Theoretical Economics 1 (2006), 95–118 1555-7561/20060095

Revenue comparisons for auctions when bidders have arbitrary types

YEON-KOO CHE
Departments of Economics, Columbia University and University of Wisconsin–Madison

IAN GALE
Department of Economics, Georgetown University

This paper develops a methodology for characterizing expected revenue from auctions when bidders' types come from an arbitrary distribution. In particular, types may be multidimensional, and there may be mass points in the distribution. One application extends existing revenue equivalence results. Another application shows that first-price auctions yield higher expected revenue than second-price auctions when bidders are risk averse and face financial constraints. This revenue ranking extends to risk-averse bidders with general forms of non-expected utility preferences.

KEYWORDS. Auctions, multidimensional types and atoms, risk aversion, Gateaux differentiable preferences.

JEL CLASSIFICATION. C70, D44.

1. INTRODUCTION

This paper develops a methodology for characterizing expected revenue from auctions in which bidders' types come from an arbitrary distribution. In particular, types may be multidimensional, and there may be mass points in the distribution. Accommodating multidimensional types is valuable because actual bidders may differ along many dimensions such as their risk attitudes and aspects of the financial constraints they face (e.g., cash holdings, sizes of credit lines, and terms of credit). Likewise, atoms may be relevant if there is positive probability that bidders will not participate, for example.

Most auction models assume a one-dimensional type space with an atomless distribution. The well-known revenue-equivalence results concern risk-neutral bidders

Yeon-Koo Che: yc2271@columbia.edu
Ian Gale: galei@georgetown.edu

The authors thank Eddie Dekel, Jeff Ely, Jinwoo Kim, Jon Levin, Bart Lipman, Mark Machina, Eric Maskin, Paul Milgrom, Marcin Pęski, Michael Peters, Larry Samuelson, Bill Sandholm, Andy Skrzypacz, Jeroen Swinkels, Joel Watson, Bob Weber, Bob Wilson, Huseyin Yildirim and seminar participants at Northwestern University, Stanford University, the University of British Columbia, the University of Virginia, the 2003 International Industrial Organization Conference, the 2003 North American Summer Meetings of the Econometric Society, and the 2004 Stony Brook Workshop on Mechanism Design with Budget Constraints for helpful comments.

Copyright © 2006 Yeon-Koo Che and Ian Gale. Licensed under the Creative Commons Attribution-NonCommercial License 2.5. Available at http://econtheory.org.

who differ only in their valuations of the good (see Myerson 1981 or Riley and Samuelson 1981, for example). Likewise, a typical model with risk aversion assumes that bidders have the same von Neumann-Morgenstern utility function, so they again differ only in their valuations (see Holt 1980, Matthews 1983, 1987 and Riley and Samuelson 1981). When bidders are completely ordered by their valuations, which are drawn from the same distribution, standard auctions yield an efficient allocation in equilibrium.[1] Hence, revenues from these auctions can be compared easily if the rents accruing to each valuation type can be compared. While this approach works in standard models with risk-neutral or risk-averse bidders, such a comparison does not work if the auctions entail different equilibrium allocations.

Suppose that bidders differ along multiple dimensions—their valuations and risk attitudes, say. If risk attitudes do not affect bidding behavior in a second-price auction, but do in a first-price auction, then the equilibrium allocations differ, making the existing methodology inapplicable. Similarly, when there are atoms in the distribution, the standard revenue equivalence argument is difficult to apply; in many cases, it does not apply.

We develop a method for characterizing equilibrium revenue in such cases. To illustrate, fix an auction form (a first- or second-price sealed-bid auction, say) with $n \geq 2$ bidders, and suppose that a symmetric equilibrium exists. Now imagine a *fictitious* risk-neutral bidder with valuation v and no financial constraints; she is henceforth referred to as the *benchmark bidder* or a *type-v bidder*. Suppose that the benchmark bidder were to participate in an auction with $n-1$ actual bidders who each employ the equilibrium bidding strategy. Now assign to each equilibrium bid a benchmark type having that bid as a best response (or some nearby type if no such type exists). This generates a cumulative distribution function (cdf), F_M, of benchmark types whose best responses mimic the actual equilibrium bidding behavior. We show that the revenue from the actual equilibrium is no less than the revenue generated when n risk-neutral bidders with valuations drawn from F_M play the same auction game. This lower bound for the revenue from the actual auction can be calculated from F_M, using a standard envelope theorem argument. An exact representation is available if two additional conditions are satisfied. Ultimately, our methodology reduces the task of comparing revenues from alternative auctions to that of comparing induced distributions of benchmark types.[2]

We present two applications of the methodology. First, we establish revenue equivalence for standard auctions when bidders are risk-neutral and face no financial constraints. In particular, our method establishes revenue equivalence for discrete types in a much broader class of auction forms than has been shown previously.[3] (Our re-

[1] Incentive compatibility makes equilibrium bidding strategies monotonic in valuations in standard auctions, which means that the allocation is efficient in all such auctions.

[2] The current method can be seen as aggregating the arbitrary type into a one-dimensional type. This aggregation method differs from other methods such as the one used for analyzing score-based auctions (see Che 1993 and Asker and Cantillon 2003). This latter method applies to quasilinear preferences; the current method applies to general preferences.

[3] Maskin and Riley (1985) and Riley (1989) demonstrate that revenue equivalence holds between sealed-bid auctions and open oral auctions when types are discrete.

sults apply to continuous and mixed distributions as well.) In the process, we identify the properties of auction forms that produce revenue equivalence. Second, we apply the methodology to generalize the result that a first-price sealed-bid auction generates higher expected revenue than does a second-price sealed-bid auction when bidders are risk-averse expected-utility maximizers and face financial constraints. The results here allow for arbitrary heterogeneity in both dimensions. The third application shows that the results hold for a broad class of non-expected utility preferences. In particular, they are shown using the Gateaux differentiable preference functional, which includes essentially all (possibly non-expected utility) preference functionals satisfying a minimal smoothness requirement.

The current model encompasses virtually all existing models that incorporate risk aversion or financial constraints. In particular, it significantly generalizes Che and Gale (1998), which considered risk-neutral bidders with private information about their valuations and (one-dimensional) financial constraints.[4] In addition to limiting attention to two-dimensional private information, the earlier paper relied crucially on the assumption that, for every equilibrium bid, there was an unconstrained type that would make that bid.[5] No such assumption is needed here.

The remainder of the paper is organized as follows. Section 2 contains the revenue characterization for general auction forms and general payoffs. Section 3 presents the three applications mentioned above. Section 4 concludes.

2. Revenue characterization

A seller is holding an auction for a single object. There are $n \geq 2$ bidders whose types (i.e., preferences and constraints) are independently and identically distributed. We make no additional assumptions about bidders' types at this point. Instead, we simply assume that the auction has a symmetric Bayesian-Nash equilibrium that yields a finite expected revenue for the seller. We study a class of auction forms that satisfy some natural conditions: every bidder makes a single bid in $\mathbb{R}_+$, the high bid wins, bidders' payments are functions of the bids, and bidders are treated symmetrically. These conditions are satisfied by first- and second-price sealed-bid auctions, as well as all-pay auctions and wars of attrition, among others.

Fix a symmetric equilibrium, which we denote "M."[6] Let B_M be the random variable

[4]Che and Gale also assumed that the cost function was submodular in the bidder's payment and budget parameter. In addition, they assumed that the distribution of types was continuous and the support of equilibrium bids had no mass points or gaps. The current paper uses a different approach, which does not require these features.

[5]This assumption is not satisfied if some equilibrium bids are so unattractive that only financially constrained types would choose them. Fang and Perreiras (2001) have shown that this possibility cannot be avoided in certain cases. They considered bidders facing absolute financial constraints, with the infimum budget strictly larger than the infimum valuation.

[6]Our revenue characterization in this section does not even require M to be an equilibrium. It only requires the profile of strategies to be symmetric. We assume an equilibrium since all subsequent applications will indeed require bidders to play their equilibrium strategies.

representing the bid made by an individual bidder in that equilibrium, and let

$$\mathscr{B}_M := \{b \in \mathbb{R} \mid \Pr\{B_M \in [b, b+\epsilon)\} > 0 \text{ and } \Pr\{B_M \in (b-\epsilon, b]\} > 0 \,\forall \epsilon > 0\}$$

be its (measurable) support.[7]

Let $x_M(b)$ and $t_M(b)$ denote an individual bidder's probability of winning and his expected payment, respectively, if he bids $b \in \mathbb{R}_+$ and all $n-1$ others employ the equilibrium strategy.

Now imagine a benchmark bidder (i.e., a risk-neutral bidder who faces no financial constraints) bidding against $n-1$ actual bidders who each employ the equilibrium strategy. We will construct a distribution of benchmark types such that the resulting distribution of best responses mimics the equilibrium bid distribution, B_M. We then characterize the seller's expected revenue using this constructed distribution of benchmark types.

Suppose that a benchmark bidder with valuation $v \in \mathscr{V} := [0, \infty)$ were to bid $b \in \mathbb{R}_+$. She would receive an expected payoff of $\pi_M(b, v) := v x_M(b) - t_M(b)$. The supremum payoff for the type-v benchmark bidder is

$$\Pi_M(v) := \sup_{b \in \mathbb{R}_+} \pi_M(b, v). \tag{1}$$

Let

$$BR_M(v) := \arg\max_{b' \in \mathbb{R}_+} \pi_M(b', v)$$

denote the set of best responses, which may be empty for a given v. Now let $(X_M(v), T_M(v))$ be a limit point of $(x_M(b), t_M(b))$ along a sequence of b that yields $\Pi_M(v)$ in the limit.[8] We then have $\Pi_M(v) = v X_M(v) - T_M(v)$.

For each $v \in \mathscr{V}$, let $\beta_M(v) := BR_M(v) \cap \mathscr{B}_M$ denote the set of best responses that are also equilibrium bids. We first show that β_M is a monotonic correspondence.

