

Selling to Intermediaries: Optimal Auction Design in a Common Value Model*

Dirk Bergemann Benjamin Brooks Stephen Morris

August 11, 2017

Abstract

We characterize revenue maximizing auctions when the bidders are intermediaries who wish to resell the good. The bidders have differential information about their common resale opportunities: each bidder privately observes an independent draw of a resale opportunity, and the highest signal is a sufficient statistic for the value of winning the good. If the good must be sold, then the optimal mechanism is simply a posted price at which all bidders are willing to purchase the good, and all bidders are equally likely to be allocated the good, irrespective of their signals. If the seller can keep the good, then under the optimal mechanism, all bidders make the same expected payment and have the same expected probability of receiving the good, independent of the signal. Conditional on the good being sold, the allocation discriminates in favor of bidders with *lower* signals. In some cases, the optimal mechanism again reduces to a posted price. The model provides a foundation for posted prices in multi-agent screening problems.

KEYWORDS: Optimal auction, intermediaries, posted price, guaranteed demand auction, common values, revenue maximization, revenue equivalence, first-price auction, second-price auction, resale, maximum value game, descending auction, local incentive constraints, global incentive constraints.

JEL CLASSIFICATION: C72, D44, D82, D83.

*Bergemann: Department of Economics, Yale University, dirk.bergemann@yale.edu; Brooks: Department of Economics, University of Chicago, babrooks@uchicago.edu; Morris: Department of Economics, Princeton University, smorris@princeton.edu. We acknowledge financial support through NSF Grant ICES 1215808. We have benefited from conversations with Gabriel Carroll, Phil Haile and Ilya Segal.

1 Introduction

1.1 Motivation and Main Results

We study the design of revenue maximizing auctions for the sale of a single indivisible good when the bidders are intermediaries. As middlemen, the bidders have no private value for the good but rather intend to sell it to one of many possible final users. The final sale opportunities are common to all of the intermediaries, and thus they share a common value for the object. Yet, at the time of the auction, the bidders have incomplete and differential information about those final sale opportunities.

In our model, each bidder receives an independent signal, which can be interpreted as the profitability of a particular sale opportunity. The bidders have a pure common value which is equal to the maximum of their signals, i.e., the best such opportunity. Thus, whoever wins the good will have access to all of the downstream opportunities observed by all of the intermediaries. More broadly, the signals can be viewed as partial information about the strength of the secondary market. The key feature of the information structure is that conditional on the highest signal, the other signals contain no additional information about the common value.

The main contribution of our paper is a characterization of the revenue maximizing auction format. Selling to intermediaries is markedly different from the auction design problems studied elsewhere in the literature. The seminal papers of Riley and Samuelson (1981) and Myerson (1981) focus on the case where bidders' values are *private*, i.e., each bidder's signal fully reveals their own value for the good. They show that the optimal auction induces the bidders to separate, with higher signal bidders being more likely to receive the good and also paying more.¹ In contrast, when the bidders are intermediaries, such screening turns out to be prohibitively costly. In many cases, it is optimal to not screen at all. Even when the seller does screen, it is more to determine whether the good should be sold to anyone, and not to identify a single exclusive winner. In fact, all bidders make the same expected payment and receive the object with the same expected probability irrespective of their private information.

Our first main result provides a necessary and sufficient condition when the optimal mechanism reduces to a remarkably simple format: a posted price. Under this condition, revenue is maximized by setting a posted price equal to the interim expectation of the value conditional on receiving the *lowest* signal. We refer to it as an *inclusive* posted price as

¹Bulow and Klemperer (1996) extend this characterization to a class of interdependent value environments that excludes the present model. We will discuss the connection with their results further below.

every type is willing to buy the object at this price. In particular, no type is excluded from the allocation.

Our second main result is that if the seller has to sell the good, that is if the probability of assigning the good has to be equal to one, then the inclusive posted price as described above is again the optimal selling mechanism. This result holds for all possible distributions of the value of the good.

These two results are closely related. Theorem 2 below shows that the highest inclusive posted price maximizes revenue among all auctions that always allocate the good. Always allocating the good is optimal for the seller when the social gains from trade always exceed the cost of screening. Intuitively, this is the case when the difference between the seller's value for the good and the smallest common value of the bidders' is sufficiently large. Our necessary and sufficient condition precisely describes when this is the case, as a function of the primitives of the model: the distribution of the common value and the number of bidders.

Our third main result is a general characterization of the optimal mechanism. Theorem 3 shows that the optimal auction can be compactly described as an indirect mechanism which we refer to as the *guaranteed demand auction* (GDA). The bidders must pay an up-front fee to participate in the auction, and after paying this fee, there are no further transfers. The allocation of the good is then determined according to the following procedure. Each bidder “demands” a minimum probability that she receive the good. It is possible to demand any probability between 0 and a maximum value that is (weakly) less than $1/N$. This maximum demand is a parameter of the mechanism that is fixed in advance by the seller. If no one makes a positive demand, then the seller keeps the good. If at least one bidder makes a positive demand, then the good is allocated with probability one. The bidder with the highest demand then receives the good with the *exact* probability that she requested and the remaining (complementary) probability is shared equally among the remaining bidders and independent of their reports. Thus, the report is a “guarantee” in the sense that each bidder will receive the good with at least the demanded probability, no matter what, even if others make very low demands, but the probability could be higher if someone else makes a larger demand.

Conditional on all bidders entering, this mechanism has an equilibrium in which bidders with very low signals pool at a demand of zero, and higher signal bidders separate according to a monotonic pure strategy, with the highest type making the highest possible demand. The optimal mechanism sets the highest transfer such that all bidders are willing to enter. A special case of the mechanism is when the maximum demand is exactly $1/N$, in which case the equilibrium collapses to all types making the maximum demand, so that the mechanism

is equivalent to setting a posted price equal to N times the entry fee (with the bidders being equally likely to win the tie break and pay the price).

1.2 Intuition and Methodology

Let us try to explain why the GDA maximizes revenue. It turns out that in equilibrium, the interim probability of receiving the good is the same for all demands, including a demand of zero. Thus, the optimal auction always shares two features in common with the posted price: the payment and interim allocation probabilities are the same for all types. However, the allocation induced by the GDA is such that conditional on the good being allocated, it is more likely to be allocated to a bidder with a type that is less than the highest type, since the highest type receives their guaranteed probability (which is always weakly less than $1/N$).

Thus, there is a sense in which the optimal auction discriminates in favor of bidders with *lower* signals, since they are more likely to win conditional on the good being allocated. This is in contrast to first- and second-price auctions, which discriminate in favor of higher signal bidders, and are optimal when bidders have independent private values (IPV). There are at least two distinct reasons why our results are so different from those obtained in the IPV setting.

First, simple accounting dictates that revenue is equal to the total surplus generated by the auction minus the sum of the bidders' surpluses. The revenue maximizing allocation trades off the surplus generated by the allocation against the surplus that bidders receive in equilibrium. In the IPV model, bidders with higher signals have higher values, so that efficiency is a strong force pushing the allocation to discriminate in favor of bidders with higher signals. The present model, in contrast, is one of common values: all of the bidders have the same value for the good, so total surplus only depends on whether the good is allocated, and not on whom the good is allocated to.

Second, the revenue maximizing auction will tend to direct the allocation towards those bidders who receive smaller *information rents*, i.e., those bidders for whom it is less costly to incentivize them to reveal their type. As is well known, one way to compute these information rents is by aggregating up discounts that bidders must receive in order to support the equilibrium strategy as a local optimum, the so-called envelope condition. This local discount is equal to the rate of change in the social surplus as a function of the bidder's type, evaluated at the equilibrium allocation. With private values, the discount reduces to the probability that the bidder receives the good, since the signal is equal to the value. With common values, or interdependent values in general, the rate at which information

rents accrue is equal to the allocation probability multiplied by the change in the value as a function of the signal. When the bidders are intermediaries, however, the value is insensitive to all signals except the highest, meaning that bidders do not accrue information rents except when they have the highest signal. Thus, it is only the highest type who can claim an information rent from such local incentive constraints. This creates a strong force for allocating the good to low signal bidders, which is precisely what happens in the equilibrium of the GDA.

Another way of interpreting the differences in information rents is the following: When the bidders are intermediaries, first- and second-price auctions tend to create lots of information rents for the bidders due to a strong *winner's curse* effect. Bidders are more likely to win when others' signals are low, which leads to an adverse inference about the value from winning, and induces bidders to shade. In contrast, under the GDA, bidders are more likely to win when others' signals are higher, resulting in a *winner's blessing*, in which winning the good leads to a positive inference about the value of the good and induces the bidders to pay higher entry fees.

This structure on the information rents forces us to use different arguments than those employed by Myerson (1981) and his methodological descendants. The standard approach to identifying the optimal auction (and one that succeeds when the signal distribution is sufficiently well-behaved) is to implement the allocation that pointwise maximizes the difference between social surplus and local information rents. But when bidders are intermediaries, the local information rents are zero for any bidder that does not have the highest signal, and moreover, allocating to any bidder would generate the same gain in social surplus. This suggests an auction in which the good is always allocated, but it is never allocated to the highest type. Such an allocation would, however, violate *global* incentive constraints, as, for example, the highest type would prefer to misreport as the lowest type.

It is thus necessary to incorporate non-local incentive constraints into the seller's optimization problem. Remarkably, it is sufficient to consider a one-dimensional family of incentive constraints corresponding to the following deviations in the normal form: instead of reporting one's true type, a bidder misreports a lower type which is randomly redrawn from the prior distribution censored at the bidder's true signal. We derive a lower bound on bidder surplus in any auction that deters such deviations, as a function of the probabilities with which the good is allocated conditional on each value. Since these probabilities also pin down total surplus, we also have an upper bound on revenue. These lower and upper bounds form the content of Theorem 1. As the optimal mechanism has to satisfy these specific incentive compatibility constraints, but others as well, such as the local constraints, the resulting upper bound in principle need not be sharp. We show, however, that the GDA described

above attains the resulting upper bound, thus verifying that it is an optimal auction. In deriving our results, we first construct the upper bound in Section 3, and then construct mechanisms that attain it in Sections 4 and 5.

