N(x) - c Q_N(x) - \lambda(v_{\max})$$

has zero expectation under $g_N(x)$.

Our transfer rule is then:

$$t_i(m) = \frac{1}{N} \sum_{n=1}^N \frac{1}{\binom{N-1}{n-1}} \sum_{\substack{I \subseteq N \setminus \{i\} \\ |I|=n-1}} \int_{x=0}^{\infty} (v_{\max} (\mu_N(\Sigma m_I + x) - \mu_N(m_i + \Sigma m_I + x)) \\ - c(Q_N(\Sigma m_I + x) - Q_N(m_i + \Sigma m_I + x))) g_{N-n+1}(x) dx. \quad (41)$$

Intuitively, bidder i pays the marginal decrease that his message has on the total allocation sensitivity and also pays his marginal increase on the total supply, where these marginal payments are averaged over all possible orders of bidder arrivals and are conditioned on the messages of bidders who arrive before i .

By Proposition 2, or simply by a direct verification, we see that the dual constraint

$$\lambda_N(v) \leq v \mu_N(\Sigma m) - \Xi_N(\Sigma m) - c Q_N(\Sigma m)$$

is satisfied for every $v \in [0, v_{\max}]$ and every $m \in M$. Therefore, by Lemma 5, for any prior H' that places zero probability on $v > v_{\max}$, the profit guarantee of the mechanism (q, t) is at least $\int \lambda_N(v) H'(dv)$. Using the central limit theorem, it is easy to see that as $N \rightarrow \infty$,

$$\max_{v \in [0, v_{\max}]} |\lambda_N(v) - (v - c)| = O(1/\sqrt{N}).$$

Thus, we get:

Proposition 9. *Suppose the allocation and transfer rules (q, t) are given by (40) and (41). Then for any prior H' that places zero probability on $v > v_{\max}$, the profit guarantee of (q, t) converges as $N \rightarrow \infty$ to the ex ante gains from trade under H' at the rate of $O(1/\sqrt{N})$.*

5.6 The value of commitment

In addition to the limiting behavior of the optimal profit guarantees, another remarkable feature of Figure 9 is that the profit guarantees of the must-sell and can-keep mechanisms are so close together, and practically coincide for $N > 5$. It turns out that both of these phenomena are general features of the model *when there is common knowledge of gains from trade*, i.e., when $H(v) = 0$ for all $v < c$.

Proposition 10 (Profit bounds). *If there is common knowledge of gains from trade, then $\widehat{\Pi}_N(H) \leq \overline{\Pi}_N(H) \leq \widehat{\Pi}_{N+1}(H)$.*

The first inequality is obvious because must-sell mechanisms are in the feasible set of can-keep mechanisms. The formal proof of the second inequality is in the Appendix, but the argument is a relatively straightforward adaptation of the arguments in Bulow and Klemperer (1996). Basically, any can-keep maxmin mechanism with N bidders can be used to construct a must-sell mechanism with $N + 1$ bidders that has the same profit guarantee: run the same mechanism among the first N bidders, and if the good does not get allocated, give it to bidder $N + 1$ at a price equal to the production cost. This offer would be accepted because of common knowledge of gains from trade, and it is obviously profit neutral. For any type space and equilibrium of the $N + 1$ -bidder auction just described, there is a corresponding type space and equilibrium of the maxmin auction with N bidders that has the same profit, thus proving that minimum profit with $N + 1$ bidders across type spaces and equilibria must be at least the optimal can-keep profit guarantee with N bidders.

We note that the interpretation of this result is somewhat different than that of the analogous result in Bulow and Klemperer (1996). They study the independent private value model with symmetric and regular value distributions, and show that the optimal must-sell mechanism (e.g., a second-price auction) with $N + 1$ bidders generates more profit than the optimal can-keep mechanism (a second-price auction with reserve price) with N bidders. Bulow and Klemperer interpret this result as a justification for using a “simple” auction format versus one that depends on the precise distribution, in order to set the reserve price. In our setting, both the can-keep and must-sell maxmin mechanisms




Figure 11: Profit guarantees as a function of the cost when $v \sim U[0, 1]$.

require knowledge of the value distribution to calibrate the mechanism (although the optimal must-sell allocation is prior independent). We therefore interpret this result more as a statement about the value of the Seller’s ability to commit to keeping the good, or alternatively, about the value of investing in flexible production technology that allows output to be decided after the revelation of private information. If there is common knowledge of gains from trade, then the value of such commitment power is smaller than the incremental profit from recruiting another bidder in the auction.

This result depends crucially on the assumption of common knowledge of gains from trade. If the gains from trade could be negative, the profit lost from having to sell the good could be significant, even with an additional bidder. To see this, consider what happens as the cost converges to the ex ante value. The optimal can-keep profit guarantee must be decreasing in the cost, and eventually it converges to zero when the ex ante gains from trade disappear, regardless of the number of bidders. The optimal must-sell profit guarantee is simply the optimal revenue guarantee minus the cost of producing a single unit, which is linearly decreasing in the cost. Since the optimal revenue guarantee is strictly less than the expected value, for sufficiently high costs, expected profit will become negative. This shows that for any $N' > N$, there exists a largest cost at which both the must-sell profit guarantee with N' bidders and the can-keep profit guarantee with N bidders coincide, and for any lower cost, the must-sell guarantee is higher, and for any higher cost, the can-keep profit guarantee is higher. Figure 11 illustrates this for the standard uniform distribution. Thus, when costs are high, flexible production as a commitment device could be much more valuable than more bidders, and a firm that is motivated by worst-case criterion would rather shut down than operate without flexible production.

Even in light of Proposition 10, it is still remarkable in Figure 9 how quickly the must-sell and can-keep profit guarantees coincide, even while both guarantees are increasing

fairly quickly in N . The following proposition addresses why this occurs (although strictly speaking it does not cover the uniform distribution).

Proposition 11 (Exact optimality of must-sell). *Suppose the value distribution has convex support $[\underline{v}, v_{max}]$ with $\underline{v} > c$, and that the distribution is absolutely continuous with a density h that is bounded away from zero. Then there exists an \underline{N} such that for all $N \geq \underline{N}$, the maxmin can-keep mechanism is a must-sell mechanism, i.e., $\widehat{\Pi}_N(H) = \overline{\Pi}_N(H)$.*

Proof of Proposition 11. The idea is closely related to Lemma 11, which showed that the pointwise limit of the graded gains function is fully revealing when there is common knowledge of gain from trade. When the distribution is absolutely continuous with convex support and a bounded derivative, then $\widehat{\gamma}_N^C$ is differentiable, and the fully-revealing virtual value under the normalized signals when $\widehat{\gamma}_N^C(x) = v - c$ is

$$\widehat{\gamma}_N^C(x) - \frac{1}{\sqrt{N}} \frac{d}{dx} \widehat{\gamma}_N^C(x) = v - c - \frac{1}{\sqrt{N}} \frac{g_N^C((G_N^C)^{-1}(H(v)))}{h(v)},$$

where we have applied the inverse function theorem to obtain the derivative of $\widehat{\gamma}_N^C$. Since h is bounded away from zero, the information rent $g_N(G_N^{-1}(H(v)))/h(v)$ is bounded above by some constant K , and we can take \underline{N} large enough so that $\underline{v} - c - K/\sqrt{\underline{N}} > 0$. For any larger N , the fully-revealing virtual-value measure is positive, so that the the must-sell maxmin mechanism $\widehat{\mathcal{M}}_N$ maximizes revenue among all can-keep mechanisms at the must-sell minmax type space $\widehat{\mathcal{T}}_N$, so that the must-sell solution is also a can-keep solution. \square

We note that the kind of competitive limit we are taking is rather different from that taken by Bulow and Klemperer (1996) or Yamashita (2016). They consider a sequence of economies with independent private values, where each new bidder has a value that is independent of those of the preceding bidders and drawn from the same distribution. In those models, exclusion will generally be optimal even as the number of bidders goes to infinity, and even if there is common knowledge of gains from trade. The reason is that as the economy grows, the total surplus is changing as well, and in fact the bidders receive the same information rent from being allocated the good, regardless of the number of bidders. For example, if the values are standard uniform, then profit is maximized by a second-price auction with reserve price of 1/2, regardless of the number of bidders. In our competitive limit, the common value distribution and total surplus are held fixed. Information about the common value becomes more dispersed in the minmax type space, and hence information rents become smaller, which is why exclusion is eventually not optimal.

5.7 Timing & bidder collusion

As we discussed in the Introduction, an important related literature concerns the optimal design of information by buyers with commitment power. This problem is studied by Roesler and Szentes (2017) for the case where there is a single buyer whose value comes

from a known distribution. The buyer can commit to a procedure for learning about the value which may leave the buyer with only partial information. After committing to this procedure, the Seller commits to a mechanism for selling the good to the buyer, which without loss of generality can be taken to be a posted price (as posted prices are optimal with a single buyer and an indivisible good). Roesler and Szentes characterize the optimal information structure for such a buyer.

Our results can be used to generalize this analysis to many buyers. Suppose that the N bidders can collusively commit to a type space *before* the Seller chooses the auction format. In contrast to the single buyer case, this information determines not just how much information the bidders have about the value, but also their higher-order beliefs. Let us suppose that the bidders must choose a regular type space, in the sense of Section 4.6. The Seller then chooses the auction mechanism. Finally, the bidders choose strategies in the mechanism selected by the Seller.

We claim that the following is an equilibrium of this game. The bidders commit to the type space $\overline{\mathcal{T}}$, after which the Seller sets a posted price equal to $\overline{\Pi}$. In equilibrium, the bidders always purchase the good. To see why this is an equilibrium, observe that no matter what beliefs the bidders choose, the Seller can always use $\overline{\mathcal{M}}$ at the second stage and guarantee himself $\overline{\Pi}$. Moreover, in the proposed equilibrium, the allocation is efficient, so this strategy profile simultaneously maximizes total surplus and minimizes profit, and therefore the bidders' surplus must be maximized. Thus, it is impossible for the bidders to deviate in a way that increases their joint payoff. Finally, setting a posted price of $\overline{\Pi}$ achieves optimal profit for the Seller when the type space is $\overline{\mathcal{T}}$, so that the Seller does not have a profitable deviation either.

This equilibrium is similar to the Roesler and Szentes solution in that the Seller is indifferent across a wide array of mechanisms, and in equilibrium the Seller uses a posted price (although any mechanism that allocates the good with probability one would also work, e.g., first- or second-price auctions without reserve price). The type space that the bidders use is in a sense a generalization of the optimal information with a single buyer, which is such that the buyer's interim expectation of the value has a Pareto distribution on a low region. While there are multiple solutions on the high region, one solution is for the buyer to learn the value exactly. Our solution is analogous, with the Pareto shaped region corresponding to the region where the value function is graded. Indeed, in the special case where $N = 1$, our solution is the same as that of Roesler and Szentes, with the exponential signal distribution simply being a different parametrization of the buyer's information.