LEMMA 1. *Suppose that $b' \in \beta_M(v')$ and $b \in \beta_M(v)$ for v, $v' \in \mathscr{V}$ with $v' > v$. Then, $b' \geq b$.*

PROOF. Incentive compatibility implies that a type-v benchmark bidder weakly prefers b to b', while a type v' does the reverse. Combining these two conditions yields

$$(v' - v)[x_M(b') - x_M(b)] \geq 0.$$

Since $v' > v$, we immediately have $x_M(b') \geq x_M(b)$. Now suppose that $b' < b$. Since b', $b \in \mathscr{B}_M$, we have $x_M(b') < x_M(b)$, which is a contradiction.[9] We conclude that $b' \geq b$. □

[7] This definition of the support differs from other possible definitions only for measure-zero sets; it simplifies the proofs of Lemmas 1 and 2.

[8] The limit point, $(X_M(v), T_M(v))$, is well defined. Let $\{b^n\}_{n=1}^{\infty}$ be a sequence such that $\Pi_M(v) = \lim_{n\to\infty}\{v x_M(b^n) - t_M(b^n)\}$. Then, since $x(b^n)$ lies in $[0,1]$, a compact set, there exists a subsequence $\{b^{k_n}\}_{n=1}^{\infty}$ such that $x_M(b^{k_n})$ converges to $X_M(v)$, say, as $n \to \infty$. Then, $t_M(b^{k_n})$ must converge to $v X_M(v) - \Pi_M(v) =: T_M(v)$ as $n \to \infty$.

[9] This follows from the definition of $\mathscr{B}_M$.

We next construct a random variable, V_M, representing the benchmark bidder's type. This is done in such a way that the resulting distribution of best responses mimics B_M. The first step is to define a function, $\phi_M : \mathscr{B}_M \mapsto \mathscr{V}$, mapping equilibrium bids into benchmark types. For each $b \in \beta_M(\mathscr{V})$,[10] let

$$\phi_M(b) := v \text{ such that } b \in \beta_M(v).$$

(If there are multiple candidates for a given b, select one of them.[11]) For each $b \in \mathscr{B}_M \setminus \beta_M(\mathscr{V})$, let

$$\phi_M(b) := \begin{cases} \inf\{\phi_M(b') \mid b' \in \beta_M(\mathscr{V}) \cap (b,\infty)\} & \text{if } \beta_M(\mathscr{V}) \cap (b,\infty) \neq \emptyset \\ \sup\{\phi_M(b') \mid b' \in \beta_M(\mathscr{V}) \cap [0,b)\} & \text{if } \beta_M(\mathscr{V}) \cap (b,\infty) = \emptyset. \end{cases}$$

In words, ϕ_M assigns to each equilibrium bid a benchmark type having b as a best response, if such a type exists; to any remaining equilibrium bid it assigns the infimum type with a best response exceeding b in $\mathscr{B}_M$ (or the supremum type with a best response less than b, if none exists). This mapping is well defined when the former types exist (i.e., $\beta_M(\mathscr{V}) \neq \emptyset$); existence is verified in the applications below. Lemma 1 implies that ϕ_M is nondecreasing. Hence, the inverse correspondence, $\phi_M^{-1}(v) := \{b \in \mathscr{B}_M \mid \phi_M(b) = v\}$, is strictly increasing in v over its range, $\mathscr{V}_M := \phi_M(\mathscr{B}_M)$.

The assigned benchmark type, $V_M := \phi_M(B_M)$, is then distributed according to the cdf

$$F_M(v) := \Pr\{\phi_M(B_M) \leq v\}. \tag{2}$$

Clearly, F_M is nondecreasing and right-continuous. The range of ϕ_M, $\mathscr{V}_M$, is the support of F_M.

Suppose, hypothetically, that n risk-neutral bidders were to draw valuations according to F_M and bid according to ϕ_M^{-1}. (In particular, a bidder with valuation v would bid in a way that matches the conditional distribution of B_M in $\phi_M^{-1}(v)$.) This scenario would replicate the equilibrium bidding behavior in the original equilibrium, M, and would yield the same revenue. We can therefore focus on the revenue generated in this scenario.

Now fix a type, $v \in \mathscr{V}_M$. For each bid $b \in \phi_M^{-1}(v)$,

$$\Pi_M(v) \geq v x_M(b) - t_M(b), \tag{3}$$

since the bid need not be a best response. If $b \in \beta_M(v)$, however, it is a best response for a type v, so (3) becomes an equality. In fact, (3) holds with equality for all $v \in \mathscr{V}_M$ if the set of best responses contains every equilibrium bid:

CONDITION (A1). $\mathscr{B}_M = \beta_M(\mathscr{V})$.

[10] Throughout, a function or correspondence defined over a set connotes the range of the function or correspondence over that set. For instance, $\beta_M(\mathscr{V}) := \bigcup_{v' \in \mathscr{V}} \beta_M(v')$.

[11] We construct ϕ_M precisely in subsequent applications.

When this condition holds, every $b \in \phi_M^{-1}(v)$ is a best response for v, for all $v \in \mathscr{V}_M$ (by construction of ϕ_M), ensuring that (3) holds with equality for all $v \in \mathscr{V}_M$. Condition (A1) is satisfied in a second-price auction since any bid b is a best response for a benchmark bidder of type $v = b$. When Condition (A1) holds, we are able to get an *exact* revenue representation. If it does not hold, we get a lower bound on revenue, based on (3).

We now characterize $\Pi_M(v)$, using $F_M(v)$. By definition,

$$\Pi_M(v) = vX_M(v) - T_M(v) \geq vX_M(v') - T_M(v')$$

for all $v' \in \mathscr{V}$. An envelope theorem argument (see Theorem 2 of Milgrom and Segal 2002) then yields

$$\Pi_M(v) = \Pi_M(0) + \int_0^v X_M(s)ds. \tag{4}$$

The supremum payoff can be characterized in terms of F_M if an additional condition is satisfied:

CONDITION (A2). *$X_M(v) = F_M(v)^{n-1}$ for almost every $v \in \mathscr{V}$.*

Given this condition, the probability that a type-v benchmark bidder wins is equal to the probability that an actual bidder bids weakly less than a type v's best response. The possibility of mass points in the equilibrium bid, B_M, and the associated (random) tie-breaking, make this condition nontrivial. Even with mass points, however, (A2) holds in equilibria of first- and second-price auctions, but it may not hold in other auctions.[12] When (A2) holds, the expected payoff in (4) takes the form seen in symmetric independent private values (IPV) auctions.

Let $V_M^{(i)}$ denote the ith order statistic of n random variables with cdf F_M (i.e., the ith highest of n realizations of V_M). We are now able to characterize the expected revenue.

THEOREM 1. *Suppose that (A2) holds. The seller's expected revenue from auction equilibrium M is greater than or equal to $\mathbb{E}[V_M^{(2)}] - n\Pi_M(0)$. If (A1) also holds, the expected revenue equals $\mathbb{E}[V_M^{(2)}] - n\Pi_M(0)$.*

PROOF. The seller's expected revenue from auction equilibrium M is given by:

$$\begin{aligned}
n\mathbb{E}_{B_M}[t_M(B_M)] &= n\mathbb{E}_{V_M}\left[\mathbb{E}_{B_M}\left[t_M(B_M) \,\middle|\, B_M \in \phi_M^{-1}(V_M)\right]\right] \\
&\geq n\mathbb{E}_{V_M}\left[\mathbb{E}_{B_M}\left[V_M x_M(B_M) - \Pi_M(V_M) \,\middle|\, B_M \in \phi_M^{-1}(V_M)\right]\right] \\
&= n\mathbb{E}_{V_M}\left[\mathbb{E}_{B_M}\left[V_M x_M(B_M) \,\middle|\, B_M \in \phi_M^{-1}(V_M)\right]\right] - n\mathbb{E}_{V_M}[\Pi_M(V_M)] \\
&= \mathbb{E}_{V_M}\left[V_M^{(1)}\right] - n\mathbb{E}_{V_M}[\Pi_M(V_M)] \\
&= \mathbb{E}_{V_M}\left[V_M^{(1)}\right] - n\mathbb{E}_{V_M}\left[\int_0^{V_M} F_M(v)^{n-1}dv\right] - n\Pi_M(0) \\
&= \mathbb{E}_{V_M}\left[V_M^{(2)}\right] - n\Pi_M(0).
\end{aligned} \tag{5}$$

[12] An example in Section 3.1 shows that it may not be satisfied in a third-price auction.

The first equality follows from the equivalence of the bids generated by the original equilibrium, M, and the bids generated in the scenario in which n risk-neutral bidders draw valuations, V_M, and bid according to $\phi_M^{-1}(V_M)$. The inequality follows from (3). The third equality follows since $n\mathbb{E}_{V_M}\left[\mathbb{E}_{B_M}\left[V_M x_M(B_M)\,\middle|\, B_M \in \phi_M^{-1}(V_M)\right]\right]$ equals the aggregate gross surplus accruing to the n risk-neutral bidders in the hypothetical scenario. Since the correspondence $\phi_M^{-1}(\cdot)$ is strictly increasing, a bidder with a higher valuation bids strictly higher than a bidder with a lower valuation in that scenario. Hence, the good is allocated efficiently among the n risk-neutral bidders, implying that the gross surplus equals the expectation of the first order statistic of V_M. The second-to-last equality follows from (4) and (A2). The last equality follows from integration by parts.

The second statement holds since (A1) implies that, for each $v \in \mathscr{V}_M$, $b \in \phi^{-1}(v)$ means $b \in \beta_M(v)$, so the inequality in (3) is an equality for all $v \in \mathscr{V}_M$, making the inequality in (5) an equality. □

3. Revenue comparisons of auctions

This section uses Theorem 1 to compare the expected revenues from different auction forms. We first impose some structure on bidders' types and preferences, the features that ultimately generate the random variable B_M. Suppose that each bidder i has a type, θ_i, drawn from an arbitrary, compact, non-empty support, Θ. Types are independently and identically distributed across bidders, and each bidder's preferences depend only on his type, which is his private information. The next two subsections compare expected revenue across standard auctions. First, we consider risk-neutral bidders without financial constraints and provide a generalized revenue-equivalence result. We then show that first-price auctions yield greater expected revenue than second-price auctions when risk aversion and financial constraints are present.