We report a number of further results. We characterize the set of optimal mechanisms, in addition to posted prices and the GDA, and we also show that the GDA can be implemented by a version of the descending clock auction. In the present version it is the probability of receiving the good that descends rather than the price. We also discuss behavior under this model of information in first- and second-price auctions, in comparison to behavior under the benchmark independent private value model. This comparison yields additional insights into the robust revenue properties of first-price auctions. Further results are discussed below as we relate our work to the literature.

1.3 Relationship to Literature

This work is most closely related to the literature on optimal auction design. In addition to the work on private values cited above, the celebrated paper of Bulow and Klemperer (1996) extends the first-order approach of Myerson (1981) to certain interdependent value models, where the bidders' values are increasing and differentiable functions of all bidders' signals. Specialized to the common value setting, their result says that variants of English auctions are optimal as long as the information rent is smaller for higher signal bidders. This condition is likely to be violated when the bidders are intermediaries who are selling in a common downstream market, since then the information of more optimistic bidders is likely to trump that of less optimistic bidders. To our knowledge, this is the first paper to extend the theory of optimal auctions in common value environments beyond the case of decreasing information rents.

The impact of the increasing information rents on the revenue across mechanisms is underscored when we revisit the comparison of mechanisms suggested by Bulow and Klemperer (1996). Their main result showed that the revenue from a standard absolute auction with $N + 1$ bidders always exceeds the revenue from an optimal auction with N bidders. With increasing rather than decreasing information rents, this results is reversed. We show that the revenue from a standard absolute auction with $N + K$ bidders is always below the revenue from an optimal auction with N bidders, for any $K > 0$.

As described above, our leading interpretation of this model is that the bidders are intermediaries who see independent draws of final sale opportunities in a common market. Whoever wins the good will obtain the best of these opportunities. A related interpretation is the following. The model is one of independent private values, so that the bidders are the

end users of the good. But in contrast to the standard model, the auction will be followed by a frictionless resale market, in which values become complete information and the interim owner of the good can make a take-it-or-leave-it offer to the bidder with the highest value. Thus, whoever wins the good in the first stage will earn revenue from resale equal to the highest of the bidders' private values. Now, as the allocation rule of the optimal mechanism favors lower signals, in effect, it induces a more active resale market. The reason is that only the bidder with the highest signal has private information that is payoff relevant, so that discriminating against this bidder reduces the total amount of information rents that bidders receive. This model of resale abstracts from the bidders' incentives to signal their values through the outcome of the auction, and instead emphasizes the common value structure that arises from the bidders' ability to resell the good in the same market. A similar informational model has been used by Gupta and LeBrun (1999) and Haile (1999, 2003) to study the effect of resale in first-price auctions. By contrast, we treat the mechanism as an endogenous object, and derive the auction format that a revenue maximizing seller would use. A similar perspective is taken in recent work by Carroll and Segal (2016) who design the optimal auction in the presence of a resale market. They derive the optimal auction as a maxmin problem where Nature chooses the information and bargaining protocol in the resale market that is least favorable to the revenue maximizing seller. Their solution and arguments are very different from the ones presented here. In particular, they establish that the least favorable resale market is then one where the *bidder with the highest value*—independent of his ownership—has the bargaining power and complete information. Thus, he can make a take it or leave it offer to the current owner of the object. By contrast, our resale interpretation implicitly requires that it is the *current owner* of the object—independent of his value—who has the bargaining power and has complete information regarding the value of the other bidders.

A version of our model, one where the final value to the bidder is exactly the maximum of N independent signals, has previously been studied by Bulow and Klemperer (2002). They were interested in the implications of the winner's curse for the revenue comparison across standard selling mechanisms and suggested that the maximum model might be a more plausible model for oil and mineral rights than the better-known model in which bidders observe the value with conditionally independent noise (i.e., *the* mineral rights model). They showed that an inclusive posted price yields higher revenue than the second-price auction without reserve prices, as an example of a case where the winner's curse effect is so severe that revenue would be greater if the seller did not screen at all.² In a sense, our constructions and arguments pursue the ideas of Bulow and Klemperer (2002) to their logical conclusion.

²See also Harstad and Bordley (1996) and Campbell and Levin (2006) for related results.

We show that sometimes the posted price mechanism constitutes the optimal mechanism, and it is not possible to improve on Bulow and Klemperer’s result. But more generally, the seller can do even better than posted prices by using the GDA.

The strength of the winner’s curse in the maximum of N independent signals model is documented in the first-price auction in our earlier work (Bergemann, Brooks, and Morris, 2017). There, we show that among all common prior type spaces for given distribution of common values, the revenue in the first-price auction is minimized by the type space that generates N independent signals such that the common value coincides with the highest signal. The symmetry in the binding constraints in the maximum value environment is noteworthy. In the first- and second-price auction, *all* upward deviation constraints are binding, whereas in the optimal auction, *all* downward deviation constraints are binding.

The rest of this paper is organized as follows. Section 2 describes our model. Section 3 presents lower bounds on bidder surplus and upper bounds on revenue. Section 4 gives necessary and sufficient conditions under which posted prices attain these bounds. Section 5 shows that the guaranteed demand auction attains the bounds more generally. Section 6 concludes with a discussion of properties of the optimal mechanism, and also draws additional connections to the literature in auction theory.

2 Model

2.1 Environment

There are N potential bidders for a single unit of a good, indexed by $i \in \mathcal{N} = \{1, \dots, N\}$. Each bidder receives a signal $s_i \in S = [\underline{s}, \bar{s}] \subseteq \mathbb{R}_+$ about the good’s value. The bidders’ signals s_i are independent draws from an absolutely continuous cumulative distribution $F(s_i)$ with density $f(s_i)$. The bidders all assign the same value to the good, which is a function of the maximum of the signals:

$$v(s_1, \dots, s_N) = v(\max\{s_1, \dots, s_N\}). \tag{1}$$

The distribution of signals, $F(s_i)$, induces a distribution $G_N(x)$ over the maximum signal from N independent draws:

$$G_N(x) = (F(x))^N.$$

We denote the associated density by:

$$g_N(x) = N(F(x))^{N-1} f(x).$$

Note that we could always normalize the model so that v is the identity by changing the distribution of the signals to

$$\widehat{F}(s) = F(v^{-1}(s)),$$

so that really there is only one parameter in the model, which is the distribution of signals. We will apply this normalization henceforth, and simply assume that the value is equal to the maximum signal. We have introduced the model in this manner to emphasize the range of possible interpretations, which includes the case where the maximum signal is a sufficient statistic for the value, but is not the value itself. For example, it could be that each bidder actually observes K independent signals, with M more being observed by the winner of the auction, and the value is the maximum of all $NK + M$ signals. This would also fall within our framework, where each bidder's signal is identified with the conditional expectation of the highest of all signals, conditional on the highest of the bidder's own K signals being greater than all of the other bidders' signals. We will return to such an interpretation when we consider comparative statics in Section 6.

The bidders are expected utility maximizers, with quasilinear preferences over the good and transfers. Thus, the ordering over pairs (q, t) of probability q of receiving the good and net transfers to the seller is represented by the utility index:

$$u(s, q, t) = v(s)q - t.$$

2.2 Mechanisms

The good is sold via an auction. For much of our analysis, and in particular for constructing bounds on revenue and bidder surplus in Theorem 1, we will restrict attention to *direct mechanisms*, whereby each bidder simply reports his own signal, and the set of possible message profiles is S^N . This is without loss of generality, by the revelation-principle arguments as in Myerson (1981). The probability that bidder i receives the good given signals $s \in S^N$ is $q_i(s) \geq 0$, with $\sum_{i=1}^N q_i(s) \leq 1$. Bidder i 's transfer is $t_i(s)$ and the interim expected transfer is denoted by:

$$t_i(s_i) = \int_{s_{-i} \in S^{N-1}} t_i(s_i, s_{-i}) f_{-i}(s_{-i}) ds_{-i}, \quad (2)$$

Bidder i 's surplus from reporting a signal s'_i when his true signal is s_i is

$$u_i(s_i, s'_i) = \int_{s_{-i} \in S^{N-1}} q_i(s'_i, s_{-i}) v(s_i, s_{-i}) f_{-i}(s_{-i}) ds_{-i} - t_i(s'_i),$$

and $u_i(s_i) = u_i(s_i, s_i)$ is the payoff from truth-telling. Ex-ante bidder surplus is simply

$$U_i = \int_{s_i=\underline{s}}^{\bar{s}} u_i(s_i) f(s_i) ds_i,$$

so that total bidder surplus is

$$U = \sum_{i=1}^N U_i.$$

A direct mechanism $\{q_i, t_i\}_{i=1}^N$ is *incentive compatible* (IC) if

$$u_i(s_i) = \max_{s'_i} u_i(s_i, s'_i), \quad (3)$$

for all i and $s_i \in S$. This is equivalent to requiring that reporting one's true signal is a Bayes Nash equilibrium. The mechanism is *individually rational* (IR) if

$$u_i(s_i) \geq 0,$$

for all i and $s_i \in S$.

Our implementation results in Theorems 2 and 3 will be stated in terms of indirect mechanisms, namely posted prices and the guaranteed demand auction. These mechanisms have simple pure strategy equilibria, and the implied direct mechanism will be clear from context.

2.3 The Seller's Problem

The seller's objective is to maximize expected revenue across all IC and IR mechanisms. Under a mechanism $\{q_i, t_i\}_{i=1}^N$, expected revenue is

$$R = \sum_{i=1}^N \int_{s_i \in S} t_i(s_i) f(s_i) ds_i.$$

Note that because values are common, total surplus only depends on whether the good is allocated, not the identity of the bidder that receives the good. Let us denote by $\bar{q}_i(v)$ the probability that the good is allocated to bidder i , conditional on the true value being v , and let

$$\bar{q}(v) = \sum_{i=1}^N \bar{q}_i(v) \quad (4)$$

be the corresponding total probability that some bidder receives the good. Total surplus is simply

$$TS = \int_{v=s}^{\bar{s}} v \bar{q}(v) g_N(v) dv,$$

and revenue is obviously $R = TS - U$.

