

Countering the Winner’s Curse: Optimal Auction Design in a Common Value Model*

Dirk Bergemann Benjamin Brooks Stephen Morris

June 1, 2019

Abstract

We characterize revenue maximizing mechanisms in a common value environment where the value of the object is equal to the highest of bidders’ independent signals. If the object is optimally sold with probability one, then the optimal mechanism is simply a posted price, with the highest price such that every type of every bidder is willing to buy the object. A sufficient condition for the posted price to be optimal among all mechanisms is that there is at least one potential bidder who is omitted from the auction. If the object is optimally sold with probability less than one, then optimal mechanisms skew the allocation towards bidders with lower signals. This can be implemented via a modified Vickrey auction, where there is a random reserve price for just the high bidder. The resulting allocation induces a “winner’s blessing,” whereby the expected value conditional on winning is higher than the unconditional expectation. By contrast, standard auctions that allocate to the bidder with the highest signal (e.g., the first-price, second-price or English auctions) deliver lower revenue because of the winner’s curse generated by the allocation rule. Our qualitative results extend to more general common value environments where the winner’s curse is large.

KEYWORDS: Optimal auction, common values, maximum game, posted price, reserve price, revenue equivalence.

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, Paul Klemperer, Ricky Vohra and Ilya Segal. We thank Christian Baker and Ian Ball for research assistance. An early version of this material circulated under the title “Selling to Intermediaries: Optimal Auction Design in a Common Value Model” (2017) and is entirely subsumed in the current version.

1 Introduction

Whenever there is interdependence in bidders' willingness to pay for a good, each bidder must carefully account for that interdependence in determining how they should bid. A classic motivating example concerns wildcatters competing for an oil tract in a first- or second-price auction. Each bidder drills test wells and forms his bids based on his sample. Richer samples suggest more oil reserves, and are associated with higher equilibrium bids. Since the high bidder wins the auction, winning means that the other bidders' samples were relatively poor. The expected value of the tract conditional on winning is therefore less than the interim expectation of the winning bidder conditional on just his signal. This *winner's curse* results in more bid shading relative to a naïve model in which bidders do not account for selection and treat interim values as ex post values.

This paper studies the design of revenue maximizing auctions in settings where there is the potential for a strong winner's curse. The prior literature on optimal auctions has largely focused on the case where values are *private*, meaning that each bidder's signal perfectly reveals his value and there is no interdependence. A notable exception is Bulow and Klemperer (1996), who generalized the revenue equivalence theorem of Myerson (1981) to models with interdependent values. They gave a condition on the form of interdependence under which revenue is maximized by an auction that, whenever the good is sold, allocates the good to the bidder with the highest signal. We will subsequently interpret the Bulow-Klemperer condition as saying that the winner's curse effect is not too strong, which roughly corresponds to a limit on how informative high signals are about the value. Aside from this work, the literature on optimal auctions with interdependent values and independent signals appears to be quite limited.^{1,2}

Our contribution is to study optimal auctions in the opposite case where the winner's curse is quite strong, while maintaining the hypothesis that signals are independent. For our main results, we focus on a simple model where the bidders have a pure common value for the good, the bidders receive independent signals, and the common value is equal to

¹Myerson (1981) includes a case where the bidders have interdependent and additively separable values, meaning that a bidder's value is a function of their own signal plus some function of the others' signals. In addition, the gains from trade between the seller and a given bidder are assumed to only depend on that bidder's private type. In contrast, we study environments where the gains from trade depends on all signals. Bulow and Klemperer (2002) study the additively separable case where the gains from trade depend on all the signals. If the winner's curse is strong, they conclude that an inclusive posted price is optimal among mechanisms that always allocate the good.

²A great deal of work on auction design with interdependent values has focused on the case where signals are correlated. For example, Milgrom and Weber (1982) show that when signals are affiliated, English auctions generate more revenue than second-price auctions, which in turn generate more revenue than first-price auctions. Importantly, this result follows from correlation in signals, and not interdependence per se.

the highest signal. We refer to this as the *maximum signal model*. For this environment, the winner’s curse in a standard auction is quite severe. Indeed, there is a precise sense in which this is the environment that has the *largest* winner’s curse: As shown by Bergemann, Brooks and Morris (2017; 2019), among all type spaces with the same distribution of a common value, this is the one that minimizes expected revenue in the first-price auction. It also minimizes revenue in second-price and English auctions if one restricts attention to affiliated-values models as in Milgrom and Weber (1982). Collectively, we refer to these as *standard auctions*. Beyond its theoretical interest, the maximum signal model captures the idea that the most optimistic signal is a sufficient statistic for the value. This would be the case if the bidders’ signals represented different ways of using the good, e.g., possible resale opportunities if the bidders are intermediaries,³ or possible designs to fulfill a procurement order, and the winner of the good will discover the best use ex post.

This model was first studied by Bulow and Klemperer (2002). They showed that the second-price auction has a “truthful” equilibrium in which each bidder submits a bid equal to their signal. This bid is less than the interim expected value for every type except the highest. Indeed, the bid shading is so large that the seller can increase revenue simply by making the highest take-it-or-leave-it offer that would be accepted by all types. We refer to this mechanism as an *inclusive posted price*. In the equilibrium of the inclusive posted price mechanism, all bidders indicate they are willing to purchase the good and are equally likely to be allocated the good. Thus, winning the good conveys no information about the value and hence the winner’s curse is completely eliminated. Importantly, while Bulow and Klemperer showed that the posted price generates more revenue than standard auctions, their analysis left open the possibility that there were other mechanisms that generated even more revenue, even in the case when the good is required to be always sold.⁴

Indeed, revenue is generally higher if the seller exercises monopoly power and rations the good when values are low. This is the case in the private value model as established by Myerson (1981), and it continues to be the case here. A simple way to do so would be to set an *exclusive posted price*, i.e., a posted price at which not all types would be willing to buy. This however turns out to be far from optimal: A high-signal bidder would face less

³One could also assume that resale takes place between the bidders, 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. The recent work of Carroll and Segal (2019) also studies optimal auction design in the presence of resale. They argue that a worst-case model of resale involves the values becoming complete information among the bidders, with the high-value bidder having all bargaining power.

⁴Bulow and Klemperer (2002) establish the optimality of the posted price among mechanisms that always allocate the good when the bidders’ common value is equal to the sum of their signals, and the distribution of signals exhibits a decreasing hazard rate, see Proposition 3. Campbell and Levin (2006) provide additional arguments for the revenue dominance of the posted price mechanism in related common value environments.

competition in a “tie break” if the others’ signals are low, thus again inducing a winner’s curse, and depressing bidders’ willingness to pay.

A first key result presents a simple mechanism that improves on any exclusive posted price. In this mechanism, the good is allocated to all bidders with equal likelihood if and only if *some* bidder’s signal exceeds a given threshold. This allocation can be implemented with the following *two-tier posted price*: The bidders express either high interest or low interest in the good. If at least one bidder expresses high interest, the good is offered to a randomly chosen bidder, and otherwise, the seller keeps the good. When a given bidder is offered the good, it is offered at the low price if all other bidders expressed low interest, and it is offered at a high price if at least one other bidder expresses high interest. In equilibrium, bidders express high interest if and only if their signal exceeds a threshold, and prices are set such that conditional on being offered the good, bidders want to accept. Curiously, rather than inducing a winner’s curse, this mechanism induces a *winner’s blessing*: if a bidder has a low signal, and therefore expressed low interest, being allocated the good indicates that others’ signals must be relatively high. This leads to a higher posterior expectation of the value, and hence greater willingness to pay even if one had expressed low interest.

While this mechanism does better than any exclusive posted price, it is possible to go even further. The optimal mechanism, it turns out, induces a winner’s blessing for *every type*. This is achieved by an allocation that—for any realized profile of signals—favors bidders with lower signals. We discuss a number of ways to implement the optimal mechanism, but one method is to use a generalization of the two-tier pricing, which is a *two-tier random reserve*: The high-signal bidder is allocated the good only if his signal exceeds a random reserve price, in which case he pays the maximum of that price and the others’ bids. Otherwise, the good is allocated to one of the other bidders the highest of the others’ bids. A concern is that the extra hurdle for the high bidder would induce bidders to underreport so as to avoid being the high bidder. The trade-off is that with a lower report, the bidder would lose surplus from the event that he still makes the high bid but that bid is less than the realized reserve price. The reserve price distribution is tuned just so that bidders are indifferent to underreporting. Indeed, this temptation to underreport is the key to deriving a tight bound on the seller’s revenue that proves that this mechanism is optimal.

Whenever the optimal two-tier random reserve mechanism allocates the good with probability one, the mechanism reduces to an inclusive posted price. This occurs whenever the lowest possible value is sufficiently large. Alternatively, a sufficient condition for an inclusive posted price to be optimal is that there is at least one bidder who is omitted from the auction. Moreover, if we restrict attention to auctions that allocate the good with probability one, then the inclusive posted price is always the revenue maximizing mechanism. We

thus strengthen the foundation for posted prices introduced in earlier work of Bulow and Klemperer (2002) by proving optimality in the maximum signal model.

