American Economic Review 2019, 109(5): 1911–1929
https://doi.org/10.1257/aer.20180773

Revenue Guarantee Equivalence†

By Dirk Bergemann, Benjamin Brooks, and Stephen Morris*

We revisit the revenue comparison of standard auction formats, including first-price, second-price, and English auctions. We rank auctions according to their revenue guarantees, i.e., the greatest lower bound of revenue across all informational environments, where we hold fixed the distribution of bidders' values. We conclude that if we restrict attention to the symmetric affiliated models of Milgrom and Weber (1982) and monotonic pure-strategy equilibria, first-price, second-price, and English auctions are revenue guarantee equivalent: they have the same revenue guarantee, which is equal to that of the first-price auction as characterized by Bergemann, Brooks, and Morris (2017). If we consider all equilibria or if we allow more general models of information, then first-price auctions have a greater revenue guarantee than all other auctions considered. (*JEL* D44, D83)

The revenue equivalence theorem is a central result in auction theory that helps us understand the relationship between the choice of auction format and the resulting revenue. In an environment with symmetric *independent private values*, it implies the surprising result that many standard auction formats, including first-price, second-price, and English auctions, all deliver the same expected revenue (Myerson 1981). By contrast, in a symmetric environment with *affiliated values*, there is a revenue ranking theorem that establishes that the first-price auction achieves less revenue than the second-price auction which in turn generates less revenue than the English auction (Milgrom and Weber 1982—henceforth, MW). And yet, there are other cases where first-price auctions will generate more revenue than second-price and English auctions, such as the independent private-value model when value distributions are asymmetric (Vickrey 1961, Maskin and Riley 2000).

Against this background, we provide a new revenue ranking theorem for environments with interdependent values. The equilibrium outcomes for a given auction

* Bergemann: Department of Economics, Yale University, Box 208281, New Haven, CT 06520-8281 (email: dirk.bergemann@yale.edu); Brooks: Department of Economics, University of Chicago, 1126 E 59th Street, Chicago, IL 60637 (email: babrooks@uchicago.edu); Morris: Department of Economics, Princeton University, 436 Prospect Avenue, Princeton, NJ 08540 (email: smorris@princeton.edu). Jeffrey Ely was the coeditor for this article. We acknowledge financial support through NSF grant SES 1459899. We would like to thank the coeditor and three anonymous referees for productive suggestions. We are grateful to Christian Baker and Ian Ball for excellent research assistance. We also thank seminar audiences at Stanford University, ESSET 2018, and Western Ontario for helpful comments. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

† Go to https://doi.org/10.1257/aer.20180773 to visit the article page for additional materials and author disclosure statements.

format are conventionally analyzed for a fixed distribution of the bidders' values and for a fixed information structure that generates the signals that bidders observe before interacting with the auction. It is known that equilibrium welfare outcomes, such as revenue, are sensitive to the specific form of information, e.g., Fang and Morris (2006) and Bergemann, Brooks, and Morris (2017)—henceforth, BBM. This presents a challenge for comparative auction theory, since it may be difficult to pin down the correct model of information, either through introspection or measurement. Given this state of affairs, we propose to rank auctions by a criterion that is less sensitive to misspecification of the informational environment. In particular, we evaluate an auction according to its *revenue guarantee*: the greatest lower bound on the auction's revenue that holds across all information structures and equilibria. Importantly, this guarantee is computed while holding fixed the payoff environment, that is, the distribution of the bidders' values.

We establish a revenue guarantee ranking for a collection of *standard auctions*, which are defined by (i) having a rule that the high-bidder wins the good and (ii) admitting an equilibrium in monotonic pure-strategies in the symmetric independent private value model. This class includes the first-price auction, the second-price auction, and the English auction. Our main result is that the first-price auction attains the highest revenue guarantee of any standard auction. In particular, the first-price auction has a greater revenue guarantee than both the second-price and English auctions.

A first step to obtain this ranking is to establish the revenue guarantee of the first-price auction. Here we appeal to an earlier result in BBM that establishes that the lowest revenue in the first-price auction arises in an information structure that we refer to as the *maximum signal model*. In this information structure, the bidders' signals are identically and independently distributed (i.i.d.) random variables with the property that the highest value bidder receives the highest signal, and the highest signal is equal to the average of the losing bidders' values.

We then ask what is the equilibrium revenue of standard auctions in the maximum signal model. Here, we obtain the first surprising result. In this specific information structure, any standard auction has an equilibrium in which bidders behave as if each bidder's expected value is equal to his individual signal rather than his true expected value, which is strictly higher. Thus, the bidders act as if they are in an independent private-value environment, even though their expected values are in fact greater than their signals. This generalizes the construction of the equilibrium of the first-price auction, given by BBM, to all standard auctions. Given this *strategic equivalence* with the independent private-value model, we can appeal to the standard revenue equivalence result to conclude that all standard auctions generate the same revenue in an equilibrium of the maximum signal model. As a result, the revenue guarantee of any standard auction can be *at most equal* to the revenue guarantee of the first-price auction.

A critical feature of our analysis is that the revenue guarantee is computed by minimizing across all information structures, including those in which bidders have only partial information about their own values. This is important to our conclusion that the first-price auction revenue-guarantee dominates the second-price auction. (A corresponding discussion applies to the English auction as well.) When bidders know their own values but may have partial information about others' values, i.e., when values are *private*, the second-price auction has a unique equilibrium in weakly undominated strategies in which bidders bid their values. Thus, with private values and the weak dominance refinement, revenue is completely insensitive to information. Indeed, Chung and Ely (2007) have given conditions under which a private-value analogue of the revenue guarantee will be maximized by such a dominant-strategy mechanism (although not necessarily the second-price auction). At the same time, these auctions have other equilibria in which revenue is very low. As we discuss later, when values are not private, such low-revenue equilibria cannot be ruled out using the weak dominance refinement. Since the revenue guarantee of the first-price auction in BBM is valid across all equilibria, we conclude that the revenue guarantee of the first-price auction is strictly higher than that of the second-price auction.

At the same time, the second-price and English auctions are known to perform well in certain interdependent value environments, namely the symmetric affiliated values models studied by MW. In that setting, these auctions have unique equilibria in monotonic pure strategies that lead to a nontrivial amount of revenue. We may ask, what is the revenue guarantee ranking if we restrict attention to symmetric affiliated values and monotonic equilibria? By modifying our methodology in this manner, we are assured that the resulting *affiliated revenue guarantees* and corresponding rankings are based on standard and uncontroversial equilibrium constructions.

Theorem 2 shows when bidders all have a *common* value for the good, then according to affiliated revenue guarantees, the first-price, second-price, and English auctions are all *revenue guarantee equivalent*. This result is established by showing that when the bidders have a common value, the critical maximum signal model is itself affiliated, so that the weak ranking of Theorem 1 is preserved. At the same time, it is a result of MW that the first-price auction generates weakly lower revenue than the second-price and English auctions in the affiliated values model. We therefore conclude that all of these mechanisms must have exactly the same affiliated revenue guarantee.