3 Bounds on Bidder Surplus and Revenue

Our first main result, Theorem 1, is a lower bound on the bidders' net utility U as a function of the allocation probability \bar{q} . This lower bound is obtained by considering a particular class of incentive constraints, which we shall explain shortly. As a result, an upper bound on revenue R can be obtained by choosing \bar{q} to maximize the difference between total surplus and bidder surplus. Our subsequent results will then verify that this upper bound is tight, by explicitly constructing IC and IR mechanisms that achieve the bound.

3.1 Revenue Equivalence and the Failure of the First-Order Approach

At a high level, this guess-and-verify approach is the same one that has been employed to great effect in the literature on auction design. The difference is the class of incentive constraints that are used. Before giving our analysis, it is helpful to briefly review the standard approach, and explain why it fails to pin down optimal revenue in the current common value setting. The standard approach uses the constraint that truthful reporting is a local maximizer of the bidder's utility. When values are private, this constraint implies that $u'_i(s_i)$, the derivative of the indirect utility, is simply the probability that the type s_i is allocated the good conditional on s_i :

$$u'_i(s_i) = q_i(s_i) \equiv \int_{s_{-i}} q_i(s_i, s_{-i}) f_{-i}(s_{-i}) ds_{-i}. \quad (5)$$

The logic is simply that of the envelope theorem: an increase in the true type must entail an increase in the reported type as well. But because of incentive compatibility, small changes in the reported type have no effect on utility, so that the change in bidder surplus is only the "direct" effect through the value at the truthful allocation. In the private value model, signals are normalized to be equal to private values, so that the value increases one for one with the signal, and the change in bidder surplus is therefore simply the interim probability of being allocated the good, $q_i(s_i)$.

In the present common value model, all bidders value the good at the maximum signal, so an increase in s_i only has a positive direct effect on the value when s_i is higher than all s_j for $j \neq i$. Thus, the derivative of the indirect utility changes to

$$u'_i(s_i) = \hat{q}_i(s_i) \equiv \int_{s_{-i}} \mathbb{I}_{s_i \geq \max_{j \neq i} s_j} q_i(s_i, s_{-i}) f_{-i}(s_{-i}) ds_{-i}. \quad (6)$$

The function $\hat{q}_i(s_i)$ is the expected probability that bidder i is allocated the good *and* bidder i has the highest signal. For future reference, we state this as a formal result:³

Lemma 1 (Revenue Equivalence).

In any incentive compatible direct mechanism,

$$u_i(s_i) = u_i(\underline{s}) + \int_{x=\underline{s}}^{s_i} \hat{q}_i(x) dx. \quad (7)$$

The standard approach would be to argue that it is without loss to take $u_i(\underline{s}) = 0$, and then use this expression to maximize the difference between total and bidder surplus across all allocations, and then construct a globally incentive compatible mechanism that attains the resulting upper bound on revenue. Unfortunately, this analysis is too permissive to derive a tight bound in our model. Consider the following allocation:

$$q_i^{\min}(s) = \mathbb{I}_{s_i < \min_{j \neq i} s_j},$$

i.e., the good is allocated to the bidder with the *lowest* signal. This results in $\hat{q}_i \equiv 0$, so Lemma 1 implies that bidders get zero information rents. But the allocation would be efficient, so that the seller obtains all of the efficient surplus as revenue!

As we shall see, the seller cannot extract all of the surplus with an incentive compatible and individually rational mechanism. This seeming paradox is resolved by the observation that while the supposedly optimal allocation q^{\min} would satisfy the local first-order constraints used to derive (6), it would violate non-local constraints in (3). Specifically, bidders would benefit by deviating to strictly lower reports. Consider, for example, the highest type \bar{s} . Under the proposed mechanism, the highest type is never allocated the good, and must therefore receive zero surplus. But by misreporting her type as \underline{s} , she is guaranteed a positive probability of being allocated the good. Moreover, with probability one, the value is strictly

³A revenue equivalence result for a general interdependent value model was proven by Bulow and Klemperer (1996) under the assumption that the value is a differentiable function of the independent signals. While our value function is not differentiable, a revenue equivalence result still holds, and it can be obtained using essentially the same steps as in Myerson (1981). For the sake of completeness, we have included a formal proof in the Appendix.

higher than what it would be if the true type were \underline{s} , so that the resulting surplus $u_i(\bar{s}, \underline{s})$ must be strictly greater than $u_i(\underline{s}) = 0$.

This example shows that in order to obtain a tight upper bound on revenue, we will have to incorporate global constraints into the optimization problem. Fortunately, our efforts to derive the revenue equivalence formula (7) are not a complete loss, as we will subsequently use this result to characterize the class of optimal mechanisms.

3.2 Bounds on Bidder Surplus and Revenue

The question is: which global constraints matter for pinning down optimal revenue? The local analysis suggests that the critical constraints might be those corresponding to downward deviations: The bidders only accrue information rents when they are allocated the good when they have the highest value, so that the seller wants to distort the allocation to lower signal bidders as much as possible. But if the allocation is too skewed, then bidders would want to deviate by reporting strictly lower types. Note that this intuition is in some sense the opposite of what happens in the private value auction model, in which the optimal auction typically discriminates in favor of higher types. An important difference is that when values are private, it is not just whether but also to whom the good is allocated that determines total surplus.

In principle we might have to consider all of the deviations from some type s_i to all lower types $s'_i \leq s_i$. We will show that it is without loss of generality to focus on the following one-dimensional family of deviations in the normal form: instead of reporting the true signal s_i , report a random $s'_i \in [\underline{s}, s_i]$ that is drawn from the truncated prior $F(s'_i)/F(s_i)$. We will refer to this deviation as *misreporting a redrawn lower signal*. Obviously, for a direct mechanism to be incentive compatible, bidders must not want to misreport in this manner.

Let us proceed by explicitly describing the incentive constraint associated with misreporting a redrawn lower signal. If a bidder with type s_i reports a randomly redrawn lower signal, their surplus is

$$\begin{aligned} & \frac{1}{F(s_i)} \int_{x=\underline{s}}^{s_i} u_i(s_i, x) f(x) dx \\ &= \frac{1}{F(s_i)} \left(\int_{x=\underline{s}}^{s_i} u_i(x) f(x) dx + \int_{x=\underline{s}}^{s_i} (s_i - x) \bar{q}_i(x) g_N(x) dx \right), \end{aligned}$$

where we recall that $\bar{q}_i(v)$ is the probability that the good is allocated to bidder i conditional on the value being v . This formula requires explanation. When a bidder of type s_i misreports a lower signal x , their surplus is higher than what the misreported type receives in equilibrium, since whenever $\max\{x, s_{-i}\} < s_i$, the true value is higher than if bidder i 's

signal had truly been x . The inner integral on the second line sums these differences across all realizations of the highest value of bidders other than i . But because the signal is redrawn from the prior, the expected difference in surplus across all misreports is simply the expected difference of $(\max\{s_i, x\} - x)$, where x is the highest of N draws from the prior F , and when bidder i is allocated the good.

Thus, a necessary condition for a mechanism to be incentive compatible is that, for all i ,

$$u_i(s_i) \geq \frac{1}{F(s_i)} \left(\int_{x=\underline{s}}^{s_i} u_i(x) f(x) dx + \int_{x=\underline{s}}^{s_i} (s_i - x) \bar{q}_i(x) g_N(x) dx \right). \quad (8)$$

Of course, if this constraint holds for each i , then it must hold on average across i , so that

$$u(s) \geq \frac{1}{F(s)} \left(\int_{x=\underline{s}}^s u(x) f(x) + \lambda(s) \right) \quad (9)$$

where

$$u(s) = \sum_{i=1}^N u_i(s)$$

and

$$\lambda(s) = \int_{x=\underline{s}}^s (s - x) \bar{q}(x) g_N(x) dx.$$

If we hold fixed $\bar{q}(v)$, we can derive a lower bound on bidder surplus (and hence an upper bound on revenue) by minimizing ex-ante bidder surplus subject to (9). Our first main result, Theorem 1, asserts that this minimum is attained by the function

$$\underline{u}(s) = \int_{x=\underline{s}}^s \lambda(x) \frac{f(x)}{(F(x))^2} dx + \frac{\lambda(s)}{F(s)},$$

which solves (9) as an equality when $\underline{u}(s) = 0$. In fact, \underline{u} is the pointwise smallest interim utility function that is non-negative and satisfies (9). Indeed, if the constraint held as a strict inequality at s , we could decrease u at that point without violating the constraint, which lowers bidder surplus. But the right-hand side is monotonic in u , so that this modification actually relaxes the constraint even further. As a result, the lower bound is attained by an indirect utility function so that all of the redrawn lower signal constraints are binding.

Thus, if a direct mechanism implements \bar{q} , total bidder surplus must be at least

$$\underline{U} = \int_{s=\underline{s}}^{\bar{s}} \underline{u}(v) f(s) ds = \int_{v=\underline{s}}^{\bar{s}} \int_{x=v}^{\bar{s}} \frac{1 - F(x)}{F(x)} dx \bar{q}(v) g_N(v) dv, \quad (10)$$

and revenue is therefore at most

$$\bar{R} = TS - \underline{U} = \int_{v=\underline{s}}^{\bar{s}} \psi(v) \bar{q}(v) g_N(v) dv, \quad (11)$$

where

$$\psi(v) = v - \int_{x=v}^{\bar{s}} \frac{1 - F(x)}{F(x)} dx$$

is the “marginal revenue” from allocating the good when the value is v . This result is stated formally as follows:

Theorem 1 (Revenue Upper Bound).

In any auction in which the probability of allocation is given by \bar{q} , bidder surplus is bounded below by \underline{U} given by (10) and expected revenue is bounded above by \bar{R} defined by (11).