5.8 On the set of possible welfare outcomes

As we argued in the introduction, a challenge for auction design with common values is the lack of a canonical type space the same stature as the IPV model when values are private. Our methodology has been to *derive* the type space using a worst-case criterion and then optimize the auction against that benchmark. While we learn a great deal from this exercise, it is still just one possible type space.




Figure 12: Possible profit and bidder surplus pairs when $c = 0$, $N = 2$, and $v \sim U[0, 1]$. Red, purple, and green denote the set of welfare outcomes for particular mechanisms. The blue set depicts the set of welfare outcomes that could obtain in some type space, when the Seller uses some profit-maximizing mechanism for that type space.

For our final topic, we ask what our solution can tell us about welfare outcomes under optimal auctions on other type spaces that are not the worst-case. Let \mathcal{W} denote the set of all pairs (U, Π) of bidder surplus and profit that could arise under some type space, if the Seller uses a profit-maximizing auction and the bidders play a Bayes Nash equilibrium. We will argue that our results tell us a great deal about this set.²³ To fix ideas, we will discuss this issue in the context of the two-bidder standard uniform example that we have used throughout the paper, with a zero cost of production. The subsequent discussion will reference various welfare outcomes that are depicted graphically in Figure 12.

Any welfare outcome $(U, \Pi) \in \mathcal{W}$ must satisfy three elementary bounds. First, our main result tells us that the Seller is guaranteed the maxmin profit, so $\Pi \geq \bar{\Pi}$. Second, it is impossible for total welfare to exceed the efficient gains from trade:

$$U + \Pi \leq \int_{v \in \mathbb{R}_+} v H(dv).$$

²³A similar analysis was conducted by Bergemann, Brooks, and Morris (2015) in the context of monopoly price discrimination, and later extended by Roesler and Szentes (2017).

Finally, the opt-out condition implies that bidders must receive non-negative rents, i.e., $U \geq 0$. These three constraints imply that the set \mathcal{W} lies within a “surplus triangle” which roughly corresponds to the blue set in Figure 12.

A full characterization of \mathcal{W} is beyond our present capabilities. We can, however, report a number of welfare outcomes that are in \mathcal{W} which will then narrow down the possible shapes the set may have. First, we construct a solution to the Seller’s Problem in which the Seller obtains maxmin profit $\bar{\Pi}$, and bidder surplus is

$$\bar{U} = \int_{x=0}^{\infty} \bar{w}(x)\bar{Q}(x)g_N(x)dx - \bar{\Pi}$$

under the truth-telling equilibrium. This is depicted at point A in Figure 12. Note that this outcome is socially inefficient because the Seller rations the good when the value is low. Also, one possible type space is the degenerate one where the bidders have no information, in which case the Seller can extract all of the surplus with a posted price of $1/2$. This is point B.

The minmax information structure $\bar{\mathcal{T}}$ allows us to identify two more points which must be in \mathcal{W} . Note that under the minmax information, the Seller is indifferent between allocating the good and not allocating when the aggregate message is less than $x^* \approx 1.79$. Under the maxmin mechanism, the Seller rations on this region. If the Seller instead breaks indifference in favor of always allocating the good, then the outcome is socially efficient, and we obtain the point C. Since \mathcal{W} must be convex,²⁴ we conclude that \mathcal{W} contains the entire efficient frontier that satisfies the elementary bounds, which is the line between B and C .

Alternatively, if the Seller breaks indifference by not allocating below x^* , we obtain the point D. Note that the bidders must still obtain positive information rents from when the aggregate signal is greater than x^* . We conjecture that it is impossible for bidder surplus to fall below this level, subject to the Seller still obtaining profit of $\bar{\Pi}$, simply because we suspect that the minmax type space is essentially unique. In particular, we conjecture that in general there does not exist a type space and optimal direct mechanism in which welfare is $(0, \bar{\Pi})$.

At the same time, we know it is possible for the bidders to receive zero surplus, as for example at point B. A natural next question is what does the “western” frontier of \mathcal{W} look like, and what is the lowest profit subject to $U = 0$? We can easily construct points below B on this line as follows. Recall the minmax value function \bar{w} , which is depicted in the left panel of Figure 13. The right-panel depicts the corresponding integrated value functions. The fully-revealing value function is in blue, and the minmax value function \bar{w} follows the red exponential curve until x^* , after which it is fully revealing. Because \bar{w} is fully-revealing on $[x^*, \infty)$, the bidders obtain positive information rents. Consider the following modification of $\bar{\mathcal{T}}$. When $x \leq x^*$, the bidders learn as before, so that the value function has the red exponential shape. When $x \geq x^*$, all the bidders learn is that

²⁴Given two type spaces, we could “randomize” between them, and there will be a direct mechanism in which the bidders truthfully report to the Seller which type space they are in.




Figure 13: Value and integrated quantile functions.

$v \geq \bar{w}(x^*)$, so that the interim expected value is constant at

$$\tilde{v} = \frac{1}{1 - G_N(x^*)} \int_{x=x^*}^{\infty} \hat{w}(x) g_N(x) dx.$$

An optimal mechanism on this type space is to allocate the good randomly if and only if $x \geq x^*$ at a price of \tilde{v} , which the buyers would be willing to accept, leaving them zero rents. This corresponds to point E in Figure 12. Note that total surplus is the same as in point D.

Can profit be pushed any lower? Indeed it can. Under the value function that generates point E, the value function jumps up discontinuously at x^* from $\bar{w}(x^*)$ to \tilde{v} . This discontinuity leaves some slack that can be further exploited. Consider a further modification of the value function whereby we extend the exponential shaped region past x^* to a new point \bar{x} , but then give no information above \bar{x} . As \bar{x} increases, the discontinuity shrinks and eventually we obtain a continuous value function. In Figure 13, this critical value function corresponds to the red curve below \bar{x} and the flat black line above \bar{x} . Because this value function induces a larger region on which the Seller is willing to withhold the good, we know that total surplus has decreased, while still maintaining the structure that gives the bidders zero surplus. This corresponds to point F in Figure 12.

Our conjecture is that this is the lowest point on the western frontier, and moreover that points D and F are connected by a class of type spaces in which the value function is exponential below some cutoff x_1 , then is flat until a cutoff x_2 , and then finally is fully-revealing above x_2 . The cutoff x_2 is determined from x_1 from the condition that the value function be continuous at both points. The dashed blue line plots the locus of welfare outcomes generated by such type spaces. Intuitively, these type spaces maximize the information rents of the bidders, subject to a given total supply, by giving the bidders information rents when values are high. Rents are higher on these values because the ratio g_N/g_{N-1} is increasing, so the contribution to social surplus (proportional to g_N) is relatively large compared to the contribution to profit (proportional to g_{N-1}).

Thus, the blue triangular region is a fair approximation of the set of welfare outcomes that can be induced by some optimal mechanism on some type space. In Figure 14, we have plotted this conjectured set for various N . The vertices represent the welfare




Figure 14: Conjectured welfare set for various N .

outcome at the saddle point. As expected, the welfare set shrinks towards the outcome $(0, 0.5)$ at which the Seller would extract all of the ex ante gains from trade.

We can use the welfare set as a heuristic benchmark against which to compare specific mechanisms. For example, consider the posted price mechanism that we discussed in the introduction. Within the class of posted price mechanisms, worst-case profit is maximized for the uniform distribution with a price of $p = 1/4$. Let us ask: what is the set of welfare outcomes that could occur for this *fixed* mechanism, as we vary information and equilibrium? We computed this set numerically,²⁵ and its boundary is depicted in purple in Figure 12. As we can see, the set of profits that can be attained by this mechanism is *strictly* below optimal profits in the strong set order, while the set of bidder surpluses under the posted price are strictly above those for the optimal mechanism. Thus, from a profit-maximization perspective, the posted price is dominated, while we can also conclude that bidders would be unequivocally better off by forcing the Seller to use this mechanism, regardless of the type space.

²⁵Specifically, we computed the set of profit and bidder surplus pairs that could arise in a Bayes correlated equilibrium of the mechanism with a posted price of $1/4$. We used a discrete approximation of the state space with 300 values, and binary actions corresponding to “buy” and “not buy.” In the event that both buyers say “buy,” the winner is decided by a uniform tie break.

We similarly computed the set of possible welfare outcomes across all BCE for the maxmin mechanism, which is depicted in red. We warn the reader that this set is more of an “artist’s impression” than the other objects in Figure 12, for the following reason. We looked at a discretized version of the maxmin mechanism with 30 reports evenly spaced between a message of 0 and 7, with the allocation and transfer being computed according to the formulae in Section 4. We also used a uniform distribution over 15 values evenly spaced between 0 and 1. Due to the discretization, minimum profit for this mechanism is only 0.24, approximately 85% of the continuum limit of about 0.28. Even in this discretized simulation, however, the point B and a significant portion of the efficient frontier was obtained in the welfare set. We have rescaled payoffs so that minimum profit coincides with the theoretical value.²⁶

This computation gives us a clear sense that the maxmin mechanism is not dominated in the same sense as the posted price. The range of possible optimal profits for the maxmin mechanism coincides with the range for the optimal mechanism, and moreover, total surplus is significantly higher than what it would be under many profit-maximizing mechanisms. Nonetheless, it is clear that there will be many type spaces in which the maxmin mechanism misses out on a significant amount of profit, e.g., those type spaces which satisfy the hypotheses of McAfee, McMillan, and Reny (1989) for full-surplus extraction.

Two final comparisons come from the must-sell maxmin mechanism, for which the welfare at the saddle point is at G.²⁷ The set of welfare outcomes for this mechanism is the entire line segment from B to G. Finally, minimum profit in the first-price auction, as computed by Bergemann, Brooks, and Morris (2017), is $1/6$ and corresponds to point H. Again, the set of possible welfare outcomes in the first-price auction is the entire line segment from B to H. While the maxmin auctions seem to have a smaller support for welfare outcomes than the first-price auction, we cannot conclude that maxmin mechanisms dominate the first-price auction as they do the posted price.

6 Conclusion

Let us briefly conclude by taking a critical eye to our results, and commenting on future directions. This paper has studied the canonical auction design problem when values are common. The novelty is to use a robust criterion for measuring the performance of an auction. The spirit of the exercise is to identify mechanism designs that are less vulnerable to misspecification of the environment. We believe that our results shed light on the nature of informationally-robust auctions. Moreover, we have demonstrated a

²⁶We computed *just* minimum profit for richer simulations with 75 values and 75 bids. The simulated value was approximately 0.26, or 93% of the theoretical limit with continuous values and messages and an unbounded message space.

²⁷When there are two bidders and a standard uniform value, the welfare outcome for the must-sell saddle point is $(0, 0.25)$. 0.25 is also the optimal maxmin posted price, if the Seller were forced to use a posted price, so that $(0, 0.25)$ is also the *best* welfare outcome for this maxmin posted price. As near as we can tell, this is pure coincidence.

methodology that we hope can be fruitfully applied to other topics in mechanism design, such as bilateral trade or voting systems.