Thus, Theorems 1 and 2 offer different perspectives of the revenue ranking result in the affiliated values model. If we are concerned with the robustness of the revenue comparison across all symmetric affiliated values information structures and monotonic equilibria, we find that the second-price and English auctions lose their advantage, as stated in the revenue guarantee equivalence theorem. Moreover, if we are at the same time concerned with the equilibrium selection, or seek to offer a revenue guarantee that is valid across all information structures and all equilibria, then we find that the first-price auction offers better guarantees than either the second-price or English auctions.

In light of our results, a natural question is: what is the mechanism with the greatest revenue guarantee? This question is answered by Bergemann, Brooks, and Morris (2016) when there are two bidders and binary common values and by Brooks and Du (2018) for general common value models. These papers construct a *saddle point*, consisting of an auction that maximizes the revenue guarantee (a maxmin auction) and a minmax information structure for which maximum revenue across all auctions and equilibria is equal to the optimal revenue guarantee. The maxmin auctions look quite different from the standard auctions considered here, and necessarily involve randomized allocations to optimally hedge ambiguity about the information structure. This paper constructs an analogous saddle point except that the seller is restricted to using standard auctions. In this case, the first-price auction is a maxmin auction and the maximum signal model is a minmax information structure. We view these results as complements: revenue guarantees are one of many criteria that could be used in selecting an auction format, and while the standard auctions considered here do not achieve the unconstrained optimal revenue guarantee, they have other desirable attributes that may rationalize their use.

The rest of the paper proceeds as follows. Section I presents the model. Section II motivates our results via a simple example. The revenue guarantee ranking of standard auctions is the subject of Section III, and revenue guarantee equivalence under affiliation is the subject of Section IV. Section V concludes.

I. Model

There are N bidders for a single unit of a good. The set of bidder indices is $\mathcal{N} = \{1, \ldots, N\}$. Bidders' values $(v_1, \ldots, v_N)$ are jointly distributed according to a probability measure $\pi(dv_1, \ldots, dv_N)$ on $\mathbb{R}_+^N$. We denote a profile, or vector of values by $v = (v_1, \ldots, v_N)$. We assume these values are non-negative and that the distribution π is exchangeable. In particular, for every bijection $\xi : \mathcal{N} \to \mathcal{N}$ and for every measurable set $X \subseteq \mathbb{R}_+^N$,

$$\pi(X) = \pi\big(f_\xi(X)\big),$$

where $f_\xi : \mathbb{R}_+^N \to \mathbb{R}_+^N$ is defined by

$$f_\xi(v) = \big(v_{\xi(1)}, \ldots, v_{\xi(N)}\big).$$

An *information structure* consists of measurable sets of signals S_i and a joint probability measure $\mu(ds_1, \ldots, ds_N, v)$ on signal and value profiles in $S_1 \times \cdots \times S_N \times \mathbb{R}_+^N$. A representative information structure is denoted $\mathcal{I}$. Any such distribution μ induces an interim expected value function

$$w^{\mathcal{I}} : S \to \mathbb{R}_+^N, \tag{1}$$

where $w^{\mathcal{I}}(s)$ is a version of the conditional expectation of the values conditional on the signal profile.

An information structure is *symmetric* if $S_1 = \cdots = S_N$ and if for all bijections $\xi : \mathcal{N} \to \mathcal{N}$, we have

$$\mu(X) = \mu\big((g_\xi, f_\xi)(X)\big),$$

where $g_\xi : S \to S$ is defined as

$$g_\xi(s) = \big(s_{\xi(1)}, \ldots, s_{\xi(N)}\big),$$

and $(g_\xi, f_\xi)(s, v) = \big(g_\xi(s), f_\xi(v)\big)$.

An information structure is *independent* if the s_i are independent random variables.

An information structure has *common values* if $v_1 = \cdots = v_N$, μ-almost surely.

An information structure has *private values* if $w^{\mathcal{I}}(s) = v_i$, μ-almost surely.

An information structure is *consistent with* π if for all measurable sets $X \subseteq \mathbb{R}_+^N$, $\mu(S \times X) = \pi(X)$, i.e., π is the marginal of μ on values.

A *mechanism* consists of measurable sets of messages M_i for each player, $M = \times_{i=1}^N M_i$, allocations $q : M \to [0,1]^N$ with $\sum_{i=1}^N q_i(m) \leq 1$ for all m, and transfers to the seller $t : m \to \mathbb{R}_+^N$. A representative mechanism is denoted $\mathcal{M}$.

A pair of an information structure $\mathcal{I}$ and mechanism $\mathcal{M}$ comprise a Bayesian game. A profile of strategies $\sigma = (\sigma_1, \ldots, \sigma_N)$, where $\sigma_i : S_i \to \Delta(M_i)$, induces ex ante expected surplus for bidder i equal to

$$U_i(\sigma; \mathcal{M}, \mathcal{I}) = \int_{s \in S, v \in \mathbb{R}^N} \int_{m \in M} (v_i q_i(m) - t_i(m)) \sigma(dm|s) \, \mu(ds, dv),$$

where we identify the profile σ with a mapping from signal profiles to conditionally independent distributions over message profiles. A *Bayes Nash equilibrium* is a strategy profile σ such that for all bidder i strategies σ_i',

$$U_i(\sigma; \mathcal{M}, \mathcal{I}) \geq U_i(\sigma_i', \sigma_{-i}; \mathcal{M}, \mathcal{I}).$$

A strategy profile induces revenue

$$R(\sigma; \mathcal{M}, \mathcal{I}) = \int_{s \in S, v \in \mathbb{R}_+^N} \int_{m \in M} \sum_{i=1}^N t_i(m) \, \sigma(dm|s) \, \mu(ds, dv).$$

We say that R is a *revenue guarantee* of the mechanism $\mathcal{M}$ with respect to π if $R(\sigma; \mathcal{M}, \mathcal{I}) \geq R$ for all $\mathcal{I}$ that are consistent with π and for all Bayes Nash equilibria σ of $(\mathcal{M}, \mathcal{I})$. Further, R is *the revenue guarantee of* $\mathcal{M}$, relative to π, if it is a revenue guarantee, and if there is no higher guarantee.

II. A Revealing Example

To illustrate our findings, let us consider an example where there are two bidders who have a pure common value v which is uniformly distributed between 0 and 1. We will compute and compare the revenue guarantees of the first- and second-price auctions. In a sense, we already know the result: the first-price auction has a nontrivial revenue guarantee, as computed by BBM. The second-price auction, on the other hand, always has equilibria in dominated strategies in which revenue is zero. Any conclusion based on these dominated equilibria would seem quite fragile. We shall see, however, that the first- and second-price auction guarantees can be ranked using equilibria that are quite natural and use undominated strategies.

To develop this result, we will compare revenue in the first- and second-price auctions on three information structures. Welfare outcomes for these examples are depicted in Figure 1, with revenue on the y-axis and total bidder surplus on the x-axis. Note that both auctions always allocate the good, so total surplus is $1/2$. Moreover, both revenue and bidder surplus must be non-negative. Thus, any welfare outcome for these mechanisms that can arise in some information structure and some equilibrium must lie on the line segment between $(1/2, 0)$ and $(0, 1/2)$.