Proof of Theorem 1. If we define a function operator

$$\Gamma(u)(s) = \frac{1}{F(s)} \left(\int_{x=\underline{s}}^s u(x) f(x) dx + \lambda(s) \right),$$

then the ex-ante constraint (9) is simply that $u \geq \Gamma(u)$. It is easily verified that \underline{u} is a fixed point of Γ . For then

$$\begin{aligned} \Gamma(\underline{u})(s) &= \frac{1}{F(s)} \left(\int_{x=\underline{s}}^s \left(\int_{y=\underline{s}}^x \lambda(y) \frac{f(y)}{(F(y))^2} dy + \frac{\lambda(x)}{F(x)} \right) f(x) dx + \lambda(s) \right) \\ &= \frac{1}{F(s)} \left(F(s) \int_{x=\underline{s}}^s \lambda(x) \frac{f(x)}{(F(x))^2} dx + \lambda(s) \right) \\ &= \underline{u}(s), \end{aligned}$$

where the second line comes from Fubini’s theorem.

We claim that \underline{u} is the lowest non-negative indirect utility function that satisfies this constraint. This follows from the following observations: First, Γ is a monotonic operator on non-negative increasing functions, so by the Knaster-Tarski fixed point theorem, it must have a smallest fixed point. Second, if Γ has another fixed point \hat{u} that is smaller than \underline{u} , then it must be that $\hat{u}(s) \leq \underline{u}(s)$ for all s , with a strict inequality for some positive measure set of s . Moreover, it must be that $\underline{u}(x) - \hat{u}(x)$ goes to zero as x goes to \underline{s} (and hence,

cannot be constant for all x). Thus, if we let $\|\cdot\|$ denote the sup norm, then

$$\begin{aligned} \|\Gamma(\underline{u}) - \Gamma(\widehat{u})\| &= \frac{1}{F(s)} \left| \int_{x=\underline{s}}^s (\underline{u}(x) - \widehat{u}(x)) f(x) dx \right| \\ &\leq \frac{1}{F(s)} \int_{x=\underline{s}}^s |\underline{u}(x) - \widehat{u}(x)| f(x) dx \\ &< \frac{1}{F(s)} \int_{x=\underline{s}}^s \|\underline{u} - \widehat{u}\| f(x) dx \\ &= \|\underline{u} - \widehat{u}\|. \end{aligned}$$

This contradicts the hypothesis that both \underline{u} and \widehat{u} are fixed points of Γ .

Finally, if \widehat{u} is any function that satisfies (9) but is not everywhere above \underline{u} , then consider the sequence $\{u^k\}_{k=0}^{\infty}$ where $u^0 = \widehat{u}$ and $u^k = \Gamma(u^{k-1})$ for $k \geq 1$. Given the base hypothesis that $u^0 \geq \Gamma(u^0) = u^1$ and that Γ is a continuous affine operator, and given that $u \geq 0$ implies that $\Gamma(u) \geq 0$ as well, we conclude that $\{u^k\}_{k=0}^{\infty}$ is monotonically decreasing, and therefore must converge pointwise to a limit that is a fixed point of Γ , which is not uniformly above \underline{u} . This implies that there exists a fixed point that is below \underline{u} , again a contradiction. Thus, \underline{u} must be the lowest fixed point of Γ . \square

This result, and the formula (11), are strongly reminiscent of the familiar revenue equivalence result of Myerson (1981), which says that revenue from any auction is the expected *virtual value* of the bidder who is allocated the good. This virtual value is the bidder's true value minus the cost of providing incentives for bidders to report truthfully. To derive that virtual value, one uses the aforementioned envelope formula (5) to compute the total rents that bidders get from a given allocation. Essentially, if the seller allocates more often to a type s_i when the profile is s , then the interim surplus of all higher types must go up at the same rate, so that the ex-ante cost of providing incentives is proportional to the mass of types higher than s_i , specifically $(1 - F_i(s_i)) f_{-i}(s_{-i})$, which can be rewritten as the familiar inverse hazard rate of s_i times the likelihood of the type profile s .

In the present setting, $\psi(v)$ represents an analogous virtual value from allocating the good conditional on the value being v : total surplus increases at a rate of v , bidder surpluses need to increase to maintain the *global* incentive constraints (9). Indeed, the rate of increase in ex-ante bidder surplus per unit increase in $\bar{q}(v)$ is precisely

$$g_N(v) \int_{x=v}^{\bar{s}} \frac{1 - F(x)}{F(x)} dx.$$

To see this, observe that increasing $\bar{q}(v)$ has two effects. There is a direct increase in $u(s)$ for all $s \geq v$ at a rate of $(v - s) g_N(v) / F(s)$. But there is also an indirect effect, in that the direct increase in $u(s)$ is passed on to all types $s' > s$ at a rate of $f(s) / F(s')$, i.e., the likelihood that the higher type s' misreports s under the given global deviation. Hence, if we let $\rho(s|v)$ denote the total rate of change in $u(s)$, then for all $s \geq v$, ρ must satisfy the integral equation

$$\rho(s|v) = \frac{1}{F(s)} \left(\int_{x=v}^s \rho(x|v) f(x) dx + (s - v) g_N(v) \right),$$

and $\rho(s|v) = 0$ for $s < v$. By integrating this equation with a suitable integrating factor, which is $f(s) / F(s)$, one can derive that for $s \geq v$,

$$\int_{x=s}^s \rho(x|v) f(x) dx = g_N(v) \int_{x=v}^s \frac{F(s) - F(x)}{F(x)} dx,$$

which gives the desired result when $s = \bar{s}$.

4 Posted Prices

It turns out that the upper bound on revenue constructed in Section 3 can be attained and is exactly the optimal revenue. We will give a full analysis of optimal mechanisms in Section 5. In many cases, however, maximum revenue can be achieved by simply setting a posted price at which all bidders (and all types) would want to purchase the good. This section characterizes exactly when such simple mechanisms are optimal.

4.1 The Optimality of Posted Prices

Our first result for this section shows that if we restrict attention to mechanisms in which the good is always allocated, then maximum revenue is achieved with posted prices. As we are in model of pure common values, the allocation is therefore guaranteed to be socially efficient.

Theorem 2 (Revenue Optimality of Posted Prices among Efficient Mechanisms).

Among all auctions that allocate the good with probability one, revenue is maximized by setting a posted price of

$$p = \int_{v=\underline{s}}^{\bar{s}} v g_{N-1}(v) dv, \tag{12}$$

i.e., the expected value of the good conditional on having a signal of zero. This posted price is inclusive, in the sense that all types are willing to purchase at this price. This mechanism results in all bidders being equally likely to receive the good, so that $\bar{q}_i(v) = 1/N$ for all i and v .

Proof of Theorem 2. It suffices to show that p is equal to the resulting upper bound on revenue, which follows from a straightforward calculation:

$$\begin{aligned} \int_{v=\underline{s}}^{\bar{s}} \psi(v) g_N(v) dv &= \psi(\bar{s}) - \int_{v=\underline{s}}^{\bar{s}} \left(1 + \frac{1-F(v)}{F(v)}\right) G_N(v) dv \\ &= \bar{s} - \int_{v=\underline{s}}^{\bar{s}} G_{N-1}(v) dv \\ &= \int_{v=\underline{s}}^{\bar{s}} v g_{N-1}(v) dv, \end{aligned}$$

where we have integrated by parts to obtain the first and the last lines. □

In fact, the revenue upper bound gives us a straightforward condition under which it is optimal to always allocate the good. The virtual value function $\psi(v)$ is strictly increasing in v , so it is everywhere non-negative if and only if $\psi(\underline{s}) \geq 0$. In that case, the revenue upper bound is maximized by setting $\bar{q}(v) \equiv 1$. As we will see in Section 5, in cases where $\psi(\underline{s}) < 0$, it is possible to strictly improve on the posted price, so that this condition is both necessary and sufficient.

Corollary 1 (Revenue Optimality of Posted Prices).

Revenue is maximized by the setting the inclusive posted price of (12) if and only if $\psi(\underline{s}) \geq 0$.

Whether or not $\psi(\underline{s}) \geq 0$ depends on the particular distribution. A necessary condition is that $\underline{s} > 0$, meaning that there is a gap between the seller's value for the good and the smallest value among the bidders. Intuitively, the larger is this gap, the easier it is to satisfy the conditions for posted prices to be optimal. We shall present an example of this shortly. However, it is possible for the information rent to become arbitrarily large when the value is close to \underline{s} , so that even if values shifted upwards by a large amount, posted prices would not be optimal.

Beyond the positive virtual value case, posted prices would be optimal if the seller were constrained to sell the good for reasons beyond our model. For example, it could be that both the seller and the intermediaries are long lived, and while the seller can commit to a mechanism for a short period of time, he cannot commit to not sell the good in the future if it is unsold. A Coasean logic suggests that the competition between the seller and his future

N	$\underline{s}(N)$
2	1.7
3	0.17
4	0.027
5	0.004
10	0.0001

Table 1: Lower bound $\underline{s}(N)$ on signal for which posted prices are optimal.

self would force immediate trade, so that the seller can do no better than run the optimal mechanism that allocates with probability one. This argument has been made formally in the context of private value auctions by Vartiainen (2013), and we suspect it can be extended to this setting, although it is beyond the scope of this paper.