A caveat is in order. We began with a critique of the standard model, in which the designer is supposed to be completely convinced that a particular type space correctly describes the world. In distancing ourselves from this view, we have taken an equally extreme position, which is that the designer is unwilling to place any restrictions on information, save for the common-prior assumption and the fixed prior on values. Verily, the truth must lie somewhere in between. Our view is that the theory will become even more useful and practically relevant if we can find ways to explore the middle ground between these two extremes, by incorporating reasonable restrictions on beliefs into the robust mechanism design problem.

References

- ABREU, D. AND H. MATSUSHIMA (1992): “Virtual Implementation in Iteratively Undominated Strategies: Complete Information,” *Econometrica*, 993–1008.
- AUMANN, R. J. AND M. MASCHLER (1985): “Game Theoretic Analysis of a Bankruptcy Problem from the Talmud,” *Journal of Economic Theory*, 36, 195–213.
- BALI, V. AND M. JACKSON (2002): “Asymptotic Revenue Equivalence in Auctions,” *Journal of Economic Theory*, 106, 161–176.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): “The Limits of Price Discrimination,” *American Economic Review*, 105, 921–57.
- (2016a): “Informationally Robust Optimal Auction Design,” Tech. rep., Princeton University and the University of Chicago and Yale University, working paper.
- (2016b): “Selling to Intermediaries: Optimal Auction Design in a Common Value Model,” Tech. rep., Princeton University and the University of Chicago and Yale University, working paper.
- (2017): “First-Price Auctions with General Information Structures: Implications for Bidding and Revenue,” *Econometrica*, 85, 107–143.
- 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.
- BLACKWELL, D. AND M. GIRSHICK (1954): *Theory of Games and Statistical Decisions.*, Oxford, England: Wiley.
- BOOS, D. D. (1985): “A Converse to Scheffe’s Theorem,” *The Annals of Statistics*, 423–427.
- BÖRGERS, T. (2017): “(No) Foundations of Dominant-Strategy Mechanisms: A Comment on Chung and Ely (2007),” *Review of Economic Design*, 1–10.
- BOSE, S., E. OZDENOREN, AND A. PAPE (2006): “Optimal Auctions with Ambiguity,” *Theoretical Economics*, 1, 411–438.
- BROOKS, B. (2013): “Surveying and Selling: Belief and Surplus Extraction in Auctions,” Working paper.
- BULOW, J. AND P. KLEMPERER (1996): “Auctions Versus Negotiations,” *The American Economic Review*, 180–194.

- CARBONELL-NICOLAU, O. AND R. P. MCLEAN (2017): “On the Existence of Nash Equilibrium in Bayesian Games,” *Mathematics of Operations Research*.
- CARRASCO, V., V. F. LUZ, N. KOS, M. MESSNER, P. MONTEIRO, AND H. MOREIRA (2017): “Optimal Selling Mechanisms Under Moment Conditions,” .
- CARROLL, G. (2016): “Robustness and Separation in Multidimensional Screening,” *Econometrica*, Forthcoming.
- CHEN, Y.-C. AND J. LI (2016): “Revisiting the Foundations of Dominant-Strategy Mechanisms,” Working paper.
- CHUNG, K.-S. AND J. C. ELY (2007): “Foundations of Dominant-Strategy Mechanisms,” *The Review of Economic Studies*, 74, 447–476.
- CRÉMER, J. AND R. P. 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*, 1247–1257.
- DU, S. (2016): “Robust Mechanisms Under Common Valuation,” Tech. rep., Simon Fraser University, working paper.
- ENGELBRECHT-WIGGANS, R., P. R. MILGROM, AND R. J. WEBER (1983): “Competitive Bidding and Proprietary Information,” *Journal of Mathematical Economics*, 11, 161–169.
- FALKNER, N. AND G. TESCHL (2012): “On the Substitution Rule for Lebesgue–Stieltjes Integrals,” *Expositiones Mathematicae*, 30, 412–418.
- HARSANYI, J. C. (1967): “Games with Incomplete Information played by Bayesian players, I–III Part I. The basic model,” *Management science*, 14, 159–182.
- HARTLINE, J. D. AND T. ROUGHGARDEN (2009): “Simple Versus Optimal Mechanisms,” in *Proceedings of the 10th ACM conference on Electronic commerce*, ACM, 225–234.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian Persuasion,” *The American Economic Review*, 101, 2590–2615.
- MASKIN, E. (1999): “Nash Equilibrium and Welfare Optimality,” *The Review of Economic Studies*, 66, 23–38.
- MCAFEE, R. P. (1993): “Mechanism Design by Competing Sellers,” *Econometrica*, 1281–1312.

- MCAFEE, R. P. AND J. MCMILLAN (1987): “Auctions with Entry,” *Economics Letters*, 23, 343–347.
- MCAFEE, R. P., J. MCMILLAN, AND P. J. RENY (1989): “Extracting the Surplus in the Common-Value Auction,” *Econometrica*, 57, 1451–1459.
- MCAFEE, R. P. AND P. J. RENY (1992): “Correlated Information and Mechanism Design,” *Econometrica*, 60, 395–421.
- MILGROM, P. R. (1979): “A Convergence Theorem for Competitive Bidding with Differential Information,” *Econometrica*, 679–688.
- MILGROM, P. R. AND R. J. WEBER (1982): “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 1089–1122.
- (1985): “Distributional Strategies for Games with Incomplete Information,” *Mathematics of operations research*, 10, 619–632.
- MYERSON, R. B. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- (1983): “Analysis of Two Bargaining Problems with Incomplete Information,” in *Game-Theoretic Models of Bargaining*, ed. by A. E. Roth, Cambridge: Cambridge University Press, 149–163.
- NEEMAN, Z. (2003): “The Effectiveness of English Auctions,” *Games and Economic Behavior*, 43, 214–238.
- RENY, P. J. (1999): “On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games,” *Econometrica*, 67, 1029–1056.
- ROESLER, A.-K. AND B. SZENTES (2017): “Buyer-Optimal Learning and Monopoly Pricing,” *American Economic Review*, 107, 2072–2080.
- ROTHSCHILD, M. AND J. E. STIGLITZ (1970): “Increasing Risk: I. A Definition,” *Journal of Economic Theory*, 2, 225 – 243.
- THOMSON, W. (1994): “Cooperative Models of Bargaining,” *Handbook of Game Theory with Economic Applications*, 2, 1237–1284.
- TULLOCK, G. (1980): “Efficient rent-seeking,” in *Toward a theory of the rent-seeking society*, ed. by J. M. Buchanan, R. D. Tollison, and G. Tullock, College Station, TX: A&M University Press, 97–112.
- VARTIAINEN, H. (2013): “Auction Design Without Commitment,” *Journal Of The European Economic Association*, 11, 316–342.
- VOHRA, R. V. (2011): *Mechanism Design: A Linear Programming Approach*, vol. 47, Cambridge University Press.

- WILSON, R. (1977): “A Bidding Model of Perfect Competition,” *The Review of Economic Studies*, 511–518.
- WOLITZKY, A. (2016): “Mechanism Design with Maxmin Agents: Theory and an Application to Bilateral Trade,” *Theoretical Economics*, 11, 971–1004.
- YAMASHITA, T. (2015): “Implementation in Weakly Undominated Strategies: Optimality of Second-Price Auction and Posted-Price Mechanism,” *The Review of Economic Studies*, 82, 1223–1246.
- (2016): “Revenue Guarantee in Auction with a (Correlated) Common Prior and Additional Information,” Working paper.

A Omitted proofs

Proof of Lemma 1. It is well known that X is a mean-preserving spread of Y if and only if for all x ,

$$\int_{y=0}^x (F_X(y) - F_Y(y))dy \geq 0 \quad (42)$$

and if the inequality holds as an equality at $x = 0$ and in the limit as $x \rightarrow \infty$. We will show that (12) holds for all α if and only if (42) holds for all x . The reason is as follows. Note that

$$\int_{y=0}^x F_X(y)dy = F_X(x)x - \int_{y=0}^{F_X(x)} F_X^{-1}(y)dy,$$

and similarly for F_Y . As a result, if $F_X(x) = F_Y(x)$, then (12) and (42) are equivalent.

Now consider an interval (\underline{x}, \bar{x}) where $F_X(x) > F_Y(x)$ for all $x \in (\underline{x}, \bar{x})$, and $F_X(\underline{x}) = F_Y(\underline{x}) = \underline{\alpha}$ and $F_X(\bar{x}) = F_Y(\bar{x}) = \bar{\alpha}$. Then $F_Y^{-1}(\alpha) > F_X^{-1}(\alpha)$ for all $\alpha \in (\underline{\alpha}, \bar{\alpha})$, and hence the integrals on the left-hand sides of (12) and (42) are monotonically increasing in x and α , respectively, and the integrals are therefore bounded between their respective values at the endpoints. Thus, if (12) and (42) are satisfied at the points at which the CDFs cross, they will be satisfied in between the crossing points as well.

The case where $F_X < F_Y$ on (\underline{x}, \bar{x}) is analogous and is omitted. \square

Proof of Lemma 2. For (a), we only need to show that $v - c$ is a mean-preserving spread of $\bar{\gamma}(x)$. But this follows from the following observations: $\bar{\Gamma}(x) \geq \hat{\Gamma}(x)$, since $\bar{\Gamma}(E^{-1}(z))$ is the concavification of $\hat{\Gamma}(E^{-1}(z))$ and E is a strictly increasing function. Thus, for all x

$$\int_{y=0}^x \bar{\gamma}(y)g_N(y)dy \geq \int_{y=0}^x \hat{\gamma}(y)g_N(y)dy,$$

which is equivalent to for all x

$$\int_{y=0}^x \bar{\gamma}(G_N^{-1}(y))dy \geq \int_{y=0}^x \hat{\gamma}(G_N^{-1}(y))dy.$$

But $\bar{\gamma}(G_N^{-1}(x))$ is the inverse of the CDF of $\bar{\gamma}(x)$, and $\hat{\gamma}(G_N^{-1}(x))$ is the inverse of the CDF of $v - c$, so by Lemma 1, $v - c$ must be a mean-preserving spread of $\bar{\gamma}(x)$.

For (b) note that both $\bar{\Gamma}(x)$ and $\hat{\Gamma}(x)$ are continuous and differentiable functions. Because we obtained $\bar{\Gamma}$ from $\hat{\Gamma}$ via a concavification, i.e., $\bar{\Gamma} \circ E^{-1}$ is the concavification of $\hat{\Gamma} \circ E^{-1}$, it must be that whenever these two functions coincide on the interior $(0, \infty)$, they have the same derivative. Thus, if $\bar{\Gamma}(x) = \hat{\Gamma}(x)$ at $x > 0$, $\bar{\gamma}(x) = \hat{\gamma}(x)$. Moreover, on a graded interval $[a, b]$ where we concavified, $\bar{\Gamma} \circ E^{-1}$ is linear, so that

$$\begin{aligned} \frac{d}{dx}\bar{\gamma}(x) &= \frac{d}{dz}(\bar{\gamma}(E^{-1}(z)))|_{z=E(x)}E'(x) \\ &= \frac{\bar{\gamma}(b) - \bar{\gamma}(a)}{b - a} \exp(x)g_N(x). \end{aligned}$$

Since $d\bar{\gamma}(x)/dx = \bar{\gamma}(x)g_N(x)$, we conclude that $\bar{\gamma}$ has an exponential shape.