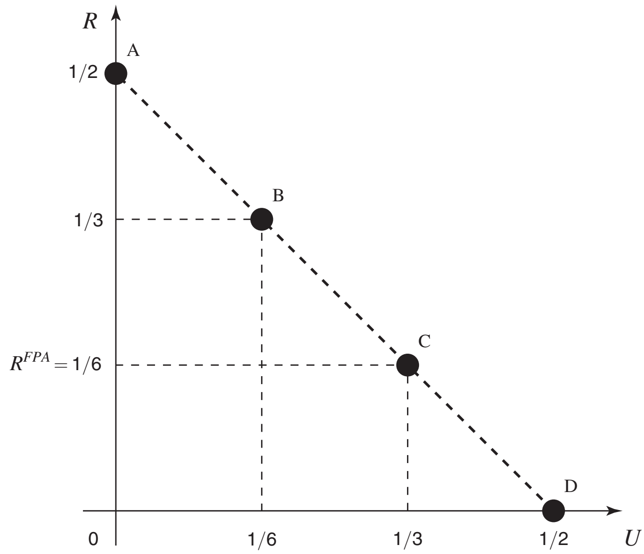

Figure 1. Welfare Outcomes for the Two-Bidder Uniform Example

Notes: Revenue is on the *y*-axis and bidder surplus is on the *x*-axis. The set of welfare outcomes which can arise in a first-price auction for some information structure and equilibrium is the line segment between A and C. The corresponding set for the second-price auction is the line segment between A and D.

Consider first the case of complete information, wherein both bidders see the value v exactly. The first-price auction has a unique equilibrium in which both bidders tie at a bid v. For the second-price auction, any bid profile would be an equilibrium as long as the high bid is at least v and the low bid is at most v. However, bidding v is a weakly dominant strategy. The ranking of first- and second-price auctions under public information therefore depends on our willingness to restrict attention to strategies that are weakly undominated. One can make arguments either way: dominated strategies are unsafe, but they could be used to facilitate collusion, as members of a bidding ring take turns as the high bidder (Rothkopf, Teisberg, and Kahn 1990). But with the weak dominance refinement, both auctions and equilibria result in expected revenue of $1/2$, i.e., full surplus extraction. This is point A in Figure 1.

Next, consider an information structure in which bidder 1 observes the value and bidder 2 receives no information, meaning that she always sees the same signal regardless of the true value. This is the "proprietary information" model of Engelbrecht-Wiggans, Milgrom, and Weber (1983). They show the first-price auction has an equilibrium in which bidder 1 bids $v/2$ and bidder 2 randomizes uniformly between 0 and $1/2$. To see that this is an equilibrium, observe if bidder 2 bids $m_2 \in [0, 1/2]$, she wins on the event that $v \in [0, 2m_2]$, in which case the expected value conditional on winning is exactly her bid. Thus, all bids between 0 and $1/2$ result in zero surplus. For bidder 1, a bid of $m_1 \in [0, 1/2]$ results in a probability $2m_1$ of winning, independent of the value, so that expected surplus is $(v - m_1)2m_1$. This expression is clearly maximized at $m_1 = v/2$. Expected revenue is the expected highest of two uniform draws from $[0, 1/2]$, which is $1/3$. This is point B in Figure 1.

In the second-price auction, bidder 1 has a weakly dominant strategy to bid v, as with complete information. But if bidder 1 follows this strategy, *all* bids in $[0,1]$ are best responses for bidder 2. Moreover, none of these bids is weakly dominated.[1] Thus, there are undominated equilibria of the second-price auction that generate all expected revenues in $[0,1/2]$. This is the entire line segment between A and D in Figure 1. This shows that even with the weak dominance refinement, we cannot escape the conclusion that there are equilibria of the second-price auction with very low revenue.[2]

The last information structure we consider is the one that minimizes revenue for the first-price auction, as computed in BBM. One bidder, chosen at random, sees the true value, and the other bidder sees a random draw from $[0,v]$ according to the cumulative distribution function $F(x|v) = \sqrt{x/v}$. Thus, the highest signal is always equal to the true value. The distribution of the low signal is chosen just so that if we integrate across values, the signals are i.i.d. draws from the distribution $F(x) = \sqrt{x}$ on $[0,1]$. The resulting distribution of the highest signal is standard uniform, the same as the value. We refer to this as the *maximum signal model*, and we denote it by $\mathcal{I}^*$.[3] In the maximum signal model, the pure common value function can be written as

$$v = w^{\mathcal{I}^*}(s_1,s_2) = \max\{s_1,s_2\}.$$

In analyzing equilibria on $\mathcal{I}^*$, we shall reverse the order of exposition and first describe an equilibrium of the second-price auction. As shown by Bulow and Klemperer (2002), there is a "truthful" equilibrium in which bidders bid their signals, i.e., $m_i = s_i$. In fact, it is an ex post equilibrium. To see that this is the case, fix a signal profile $s = (s_i,s_j)$. Any bid $m_i < s_i$ is weakly dominated by bidding s_i, since the true value is at least s_i with probability 1. On the other hand, a deviation to a bid $m_i > s_i$ would result in bidder i additionally winning when $s_i < m_j = s_j \leq m_i$. On this event, s_j is the high signal and is therefore the value. But it is also equal to the second-highest bid, so that ex post surplus is zero. Thus, there is no change in surplus from deviating to a higher bid. Expected revenue in this equilibrium is $1/6$, i.e., the expected second-highest signal.

Note that in this equilibrium, each bidder behaves as if their signal were the true value, even though the signal is only a lower bound. More precisely, the bidders

[1] For example, the bid $m_2 \in [0,1]$ is the unique best response when bidder 1 bids $m_1 = v^2/m_2$.

[2] Abraham et al. (2014) argue that all of these equilibria will survive trembling hand perfection. They propose a stronger refinement that they term "tremble robust equilibrium" which selects $m_2 = 0$ as the unique equilibrium strategy for bidder 2. See also Milgrom and Mollner (2018) for a discussion of refinements in this game. The multiplicity of undominated equilibria and the existence of equilibria with extremely low revenue is hardly limited to this information structure. Milgrom (1981) gives other examples of interdependent values information structures in which the second-price auction has a continuum of undominated equilibria, some of which generate extremely low revenue. Milgrom (1981, p. 939), writes, "Vickrey auctions generally have a plethora of equilibria, many of which are quite pathological… I can find no completely convincing way to rule out these strange equilibria." This is in stark contrast to the private value case, where weak dominance selects a unique equilibrium, regardless of the information structure.

