Countering the winner's curse:
Optimal auction design in a common value model

DIRK BERGEMANN
Department of Economics, Yale University

BENJAMIN BROOKS
Department of Economics, University of Chicago

STEPHEN MORRIS
Department of Economics, Massachusetts Institute of Technology

We characterize revenue maximizing mechanisms in a common value environment where the value of the object is equal to the highest of the bidders' independent signals. If the revenue maximizing solution is to sell the object with probability 1, then an optimal mechanism is simply a posted price, namely, the highest price such that every type of every bidder is willing to buy the object. If the object is optimally sold with probability less than 1, then optimal mechanisms skew the allocation toward bidders with lower signals. 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. Our qualitative results extend to more general common value environments with a strong winner's curse.

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

JEL CLASSIFICATION. C72, D44, D82, D83.

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's value 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 generalize the revenue equivalence theorem of Myerson (1981) to models with interdependent values. They give a condition on the form of interdependence under which revenue is maximized by an auction that allocates the good to the bidder with the highest signal whenever the good is sold. We 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 Bayesian optimal auctions with interdependent values and independent signals appears to be quite limited.

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.[1] 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 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 et al. (2017a, 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 to use the good, e.g., possible resale opportunities if the bidders are intermediaries or possible designs to fulfill a procurement order, and the winner of the good will discover the best use ex post.[2]

---

[1]A 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, not interdependence per se.

[2]One could also assume that resale takes place between the bidders, the values 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 is used by Gupta and LeBrun (1999) and Haile (2003) to study asymmetric

The maximum signal model was first introduced 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 his 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 inclusive 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.[3]

Moreover, revenue is generally higher if the seller exercises monopoly power and rations the good when values are low. This is the case in the independent 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 highest-signal bidder would face less competition in a "tie break" if the others' signals are low, thus again inducing a winner's curse and depressing bidders' willingness to pay.

Our first 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; 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 expressed high interest. In equilibrium, bidders express high interest if and only if their signal exceeds a threshold, and prices are set as high as possible 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

---

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.

[3]In Bulow and Klemperer (2002), the main focus is on the additively separable model where the common value is equal to the sum of the individual signals. They conclude that an inclusive posted price is optimal among mechanisms that always allocate the good when the distribution of signals exhibits a decreasing hazard rate; see their Proposition 3. Campbell and Levin (2006) provide additional arguments for the revenue dominance of the posted price mechanism in related common value environments.

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 two-tier pricing, which is a *random two-tier posted price*: The highest-signal bidder is allocated the good only if his signal exceeds a random posted 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 at 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 posted price. The posted 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 random two-tier posted price allocates the good with probability 1, 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 1, 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 random two-tier price mechanism is optimal utilizes a novel argument. The standard approach to optimal mechanism design, based on the seminal contribution of Myerson (1981), relies on two key results: the revenue equivalence theorem, which says that revenue is the expected "virtual value" of the bidder who is allocated the good, and also the result that an allocation is implementable if and only if each bidder's interim allocation is weakly increasing in his signal. Myerson (1981) establishes these results for independent private values and, more generally, when utility functions are additively separable in the signals.[4] When the virtual value is monotonically increasing, dubbed the *regular case* in Myerson (1981), the optimal allocation subject to local incentive constraints is interim monotone and, hence, is implementable. Otherwise, an "ironing" construction can be used to solve for the optimal allocation subject to interim monotonicity. As we discuss below, the revenue equivalence theorem extends to our setting. It turns out, however, that the optimal allocation subject to only local incentive constraints is not implementable. In addition, we show by example that interim monotonicity of the allocation is neither necessary nor sufficient for implementability, and there is no analogue of ironing that can be used to identify the optimal allocation.

We, therefore, proceed differently. We identify a class of global deviations, one for each type of bidder, that pin down optimal revenue: Instead of reporting their true signal, a bidder misreports a lower signal that is drawn with probabilities that are proportional to the prior. We refer to this as *misreporting a redrawn lower signal*. This class

---

[4]Myerson also assumes that the gains from trade between the seller and a given bidder are assumed to depend only on that bidder's private type. In contrast, we study environments where the gains from trade depend on all signals.

of deviations is motivated by several considerations, which we discuss in detail in Section 4. We solve for optimal revenue subject to bidders not wanting to misreport redrawn lower signals. The solution is necessarily an upper bound on revenue when we respect all incentive constraints. Theorem 1 then establishes that the random two-tier price mechanism attains the upper bound and, hence, is an optimal mechanism.

Finally, we argue that our key qualitative results extend beyond the specific maximum signal model to a wide range of common value environments that exhibit *increasing information rents*. The condition of increasing information rent captures the idea that higher signals are 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. 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.

A significant feature of the direct mechanisms we construct is that they do not exclusively allocate the object to the bidder with the highest signal. Mechanisms with this property are optimal in the case of the maximum-signal model, and they improve over standard auctions in more general environments. This is in spite of the fact that bidders are ex ante symmetric. Similar features also emerge in the recent literature studying robust optimal auction design with common values, in particular when the seller takes a worst case over the bidders' information and the equilibrium. Brooks and Du (2019) define a class of "proportional auctions" and show that they solve the unrestricted max-min revenue maximization problem, and are Bayesian optimal on the endogenous worst-case informational environment. In related work, Du (2018) constructs robust mechanisms that asymptotically extract all of the surplus as the number of bidders go to infinity. The theme linking these papers and ours is that randomization in the allocation can limit the effect of a winner's curse in depressing bids. The continued study of auctions with randomized allocations is a promising direction for future research.

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 a winner's curse to two-tier price mechanisms that generate winner's blessing. Section 4 proves the optimality of the latter 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 $i \in \mathcal{N} \triangleq \{1, \ldots, N\}$. Each bidder i receives a real signal $s_i \in S \triangleq [\underline{s}, \overline{s}]$, 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 $F_{\mathcal{N}}(s) \triangleq \times_{i=1}^{N} F(s_i)$, $F_{-i}(s_{-i}) \triangleq \times_{j \neq i} F(s_j)$, and $F^k(x) \triangleq (F(x))^k$ for positive integers k.

These conventions are extended to the density in the obvious way. The bidders all assign the same value to the good, which is the maximum of the signals:

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

We frequently use the shorter expression $\max s$, which selects the maximal element from the vector $s = (s_1, \ldots, s_N)$. 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 $v(s)q - t$.

2.2 *Direct mechanisms*

We 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 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 (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

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

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

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

and $u_i(s_i) \triangleq u_i(s_i, s_i)$ is the payoff from reporting truthfully. Ex ante bidder surplus is

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

and total bidder surplus is

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

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

$$R \triangleq \sum_{i=1}^{N} \int_{s_i=\underline{s}}^{\overline{s}} t_i(s_i) f(s_i) \, ds_i.$$

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

$$\overline{q}(v) \triangleq \sum_{i=1}^{N} \overline{q}_i(v)$$

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

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

and revenue is obviously $R = \mathrm{TS} - U$.