The proof that the two-tier random reserve mechanism is optimal utilizes a novel argument. Standard optimal mechanism design in the additively separable case, as in Myerson (1981), relies on the well-known result that an allocation is implementable if and only if each bidder's interim allocation probability is weakly increasing in his signal. Even when interim monotonicity fails to characterize implementability, it is sometimes possible to show that the optimal allocation subject to only local incentive compatibility is also implementable, as it is the case in the model of Bulow and Klemperer (1996). By contrast, in the maximum signal model, interim monotonicity is neither necessary nor sufficient for implementability, and the optimal allocation subject to the local relaxation only is not incentive compatible. Our novel argument involves using non-local incentive constraints to establish a lower bound on bidder surplus. We then construct a mechanism that attains a corresponding upper bound on revenue.

Finally, we argue that our key qualitative results extend beyond the pure common-value maximum signal model to a wide range of interdependent-value environments that exhibit *increasing information rents*. In the case of common values, this condition captures the idea that higher signals are substantially more informative about the value than lower signals. We describe natural mechanisms with exclusion that generate more revenue than either the inclusive or exclusive posted price mechanisms. We also describe implementable allocations that generate even more revenue under weak conditions. It remains an open question whether posted prices continue to be optimal among efficient mechanisms in the presence of increasing information rents. We suspect that the pattern of binding incentive constraints at the optimal mechanism could in general be quite complicated. This presents a major challenge for future research on optimal auctions in general interdependent value settings.

The rest of this paper proceeds as follows. Section 2 describes the model. Section 3 shows how to increase revenue by moving from standard auctions that generate winner's curse to posted price mechanisms that generate winner's blessing. Section 4 proves the optimality of these mechanisms. Section 5 generalizes our analysis to the case of increasing information rents. Section 6 concludes.

2 Model

2.1 Environment

There are N bidders for a single unit of a good, indexed by $i \in \mathcal{N} = \{1, \dots, N\}$. Each bidder i receives a real signal $s_i \in S = [\underline{s}, \bar{s}]$, where $\underline{s} \geq 0$, about the good's value. The bidders' signals s_i are independent draws from an absolutely continuous cumulative distribution F with strictly positive density f . We adopt the shorthand notation that $F_{-i}(s_{-i}) = \times_{j \neq i} F_j(s_j)$ and $F^k(x) = (F(x))^k$ for positive integers k . The bidders all assign the same value to the good, which is the maximum of the signals:

$$v(s_1, \dots, s_N) \triangleq \max\{s_1, \dots, s_N\} = \max s.$$

We frequently use the shorter expression $\max s$ which selects the maximal element from the vector $s = (s_1, \dots, s_N)$. In Section 5, we discuss corresponding results for general common value environments.

The distribution of signals, F , induces a distribution $G(x) \triangleq F^N(x)$ over the maximum signal from N independent draws. We denote the associated density by:

$$g(x) \triangleq N F^{N-1}(x) f(x).$$

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 t to the seller is represented by the utility index:

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

2.2 Direct Mechanisms

We will model a seller who can commit to a mechanism and select the equilibrium played by the bidders. For much of our analysis, and in particular for constructing bounds on revenue and bidder surplus in Theorem 3, 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:

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

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

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

and 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

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

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 $s_i \in S$.

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

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

Since values are common, total surplus only depends on whether the good is allocated, not the identity of the bidder that receives the good. Moreover, the surplus depends only on the value $v(s) = \max_i \{s_i\}$, and not the entire vector s of signals. Let us thus denote by $\bar{q}_i(v)$ the probability that the good is allocated to bidder i , conditional on the value being v , and let

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

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

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

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

3 Countering the Winner's Curse

We start by reviewing the behavior of standard auctions. We then progressively improve the mechanism and allocation to arrive at the optimal mechanism. Optimality is proven in the next section.

3.1 Standard Auctions and the Winner's Curse

First-price, second-price, and English auctions all admit monotonic pure-strategy equilibria, which result in the highest-signal bidder being allocated the good. For ease of discussion, we describe the outcome in terms of the second-price auction. A bidder with a signal s_i forms his *interim* expectation of the common value $\mathbb{E}[v(s_i, s_{-i}) | s_i]$, and then submits a bid $b_i(s_i)$. In the maximum signal model, the signal s_i is a sharp lower bound on the ex post value of the object: given any signal s_i , bidder i knows that the true value of the object has to be in the interval $[s_i, \bar{s}]$. Thus, the interim expectation of the bidder i satisfies

$$\mathbb{E}[v(s_i, s_{-i}) | s_i] > s_i,$$

for all $s_i < \bar{s}$. Yet, given the interim expectation, in the second-price auction, there is an equilibrium in which each bidder i bids only

$$b_i^*(s_i) = s_i.$$

Thus, even though the winning bidder only pays the second-highest bid, the equilibrium bid is equal to the lowest possible realization of the common value given the interim information s_i .

In the monotonic pure-strategy equilibrium b^* , the bidder with the highest signal submits the highest bid. Thus, the signal s_i which provided a sharp *lower* bound on the common value at the bidding stage, becomes a sharp *upper* bound conditional on winning. In fact, it

coincides with the true common value:

$$\mathbb{E}[v(s_i, s_{-i}) | s_i, s_j \leq s_i, \forall j \neq i] = s_i.$$

The expectation of the value conditional on knowing that s_i is the highest signal is simply s_i !

The resulting allocation is:

$$q_i(s) = \begin{cases} \frac{1}{|\arg \max s_j|} & \text{if } s_i = \max s; \\ 0 & \text{otherwise.} \end{cases}$$

Thus, the equilibrium of the second-price auction exhibits a winner's curse: For a realized profile of signals s , it is the bidder with the highest signal who receives the good. The winner therefore learns that his signal was more favorable than all the other signals. In turn, each bidder lowers his equilibrium bid from the interim estimate of the value $\mathbb{E}[v(s_i, s_{-i}) | s_i]$ all the way to the *lowest possible* value in the support of this posterior probability distribution, namely s_i . In this sense, the winner's curse is as large as it can possibly be. In Bergemann, Brooks, and Morris (2019), we show that there is revenue equivalence between the first-price, second-price and the English auction in the specific common value setting, and hence the same winner's curse arise across these standard auction formats.

3.2 Inclusive Posted Prices and Zero Winner's Curse

Given the strength of the winner's curse and the extent of bid shading, it is natural to ask whether other mechanisms can mitigate the winner's curse and thus increase revenue. Bulow and Klemperer (2002) establish that a simple but very specific posted price mechanism can attain higher revenue than the standard auctions with their monotonic equilibria.

In a posted price mechanism with price p , the object is allocated with uniform probability among those bidders who declared their willingness to pay p to receive the object. The specific posted price suggested by Bulow and Klemperer (2002) is the expectation of the highest of $N - 1$ independent draws from the signal distribution F :

$$p_I \triangleq \int_{x=\underline{s}}^{\bar{s}} x d(F^{N-1}(x)). \quad (2)$$

We refer to p_I as the *inclusive posted price*. To wit, p_I is exactly equal to the interim expectation that a bidder i with the lowest possible signal realization $s_i = \underline{s}$ has about the common value of the object. The price p_I is the maximal price with the property that

every type is willing to buy the object. Thus, all types are “included” in the allocation. The inclusive posted price can also be interpreted as the expectation of a particular Vickrey price. That is, for a random bidder to receive the object, he would have to pay the highest signal among the competing bidders, $\max s_{-i}$, the expectation of which is simply p_I .

Proposition 1 (Inclusive Posted Price).

The inclusive posted price yields a higher revenue than the monotonic pure strategy equilibrium of any standard auction.

Bulow and Klemperer (2002), Section 9, established the revenue ranking regarding the English auction. Bergemann, Brooks, and Morris (2019) established a revenue equivalence result for the maximum signal model that establishes the above proposition. The source of the revenue ranking can be explained as follows. Revenue under the inclusive posted price is equal to the expectation of the highest of $N - 1$ independent and identical signals from F . By contrast, the revenue in the standard auctions is equal to the expectation of the second order statistic of N signals. The former must be greater than the latter, since the inclusive posted price revenue can be obtained by throwing out one of N draws at random and then taking the highest of the remaining realizations, whereas the standard auction revenue is obtained by systematically throwing out the highest of the N draws, and then taking the highest remaining.

Note that the allocation induced by the inclusive posted price assigns an equal probability to every bidder i : $q_i(s) \equiv 1/N$. As a result, the event of winning conveys no additional information about the value of the object to any of the winning bidders. Thus, in sharp contrast to the standard auctions, a bidder’s expected value conditional on receiving the object is the same as the unconditional expectation, i.e., there is zero winner’s curse.