3.1 *Risk neutral bidders without financial constraints*

Many auction forms yield the same expected revenue when bidders are risk neutral and ex ante identical. Revenue equivalence results in the IPV setting typically depend on assumptions such as connectedness or absolute continuity of the distribution of types (see Myerson 1981 or Riley and Samuelson 1981, for example). Maskin and Riley (1985) and Riley (1989) extend the revenue equivalence between first-price sealed-bid auctions and oral ascending (or second-price sealed-bid) auctions to discrete types. Theorem 1 enables us to generalize those results for arbitrary type distributions. In the process, we identify features that make revenue equivalence possible.

Let a bidder of type $\theta \in \Theta$ have a valuation θ.[13] A bidder's valuation has a nondecreasing and right-continuous cdf, $F : \Theta \mapsto [0,1]$, which may have mass points and gaps. As above, a bid b wins with some probability $x_M(b)$ in equilibrium M and entails an expected payment $t_M(b)$. Given risk neutrality and no financial constraints, a bidder with valuation θ receives an expected payoff of $\theta x_M(b) - t_M(b)$ when bidding b.

[13] Without loss of generality we assume $\Theta \subset \mathbb{R}_+$ here. If actual valuations are a function of multiple components, what ultimately matters is just the value of that function.

We again consider auctions in which the high bid wins and the bidders are treated symmetrically (ties are broken randomly). In addition, we assume that a bidder's payment depends only on his own bid and the highest competing bid. Formally, bidder i's payment is $\tau_w(b_i, b_{m(i)}) \in \mathbb{R}_+$ if he wins and $\tau_l(b_i, b_{m(i)}) \in \mathbb{R}_+$ if he loses, where b_i and $b_{m(i)} := \max_{j \neq i} b_j$ denote bidder i's bid and the highest competing bid, respectively.[14] (Since the high bidder wins, τ_w is defined for $b_i \geq b_{m(i)}$ and τ_l is defined for $b_i \leq b_{m(i)}$.) An auction form that satisfies these conditions is called a *standard auction*. A standard auction displays *continuous payments* if the following conditions hold:

$$\tau_w(0,0) = \tau_l(0,\cdot) = 0 \text{ and } \tau_k(\cdot, b_{m(i)}) \text{ is continuous for } k = w, l, \text{ in the relevant domain.}$$

Many familiar auction forms have all of these features: In a first-price auction, the winner pays $\tau_w(b_i, b_{m(i)}) = b_i$ and a loser pays $\tau_l(b_i, b_{m(i)}) = 0$; in a second-price auction, $\tau_w(b_i, b_{m(i)}) = b_{m(i)}$ and $\tau_l(b_i, b_{m(i)}) = 0$; in an all-pay auction, $\tau_w(b_i, b_{m(i)}) = b_i$ and $\tau_l(b_i, b_{m(i)}) = b_i$; and in a war of attrition, $\tau_w(b_i, b_{m(i)}) = b_{m(i)}$ and $\tau_l(b_i, b_{m(i)}) = b_i$. Many other auctions forms are allowed. For instance, nothing in the definition precludes non-monotonic portions in the payment functions.

The restriction to standard auctions is appealing, but it does preclude mechanisms such as third-price auctions.[15] The role of the various conditions is made precise later.

We now demonstrate revenue equivalence for symmetric equilibria of standard auctions with continuous payments.[16] A preliminary result enumerates some useful properties of equilibria. Let $\delta_M(b) := \tau_w(b,b) - \tau_l(b,b)$ denote the difference between what a winner and a loser pay when tying with a high bid of b.

LEMMA 2. *Suppose that the bidders are risk neutral and face no financial constraints. A symmetric equilibrium in a standard auction with continuous payments has the following properties. (a) If B_M has a mass point at $b \in \mathscr{B}_M$, and if $b \in \beta_M(v)$ for some $v \in \mathscr{V}$, then $v = \delta_M(b)$ and $v \in \Theta$. (b) For any $v, v' \in \Theta$ with $v < v'$, if $b \in \beta_M(v)$ and $b' \in \beta_M(v')$, then $b < b'$.*

The proofs of this result and of several subsequent results are in the Appendix.

The second part of the lemma means that a symmetric equilibrium of a standard auction with continuous payments admits an efficient allocation. While efficiency of IPV auctions is a familiar result, the result here is significant because the class of auction forms considered is broad, and we allow for atoms in the distribution of types. The property concerning mass points is crucial for efficiency and revenue equivalence. (Example 2 below shows that efficiency is not guaranteed in a third-price auction—which is not a

[14] The need to define two payment functions arises because ties may occur with non-zero probability, and the bidder's payment may depend on whether she wins or loses.

[15] The allowed payment functions also exclude strictly positive entry fees and reserve prices, but these exclusions are more innocuous. The analysis can be extended to incorporate these features since the equilibrium with a reserve price is observationally equivalent to our model with a particular cdf.

[16] Again, we assume existence of a symmetric equilibrium. Some auction forms with continuous payment functions may fail to admit an equilibrium. For instance, the degenerate case of $\tau_w = \tau_l := 0$ produces unbounded bids.

standard auction—if there are mass points.) Let $\theta^{(2)}$ denote the second order statistic of n random variables with cdf F. We now present the revenue equivalence result.

PROPOSITION 1. *Suppose that the bidders are risk neutral and face no financial constraints. A symmetric equilibrium of a standard auction with continuous payments yields expected revenue of* $\mathbb{E}[\theta^{(2)}]$.

PROOF. Fix a symmetric equilibrium. We first show that $\Pi_M(0) = 0$. The continuous payments property implies that a bidder of type $v = 0$ can get a payoff of zero by bidding zero. The payoff cannot be strictly positive, however, since payments are nonnegative, so $\Pi_M(0) = 0$.

The next step is to pin down F_M. To that end, we first construct ϕ_M. Condition (A1) holds since $\beta_M(\Theta) = \mathscr{B}_M$. Hence, for each $b \in \mathscr{B}_M$, we can set $\phi_M(b) = v \in \Theta$ for v such that $b \in \beta_M(v)$. Such a v is unique since, by Lemma 2(b), any selection from $\beta_M(v)$ is strictly increasing in v for $v \in \Theta$. Since an actual bidder with $\theta \in \Theta$ chooses a bid in $\beta_M(\theta)$ in equilibrium, and since ϕ_M picks $v \in \Theta$ for $b \in \beta_M(\theta)$, and such assignment is unique, we must have $\theta = \phi_M(B_M)$. Consequently,

$$F_M(v) = \Pr\{\phi_M(B_M) \le v\} = \Pr\{\theta \le v\} = F(v) \tag{6}$$

for every $v \in \mathscr{V}$.

The last step is to show that (A2) holds. By Lemma 2(b), any $v \in \Theta$ wins with probability $F(v)^{n-1}$ unless F jumps at v. The set of valuations in Θ with mass points has measure zero. Since $F(v)^{n-1} = F_M(v)^{n-1}$, we have $X_M(v) = F_M(v)^{n-1}$ for almost every $v \in \Theta$. Hence, it now suffices to show that $X_M(v) = F_M(v)^{n-1}$ for each $v \in \mathscr{V} \backslash \Theta$.

Fix $v \in \mathscr{V} \backslash \Theta$. Either $X_M(v) = x_M(b)$ for some mass point $b \in \mathscr{B}_M$, or $X_M(v) \in \{F(\tilde{v}_-)^{n-1}, F(\tilde{v})^{n-1}\}$ for some $\tilde{v} \in \Theta$, where $F(\tilde{v}_-)$ denotes the left-hand limit of F at $\tilde{v}$. The former cannot be true; otherwise, Lemma 2(a) would imply $v \in \Theta$, which contradicts $v \in \mathscr{V} \backslash \Theta$. Hence, we conclude that $X_M(v) \in \{F(\tilde{v}_-)^{n-1}, F(\tilde{v})^{n-1}\}$ for some $\tilde{v} \in \Theta$.

It remains to show that $X_M(v) = F_M(v)^{n-1}$. This requires a preliminary step. Consider an arbitrary $v' \in \Theta$ and some $b' \in \beta_M(v')$. Incentive compatibility for a type v means

$$v X_M(v) - T_M(v) = \sup_{b \in \mathbb{R}_+} \pi_M(b, v) \ge \pi_M(b', v) = v x_M(b') - t_M(b'). \tag{7}$$

For the type v', it implies

$$v' x_M(b') - t_M(b') = \sup_{b \in \mathbb{R}_+} \pi_M(b, v') \ge v' X_M(v) - T_M(v). \tag{8}$$

Combining (7) and (8), we obtain

$$(v - v')[X_M(v) - x_M(b')] \ge 0. \tag{9}$$

Suppose that $X_M(v) \ne F_M(v)^{n-1}$. If $F_M(v)^{n-1} < X_M(v) \le F(\tilde{v})^{n-1}$, there exists $v' \in (v, \tilde{v}) \cap \Theta$, or else Θ has a mass point at $v' = \tilde{v}$ and $X_M(v) = F(\tilde{v})^{n-1}$. Either way, $x_M(b') <$

$X_M(v)$, for some $b' \in \beta_M(v')$.[17] We thus have a contradiction to (9). If $F_M(v)^{n-1} > X_M(v) \geq F(\tilde{v}_-)^{n-1}$, there must exist $v' \in (\tilde{v}, v) \cap \Theta$. Since $F(\tilde{v})^{n-1} < x_M(b')$ for all $b' \in \beta_M(v')$, given Lemma 2(b) and $X_M(v) \leq F(\tilde{v})^{n-1}$, we again have a contradiction to (9). We conclude that $X_M(v) = F(\tilde{v})^{n-1} = F_M(v)^{n-1}$, proving that (A2) holds.