3. Countering the winner's curse

We start by reviewing the winner's curse in standard auctions. We then progressively improve the revenue with a sequence of mechanisms to arrive at the optimal mechanism. Optimality is proven in the next section.

3.1 *Standard auctions and the winner's curse*

In the maximum signal model, 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 $[s_i, \overline{s}]$. Note that the interim expectation of bidder i satisfies

$$\mathbb{E}[v(s)|s_i] > s_i$$

for all $s_i < \overline{s}$. Even so, there is an equilibrium of the second-price auction in which each bidder i bids only $b_i^*(s_i) = s_i$. Expected revenue in this equilibrium is the expected second-highest signal. The equilibrium bid is equal to the *lowest* possible realization of the common value given the interim information s_i. But 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,

$$\mathbb{E}\big[v(s)|s_i, s_j \le s_i \, \forall j \ne i\big] = s_i,$$

as the expectation of the value conditional on s_i being the highest signal is simply s_i. The resulting equilibrium allocation is

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

Thus, the second-price auction exhibits a winner's curse: The bidder with the highest signal receives the good and 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 $\mathbb{E}[v(s_i, s_{-i})|s_i]$ down 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 et al. (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 arises across these standard auction formats.

3.2 *Inclusive posted prices and no 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 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:

$$p_I \triangleq \int_{x=\underline{s}}^{\overline{s}} x\, d\big(F^{N-1}(x)\big). \tag{1}$$

We refer to p_I as the *inclusive posted price*. This is 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, which is the largest price such that every type is willing to buy the object. At this price, all types are "included" in the allocation. Another interpretation is that p_I is the

ex ante expectation of a particular Vickrey price, which is the highest signal among the competing bidders,[5] $\max s_{-i}$.

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 et al. (2019) established a revenue equivalence result between standard auctions for the maximum signal model. This equivalence completes the proof of the above proposition. The revenue ranking can be understood 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 any of the standard auctions is equal to the expectation of the second highest of N draws. 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 draw.

The allocation induced by the inclusive posted price assigns the object with equal probability to every bidder i: $q_i(s) = 1/N$. As a result, the event of winning conveys no additional information about the value of the object to any of the winning bidders. In sharp contrast to the standard auctions, a bidder's expected value conditional on receiving the object is the same as his unconditional expectation, i.e., there is no winner's curse.

A different perspective on the revenue ranking result is offered by the revenue equivalence theorem. Specifically, using only 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

$$\pi_i(s) \triangleq v(s) - \frac{1 - F(s_i)}{f(s_i)} \frac{\partial v(s)}{\partial s_i}. \tag{2}$$

In the maximum signal model, this simplifies to

$$\pi_i(s) = \begin{cases} \max s & \text{if } s_i < \max s \\ \max s - \dfrac{1 - F(s_i)}{f(s_i)} & \text{if } s_i = \max s. \end{cases} \tag{3}$$

The partial derivative of the value with respect to the signal s_i is now simply the indicator $\mathbb{I}_{s_i = \max s}$ for whether bidder i has the highest signal. As a result, only the highest-signal

---

[5]The existence of the Vickrey prices also demonstrates that the inclusive posted price could alternatively be implemented by a set of personalized prices that would satisfy the stronger ex post individual rationality and ex post incentive constraints, as we argue in more detail in Section 3.4.

bidder receives an information rent, equal to the inverse hazard rate $(1 - F(s_i))/f(s_i)$. Thus, those bidders who do *not* have the highest signal have a higher virtual value—equal to the common value v—than the highest-signal bidder. In consequence, the lower is the probability that the highest-signal bidder receives the good, the higher is revenue. 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 highest-signal bidder with probability 1.[6]

3.3 *Two-tier posted price 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, the object is awarded even if the virtual value of some bidder or even the average virtual value across all bidders is negative. In these circumstances, not allocating the good can typically raise revenue.

A first attempt at raising revenue is to post a price, denoted by p_E, that is strictly higher than the inclusive posted price p_I; thus, $p_E > p_I$. By definition, the price p_E exceeds the interim expectation of the bidder with the lowest possible signal $s_i = \underline{s}$. Any such price p_E therefore induces some threshold $r \in (\underline{s}, \overline{s}]$, so that every bidder i with a signal $s_i \geq r$ accepts the offer and with signal $s_i < r$ rejects the offer. The resulting assignment probabilities are

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

We refer to a posted price $p_E > p_I$ as an *exclusive posted price*, and refer to the threshold r as the *exclusion level*. The posted price p_E, which implements the exclusion level r, is the expectation of the common value for the type $s_i = r$ conditional on receiving the good:

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

The exclusive posted price assigns the object with uniform probability among all those bidders whose signal s_i exceeds the threshold r. But in contrast to the inclusive posted price, the exclusive posted price again tilts the allocation toward higher-signal bidders, $s_i > r$, by excluding the lower-signal bidders, $s_i < r$. Thus, while the exclusive posted price does ration the object, it reintroduces the winner's curse. The formula of the virtual value, (3), suggests that revenue would be higher if we shift the allocation

---

[6]Note that the value function in the maximum signal model 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.

toward lower-signal bidders. For example, revenue is higher if we instead implement a uniform allocation conditional on assigning the object:

$$q_i(s) = \begin{cases} \dfrac{1}{N} & \text{if } \max s \geq r \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

This mechanism achieves the same exclusion level r and, hence, maintains the same ex post social surplus, but it generates less bidder surplus and more revenue.

One way to implement this allocation is with the following *two-tier posted price*: Let

$$p_L \triangleq r, \qquad p_H \triangleq \frac{\displaystyle\int_{x=r}^{\bar{s}} x\, d\big(F^{N-1}(x)\big)}{1 - F^{N-1}(r)}.$$

Thus, p_H is the expected value of a bidder with signal $s_i \leq r$ conditional on the highest signal among the remaining $N-1$ bidders weakly exceeding r and, thus, $p_H > p_L$. In this mechanism, every bidder is asked to express either a *low* or a *high* 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.

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. To see this, notice 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 results 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, it is a best reply 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 lower-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.