Proposition 1 establishes that the inclusive posted price generates higher revenue than the standard auction. An alternative perspective to understand this result is by using the revenue equivalence theorem. Specifically, using local incentive constraints, one can solve for the transfers in terms of the allocation to conclude that expected revenue is equal to the expected *virtual value* of the buyer who is allocated the good. Bulow and Klemperer (1996) derive the virtual value for a general interdependent values model. When the bidders have a common value that is a monotonic and differentiable function $v(s)$ of all bidders’ signals, bidder i ’s virtual value is:

$$\pi_i(s_i, s_{-i}) = v(s_i, s_{-i}) - \frac{1 - F(s_i)}{f(s_i)} \frac{\partial v(s_i, s_{-i})}{\partial s_i}. \quad (3)$$

The first term on the right-hand side is simply the common value of the object, i.e., the social surplus generated by allocating the good. The second term is the inverse hazard rate,

which is a measure of the relative number of higher types who gain an information rent by being able to mimic type s_i . The final term is the sensitivity of the value to bidder i 's signal. Clearly there must be some component of this form, for if the value does not depend on bidder i 's signal, then bidder i has no valuable private information and should not obtain an information rent.

In the maximum signal model, equation (3) simplifies to:

$$\pi_i(s_i, s_{-i}) = \begin{cases} \max s, & \text{if } s_i < \max s; \\ \max s - \frac{1-F(s_i)}{f(s_i)}, & \text{if } s_i = \max s. \end{cases} \quad (4)$$

Now the partial derivative of the valuation function with respect to the signal s_i is simply the indicator function $\mathbb{I}_{\{s_i = \max s\}}$ for whether bidder i has the highest signal or not. A significant implication is that only the high signal bidder receives an information rent, equal to the inverse hazard rate $(1 - F(s_i))/f(s_i)$, so that all bidders other than the highest-signal bidder have a higher virtual utility. In consequence, revenue is higher the lower is the probability that the high-signal bidder is allocated the good. Relative to the standard auctions, the inclusive posted price attains higher revenue because it allocates the object uniformly across all bidders, whereas the standard auctions give it to the high signal bidder with probability one.⁵

3.3 Exclusive Auctions and the Winner's Blessing

A notable feature of the inclusive posted price is that the object is awarded with uniform probability for *every* type profile realization s . In particular, this implies that the object is awarded even if the virtual value of some bidder, or even the average virtual value across all bidders is negative for a given signal profile s .

A reasonable first attempt at raising revenue would be to post a price p that is strictly higher than the inclusive posted price p_I . By definition then, the price p would exceed the interim expectation of the bidder with the lowest possible signal $s_i = \underline{s}$. Any such price p would then induce some threshold $r \in (\underline{s}, \bar{s}]$, so that every bidder i with a signal $s_i \geq r$ would accept the offer and all types below would reject the offer. The resulting assignment probabilities would be:

$$q_i(s) = \begin{cases} \frac{1}{|\{j|s_j \geq r\}|}, & \text{if } s_i \geq r; \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

⁵Note that the value function in the maximum signal is not differentiable, so that the theorem of Bulow and Klemperer (1996) does not apply. It is, however, straightforward to extend their theorem to cover the maximum signal model.

Consequently, we refer to the threshold r as the *exclusion level*. The posted price that implements the exclusion level r is the expectation of the common value for the type $s_i = r$ conditional on receiving the good:

$$p_E \triangleq \frac{\int_{\{s_{-i} | \max s_{-i} \geq r\}} \max\{r, s_{-i}\} q_i(r, s_{-i}) dF_{-i}(s_{-i})}{\int_{\{s_{-i} | \max s_{-i} \geq r\}} q_i(r, s_{-i}) dF_{-i}(s_{-i})}.$$

By extension, we refer to a posted price $p_E > p_I$ as an *exclusive posted price*. But in contrast to the inclusive posted price, the resulting allocation again tilts the allocation towards the higher signal bidders, $s_i > r$, by excluding the low signal bidders, $s_i < r$. Thus, while the exclusive posted price does ration the object, it reintroduces a winner's curse and a results in depressed willingness to pay.

As we realized with the inclusive posted price, revenue would be higher if we reallocated the good to lower signal bidders. For example, revenue would be higher if were instead to implement a uniform allocation conditional on assigning the object:

$$q_i(s) = \begin{cases} \frac{1}{N} & \text{if } \max s \geq r; \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

This mechanism achieves the same exclusion level r and hence maintains the same ex post surplus, but it attains a higher revenue.

One way to effect this allocation is with a *two-tier posted price* (p_L, p_H) . Every bidder is asked to express a high interest or a low interest in the good. If all bidders express low interest, then the seller keeps the good. If at least one bidder expresses high interest, then all bidders are offered a chance to purchase, with equal probability. When bidder i is offered the good, the associated price is either a low price p_L if all other bidders expressed low interest or a high price $p_H > p_L$ if at least one other bidder expresses high interest. The specific prices are

$$p_L \triangleq r,$$

and

$$p_H \triangleq \frac{\int_r^{\bar{s}} x d(F^{N-1}(x))}{1 - F^{N-1}(r)}.$$

Thus, p_H is the expected value of a bidder with signal $s_i \leq r$ conditional on knowing that the highest signal among the remaining $N - 1$ bidders weakly exceeds r .

We claim that there is an equilibrium of this mechanism where bidders express high interest if $s_i \geq r$ and express low interest otherwise. Bidders always agree to buy the good at the offered price, whatever that may be. In fact, this strategy is optimal even if a

bidder i were to know whether $\max s_{-i}$ is less than or greater than r , i.e., whether all of the other bidders express low interest or at least one expresses high interest. If we condition on $\max s_{-i} < r$, there is effectively a posted price of r , and the value is at least the price if and only if $s_i \geq r$. It is a best reply to express high interest and accept the low price when $s_i \geq r$ and to express low interest otherwise. If we condition on $\max s_{-i} \geq r$, then expressing high or low interest result in the same outcome, which is a probability $1/N$ of being offered the good at the high price. The expected value across all s_{-i} is always at least p_H , since the true value is the maximum of s_i and s_{-i} . Thus, one best reply is to express high interest if $s_i \geq r$ and express low interest otherwise.

Note that this mechanism implements the same ex post total surplus as the exclusive posted price, but low-signal bidders are more likely to receive the good for every signal profile. As a result, information rents are reduced relative to the exclusive posted price, and we have proven the following.

Proposition 2 (Two-Tier Posted Price).

The two-tier posted price (p_L, p_H) yields a weakly higher revenue than the exclusive posted price that implements the same exclusion level.

Note that the two-tier posted price again induces a winner's blessing for types $s_i < r$: being allocated the good implies that $\max s_{-i} \geq r$. Thus, we again see the relationship between higher revenue and the presence of a winner's blessing.

3.4 The Optimal Mechanism

The inclusive posted price or the two-tier posted price depress the probability of winning to $1/N$ for the bidder with the highest signal whereas he would have won with probability one in any standard auction, such as the second-price auction. The natural next question is whether there exist mechanisms that reduce the high-signal bidder's probability of winning even further, below the uniform probability $1/N$. We might think this is impossible, based on intuition from the private-value case where higher types must have higher interim allocations. With interdependent values, however, it is possible to skew the allocation against the high-signal bidder, as we now explain.

As a segue, let us first observe that there is another implementation of the allocation (6) induced by the two-tier posted price: bidders report their signals, the good is allocated with uniform probability if the highest report exceeds r , and any bidder who is allocated the good makes the Vickrey payment $\max\{r, s_{-i}\}$. We refer to this as a *two-tier reserve* mechanism: the high bidder faces a non-trivial reserve price r , whereas the low-signal bidders face no such reserve price (although they are still not allocated the good if all of the bids are below r). It

is straightforward to verify that this mechanism is ex-post incentive compatible and ex-post individually rational, i.e., even if the realized signals are complete information among the bidders.

Now consider the following modification of this revelation game. The object is allocated if at least one of the bidders reports a signal exceeding the threshold r . We give the bidder i with the highest reported signal the priority to purchase the object, but we ask him to pay a posted price that is the maximum of the reported signals of the others, and an additional random variable x . Thus, bidder i faces a posted price of $\max\{x, s_{-i}\}$. The distribution of x is denoted by $H(x)$ and has support in $[r, \infty]$. In particular, it is possible for this reserve price to be infinite, in which case it is impossible for the high bid to meet the threshold. Bidder i is allocated the good at the realized price if it is less than his reported signal. Otherwise, one of the other bidders is offered the good at a price equal to the highest bid.

Note that the allocation and transfer rules reduce to the two-tier posted price mechanism when H puts probability $1/N$ on $x = r$ and probability $(N - 1)/N$ on $x = \infty$, in which case we have already showed that truthful bidding is an equilibrium. However, if we choose H to put less probability on $x = r$, then the allocation is effectively skewed towards low-signal bidders, as long as bidding is truthful. This begs the question, for which distributions H is truthful bidding an equilibrium?