For (c), we have proven that on an interval where $\bar{\Gamma} \neq \hat{\Gamma}$, $\bar{\gamma}(x)$ is exponential, and hence continuous and increasing. Moreover, it cannot be that $\hat{\gamma}$ has a discontinuity at a point x where $\bar{\Gamma}(x) = \hat{\Gamma}(x)$, because then $\hat{\Gamma}(E^{-1}(z))$ would have a convex kink at $z = E(x)$. Since $\hat{\gamma}$ is increasing, this implies that $\bar{\gamma}$ and $\hat{\gamma}$ coincide and are both continuous and increasing wherever $\bar{\Gamma}$ coincides with $\hat{\Gamma}$. Moreover, at points where the concavified and non-concavified regions of $\bar{\Gamma} \circ E^{-1}$ meet, $\bar{\Gamma} \circ E^{-1}$ must be continuously differentiable, so that $\bar{\gamma}$ is continuous.

For (d), it must be that $d\bar{\gamma}(x)/dx = \bar{\gamma}(x)$ on any concavified region. Now suppose that $\bar{\Gamma}(x) = \hat{\Gamma}(x)$ at a point where $\hat{\Gamma}$ is differentiable but that the derivative from the right is $\hat{\Gamma}'(x) > \hat{\gamma}(x)$. But the derivative of $\hat{\Gamma}(E^{-1}(z))$ from the right, evaluated at $z = E(x)$, is

$$\frac{\hat{w}(E^{-1}(z))g_N(E^{-1}(z))}{E'(E^{-1}(z))} = \frac{\hat{w}(E^{-1}(z))}{\exp(E^{-1}(z))}.$$

So if $\hat{\gamma}$ increases faster than exponential at x , then the derivative of $\hat{\Gamma} \circ E^{-1}$ is increasing at $z = E(x)$, and hence $\hat{\Gamma} \circ E^{-1}$ is convex at $E(x)$. This contradicts the hypothesis that $\bar{\Gamma}(x) = \hat{\Gamma}(x)$. \square

Proof of Lemma 3. Let $U_i(s_i, m_i)$ denote the utility of a type s_i that misreports as m_i to the direct mechanism (q, t) , when others report truthfully:

$$U_i(s_i, m_i) = \int_{s_{-i}} (w(s_i, s_{-i})q_i(s_i, s_{-i}) - t_i(s_i, s_{-i}))f_{-i}(s_{-i})ds_{-i},$$

with $U_i(s_i) \equiv U_i(s_i, s_i)$. The incentive constraints say that $U_i(s_i) = \max_{m_i} U_i(s_i, m_i)$. As a result,

$$\begin{aligned} U_i(s_i) - U_i(m_i) &\geq U_i(s_i, m_i) - U_i(s_i) \\ &= \int_{s_{-i}} (w(s_i, s_{-i}) - w(m_i, s_{-i}))q_i(m_i, s_{-i})f_{-i}(s_{-i})ds_{-i}. \end{aligned}$$

Since w is increasing in s_i , this implies that U_i is increasing as well.

Now, it must be that

$$U_i(s_i) - \underline{U}_i = \int_{x=0}^{s_i} dU_i(x),$$

where this is a Riemann-Stieltjes integral with respect to the monotonic and right-continuous function U_i (due to right-continuity of w in s_i). This is the limit of finite sums of the form

$$\sum_{k=1}^K (U_i(s_i^k) - U_i(s_i^{k-1}))$$

over partitions

$$0 = s_i^0 \leq s_i^1 \leq \dots \leq s_i^K = s_i,$$

as we take the absolute distance between adjacent points going to zero. Using the incentive constraint, we can bound each of these sums below by

$$\int_{s_{-i}} \left[\sum_{k=1}^K (w(s_i^k, s_{-i}) - w(s_i^{k-1}, s_{-i})) q_i(m_i, s_{-i}) \right] f_{-i}(s_{-i}) ds_{-i}. \quad (43)$$

Note that the Riemann sum inside the integral is bounded above by $w(s_i, s_{-i}) - w(0, s_{-i})$, which is an integrable function. Lebesgue's Dominated Convergence Theorem therefore implies that we have that the limit of the integrals (43) is equal to the integral of the limit function (of s_i), which is

$$\int_{s_{-i}} \left[\int_{x=0}^{s_i} q_i(x, s_{-i}) \omega_i(dx; s_{-i}) \right] f_{-i}(s_{-i}) ds_{-i},$$

where ω_i is the measure defined in (16). This expression is therefore a lower bound on $U_i(s_i)$, and hence a lower bound on bidder i 's surplus is

$$\int_{s_i} \int_{s_{-i}} \left[\int_{x=0}^{s_i} q_i(x, s_{-i}) \omega_i(dx; s_{-i}) \right] f_{-i}(s_{-i}) ds_{-i} f_i(s_i) ds_i.$$

By Tonelli's Theorem, and the fact that the integrand is bounded and product measurable, we can exchange the order of integration, so that this lower bound rearranges to

$$\begin{aligned} & \int_{s_{-i}} \left[\int_{s_i=0}^{\infty} \int_{x=0}^{s_i} q_i(x, s_{-i}) \omega_i(dx; s_{-i}) f_i(s_i) ds_i \right] f_{-i}(s_{-i}) ds_{-i} \\ &= \int_{s_{-i}} \left[\int_{s_i=0}^{\infty} \int_{x=s_i}^{\infty} f_i(x) dx q_i(s_i, s_{-i}) \omega_i(ds_i; s_{-i}) \right] f_{-i}(s_{-i}) ds_{-i} \\ &= \int_{s_{-i}} \left[\int_{s_i=0}^{\infty} (1 - F_i(s_i)) q_i(s_i, s_{-i}) \omega_i(ds_i; s_{-i}) \right] f_{-i}(s_{-i}) ds_{-i} \\ &= \int_s \frac{1 - F_i(s)}{f(s)} q_i(s, s_{-i}) f(s) \omega_i(ds; s_{-i}) ds_{-i}. \end{aligned}$$

Using the fact that total surplus is

$$\int_s w(s) \sum_{i=1}^N q_i(s) f(s) ds$$

and that profit is total surplus minus total bidder surplus, we have the result. \square

Proof of Lemma 8. In the discussion after the statement of Lemma 8, we have already proved

$$\begin{aligned} \int_{m_{-i}} \frac{\partial \bar{t}_i}{\partial m'_i}(m'_i, m_{-i}) \exp(-\Sigma m_{-i}) dm_{-i} &= \int_{x=0}^{\infty} \bar{\Xi}(m'_i + x)(g_{N-1}(x) - g_N(x)) dx \\ &= \int_{x=0}^{\infty} \bar{\Xi}(m'_i + x) g'_N(x) dx. \end{aligned}$$

From the profit-incentive alignment, we have

$$d\bar{\Xi}(x) = d(\hat{w}(x)\bar{\mu}(x)) - \bar{\mu}(x)d\hat{w}(x) - c\bar{Q}'(x)dx = \hat{w}(x)d\bar{\mu}(x) - c\bar{Q}'(x)dx.$$

This equality depends on the fact that \hat{w} and $\bar{\mu}$ never have discontinuities at the same point. This must be the case, because as the proof of Lemma 2 shows, whenever $\hat{w} = \hat{\gamma} + c$ is discontinuous, $\hat{\Gamma}$ is strictly convex, so that the region around the discontinuity is concavified and $\bar{\gamma}$ has an exponential shape. Moreover, since $\bar{\mu}$ is constant on any region where $\hat{\gamma} \neq \bar{\gamma}$, we conclude that

$$d\bar{\Xi}(x) = \bar{w}(x)d\bar{\mu}(x) - c\bar{Q}'(x)dx.$$

Applying the integration-by-parts formula for the Riemann-Stieltjes integral, we finally conclude that

$$\begin{aligned} \int_{m_{-i}} \frac{\partial \bar{t}_i}{\partial m'_i}(m'_i, m_{-i}) \exp(-\Sigma m_{-i}) dm_{-i} &= - \int_{x=0}^{\infty} g_N(x) \bar{w}(m'_i + x) d\mu(m'_i + x) \\ &\quad + c \int_{x=0}^{\infty} g_N(x) \bar{Q}'(m'_i + x) dx. \end{aligned}$$

We next derive the rate of change in surplus from being allocated the good as a function of the report; we claim:

$$\begin{aligned} &\int_{m_{-i}} \bar{w}(m_i + \Sigma m_{-i}) \frac{\partial \bar{q}_i}{\partial m'_i}(m'_i, m_{-i}) \exp(-\Sigma m_{-i}) dm_{-i} \tag{44} \\ &= - \int_{x=0}^{\infty} g_N(x) \bar{w}(m_i + x) d\bar{\mu}(m'_i + x) + \int_{x=0}^{\infty} \left(\bar{w}(m_i + x) - \frac{d\bar{w}}{dx}(m_i + x) \right) \bar{Q}'(m'_i + x) g_N(x) dx. \end{aligned}$$

The second line of the equation above can be written as:

$$\begin{aligned} &\int_{x=0}^{\infty} \left(g_N(x) \frac{d\bar{w}}{dx}(m_i + x) + (g_{N-1}(x) - g_N(x)) \bar{w}(m_i + x) \right) \bar{\mu}(m'_i + x) dx \\ &\quad + \int_{x=0}^{\infty} \left(\bar{w}(m_i + x) - \frac{d\bar{w}}{dx}(m_i + x) \right) \bar{Q}'(m'_i + x) g_N(x) dx \\ &= \int_{x=0}^{\infty} g_{N-1}(x) \bar{w}(m_i + x) \bar{\mu}(m'_i + x) dx \\ &\quad + \int_{x=0}^{\infty} \left(\bar{w}(m_i + x) - \frac{d\bar{w}}{dx}(m_i + x) \right) (\bar{Q}'(m'_i + x) - \bar{\mu}(m'_i + x)) g_N(x) dx \\ &= \int_{x=0}^{\infty} g_{N-1}(x) \bar{w}(m_i + x) \bar{\mu}(m'_i + x) dx \\ &\quad - \int_{x=0}^{\infty} \left(\bar{w}(m_i + x) - \frac{d\bar{w}}{dx}(m_i + x) \right) \bar{Q}'(m'_i + x) \frac{N-1}{m'_i + x} g_N(x) dx, \end{aligned}$$

since $\bar{\mu}(x) = (N-1)\bar{Q}(x)/x + \bar{Q}'(x)$. Moreover, an additional integration by parts and a rearrangement give:

$$\begin{aligned} & \int_{x=0}^{\infty} \left(\bar{w}(m_i + x) - \frac{d\bar{w}}{dx}(m_i + x) \right) \bar{Q}(m'_i + x) \frac{N-1}{m'_i + x} g_N(x) dx \\ &= \int_{x=0}^{\infty} \bar{w}(m_i + x) \left(\bar{Q}(m'_i + x) \frac{N-1}{m'_i + x} g_N(x) + \frac{d}{dx} \left(\bar{Q}(m'_i + x) \frac{N-1}{m'_i + x} g_N(x) \right) \right) dx \\ &= \int_{x=0}^{\infty} \bar{w}(m_i + x) \left(\bar{Q}(m'_i + x) \frac{N-1}{m'_i + x} + \bar{Q}'(m'_i + x) \frac{x}{m'_i + x} - \bar{Q}(m'_i + x) \frac{x}{(m'_i + x)^2} \right) g_{N-1}(x) dx, \end{aligned}$$

which proves (44) since

$$\begin{aligned} \frac{\partial \bar{q}_i}{\partial m'_i}(m'_i, m_{-i}) &= \frac{m'_i}{m'_i + \Sigma m_{-i}} \bar{Q}'(m'_i + \Sigma m_{-i}) + \frac{\Sigma m_{-i}}{(m'_i + \Sigma m_{-i})^2} \bar{Q}(m'_i + \Sigma m_{-i}) \\ &= \bar{\mu}(m'_i + \Sigma m_{-i}) - \bar{Q}(m'_i + \Sigma m_{-i}) \frac{N-1}{m'_i + \Sigma m_{-i}} - \bar{Q}'(m'_i + \Sigma m_{-i}) \frac{\Sigma m_{-i}}{m'_i + \Sigma m_{-i}} \\ &\quad + \bar{Q}(m'_i + \Sigma m_{-i}) \frac{\Sigma m_{-i}}{(m'_i + \Sigma m_{-i})^2}. \end{aligned}$$

Therefore, we conclude that

$$\begin{aligned} & \int_{m_{-i}} \left(\bar{w}(m_i + \Sigma m_{-i}) \frac{\partial \bar{q}_i}{\partial m'_i}(m'_i, m_{-i}) - \frac{\partial \bar{t}_i(m'_i, m_{-i})}{\partial m'_i} \right) \exp(-\Sigma m_{-i}) dm_{-i} \\ &= \int_{x=0}^{\infty} g_N(x) (\bar{w}(m'_i + x) - \bar{w}(m_i + x)) d\mu(m'_i + x) \\ &\quad + \int_{x=0}^{\infty} \left(\bar{\gamma}(m_i + x) - \frac{d\bar{\gamma}}{dx}(m_i + x) \right) \bar{Q}'(m'_i + x) g_N(x) dx. \end{aligned}$$

Thus,

$$\begin{aligned} & \int_{m_{-i}} (\bar{w}(m_i + \Sigma m_{-i}) (\bar{q}_i(m'_i, m_{-i}) - \bar{q}_i(m_i, m_{-i})) \\ &\quad - (\bar{t}_i(m'_i, m_{-i}) - \bar{t}_i(m_i, m_{-i}))) \exp(-\Sigma m_{-i}) dm_{-i} \\ &= \int_{y=m_i}^{m'_i} \int_{x=0}^{\infty} g_N(x) (\bar{w}(y+x) - \bar{w}(m_i+x)) d\mu(y+x) dy \\ &\quad + \int_{x=0}^{\infty} \left(\bar{\gamma}(m_i+x) - \frac{d\bar{\gamma}}{dx}(m_i+x) \right) (\bar{Q}(m'_i+x) - \bar{Q}(m_i+x)) g_N(x) dx, \end{aligned}$$

where we have used the Fubini theorem for the second line. We claim that this expression must be non-positive. For the first integral, because \bar{w} is increasing and $\bar{\mu}$ is decreasing, the integrand is negative if $y > m_i$ and positive if $y < m_i$, so that in fact the first integral is single-peaked at $m'_i = m_i$. For the second integral, recall that by construction, $\bar{\gamma}(m_i+x) - d\bar{\gamma}/dx(m_i+x) \geq 0$, and if it is strictly positive, then the good must be allocated with probability one, so $\bar{Q}(m_i+x) = 1$. But this is the maximum possible supply, so

$\overline{Q}(m'_i + x) \leq 1$, and hence the integrand is non-positive. This concludes the proof of global incentive compatibility.

We remark that this proof is the same for both the non-compactified and compactified versions of $\overline{\mathcal{M}}$. The reason is that under the truth-telling strategies, infinite messages are sent with zero probability, and conditional on others sending finite messages, the allocation and transfer rules are both continuous in one's own message. As a result, if no finite deviations are profitable, then a deviation to an infinite message is not profitable either. □

Incidentally, the only feature of \overline{Q} that was used in this proof was that the maximum supply is attained whenever the value function is fully revealing. Truth-telling would still be an equilibrium for any supply rule that shares this feature.

Proof of Proposition 5. We first verify that Condition 1 is satisfied in Carbonell-Nicolau and McLean (2017). For any i and $\epsilon > 0$, we will find a measurable mapping $\zeta : M_i \rightarrow M_i$ so that for all pairs (s, m) of a signal profile and a message profile, there exists a neighborhood U of m_{-i} so that for all $m'_{-i} \in U$,

$$w(s)q_i(\zeta(m_i), m'_{-i}) - t_i(\zeta(m_i), m'_{-i}) \geq w(s)q_i(m_i, m_{-i}) - t_i(m_i, m_{-i}) - \epsilon. \quad (45)$$

This is trivially satisfied by the mapping

$$\zeta_i(m_i) = \begin{cases} m_i & \text{if } m_i > 0; \\ \delta & \text{otherwise,} \end{cases}$$

where δ will be defined shortly. Let us first deal with the profiles for which $m_i > 0$. Since t_i is continuous, we can define a neighborhood V so that $t_i(m_i, \cdot)$ varies by no more than $\epsilon/2$ within U . Also, we can define a neighborhood V'_j according to the following cases. If $\Sigma m = \infty$, then we set

$$V'_j = \begin{cases} [0, \infty) & \text{if } m_j \in [0, \infty); \\ [0, \infty] & \text{if } m_j = \infty, \end{cases}$$

and set $U = V \cap (\times_{j \neq i} V'_j)$. Now, as we vary $m_{-i} \in U$, the number of bidders sending infinite messages can only decrease, so the allocation is unchanged, so bidder i 's allocation can only increase, but the transfer varies by no more than $\epsilon/2$, so (45) is satisfied. On the other hand, if $\Sigma m < \infty$, then the allocation q_i is continuous at m , so we can define V' so that q_i varies by no more than $\epsilon/(2w(s))$ within V' , and set $U = V \cap V'$, and (45) will again be satisfied.

If $m_i = 0$, we can take δ to be small enough so that for all $x \leq \delta$, $t_i(x, m_{-i}) < \epsilon/2$ for all m_{-i} . Also, since Q is continuous, we can take δ to be small enough so that $Q(x) \geq Q(0) - \epsilon/(4v_{\max})$ for all $x \leq \delta$. As q_i is zero unless $m_j = 0$ for all $j \neq i$, we have already proven that the definition is satisfied for profiles such that $\Sigma m_{-i} \neq 0$. If $\Sigma m_{-i} = 0$, then we can define a δ' such that $Q(x) \geq Q(0) - \epsilon/(2v_{\max})$ for all $x \leq \delta + \delta'$,

and let the ball V' be the set of m_{-i} such that $\sum m_{-i} < \min\{\delta', \delta/(N-1)\}$. Then for all such m_{-i} ,

$$w(s)q_i(\delta, m_{-i}) - t_i(\delta, m_{-i}) \geq w(s)\delta \left(Q(0) - \frac{\epsilon}{2v_{\max}} \right) - \frac{\epsilon}{2}.$$

Finally, we verify that the rest of the conditions of Theorem 1 of Carbonell-Nicolau and McLean (2017) are satisfied. Product continuity of the type space has been already assumed. The remaining condition is that for every s , the aggregate utility

$$\sum_{i=1}^N [w(s)q_i(m) - t_i(m)]$$

is upper semi-continuous. But this follows from the fact that the transfers are continuous, and the sum of the allocations is upper semi-continuous. \square

Proof of Lemma 11. The first statement follows from the fact that $\widehat{\gamma}_N^C(x) = \widehat{\gamma}_\infty^C(\Phi^{-1}(G_N^C(x)))$, and the fact that G_N^C converges pointwise to Φ .

The second statement requires a more careful proof. Recall that the graded gains is defined via a multi-step procedure, which we describe again in the normalized signals: we define the integrated fully-revealing gains function

$$\widehat{\Gamma}_N^C(x) = \int_{x'=-\sqrt{N}}^x \widehat{\gamma}_N^C(x') g_N^C(x') dx'$$

and the integrated exponential growth gains function:

$$E_N(x) = \int_{x'=-\sqrt{N}}^x \exp(\sqrt{N}x' + N) g_N^C(x') dx'.$$

We then define $\text{cav}(\widehat{\Gamma}_N^C \circ E_N^{-1})$ to be the smallest concave function that is everywhere above $\widehat{\Gamma}_N^C \circ E_N^{-1}$, i.e., its concavification, and we set $\overline{\Gamma}_N^C = \text{cav}(\widehat{\Gamma}_N^C \circ E_N^{-1}) \circ E_N$. Finally, we define $\overline{\gamma}_N^C(x)$ so that $\overline{\gamma}_N^C(x) g_N^C(x)$ is the right-derivative of $\overline{\Gamma}_N^C(x)$.

Note that there is always a concavified region at the bottom, which may be degenerate, which covers the region $[0, (G_N^C)^{-1}(\Phi(x^*))]$. The reason is that $\widehat{\Gamma}_N^C$ is non-positive on this region, and strictly positive for higher values (from our assumption that ex ante gains from trade are strictly positive). Thus, the point 0 will always be concavified with some point above this interval.

Next, we will show that the convergence occurs at all points $\widehat{x} > x^*$ where $\overline{\gamma}_\infty^C$ is differentiable. From monotonicity, we know that the set of points where this function is *not* differentiable have Lebesgue measure zero, so that right-continuity implies convergence everywhere. So, let \widehat{x} be such a point, and define the sequence

$$\widehat{x}_N = (G_N^C)^{-1}(\Phi(\widehat{x})).$$

We will show that for N sufficiently large, $\overline{\gamma}_N^C(\widehat{x}_N) = \widehat{\gamma}_N^C(\widehat{x}_N)$, i.e., the graded gains function becomes fully-revealing at \widehat{x}_N . This is equivalent to showing that $\overline{\Gamma}_N^C(\widehat{x}_N) =$

$\widehat{\Gamma}_N^C(\widehat{x}_N)$ for N sufficiently large. This is equivalent to showing that the concavification of $\widehat{\Gamma}_N^C \circ E_N^{-1}$ at $y_N = E_N(\widehat{x}_N)$ is equal to $\widehat{\Gamma}_N^C \circ E_N^{-1}(y_N)$.