[3] Bulow and Klemperer (2002) were the first to study the common-value maximum signal model in the context of the English auction. They showed that bidding one's signal is an equilibrium and that the resulting revenue is less than what the seller would obtain with a posted price. Bergemann, Brooks, and Morris (2018) characterize the optimal auction in the maximum signal model. They show that the posted price is optimal when the good must be sold, but otherwise the optimal mechanism has a different form.

play the standard equilibrium of the second-price auction for the information structure in which values are independent, private, and symmetric draws from F, which is denoted $\mathcal{I}^{IPV}$. Note also that this information structure is not consistent with the original uniform common value distribution. Now, it is well known that in the symmetric IPV model, the first-price auction has an equilibrium in which the ex post allocation and each type's expected payment are the same as in the truthful equilibrium of the second-price auction. Indeed, these properties pin down the equilibrium bid in the first-price auction as $b_i = s_i/3$. Returning to the first-price auction in the common value environment $\mathcal{I}^*$, if both bidders follow the strategies of bidding $s_i/3$, then the interim payment and ex post allocation from bidding as the type s_i' are the same as they were in the second-price auction equilibrium on $\mathcal{I}^*$. Since bidders did not want to mimic other types in the second-price auction, we conclude that the same holds in this equilibrium of the first-price auction. Moreover, any non-equilibrium bid is dominated by a lower equilibrium bid with the same winning probability. We conclude that these strategies are an equilibrium of the first-price auction on $\mathcal{I}^*$, with the same revenue of $1/6$. This welfare outcome corresponds to point C in Figure 1.

As we comment on further below, this is in fact the revenue guarantee of the first-price auction, meaning that there is no information structure and equilibrium with lower expected revenue. Some intuition for this can be gleaned from the fact that in this equilibrium, bidders are indifferent between their equilibrium bids and all higher bids. This is suggestive of the difficulty of achieving a lower bid distribution, without some bidder being tempted to deviate up. This intuition is incomplete, however, since the same indifference also obtains in the truthful equilibrium of the second-price auction on $\mathcal{I}^*$.

The equilibrium analysis on $\mathcal{I}^*$ reaffirms what we already knew: first-price auction has a nontrivial revenue guarantee, attained at $\mathcal{I}^*$. Moreover, the second-price auction has an equilibrium on $\mathcal{I}^*$ with the same revenue, so that the second-price auction's revenue guarantee is weakly lower. But it is conceptually important that we can rank the mechanisms using this environment and these equilibria. Notice that $\mathcal{I}^*$ is symmetric, has independent signals, and the value function

$$w^{\mathcal{I}^*}(s_1, s_2) = \max\{s_1, s_2\}$$

is increasing in the signals. It is therefore an affiliated values model in the sense of MW. In such environments, both first-price auction and second-price auction will always have nontrivial symmetric and monotonic pure-strategy equilibria, which reduce to the equilibria described above.

Affiliated information structures represent a class of environments in which there is a principled way to select a nontrivial and appealing equilibrium of the second-price auction. Thus, the revenue guarantees of first- and second-price auctions are ranked without resorting to exotic or dominated equilibria. Theorem 2 pursues this idea to its logical conclusion, and argues that if we restrict attention to affiliated values models in computing the revenue guarantee, the two mechanisms are *revenue guarantee equivalent*. For the equilibrium comparison in $\mathcal{I}^*$ shows that the affiliated guarantee of the first-price auction must still exceed that of the second-price auction, and MW show that the second-price auction generate weakly more revenue than the first-price auction in any affiliated values model, so that $\mathcal{I}^*$ is the worst-case for the second-price auction as well.

III. Revenue Guarantee Ranking

We now establish a revenue guarantee ranking across a number of classic auction formats, including the first-price, second-price, and English auctions. We begin the analysis by reviewing the revenue guarantee of the first-price auction.

A. *Revenue Guarantee of the First-Price Auction*

For a real-vector $x \in \mathbb{R}^N$, we let $x^{(k)}$ denote the kth highest element of the vector. Thus, $x^{(1)}$ is the first-order statistic, $x^{(2)}$ is the second-order statistic, etc.

The first-price auction, denoted $\mathcal{M}^{FPA}$, is formally defined as follows: $M_i = \mathbb{R}_+$,

$$q_i^{FPA}(m) = \begin{cases} \frac{1}{|\mathrm{argmax}_j m_j|}, & \text{if } i \in \mathrm{argmax}_j m_j, \\ 0, & \text{otherwise;} \end{cases}$$

and

$$t_i^{FPA}(m) = q_i^{FPA}(m) m^{(1)}.$$

For any vector of values v, we define the average among the $N-1$ lowest values by $\alpha : \mathbb{R}^N \to \mathbb{R}$:

$$\alpha(v) = \frac{1}{N-1}\left(\sum_{i=1}^{N} v_i - \max_i v_i\right). \tag{2}$$

This statistic has a cumulative distribution function

$$H(x) = \int_{\{v \in \mathbb{R}^N | \alpha(v) \leq x\}} \pi(dv). \tag{3}$$

Thus, H is the distribution of the average of the $N-1$ lowest values, $\alpha(v)$, when the value profile v is distributed according to π. For technical convenience, we assume throughout the rest of our analysis that π is such that the induced H is non-atomic.[4]

The critical information structure, which we denote by $\mathcal{I}^*$, is as follows. The sets of signals are $S_i = \mathbb{R}_+$ and are independent draws from the distribution

$$G(x) = \left(H(x)\right)^{1/N}.$$

This distribution is chosen such that the first-order statistic of these N signals coincides with the distribution of $\alpha(v)$. The signals are correlated with the value so that the bidder with the highest ex post value v_i observes the highest signal, breaking ties uniformly, and the remaining signals are randomly assigned to lower value bidders.

[4] See BBM for a discussion of relaxing this assumption.

An equivalent description of this information structure can be given in terms of the conditional distribution of signals given the values, $\mu(ds|v)$. Conditional on the value profile v, we pick one of the high-value bidders at random and give that bidder a signal $s_i = \alpha(v)$. The remaining bidders are shown signals that are independent draws from the distribution

$$\frac{G(s_j)}{G(\alpha(v))}$$

over the range $[0, \alpha(v)]$.

Theorem 1 in BBM establishes that the revenue guarantee in the first-price auction, R^{FPA} is the expected second-highest of N draws from the cumulative distribution G:

$$R^{FPA} = E_{(s_1,\dots,s_N)\overset{iid}{\sim}G}\left[s^{(2)}\right]. \tag{4}$$

This level of revenue is attained in a Bayes Nash equilibrium of information structure $\mathcal{I}^*$, as we now explain. By construction, the high-value bidder always receives a signal which is equal to $\alpha(v)$, the average of the $N-1$ lowest value bidders. Thus, conditional on the signal profile s, the average of the $N-1$ lowest values is known to be $\max_i s_i$. Moreover, due to exchangeability of π, low-value bidders believe they are equally likely to be anywhere in the distribution, so their expected value is exactly equal to the highest signal. On the other hand, the high bidder knows $\alpha(v)$ and knows that she has the high value, but does not know what it is, so the expected value of the high signal bidder is

$$\eta(s_i) = E_\pi\left[\max_j v_j \middle| \alpha(v) = s_i\right],$$

which is necessarily greater than s_i. Again, this expression is independent of i, due to the symmetry of π. To summarize, the induced interim expected value function is

$$w_i^{\mathcal{I}^*}(s) = \begin{cases} \max_j s_j, & \text{if } s_i < \max_j s_j; \\ \eta(s_i), & \text{if } s_i = \max_j s_j. \end{cases}$$

This is the generalization of the *maximum signal model* described in the previous section.