Note that for any H , at a truthful strategy, bidders have no incentive to overreport: This can only result in being allocated the good at a price that exceeds the value. Also, reporting any signal less than r is equivalent to reporting a signal of r . Thus, for incentive compatibility, it suffices to check that a bidder i with signal $s_i \geq r$ does not want to misreport $s'_i \in [r, s_i]$. To that end, consider the surplus of such a bidder, assuming that all other bidders report truthfully. This is

$$u_i(s_i, s'_i) = \int_{x=r}^{s'_i} (s_i - x) d(H(x) F^{N-1}(x)) \\ + \int_{x=s'_i}^{\bar{s}} (\max\{s_i, x\} - x) \frac{1 - H(x)}{N - 1} d(F^{N-1}(x)).$$

The derivative of this expression with respect to s'_i is

$$(s_i - s'_i) \left[d(H(s'_i) F^{N-1}(s'_i)) - \frac{1 - H(s'_i)}{N - 1} d(F^{N-1}(s'_i)) \right].$$

So, a sufficient condition for downward deviations to not be attractive is that the term inside the brackets is non-negative for all s'_i , which reduces to

$$\frac{dH(x)}{1 - NH(x)} \geq \frac{1}{N-1} \frac{d(F^{N-1}(x))}{F^{N-1}(x)}.$$

If we solve the above inequality as an equality, with the boundary condition $H(r) = 0$, we obtain

$$H(x) \triangleq \frac{1}{N} \left(1 - \left(\frac{F(r)}{F(x)} \right)^N \right). \quad (7)$$

With this particular distribution for the high-bidder's reserve, we refer to this game form as the *two-tier random reserve mechanism*. We have just verified that this mechanism is incentive compatible. In fact, bidders are indifferent between truthful reporting and all downward misreports.

Note that by construction $H(r) = 0$, so that a bidder with the highest signal close to the exclusion threshold r is unlikely to receive the object. Moreover, even the bidder with the highest possible signal \bar{s} receives the object with probability less than $1/N$ since

$$H(\bar{s}) = \frac{1}{N} (1 - F^N(r)) < \frac{1}{N}.$$

We have therefore completed the proof of the following result:

Proposition 3 (Two-Tier Random Reserve).

The two-tier random reserve mechanism yields a higher revenue than the two-tier posted price that implements the same exclusion level.

The two-tier random reserve mechanism has the feature that the resulting interim probability $q_i(s_i)$ of receiving the object is constant in the signal s_i . Specifically,

$$\begin{aligned} q_i(s_i) &= F^{N-1}(s_i)H(s_i) + \int_{x=s_i}^{\bar{s}} \left(\frac{1 - H(x)}{N-1} \right) d(F^{N-1}(x)) \\ &= \frac{1 - F^N(r)}{N}. \end{aligned}$$

The interim allocation probability is the product of the probability that the object is allocated to some bidder and the probability that bidder i receives the object conditional on it being allocated at all. The two-tier random reserve in fact favors bidders with lower signals. In particular, the ex post probability $q_i(s)$ of receiving the object in the two-tier random reserve

mechanism can be computed to be:

$$q_i(s) = \begin{cases} H(\max s) & \text{if } s_i > s_j \forall j \neq i \text{ and } s_i \geq r; \\ \frac{1}{N-1} (1 - H(\max s)) & \text{if } s_i < \max s \text{ and } \max s \geq r; \\ 0 & \text{otherwise.} \end{cases}$$

Conditional on the realized signal profile, high-signal bidders are strictly *less* likely to receive the good. Thus, conditioning on winning results in a higher expected value for all types. In effect, the random reserve price turns the winner's curse into a *winner's blessing*. This results in an increased willingness-to-pay in equilibrium, and an increase in the revenue generated by the auction.

We note that the two-tier random reserve mechanism, while interim incentive compatible, is not ex-post incentive compatible anymore, as was the inclusive or two-tier posted price. If more than one bidder reported a signal s_i that exceeded the threshold r , then ex-post the high signal bidder would prefer to report a lower signal. Since the report of the other bidder would already guarantee that the object is allocated, a downward report would increase the probability of receiving the object, and lower the expected price to be paid.

Note that because the interim allocation is constant, the interim transfer must be constant as well. The highest type \bar{s} is certain that the value is \bar{s} and by construction is indifferent to all downward deviations, so that the payoff $\bar{s}q_i(s_i) - t_i(s_i)$ must be independent of s_i . But since $q_i(s_i)$ is constant, $t_i(s_i)$ must be constant as well. Thus, another implementation of this allocation is that every bidder pays the constant interim transfer as an entry fee, after which they make their reports, and the optimal allocation is implemented.

Until now, we have treated the exclusion threshold r as a fixed parameter. Thus, there is actually a one-dimensional family of two-tier random reserve mechanisms, indexed by r . At the extreme where $r = 0$, the associated allocation reduces to the uniform probabilities implemented by the inclusive posted price mechanism.

The revenue maximizing threshold can be understood as follows. Expected revenue is the difference between total surplus and bidder surplus. The effect of increasing the exclusion threshold on total surplus is immediate: surplus is lost from the good not being allocated when the value is r . Next, since a bidder receives positive surplus only if he has the highest signal, bidder surplus in the two-tier random reserve mechanism is:

$$\begin{aligned} U &= \int_{s=r}^{\bar{s}} \int_{x=r}^s (s-x) d(H(x) F^{N-1}(x)) dF(s) \\ &= \int_{x=r}^{\bar{s}} \frac{1}{N} \frac{1-F(x)}{F(x)} (F^N(x) - F^N(r)) dx, \end{aligned}$$

where we have simply plugged in the definition (7). Thus, the effect of an increase in r on U is

$$\frac{dU}{dr} = -\frac{1}{N} \int_{x=r}^{\bar{s}} \frac{1-F(x)}{F(x)} dx \frac{d(F^N(r))}{dr}.$$

The overall effect of increasing r on revenue is therefore

$$\frac{dR}{dr} = -\psi(r) \frac{d(F^N(r))}{dr},$$

where

$$\psi(r) \triangleq r - \int_{x=r}^{\bar{s}} \frac{1-F(x)}{F(x)} dx. \quad (8)$$

Note that $\psi(r)$ is continuous and strictly increasing in r , and it is positive when r is sufficiently large. As a result, revenue is single peaked in the reserve price, and the optimal reserve price r^* is the smallest r such that $\psi(r) \geq 0$.

Thus, we have established that the two-tier random reserve mechanism with threshold r^* generates more revenue than standard auctions, inclusive and exclusive posted prices, and two-tier posted prices. Indeed, our main result is that it is revenue maximizing among all incentive compatible and individually rational mechanisms:

Theorem 1 (Optimality of Two-Tier Random Reserve).

The two-tier random reserve mechanism with cutoff r^ maximizes revenue across all IC and IR direct mechanisms.*

When the gains from the bias toward low-signal bidders is small relative to the cost of restricting supply, the inclusive posted prices indeed emerges as the optimal mechanism.

Corollary 1 (Optimality of Inclusive Posted Price).

The inclusive posted price maximizes revenue across all IC and IR direct mechanisms if and only if $\psi(\underline{s}) \geq 0$.

We will also show that the inclusive posted is *always the* optimal mechanism if one restricts attention to mechanisms where the object has always to be allocated. We refer to this class of mechanism as *must-sell mechanisms*.

Theorem 2 (Must-sell Optimality of Inclusive Posted Prices).

If the object is required to be allocated with probability one, then the inclusive posted price maximizes expected revenue across all IC and IR mechanisms.

We prove these theorems in the next section. We emphasize that the arguments are novel and require the explicit consideration of global incentive constraints. In particular,

the optimality of the posted price within must-sell mechanisms does not follow from the arguments reported in Bulow and Klemperer (2002).⁶

4 Optimal Mechanisms

The broad strategy in proving Theorems 1 and 2 is to show that the allocations described above attain an upper bound on revenue, where that upper bound is derived using a subset of the bidders' incentive constraints. Before developing this argument, we briefly review existing approaches and explain why they are inadequate for our purposes.

4.1 Local Versus Global Incentive Compatibility

The standard approach in auction theory is to use local incentive constraints to solve for transfers in terms of allocations, and rewrite the expected revenue in terms of the expected virtual value of the bidder who is allocated the good. Note that the formula for the virtual value (4) tells us what revenue must be as a function of the allocation if local incentive constraints are satisfied, but it does not tell us which allocations can be implemented subject to all incentive constraints.

In the case studied by Bulow and Klemperer (1996) where the winner's curse effect is weak, the virtual value is pointwise maximized by allocating the good to the bidder with the highest signal (that is, whenever allocating the good is better than withholding it). One can then appeal to existing characterizations of equilibria of English auctions with interdependent values à la Milgrom and Weber (1982) to show that such an allocation is implementable. This proof strategy will not work in the maximum signal model. As we have argued, the bidder with the highest signal always has the lowest virtual value, so pointwise maximization of $\pi_i(s)$ would *never* allocate the object to the high-signal bidder. Moreover, it is straightforward to argue that such an allocation would not be incentive compatible. If it were, then the highest type would receive the good with probability zero, and the lower types with probability one. The high type must therefore be *paid* by the mechanism an amount equal to the positive surplus that could be obtained by pretending to be the lowest type. But this surplus must be strictly greater than that obtained by the lowest type, thus tempting the lowest type to misreport as the highest.

⁶As we mentioned earlier in Footnote 1, Bulow and Klemperer (2002) establish the optimality of the inclusive posted price mechanism among efficient mechanisms in a different environment (the "wallet game") where the value is the sum of independent signals. This case is "additively separable," so that the usual monotonicity condition on the interim allocation is necessary and sufficient for implementability. However, the maximum signal model is not additively separable, and therefore necessitates new arguments.