4.2 A Uniform Example

We will illustrate these results with a simple uniform example. Suppose that the value is uniformly distributed on $[\underline{s}, \underline{s} + 1]$. Then it follows that the distribution of the individual signal is given by:

$$F(x) = \begin{cases} 0 & \text{if } x < \underline{s}; \\ (x - \underline{s})^{\frac{1}{N}} & \text{if } \underline{s} \leq x < \underline{s} + 1; \\ 1 & \text{otherwise.} \end{cases}$$

We can then compute the generalized virtual value:

$$\begin{aligned} \psi(v) &= v - \int_{x=v}^{\underline{s}+1} \frac{1 - (x + \underline{s})^{\frac{1}{N}}}{(x - \underline{s})^{\frac{1}{N}}} dx \\ &= \underline{s} + 1 - \frac{N}{N-1} \left[(\underline{s} + 1)^{\frac{N-1}{N}} - v^{\frac{N-1}{N}} \right]. \end{aligned}$$

Thus, we know that posted prices are optimal if

$$\underline{s} + 1 - \frac{N}{N-1} \left[(\underline{s} + 1)^{\frac{N-1}{N}} - \underline{s}^{\frac{N-1}{N}} \right] \geq 0.$$

This equation is satisfied as long as \underline{s} is above a critical value $\underline{s}(N)$, which is decreasing in N . A few sample values are given in Table 1. The optimal inclusive posted price is

$$\begin{aligned}
p &= \int_{v=\underline{s}}^{\underline{s}+1} v g_{N-1}(v) dv \\
&= \frac{N-1}{2N-1} \left[(\underline{s}+1)^{\frac{2N-1}{N}} - \underline{s}^{\frac{2N-1}{N}} \right].
\end{aligned}$$

This is also the optimal revenue when $\underline{s} \geq \underline{s}(N)$, and for all \underline{s} when the good must be sold.

4.3 Other Optimal Mechanisms

While Theorem 2 shows that the posted price is an optimal efficient mechanism, we may very well wonder if there are other optimal mechanisms. While there are some degrees of freedom in designing the optimal auction, any optimal mechanism must share a number of properties with the posted price. First, the posted price has the property that all types are equally likely to be allocated the good and all types make the same transfer. A similar property holds for all optimal mechanisms.

Proposition 1 (Type Independent Transfers and Allocations).

In any incentive compatible and individually rational mechanism such that (9) holds as an equality for all s_i , then $q_i(s_i)$ and $t_i(s_i)$ are independent of type, i.e., there exist constants q_i and t_i such that $q_i(s_i) \equiv q_i$ and $t_i(s_i) \equiv t_i$.

Proof of Proposition 1. Since (9) binds, it must be that (8) binds as well, and each bidder i is indifferent to downward deviations. Moreover, since each type s_i is indifferent to reporting a randomly redrawn lower signal, it must be that s_i is also indifferent to reporting any type $s'_i \leq s_i$. For if there were a positive measure of types for which $u_i(s_i) > u_i(s_i, s'_i)$, and if there is indifference on average, then there must be some other type s'_i such that $u_i(s_i, s'_i) > u_i(s_i)$.

Now consider the highest type \bar{s} who knows that the value is \bar{s} . Then for all s_i and s'_i , $u_i(\bar{s}, s_i) = u_i(\bar{s}, s'_i)$ implies that

$$t_i(s_i) - t_i(s'_i) = \bar{s}(q_i(s_i) - q_i(s'_i)).$$

Notice that if this difference is strictly positive, then since the value conditional on a signal of s'_i is strictly less than \bar{s} with probability 1, we have that

$$u(s'_i) - u(s'_i, s_i) < \bar{s}(q_i(s'_i) - q_i(s_i)) - (t_i(s'_i) - t_i(s_i)) = 0,$$

which contradicts the indifference of type s'_i to reporting s_i . □

For example, it would also be optimal to implement a biased posted price that breaks ties non-uniformly.

In principle, this leaves open the possibility that the allocation of the good depends in a complicated manner on the realized type profile. However, the nature of such correlation is considerably restricted by the revenue equivalence formula (7). In any optimal efficient mechanism, we know that interim bidder surplus must be at its lower bound, which is

$$\underline{u}_i(s_i) = \int_{x=s}^{s_i} \frac{1}{F(x)} \int_{y=s}^x \bar{q}_i(y) g_N(y) dy dx.$$

Hence, in any incentive compatible mechanism, we must have

$$\widehat{q}_i(s_i) = \frac{1}{F(s_i)} \int_{x=s}^{s_i} \bar{q}_i(x) g_N(x) dx.$$

In a symmetric mechanism, we must have $\bar{q}_i(x) \equiv 1/N$, so that

$$\widehat{q}_i(s_i) = \frac{1}{NF(s_i)} G_N(s_i) = \frac{1}{N} G_{N-1}(s_i).$$

In other words, bidder i is no more or less likely to be allocated the good whether he has the highest signal or not.

Thus, while the optimal mechanism need not be a posted price, it must share some crucial features with the inclusive posted price: the interim probability of getting the good and the interim transfer must be independent of type, and the interim probability of getting the good must be the same whether or not one has the highest signal.

5 Optimal Mechanisms in the General Case

Essentially all of the previous analysis can be generalized to construct a revenue maximizing mechanism for any distribution of values. The characterization consists of two steps. First, deriving an upper bound on revenue using the bound from Theorem 1. Second, constructing an indirect mechanism that achieves that upper bound.

For the first step, it is straightforward to see that the \bar{q} that maximizes the revenue upper bound (11) is bang-bang: $\bar{q}(v)$ is one whenever $\psi(v) > 0$ and is zero when $\psi(v) < 0$. Since $\psi(v)$ is strictly increasing, there is some smallest r such that $\psi(r) \geq 0$, with $\psi(v) > 0$ for all $v > r$, so that the optimal allocation is

$$\bar{q}(v) = \mathbb{I}_{v \geq r}. \tag{13}$$

In other words, the seller keeps the good when the value is less than r , and wants to sell the good with probability one when the value is above r .

We will construct a mechanism, which we will refer to as *guaranteed demand auction* (GDA), that implements this allocation and holds bidders down to their lower bound surplus. In any such mechanism, all of the downward global constraints must bind. As we observed in Proposition 1, this implies that interim transfers and interim allocation probabilities are independent of type. Coming from the tradition of auction design with private values, this seems almost paradoxical, since monotonically increasing allocations and transfers are the standard tools for screening in that environment. In the present interdependent value setting, we can, however, screen without such devices. The bidders care about not just the likelihood of receiving the good but also about how their allocation is correlated with the types of others. Moreover, different types have different preferences over such correlation. For example, the highest type is sure of the value, and only cares about the interim allocation probability, but not about how the allocation is correlated with others' types. A type $s_i < \bar{s}$, however, is unsure of the value, and prefers to receive the good more often when others' signals are higher than s_i , since this is when the value is highest. Thus, a menu of allocation rules with different kinds of correlation could induce bidders to separate.

5.1 The Guaranteed Demand Auction

Let us demonstrate how this can be done. Consider the following *guaranteed demand game* for determining the allocation of the good. Each bidder makes a demand $d_i \in [0, \bar{d}]$, where \bar{d} is a parameter of the game and is a number between 0 and $1/N$. The allocation is determined as follows. Let i^* denote the identity of the bidder with the highest demand, chosen randomly if there are multiple high demanders. If $d_{i^*} > 0$, then bidder i^* is allocated the good with probability d_{i^*} and each bidder $j \neq i^*$ receives the good with probability $(1 - d_{i^*}) / (N - 1)$. If $d_{i^*} = 0$, then the seller keeps the good. Thus, conditional on the highest demand being positive, a bidder is more likely to be allocated the good when they do not have the high demand, so that a bidder's probability of receiving the good is always at least their demand.

We claim that there is an equilibrium in which each bidder follows the monotonic pure strategy

$$\sigma(s_i) = \begin{cases} \frac{1}{N} \left(1 - \frac{G_N(\hat{r})}{G_N(s_i)}\right) & \text{if } s_i \geq \hat{r}; \\ 0 & \text{if } s_i < \hat{r}, \end{cases} \quad (14)$$

where \hat{r} solves

$$G_N(\hat{r}) = 1 - N\bar{d}. \quad (15)$$

This cutoff \hat{r} is chosen so that

$$\sigma(\bar{s}) = (1 - (1 - N\bar{d})) / N = \bar{d},$$

i.e., the highest type makes the highest possible demand.

Let us verify that these strategies are indeed an equilibrium. If the other bidders follow this strategy, then expected payoff to a type s_i that demands $\sigma(s'_i)$ for $s'_i \geq \hat{r}$ is

$$\tilde{u}_i(s_i, s'_i) = \sigma(s'_i) \int_{x=\hat{r}}^{s'_i} \max\{s_i, x\} g_{N-1}(x) dx + \int_{x=s'_i}^{\bar{s}} \max\{s_i, x\} \frac{1 - \sigma(x)}{N-1} g_{N-1}(x) dx.$$

The derivative with respect to s'_i is

$$\begin{aligned} \frac{\partial \tilde{u}_i(s_i, s'_i)}{\partial s'_i} &= \sigma'(s'_i) \int_{x=\hat{r}}^{s'_i} \max\{s_i, x\} g_{N-1}(x) dx + \max\{s_i, s'_i\} \left(\frac{N}{N-1} \sigma(s'_i) - \frac{1}{N-1} \right) g_{N-1}(s'_i) \\ &= \frac{1}{N} \frac{G_N(\hat{r})}{(G_N(s'_i))^2} g_N(s'_i) \int_{x=\hat{r}}^{s'_i} \max\{s_i, x\} g_{N-1}(x) dx - \max\{s_i, s'_i\} \frac{G_N(\hat{r})}{G_N(s'_i)} \frac{1}{N-1} g_{N-1}(s'_i) \\ &= \frac{G_N(\hat{r})}{G_N(s'_i)^2} f(s'_i) \frac{1}{G_{N-1}(s'_i)} \int_{x=\hat{r}}^{s'_i} (\max\{s_i, x\} - \max\{s_i, s'_i\}) g_{N-1}(x) dx. \end{aligned}$$

When $s'_i > s_i$, the integrand is negative (since the variable of integration is always less than s'_i) so that the derivative is negative. And when $s'_i < s_i$, the derivative is zero! Thus, no type has an incentive to deviate, and the proposed strategies are an equilibrium.

Now, we can turn the guaranteed demand game into a *guaranteed demand auction* (GDA) by adding entry fees f_i . In other words, each bidder's message consists of a pair of an entry decision and a demand. If bidder i decides to enter, they pay f_i to the seller, and the good is allocated among the bidders who enter according to the rules of the guaranteed demand game. As long as $f_i \leq \tilde{u}_i(\underline{s})$ for all i , all bidders will be willing to enter the auction, and play the previously described equilibrium of the guaranteed demand game. We summarize this discussion with the following proposition:

Proposition 2 (Equilibrium of the Guaranteed Demand Auction).