There is a monotonic pure-strategy equilibrium of the first-price auction on the maximum signal model in which a type s_i bids

$$\beta^{FPA}(s_i) = E_{s_{-i}\overset{iid}{\sim}G}\left[s_{-i}^{(1)} | s_{-i}^{(1)} \le s_i\right]. \tag{5}$$

By this, we mean that the strategy $\sigma(\cdot|s_i)$ puts probability 1 on $\beta^{FPA}(s_i)$. We hereafter adopt this notation for pure strategies. Proposition 1 below restates Theorem 1 of BBM in the current language and notation.

PROPOSITION 1 (BBM): *The revenue guarantee of the first-price auction is* R^{FPA}*. Moreover, the strategies* β^{FPA} *are an equilibrium of* $(\mathcal{M}^{FPA}, \mathcal{I}^*)$ *and* $R^{FPA} = R(\beta^{FPA}; \mathcal{M}^{FPA}, \mathcal{I}^*)$.

The first step in the proof of this result establishes that R^{FPA} is a lower bound on revenue of $\mathcal{M}^{FPA}$ in any equilibrium in any information structure. The lower bound is attained by solving a relaxed program in which revenue is minimized over winning bid distributions that satisfy a one-dimensional class of incentive constraints, which correspond to bidders deviating by bidding the maximum of their equilibrium bid and some fixed cutoff b. At the minimum, all of these incentive constraints bind.

The second step shows that β^{FPA} is an equilibrium under $\mathcal{I}^*$ in which revenue is R^{FPA}, so that the lower bound is attained. It is this second step that is the most relevant for the new results in our paper. The information structure $\mathcal{I}^*$ is strategically very similar to another information structure, which we denote by $\mathcal{I}^{IPV}$: as before, signals are independent draws from G, but now the joint distribution of values is such that $s_i = v_i$ with probability one for all i. In other words, $\mathcal{I}^{IPV}$ is the *independent private values* information structure in which the individual values are distributed according to G. Importantly, the highest value among the N values has the same distribution H as the value in maximum signal model. We note that all of the bidders, except possibly the bidder with the highest signal, have a strictly lower expected value for the object in $\mathcal{I}^{IPV}$ than they do in $\mathcal{I}^*$.

It is a standard result that there is a monotonic pure-strategy equilibrium of $(\mathcal{M}^{FPA}, \mathcal{I}^{IPV})$ in which a bidder with value s_i bids the expected highest of the others' signals, conditional on others' signals being less than s_i, i.e., $\beta^{FPA}(s_i)$. We can use this to show that β^{FPA} is also an equilibrium of $(\mathcal{M}^{FPA}, \mathcal{I}^*)$. First, consider a deviation in which a type s_i bids $\beta^{FPA}(s_i')$ for some $s_i' \leq s_i$. Then the bidder only wins when the highest of the others' signals is less than s_i. The deviator's surplus is therefore the same as what it would be in $\mathcal{I}^{IPV}$ but with a value $\eta(s_i)$ instead of s_i. It is well known that a player's surplus is single-peaked in their bid (Matthews 1995). It follows that a bidder's surplus is increasing in his bid as long as he is bidding as a type $s_i' < \eta(s_i)$. Thus, we conclude that downward deviations are not attractive.

On the other hand, by deviating to a higher bid, the deviator's surplus would be

$$\begin{aligned}
&\eta(s_i)\, G^{N-1}(s_i) + \int_{x=s_i}^{s_i'} x\, d\big(G^{N-1}(x)\big) - \beta^{FPA}(s_i')\big(G^{N-1}(s_i')\big) \\
&\quad = \int_{x=0}^{s_i} \eta(s_i)\, d\big(G^{N-1}(x)\big) + \int_{x=s_i}^{s_i'} x\, d\big(G^{N-1}(x)\big) - \int_{x=0}^{s_i'} x\, d\big(G^{N-1}(x)\big) \\
&\quad = \int_{x=0}^{s_i} \big(\eta(s_i) - x\big)\, d\big(G^{N-1}(x)\big),
\end{aligned}$$

which is independent of s_i'. Finally, it is clear that bidding above the highest or below the lowest equilibrium bid is dominated.

From the revenue equivalence theorem (Myerson 1981), we know that revenue in this equilibrium must be equal to that of the second-price auction, which is the expected second-highest signal, thus giving us the formula (4).

While a complete proof that $\mathcal{I}^*$ minimizes revenue for the first-price auction is beyond the scope of this paper, we observe that $\mathcal{I}^*$ has two critical properties that make revenue especially low. First, we note it is informative enough so that the resulting equilibrium allocation is efficient. This is important, because bidders incentive to deviate up depends on their values when they *lose* the auction. When the allocation is efficient, losing bidders' values are lowest, so that the incentive to deviate up is weakest. Second, it is well known that equilibrium bid shading in common value environments depends on how bidders update their expected values when they condition on winning the auction. Roughly speaking, the more that bidders revise their expectations down, the less they bid in equilibrium. In the maximum signal model, this updating is extreme. When a bidder just conditions on their signal s_i, the signal is a lower bound on the ex post value. But on the marginal event where a bidder wins by slightly increasing their bid, s_i is *equal* to the expected value. In equilibrium, bid shading is so large that bidders behave as if their true value were the lowest *possible* value.

B. *Revenue Ranking*

With the revenue guarantee of the first-price auction in hand, we now turn to ranking the revenue guarantees of standard mechanisms. A mechanism is *standard* if (i) messages are one-dimensional bids and the high bidder is allocated the good and (ii) there is a symmetric and strictly monotonic pure-strategy equilibrium when values are symmetric, independent, and private. The first-price auction, second-price auction, all-pay auction, and the war-of-attrition[5] are all examples of standard auctions.

Our first main result is a ranking of revenue guarantees of standard auctions.

THEOREM 1 (Revenue Guarantee Ranking): *If $\mathcal{M}$ is standard, then R^{FPA} is greater than any revenue guarantee of $\mathcal{M}$. Thus, the first-price auction maximizes the revenue guarantee among standard auctions.*

To prove the theorem, we first establish the following result.

PROPOSITION 2 (Strategic Equivalence): *Suppose that $\mathcal{M}$ is a standard mechanism and β is a symmetric and strictly monotonic pure-strategy equilibrium of $(\mathcal{M}, \mathcal{I}^{IPV})$. Then β is also an equilibrium of $(\mathcal{M}, \mathcal{I}^*)$.*

PROOF OF PROPOSITION 2:

When others use the strategy β, bidding $\beta(s_i)$ must result in the bidder winning when $s_i \geq \max_{j \neq i} s_j$ and making an interim payment $T(s_i)$. This "direct" allocation is precisely the one that is induced by the first-price auction. Moreover, from the revenue equivalence theorem, we know that the interim expected payment must be the same as that induced by the first-price auction, denoted $T^{FPA}(s_i)$, up to a constant that depends on i but not on s_i:

$$T(s_i) = T^{FPA}(s_i) + c_i.$$

[5] By war of attrition, we mean the sealed-bid mechanism defined by Krishna and Morgan (1997): each bidder submits a real number, the high bidder wins, and all bidders pay the minimum of their bid and the second-highest bid.