To that end, let

$$\zeta_N(y) \equiv \frac{d}{dy} \widehat{\Gamma}_N^C(E_N^{-1}(y)) = \frac{\widehat{\gamma}_N^C(E_N^{-1}(y))}{\exp(\sqrt{N}E_N^{-1}(y) + N)}.$$

We will show that for N sufficiently large, for all $y \in \mathbb{R}$,

$$\widehat{\Gamma}_N^C(E_N^{-1}(y)) \leq \widehat{\Gamma}_N^C(E_N^{-1}(y_N)) + \zeta_N(y_N)(y - y_N). \quad (46)$$

Since $\widehat{\gamma}_N^C$ is differentiable at x_N , we can find an ϵ sufficiently small so that for all $N \geq \widetilde{N}$, if $|x - \widehat{x}_N| < \epsilon$, then

$$|\widehat{\gamma}_N^C(x) - \widehat{\gamma}_N^C(\widehat{x}_N)| < (1 + \frac{d}{dx} \widehat{\gamma}_N^C(\widehat{x}_N)) |x - \widehat{x}_N|.$$

Note that

$$\begin{aligned} \zeta_N(E_N(x)) - \zeta_N(E_N(\widehat{x}_N)) &= \frac{\widehat{\gamma}_N^C(x)}{\exp(\sqrt{N}x + N)} - \frac{\widehat{\gamma}_N^C(\widehat{x}_N)}{\exp(\sqrt{N}\widehat{x}_N + N)} \\ &= \frac{\widehat{\gamma}_N^C(x) - \widehat{\gamma}_N^C(\widehat{x}_N) \exp(\sqrt{N}(x - \widehat{x}_N))}{\exp(\sqrt{N}x + N)}. \end{aligned}$$

So, for $x \in [\widehat{x}_N, \widehat{x}_N + \epsilon]$,

$$\begin{aligned} \zeta_N(E_N(x)) - \zeta_N(E_N(\widehat{x}_N)) &\leq \frac{\widehat{\gamma}_N^C(\widehat{x}_N) + \frac{d}{dx} \widehat{\gamma}_N^C(\widehat{x}_N)(x - \widehat{x}_N) - \widehat{\gamma}_N^C(\widehat{x}_N) \exp(\sqrt{N}(x - \widehat{x}_N))}{\exp(\sqrt{N}x + N)} \\ &= \frac{\frac{d}{dx} \widehat{\gamma}_N^C(\widehat{x}_N)(x - \widehat{x}_N) + \widehat{\gamma}_N^C(\widehat{x}_N)(1 - \exp(\sqrt{N}(x - \widehat{x}_N)))}{\exp(\sqrt{N}x + N)}, \end{aligned}$$

which is non-positive for N sufficiently large, since the numerator is zero at $x = \widehat{x}_N$ and its derivative with respect to x is

$$\frac{d}{dx} \widehat{\gamma}_N^C(\widehat{x}_N) - \sqrt{N} \widehat{\gamma}_N^C(\widehat{x}_N) \exp(\sqrt{N}(x - \widehat{x}_N)),$$

which is negative for \sqrt{N} sufficiently large, since $\widehat{\gamma}_N^C(\widehat{x}_N) \rightarrow \widehat{\gamma}_\infty^C(\widehat{x}) > 0$ and $\frac{d}{dx} \widehat{\gamma}_N^C(\widehat{x}_N) \approx \frac{d}{dx} \widehat{\gamma}_\infty^C(\widehat{x}) > 0$. Similarly, if $x \in [\widehat{x}_N - \epsilon, \widehat{x}_N]$, then

$$\zeta_N(E_N(x)) - \zeta_N(E_N(\widehat{x}_N)) \geq \frac{-\frac{d}{dx} \widehat{\gamma}_N^C(\widehat{x}_N)(x - \widehat{x}_N) + \widehat{\gamma}_N^C(\widehat{x}_N)(1 - \exp(\sqrt{N}(x - \widehat{x}_N)))}{\exp(\sqrt{N}x + N)},$$

which is non-negative for N sufficiently large, since the numerator is zero at $x = \widehat{x}_N$ and the derivative with respect to $x - \widehat{x}_N$ is

$$-\frac{d}{dx} \widehat{\gamma}_N^C(\widehat{x}_N) - \sqrt{N} \widehat{\gamma}_N^C(\widehat{x}_N) \exp(\sqrt{N}(x - \widehat{x}_N)),$$

which is negative for all \sqrt{N} .

Thus, we conclude that locally around \hat{x}_N , $\zeta_N(x) - \zeta_N(\hat{x}_N)$ is single-crossing from above. Moreover, we can always pick \sqrt{N} large enough so that

$$\hat{\gamma}_N^C(\hat{x}_N) \exp(\sqrt{N}\epsilon + N) \geq v_{\max} \geq \hat{\gamma}_N^C(x)$$

for all x , so that $\zeta_N(x) - \zeta_N(\hat{x}_N) \leq 0$ for all $x \geq \hat{x}_N$. This proves the inequality (46) for $x \geq \hat{x}_N - \epsilon$. Finally, observe that

$$\begin{aligned} & \hat{\Gamma}_N^C(\hat{x}_N) + \zeta_N(E_N(\hat{x}_N))(E_N(\hat{x}_N - \epsilon) - E_N(\hat{x}_N)) \\ &= \hat{\Gamma}_N^C(\hat{x}_N) + \hat{\gamma}_N^C(\hat{x}_N) \int_{x'=\hat{x}_N-\epsilon}^{\hat{x}_N} \exp(\sqrt{N}(x' - \hat{x}_N)) g_N^C(x') dx'. \end{aligned}$$

Since the integrand converges to zero pointwise as $N \rightarrow \infty$, the whole expression converges to $\hat{\Gamma}_\infty^C(\hat{x})$, which is necessarily greater than $\hat{\Gamma}_N^C(x)$ for $x < \hat{x}_N - \epsilon$ for N sufficiently large, since the latter converges to a quantity less than $\hat{\Gamma}_\infty^C(x)$ for $x < \hat{x} - \epsilon$. This completes the proof. \square

Proof of Proposition 10. Only the second inequality requires a more detailed proof. Consider the must-sell mechanism $\overline{\mathcal{M}}$ with N bidders, and construct a new $N + 1$ -bidder must-sell mechanism \mathcal{M} with $M_i = \mathbb{R}_+$ and allocation

$$q_i(m_1, \dots, m_{N+1}) = \begin{cases} \bar{q}_i(m_1, \dots, m_N) & \text{if } i \leq N; \\ 1 - \sum_{i=1}^N \bar{q}_i(m_1, \dots, m_N) & \text{otherwise} \end{cases}$$

and transfers

$$t_i(m_1, \dots, m_{N+1}) = \begin{cases} \bar{t}_i(m_1, \dots, m_N) & \text{if } i \leq N; \\ 0 & \text{otherwise.} \end{cases}$$

Clearly, this is a regular must-sell mechanism, and we claim that it provides a profit guarantee of $\overline{\Pi}_N(H)$. This claim will complete the proof, since the optimal profit guarantee among regular must-sell mechanisms is $\widehat{\Pi}_{N+1}(H)$, which must be weakly higher.

To prove the claim, we will show that any type space and equilibrium must generate profit of at least $\overline{\Pi}_N(H)$. Consider any $N + 1$ -bidder type space $\mathcal{T} = (\{S_i\}_{i=1}^{N+1}, \pi, w)$ and equilibrium β of the mechanism \mathcal{M} . Let $\mathcal{T}' = (\{S_i\}_{i=1}^N, \pi', w')$ be the N -bidder type space where the signal spaces are the same as those of the first N bidders in \mathcal{T} , the measure π' is the marginal of π on $S_1 \times \dots \times S_N$, and w' is a version of the conditional expectation of w given (s_1, \dots, s_N) . Also, let β' be the strategy profile consisting of $(\beta_1, \dots, \beta_N)$, which is a strategy profile in the N -bidder maxmin mechanism $\overline{\mathcal{M}}$. Since bidder $N + 1$'s message did not affect the allocations or transfers of bidders $i \leq N$, it is obvious that each bidder $i \leq N$'s ex ante utility from a strategy profile (β_i'', β'_{-i}) under mechanism $\overline{\mathcal{M}}$ is the same as from the strategy profile (β_i'', β_{-i}) under mechanism \mathcal{M} , and moreover, since only bidders $i \leq N$ make non-zero transfers, profit is the same under these equilibria as well. But since $\overline{\mathcal{M}}$ is the maxmin mechanism, we know that profit is at least $\overline{\Pi}_N(H)$. \square

B The Shapley rule

Let us argue that the Shapley rule is part of a maxmin mechanism when the support of values is $\{0, 1\}$, and the transfers are given by (28) and (8). The minmax type space in this case has a value function of the form

$$w(s) = \min \{1, a \exp(\Sigma s)\},$$

where $a \in (0, 1)$ is a parameter chosen so that the expected value is equal to the prior mean, which is just the probability that the value v is equal to 1. The expected value, as a function of the number of players and a , is

$$V_N(a) = a \sum_{k=0}^N \frac{(-\log(a))^k}{k!}.$$

This calculation is described in Bergemann et al. (2016a), in which it was notated as $E_{n,a,1}$. Also, maximum profit is just the expectation of the value from the perspective of a bidder i with $s_i = 0$, which it turns out is just $V_{N-1}(a)$ (player i 's signal drops out of the product).

There is a convenient reparametrization of this type space, where the s_i are i.i.d. draws from an exponential with arrival rate $-1/\log(a)$, i.e.,

$$s_i \sim F(s_i) = 1 - \exp\left(-\frac{1}{\log(a)}(-s_i)\right),$$

and the value function is

$$v(s) = a^{\max\{0, 1 - \sum_i s_i\}},$$

and a is the same parameter between the two models. The reason for reparametrizing is because it will turn out to be the BCE for the maxmin mechanism with the Shapley allocation rule. Specifically, if the buyers submit demands s , then the Seller fills the demands in a random order until the ‘‘supply’’ runs out.

The value multipliers for this model are

$$\begin{aligned} \lambda(0) &= V_{n-1}(a) + \frac{nV_n(a)}{\log(a)} \\ \lambda(1) &= V_{n-1}(a) - \frac{n(1 - V_n(a))}{\log(a)}. \end{aligned}$$

If we choose a local multiplier of

$$\alpha = -\frac{1}{\log(a)},$$

then the dual constraints will be binding for both $v = 0$ and $v = 1$ when $\sum_i s_i < 1$, and the $v = 1$ constraints will imply the $v = 0$ constraints when $\sum_i s_i > 1$ (assuming

we use the Shapley allocation rule described above). Again, this calculation is done in Bergemann et al. (2016a).

So, we just have to argue that the induced Ξ has zero expectation:

$$\Xi(s) = \lambda(1) - \alpha \sum_{i=1}^n \frac{\partial q_i(s)}{\partial s_i}.$$

This will imply that there exist transfers that are feasible for the dual.