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* p_E *with the same exclusion level.*

The two-tier posted price induces a winner's blessing for types $s_i < r$: Being allocated the good implies that $\max s_{-i} \geq r$. Formally, the interim expectation of a bidder with

$s_i < r$ is smaller than the expected value conditional on receiving the object:

$$\mathbb{E}\big[v(s_i, s_{-i})|s_i\big] = \frac{\int_{x=\underline{s}}^{\overline{s}} \max\{s_i, x\}\,d\big(F^{N-1}(x)\big)}{\int_{x=\underline{s}}^{\overline{s}} d\big(F^{N-1}(x)\big)} < \frac{\int_{x=r}^{\overline{s}} x\,d\big(F^{N-1}(x)\big)}{\int_{x=r}^{\overline{s}} d\big(F^{N-1}(x)\big)}.$$

By contrast, for all types $s_i \geq r$, the expected value conditional on receiving the good equals the interim expectation, as with the inclusive posted price. Thus, for these types, there is neither a winner's curse nor a winner's blessing.[7]

3.4 *Random two-tier posted price*

The inclusive posted price and the two-tier posted price depresses the probability of winning to $1/N$ for the bidder with the highest signal, whereas that bidder would have won with probability 1 in any standard auction. Do there exist mechanisms that reduce the highest-signal bidder's probability of winning even further? One 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 highest-signal bidder, as we now explain.

As a segue, let us first observe that there is another implementation of the allocation (4) 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 $\max\{r, s_{-i}\}$. Thus, the high bidder faces a nontrivial posted price r, whereas the lower-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., truthful reporting is an equilibrium 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. Breaking ties randomly, 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 posted 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

---

[7]In the context of efficient trade with interdependent values, Fieseler et al. (2001) refer to a "blessing" when the value of the winning bidder is a decreasing function of the other agents' signal (rather than increasing as in common value models of auctions). In related work on efficient bargaining, Segal and Whinston (2011) refer to this condition as a "winner's blessing" in contrast to the winner's curse. In the current analysis, the presence of a winner's curse or winner's blessing is not determined by an exogenous statistical feature of the value model, but by the endogenous choice of the allocation through the revenue maximizing mechanism.

is offered the good at a price equal to the highest reported signal. The allocation just
described is

$$
q_i(s) = \begin{cases}
\dfrac{1}{|\arg\max s|}H(\max s) + \left(1 - \dfrac{1}{|\arg\max s|}\right)\dfrac{1}{N-1}(1 - H(\max s)) \\
\quad \text{if } s_i = \max s \text{ and } \max s \geq r \\[2mm]
\dfrac{1}{N-1}(1 - H(\max s)) \\
\quad \text{if } s_i < \max s \text{ and } \max s \geq r \\[2mm]
0 \quad \text{if } \max s < r.
\end{cases}
$$

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 shown that truthful bidding is an equilibrium. However, if we
choose H to put less probability on $x = r$, then the allocation is effectively skewed to-
ward lower-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, bidders have no incentive to overreport when others report
truthfully: 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,
to verify 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:

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

The derivative of this expression with respect to s_i' is

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

So, a sufficient condition for downward deviations not to be attractive is that the term
inside the brackets is nonnegative for all s_i', which reduces to

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

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

$$
H(x) \triangleq \frac{1}{N}\left(1 - \left(\frac{F(r)}{F(x)}\right)^N\right). \tag{5}
$$

With this distribution for x, we refer to the resulting game form as a *random two-tier posted* price. We have just verified that this mechanism is incentive compatible. In fact, bidders are indifferent between truthful reporting and all downward misreports.

By construction, $H(r) = 0$, so 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

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

We have, therefore, completed the proof of the following result.

PROPOSITION 3 (Random two-tier posted price). *The random two-tier posted price yields a higher revenue than the two-tier posted price for a given exclusion level r.*

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

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

We return to this property and its further implications in Section 4.3.

Until now, we have treated the exclusion threshold r as a fixed parameter. Thus, there is actually a one-dimensional family of random two-tier posted price 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 r^* can be determined 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 random two-tier posted price is

$$U = \int_{s=r}^{\bar{s}} \int_{x=r}^{s} (s - x) d\left(H(x)F^{N-1}(x)\right) dF(s)$$

$$= \int_{x=r}^{\bar{s}} \frac{1}{N} \frac{1 - F(x)}{F(x)} \left(F^N(x) - F^N(r)\right) dx,$$

where we have simply plugged in the definition of $H(x)$ from (5). Thus, the effect of an increase in r on U is

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

The overall effect of increasing r on revenue is, therefore,

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

where

$$\psi(r) \triangleq r - \int_{x=r}^{\overline{s}} \frac{1 - F(x)}{F(x)} \, dx. \tag{6}$$

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 posted price, and the optimal reserve price r^* is the smallest r such that $\psi(r) \geq 0$:

$$r^* \triangleq \min\left\{r \middle| r - \int_{x=r}^{\overline{s}} \frac{1 - F(x)}{F(x)} \, dx \geq 0\right\}.$$

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. We expand on this interpretation in the next section.

Thus, we have established that the random two-tier posted price 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 maximizes revenue among all incentive compatible and individually rational mechanisms.

THEOREM 1 (Optimality of random two-tier posted price). *The random two-tier posted price with cutoff r^* maximizes revenue across all IC and IR direct mechanisms.*

The random two-tier posted price mechanism tilts the allocation toward lower-signal bidders and away from the highest-signal bidder. In consequence, every type of bidder is more likely to receive the object when their signal is not the highest, and this mechanism extends the winner's blessing to all types. Formally, the interim expectation for all bidders $s_i \in [\underline{s}, \overline{s})$ is strictly smaller than the expected value conditional on receiving the object,

$$\mathbb{E}\big[v(s_i, s_{-i})|s_i\big] = \int_{x=\underline{s}}^{\overline{s}} \max\{s_i, x\} \, d\big(F^{N-1}(x)\big)$$

$$< s_i F^{N-1}(s_i) H(s_i) + \int_{x=s_i}^{\overline{s}} x\left(\frac{1 - H(x)}{N - 1}\right) d\big(F^{N-1}(x)\big),$$

where the inequality can be verified after replacing the distribution $H(x)$ by the prior distribution of signals, $F(x)$, as defined above in (5). Thus, the optimal mechanism reverses the winner's curse and induces a winner's blessing for all types (except the highest type).

Note that as r goes to $\underline{s}$, H converges pointwise to the constant function $1/N$. As a result, when the gains from the bias toward lower-signal bidders is small relative to the cost of restricting supply, the inclusive posted price 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 also show that the inclusive posted is *always the* optimal mechanism if one restricts attention to mechanisms where the object must always 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* 1*, then the inclusive posted price maximizes expected revenue across all IC and IR mechanisms.*

We prove Theorems 1 and 2 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).[8]

4. OPTIMAL MECHANISMS

The broad strategy in proving Theorems 1 and 2 is to show that the allocations induced by random two-tier posted prices 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 to rewrite expected revenue as the expected virtual value of the bidder who is allocated the good. Note that the formula for the virtual value (3) 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 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, namely that of Milgrom and Weber (1982), to show that such an allocation is implementable. This proof strategy does not work in the maximum signal model. As we have shown, 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 highest-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 0 and the lower types with probability 1. The high type must,

---

[8]As we mentioned earlier in footnote 3, 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. The maximum signal model is not additively separable and, therefore, necessitates new arguments.