As long as $f_i \leq \tilde{u}_i(\underline{s})$ for all i , it is an equilibrium for all bidders to enter the GDA and make demands according to (14). In equilibrium, bidders are indifferent between their equilibrium demands and all lower demands.

5.2 Optimality of the Guaranteed Demand Auction

Given the characterization of Proposition 2, it is fairly easy to see that for appropriately chosen parameters, the GDA maximizes expected revenue. To be precise, consider the GDA where we set

$$\bar{d} = \frac{1 - G_N(r)}{N},$$

where r is the screening level that maximizes the upper bound on revenue, and set the entry fee to be $f_i = \tilde{u}_i(\underline{s})$ for all i . Then by Proposition 2, it is an equilibrium for all bidders to enter and bid according to (14). Our choice of \bar{d} ensures that some bidder makes a positive demand if and only if some bidder has a signal greater than r , which occurs precisely when the value is greater than r . Thus, the allocation implemented by this auction and equilibrium is the optimal \bar{q} given by (13). Our choice of entry fee also ensures that the lowest type gets surplus of zero from participating in the auction.

Moreover, under these strategies, bidders are indifferent between their equilibrium demands and all lower demands. Thus, they are indifferent to misreporting randomly redrawn lower signals, so that the interim utility function is precisely the \underline{u} described in Section 3. Bidder surplus must therefore be minimized subject to implementing the optimal \bar{q} , so that revenue achieves its upper bound. This completes the proof of the following result:

Theorem 3 (Optimality of the Guaranteed Demand Auction).

Revenue is maximized by running a guaranteed demand auction with maximum demand $\bar{d} = (G_N(r) - 1) / N$ and a symmetric entry fee which is equal to $\tilde{u}(\underline{s})$.

Let us revisit the uniform example. When $\underline{s} < s(N)$, the optimal auction is not a posted price, but rather withholds the good when the value is sufficiently low. For example, when $\underline{s} = 0$, the virtual value reduces to

$$\psi(v) = v - \int_{x=v}^1 \frac{1 - x^{1/N}}{x^{1/N}} dx = \frac{1}{N-1} \left(Nv^{\frac{N-1}{N}} - 1 \right)$$

The optimal exclusion level solves $\psi(r) = 0$, which is

$$r(N) = \frac{1}{N^{\frac{N-1}{N}}}.$$

Optimal revenue in the can-keep case can be computed using $F(x) = x^{1/N}$ and the optimal cut-off above as:

$$\begin{aligned} R(N) &= \frac{1}{N-1} \int_{x=r}^1 \left(Nx^{\frac{N-1}{N}} - 1 \right) dx \\ &= \frac{N}{N-1} \frac{N}{2N-1} \left(1 - \frac{1}{N^{\frac{2N-1}{N-1}}} \right) - \frac{1}{N-1} \left(1 - \frac{1}{N^{\frac{N}{N-1}}} \right). \end{aligned}$$

In contrast, the optimal inclusive posted price is

$$\int_{v=0}^1 v g_{N-1}(v) dv = \frac{N-1}{2N-1}.$$

As $N \rightarrow \infty$, both expressions converge to the efficient surplus of $1/2$.

5.3 Alternative Implementations

As a final topic in this section, we discuss two alternative implementations to the guaranteed demand auction. First, when $r = \underline{s}$, we have already proven in Theorem 2 that an inclusive posted price maximizes revenue. Indeed, when this happens, the strategies described in (14) reduce to all bidders demanding $1/N$ and being equally likely to receive the good. This is the same outcome as would be achieved by simply posting a price equal to the optimal entry fee, which is $\tilde{u}(\underline{s})$.

More generally, there is an interesting alternative implementation of the GDA using a “descending clock,” in a manner that is in a sense dual to a Dutch auction. In the Dutch auction, the value of the clock represents the *price* at which the bidder who stops the clock will purchase the good. Here, the value of the clock represents the *probability* with which the bidder who stops the clock gets allocated the good. As before, the bidders must pay a fee to enter the auction, after which they play a descending clock game. Let d denote the value on the clock, which starts at \bar{d} and gradually descends to zero. The game ends when some bidder stops the clock. If bidder i stops the clock at $d > 0$, then that bidder receives the good with probability d , and each bidder $j \neq i$ receives the good with probability $(1-d)/(N-1)$. If the clock hits zero, the game ends and the seller keeps the good.

It is not hard to see that there is an equilibrium of this descending clock auction in which each bidder uses the cutoff strategy stopping the clock as soon as $d \leq \sigma(s_i)$. For if a bidder waits until time $d = \sigma(\hat{s})$ before stopping the clock, then surplus from stopping at time

$\sigma(s'_i)$ for $s'_i \leq s_i$, net of the entry fee, is just

$$\sigma(s'_i) \int_{x=\hat{r}}^{s'_i} \max\{s_i, x\} g_{N-1}(x) dx + \int_{x=s'_i}^{\hat{s}} \max\{s_i, x\} \frac{1 - \sigma(x)}{N-1} g_{N-1}(x) dx.$$

This is nearly the same expression as we had for the interim utility $\tilde{u}_i(s_i, s'_i)$ in Section 5.1, except that the second term excludes very high types for the other bidders. The same steps as we followed before would show that this expression is decreasing in s'_i for $s'_i > s_i$ and is constant for $s'_i \leq s_i$. So, an optimal strategy is to wait and stop the clock at $d = \sigma(s_i)$, or stop immediately if $d < \sigma(s_i)$.

Interestingly, even as d gets arbitrarily close to zero, bidders are still willing to wait and see if someone else stops the auction, with the reason being that the probability of being allocated the good as the bidder who stops the auction is always small compared to the corresponding probability when someone else stops the auction. Thus, bidders with low signals are willing to wait and hope that someone else stops the auction before it is too late.

An advantage to the descending implementation is that it elicits only as much information as is needed to determine the allocation. In particular, the outcome of the auction only reveals the highest of the bidders signals, i.e., the value, and the identity of the bidder who had the high signal. In contrast, the guaranteed demand auction would elicit the signals of all of the bidders.

Finally, let us return to the question of uniqueness of the optimal auction. Clearly, any mechanism that achieves the bound must allocate the good if and only if the value is at least r , so that the optimal \bar{q} is pinned down. And in order to achieve the lower bound on bidder surplus, all of the incentive constraints for misreporting a redrawn lower signal must bind, so that Proposition 1 implies that the expected transfer and expected allocation probabilities are independent of type. Also, the envelope formula of Lemma 1 again implies that $\hat{q}_i(v)$ is pinned down from \bar{q}_i , and indeed, in any symmetric optimal mechanism, we must have that $\hat{q}_i(v) = \sigma(v)$, the equilibrium demand under the GDA, since when bidder i has the highest type $s_i = v$, she will receive the good with probability equal to her demand. Thus, while there may be other optimal auctions, all must share several key features in common with the GDA.

6 Discussion

6.1 Strategic Equivalence with Independent Private Value Environment

We now turn our attention to several extensions and applications of our model. First, there is a remarkable connection between incentive compatible allocations in the maximum common value model and those in the independent private value model, which has implications for the theory of robust mechanism design. We can associate each maximum common value model with an independent private value (IPV) model that has the same distribution of signals. Under the former, the value is common and equal to the maximum of the signals, and in the latter, each bidder's signal is equal to his private value. In principle, equilibrium behavior could be quite different under the two models. However, for mechanisms like first- and second-price auctions that discriminate in favor of more optimistic bidders, bidders will behave exactly the same under the two models. In other words, bidders will behave in the common value model as if their signals are their true values.

This result generalizes an observation of Bulow and Klemperer (2002) that bidding one's signal is an equilibrium of the second-price auction in the maximum common value model, as it would be with independent private values. In this equilibrium, the bidder with the highest signal wins the auction and pays the second-highest signal. To see this, suppose that bidders $j \neq i$ are bidding their signals. The surplus to a bidder with signal s_i from bidding $b < s_i$ (ignoring ties) is

$$\int_{s_{-i} \in X([\underline{s}, b])} \left(s_i - \max_{j \neq i} s_j \right) f_{-i}(s_{-i}) ds_{-i},$$

which is clearly increasing in b . On the other hand, if $b > s_i$, then the surplus for the bidder i is

$$\int_{s_{-i} \in X([\underline{s}, s_i])} \left(s_i - \max_{j \neq i} s_j \right) f_{-i}(s_{-i}) ds_{-i} + \int_{s_{-i} \in X([s_i, b])} \left(\max_{j \neq i} s_j - \max_{j \neq i} s_j \right) f_{-i}(s_{-i}) ds_{-i}$$

which is equal to the surplus from bidding s_i ! It is therefore optimal to bid any amount which is at least your signal, and, in particular, it is optimal to bid your signal. Thus, we find that in the second-price auction of this pure common value environment, each bidder behaves as if his signal is his true private value rather than a signal and in particular a lower bound on the pure common value.

We generalize this observation in the following manner. Consider the alternative model in which each bidder's signal is again drawn from F , but instead of the value being the highest of the signals, the value is the bidder's own signal. In other words, this is the independent private value model, where bidder i 's value is

$$v_i(s_1, \dots, s_N) = s_i.$$

Let $H(s) = \{i | s_i = \max_j s_j\}$ denote the set of bidders with high signals. We will say that the direct mechanism $\{q_i, t_i\}$ is *conditionally efficient* if (i) $q_i(s) > 0$ if and only if $s_i \in H(s)$ and (ii) there exists a cutoff r such that the good is allocated whenever $\max_i s_i > r$.

Proposition 3 (Strategic Equivalence).

Suppose a direct mechanism $\{q_i, t_i\}$ is incentive compatible and individually rational for the independent private value model in which $v_i(s) = s_i$ and that the allocation is conditionally efficient. Then $\{q_i, t_i\}$ is also incentive compatible and individually rational for the maximum common value model in which $v_i(s) = \max_j \{s_j\}$.