When the pointwise maximization approach fails, one needs to explicitly include global incentive constraints in the optimization problem, in addition to the local incentive constraints that are implicit in the revenue equivalence formula. In the additively separable case, e.g., where the value is the sum of the bidders' signals, global incentive constraints are equivalent to the interim allocation being non-decreasing. But in general interdependent value models, interim monotonicity is neither necessary nor sufficient for incentive compatibility, and we know of no general characterization of which allocations are implementable in these environments.⁷

Thus, we must find a new way of incorporating global constraints into the seller's optimization problem. The key question is: which global constraints pin down optimal revenue? The analysis of the preceding section suggests that the critical constraints might be those corresponding to downward deviations: Each bidder accrues information rents only when he is allocated the good and has the highest signal. Thus, 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. Moreover, *all* of the downward constraints are binding in the putative optimal allocations, thus suggesting that they all must be used to obtain a tight upper bound on revenue.

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 not common, it is not just whether but also to whom the good is allocated that determines total surplus.

⁷The following two allocation rules—within the maximum signal model—show that interim monotonicity of the allocation is neither necessary nor sufficient for incentive compatibility. Consider the case of two bidders, $i = 1, 2$ who have binary signals $s_i \in \{0, 1\}$, which are equally likely. We consider two allocation rules for bidder 1, q_1 , as given by one of the following tables. The allocation for bidder 2 is constant across signal realizations and is simply $q_2 = 0$.

q_1	s_2	0	1	q_1	s_2	0	1
		1	0			0	1
s_1	1	1	0	s_1	1	0	1
	0	1	1		0	1	0

The allocation on the left is not interim monotone in s_1 but is easily implemented by charging a price of s_2 whenever the good is allocated to bidder 1. The allocation on the right is interim monotone but cannot be implemented: The low type must pay an interim transfer which is at least that of the high type in order to prevent the high type from misreporting. But this implies the low type would prefer to misreport, to pay weakly less and get the good when it is worth 1 rather than 0. These examples could be made efficient by adding a third bidder, who receives the good when it would not be allocated to bidder 1, with zero transfer. Note that the third bidder's allocation probability is independent of their signal. As a result, the example can be made symmetric simply by randomly permuting the roles of the bidders.

4.2 Proof of Theorems 1 and 2

We now begin our formal proof. Consider the following one-dimensional family of deviations in the direct mechanism: 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

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

where we recall that $\bar{q}_i(v)$ defined in (1) 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 , his 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 second 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 ,

$$u_i(s_i) \geq \frac{1}{F(s_i)} \left(\int_{x=\underline{s}}^{s_i} u_i(x) f(x) dx + \int_{x=\underline{s}}^{s_i} (s_i - x) \bar{q}_i(x) g(x) dx \right). \quad (9)$$

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

$$u(y) \geq \frac{1}{F(y)} \left(\int_{x=\underline{s}}^y u(x) f(x) dx + \lambda(y) \right), \quad (10)$$

where

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

and

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

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 (10). Our first main result, Theorem 3, asserts that this minimum is attained by the function

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

which solves (10) as an equality when $\underline{u}(\underline{s}) = 0$. In fact, \underline{u} is the pointwise smallest interim utility function that is non-negative and satisfies (10). Indeed, if the constraint held as a strict inequality at some y , then 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

$$\underline{U} \triangleq \int_{y=\underline{s}}^{\bar{s}} \underline{u}(y) f(y) dy = \int_{y=\underline{s}}^{\bar{s}} \int_{x=y}^{\bar{s}} \frac{1-F(x)}{F(x)} dx \bar{q}(y) d(F^N(y)) dy, \quad (11)$$

and revenue is therefore at most

$$\bar{R} \triangleq TS - \underline{U} = \int_{v=\underline{s}}^{\bar{s}} \psi(v) \bar{q}(v) dF^N(v) dv, \quad (12)$$

where $\psi(v)$ was defined in (8) as the virtual value from allocating the good when the value is v . This result is stated formally as follows:

Theorem 3 (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} and expected revenue is bounded above by \bar{R} .

Proof of Theorem 3. It remains to prove formally that \underline{u} is the lowest u that satisfies (10). Define the function operator

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

on the space of non-negative integrable utility functions. From the argument in the text, it is clear that any average indirect utility function u that is induced by an IC and IR mechanism must satisfy (10), which is equivalent to $u \geq \Gamma(u)$. It is easily verified that \underline{u} is a fixed point of Γ . For then

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

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 is a consequence of 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 \widehat{u} that is smaller than \underline{u} , then it must be that $\widehat{u}(s) \leq \underline{u}(s)$ for all s , with a strict inequality for some positive measure set of s . Moreover, it must be that $\underline{u}(x) - \widehat{u}(x)$ goes to zero as x goes to \underline{s} (and hence, cannot be constant for all x). Let $\|\cdot\|$ denote the sup norm, and suppose that $\|\Gamma(\underline{u}) - \Gamma(\widehat{u})\|$ is attained at s . 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 (10) 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

We can now complete the proofs of our main theorems.

Proof of Theorems 1 . If the seller can withhold the good, then we can derive an upper bound on optimal revenue by maximizing the bound (12) pointwise. Since $\psi(v)$ is monotonic, the pointwise maximum is attained by the allocation

$$\bar{q}(v) = \begin{cases} 1 & \text{if } v \geq r^*; \\ 0 & \text{if } v < r^*, \end{cases} \quad (13)$$

where

$$r^* = \min \{v | \psi(v) \geq 0\}.$$

Clearly, this is the allocation that is implemented by the two-tier random reserve mechanism. Moreover, we have already verified that all downward incentive constraints bind, so that the revenue upper bound is attained. \square

Proof of Theorem 2. If the good must be allocated, then $\bar{q}(v) = 1$ for all v , which completely determines the upper bound on revenue from misreporting a redrawn lower signal. The upper bound will be attained by any mechanism that implements this allocation and makes all of the downward incentive constraints bind. But all types are treated the same way by the inclusive posted price, so that all downward constraints bind, and the upper bound on revenue is attained. \square

The function $\psi(v)$ can be interpreted as the virtual value from allocating the good conditional on the value being v , albeit a different virtual value than the one obtained from only local incentive constraints: Both virtual values have a term equal to the value of the good, which is simply the change in social surplus from allocating versus not allocating. They differ in the second part, which is the total information rent from the allocation. The local incentive constraints indicate that for every unit probability that a bidder of type s_i is allocated the good when they have the highest type, all higher types get a unit information rent, so that the relative information rent is $(1 - F(s_i)) / f(s_i)$, the inverse hazard rate. In contrast, the *global* incentive constraints (10) indicate that the rate of increase in ex-ante bidder surplus per unit increase in $\bar{q}(v)$ is precisely

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

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 $(v - s) d(F^N(v)) / F(s)$. 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

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

and $\rho(s|v) = 0$ for $s < v$. By integrating this equation with the integrating factor $f(s) / F(s)$, we conclude that

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

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

We can compare the pattern of binding incentive constraints for the standard auctions (first-price, second-price, and English) and the optimal mechanism visually in Figure 1. Here we consider an example with two bidders where the value is standard uniform, so $G(v) = v$ and $F(s) = \sqrt{s}$. Each graph describes the indirect utility for three types, $s_i \in \{1/4, 1/2, 3/4\}$ in the second-price auction and the optimal mechanism, respectively. Each curve describes for each type s_i the indirect utility the type would receive from reporting any other signal $s'_i \in [0, 1]$. The equilibrium utility supported by truth-telling is indicated by the corresponding vertical line. The first observation is that the equilibrium utility—the information rent of each bidder—drops by at least a factor of four by moving from the second-price auction to the optimal mechanism. Thus, the revenue gain from eliminating the winner's curse is substantial. Second, in moving from a standard auction to the optimal mechanism, the structure of the binding incentive constraints reverses completely. In standard auctions, the winner's curse is so strong, and consequently the equilibrium bid is so low, that each bidder is indifferent between his equilibrium bid $b_i(s_i)$ and any higher bid on the entire unit interval! Thus, all *upward* incentive constraints are binding. By contrast, in the optimal mechanism, all *downward* incentive constraints are binding. That is, the information rent of each bidder is lowered so far that each bidder is indifferent between reporting truthfully and offering any misreport between 0 and the true signal s_i . We note that the contrast in the structure of the incentive constraints holds true for all continuous signal distributions in the maximum signal model. That is, all upward constraints bind in standard auctions, and all downward constraints bind in the optimal mechanism.