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.

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 nondecreasing. 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.[9]

Thus, we must incorporate 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 towards 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.[10] 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.

In the next section, we formally establish an upper bound on revenue that is attained by the random two-tier posted price described in Theorem 1. The bound is obtained by solving a relaxation of the revenue maximization program, where we drop all of the incentive constraints except those that correspond to a one-dimensional family of deviations in the normal form, one for each type s_i, wherein s_i misreports as a type $s_i' \in [\underline{s}, s_i]$

---

[9]The 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			q_1		s_2	
		0	1				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 that is at least that of the high type so as 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 his signal. As a result, the example can be made symmetric simply by randomly permuting the roles of the bidders.

[10]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.

with likelihood $\tilde{f}(s_i'|s_i)$. We now informally motivate this relaxation. First, the fact that the deviation likelihood is independent of i is suggested by the symmetry of the model. Next, a dual interpretation of these constraints is that $\tilde{f}(s_i'|s_i)$ is proportional to the optimal Lagrange multiplier on the incentive constraint for s_i misreporting as s_i' (so that replacing incentive constraints with weighted sums, where the weights are equal to the optimal Lagrange multipliers, does not change the value of the program). The fact that types misreport only as lower types corresponds to our observation in the previous paragraph that bidders need to be indifferent to all downward deviations.

The functional form for $\tilde{f}$ can be understood as follows. Ultimately we argue that the random two-tier posted price is an optimal mechanism. Recall that this mechanism has the feature that $u_i(s_i) = u_i(s_i, s_i')$ for all $s_i' \leq s_i$. Thus, a feature of this mechanism is that all downward deviations result in the same payoff, so that $u_i(s_i, s_i') - u_i(s_i, s_i'') = 0$ as long as $s_i > \max\{s_i', s_i''\}$. Thus, the relative value of misreporting as lower types does not depend on which type is misreporting. This suggests a crucial property of the deviation likelihoods: The relative likelihood of deviating to lower types does not depend on which type is deviating, i.e., $\tilde{f}(s_i'|s_i)/\tilde{f}(s_i''|s_i)$ is independent of s_i (again, as long as $s_i > \max\{s_i', s_i''\}$). Equivalently, there exists a function $\tilde{f}(s_i')$ (which no longer depends on s_i) such that $\tilde{f}(s_i'|s_i) = \tilde{f}(s_i')/\tilde{F}(s_i)$, where $\tilde{F}(s_i) \triangleq \int_{x=\underline{s}}^{s_i} \tilde{f}(x)\, dx$.

Now, given this functional form, the aggregated incentive constraint is

$$
\begin{aligned}
u_i(s_i)\tilde{F}(s_i) &\geq \int_{s_i'=\underline{s}}^{s_i} \tilde{f}(s_i') u_i(s_i, s_i')\, ds_i' \\
&= \int_{s_i'=\underline{s}}^{s_i} \tilde{f}(s_i') \bigg[u_i(s_i') \\
&\quad + \int_{s_{-i}\in S^{N-1}} \max\{0, s_i - \max\{s_i', s_{-i}\}\} q_i(s_i', s_{-i}) f_{-i}(s_{-i})\, ds_{-i} \bigg] ds_i'.
\end{aligned}
$$

This integral inequality ultimately is used to bound bidder utility from below. It is intuitive that given the allocation q_i, there is a smallest indirect utility that satisfies this inequality for all s_i, which is the function $\underline{u}_i$ that makes the constraint hold as equality everywhere, with $\underline{u}_i(\underline{s}) = 0$. (This is established rigorously for the optimal $\tilde{f}$ in Theorem 3 below.) Treating the aggregated incentive constraint as an equality and differentiating both sides, we obtain the differential equation

$$
u_i'(s_i)\tilde{F}(s_i) = \int_{s'\in[\underline{s}, s_i]^N} q_i(s')\tilde{f}(s_i') f_{-i}(s_{-i}')\, ds'.
$$

Subject to the boundary condition, the solution is

$$
\underline{u}(s_i) \triangleq \int_{s_i''=\underline{s}}^{s_i} \frac{1}{\tilde{F}(s_i'')} \int_{s'\in[\underline{s}, s_i'']^N} q_i(s')\tilde{f}(s_i') f_{-i}(s_{-i}')\, ds'\, ds_i''.
$$

Bidder i's ex ante expected utility U_i is then at least

$$
\int_{s_i=\underline{s}}^{\bar{s}} \underline{u}(s_i) f(s_i)\, ds_i = \int_{s_i=\underline{s}}^{\bar{s}} \int_{s_i''=\underline{s}}^{s_i} \frac{1}{\tilde{F}(s_i'')} \int_{s'\in[\underline{s}, s_i'']^N} q_i(s')\tilde{f}(s_i') f_{-i}(s_{-i}')\, ds'\, ds_i''\, f_{\mathcal{N}}(s_i)\, ds_i
$$

$$= \int_{s' \in S^N} \int_{s_i = \underline{s}}^{\overline{s}} \int_{s_i'' = \underline{s}}^{s_i} \frac{1}{\tilde{F}(s_i'')} \mathbb{I}_{s_i'' \geq \max(s')} \, ds_i'' f(s_i) \, ds_i q_i(s') \frac{\tilde{f}(s_i')}{f(s_i')} f_{\mathcal{N}}(s') \, ds'$$

$$= \int_{s' \in S^N} \int_{s_i'' = \max(s')}^{\overline{s}} \frac{1 - F(s_i'')}{\tilde{F}(s_i'')} \, ds_i'' \frac{\tilde{f}(s_i')}{f(s_i')} q_i(s') f_{\mathcal{N}}(s') \, ds'.$$

As revenue is simply expected surplus less the bidders rents, the previous equation (with a change of variables) indicates that

$$\sum_{i=1}^{N} \left(\int_{s \in S^N} \max(s) q_i(s) f_{\mathcal{N}}(s) \, ds - U_i \right)$$

$$\leq \sum_{i=1}^{N} \int_{s \in S^N} \left(\max(s) - \int_{s_i' = \max(s)}^{\overline{s}} \frac{1 - F(s_i')}{\tilde{F}(s_i')} \, ds_i' \frac{\tilde{f}(s_i)}{f(s_i)} \right) q_i(s) f_{\mathcal{N}}(s) \, ds. \qquad (7)$$

This equation represents an analogue of the Myersonian revenue equivalence formula, except that the term in parentheses is a generalized virtual value that takes into account this particular class of global incentive constraints.