Since (A1) and (A2) hold, $\Pi_M(0) = 0$, and $F_M(\cdot) = F(\cdot)$, Theorem 1 indicates that the expected revenue in a symmetric equilibrium of auction form M equals $\mathbb{E}[\theta^{(2)}]$. □

In the usual case with an atomless distribution, the equilibrium allocation pins down the rents for all types, up to a constant. The efficiency result in Lemma 2(b), along with the property that the infimum type receives a payoff of zero, then yields revenue equivalence.

It is important to note that equality of rents may not hold if the types are discrete. To see why, consider a two-point support, $\Theta = \{\theta^L, \theta^H\}$, and two auction equilibria, A and B. Then, (4) implies that an actual bidder of type θ^H receives rents equal to

$$\Pi_M(\theta^H) = \Pi_M(\theta^L) + \int_{\theta^L}^{\theta^H} X_M(s)\,ds$$

in $M = A, B$. Now suppose that $X_A(\theta) = X_B(\theta)$ for $\theta \in \{\theta^L, \theta^H\}$ (i.e., the equilibrium allocation is the same for the actual types) and $\Pi_A(\theta^L) = \Pi_B(\theta^L)$. The rents accruing to a type θ^H may differ across auction equilibria if $X_A(v) \neq X_B(v)$ for $v \in (\theta^L, \theta^H)$. In other words, the incentives of benchmark types that are not actual types affect the calculation of the equilibrium rents accruing to the actual types. One must therefore keep track of the incentives of all benchmark types in order to compare revenue, even though the actual types are discrete.

Revenue equivalence obtains in standard auctions with continuous types because benchmark types have the same incentives across auction forms. In particular, a benchmark type does not mimic a neighboring actual type that is a mass point of the distribution (Lemma 2); this non-mimicking behavior implies that $X_M(v) = F(v)^{n-1}$ for almost every v, as required by Condition (A2). In fact, the restriction to standard auctions with continuous payments is necessary for the revenue equivalence result, as is illustrated next.

EXAMPLE 1 (Discontinuous payments). Suppose that $n = 2$ bidders draw valuations from $\Theta = \{1, 2\}$ with probability $\frac{1}{2}$ each. In a second-price auction, the allocation is efficient since $\theta = 2$ outbids $\theta = 1$. Benchmark types with $v \in \{1, 2\}$ have best responses in the equilibrium support. Those with $v \in (1, 2)$ would strictly outbid an actual bidder with $\theta = 1$, implying $X_s(v) = \frac{1}{2}$ for these types. By Proposition 1, the seller receives expected revenue of $\mathbb{E}[\theta^{(2)}] = \frac{5}{4}$.

[17] In case of a mass point at $v' = \tilde{v}$, the probability of winning is strictly less than $F(\tilde{v}_+)^{n-1}$ if a type $\tilde{v}$ bids $b' \in \beta_M(\tilde{v})$ (either because a tie occurs at b' with positive probability or because a mixed strategy is adopted by the type $\tilde{v}$ in equilibrium), so the statement holds with $v' = \tilde{v}$. If there is not a mass point at $\tilde{v}$, there exists $v' \in (v, \tilde{v}) \cap \Theta$, and the statement follows from Lemma 2(b).

Now consider an optimal auction (denoted "$M = o$").[18] It has $\tau_l := 0$ and

$$\tau_w(b_i, b_{m(i)}) = \begin{cases} b_i & \text{if } b_i \leq 1 \\ \frac{5}{3} & \text{if } b_i \in (1,2] \\ b_i & \text{if } b_i > 2. \end{cases}$$

An actual bidder with $\theta = 1$ bids 1 in the symmetric equilibrium, while a bidder with $\theta = 2$ bids 2, so the allocation is efficient. But, the seller's expected revenue is $\frac{3}{2} > \frac{5}{4}$ here, so revenue equivalence fails even though the equilibrium allocation and the expected payoff to the lowest type ($\theta = 1$) are the same as in the second-price auction. This result follows from the failure of the continuous payments property, as the winner's payment jumps up at $b = 1$ and again at $b = 2$.[19] A benchmark bidder with valuation $v \in (1,2)$ would not wish to outbid an actual bidder with $\theta = 1$, so $X_o(v) = \frac{1}{4} < \frac{1}{2} = F_o(v)$ for $v \in (1,2)$, violating (A2). The rent accruing to an actual bidder with $\theta = 2$ is therefore smaller in the optimal auction, leading to higher revenue there. ◊

EXAMPLE 2 (Dependence on other bids). Suppose that there are $n = 3$ bidders with the same two-point type distribution as in Example 1. In a second-price auction, the second-highest bid is equally likely to be 1 or 2, so the seller's expected revenue is $\frac{3}{2}$. A benchmark type with $v \in (1,2)$ would submit a bid in $(1,2)$, so $X_s(v) = \frac{1}{4} = F(v)^2$.

Now consider a third-price auction (denoted "$M = t$"); payments clearly depend on bids other than the own bid and the highest competing bid. There is a continuum of equilibria indexed by $\gamma \in [3.5, 5]$; for γ in this interval, it is symmetric equilibrium behavior for each bidder to bid 1 if $\theta = 1$ and to bid γ if $\theta = 2$. The seller receives γ if all three bidders have $\theta = 2$, and 1 otherwise, so the expected revenue is $\frac{\gamma+7}{8}$. In particular, the equilibria with $\gamma < 5$ all yield expected revenue strictly less than $\frac{3}{2}$, so revenue equivalence fails.

The revenue nonequivalence is again explained by the incentives of the benchmark types in $(1,2)$. Consider the equilibrium in which an actual type $\theta = 2$ bids $x = 4$. In this equilibrium, benchmark types with $v \in (\frac{7}{4}, \frac{11}{5})$ would bid 4, just as the type $\theta = 2$ would.[20] Thus, $X_t(v) = \frac{7}{12} > \frac{1}{4} = F(v)^2$ for $v \in (\frac{7}{4}, 2)$, again violating (A2). This time there is more rent for the actual type $\theta = 2$—and lower expected revenue for the seller—than in the second-price auction.

Modifying this example also shows that efficiency is not guaranteed in a third-price auction. Suppose that three bidders each have a valuation drawn from $\Theta = \{1, 2-\epsilon, 2+\epsilon\}$ with probabilities $\frac{1}{2}$, $\frac{1}{4}$, and $\frac{1}{4}$, respectively. The above argument implies that there is a symmetric equilibrium in which actual bidders with $\theta = 2-\epsilon$ and $\theta = 2+\epsilon$ bid 4, for sufficiently small $\epsilon > 0$. Given random tie-breaking, this pooling produces an

[18] It is straightforward to confirm directly that this auction implements the optimal mechanism.

[19] While $x_o(b)$ jumps from $\frac{1}{4}$ to $\frac{1}{2}$ when b exceeds 1, τ_w also jumps, from 1 to $\frac{5}{3}$.

[20] The expected payoff is $\frac{1}{4}(v-1)$, $\frac{1}{12}(7v-10)$, or $v - \frac{7}{4}$ when bidding $b \in (1,4)$, $b = 4$, or $b \in (4,\infty)$, respectively. Benchmark types $v \in (1, \frac{7}{4})$, $v \in (\frac{7}{4}, \frac{11}{5})$, and $v \in (\frac{11}{5}, \infty)$, strictly prefer the first, second, and third alternatives, respectively.

inefficient allocation with positive probability, which confirms that Lemma 2(b) relies on the dependence on just the two bids. ◊

3.2 *Bidders with risk aversion and financial constraints*

We next compare the expected revenue from first- and second-price auctions when bidders' payoffs are strictly concave in the payments they make. Risk aversion and financial constraints constitute two possible sources of concavity. Risk aversion has long been considered an important determinant of bidder behavior in the theoretical and experimental literatures.[21] The importance of financial constraints, which arise when the marginal cost of expenditure is increasing, has been recognized in a growing literature.[22]

Suppose that each bidder i has a type drawn independently and identically from a nonempty, measurable set, Θ. A bidder of type θ gets von Neumann-Morgenstern utility of $u(x;\theta)$ if he wins the object and pays $x \in \mathbb{R}_+$. He receives utility of zero if he does not win. We make two assumptions concerning the utility function:

(U1) *For each $\theta \in \Theta$, $u(\cdot;\theta)$ is continuous, strictly decreasing and (weakly) concave.*

(U2) *For each $\theta \in \Theta$, $u(0;\theta) \geq 0$, with $u(0;\theta) > 0$ for a set with positive measure. Conversely, there exists $K > 0$ such that $u(K,\theta) < 0$ for all $\theta \in \Theta$.*

The concavity requirement of Assumption (U1) is consistent with risk aversion and financial constraints. Assumption (U2) ensures that some types have an incentive to participate and that the equilibrium bids are bounded. There may be atoms in preferences since $u(\cdot,\theta)$ could be constant over an interval in Θ, or there may simply be atoms in Θ; our model is general enough to accommodate many scenarios.

Bidders' financial constraints fit easily into our model. Suppose $u(x;\theta) = \theta_1 - c(x,\theta)$, where θ_1 is the valuation and $c(\cdot,\theta)$ is a strictly increasing and convex *cost-of-expenditure function*.[23] Allowing Θ to be arbitrary enables us to capture different aspects of financial constraints such as the size of cash holdings and the terms of credit lines. For instance, suppose that

$$c(x,\theta) = \begin{cases} x & \text{if } x \leq \theta_2 \\ \theta_2 + (x-\theta_2)[1+\theta_3] & \text{if } \theta_2 < x \leq \theta_2 + \theta_4 \\ \theta_2 + \theta_4[1+\theta_3] + (x-\theta_2-\theta_4)[1+\theta_3+\theta_5] & \text{if } x > \theta_2 + \theta_4. \end{cases}$$

A buyer of type θ has a valuation θ_1 and cash holdings of θ_2. He can borrow up to θ_4 at the interest rate θ_3, and he faces a higher interest rate of $\theta_3 + \theta_5$ when exceeding the

[21] Holt (1980), Maskin and Riley (1984), Matthews (1983, 1987) and Riley and Samuelson (1981) are some of the major theoretical contributions. Kagel (1995) discusses the possible role of risk aversion in explaining certain anomalies in auction experiments.