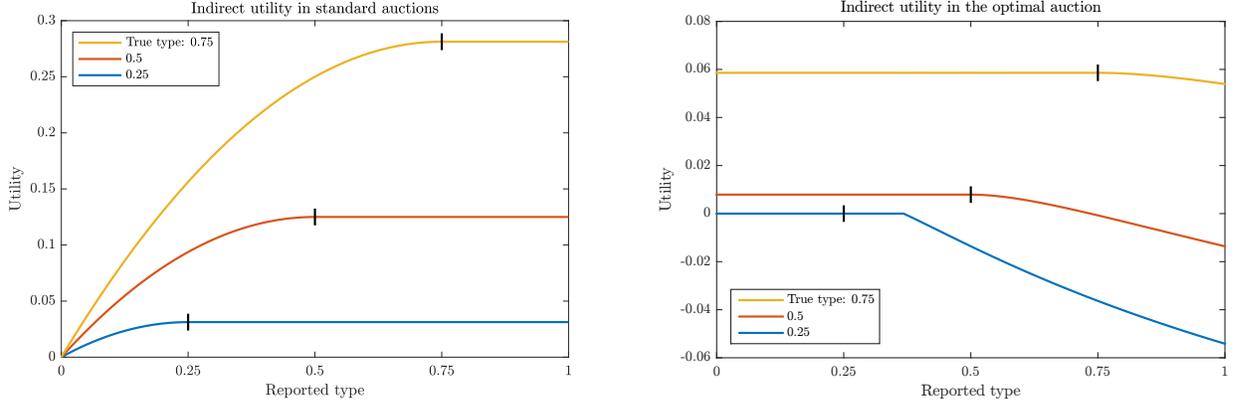


Figure 1: Uniform Downward vs Upward Incentive Constraints

4.3 Uniqueness of the Optimal Allocation

Theorem 3 gives us a bound on revenue that is attained by the two-tier random reserve mechanism, thus proving its optimality. A natural next question is whether there are other optimal allocations. The answer is by-and-large no: any optimal allocation must share a number of key properties with that induced by the one we have constructed.

First, the allocation we constructed has the property that all types have the same interim allocation probability and interim transfer. This must also be true in any optimal mechanism. The reason is as follows. In any optimal mechanism, we must have \bar{q} be the step function given in (13). Moreover, the average downward deviation constraint (10) must hold as an equality for all types in order for bidder surplus to be at its lower bound. It must also be that (9) binds as well, and each bidder i is indifferent to downward deviations. Otherwise, if one of the individual constraints were slack, some other constraint must be violated, in order for the bidders to be indifferent on average. 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 misreporting any particular $s'_i \leq s_i$. For if there were a positive measure of types for which $u_i(s_i) > u_i(s_i, s'_i)$, and if there were indifference on average, then there would 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 were strictly positive, then since the value conditional on a signal of s'_i is strictly less than \bar{s} with probability 1, we would have

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

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

In addition, in any optimal mechanism, the bidders' indirect utility function is precisely that induced by the two-tier random reserve, which is

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

Now, the interim incentive constraint says that

$$\underline{u}_i(s_i) = \max_{s'_i} \int_{s_{-i}} [\max\{s_i, s_{-i}\} q_i(s'_i, s_{-i}) - t_i(s'_i, s_{-i})] dF_{-i}(s_{-i}).$$

The envelope condition then implies

$$\frac{d\underline{u}_i(s_i)}{ds_i} = \hat{q}_i(s_i),$$

where

$$\hat{q}_i(s_i) \triangleq \int_{s_{-i}} q_i(s_i, s_{-i}) \mathbb{I}_{s_i \geq s_{-i}} dF_{-i}(s_{-i})$$

is the interim probability that bidder i is allocated the good *and* has the highest signal. Putting these two together, we conclude that

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

Thus, any optimal allocation must share some crucial features with that induced by the two-tier random reserve mechanism: the interim probability of getting the good and the interim transfer must be independent of type, and the interim probability of getting the good and having the high signal must be the same in all optimal allocations. In a symmetric optimal mechanism, $\bar{q}_i = \bar{q}/N$, and these objects are pinned down even more tightly. However, there is still some flexibility in how the good is allocated among low-signal bidders. In some sense, the two-tier random reserve mechanism takes the simplest approach by treating all low-signal bidders symmetrically.

4.4 Alternative Implementations

While the two-tier random reserve mechanism nominally requires detailed knowledge of the environment in order to calibrate the distribution H , there exist other implementations that “discover” the optimal distribution in equilibrium. Consider the following mechanism, which we refer to as the *guaranteed demand auction* (GDA): Each bidder first decides whether to

pay an entry fee ϕ to enter the auction. Upon entering, the bidder then makes demand $\delta_i \in [0, \bar{\delta}]$ of a probability of receiving the good. The only parameters of the auction are the entry fee $\phi \geq 0$ and the upper bound $\bar{\delta} \in [0, 1/N]$. There are no payments beyond the entry fee. If bidder i decides not to enter, then the auction proceeds without him, and the payment and assignment probability of bidder i are both zero.

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 $\delta_{i^*} > 0$, then bidder i^* is allocated the good with probability δ_{i^*} and each bidder $j \neq i^*$ receives the good with probability $(1 - \delta_{i^*}) / (N - 1)$. Thus, a bidder is more likely to be allocated the good when he does not have the highest demand as

$$\delta_{i^*} < \frac{1 - \delta_{i^*}}{N - 1},$$

as long as $\delta_{i^*} \leq \bar{\delta} < 1/N$. In consequence, a bidder is guaranteed to receive the good with probability at least their demand.

We claim that there is an equilibrium in which each bidder simply demands a quantity that mimics the earlier random reserve distribution $H(x)$, see (7):

$$\delta(s_i) = \begin{cases} \frac{1}{N} \left(1 - \left(\frac{F(r)}{F(s_i)} \right)^N \right) & \text{if } s_i \geq r, \\ 0 & \text{if } s_i < r; \end{cases}$$

and where r solves

$$F^{N-1}(r) = 1 - N\bar{\delta}.$$

It is easily verified that the induced interim allocation is exactly the same as that induced by the random price mechanism that implements the exclusion threshold r . Since the interim transfers are constant as well, we conclude that conditional on entering, the proposed strategies are an equilibrium. Moreover, if $\bar{\delta}$ is chosen so that the exclusion threshold is the optimal r^* , and if ϕ is the highest entry fee such that all types are willing to enter, then the induced allocation and bidder surplus will be precisely those of the two-tier random reserve mechanism.

4.5 Omitted Bidders and Optimality of the Inclusive Posted Price

In Corollary 1 we gave a necessary and sufficient condition for the inclusive posted price to be the optimal mechanism. When this condition is not met, the seller can achieve greater revenue using the two-tier random reserve mechanism to withhold the good when the value

is low. A trade-off is that the exclusive two-tier random reserve construction is significantly more complicated than the inclusive posted price. A classic result of Bulow and Klemperer (1996) demonstrates that the value of such exclusion may be quite limited. They argue that the difference between optimal revenue and optimal must-sell revenue is bounded above by the additional revenue from the optimal must-sell mechanism when an additional bidder is added to the auction in a natural way. In the particular context of Bulow and Klemperer (1996), which excludes the maximum signal model, the optimal must-sell mechanism is an English auction. But we now argue that in the context of the maximum signal model of Bulow and Klemperer (2002) and this paper, this means the posted price mechanism is optimal.

Let us suppose that there are N potential bidders. As before, they receive independent signals drawn from F , and the common value of all bidders is the maximum of these signals. Only the first $N' \leq N$ of the bidders participate in the auction. We say that there are *omitted bidders* if $N' < N$. Following Bulow and Klemperer (1996), the expected value of a bidder $i \leq N'$ conditional on $(s_1, \dots, s_{N'})$ is the expectation of the maximum of all N signals, integrated across $(s_{N'+1}, \dots, s_N)$. If we let

$$w(x) \triangleq \int_{y=\underline{s}}^{\bar{s}} \max\{x, y\} d\left(F^{N'-N}(y)\right),$$

then the expected value conditional on $(s_1, \dots, s_{N'})$ is simply $w(\max_{i \leq N'} s_i)$.

Proposition 4 (Omitted Bidders).

If there are omitted bidders, then the inclusive posted price with price p^I as in (2) is an optimal mechanism.

Proof of 4. Suppose that there is an IC and IR mechanism that generates revenue R when only bidders $i \leq N' < N$ participate. Then there is an IC and IR must-sell mechanism with all N bidders in which the seller simply runs the same mechanism as with N' , and gives the good away for free to bidder $N' + 1$ whenever it would not have been allocated to a bidder $i \leq N'$. Clearly this must-sell mechanism generates revenue of R , which must be less than p^I , which is maximum revenue across all must-sell mechanisms. As a result, any achievable revenue with N' bidders must be less than p^I . But revenue of p^I can be obtained when there are only N' bidders by, say, making a take-it-or-leave-it offer to bidder $i = 1$ at price p^I , which would be accepted with probability one. We conclude that optimal revenue with N' bidders is p^I . \square

In particular, note that with omitted bidders, the optimal revenue is equal to p^I for all $N' < N$, and optimal must-sell revenue is equal to p^I for all $N' \leq N$. An alternate

proof of this result could be given using Corollary 1 and showing that $\psi(\underline{s}) \geq 0$. To do so, it is necessary to change the units of messages so that the expected value conditional on the highest of the first N' signals is exactly the highest of the first N' signals, i.e., a signal s_i must be relabeled $\eta(s_i)$. Using the change of variables formula, and the fact that $d\eta(x) = F^{N-N'}(x) dx$, we conclude that

$$\begin{aligned} \psi(\underline{s}) &= w(\underline{s}) - \int_{x=\underline{s}}^{\bar{s}} \frac{1 - F(x)}{F(x)} dw(x) \\ &= \int_{x=\underline{s}}^{\bar{s}} x d\left(F^{N-N'}(x)\right) - \int_{x=\underline{s}}^{\bar{s}} (1 - F(x)) F^{N-N'-1}(x) dx \\ &= \int_{x=\underline{s}}^{\bar{s}} x d\left(F^{N-N'-1}(x)\right), \end{aligned}$$

which is positive.