Now we may ask, which $\tilde{f}$ is optimal, i.e., minimizes the bound? Without global incentive constraints, the seller will wish to allocate the good to the lower-signal bidders. This suggests that the optimal $\tilde{f}$ should be such that the seller no longer wishes to exclusively allocate to lower-signal buyers. In fact, there is a particular choice for $\tilde{f}$ such that the term multiplying $q_i(s)$ in (7) is independent of i, which is $\tilde{f}(s_i) = f(s_i)$, i.e., deviations are proportional to the prior. As a result, the seller will be *indifferent between allocating to all bidders*. The resulting revenue bound is

$$\sum_{i=1}^{N} \int_{s \in S^N} \left(\max(s) - \int_{s_i' = \max(s)}^{\overline{s}} \frac{1 - F(s_i')}{F(s_i')} \, ds_i' \right) q_i(s) f_{\mathcal{N}}(s) \, ds$$

$$= \sum_{i=1}^{N} \int_{s \in S^N} \psi(\max s) q_i(s) f_{\mathcal{N}}(s) \, ds$$

$$= \int_{v = \underline{s}}^{\overline{s}} \psi(v) \overline{q}(v) g(v) \, dv.$$

We have already argued that this expression is maximized by setting $\overline{q} = 1$ if $\max s \geq r^*$ and $q_i = 0$ otherwise, and, moreover, this revenue upper bound is attained by the optimal random two-tier posted price.

4.2 *Proof of optimality: Theorems 1 and 2*

We now formally state and prove the upper bound on revenue.

THEOREM 3 (Revenue upper bound). *In any IC and IR mechanism in which the probability of allocating the good conditional on the value is given by $\overline{q}$, expected revenue is*

bounded above by

$$\overline{R} \triangleq \int_{v=\underline{s}}^{\overline{s}} \psi(v)\overline{q}(v)g(v)dv, \tag{8}$$

where $\psi(v)$ is defined in (6).

PROOF. Consider the deviation in which type s_i misreports as $s_i' \in [\underline{s}, s_i]$ according to the cumulative distribution $F(s_i')/F(s_i)$. We refer to this as *misreporting a redrawn lower signal.* The analysis of the preceding section shows that a bidder with signal s_i will not want to deviate in this manner only if

$$u_i(s_i) \geq \frac{1}{F(s_i)} \int_{s' \in [\underline{s}, s_i]^N} [u_i(s_i') + \max\{0, s_i - \max\{s'\}\}q_i(s')]f_{\mathcal{N}}(s')\,ds'$$

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

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

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

where

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

and

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

Now define the function operator

$$\Gamma(u)(y) \triangleq \begin{cases} \dfrac{1}{F(y)} \left(\displaystyle\int_{x=\underline{s}}^{y} u(x)f(x)\,dx + \lambda(y) \right) & \text{if } y > \underline{s} \\ 0 & \text{if } y = \underline{s} \end{cases}$$

on the space of nonnegative integrable utility functions such that $u(\underline{s}) = 0$. Then (9) is equivalent to $u \geq \Gamma(u)$ in the pointwise order.

It is easily verified that the function

$$\underline{u}(y) \triangleq \begin{cases} \displaystyle\int_{x=\underline{s}}^{y} \lambda(x)\frac{f(x)}{(F(x))^2}\,dx + \frac{\lambda(y)}{F(y)} & \text{if } y > \underline{s} \\ 0 & \text{if } y = \underline{s} \end{cases} \tag{10}$$

is a fixed point of Γ and, therefore, solves (9) as an equality:

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

$$= \frac{1}{F(y)}\left(F(y)\int_{x=\underline{s}}^{y} \lambda(x)\frac{f(x)}{\left(F(x)\right)^2}\,dx + \lambda(y)\right)$$

$$= \underline{u}(y),$$

where the second line comes from Fubini's theorem.

We now argue that $\underline{u}$ is the lowest nonnegative u that satisfies (9). This is a consequence of the following observations. First, Γ is a monotonic operator on nonnegative increasing functions in the pointwise order, 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

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

$$\leq \frac{1}{F(s)}\int_{x=\underline{s}}^{s}\left|\underline{u}(x) - \widehat{u}(x)\right|f(x)\,dx$$

$$< \frac{1}{F(s)}\int_{x=\underline{s}}^{s}\|\underline{u} - \widehat{u}\|f(x)\,dx$$

$$= \|\underline{u} - \widehat{u}\|.$$

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

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

As a result, if a direct mechanism implements $\overline{q}$, total bidder surplus must be at least

$$\underline{U} \triangleq \int_{y=\underline{s}}^{\overline{s}} \underline{u}(y)f(y)\,ds = \int_{y=\underline{s}}^{\overline{s}}\int_{x=y}^{\overline{s}} \frac{1 - F(x)}{F(x)}\,dx\,\overline{q}(y)\,d\left(F^N(y)\right),$$

and revenue is $\mathrm{TS} - \underline{U}$, which is equal to $\overline{R}$.   □

We can now complete the proofs of our main theorems.

PROOF OF THEOREM 1. If the seller can withhold the good, then we can derive an upper bound on optimal revenue by maximizing the bound (8) pointwise. Since $\psi(v)$ is

monotonic, the pointwise maximum is attained by the allocation

$$\bar{q}(v) = \begin{cases} 1 & \text{if } v \geq r^* \\ 0 & \text{if } v < r^*, \end{cases} \tag{11}$$

where

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

This is the allocation that is implemented by the random two-tier posted price. Moreover, we have already verified that bidders are indifferent to all downward deviations, so that the revenue upper bound is attained.   □

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 is attained by any mechanism that implements this allocation and makes bidders indifferent to all downward deviations. But all types are treated the same way by the inclusive posted price, so that all downward incentive constraints hold as equalities and the upper bound on revenue is attained.   □

4.3 *Features of the optimal auction*

We now review some key properties of optimal mechanisms, including the structure of optimal allocations and transfers, the pattern of binding incentive constraints, and the nature of updating from being allocated the good.

Features of optimal direct mechanisms   Theorem 1 establishes that a random two-tier posted price maximizes revenue. A natural question is whether there are other optimal allocations. While there may be other optimal auctions, the induced allocation must share a number of key properties with the one we have constructed.

First, the logic of Theorem 3 implies that in any optimal mechanism, we must have $\overline{q}$ be the step function given in (11).