[22] See Benoît and Krishna (2001), Che and Gale (1996, 1998), Fang and Perreiras (2003), Laffont and Robert (1996) Maskin (2000), Rhodes-Kropf and Viswanathan (forthcoming), and Zheng (2001).

[23] As noted, Che and Gale (1998) considered a case in which Θ was two-dimensional, with $u(x;\theta) = \theta_1 - c(x,\theta_2)$; they made several additional assumptions.

credit limit, θ_4. This example allows for non-nested constraints as a buyer could face a tighter constraint than other buyers do in one dimension (e.g., the size of the credit line), but a looser constraint in another (e.g., the interest rate).

Our model also allows for bidder risk aversion with more complex risk characteristics than is the case in existing models. Bidders may differ in both valuations and the degree of risk aversion, as with the CARA utility function:

$$u(x;\theta) = 1 - \exp[-\theta_2(\theta_1 - x)],$$

where θ_1 represents the valuation and θ_2 represents the degree of absolute risk aversion. More general preferences are also possible, with non-CARA utility functions and a general θ. For example, bidders could differ in their attitudes toward risk, and risk aversion could vary with income. There could also be financial constraints in addition to risk aversion.

The revenue comparisons rely on certain properties of symmetric equilibria in each auction form, given (U1)–(U2). We begin with a second-price auction. In a symmetric equilibrium of a second-price auction, it is optimal for a bidder to raise b until $u(b;\theta) = 0$ since he gets utility of zero if he does not win. More precisely, it is a weakly dominant strategy for a bidder of type θ to bid[24]

$$B_s(\theta) := \max\{x \mid u(x;\theta) \geq 0\}. \tag{10}$$

Given (U2), $\mathscr{B}_s$ is bounded since the supremum bid is $\overline{b}_s := \sup_{\theta\in\Theta} B_s(\theta) < K$. In fact, the next lemma shows that this is the unique symmetric equilibrium strategy. Moreover, (A1) holds since each equilibrium bid, $b \in \mathscr{B}_s$, is a best response for a benchmark bidder of type $v = b$. To apply our methodology, we construct the mapping ϕ_s such that $\phi_s(b) = b$ for each $b \in \mathscr{B}_s$. With ϕ_s constructed this way, (A2) holds.

LEMMA 3. *Assume (U1)–(U2) hold. In any symmetric equilibrium of a second-price auction, each bidder bids according to the strategy $B_s(\cdot)$ with probability one. In addition, with $\phi_s(b) = b$ for all $b \in \mathscr{B}_s$, (A2) holds, and $F_s(v) = \Pr\{B_s(\theta) \leq v\}$, for all $v \in \mathscr{V}$.*

For a first-price auction, we assume existence of a symmetric equilibrium in *pure* strategies in which a bidder with $\theta \in \Theta$ bids $B_f(\theta)$. While assuming existence of a symmetric equilibrium is a restriction, any *mixed-strategy* equilibrium can essentially be rendered pure by introducing artificial types with the same preferences as existing types that employ mixed strategies. That is, one can generate the same distribution of bids in a pure-strategy equilibrium with artificial types as in the original mixed-strategy equilibrium.[25] In that sense, our comparison applies to a general distribution that could

[24] The maximum is well defined, given continuity of $u(\cdot;\theta)$.

[25] Let Θ_m be the set of types playing mixed strategies in the symmetric equilibrium. For each $\theta \in \Theta_m$, let $\mathscr{B}_\theta$ denote the support of bids for that type. Now augment the types for $\theta \in \Theta_m$. Specifically, create types of the form (θ, b), with $b \in \mathscr{B}_\theta$, such that the distribution of types coincides with the distribution of bids in the original equilibrium. It is now a symmetric *pure-strategy* equilibrium for each $\theta \in \Theta \setminus \Theta_m$ to make its original equilibrium bid and for each (θ, b), with $\theta \in \Theta_m$ and $b \in \mathscr{B}_\theta$, to bid b. This pure-strategy equilib-

involve atoms (in the usual sense) and multiple dimensions. While equilibria of first-price auctions cannot be explicitly characterized, we can establish several properties of the equilibria that help us to apply Theorem 1.

LEMMA 4. *Given (U1)–(U2), any symmetric equilibrium of a first-price auction has the following properties. (a)* $\inf \mathscr{B}_f = \inf \mathscr{B}_s =: \underline{b}$, $\Pr\{B_f(\theta) \le \underline{b}\} = \Pr\{B_s(\theta) \le \underline{b}\}$, *and* $\overline{b}_f := \sup \mathscr{B}_f < K$. *(b)* $\mathscr{B}_f$ *is an interval and it has no mass points at any* $b > \underline{b}$. *(c) There exist* $\hat{v} \ge \underline{b}$ *and* $\hat{v}' \ge \hat{v}$ *such that* $v < \underline{b}$ *implies* $X_f(v) = 0$ *and* $\beta_f(v) = \emptyset$; $v \in (\underline{b}, \hat{v})$ *implies* $BR_f(v) = \emptyset$; $v \in (\hat{v}, \hat{v}')$ *implies* $\beta_f(v) \setminus \{\underline{b}\} \ne \emptyset$; *and* $v > \hat{v}'$ *implies* $BR_f(v) = \{\overline{b}_f\}$. *(d) The set* $\beta_f(\mathscr{V})$ *is nonempty, so* ϕ_f *is well-defined. (e) (A2) holds.*

Our revenue comparison then follows.

PROPOSITION 2. *Given (U1)–(U2), a symmetric equilibrium of a first-price auction yields (weakly) higher expected revenue than the symmetric equilibrium of the second-price auction in (10). The revenue ranking is strict if* $\overline{b}_s > \underline{b}$ *and* $u(\cdot, \theta)$ *is strictly concave for all* $\theta \in \Theta$.

PROOF. Fix symmetric equilibria for the first-price auction ($M = f$) and the second-price auction ($M = s$). It is straightforward to establish that $\Pi_f(0) = \Pi_s(0) = 0$ (just as in the proof of Proposition 2). Lemmas 3 and 4 have shown that (A1) and (A2) hold in the equilibrium of a second-price auction and that (A2) holds in the equilibrium of a first-price auction. Hence, by Theorem 1, to get the revenue ranking it suffices to show that $F_f(v^*) \le F_s(v^*)$ for every $v^* \in \mathscr{V}$. Note that $F_s(v^*) = \Pr\{B_s(\theta) \le v^*\}$ for each $v^* \in \mathscr{V}$, by Lemma 3.

Now turn to the first-price auction. Once again, $\beta_f(\cdot)$ is nondecreasing, by Lemma 1, and it is bounded above by K (by Lemma 4(a)), so $\beta_f(v)$ is a singleton for almost every v for which it is nonempty. By Lemma 4(c), $\beta_f(v) \ne \emptyset$ for every $v \in (\hat{v}, \hat{v}')$, so $\beta_f(v)$ contains a unique best response for almost every $v \in (\hat{v}, \hat{v}')$. Fix a valuation, $v^* \in (\hat{v}, \hat{v}')$, with a unique best response, which we denote $b_f(v^*)$. Then,

$$F_f(v^*) = \Pr\{\phi_f(B_f(\theta)) \le v^*\} = \Pr\{B_f(\theta) \le b_f(v^*)\} = \Pr\{B_f(\theta) < b_f(v^*)\},$$

where the first and the second equalities follow from (A.6), and the last follows from there being no mass at $b_f(v^*) > \underline{b}$ (by Lemma 4(b)). Hence, a sufficient condition for $F_f(v^*) \le F_s(v^*)$ to hold is that $B_s(\theta) \le v^*$ whenever $B_f(\theta) < b^* := b_f(v^*)$.

Fix θ such that $B_f(\theta) =: b < b^*$. We will show that this implies $B_s(\theta) \le v^*$. Observe first that $v^* \in (\hat{v}, \hat{v}')$ and $\hat{v} \ge \underline{b}$. This means that $v^* > b^*$ and $x_f(b^*) > 0$ since a benchmark bidder of type $v^* > \underline{b}$ can get a strictly positive expected payoff by bidding

rium has the same equilibrium bid distribution and the same expected revenue as the original equilibrium. Meanwhile, this change to the type space would have no effect on the equilibrium of the second-price auction since the last component is not payoff-relevant. Hence, the revenue comparison between the first- and second-price auctions for the original type space follows trivially from the revenue comparison for the extended type space, which is established in Proposition 2.

$b' \in (\underline{b}, v^*)$. If a type-$\theta$ actual bidder finds it optimal to bid b in a first-price auction, then

$$x_f(b)u(b;\theta) \geq x_f(b^*)u(b^*;\theta). \tag{11}$$

Meanwhile, a type-v^* benchmark bidder has b^* as a best response, so

$$x_f(b^*)[v^* - b^*] \geq x_f(b)[v^* - b]. \tag{12}$$

There are now two subcases. Suppose, first, that either $x_f(b) = 0$ or $u(b^*;\theta) \leq 0$. In this case, (11) implies $u(v^*;\theta) < u(b^*;\theta) \leq 0$ since $v^* > b^*$. Hence, $B_s(\theta) \leq v^*$, as was to be shown.