As a corollary, consider any mechanism that admits an equilibrium in which the allocation is conditionally efficient in the independent private value model, and such that all actions are used in equilibrium. We will refer to such a mechanism as *standard*. Then the same strategies will also be an equilibrium of that mechanism in the corresponding maximum common value model with the same distribution of signals but in which all bidders have a common value equal to the maximum signal. Moreover, if two such mechanisms are revenue equivalent in the IPV setting, they will also be revenue equivalent when the value is the maximum signal.

This strategic equivalence is surprising. When we transform the independent private value model into the common value model, the bidders' interim expectations of their values increase substantially, since their value is now the maximum of their own signal and the highest signal of others. In principle, one might think that the higher values would induce the bidders to bid more aggressively. However, this turns out not to be the case. The reason is that others' bidding strategies are correlated with the value in such a way that all of the potential surplus gains from more aggressive bidding are *exactly* dissipated by higher sales prices. In fact, while bidders would strictly prefer their equilibrium bids over higher bids in the IPV model, they become *indifferent* between their equilibrium bids and all higher bids when the value is the maximum signal. In other words, the winner's curse is exactly strong enough to make any bid beyond the realized signal just as good as bidding the realized signal.

6.2 Application to Robust Auction Design

We can combine Proposition 3 with results from our prior work on informationally robust predictions in first-price auctions (Bergemann, Brooks, and Morris, 2017) to understand the robust revenue properties of the first-price auction. In the earlier paper, we analyzed the range of possible revenue outcomes of the first-price auction, where we held fixed a general symmetric distribution of bidders' values but varied the bidders' information and equilibrium strategies. An implication of the main result in (Bergemann, Brooks, and Morris, 2017) to the setting where bidders' values are common is that the information structure that minimizes revenue is precisely the one in which the bidders have independent signals, and the value is the maximum signal. Thus we show that the model of the maximum of independent signals attains the minimum revenue for a first-price auction, across all type spaces with a fixed marginal distribution over a pure common value.

In combination with the above strategic equivalence result, Proposition 3, we conclude that first-price auctions generate greater *worst-case* revenue across all type spaces and equilibria than does any mechanism that implements (conditionally) efficient allocations in the corresponding IPV setting. To wit, by Proposition 3, we can extend the revenue equivalence result from the independent private value model to the maximum common value model for standard auctions. By Theorem 1 of Bergemann, Brooks, and Morris (2017), the lowest revenue in a first-price auction is achieved in the maximum type space. By Proposition 3, this revenue is also achieved in all other standard auctions, but in any such auction, the worst case type space could possibly be different, and hence yield even lower revenues.

Corollary 2 (Robust Performance of First-Price Auction).

Suppose there is a pure common value v with fixed distribution $G(v)$. Then the first-price auction generates greater minimum revenue than any other standard mechanism, where the minimum is taken across all Bayes Nash equilibria and across all common value common prior type spaces where the distribution of the common value is G .

Note that while our proposition only asserts that first-price auctions are weakly better than other standard mechanisms in the maxmin sense, in some cases we know that the ordering is strict. For example, second-price auctions admit “bidding ring” equilibria in which one bidder bids a large amount while the others bid zero. This revenue ranking contrasts sharply with the ranking of mechanism suggested by “linkage principle” of Milgrom and Weber (1982), who find that when values are affiliated, English auctions generate more revenue than second-price auctions, which in turn generate greater revenue than first-price auctions. (Note that the maximum common value model is an affiliated values model.) The different conclusion here stems from the fact that our worst-case criterion varies the

information structure as we vary the mechanism, whereas Milgrom and Weber’s comparison holds the information structure constant while comparing mechanisms under a particular equilibrium selection.

One can therefore interpret Corollary 2 as an explanation as to why first-price auctions are so much more prevalent than second-price auctions, or any other auction format. The reason is that while all of these mechanisms generate the same revenue in the IPV environment, standard mechanisms other than the first-price auction are more susceptible to low revenue in other informational environments and when equilibrium selection is less favorable to the seller.

6.3 Auctions versus Optimal Mechanisms

We derived the optimal mechanism in an environment with independent signals and common values. In a seminal paper, Bulow and Klemperer (1996), established the limited power of optimal mechanisms as opposed to standard auction formats. They showed that the revenue of the optimal auction with N bidders is strictly less than that of a standard auction without a reserve price with $N + 1$ bidders. The pure common value environment analyzed here is an instance of their more general interdependent value environment with one exception. The virtual utility function—or marginal revenue function in the language of Bulow and Klemperer (1996)—is not monotone due the maximum operator in the common value model. We saw that this aspect of the environment lead to an optimal mechanism with features distinct from the standard first- or second-price auction. Namely, the optimal mechanism elicits the information from the bidder with the highest signal but minimizes the probability of assigning him the object subject to incentive constraints. This raises the question whether the revenue comparison suggested by Bulow and Klemperer (1996) still resolves in favor of the standard auction.

In the pure common value environment considered here, the value of the object is the same for all bidders. However, only the bidder with the highest signal can guarantee himself an information rent. Indeed, the *virtual utility* of each bidder, $\pi_i(s_i, s_{-i})$ is constant in s_i and equal to the utility until s_i becomes the largest signal. At this critical point, the virtual utility of bidder i displays a downward jump, and thereafter has the standard expression of the virtual utility:

$$\pi_i(s_i, s_{-i}) = \begin{cases} \max_j \{s_j\}, & \text{if } s_i \leq \max\{s_{-i}\}; \\ \max\{s_j\} - \frac{1-F_i(s_i)}{f_i(s_i)}, & \text{if } s_i > \max\{s_{-i}\}. \end{cases} \quad (16)$$

The downward discontinuity in the virtual utility indicates why the seller wishes to minimize the probability of assigning the object to the bidder with the high signal. We notice that the downward discontinuity is due to the value function of the bidders, and arises independent of the nature of the distribution function. The virtual utility of bidder i therefore fails the monotonicity assumption even when the hazard rate of the distribution function is increasing everywhere. Bulow and Klemperer (1996) required the monotonicity of the virtual utility when establishing their main result that an absolute English auction with $N + 1$ bidders is more profitable than any optimal mechanism with N bidders.

Indeed, we can show that the revenue ranking established in Bulow and Klemperer (1996) does not extend to the current auction environment in a surprisingly strong sense. We compare the revenue from the optimal auction with N bidders to an absolute second-price, or equivalently, absolute second-price auction with $N + K$ bidders. The term absolute refers to the fact that there is no reserve price imposed, thus the object is sold with probability one in the corresponding absolute auction.

Proposition 4 (Revenue Comparison).

For every $N \geq 1$ and every $K \geq 1$, the revenue from an absolute second-price auction with $N + K$ bidders is strictly dominated by the revenue of an optimal auction with N bidders.

In view of the characterization of the optimal auction in Theorem 1, the proof of the above result then follows from a simple comparison of the *second order statistic* of $N + K$ independent and identically distributed signals and the *first order statistic* of $N + K - 1$ independent and identically distributed signals. An elementary calculation, or equivalently a moment of reflection, establishes that the latter exceeds the former. Why then is this enough to establish the result? The second order statistic of $N + K$ signals simply represents the revenue of the absolute second-price auction with $N + K$ bidders. By Theorem 1, the optimal mechanism (weakly) exceeds result the revenue from a posted price set equal to the maximum of $N + K - 1$ signals. Now, if instead of $N + K$ bidders, the optimal auction only has N bidders, then it is as if only N independent and identical distributed signals are revealed to the N bidders. But as the pure common value of the object is not affected by the number of bidders, it is as if the remaining K signals are simply not disclosed, but the N participating bidders still form the expectation over the $N + K$ signals. Thus an attainable revenue for the seller is to offer the object at random to a bidder at a posted price set equal to the maximum of $N + K - 1$ signals. The optimal auction will typically be able to do even better. But as this lower bound on the revenue is already enough to dominate the second-price auction, the argument is complete.

6.4 Auctions with Resale

We conclude by revisiting the resale interpretation of our model. We observed in the Introduction that a leading interpretation is that the bidders have independent private values, but that the auction will be followed by a resale market. In particular, the values will exogenously become complete information, and the winner of the good can make a take-it-or-leave-it offer to one of the other bidders. Such a model of resale has been used by Gupta and LeBrun (1999) and Haile (2003) to study asymmetric first-price auctions, whereas we are interested in the optimal auction in such a setting.

The recent work of Carroll and Segal (2016) also studies optimal auction design in the presence of resale. In their model of resale, the least favorable situation for the seller arises when the values become complete information and the high-value bidder gets to make a take-it-or-leave-it offer to whoever wins the good in the primary auction. By contrast, in the resale interpretation of our model, it is the winner in the primary auction that makes the subsequent offer. In both models, non-local constraints are also very important to the analysis. However, in their model they arise only when the distribution of values is asymmetric, whereas they arise in our setting already with symmetric distribution of values. The source of the information rents and the corresponding virtual utilities for the bidders are also very different in the two models. In their model, every bidder—except the high-value bidder—gets his private-value information rent. If he wins the good, then he gets to sell it to the high-value bidder, but at a price equal to his value. The high-value bidder gets an extra information rent regardless of whether or not he wins the good at auction, since he can obtain positive surplus from purchasing the good in the resale market, and that surplus is again linearly increasing in his private value. In contrast, in our model, only the high-value bidder has locally relevant private information about the value of the good being allocated, so it is only this bidder that gets an information rent. With this, in our setting, the optimal allocation is distorted away from the high signal bidder “as much as possible” subject to the non-local constraints. In their work, to avoid giving double rents, the good is allocated to the high-signal bidder “as much as possible” subject to the non-local constraints. This distinction in terms of the information rent leads to different deviations that matter. In our problem they are of the form: a type s_i misreports a type drawn on $[\underline{s}, s_i]$ with weights proportional to the prior. In their problem, the deviations that matter are of the form: a type $s_i \in [\hat{s}, \tilde{s}]$ misreports a type drawn from $[\underline{s}, \hat{s}]$, and the probability weights are proportional to the prior. Thus, while the constraints differ across these two models, they fall within the larger class of interim constraints that correspond to deviating from a true value s_i to a stochastic misreport that is drawn from some lower support $[\underline{s}, \hat{s}]$ where $\hat{s} \leq s_i$.