Now, Ξ has a complicated expression for each s . But its expectation is quite simple. Under the Shapley rule, a player i is equally likely to have any place in line. If he is the k th bidder served, then the probability that increasing demand will increase his allocation is simply the probability that the sum of the first k demands is less than 1, which is $G_k(-\log(a))$ (since the demands have an arrival rate of $-1/\log(a)$). Thus, the expectation of $\Xi(s)$ is

$$\lambda(1) - \alpha \sum_{i=1}^N \frac{1}{N} \sum_{k=1}^N G_k(-\log(a)).$$

Now observe that

$$G_k(-\log(a)) = 1 - V_{k-1}(a).$$

Hence,

$$\begin{aligned} \alpha \sum_{i=1}^N \frac{1}{N} \sum_{k=1}^N G_k(-\log(a)) &= -\frac{1}{\log(a)} \sum_{k=1}^N (1 - V_{k-1}(a)) \\ &= -\frac{N}{\log(a)} + \frac{1}{\log(a)} \sum_{k=0}^{N-1} V_k(a) \\ &= -\frac{N}{\log(a)} + \frac{a}{\log(a)} \sum_{k=0}^{N-1} \sum_{m=0}^k \frac{(-\log(a))^m}{m!} \\ &= -\frac{N}{\log(a)} + \frac{a}{\log(a)} \sum_{m=0}^{N-1} \frac{N-m}{m!} (-\log(a))^m \\ &= -\frac{N}{\log(a)} + \frac{N}{\log(a)} V_{N-1}(a) + V_{N-2}(a) \\ &= -\frac{N}{\log(a)} + \frac{N}{\log(a)} V_{N-1}(a) + V_{N-2}(a) \\ &\quad + \frac{N}{\log(a)} a \frac{(-\log(a))^N}{N!} + a \frac{(-\log(a))^{N-1}}{(N-1)!} \\ &= -\frac{N}{\log(a)} + \frac{N}{\log(a)} V_N(a) + V_{N-1}(a) \\ &= \lambda(1). \end{aligned}$$

So the expectation of Ξ is zero, and the transfers exist.

Next, let us argue that such transfers and the Shapley allocation are an incentive compatible direct mechanism on the minmax type space, so that a BCE exists. For this purpose, it seems easiest to go back to our original parametrization of the type space, with

$$w(s) = \min\{1, a \exp(\Sigma s)\}.$$

It is fairly straightforward to show that this is the same type space we had before, where the old signal t_i maps to $s_i = -\log(a)t_i$. The multiplier on local incentives is now 1, but importantly, the allocation rule is the Shapley rule applied to “demands” $-s_i/\log(a)$, so we have to be mindful of the extra term when we differentiate the allocation.

Let us compute the interim utility of a misreport of s'_i when the true signal is s_i at the minmax type space. This is

$$U_i(s_i, s'_i) = \int_{s_{-i}} (\min\{1, a \exp(s_i + \Sigma s_{-i})\} q_i(s'_i, s_{-i}) - t_i(s'_i, s_{-i})) \exp(-\Sigma s_{-i}) ds_{-i}. \quad (47)$$

We will show that this expression is single-peaked in s'_i .

The transfers are such that

$$\begin{aligned} \int_{s_{-i}} t_i(s'_i, s_{-i}) \exp(-\Sigma s_{-i}) ds_{-i} &= -\exp(s'_i) \int_{x_i=0}^{s'_i} \int_{s_{-i}} \xi_i(x_i, s_{-i}) \exp(-\Sigma s_{-i}) ds_{-i} \exp(-x_i) dx_i \\ &= -\exp(s'_i) \int_{x_i=0}^{s'_i} \int_{s_{-i}} \Xi(x_i, s_{-i}) \exp(-x_i - \Sigma s_{-i}) ds_{-i} dx_i, \end{aligned}$$

where

$$\Xi(s) = \lambda(1) - \sum_{j=1}^N \frac{\partial q_j(s)}{\partial s_j}.$$

(Here we have used the unitary multiplier on incentive constraints.) Using the earlier tricks, we can simplify part of the expectation of Ξ , which we denote by A for future

reference:

$$\begin{aligned}
A &= \int_{x_i=0}^{s'_i} \int_{s_{-i}} \sum_{j=1}^N \frac{\partial q_j(x_i, s_{-i})}{ds_j} \exp(-x_i - \Sigma s_{-i}) ds_{-i} dx_i \\
&= \int_{x_i=0}^{s'_i} \int_{s_{-i}} \frac{\partial q_i(x_i, s_{-i})}{\partial s_i} \exp(-x_i - \Sigma s_{-i}) ds_{-i} dx_i \\
&\quad + \int_{x_i=0}^{s'_i} \int_{s_{-i}} \sum_{j \neq i} \frac{\partial q_j(x_i, s_{-i})}{\partial s_j} \exp(-x_i - \Sigma s_{-i}) ds_{-i} dx_i \\
&= -\frac{1}{\log(a)} \int_{x_i=0}^{s'_i} \frac{1}{N} \sum_{k=1}^N G_{k-1}(-\log(a) - x_i) \exp(-x_i) dx_i \\
&\quad - \frac{1}{\log(a)} \int_{x_i=0}^{s'_i} \frac{N-1}{N} \sum_{k=1}^N \frac{k-1}{N-1} G_{k-1}(-\log(a) - x_i) \exp(-x_i) dx_i \\
&\quad - \frac{1}{\log(a)} \int_{x_i=0}^{s'_i} \frac{N-1}{N} \sum_{k=1}^N \frac{N-k}{N-1} G_k(-\log(a)) \exp(-x_i) dx_i.
\end{aligned}$$

The third-to-last line comes from integrating over all the events where the report x_i is pivotal to i 's allocation, which is when the sum of the $k-1$ reports of those served before i is not larger than $-\log(a) - s_i$ (the sum of demands hits one when the sum of the reported types hits $-\log(a)$). The second-to-last and last lines compute something similar, but here the calculation depends on whether player i (whose signal is fixed) is served before or after player j . We can combine terms to simplify this expression to

$$\begin{aligned}
&= -\frac{1}{\log(a)} (1 - \exp(-s'_i)) \sum_{k=1}^N \frac{N-k}{N} G_k(-\log(a)) \\
&\quad - \underbrace{\frac{1}{\log(a)} \sum_{k=1}^N \frac{k}{N} \int_{x_i=0}^{s'_i} G_{k-1}(-\log(a) - x_i) \exp(-x_i) dx_i}_{B}.
\end{aligned}$$

Let us focus on this last expression:

$$\begin{aligned}
B &= \int_{x_i=0}^{s'_i} \sum_{k=1}^N \frac{k}{N} \left(1 - \sum_{l=0}^{k-1} g_l(-\log(a) - x_i) \right) \exp(-x_i) dx_i \\
&= \sum_{k=1}^N \frac{k}{N} \left(1 - \exp(-s'_i) - \sum_{l=0}^{k-1} \int_{x_i=0}^{s'_i} \frac{(-\log(a) - x_i)^{l-1}}{(l-1)!} \exp(\log(a)) dx_i \right) \\
&= \sum_{k=1}^N \frac{k}{N} \left(1 - \exp(-s'_i) + \sum_{l=0}^k \left(\frac{(-\log(a) - s'_i)^{l-1}}{(l-1)!} - \frac{(-\log(a))^l}{l!} \right) \exp(\log(a)) \right) \\
&= \sum_{k=1}^N \frac{k}{N} \left(1 - \exp(-s'_i) + \sum_{l=1}^k (\exp(s'_i) g_l(-\log(a) - s'_i) - g_l(-\log(a))) \right) \\
&= \sum_{k=1}^N \frac{k}{N} (G_k(-\log(a)) - \exp(-s'_i) G_k(-\log(a) - s'_i)).
\end{aligned}$$

The tricky line here is the third one. The integration increases the index of all of the integrands by one, so that really the index l should start at 1, but then we can add the $l = 0$ terms which are the same for both summands and hence net to zero. So, substituting this back in, we obtain

$$\begin{aligned}
A &= -\frac{1}{\log(a)} \sum_{k=1}^N G_k(-\log(a)) - \exp(-s'_i) \frac{1}{\log(a)} \sum_{k=1}^N \frac{N-k}{N} G_k(-\log(a)) \\
&\quad - \frac{1}{\log(a)} \exp(-s'_i) \sum_{k=1}^N \frac{k}{N} G_k(-\log(a) - s'_i) \\
&= (1 - \exp(-s'_i)) \lambda(1) + \exp(-s'_i) \frac{1}{\log(a)} \sum_{k=1}^N \frac{k}{N} G_k(-\log(a)) \\
&\quad - \frac{1}{\log(a)} \exp(-s'_i) \sum_{k=1}^N \frac{k}{N} G_k(-\log(a) - s'_i),
\end{aligned}$$

where we have used our formula for $\lambda(1)$ from the previous section. Substituting back into the expected transfer, we conclude that

$$\int_{s_{-i}} t_i(s'_i, s_{-i}) \exp(-\Sigma s_{-i}) ds_{-i} = -\frac{1}{\log(a)} \sum_{k=1}^N \frac{k}{N} G_k(-\log(a)) + \frac{1}{\log(a)} \sum_{k=1}^N \frac{k}{N} G_k(-\log(a) - s'_i).$$

This implies that

$$C = \int_{s_{-i}} \frac{\partial t_i(s'_i, s_{-i})}{\partial s'_i} \exp(-\Sigma s_{-i}) ds_{-i} = -\frac{1}{\log(a)} \sum_{k=1}^N \frac{k}{N} g_k(-\log(a) - s'_i).$$

Thus, to evaluate the derivative of (47), we just have to evaluate

$$D = \int_{s_{-i}} \min\{1, a \exp(-s_i - \Sigma s_{-i})\} \frac{\partial q_i(s'_i, s_{-i})}{\partial s'_i} \exp(-\Sigma s_{-i}) ds_{-i}.$$

When $s_i \leq -\log(a)$, we can decompose this expression as follows:

$$\begin{aligned} D &= \int_{\{s_{-i} | \Sigma s_{-i} \leq -\log(a) - s'_i\}} a \exp(-s_i) \frac{\partial q_i(s'_i, s_{-i})}{\partial s_i} ds_{-i} \\ &\quad + \int_{s_{-i}} \frac{\partial q_i(s'_i, s_{-i})}{\partial s'_i} \exp(-\Sigma s_{-i}) ds_{-i} \\ &\quad - \int_{\{s_{-i} | \Sigma s_{-i} \leq -\log(a) - s'_i\}} \frac{\partial q_i(s'_i, s_{-i})}{\partial s_i} \exp(-\Sigma s_{-i}) ds_{-i}. \end{aligned}$$

If $s'_i \leq s_i$, then $\partial q_i / \partial s'_i = -1 / \log(a)$, and we can simplify this further to

$$\begin{aligned} D &= -\frac{1}{\log(a)} a \exp(-s_i) \int_{\{s_{-i} | \Sigma s_{-i} \leq -\log(a) - s'_i\}} ds_{-i} \\ &\quad - \frac{1}{\log(a)} \frac{1}{N} \sum_{k=1}^N G_{k-1}(-\log(a) - s'_i) - G_{N-1}(-\log(a) - s'_i). \end{aligned}$$