Now suppose that $x_f(b) > 0$ and $u(b^*;\theta) > 0$. Multiplying the respective sides of (11) and (12), and dividing through by $x_f(b)x_f(b^*)$, we get

$$u(b;\theta)[v^* - b^*] \geq u(b^*;\theta)[v^* - b]. \tag{13}$$

Concavity of $u(\cdot;\theta)$ implies

$$u(b;\theta) \leq u(b^*;\theta) + u_1(b^*;\theta)[b - b^*], \tag{14}$$

where $u_1(y;\theta)$ denotes an arbitrary sub-derivative with respect to the first argument, evaluated at $(y;\theta)$. Substituting this bound for $u(b;\theta)$ into (13) yields

$$[u(b^*;\theta) + u_1(b^*;\theta)(v^* - b^*)](b^* - b) \leq 0. \tag{15}$$

Concavity of $u(\cdot;\theta)$ also gives

$$u(v^*;\theta) \leq u(b^*;\theta) + u_1(b^*;\theta)(v^* - b^*). \tag{16}$$

Since $b < b^*$, (15) then implies that $u(v^*;\theta) \leq 0$, which again means $B_s(\theta) \leq v^*$. We conclude that $F_f(v^*) \leq F_s(v^*)$ for almost every $v^* \in (\hat{v}, \hat{v}')$.

Now consider any $v^* > \hat{v}'$. Clearly,

$$F_f(v^*) \leq 1 = \Pr\{B_f(\theta) \leq \overline{b}_f\}.$$

Again, it suffices to show that $B_s(\theta) \leq v^*$ whenever $B_f(\theta) \leq \overline{b}_f$. Since $BR_f(v^*) = \{\overline{b}_f\}$, by Lemma 4(c), and $v^* > \overline{b}_f$, the same argument as before proves the result.

We next consider $v^* \in (\underline{b}, \hat{v})$. By Lemma 4(c), $v \in (\underline{b}, v^*]$ implies $BR_f(v) = \emptyset$, so $v \notin \mathscr{V}_f$. Hence, $F_f(v^*) = F_f(\underline{b})$. In addition, if this region exists, there must be a mass point at $\underline{b}$, implying $\underline{b} \in \mathscr{B}_f$. Furthermore, a benchmark bidder of type $v = \underline{b}$ has $\beta_f(v) = \{\underline{b}\}$. This means that, for any $v^* \in (\underline{b}, \hat{v})$, we have

$$F_f(v^*) = F_f(\underline{b}) = \Pr\{B_f(\theta) \leq \underline{b}\},$$

where the last equality follows from (A.5). But, Lemma 4(a) implies $\Pr\{B_f(\theta) \leq \underline{b}\} = \Pr\{B_s(\theta) \leq \underline{b}\}$, so

$$F_f(v^*) = \Pr\{B_s(\theta) \leq \underline{b}\} = F_s(\underline{b}) \leq F_s(v^*),$$

where the inequality follows from $v^* > \underline{b}$.

Finally, for almost every $v^* \in [0, \underline{b})$, Lemmas 4(c) and 4(e) imply that $F_f(v^*) = X_f(v^*) = 0$. Since $F_s(v^*) \geq 0$, we have $F_f(v^*) \leq F_s(v^*)$.

The analysis has shown that $F_f(v^*) \leq F_s(v^*)$ for almost every $v^* \in \mathscr{V}$. Since F_s is right continuous, $F_s(v) \geq F_f(v)$ for every $v \in \mathscr{V}$, as was to be shown. The second statement (the strict ranking) follows since if $\overline{b}_s > \underline{b}$, then $\overline{b}_f > \underline{b}$, and strict concavity causes (14) and the corresponding inequality for $v^* > \hat{v}'$ to be strict, proving that $F_f(v) < F_s(v)$ for a positive measure of v. □

The proof has established a stochastic dominance relationship between the induced distributions (i.e., $F_f(\cdot) \leq F_s(\cdot)$), which has an intuitive interpretation: a fictitious risk-neutral bidder would be more likely to lose to bidders who are risk-averse (or financially constrained) in a first-price auction than in a second-price auction. This intuition parallels the familiar one that risk aversion makes bidders more aggressive in a first-price auction than in a second-price auction.[26]

REMARK (Payoff nonequivalence). The bidders themselves may have a strict preference for one auction form over another. To see this easily, suppose that Θ contains a risk-neutral type. Such a bidder would win with a lower probability in a first-price auction than in a second-price auction (see the proof of Proposition 2), so $\Pi_f(v) = \int_0^v F_f(z)^{n-1} dz \leq \Pi_s(v) = \int_0^v F_s(z)^{n-1} dz$. This payoff nonequivalence holds even when all types have CARA utility if there is heterogeneity in the coefficient of absolute risk aversion.[27]

3.3 *Risk averse bidders with non-expected utility preferences*

To this point we have assumed that bidders are expected utility maximizers. We now show that our ranking of first- and second-price auctions continues to hold for risk-averse bidders who are not expected utility maximizers.[28] This robustness is important given the well-documented violations of the predictions of the expected utility model.

The result is shown using the Gateaux differentiable preference functional, which replaces the independence axiom with a minimal condition of smoothness. Gateaux differentiability does not require continuity of the preference functional in the distribution, and is thus weaker than Fréchet differentiability (see Machina 1982). It encompasses many well-known non-expected utility preferences, such as those satisfying the

[26] With multi-dimensional types, the standard intuition may not work in an absolute sense since risk aversion in our general form may affect the bidding in a second-price auction.

[27] Matthews (1987) found payoff equivalence for bidders with identical CARA utility. The reason for nonequivalence here is consistent with the logic for nonequivalence with (one-dimensional) non-CARA preferences there since bidders with different valuations may have different levels of absolute risk aversion in this latter case.

[28] Auctions in which bidders are not expected utility maximizers have been studied by Karni and Safra (1989), and Neilson (1994), for example. Those papers compare second-price and ascending-bid auctions. In both cases, private information is one-dimensional, and the prize is a lottery. The issues they study do not arise in our setting, where the object's value is deterministic. Volij (2002) found payoff equivalence using the *dual theory of choice* model.

betweenness axiom (Dekel 1986) and rank-dependent expected utility (Quiggin 1982 and Wakker 1994), given some additional restrictions.[29]

To begin, let an actual bidder with type $\theta = (\theta_1, \theta_2) \in [0,1] \times \Theta_2 =: \Theta$ earn a random net surplus $Y = \theta_1 \cdot \mathbb{I}_{\{\text{win}\}} - c(Z, \theta_2)$ where Z is a random variable representing the payment and $\mathbb{I}_{\{\text{win}\}}$ is an indicator function that equals 1 when the bid wins, and zero otherwise; $c(\cdot, \theta_2)$ is increasing and continuous, with $c(0, \theta_2) = 0$; and Θ_2 is arbitrary. We make no particular assumptions concerning θ except that, for simplicity, θ_1 does not have an atom at zero. Let $\Delta([0,1])$ denote the set of all probability distributions of the net surplus. We assume that the type-θ bidder's preference functional, $\mathscr{U}(\cdot; \theta_2) : \Delta([0,1]) \mapsto \mathbb{R}$, is *Gateaux differentiable*: *For each* $\mathscr{F} \in \Delta([0,1])$, *there exists* $\xi(\cdot, \mathscr{F}; \theta_2) : [0,1] \mapsto \mathbb{R}$ *such that,* $\forall \mathscr{G} \in \Delta([0,1])$ *and* $\alpha \in [0,1]$,

$$\mathscr{U}((1-\alpha)\mathscr{F} + \alpha\mathscr{G}; \theta_2) - \mathscr{U}(\mathscr{F}; \theta_2) = \alpha \int_{[0,1]} \xi(\cdot, \mathscr{F}; \theta_2) d[\mathscr{G} - \mathscr{F}] + o(\alpha).$$

Note that this functional collapses to an expected utility representation if the Gateaux derivative, $\xi(\cdot, \mathscr{F}; \theta_2)$, does not depend on $\mathscr{F}$.[30] In general, its dependence on $\mathscr{F}$ means that the preferences do not conform to the expected utility representation, although such a representation is valid for local directional shifts of the distribution. With Gateaux differentiable preferences, monotonicity of preferences and risk aversion are represented by $\xi(\cdot, \mathscr{F}; \theta_2)$ being strictly increasing and concave, respectively, for all $\mathscr{F} \in \Delta([0,1])$ (see Chew and Mao 1995).

The monotonicity of $\xi(\cdot, \mathscr{F}; \theta_2)$ means that an actual bidder prefers a (first-order) stochastically dominating shift of the distribution of Y. This implies that the equilibrium of the second-price auction takes the same form as before. That is, a type-θ bidder bids

$$B_s(\theta) = \max\{x \mid \theta_1 \geq c(x, \theta_2)\},$$

which is equivalent to (10), thus satisfying (A1)–(A2). The associated random payoff stochastically dominates the random payoff associated with any other bid. A symmetric (pure-strategy) equilibrium of a first-price auction, assuming it exists, is characterized as in Lemma 4.

The benchmark bidder has the same characteristics as in the previous section (i.e., a risk-neutral, expected-utility maximizer with no financial constraints). The previous revenue ranking then extends to this environment.

PROPOSITION 3. *Given a Gateaux derivative* $\xi(\cdot, \mathscr{F}; \theta_2)$ *that is strictly increasing and (weakly) concave for all* $(\mathscr{F}, \theta_2) \in \Delta([0,1]) \times \Theta_2$, *a symmetric equilibrium of a first-price*

[29]The betweenness axiom gives rise to an implicit function representation, which is Gateaux differentiable if the partial derivative of the equation with respect to its second argument is bounded. The rank-dependent expected utility representation is Gateaux differentiable if its probability transformation function is differentiable. See Chew and Mao (1995) for details.

[30]The current framework does not include the one in Section 3.2 as a special case, however. In the current approach, we are treating the good and the money as perfect substitutes, which is why we focus simply on the stream of "net surplus." We do not impose such a condition in Section 3.2.

auction yields (weakly) higher expected revenue than the symmetric equilibrium of the second-price auction in (10). The revenue ranking is strict if $\xi(\cdot,\mathscr{F};\theta_2)$ is strictly concave for all $(\mathscr{F},\theta_2)$.