Thus, when there are omitted bidders, the seller does not benefit at all from exclusion, and posted prices are optimal. Bringing omitted bidders into the auction does not increase optimal revenue unless *all* potential bidders are included. We regard this as a further argument in favor of the inclusive posted price as a simple and robust mechanism for revenue extraction.

This finding may be contrasted with a more literal and naive interpretation of the result of Bulow and Klemperer (1996), which is that an English auction with $N' + 1$ bidders generates more revenue than the optimal auction with N' bidders. This result crucially relies on the hypothesis that bidders with higher signals have higher virtual values, which is violated in the maximum signal model. Indeed, as long as there are $N' \geq 2$ bidders, revenue from an English auction is actually *decreasing* in the number of bidders. The reason is that as long as $N' \geq 2$, competition between the bidders will make the participation constraints bind, so that revenue is equal to the expected highest virtual value among the first N' bidders. But when more bidders are included in the auction, it becomes more and more likely that the bidder who is allocated the good has the highest signal among all N potential bidders, which is the only case in which a bidder receives an information rent according to (4). This is consistent with the results of Bulow and Klemperer (2002), Section 7, that when high signal bidders have lower virtual values, excluding bidders in standard auctions will raise revenue.

5 General Common Values

We now broaden our analysis beyond the maximum signal model, and ask which of our results will generalize to environments with common values and a strong winner's curse. We

shall shortly define a class of such environments. It is always possible to achieve an efficient allocation with a posted price. We will construct mechanisms that implement uniform allocations across bidders while also withholding the good when the value is low. This can lead to strictly higher revenue. We finally give conditions under which it is possible to implement mechanisms that skew the allocation even further away from the high signal bidder and further increasing revenue. Note that we stop short of characterizing optimal auctions for these environments. As we indicated in the introduction, the pattern of binding incentive constraints at the optimal mechanism could in general be quite complicated, and will depend on the fine details of the information structure.

5.1 Increasing Information Rents

We continue to assume that bidders receive independent real signals, but bidder i 's signal is now drawn from an idiosyncratic distribution F_i . We continue assume that there is a common value whose expectation given the signals is given by the function $v(s_1, \dots, s_N)$ that is weakly increasing in each signal s_i . The virtual value of a bidder is still given by the general formula (3).

We say that the common value model displays *increasing information rents* if for all signal profiles s and for all i, j :

$$s_i > s_j \implies \frac{1 - F_i(s_i)}{f_i(s_i)} \frac{\partial v(s_i, s_{-i})}{\partial s_i} \geq \frac{1 - F_j(s_j)}{f_j(s_j)} \frac{\partial v(s_j, s_{-j})}{\partial s_j}. \quad (14)$$

Conversely, the common value model has *decreasing information rents* if for all signal profiles s and all i, j :

$$s_i > s_j \implies \frac{1 - F_i(s_i)}{f_i(s_i)} \frac{\partial v(s_i, s_{-i})}{\partial s_i} \leq \frac{1 - F_j(s_j)}{f_j(s_j)} \frac{\partial v(s_j, s_{-j})}{\partial s_j}.$$

The notion of increasing information rents compares information rents across bidder i and j but not across signals of any given bidder i . For example, it does not require that each bidder i has an increasing or decreasing virtual value in his own signal s_i .

In the maximum signal model, the increasing information rents condition is satisfied for any distribution function F as the term $\partial v(s_i, s_{-i}) / \partial s_i$ is positive only for the bidder with the maximum signal, and it is zero for all other bidders.

A prominent example of a common value model is the *wallet model* where the common value is the sum of the signals:

$$v(s_1, \dots, s_N) = \sum_{i=1}^N s_i.$$

This model was the focus of the analysis of Bulow and Klemperer (2002). Here, the marginal value of signal i is constant. The environment satisfies increasing information rents if and only if the inverse hazard rate is increasing, or equivalently if the hazard rate is decreasing. Thus, in the wallet game, whether the information rent is increasing or decreasing is entirely a matter of the monotonicity of the hazard rate. With the exponential distribution, the wallet model displays weakly increasing information rent. If the value function is given by the sum of nonlinear elements, for example

$$v(s_1, \dots, s_N) = \sum_{i=1}^N (s_i)^\alpha,$$

with $\alpha > 1$, then the wallet game with exponential signals displays strictly increasing information rents.

The increasing information rent condition implies that the revenue-maximizing allocation should be biased towards low-signal bidders. But the additional generality of the common value model complicates our earlier analysis in two respects. First, the common value of the object now depends on the entire profile of signals rather than just the highest signal s_i . This complicates our constructions in Section 3, as the bidders' payments will now have to depend on the entire signal profile, rather than just the highest of the others' signals. Second, the virtual value of the bidders with lower signals may now differ across bidders. Thus, while the optimal mechanism in the maximum signal model could be described just in terms of the allocation of the high-signal bidder and a representative low-signal bidder, the optimal mechanism in the general model might be significantly more complicated and must explicitly specify the allocations of all low-signal bidders.

While the exact characterization of revenue-maximizing mechanisms remains an open question, we can show two senses in which our results generalize to all increasing information rent environments. First, we show that allocating to bidders with lower signals must increase revenue, if it is incentive compatible to do so. Second, we show that we can generalize the constructions from Section 3 to be incentive compatible. Thus without establishing optimality, we describe simple revenue enhancing mechanisms building on the insights of the maximum signal model.

5.2 More Advantageous Selection

The sequence of mechanisms constructed in Section 3 lead to progressively higher revenue because they progressively skew the allocation away from high-signal bidders, who have high information rents, and towards low-signal bidders, who have low information rents. We now give a general formulation of this comparative static. Fix two allocations $q, q' : S^N \rightarrow [0, 1]$ such that

$$\sum_{i=1}^N q_i(s) = \sum_{i=1}^N q'_i(s). \quad (15)$$

In words, the allocations have the same total probability of allocating the good conditional on the signal profile s , and hence induce the same social surplus. We say that q has *more advantageous selection than* q' if for all s and x ,

$$\sum_{\{i|s_i \leq x\}} q_i(s) \geq \sum_{\{i|s_i \leq x\}} q'_i(s). \quad (16)$$

Thus, the more advantageously selective allocation q places more probability on low-signal bidders being allocated the good than does q' . It is more advantageously selective because, for every s_i , the interim expectation of the value conditional on receiving the good is higher under q than under q' :

$$\mathbb{E}[v(s) | s_i, q_i(s) > 0] \geq \mathbb{E}[v(s)' | s_i, q'_i(s) > 0]. \quad (17)$$

Our first formal result for this section shows that if information rents are increasing, then more advantageous selection increases revenue.

Theorem 4 (More Advantageous Selection).

Suppose that information rents are increasing and that q and q' are implementable allocations. If q has more advantageous selection than q' , then the revenue maximizing revenue is greater under q than under q' .

Crucially,

Proof of Theorem 4. Since the two allocations have the same total probability of allocating the good, for a given signal profile, they must induce the same social surplus. At the same time, by shifting the allocation to lower signal buyers, the bidders' information rents are reduced. Let

$$Z(x) = \sum_{\left\{i \mid \frac{1-F_i(s_i)}{f_i(s_i)} \frac{\partial v(s_i, s_{-i})}{\partial s_i} < x\right\}} q_i(s_i),$$

and define Z' analogously in terms of q' . Then increasing information rents implies that Z' first-order stochastically dominates Z , and hence

$$\begin{aligned} \sum_{i=1}^N q_i(s) \frac{1 - F_i(s_i)}{f_i(s_i)} \frac{\partial v(s_i, s_{-i})}{\partial s_i} &= \int_{x=-\infty}^{\infty} x dZ(x) \\ &\leq \int_{x=-\infty}^{\infty} x dZ'(x) \\ &= \sum_{i=1}^N q'_i(s) \frac{1 - F_i(s_i)}{f_i(s_i)} \frac{\partial v(s_i, s_{-i})}{\partial s_i}. \end{aligned}$$

Since total surplus is the same, and information rents are weakly lower with q , revenue must be weakly larger. \square

Similarly, if the information rents are decreasing, and if q has less advantageous selection than q' —in the sense that the reverse inequalities in 16 hold for all x —then maximum revenue across mechanisms that implement q is lower than maximum revenue across mechanisms that implement q' .

5.3 Revenue Improving Mechanisms

Thus, with increasing information rents, more advantageous selection increases revenue. The question remains how much advantageous selection can be achieved subject to incentive compatibility. While it is always possible to implement an allocation in which the high-signal bidder always receives the good, e.g., with standard auctions, there are generally non-trivial bounds on how much advantageous selection can be created, as in the maximum signal model. We do not have a general characterization of exactly how much advantageous selection can be attained. We can, however, describe some simple allocations that can always be implemented and significantly reduce adverse selection, and hence the winner's curse.

First, it is always possible to implement a range of *neutrally selective* allocations, in which the ex post allocation probability is the same for all bidders. In particular, the efficient neutrally selective allocation is always implementable via an inclusive posted price, as previously defined in (2).