Second, so as to attain the revenue bound from Theorem 3 exactly, all of the average downward constraints (9) must hold as equalities. This implies that all *pointwise* downward incentive constraints hold as equalities as well, i.e., for all s_i and $s_i' < s_i$, $u_i(s_i) = u_i(s_i, s_i')$. Otherwise, if some subset of these constraints were slack, then since they are equalities on average, some other subset must be strictly violated. Now, fix two types s_i and s_i'. If $\hat{s} \geq \max\{s_i, s_i'\}$, then because bidders are indifferent to all downward deviations,

$$0 = u_i(\hat{s}, s_i) - u_i(\hat{s}, s_i')$$

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

where $t_i(\cdot)$ is the interim transfer function. Thus, the derivative of the last line with respect to $\hat{s}$ must be zero:

$$0 = \int_{s_{-i} \in [\underline{s}, \hat{s}]^{N-1}} \big(q_i(s_i, s_{-i}) - q_i(s_i', s_{-i})\big) f_{-i}(s_{-i})\, ds_{-i}.$$

Evaluated at $\hat{s} = \bar{s}$, this implies that $q_i(s_i) = q_i(s_i')$, where $q_i(\cdot)$ is the interim allocation. Thus, in any optimal allocation, all types receive exactly the same interim allocation.

The fact that interim allocations are constant implies that 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 he make his reports, and the optimal allocation is implemented. We next describe such an implementation in greater detail.

Finally, the argument in Theorem 3 actually establishes a much stronger result, that the sum of the indirect utilities $u_i(y)$ is at least $\underline{u}(s_i)$, and this level must be attained in any optimal mechanism. Recall that another way to compute utilities is with the envelope formula, which says that $u_i'(s_i) = \hat{q}_i(s_i)$, where the right-hand side is the interim probability that bidder i is allocated the good *and* has the highest signal:

$$\hat{q}_i(s_i) \triangleq \int_{s_{-i} \in S^{N-1}} q_i(s_i, s_{-i}) \mathbb{I}_{s_i \geq s_{-i}} \, dF_{-i}(s_{-i}).$$

Putting these two observations together, we conclude that

$$\sum_{i=1}^{N} \widehat{q}_i(s_i) = \underline{u}'(s_i) = \frac{1}{F(s_i)} \int_{x=\underline{s}}^{s_i} \overline{q}_i(x) g(x) \, dx,$$

where the term on the right-hand side is a direct calculation from (10).

Thus, any optimal allocation must have the same aggregate allocation, the same constant interim allocation, and the same interim probability of getting the good and having the high signal. There is still some flexibility in how the good is allocated among lower-signal bidders. The random two-tier posted price simply treats all lower-signal bidders symmetrically.

Global incentive constraints   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. 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 truthful reporting 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 4 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 support of the

FIGURE 1. Uniform downward versus upward incentive constraints.

signal distribution. 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.

We emphasize that this holds true for all continuous value (and signal) distributions, not just the uniform distribution plotted here.

Interim and ex post allocation probability   While the interim allocation probability is constant, the ex post probability $q_i(s)$ induced by the random two-tier posted price mechanism depends on the realized signals:

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

Thus, conditional on the realized signal profile, highest-signal bidders are strictly *less* likely to receive the good than lower-signal bidders, since

$$H(\max s) \leq H(\overline{s}) = \frac{1 - F^N(r)}{N} < \frac{1}{N} < \frac{1}{N} + \frac{1}{N^2} F^N(r) = \frac{1 - H(\overline{s})}{N-1} \leq \frac{1 - H(\max s)}{N-1}.$$

Conditioning on winning therefore results in a higher expected value for all types. In effect, the random posted 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 while the allocation is interim incentive compatible, it cannot be made ex post incentive compatible. The reason is that q_i is not monotonic in s_i, which is necessary for ex post implementation: If a bidder $j \neq i$ reports a signal above r, then bidder i receives the good with higher probability with a report less than $\max s_{-i}$ than he does with a report greater than $\max s_{-i}$. The allocations associated with the inclusive posted price and the two-tier posted price can, however, be implemented ex post.

4.4 *Alternative implementation: Guaranteed demand auction*

While the random two-tier posted price nominally requires detailed knowledge of the environment so as 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 $\delta_i \in [0, \overline{\delta}]$ of a probability of receiving the good. The only parameters of the auction are the entry fee $\phi \geq 0$ and the upper bound $\overline{\delta} \in [0, 1/N]$. 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 0.

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

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

because $\delta_{i*} \leq \overline{\delta} < 1/N$. In consequence, a bidder is guaranteed to receive the good with probability at least his demand.

We claim that there is an equilibrium in which each bidder simply demands a quantity

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

and where r solves

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

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 $\overline{\delta}$ 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 random two-tier posted price.

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 random two-tier posted price to withhold the good when the value is low. A trade-off is that the exclusive random two-tier posted price 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, whereas in the maximum signal model, the posted price mechanism is optimal among must-sell mechanisms. But as we now argue, when there are omitted bidders, the inclusive posted price is in fact optimal among *all* mechanisms, including those that ration the good.[11]

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 $(s_1, \ldots, s_{N'})$ is the expectation of the maximum

---

[11]Kirkegaard (2006) also emphasizes the broader applicability of the Bulow and Klemperer (1996) limited value of exclusion result.

of all N signals, integrated across $(s_{N'+1}, \ldots, s_N)$. If we let

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

then the expected value conditional on $(s_1, \ldots, s_{N'})$ 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 (1) is an optimal mechanism.*

PROOF. 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 1. We conclude that optimal revenue with N' bidders is p^I. □

In particular, note that with omitted bidders, optimal revenue is equal to p^I for all $N' < N$ and the optimal mechanism always allocates the good. 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 the narrow 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 (3). This is consistent with the results of Bulow and Klemperer (2002), Section 7, that when highest-signal bidders have lower virtual values, excluding bidders in standard auctions will raise revenue.[12]

---

[12]In the current analysis, as in Bulow and Klemperer (1996), the entry of new bidders is taken as an exogenous event. An anonymous referee suggested that with endogenous entry of new bidders, the optimal

5. Common values and increasing information rents

We now broaden our analysis beyond the maximum signal model and ask which of our results generalizes to environments with common values and a strong winner's curse. We shortly define a class of such environments. It is always possible to achieve an efficient allocation with a posted price. We 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 highest-signal bidder and further increase 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 depends on the fine details of the information structure.

5.1 *Increasing information rents*

We continue to assume that bidders receive independently distributed signals, but bidder i's signal can now be drawn from an idiosyncratic distribution F_i. We also continue to assume that there is a common value whose expectation, given the signals $(s_1, \ldots, s_N)$, is given by a function $v(s_1, \ldots, s_N)$ that is weakly increasing in each signal s_i. The virtual value of a bidder is still given by the general formula (2).