PROOF. The characterizations of Lemmas 3 and 4 follow, except that we now have $\underline{b}=0$, and there exists $\hat{v}'$ such that $\beta_f(v)\setminus\{\underline{b}\}\neq\emptyset$ if $v\in(0,\hat{v}')$ and $\beta_f(v)=\{\overline{b}_f\}$ if $v>\hat{v}'$. (Since θ_1 does not have an atom at 0, $\mathscr{B}_f$ forms an open interval and has no atoms, which implies that $BR_f(v)$ is nonempty for all v; the specific characterization follows from the proof of Lemma 4(c).) Since the proof mirrors that of Proposition 2, we simply highlight the differences.

Let $B_f(\theta)$ denote the equilibrium strategy under a first-price auction. As in the proof of Proposition 2, the weak ranking holds if the following condition holds: For each $v^*\in(0,\hat{v}')$ such that $\beta_f(v^*)$ has a singleton element, $b_f(v^*)$, $B_f(\theta)<b_f(v^*)$ implies $B_s(\theta)\leq v^*\Leftrightarrow\theta_1\leq c(v^*,\theta_2)$.[31]

As before, fix any $\theta=(\theta_1,\theta_2)$ such that $B_f(\theta)=:b<b_f(v^*)=:b^*$, and let $x^*:=x_f(b^*)$ and $x:=x_f(b)$. Then, incentive compatibility for a type-v^* benchmark bidder implies

$$x^*[v^*-b^*]\geq x[v^*-b]. \tag{17}$$

Next, consider an actual bidder with type θ, and let $\mathscr{F}_{b'}$ denote the distribution of his surplus when he makes some bid b'. Since $B_f(\theta)=b$, a bid of b must be (weakly) preferred to any other single bid or any mixed strategy over bids.

Suppose that this bidder randomizes between b and b^* with probabilities $1-\alpha$ and α, respectively. His payoff is $\mathscr{U}((1-\alpha)\mathscr{F}_b+\alpha\mathscr{F}_{b^*};\theta_2)$. Since bidding b with probability one is optimal, we have

$$\begin{aligned}0&\geq\left.\frac{\partial\mathscr{U}((1-\alpha)\mathscr{F}_b+\alpha\mathscr{F}_{b^*};\theta_2)}{\partial\alpha}\right|_{\alpha=0}\\&=\int_{[0,1]}\xi(\cdot,\mathscr{F}_b;\theta_2)d[\mathscr{F}_{b^*}-\mathscr{F}_b]\\&=x^*\xi(\theta_1-c(b^*,\theta_2),\mathscr{F}_b;\theta_2)+(1-x^*)\xi(0,\mathscr{F}_b;\theta_2)\\&\qquad-[x\xi(\theta_1-c(b,\theta_2),\mathscr{F}_b;\theta_2)+(1-x)\xi(0,\mathscr{F}_b;\theta_2)],\end{aligned} \tag{18}$$

where the first equality follows from Gateaux differentiability and the second follows from the fact that both $\mathscr{F}_b$ and $\mathscr{F}_{b^*}$ involve two-point distributions. We can rewrite (18) as

$$x[\xi(\theta_1-c(b,\theta_2),\mathscr{F}_b;\theta_2)-\xi(0,\mathscr{F}_b;\theta_2)]\geq x^*[\xi(\theta_1-c(b^*,\theta_2),\mathscr{F}_b;\theta_2)-\xi(0,\mathscr{F}_b;\theta_2)]. \tag{19}$$

As in the proof of Proposition 2, combining (17), (19), and concavity of $\xi(\cdot,\mathscr{F}_b;\theta_2)$ gives

$$[\xi(\theta_1-c(b^*,\theta_2),\mathscr{F}_b;\theta_2)+\xi_1(\theta_1-c(b^*,\theta_2),\mathscr{F}_b;\theta_2)(v^*-b^*)-\xi(0,\mathscr{F}_b;\theta_2)](b^*-b)\leq 0. \tag{20}$$

[31] Recall that almost every $v\in\mathscr{V}_f$ has a singleton element.

By concavity of $\xi(\cdot, \mathscr{F}_b; \theta_2)$,

$$\xi(\theta_1 - c(v^*, \theta_2), \mathscr{F}_b; \theta_2) \leq \xi(\theta_1 - c(b^*, \theta_2), \mathscr{F}_b; \theta_2) + \xi_1(\theta_1 - c(b^*, \theta_2), \mathscr{F}_b; \theta_2)(v^* - b^*),$$

so (20) implies

$$\xi(\theta_1 - c(v^*, \theta_2), \mathscr{F}_b; \theta_2) \leq \xi(0, \mathscr{F}_b; \theta_2),$$

from which it follows that

$$\theta_1 \leq c(v^*, \theta_2),$$

or $B_s(\theta) \leq v^*$, as was to be shown. This gives the weak ranking. Strict concavity makes the inequalities strict, resulting in a strict ranking. □

4. Conclusion

This paper develops a methodology for characterizing a seller's expected revenue when bidders' types come from an arbitrary distribution. In particular, types may be multidimensional, with mass points and gaps in the distribution, and the support of equilibrium bids may have mass points and gaps itself. The revenue characterization result is used to generalize existing revenue equivalence results and to show that first-price auctions revenue-dominate second-price auctions when bidders are risk averse and face financial constraints.

By considering arbitrary distributions, this paper greatly expands the range of cases for which revenue comparisons can be made. Our method may therefore have useful applications for other cases with multidimensional types and general forms of nonlinear payoffs.

Appendix

Proof of Lemma 2. To prove property (a), suppose that B_M has a mass point at $b \in \mathscr{B}_M$, with $x_M(b_+) - x_M(b_-) =: m_b > 0$, where $x_M(b_-)$ and $x_M(b_+)$ denote the left and right limit of $x_M(\cdot)$ at b, respectively. Consider a benchmark bidder with valuation $v \in \mathscr{V}$. By raising the bid infinitesimally above b, a bidder obtains an expected payoff of

$$\begin{aligned}\pi_M(b_+, v) = x_M(b_-)(v - \mathbb{E}[\tau_w(b, B_M^{(1:n-1)}) \mid B_M^{(1:n-1)} < b]) + m_b(v - \tau_w(b, b)) \\ - [1 - x_M(b_+)]\mathbb{E}[\tau_l(b, B_M^{(1:n-1)}) \mid B_M^{(1:n-1)} > b], \quad \text{(A.1)}\end{aligned}$$

where $B_M^{(1:n-1)}$ is the first order statistic of $n-1$ independent draws of B_M. (The continuous payments property gives $\tau_k(b_+, B_M^{(1:n-1)}) = \tau_k(b, B_M^{(1:n-1)})$ for $k = w, l$.) A bid of b would give the benchmark bidder an expected payoff of

$$\begin{aligned}\pi_M(b, v) = x_M(b_-)(v - \mathbb{E}[\tau_w(b, B_M^{(1:n-1)}) \mid B_M^{(1:n-1)} < b]) + m_b \rho_b(v - \tau_w(b, b)) \\ - m_b(1 - \rho_b)\tau_l(b, b) - [1 - x_M(b_+)]\mathbb{E}[\tau_l(b, B_M^{(1:n-1)}) \mid B_M^{(1:n-1)} > b], \quad \text{(A.2)}\end{aligned}$$

where ρ_b denotes the probability of winning conditional on bidding b and tying for the high bid (i.e., when $(b_i, b_{m(i)}) = (b, b)$). The expected gain from raising b_i above b is therefore

$$\pi_M(b_+, v) - \pi_M(b, v) = m_b(1-\rho_b)[v - \delta_M(b)]. \tag{A.3}$$

Likewise, we have

$$\pi_M(b, v) - \pi_M(b_-, v) = m_b \rho_b [v - \delta_M(b)], \tag{A.4}$$

for $b > 0$.

Equations (A.3) and (A.4) imply that a benchmark bidder with a best response of b must have a valuation $v = \delta_M(b)$; otherwise, the expected payoff would jump when raising or lowering the bid marginally.[32] Since no other actual type can contribute to the mass, we must have $v \in \Theta$, as was claimed in (a).

To prove property (b), fix $v', v \in \Theta$ with $v' > v$. Lemma 1 shows that we must have $b' \geq b$ if $b' \in \beta_M(v')$ and $b \in \beta_M(v)$. If $b' = b$, an interval of valuations must have a best response of b, but this contradicts (a).[33] Hence, we must have $b' > b$. □

Proof of Lemma 3. Fix a constant, $\epsilon > 0$. A bidder ϵ-*overbids* if she bids more than $B_s(\theta) + \epsilon$ when her type is $\theta \in \Theta$; she ϵ-*underbids* if she bids less than $B_s(\theta) - \epsilon$. To prove uniqueness, suppose that there exists a symmetric equilibrium in which a bidder ϵ-overbids or ϵ-underbids with strictly positive probability. In particular, suppose that the probability of ϵ-overbidding is positive. Let $\mathscr{B}^+$ be the support of equilibrium bids that entail ϵ-overbidding, and let $\overline{b}^+ := \sup \mathscr{B}^+$.[34] Then, for a fixed $\delta < \epsilon$, a given bidder ϵ-overbids in $(\overline{b}^+ - \delta, \overline{b}^+]$ with positive probability. Since all bidders adopt the same strategy, the highest rival bid lies in that same interval with positive probability. By deviating to $B_s(\cdot)$ whenever she would have ϵ-overbid, a bidder is strictly better off. This is immediate if $\mathscr{B}^+$ has a mass point at $\overline{b}^+$, since ϵ-overbidding at $\overline{b}^+$ would result in a positive probability of winning, which would give the bidder a payoff of $u(\overline{b}^+, \theta) \leq u(B_s(\theta) + \epsilon, \theta) < u(B_s(\theta), \theta) = 0$. If $\mathscr{B}^+$ does not have a mass point at $\overline{b}^+$, then $\overline{b}^+ = \sup \mathscr{B}^+$ means that the highest rival bid lies in any nonempty subset of $(\overline{b}^+ - \delta, \overline{b}^+)$ with positive probability. Then, deviating to $B_s(\cdot)$ whenever she would have ϵ-overbid in $(\overline{b}^+ - \frac{1}{2}\delta, \overline{b}^+)$, say, is strictly profitable. The argument is analogous when the probability of ϵ-underbidding is positive.