Proposition 5 (Inclusive Posted Price).

The inclusive posted price mechanism yields a higher revenue than the standard auctions in every environment with increasing information rents.

Proof. The inclusive posted price and any standard auction assign the object with probability one at every type profile s . The inclusive posted price is an incentive compatible and neutrally

selective mechanism. It thus offers a more advantageous selection than any standard auction that always selects the bidder with the highest signal. The revenue ranking now follows directly from Theorem 4. \square

Similarly, it is always possible to implement an allocation that allocates the good if and only if the value exceeds a threshold r and, conditional on the signal profile, all bidders are equally likely to be allocated the good. Such an allocation necessarily leads to higher revenue. The ex post incentive compatible Vickrey price for agent i now depends on the entire signal profile s_{-i} of all the other agents. For a given screening level r for the common value that the seller wishes to select, we can define a personalized price for agent i as follows:

$$p_i(r, s_{-i}) \triangleq \max \{r, v(\underline{s}, s_{-i})\}. \quad (18)$$

The payment $p_i(r, s_{-i})$ represents the Vickrey payment of bidder i and thus can vary across bidders. The revelation game now asks each bidder for his signal s_i and allocates the object uniformly across the bidders if the reported signal profile s generates a value $v(s) \geq r$:

$$q_i(s) = \begin{cases} \frac{1}{N} & \text{if } v(s) \geq r, \\ 0 & \text{otherwise;} \end{cases} \quad (19)$$

and asks for the Vickrey payment $p_i(r, s_{-i})$ if the object is assigned to agent i .

Proposition 6 (Personalized Price).

The optimal personal price mechanism yields a (weakly) higher revenue than the inclusive posted price.

Proof. The personalized price mechanism is clearly ex-post incentive compatible for every r . It is neutrally selective for every r . For $r = v(\underline{s})$, the ex-ante expected payment of each bidder equals the inclusive posted price. Thus, the optimal personalized price mechanism must deliver a (weakly) higher revenue than the inclusive posted price. In particular, the optimal personalized price mechanism attains a strictly higher revenue if the average virtual utility at the lowest type profile is negative, and thus exclusion becomes strictly beneficial. \square

If the common value model is given by the maximum signal, then the optimal personalized price mechanism can be implemented by the two-tier price mechanism

In the general common value setting, though, excluding at a given value threshold may not be revenue maximizing, even if we restrict to neutrally selective allocations, and the optimal neutrally selective allocation could be quite complicated. There is a simple condition,

however, under which we can say what the optimal such allocation is. Let us say that the environment is *mean-regular* if the average virtual value

$$\pi(s) \triangleq \frac{1}{N} \sum_{i=1}^N \pi_i(s)$$

is monotonically increasing in the signal profile s . If the environment is mean-regular, then it is possible to implement the following allocation

$$q_i(s) = \begin{cases} \frac{1}{N}, & \text{if } \pi(s) \geq 0; \\ 0, & \text{otherwise.} \end{cases} \quad (20)$$

For under mean-regularity, the allocation defined by (20) is monotonic, so that it can be implemented by an analogous pricing rule to (18). In particular, a bidder who is allocated the good must pay

$$p_i(s_{-i}) = \min \{v(s'_i, s_{-i}) \mid \pi(s'_i, s_{-i}) \geq 0\}.$$

But just as in the maximum signal model, under increasing information rents, there is further scope to increase revenue, namely, with a generalization of the two-tier random reserve mechanism. As in Proposition 3 the good is withheld if the high type s_i is below the exclusion threshold r . If $s_i \geq r$, we draw a threshold type x for the highest type according to a distribution H , which has a density h . We shall shortly describe a class of such distributions can be implemented. The high type s_i is allocated the good if and only if $s_i \geq x$, and otherwise we randomly allocate the good to one of the low bidders.

The complication relative to the maximum signal model is to determine transfers such that this allocation is incentive compatible. They are now constructed on the basis of the Vickrey prices which depend on the entire profile s rather than the high signal s_i only.

Proposition 7 (Generalized Two-Tier Random Reserve).

For symmetric signal distributions $F_i = F$, there is a generalized two-tier random reserve mechanism that yields a higher revenue than any personalized price mechanism for the same exclusion level r .

Proof. Let us define

$$\begin{aligned} \hat{v}(x, y) &= \mathbb{E}[v(s_i, s_{-i}) \mid s_i = x, \max s_{-i} = y], \\ \tilde{v}(x, y) &= \mathbb{E}[v(s_i, s_{-i}) \mid s_i = x, \max s_{-i} \leq y]. \end{aligned}$$

We further denote by $\Gamma(x) = F^{N-1}(x)$ the distribution of the highest of the others' signals and γ its density. The prices will be set according to two different cases. First, if the high-bidder i is allocated the good, then the price is $\hat{v}(\max s_{-i}, \max s_{-i})$; second when $s_i > \max s_{-i} > x$, and the high-bidder pays

$$p(x) = \tilde{v}(x, x) - (\hat{v}(x, x) - \hat{v}(0, x)) \frac{1}{N-1} \frac{1-H(x)}{h(x)} \frac{\gamma(x)}{\Gamma(x)}$$

if $s_i \geq x > y$. Finally, if one of the low-signal bidders is allocated the good, they pay $\hat{v}(0, \max s)$.

The surplus from a report s'_i when the type is s_i is

$$\begin{aligned} & \int_{y=\underline{s}}^{s'_i} \int_{x=r}^{s'_i} (\hat{v}(s_i, y) - \mathbb{I}_{x>y} p(x) - \mathbb{I}_{y>x} \hat{v}(y, y)) h(x) dx \gamma(y) dy \\ & + \int_{y=s'_i}^{\bar{s}} (\hat{v}(s_i, y) - \hat{v}(0, y)) \frac{1-H(y)}{N-1} \gamma(y) dy. \end{aligned}$$

The derivative with respect to s'_i is

$$\begin{aligned} & (\tilde{v}(s_i, s'_i) - p(s'_i)) h(s'_i) \Gamma(s'_i) + (\hat{v}(s_i, s'_i) - \hat{v}(s'_i, s'_i)) H(s'_i) \gamma(s'_i) \\ & - (\hat{v}(s_i, s'_i) - \hat{v}(0, s'_i)) \frac{1-H(s'_i)}{N-1} \gamma(s'_i). \end{aligned}$$

Plugging in the formula for p , the derivative of the indirect utility reduces to

$$\begin{aligned} & (\tilde{v}(s_i, s'_i) - \tilde{v}(s'_i, s'_i)) h(s'_i) \Gamma(s'_i) \\ & - (\hat{v}(s_i, s'_i) - \hat{v}(s'_i, s'_i)) \left(\frac{1-NH(s'_i)}{N-1} \right) \gamma(s'_i). \end{aligned}$$

Thus, as long as H satisfies

$$\frac{h(x)}{1-NH(x)} \geq \frac{1}{N-1} \max_y \frac{\hat{v}(y, x) - \hat{v}(x, x)}{\tilde{v}(y, x) - \tilde{v}(x, x)} \frac{\gamma(x)}{\Gamma(x)}, \quad (21)$$

bidder surplus will be single-peaked at $s'_i = s_i$, and truthful reporting will be incentive compatible. If we assume that the right-hand side is bounded for all x , then there exist H functions that satisfy $H(0) = 0$ and asymptote to $H(\infty) \leq 1/N$, and satisfy the differential inequality. Such is the case for the maximum signal model, where the right-hand side reduces to $f(x)/F(x)$. For such an H , the proposed mechanism is incentive compatible.

Finally, as the generalized two-tier random reserve mechanism induces an allocation that is more advantageously selective than that induced by personalized prices, by Theorem 4, it generates more revenue. \square

With the above result, we have extended the revenue ranking result from the maximum signal model to all common value environments with increasing information rents.

6 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 his own value. 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 value is more sensitive to the private information of bidders with more optimistic beliefs. This seems like a natural feature in economic environments where the most optimistic bidder has the most information for determining the best-use value of the good, and therefore has a greater information rent. This case is not covered by the characterizations of optimal revenue that exist in the literature, which 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 lower 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 his 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 bidders' values.

Bulow and Klemperer (2002) argued that it may be difficult to tell whether information rents are increasing or decreasing, and that with interdependent values, the inclusive posted price may not be as naïve as auction theorists are tempted to assume. We agree with this conclusion and add the observation that we do not have to give up on using monopoly ex-

clusionary power. We can construct simple exclusive mechanisms which can be implemented in a wide range of environments and mitigate the loss in revenue due to the winner's curse.

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.

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.
- (2019): “Revenue Guarantee Equivalence,” *American Economic Review*, 109, 1911–1929.
- 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 (2019): “Robustly Optimal Auctions with Unknown Resale Opportunities,” *Review of Economic Studies*, forthcoming.
- GUPTA, M. AND B. LEBRUN (1999): “First Price Auctions with Resale,” *Economics Letters*, 64, 181–185.
- HAILE, P. (2003): “Auctions with Private Uncertainty and Resale Opportunities,” *Journal of Economic Theory*, 108, 72–110.
- 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.