Thus, a bidder with signal s_i profits from a deviation from $\beta(s_i)$ to $\beta(s_i')$ in the game $(\mathcal{M},\mathcal{I}^*)$ if and only if s_i profits from a deviation from $\beta^{FPA}(s_i)$ to $\beta^{FPA}(s_i')$ in the game $(\mathcal{M}^{FPA},\mathcal{I}^*)$. Since the latter deviation is unprofitable, the former must be as well. Finally, it cannot be that there is any type that wants to deviate to a message that is not sent in equilibrium. The fact that there are no atoms in the bid distribution induced by β implies that for any message, there is an equilibrium message which induces the same allocation. If an out-of-equilibrium message were a profitable deviation, it would have to entail a lower expected transfer than the equilibrium message with the same winning probability, which contradicts the hypothesis that β is an equilibrium of $(\mathcal{M},\mathcal{I}^{IPV})$. ∎

We now complete the proof of Theorem 1.

PROOF OF THEOREM 1:

We will simply exhibit an information structure and equilibrium in which revenue is less than or equal to R^{FPA}. The information structure is $\mathcal{I}^*$. From the hypothesis that $\mathcal{M}$ is standard, we know that $(\mathcal{M},\mathcal{I}^{IPV})$ must have a symmetric equilibrium in strictly monotonic pure-strategies, which we denote by β. From Proposition 2, β is also an equilibrium of the game $(\mathcal{M},\mathcal{I}^*)$. This implies the result, since the revenue-equivalence theorem implies that

$$(6) \qquad R(\beta;\mathcal{M},\mathcal{I}^*) = R(\beta;\mathcal{M},\mathcal{I}^{IPV}) \leq R(\beta^{FPA};\mathcal{M}^{FPA},\mathcal{I}^{IPV}) = R^{FPA},$$

where the inequality follows from the fact that the first-price auction maximizes revenue among all efficient auctions when values are symmetric, private, and independent. In particular, the monotonic equilibrium of the first-price auction induces an efficient allocation and makes the participation constraint of the lowest type bind. ∎

Thus, the first-price auction has a greater revenue guarantee than second-price auction, all-pay auction, war of attrition, and all combinations of these mechanisms. Note that the English auction is not a standard mechanism in the sense described above, since in its normal form, messages are functions that indicate when a bidder would drop out of the auction, depending on the prices at which other bidders drop out. However, the revenue guarantee ranking can be easily extended to the English auction, since that auction is strategically identical to the second-price auction on $\mathcal{I}^*$, and there is an equilibrium in which bidders drop out when the price reaches their signal. We comment on this further in the next section.

While Theorem 1 only shows a weak ranking, in the case of the second-price and English auctions, the ranking is clearly strict, since these mechanisms have "bidding ring" equilibria with strategies different from β in which one bidder makes a high bid and the others effectively do not participate in the auction. Such equilibria would be undominated if we use a generalization of the Engelbrecht-Wiggans, Milgrom, and Weber (1983) construction: if bidder 1 sees the average value and the other bidders see nothing, then every bid in the support of the average is undominated.[6]

C. *Reserve Prices*

Theorem 1 pertains to revenue in standard mechanisms, which by definition must always allocate the good. It is well known that a seller may be able to increase revenue by rationing the good, e.g., by imposing a minimum bid r in the first-price auction. Under the additional hypothesis that values are common, we can extend Theorem 1 to standard mechanisms with rationing in the following manner. BBM derive the worst-case information structure $\mathcal{I}^*(r)$ for the first-price auction with a reserve price $r \in \mathbb{R}_+$. This information structure can be constructed from $\mathcal{I}^*$ by identifying a cutoff $\hat{s}(r)$, and letting bidders observe

$$s_i' = \begin{cases} s_i, & \text{if } s_i \geq \hat{s}(r); \\ r, & \text{otherwise.} \end{cases}$$

Thus, signals below the cutoff are pooled into a signal of r. The cutoff is chosen so that the expected value is r conditional on r being the highest signal. There is a monotonic pure-strategy equilibrium of the first-price auction under $\mathcal{I}^*(r)$ which attains the revenue guarantee, which is once again the monotonic equilibrium that would obtain if the bidders interpreted their signals as private values. We denote this information structure by $\mathcal{I}^{IPV}(r)$. The induced revenue guarantee $R^{FPA}(r)$ is the expected second highest signal from $\mathcal{I}^{IPV}(r)$.

Let us extend the notion of a standard mechanisms to those with one-dimensional bids, and such that conditional on the good being allocated, it is the high bidder who receives the good. We also assume that if the good is allocated at some bid profile, then it is also allocated at all weakly higher bid profiles. Importantly, however, the auction need not always allocate the good. Now, consider a standard mechanism $\mathcal{M}$ which induces the same allocation in $\mathcal{I}^{IPV}(r)$ as does the first-price auction with reserve price r. Examples of such mechanisms are the second-price or English auctions with the same reserve price. Then $\mathcal{M}$ has a revenue guarantee which is weakly less than $R^{FPA}(r)$.

The proof is nearly the same as before. Such a mechanism must be revenue equivalent to the first-price auction on $\mathcal{I}^{IPV}(r)$. The same logic as in Proposition 2 regarding the strategic equivalence implies that the equilibrium of $\mathcal{M}$ on $\mathcal{I}^{IPV}(r)$ would continue to be an equilibrium on $\mathcal{I}^*(r)$. Moreover, this equilibrium has revenue less than or equal to $R^{FPA}(r)$, so the revenue guarantee for $\mathcal{M}$ is weakly below that of the first-price auction. Thus, the first-price auction with a reserve price revenue-guarantee dominates a broad class of standard mechanisms that induce similar rationing.

[6] In personal communication, Ziwei Wang has given an example of an information structure and equilibrium in which revenue from the all-pay auction is strictly lower than R^{FPA}. Thus, the ranking with the all-pay auction is also strict.

IV. Revenue Guarantee Equivalence

A. *Equivalence with Common Values*

The revenue guarantee as defined in Section I is a tight lower bound on revenue in the first-price auction that holds across all information structures and all equilibria. As we mentioned previously, we could therefore have quite easily concluded that the second-price and English auctions have lower revenue guarantees than the first-price auction, since the former have "bidding ring" equilibria in which one bidder bids a large amount and the others bid zero. It is insightful to consider an alternative definition of the revenue guarantee, that restricts attention to information structures in which the second-price and English auctions have compelling and nontrivial equilibria. A natural class of information structures are those with affiliated values, as defined by MW. Our next result shows that even if we restrict attention to affiliated and common value information structures and if we select the monotonic pure-strategy equilibrium, the first-price auction still outperforms the second-price and English auctions. In fact, they all perform equally well.