We say that the model has *increasing information rents* if for all signal profiles s and for all i, j,

$$s_i > s_j \quad \Longrightarrow \quad \frac{1 - F_i(s_i)}{f_i(s_i)} \frac{\partial v(s)}{\partial s_i} \geq \frac{1 - F_j(s_j)}{f_j(s_j)} \frac{\partial v(s)}{\partial s_j}.$$

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

$$s_i > s_j \quad \Longrightarrow \quad \frac{1 - F_i(s_i)}{f_i(s_i)} \frac{\partial v(s)}{\partial s_i} \leq \frac{1 - F_j(s_j)}{f_j(s_j)} \frac{\partial v(s)}{\partial s_j}.$$

Note that the notion of increasing information rents compares information rents across bidders i and j, but not across signals s_i and s_i' 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 $\partial v(s)/\partial s_i$ is positive only for the bidder with the maximum signal, and it is zero for all other bidders.

Another prominent example is the *wallet model*, where the common value is the sum of the signals:

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

This model was the focus of the analysis in 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 has weakly increasing information rent. If the value function is given by the sum of nonlinear elements, for example,

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

with $\alpha > 1$, then the wallet game with exponential signals has strictly increasing information rents.[13]

The increasing information rent condition implies that the revenue maximizing allocation should be biased toward lower-signal bidders. But the 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 s of signals rather than just the highest signal $\max s$. This complicates our constructions in Section 3, as the bidders' payments now have to depend on the entire signal profile, rather than just the highest of the others' signals. Second, the virtual values of the bidders with lower signals may now differ across bidders. Thus, whereas the optimal mechanism in the maximum signal model could be described just in terms of the allocation of the highest-signal bidder and a representative lower-signal bidder, the optimal allocation in the general model might differentiate between lower-signal bidders.

While the exact characterization of revenue maximizing mechanisms remains an open question, we generalize our results in two steps 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 revenue improving constructions from Section 3. Thus, without establishing optimality, we describe simple revenue enhancing mechanisms building on the insights of the maximum signal model.

---

[13]Povel and Singh (2004, 2006) consider a variant of the wallet game with two bidders. While the distribution of the signals is assumed to be independent and identical across bidders, the signal of each bidder enters the common value function with different weights. Thus, the agent with the larger weight can be said to have more private information. They show that the asymmetry across bidders leads to asymmetric allocation rules that ameliorate the winner's curse. The analysis is driven by the ex ante asymmetry in the scale of the bidders' information, but otherwise follows the local arguments of Myerson (1981) and Bulow and Klemperer (1996).

5.2 *More advantageous selection*

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

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

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,

$$\sum_{\{i|s_i \leq x\}} q_i(s) \geq \sum_{\{i|s_i \leq x\}} q'_i(s). \tag{12}$$

Thus, the more advantageously selective allocation q places more probability on lower-signal bidders being allocated the good than does q'. Given a signal profile s, it results in a lower unconditional expected value for the bidder who is allocated the good and, correspondingly, more positive updating about the value from the event of winning the auction. Generally speaking, this results in a weaker winner's curse.

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 maximal revenue across all IC and IR direct mechanisms that implement q is greater than maximal revenue across all IC and IR direct mechanisms that implement q'.*

PROOF. Since the two allocations have the same total probability of allocating the good for each 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

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

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

$$\sum_{i=1}^{N} q_i(s) \frac{1 - F_i(s_i)}{f_i(s_i)} \frac{\partial v(s)}{\partial s_i} = \int_{x=-\infty}^{\infty} x \, dZ(x)$$

$$\leq \int_{x=-\infty}^{\infty} x \, dZ'(x)$$

$$= \sum_{i=1}^{N} q'_i(s) \frac{1 - F_i(s_i)}{f_i(s_i)} \frac{\partial v(s)}{\partial s_i}.$$

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

Similarly, if information rents are decreasing and if q has less advantageous selection than q'—in the sense that the reverse inequalities in (12) hold for all x—then the 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 as to how much advantageous selection can be achieved subject to incentive compatibility constraints. While it is always possible to implement an allocation in which the highest-signal bidder always receives the good, e.g., with standard auctions, there are generally nontrivial bounds on how much the allocation can be skewed toward lower-signal bidders, as in the maximum signal model. We do not have a general characterization of how much the allocation can be skewed. We can, however, describe some simple allocations that are always implementable and significantly reduce 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 (1).

PROPOSITION 5 (Inclusive posted price). *The inclusive posted price mechanism yields a higher revenue than any IC and IR mechanism that always allocates the good to a highest-signal bidder.*

PROOF. The inclusive posted price and any standard auction assign the object with probability 1 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.   □

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. 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:

$$p_i(r, s_{-i}) \triangleq \max\{r, v(\underline{s}, s_{-i})\}. \tag{13}$$

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 $v(s) \geq r$,

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

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

PROPOSITION 6 (Personalized price). *The personalized price mechanism yields 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 value at the lowest-type profile is negative and, thus, exclusion becomes strictly beneficial.   □

In the maximum signal model, the optimal personalized price mechanism is revenue equivalent to the two-tier price mechanism. In general common value settings, however, excluding at a given value threshold may no longer be revenue maximizing, even if we restrict attention 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 allocation

$$q_i(s) = \begin{cases} \dfrac{1}{N} & \text{if } \pi(s) \geq 0 \\ 0 & \text{otherwise.} \end{cases} \tag{14}$$

Under mean regularity, the allocation defined by (14) is monotonic, so that it can be implemented by a pricing rule analogous to (13). In particular, a bidder who is allocated the good must pay

$$p_i(s_{-i}) \triangleq \min\{v(s_i', s_{-i}) | \pi(s_i', s_{-i}) \geq 0\}.$$

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 random two-tier posted price. 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 shortly describe a class of such distributions that can be implemented. The high type s_i is allocated the good if and only if $s_i \geq x$; 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 on the high signal s_i only.

Let us introduce the notation

$$\hat{v}(x, y) \triangleq \mathbb{E}\big[v(s_i, s_{-i})|s_i = x, \max s_{-i} = y\big]$$

$$\tilde{v}(x, y) \triangleq \mathbb{E}\big[v(s_i, s_{-i})|s_i = x, \max s_{-i} \leq y\big].$$

PROPOSITION 7 (Generalized random two-tier posted price). *Assume that the signal distributions are symmetric, i.e., $F_i = F$, and that for all $x \in S$, the expression*

$$\frac{\hat{v}(y, x) - \hat{v}(x, x)}{\tilde{v}(y, x) - \tilde{v}(x, x)} \tag{15}$$

is uniformly bounded across x and y. Then there is a generalized random two-tier posted price that yields a higher revenue than any personalized price mechanism for the same exclusion level r.