Now, the first integral is simply the area of the $n - 1$ -dimensional regular simplex, where the length of the sides is $L = -\log(a) - s'_i$, which is simply $L^n / n!$. So that this simplifies even further to

$$\begin{aligned} D &= -\frac{1}{\log(a)} \left(g_N(-\log(a) - s_i) + \frac{1}{N} \sum_{k=1}^N G_{k-1}(-\log(a) - s'_i) - G_{N-1}(-\log(a) - s'_i) \right) \\ &= -\frac{1}{\log(a)} \left(g_N(-\log(a) - s_i) - \frac{1}{N} \sum_{k=1}^N g_{k-1}(-\log(a) - s'_i) + \sum_{k=0}^{N-1} g_k(-\log(a) - s'_i) \right) \\ &= -\frac{1}{\log(a)} \left(g_N(-\log(a) - s_i) - \frac{1}{N} \sum_{k=1}^N \sum_{l=0}^{k-1} g_l(-\log(a) - s'_i) + \sum_{k=0}^{N-1} g_k(-\log(a) - s'_i) \right) \\ &= -\frac{1}{\log(a)} \left(g_N(-\log(a) - s_i) + \sum_{k=0}^{N-1} \frac{k}{N} g_k(-\log(a) - s'_i) \right). \end{aligned}$$

Thus, in the case where $s'_i \leq s_i \leq -\log(a)$, we have that

$$\begin{aligned} \frac{\partial U_i(s_i, s'_i)}{\partial s'_i} &= D - C \\ &= -\frac{1}{\log(a)} \left(g_N(-\log(a) - s_i) + \sum_{k=0}^{N-1} \frac{k}{N} g_k(-\log(a) - s'_i) - \sum_{k=1}^N \frac{k}{N} g_k(-\log(a) - s'_i) \right) \\ &= -\frac{1}{\log(a)} (g_N(-\log(a) - s_i) - g_N(-\log(a) - s'_i)) \geq 0 \end{aligned}$$

since g_N is monotonically decreasing and $s_i \geq s'_i$.

It is straightforward to prove the other cases as well. When $s_i \leq -\log(a)$ and $s'_i \geq s_i$, we have that $\partial q_i(s'_i, s_{-i})/\partial s'_i \leq -1/\log(a)$, so that

$$\frac{\partial U_i(s_i, s'_i)}{\partial s'_i} \leq D - C \leq 0$$

since $s_i \leq s'_i$. The final case is when $s_i > -\log(a)$. In this case,

$$\begin{aligned} D &= \int_{s_{-i}} \frac{\partial q_i(s'_i, s_{-i})}{\partial s'_i} \exp(-\Sigma s_{-i}) ds_{-i} \\ &= -\frac{1}{\log(a)} \frac{1}{N} \sum_{k=1}^N G_{k-1}(-\log(a) - s'_i) \\ &= -\frac{1}{\log(a)} \left(1 - \frac{1}{N} \sum_{k=1}^N \sum_{l=1}^{k-1} g_l(-\log(a) - s'_i) \right) \\ &= -\frac{1}{\log(a)} \left(1 - \sum_{k=1}^N \frac{N-k}{N} g_k(-\log(a) - s'_i) \right). \end{aligned}$$

So, we conclude that

$$\begin{aligned} \frac{\partial U_i(s_i, s'_i)}{\partial s'_i} &= -\frac{1}{\log(a)} \left(1 - \sum_{k=1}^N g_k(-\log(a) - s'_i) \right) \\ &= -\frac{1}{\log(a)} G_N(-\log(a) - s'_i) \end{aligned}$$

which is positive as long as $s'_i \leq -\log(a)$, and is zero for $s'_i > -\log(a)$. So again, we conclude that truth-telling is optimal, and the mechanism with the Shapley allocation rule is an incentive compatible direct mechanism. As a result, the primal is feasible.

C Redux under CLT-normalized messages

The critical difference from Section 4 is that the multiplier on local incentive compatibility becomes $1/\sqrt{N}$, rather than 1. The dual constraints therefore become

$$\Xi^C(m) = \sum_{i=1}^N \left[\frac{1}{\sqrt{N}} \frac{\partial t_i^C(m)}{\partial m_i} - t_i^C(m) \right] = \frac{1}{\sqrt{N}} \sum_{i=1}^N \frac{\partial q_i^C(m)}{\partial m_i} \widehat{w}_N^C(\Sigma m) - \lambda^C(\widehat{w}_N^C(\Sigma m)) - cQ^C(m).$$

The usual envelope argument implies that $\frac{d}{dx} \lambda^C(x) = \mu^C(x)$, where

$$\mu^C(x) = \frac{1}{\sqrt{N}} \sum_{i=1}^N \frac{\partial q_i^C(m)}{\partial m_i}.$$

Note that in the fully-revealing case,

$$q_i^C(m) = \frac{\sqrt{N}m_i + 1}{\sqrt{N}\Sigma m + N},$$

so that

$$\begin{aligned} \mu^C(x) &= \frac{1}{\sqrt{N}} \sum_{i=1}^N \sqrt{N} \frac{m_{-i} + N - 1}{(\sqrt{N}\Sigma m + N)^2} \\ &= \frac{N - 1}{\sqrt{N}x + N}. \end{aligned}$$

On the other hand, when the good is being linearly rationed on the low region $[-\sqrt{N}, x^*]$,

$$q_i^C(m) = \frac{\sqrt{N}m_i + 1}{\sqrt{N}x^* + N},$$

so

$$\begin{aligned} \mu^C(x) &= \frac{N}{\sqrt{N}} \frac{\sqrt{N}}{\sqrt{N}x^* + N} \\ &= \frac{N}{\sqrt{N}x^* + N}. \end{aligned}$$

It is easily verified that if

$$\xi_i^C(m) = t_i^C(m) - \frac{1}{\sqrt{N}} \frac{\partial t_i^C(m)}{\partial m_i},$$

then

$$t_i^C(m) = \frac{1}{\sqrt{N}} \exp(\sqrt{N}m_i) \int_{x=0}^{m_i} \exp(-\sqrt{N}x) \xi_i^C(x, m_{-i}) dx.$$

The conditions on the growth functions are now that

$$\sum_{i=1}^N \xi_i^C(m) = \Xi^C(\Sigma m);$$

$$\int_{x=0}^{\infty} \exp(-\sqrt{N}x) \xi_i^C(x, m_{-i}) dx = 0.$$

We solve these equations as before, except that now we use the distributions g_n^C of the expectations of n independent draws from the distribution $1 - \exp(\sqrt{N}(z+1))$ on $[-1, \infty)$, which has hazard rate \sqrt{N} rather than 1. Note that the distribution g_n^C is approximately a normal distribution when n is large. Thus, the transfers end up being

$$t_i^C(m) = \frac{1}{\sqrt{N}} \frac{1}{2^{N-1}} \sum_{I \subseteq \{1, \dots, N\} \setminus \{i\}} \int_{x=0}^{\infty} (\Xi^C(x + \Sigma m_I) - \Xi^C(x + \Sigma m_I + m_i)) g_{N-|I|}^C(x) dx.$$

We are now able to look at the limit transfers. Consider the object

$$\lim_{N \rightarrow \infty} \frac{N t_i^C(m)}{m_i},$$

which is the transfer per unit message, and scaled up by N so that it is in per capita terms. This object is approximately

$$\begin{aligned} \tau(F) &\approx \lim_{N \rightarrow \infty} \sqrt{N} \int_{\delta=0}^1 \int_{\alpha=0}^1 \int_{x=0}^{\infty} \frac{\Xi^C(\delta x + \alpha) - \Xi^C(\delta + \alpha + m_i)}{m_i} \phi(x) dx F(d\alpha|\delta) d\delta \\ &\approx - \lim_{N \rightarrow \infty} \sqrt{N} \int_{\delta=0}^1 \int_{\alpha=0}^1 \int_{x=0}^{\infty} \frac{d}{dx} \Xi^C(\delta x + \alpha) \phi(x) dx F(d\alpha|\delta) d\delta, \end{aligned}$$

where Ξ^C is the total excess growth in CLT-normalized units. This is

$$\frac{d}{dx} \Xi^C(x) = \frac{d}{dx} \mu^C(x) \widehat{w}_N^C(x) - c \frac{d}{dx} Q^C(x).$$

When the good is fully supplied, only the first term is non-zero and we have

$$\begin{aligned} \sqrt{N} \frac{d}{dx} \Xi^C(x) &= -\sqrt{N} \frac{(N-1)\sqrt{N}}{(\sqrt{N}x + N)^2} \widehat{w}_N(x) \\ &\approx -\widehat{w}_\infty(x) \end{aligned}$$

when N is large, whereas if the good is linearly rationed on the low region, only the second term is non-zero, and since

$$Q^C(x) = \frac{x + \sqrt{N}}{x^* + \sqrt{N}},$$

we obtain

$$\begin{aligned}\sqrt{N} \frac{d}{dx} \Xi^C(x) &= -\sqrt{N} \frac{c}{x^* + \sqrt{N}} \\ &\approx -c\end{aligned}$$

when N is large. Thus, in both cases, we find that $\sqrt{N} \frac{d}{dx} \Xi^C(x)$ converges to $\bar{w}_\infty(x)$, the graded value function in the limit. And hence

$$\tau(F) = \int_{\delta=0}^1 \int_{\alpha=0}^1 \int_{x=0}^{\infty} \bar{w}(\delta x + \alpha) \phi(x) dx F(d\alpha|\delta) d\delta.$$

D Notation

Symbol	Meaning
N	Number of bidders
Σx	Sum of elements of x
H	Prior distribution of the common value
G_k	Erlang CDF with k draws and unit arrival rate
g_k	Erlang PDF with k draws and unit arrival rate
c	Seller's cost
U	Bidder surplus
Π	Profit
$\bar{\Pi}$	Can-keep optimal profit guarantee
$\hat{\Pi}$	Must-sell optimal profit guarantee
\mathcal{T}	Type space
S_i	Set of signals
π	Joint distribution of types
w	Value function
γ	Gains function
Γ	Integrated gains function
\mathcal{M}	Mechanism
q_i	Allocation rule
Q	Total supply
t_i	Transfer rule
β	Bidder strategy
$(\bar{\mathcal{M}}, \bar{\mathcal{T}}, \bar{\beta})$	Can-keep solution
$(\hat{\mathcal{M}}, \hat{\mathcal{T}}, \bar{\beta})$	Must-sell solution
ψ_i	Virtual gains measure
$\hat{\gamma}$	Fully-revealing gains function
$\bar{\gamma}$	Graded gains function
\hat{w}	Fully-revealing value function
\bar{w}	Graded value function
μ	Total allocation sensitivity
ξ_i	Excess growth
Ξ	Total excess growth
λ	Linear approximation to the optimal profit guarantee
ζ	A permutation of the bidders
Z	Set of permutations
p	Price function

Table 1: Notation