Proceeding more formally, we say that a common value information structure $\mathcal{I}$ is *affiliated* if (i) $S_i = \mathbb{R}$ for all i, (ii) the marginal of μ on S is absolutely continuous with respect to Lebesgue measure and has a density $f(s)$, (iii) f is log supermodular (i.e., affiliated), and (iv) the induced interim expected value function $w^{\mathcal{I}}(s)$ is continuous and weakly increasing in all coordinates.

MW show that when the information structure is symmetric and affiliated, the first-price, second-price, and English auctions all have unique symmetric equilibria in strictly monotonic pure strategies. For the first-price auction, the equilibrium bidding function is

$$\beta^{FPA}(s_i) = E_{\mu}\left[w^{\mathcal{I}}(s_i, s_{-i}) \middle| s_{-i}^{(1)} \leq s_i\right].$$

We now formally define the second-price and English auctions. The second-price auction, denoted $\mathcal{M}^{SPA}$, has an allocation rule $q^{SPA} = q^{FPA}$ that is the same as that of the first-price auction, but the payment is the second-highest bid, i.e.,

$$t_i^{SPA}(m) = q_i^{SPA}(m)\, m^{(2)}.$$

MW show that the symmetric monotonic equilibrium of the second-price auction is

$$\beta^{SPA}(s_i) = E_{\mu}\left[w^{\mathcal{I}}(s_i, s_{-i}) \middle| s_{-i}^{(1)} = s_i\right].$$

The English auction, denoted $\mathcal{M}^{EA}$, is more complicated to define. It has messages that are collections of mappings $m_i^I : \mathbb{R}_+ \times \mathbb{R}_+^I \to \mathbb{R}_+$ for all $I \subseteq \mathcal{N}\backslash\{i\}$ that indicate, as a function of the prices at which bidders in I dropped out of the auction, at which price bidder i should drop out. A profile of messages induces an outcome wherein the first bidder i to drop is the one with the smallest $m_i^{\varnothing}(s_i)$, the second bidder j drops out at price $m_j^i(s_j, m_i^{\varnothing}(s_i))$, etc. The auction ends when only one bidder remains, and the remaining bidder gets the good and pays the price at which the penultimate bidder dropped out. We break ties uniformly if more than one bidder drops out simultaneously to end the auction. For a more formal description of the English auction, see MW. They show that there is an equilibrium of this game in which, conditional on the first K bidders dropping out at prices $y^{N-1} \leq \cdots \leq y^{N-1-K}$, a bidder with signal $s_i \geq y^{N-1-K}$ drops out at price

$$\beta^{EA}(s_i, y^{N-1}, \ldots, y^{N-1+K}) = E_\mu\left[w^{\mathcal{I}}(s, s_{-i})\middle| s_{-i}^{(k)} = y^k \; \forall k \geq N-1+K, \; s_{-i}^{(k)} = s_i \; \forall k \leq N-K\right].$$

For each $\mathcal{M} \in \{\mathcal{M}^{FPA}, \mathcal{M}^{SPA}, \mathcal{M}^{EA}\}$, we say that R is an *affiliated revenue guarantee for* $\mathcal{M}$ if for any symmetric affiliated values information structure $\mathcal{I}$ and symmetric strictly monotonic pure-strategy equilibrium β of $(\mathcal{M}, \mathcal{I})$, $R(\beta; \mathcal{M}, \mathcal{I}) \geq R$. As before, R is *the* affiliated revenue guarantee of $\mathcal{M}$ if it is an affiliated revenue guarantee and it is greater than any other affiliated revenue guarantee.

THEOREM 2 (Revenue Guarantee Equivalence): *The first-price*, *second-price*, *and English auctions all have the same affiliated revenue guarantee*, *which is equal to* R^{FPA}.

PROOF OF THEOREM 2:

The proof proceeds by two short steps.

Step 1: It is easy to verify that when values are common, $\mathcal{I}^*$ is such that signals are real-valued, there is a symmetric and independent signal distribution (which is a special case of affiliation), and the value function is $w^{\mathcal{I}^*}(s) = \max_i s_i$, which is continuous and monotonic. Thus, $\mathcal{I}^*$ is a symmetric and affiliated information structure. Moreover, the monotonic equilibria of the first-price, second-price, and English auctions on $\mathcal{I}^*$ all reduce the previously described equilibria on $\mathcal{I}^{IPV}$. All of these equilibria induce revenue equal to the expected second-highest signal, which is R^{FPA}. The lower bound from BBM still applies, so we conclude that this is the affiliated revenue guarantee of the first-price auction, and it continues to be an upper bound on the affiliated revenue guarantees of the second-price and English auctions.

Step 2: MW show that for any symmetric affiliated information structure, the first-price, second-price, and English auctions have unique symmetric strictly monotonic equilibria, which are described above, and that

$$R(\beta^{FPA}; \mathcal{M}^{FPA}, \mathcal{I}) \leq R(\beta^{SPA}; \mathcal{M}^{SPA}, \mathcal{I}) \leq R(\beta^{EA}; \mathcal{M}^{EA}, \mathcal{I}).$$

This implies that the affiliated revenue guarantee of the English auction must weakly exceed that of the second-price auction, which in turn weakly exceeds that of the first-price auction. The theorem follows from combining this with the opposite revenue guarantee ranking from Step 1. ∎

Thus, we find that when we restrict attention to well-behaved (symmetric and affiliated value) environments, the second-price and English auctions do no better than the first-price auction in the worst case. At the same time, if we relax these hypotheses (symmetry, affiliated values, favorable equilibrium selection), the worst-case for the first-price auction must remain the same, while for the other mechanisms it can only decrease.

In a sense, Theorem 2 shows that the information structure $\mathcal{I}^*$ has the strongest winner's curse of any affiliated values environment with the given distribution over the common value. By winner's curse, we mean the adverse selection from winning the good under a mechanism and equilibrium in which the high-signal bidder is allocated the good. It is well known that the presence of a winner's curse induces the bidders to shade their bids, so that they bid based on their pivotal value on the marginal event that they win. This updating is particularly severe in the maximum signal information structure $\mathcal{I}^*$: learning that one has the highest signal means that the value is *exactly* equal to one's own signal, whereas at the moment when the bidder only knows his own signal, it is only a *lower bound* on the true value of the object. Holding fixed the auction, a crude way of measuring the amount of adverse selection is the difference between expected value and expected revenue. Theorem 2 shows that this difference is the largest under the monotonic equilibrium of $\mathcal{I}^*$, regardless whether it is the first-price, second-price, or English auction that is used to measure the effect.