PROOF. The prices are 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 $s_i > \max s_{-i} > x$, the high bidder pays

$$p(x) \triangleq \tilde{v}(x, x) - \big(\hat{v}(x, x) - \hat{v}(\underline{s}, x)\big)\frac{1 - H(x)}{h(x)}\frac{f(x)}{F(x)} \tag{16}$$

if $s_i \geq x > y$. Finally, if a lower-signal bidder is allocated the good, he pays $\hat{v}(\underline{s}, \max s_{-i})$.

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

$$\int_{y=\underline{s}}^{s_i'} \int_{x=r}^{s_i'} \big(\hat{v}(s_i, y) - \mathbb{I}_{x>y}p(x) - \mathbb{I}_{y>x}\hat{v}(y, y)\big)h(x)\,dx\,d\big(F^{N-1}(y)\big)$$

$$+ \int_{y=s_i'}^{\overline{s}} \big(\hat{v}(s_i, y) - \hat{v}(\underline{s}, y)\big)\frac{1 - H(y)}{N - 1}\,d\big(F^{N-1}(y)\big).$$

The derivative with respect to s_i' is

$$\big(\tilde{v}(s_i, s_i') - p(s_i')\big)h(s_i')\Gamma(s_i') + \big(\hat{v}(s_i, s_i') - \hat{v}(s_i', s_i')\big)H(s_i')\frac{d\big(F^{N-1}(s_i')\big)}{ds_i'}$$

$$- \big(\hat{v}(s_i, s_i') - \hat{v}(\underline{s}, s_i')\big)\frac{1 - H(s_i')}{N - 1}\frac{d\big(F^{N-1}(s_i')\big)}{ds_i'}.$$

By inserting the formula (16) for p, the derivative of the indirect utility reduces to

$$\big(\tilde{v}(s_i, s_i') - \tilde{v}(s_i', s_i')\big)h(s_i')F^{N-1}(s_i')$$

$$- \big(\hat{v}(s_i, s_i') - \hat{v}(s_i', s_i')\big)\left(\frac{1 - NH(s_i')}{N-1}\right)\frac{d\big(F^{N-1}(s_i')\big)}{ds_i'}.$$

Thus, as long as H satisfies

$$\frac{h(x)}{1 - NH(x)} \geq \max_y \frac{\hat{v}(y, x) - \hat{v}(x, x)}{\tilde{v}(y, x) - \tilde{v}(x, x)}\frac{f(x)}{F(x)}, \tag{17}$$

bidder surplus is single-peaked at $s_i' = s_i$ and truthful reporting is incentive compatible. Since (15) is uniformly bounded by some C, inequality (17) is satisfied as long as

$$\frac{h(x)}{1 - NH(x)} \geq C\frac{f(x)}{F(x)}.$$

This equation is satisfied as an equality by the function

$$H(x) = \frac{1}{N}\left(1 - \left(\frac{F(x)}{F(r)}\right)^{\frac{C}{N-1}}\right).$$

As a result, there exist H functions that satisfy $H(\underline{s}) = 0$ and asymptote to $H(\infty) \leq 1/N$, and satisfy (17). The associated generalized random two-tier posted price is incentive compatible. As the generalized random two-tier posted price induces an allocation that is more advantageously selective than that induced by personalized prices, by Theorem 4, it generates more revenue. □

Note that the hypotheses for Proposition 7 are satisfied for the maximum signal model, since then (15) is equal to 1 for all x and y.

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 feature that the 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 relevant 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 exclusionary power. We can construct simple exclusive mechanisms that 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, Dirk, Benjamin Brooks, and Stephen Morris (2017a), "First price auctions with general information structures: Implications for bidding and revenue." *Econometrica*, 85, 107–143. [1400]

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris (2017b), *Selling to Intermediaries: Optimal Auction Design in a Common Value Model*. 2064R, Cowles Foundation for Research in Economics.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris (2019), "Revenue Guarantee Equivalence." *American Economic Review*, 109, 1911–1929. [1400, 1406, 1407]

Brooks, Benjamin and Songzi Du (2019), *Optimal Auction Design With Common Values: An Informationally Robust Approach*. University of Chicago and University of California at San Diego. [1403]

Bulow, Jeremy and Paul Klemperer (1996), "Auctions Versus Negotiations." *American Economic Review*, 86, 180–194. [1400, 1407, 1408, 1414, 1424, 1425, 1427]

Bulow, Jeremy and Paul Klemperer (2002), "Prices and the Winner's Curse." *RAND Journal of Economics*, 33, 1–21. [1401, 1402, 1406, 1407, 1414, 1425, 1427, 1433]

Campbell, Colin and Dan Levin (2006), "When and why not to auction." *Economic Theory*, 27, 583–596. [1401]

Carroll, Gabriel and Ilya Segal (2019), "Robustly Optimal Auctions with Unknown Resale Opportunities." *Review of Economic Studies*, 86, 1527–1555. [1401]

Du, Songzi (2018), "Robust Mechanisms Under Common Valuation." *Econometrica*, 86, 1569–1588. [1403]

Fieseler, Karsten, Thomas Kittsteiner, and Benny Moldovanu (2001), *Partnerships, Lemons and Efficient Trade*. University of Mannheim. [1410]

Gupta, Madhurima and Bernard LeBrun (1999), "First price auctions with resale." *Economics Letters*, 64, 181–185. [1400]

Haile, Phil (2003), "Auctions with private uncertainty and resale opportunities." *Journal of Economic Theory*, 108, 72–110. [1400]

Jehiel, Philippe and Laurent Lamy (2015), "On Discrimination in Auctions with Endogenous Entry." *American Economic Review*, 105, 2595–2643. [1426]

Kirkegaard, Rene (2006), "A short proof of the Bulow–Klemperer auctions vs. negotiations result." *Economic Theory*, 28, 449–452. [1424]

Levin, Dan and James Smith (1994), "Equilibrium in Auctions with Entry." *American Economic Review*, 84, 585–599. [1426]

Milgrom, Paul and Robert Weber (1982), "A Theory of Auctions and Competitive Bidding." *Econometrica*, 50, 1089–1122. [1400, 1414]

Myerson, Roger (1981), "Optimal Auction Design." *Mathematics of Operations Research*, 6, 58–73. [1400, 1401, 1402, 1404, 1427, 1432]

Povel, Paul and Rajdeep Singh (2004), "Using bidder asymmetry to increase seller revenue." *Economics Letters*, 84, 17–20. [1427]

Povel, Paul and Rajdeep Singh (2006), "Takeover Contests with Asymmetric Bidders." *Review of Financial Studies*, 19, 1399–1431. [1427]

Segal, Ilya and Michael Whinston (2011), "A simple status quo that ensures participation (with application to efficient bargaining)." *Theoretical Economics*, 6, 109–125. [1410]