The resale interpretation has some limitations. Truth-telling is an equilibrium of this mechanism under the assumption that values automatically become complete information in the secondary market, so that the resale price is exactly the highest value. Truth-telling would no longer be incentive compatible if bidders had to infer one another's values from the outcome of the auction. To see this, recall that bidders are indifferent between reporting their true type and reporting any lower type. Intuitively, when a bidder has the highest signal, they only receive rents from being allocated the good outright since otherwise they have to buy it at its value in the secondary market. The optimal mechanism makes it so that the probability of being allocated the good is independent of the report $s'_i < s_i$, conditional on s_i being the highest signal, so that bidders do not benefit from deviating down. It is essential for this logic that the bidder not make any rents in the resale market, even after they report a lower signal. On the other hand, if by reporting a lower signal a bidder could signal a lower willingness to pay, then the deviator could buy the good in the resale market at a price strictly less than its value, so that the downward deviator would be strictly better off. It remains an open question what would be the form of the revenue maximizing mechanism if resale prices were influenced by the auction format.

7 Conclusion

This paper contributes to the theory of revenue maximizing auctions when the bidders have a common value for the good being sold. In the classic treatment of revenue maximization, due to Myerson (1981), the potential buyers of the good have independent signals about the value. While the standard model does encompass some common value environments, the leading application is to the case of *independent private values*, wherein each bidder observes their own value, which is distributed independently of the others' values. In benchmark settings, the optimal auction is simply a first- or second-price auction with a reserve price. More broadly, the optimal auction induces an allocation that discriminates in favor of more optimistic bidders, i.e., bidders whose expectation of the value is higher. By contrast, the class of common value models we have studied have the qualitative feature that values are more sensitive to the private information of bidders with more optimistic beliefs. This seems like a natural feature of many economic environments, in which the most optimistic bidder has the most useful information for determining the best-use value of the good, and therefore has a greater information rent. One class of models for which this is the case is when the bidders are intermediaries who wish to sell the good in some secondary market. In contrast, the characterizations of optimal revenue that exist in the literature depend on information rents being smaller for bidders who are more optimistic about the value.

The qualitative impact is that while earlier results found that optimal auctions discriminate in favor of more optimistic bidders, we find that optimal auctions discriminate in favor of less optimistic bidders, since they obtain less information rents from being allocated the good. In certain cases, the optimal auction reduces to a fully inclusive posted price, under which the likelihood that a given bidder wins the good is independent of their private information. In many cases, however, the optimal auction strictly favors bidders whose signals are not the highest. This is necessarily the case when there is no gap between the seller's cost and the support of bidder's values.

More broadly, we have extended the theory of optimal auctions to a new class of common value models. The analysis yields substantially different insights than those obtained by the earlier literature. We are hopeful that the methodologies we have developed can be used to understand optimal auctions in other as-yet unexplored interdependent value environments.

A Appendix

Proof of Lemma 1. The proof follows closely that of Lemma 2 in Myerson (1981). Let us define $\hat{q}_i(x, y)$ to be the likelihood that bidder i is allocated the good conditional on bidder i 's signal being x and on the highest of the other's signals being y , i.e.,

$$\hat{q}_i(x, y) = \frac{1}{g_{N-1}(y)} \int_{\{s_{-i} | \max_{j \neq i} s_j = y\}} q_i(x, s_{-i}) f_{-i}(s_{-i}) ds_{-i}.$$

This is related to \hat{q}_i by

$$\hat{q}_i(x) \equiv \int_{y=\underline{s}}^x \hat{q}_i(x, y) g_{N-1}(y) dy,$$

that is the probability that bidder i is allocated the good *and* that bidder i has the high signal, conditional on i 's signal being x . If $s_i \leq s'_i$, then

$$\begin{aligned} u_i(s'_i, s_i) &= s'_i \hat{q}_i(s_i) + \int_{y=s_i}^{\bar{s}} \max\{s'_i, y\} \hat{q}_i(s_i, y) g_{N-1}(y) dy - t_i(s_i) \\ &\geq s'_i \hat{q}_i(s_i) + \int_{y=s_i}^{\bar{s}} y \hat{q}_i(s'_i, y) g_{N-1}(y) dy - t_i(s'_i). \end{aligned}$$

Thus,

$$u_i(s'_i, s_i) - u_i(s_i) \geq (s'_i - s_i) \hat{q}_i(s_i),$$

and hence

$$u_i(s'_i) \geq u_i(s_i) + (s'_i - s_i) \hat{q}_i(s_i).$$

We therefore conclude that u_i is monotonically increasing. Similarly, if $s_i \leq s'_i$, we have:

$$\begin{aligned} u_i(s_i, s'_i) &= s_i \widehat{q}_i(s'_i) + \int_{y=s_i}^{s'_i} (y - s_i) \widehat{q}_i(s'_i, y) g_{N-1}(y) dy \\ &\quad + \int_{y=s'_i}^{\bar{s}} y \widehat{q}_i(s'_i, y) g_{N-1}(y) dy - t_i(s'_i) \\ &\geq s_i \widehat{q}_i(s'_i) + \int_{y=s'_i}^{\bar{s}} y \widehat{q}_i(s'_i, y) g_{N-1}(y) dy - t_i(s'_i) \end{aligned}$$

Thus,

$$u_i(s_i) + (s'_i - s_i) \widehat{q}_i(s'_i) \geq u_i(s'_i).$$

Letting $\Delta = s'_i - s_i$, we conclude that

$$u_i(s_i) - u_i(s_i - \Delta) \leq \Delta \widehat{q}_i(s_i) \leq u_i(s_i + \Delta) - u_i(s_i),$$

so that $u_i(s_i)$ is differentiable and $u'_i(s_i) = \widehat{q}_i(s_i)$. \square

Proof of Proposition 4. The value of the object is given by

$$v = \max \{s_1, \dots, s_N, s_{N+1}, \dots, s_{N+K}\}. \quad (17)$$

The revenue of the (absolute) second-price auction with $N + K$ bidders is the second order statistic of $N + K$ independent and identical signals.

In the optimal auction with only N bidders, the value of the object is still given by (17). The seller could still offer the object to a random bidder i at a posted price equal to the maximum of $N + K - 1$ signal realization, and thus implicitly assuming that bidder i has the lowest possible signal realization:

$$p = \mathbb{E} [\max \{s_1, \dots, s_i = \underline{s}, \dots, s_{N+K}\}].$$

The resulting revenue is equal to the first order statistic of $N + K - 1$ independent and identical signals. As a moment of reflection should convince the reader, that the latter is always larger than the former, this establishes the result. After all, the first order statistic of $N + K - 1$ is equal to the first order statistic of $N + K$ random variables after we randomly remove one of variables. But in second order statistic of $N + K$, we remove by design the highest realization, thus leading to a lower expected value.

Formally, the expectation of the second order statistic for $Z = N + K$ independent and identically distributed random variables is given by

$$\int s(Z(Z-1)(1-F(s))F(s)^{Z-2}f(s)ds,$$

and the first order statistic of $Z-1 = N + K - 1$ random variables is given by

$$\int s((Z-1)F(s)^{Z-2}f(s)ds.$$

The difference is given by

$$\begin{aligned} & \int s(Z(Z-1)(1-F(s))F(s)^{Z-2}f(s)ds - \int s((Z-1)F(s)^{Z-2}f(s)ds \\ &= \int (1-(Z-1)(1-F(s)))s(Z-1)F(s)^{Z-2}f(s)ds \\ &= (Z-1) \left(\int sf(s) ((Z-1)F(s)^{Z-1} - (Z-2)F(s)^{Z-2}) ds \right) \end{aligned}$$

From integration by parts we get that

$$\int s (kf(s)F(s)^k) ds = 1 - \int (F(s) + sf(s)) (F(s)^k) ds,$$

and thus

$$\begin{aligned} & (Z-1) \left(\int sf(s) ((Z-1)F(s)^{Z-1} - (Z-2)F(s)^{Z-2}) ds \right) \\ &= (Z-1) \int (F(s) + sf(s)) (F(s)^{Z-2} - F(s)^{Z-1}) ds \\ &> 0, \end{aligned}$$

which completes the argument. □

References

- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2017): “First Price Auctions with General Information Structures: Implications for Bidding and Revenue,” *Econometrica*, 85, 107–143.
- BULOW, J. AND P. KLEMPERER (1996): “Auctions vs Negotiations,” *American Economic Review*, 86, 180–194.
- (2002): “Prices and the Winner’s Curse,” *RAND Journal of Economics*, 33, 1–21.
- CAMPBELL, C. AND D. LEVIN (2006): “When and why not to auction,” *Economic Theory*, 27, 583–596.
- CARROLL, G. AND I. SEGAL (2016): “Robustly Optimal Auctions with Unknown Resale Opportunities,” Tech. rep., Stanford University.
- GUPTA, M. AND B. LEBRUN (1999): “First Price Auctions with Resale,” *Economics Letters*, 64, 181–185.
- HAILE, P. (1999): “Auctions with Resale,” Tech. rep., University of Wisconsin.
- (2003): “Auctions with Private Uncertainty and Resale Opportunities,” *Journal of Economic Theory*, 108, 72–110.
- HARSTAD, R. AND R. BORDLEY (1996): “Lottery Qualification Auctions,” in *Advances in Applied Microeconomics: Auctions*, ed. by M. Baye, JAI Press, vol. 6.
- MILGROM, P. AND R. WEBER (1982): “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 50, 1089–1122.
- MYERSON, R. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- RILEY, J. AND W. SAMUELSON (1981): “Optimal Auctions,” *American Economic Review*, 71, 381–392.
- VARTIAINEN, H. (2013): “Auction Design without Commitment,” *Journal of the European Economic Association*, 11, 316–342.