B. *Interdependent Values and Reserve Prices*

Theorem 1 is presented for symmetric interdependent values, whereas Theorem 2 is stated only for common values. Some subtleties arise in extending the latter theorem, due to the non-monotonicity of the conditional expectation $w_i^{\mathcal{I}^*}(s)$ at the information structure $\mathcal{I}^*$, so that it is not affiliated in the sense of MW. To see why, consider a model with two bidders and values are independent standard uniform random variables. Under $\mathcal{I}^*$, the highest signal is equal to the smallest value:

$$w_i(s) = \begin{cases} s_j, & \text{if } s_i < s_j; \\ \frac{1+s_i}{2}, & \text{otherwise.} \end{cases}$$

This function is discontinuous along the diagonal, when $s_i = s_j$.

This discontinuity gives rise to multiple monotonic pure-strategy equilibria in standard auctions: for the second-price auction, bidding s_i is an equilibrium, but so is bidding $(1+s_i)/2$. Under both strategy profiles, the high signal bidder wins and pays less than his expected value, so that downward deviations are not attractive. But on the marginal event where increasing one's bid changes the outcome, the bidder's expected value jumps down, so that the price paid is weakly greater than the expected value. One can similarly construct multiple equilibria of the first-price auction, with the one described in equation (5) being the lowest.

So, to generalize Theorem 2 to interdependent values, we have to both expand the set of information structures and specify which equilibria we compare. This can be done as follows. An information structure has *generalized affiliated values* if signals are affiliated and the value function is of the form

$$w_i^{\mathcal{I}}(s) = \hat{w}_i(s) + \mathbf{1}_{s_i \geq \max_{j \neq i} s_j} \tilde{w}_i(s)$$

where $\hat{w}_i$ is monotonic and $\tilde{w}_i$ is non-negative. It is easily verified that when information is symmetric and generalized affiliated, the first-price, second-price, and English auctions all have monotonic pure-strategy equilibria in which bidders act "as if" the value function were $\hat{w}_i$. We refer to this as the *minimal monotonic equilibrium*. Moreover, the linkage principle of MW applies to these equilibria, so that English auction generates more revenue than second-price auction generates more revenue than first-price auction. The *generalized affiliated revenue guarantee* is minimum revenue in the minimal monotonic equilibrium across all generalized affiliated values information structures that are consistent with the given prior. The proof of Theorem 2 then applies to this setting, and we conclude that the first-price, second-price, and English auctions are all generalized affiliated revenue guarantee equivalent.

Theorem 2 can be extended to cover auctions with reserve prices as well when values are common. The information structure $\mathcal{I}^*(r)$ described in Section IIIC is one of affiliated values, and the revenue ranking of MW in affiliated environments extends to auctions with reserve prices. Moreover, the equilibrium $\beta^{FPA}(\cdot\,;\,r)$ coincides with the one described by MW. Thus, the first-price, second-price, and English auctions with reserve price r are revenue guarantee equivalent, with a guarantee of $R^{FPA}(r)$.

V. Conclusion

We presented a novel version of the revenue ranking and revenue equivalence theorems. We compared auction formats in terms of their revenue guarantees across all information environments, rather than in terms of revenue in a specific information environment. This analysis yields a powerful new argument in favor of the first-price auction as achieving a greater revenue guarantee than other standard mechanisms, such as the second-price and English auctions.

Importantly, the dominance of the first-price auction among standard auctions depends on the possibility that bidders are uncertain about their own valuations. If we were to assume that bidders at least know their own values, then there is a simple argument that the second-price auction should revenue guarantee dominate the first-price auction. For if bidders know their values, the second-price auction has a unique equilibrium in weakly undominated strategies, in which a high-value bidder wins the auction and pays the second-highest value. On the other hand, if values were complete information among the bidders, the same ex post outcome would obtain in the first-price auction: a high-value bidder would win the auction, but the sale price would be competed up to the second-highest value. The private-value revenue guarantee of the first-price auction must therefore be weakly less than revenue in the undominated equilibrium of the second-price auction. Moreover, Fang and Morris (2006) and BBM gave examples of private-value information structures with binary values in which the revenue from an first-price auction is strictly lower than that of the second-price auction. This suggests that one could always find an information structure that with positive probability gives the bidders common knowledge that values are in a binary value subset with information of the form constructed by Fang and Morris. In such an information structure, the first-price auction would generate strictly lower revenue than the weakly undominated equilibrium of the second-price auction.

Thus, which mechanisms have better revenue guarantees depends on what assumptions we are willing to make about information and behavior. Indeed, the revenue rankings of Maskin and Riley (2000) and MW can be viewed as a limiting case where a single information structure and equilibrium is considered. A promising direction for future research is to identify classes of information structures and corresponding revenue guarantee rankings that interpolate between our findings and the classical results.

REFERENCES

Abraham, Ittai, Susan Athey, Moshe Babaioff, and Michael D. Grubb. 2014. "Peaches, Lemons, and Cookies: Designing Auction Markets with Dispersed Information." Unpublished.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2016. "Informationally Robust Auction Design." Cowles Foundation Discussion Paper 2065.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2017. "First-Price Auctions with General Information Structures: Implications for Bidding and Revenue." *Econometrica* 85 (1): 107–43.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2018. "Countering the Winner's Curse: Optimal Auction Design in a Common Value Model." Cowles Foundation Discussion Paper 2147.

Brooks, Benjamin, and Songzi Du. 2018. "Optimal Auction Design with Common Values: An Informationally-Robust Approach." Unpublished.

Bulow, Jeremy, and Paul Klemperer. 2002. "Prices and the Winner's Curse." *RAND Journal of Economics* 33 (1): 1–21.

Chung, Kim-Sau, and J. C. Ely. 2007. "Foundations of Dominant-Strategy Mechanisms." *Review of Economic Studies* 74 (2): 447–76.

Engelbrecht-Wiggans, Richard, Paul R. Milgrom, and Robert J. Weber. 1983. "Competitive Bidding and Proprietary Information." *Journal of Mathematical Economics* 11 (2): 161–69.

Fang, Hanming, and Stephen Morris. 2006. "Multidimensional Private Value Auctions." *Journal of Economic Theory* 126 (1): 1–30.

Krishna, Vijay, and John Morgan. 1997. "An Analysis of the War of Attrition and the All-Pay Auction." *Journal of Economic Theory* 72 (2): 343–62.

Maskin, Eric, and John Riley. 2000. "Asymmetric Auctions." *Review of Economic Studies* 67 (3): 413–38.

Matthews, Steven. 1995. "A Technical Primer on Auction Theory I: Independent Private Values." Unpublished.

Milgrom, Paul R. 1981. "Rational Expectations, Information Acquisition, and Competitive Bidding." *Econometrica* 49 (4): 921–43.

Milgrom, Paul, and Joshua Mollner. 2018. "Equilibrium Selection in Auctions and High Stakes Games." *Econometrica* 86 (1): 219–61.

Milgrom, Paul, and Robert J. Weber. 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50 (5): 1089–1122.

Myerson, Roger B. 1981. "Optimal Auction Design." *Mathematics of Operations Research* 6 (1): 58–73.

Rothkopf, Michael H., Thomas J. Teisberg, and Edward P. Kahn. 1990. "Why Are Vickrey Auctions Rare?" *Journal of Political Economy* 98 (1): 94–109.

Vickrey, William. 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance* 16 (1): 8–37.