To prove the second statement, note that almost every $v \in \mathscr{V}$ is not a mass point in $\mathscr{B}_s$. Fix any such v. Since it is a best response to bid v, a benchmark bidder with v wins

[32] When $b = 0$, we have $\tau_w(0,0) = \tau_l(0,0) = 0$, so $\delta_M(0) = 0$. Given (A.3), $b = 0$ implies $v = 0$, so $v = \delta_M(b)$ again.

[33] Following footnote 7, if $v, v' \in \Theta$ with $v < v'$, there must be atoms at v and v' or a positive measure of types in $(v, v') \cap \Theta$.

[34] Let $P_\delta^+(b)$ be the probability that a type-θ bidder bids $b' \in [b, b+\delta)$ such that $b' \geq B_s(\theta) + \epsilon$ when $\theta \in \Theta$, and let $P_\delta^-(b)$ be the probability that $b' \in (b-\delta, b]$ such that $b' \geq B_s(\theta) + \epsilon$ when $\theta \in \Theta$. Then

$$\mathscr{B}^+ := \{b \mid P_\delta^-(b) > 0 \text{ and } P_\delta^+(b) > 0, \forall \delta > 0\}.$$

with probability $X_s(v) = (\Pr\{B_s \le v\})^{n-1}$. Meanwhile, since $\phi_s(B_s) = B_s$, we have

$$F_s(v) = \Pr\{\phi_s(B_s) \le v\} = \Pr\{B_s \le v\},$$

where the first equality follows from (2). Combining the two preceding facts, we have $X_s(v) = F_s(v)^{n-1}$, proving (A2). Further, since $B_s = B_s(\theta)$ in the unique symmetric equilibrium, we have $F_s(v) = \Pr\{B_s(\theta) \le v\}$, for all $v \in \mathscr{V}$. □

Proof of Lemma 4. The first part of (a) claims that $\underline{b}_f = \underline{b}_s$, where $\underline{b}_M := \inf \mathscr{B}_M$ for $M = f, s$. Suppose that $\underline{b}_s < \underline{b}_f$ instead. Then, any type θ with $B_s(\theta) \in (\underline{b}_s, \underline{b}_f)$ could get a strictly positive expected payoff in a first-price auction by bidding $B_s(\theta) - \epsilon > \underline{b}_s$, for small $\epsilon > 0$. This means that $\underline{b}_f \le \underline{b}_s$, which gives a contradiction. Now suppose that $\underline{b}_f < \underline{b}_s$. Then, the equilibrium payoff for every $\theta \in \Theta$ must be bounded away from zero in the first-price auction. This would mean that there is a mass point at $\underline{b}_f$, yet those types putting mass at $\underline{b}_f$ could profitably deviate by raising the mass slightly above $\underline{b}_f$. (The deviation raises the probability of winning discontinuously but the (strictly positive) payoff upon winning decreases continuously, by (U1).) Hence, we have a contradiction, so $\underline{b}_f = \underline{b}_s$. In addition, the same actual types bid $\underline{b}$ in both formats, so $\Pr\{B_f(\theta) \le \underline{b}\} = \Pr\{B_s(\theta) \le \underline{b}\}$. The last claim in (a) holds since $\overline{b}_f \ge K$ would mean that a positive measure of θ has a non-zero probability of winning with a bid that entails a strictly negative payoff, a situation avoided by bidding zero.

We next prove (b). It is easy to see that $\mathscr{B}_f$ must be an interval. If there were a gap in $\mathscr{B}_f$, any bid within ϵ of the supremum of the gap could profitably be lowered to ϵ above the infimum of the gap, for small $\epsilon > 0$. (The drop in the probability of winning would be of order ϵ, but the increase in the payoff upon winning is roughly proportional to the length of the gap.) Consequently, there cannot be a gap in $\mathscr{B}_f$.

To see that there cannot be a mass point above $\underline{b}$, suppose that the equilibrium strategy called for a particular bid, $b > \underline{b}$, to be submitted with positive probability. That bid exceeds the infimum, so it wins with strictly positive probability, and almost every type that bids b has $u(b;\theta) > 0$.[35] Since the winning probability jumps at b and $u(\cdot,\theta)$ is continuous, a profitable deviation exists when a bidder draws one of these types. Thus, there cannot be a mass point at any $b > \underline{b}$ in equilibrium.

In order to prove (c) we make several preliminary observations. First, a benchmark bidder with $v < \underline{b}$ gets a negative payoff if she wins. This means that $X_f(v) = 0$ and $\beta_f(v) = BR_f(v) \cap \mathscr{B}_f = \emptyset$ for $v < \underline{b}$. Second, for all $b > \overline{b}_f$ and all v, $\pi_f(b,v) - \pi_f(\overline{b}_f, v) = \overline{b}_f - b < 0$ since $x_f(b) = x_f(\overline{b}_f) = 1$, given the absence of mass points shown in (b). Third, for any $v > \underline{b}$ and $b \le \underline{b}$, there exists $b' > \underline{b}$ such that $\pi(b',v) > \pi(b,v)$.[36] These observations imply that, for any $v > \underline{b}$, whenever $BR_f(v)$ is non-empty it must be a subset of $(\underline{b}, \overline{b}_f]$. (That is, for each $v > \underline{b}$, either $BR_f(v) = \emptyset$ or $BR_f(v) \cap (\underline{b}, \overline{b}_f] \ne \emptyset$.) Since $\pi_f(b,v)$ has increasing differences in $(b;v)$ and is continuous in b for $b \in (\underline{b}, \overline{b}_f]$, we have two additional observations: (i) if $BR_f(v) \cap (\underline{b}, \overline{b}_f] \ne \emptyset$, then $BR_f(v') \cap (\underline{b}, \overline{b}_f] \ne \emptyset$ for $v' > v$; and

[35] If not, a slightly lower bid would win with positive probability, but the payoff upon winning would be strictly greater, so the expected payoff from this lower bid would be strictly positive.

[36] This holds since $b \le \underline{b}$ implies that either $x_f(b) = 0$ or else $b = \underline{b}$ and there is a mass point at $\underline{b}$; raising the bid slightly above $\underline{b}$ is profitable in the latter case.

(ii) if $BR_f(v)\cap(\underline{b},\overline{b}_f)\neq\emptyset$ and $BR_f(v')\cap(\underline{b},\overline{b}_f)\neq\emptyset$, for $v'>v$, then $BR_f(v'')\cap(\underline{b},\overline{b}_f)\neq\emptyset$ for all $v''\in(v,v')$. The earlier observations, combined with (i), imply that there exists $\hat{v}\geq\underline{b}$ such that $BR_f(v)$ is empty for $v\in(\underline{b},\hat{v})$ and $BR_f(v)\cap(\underline{b},\overline{b}_f]\neq\emptyset$ for $v>\hat{v}$.[37] This latter conclusion, along with (ii), implies that there exists $\hat{v}'\geq\hat{v}$ such that $BR_f(v)\cap(\underline{b},\overline{b}_f)\neq\emptyset$ if $v\in(\hat{v},\hat{v}')$ and $BR_f(v)=\{\overline{b}_f\}$ if $v>\hat{v}'$. Since $(\underline{b},\overline{b}_f)\subset\mathscr{B}_f$, $\beta_f(v)\setminus\{\underline{b}\}\neq\emptyset$ if $v\in(\hat{v},\hat{v}')$, so the proof is complete.

We next prove (d). There are two cases, depending on whether there is a mass point at $\underline{b}$. If there is a mass point, then $\underline{b}\in\mathscr{B}_f$ and $BR_f(\underline{b})$ contains $\underline{b}$, so $\beta_f(v)\neq\emptyset$ for $v=\underline{b}$. If there is not a mass point at $\underline{b}$, any $v\in(\underline{b},x_f(\overline{b}_f)/x_f'(\overline{b}_{f-})+\overline{b}_f)$ has $BR_f(v)\cap(\underline{b},\overline{b}_f)\neq\emptyset$, where $x_f'(\overline{b}_{f-})$ denotes a left-hand derivative of x_f at $\overline{b}_f$, which is positive, by (b). Since $(\underline{b},\overline{b}_f)\subset\mathscr{B}_f$, $\beta_f(v)\neq\emptyset$ for such v, which completes the proof.

Last, we turn to (e). Fix any $v<\underline{b}$. By (c), $X_f(v)=0$ and $\beta_f(v)=BR_f(v)\cap\mathscr{B}_f=\emptyset$. The latter fact means that $v'\notin\mathscr{V}_f=\phi_M(\mathscr{B}_f)$ if $v'\leq v$. Hence, $F_f(v)^{n-1}=\Pr\{\phi_M(\mathscr{B}_f)\leq v\}^{n-1}=0=X_f(v)$.

Now consider any $v\in(\underline{b},\hat{v})$. Then, (c) implies that $BR_f(v)=\emptyset$, which can only arise when $X_f(v)=(\Pr\{B_f(\theta)\leq\underline{b}\})^{n-1}$ and $\underline{b}$ is a mass point. The latter fact implies $\underline{b}\in\mathscr{B}_f$. Further, by (c), $BR_f(v')=\emptyset$ for each $v'\in(\underline{b},v]$, so v' does not belong to the support of V_f. It follows that $F_f(v)=F_f(\underline{b})$. A benchmark bidder of type $v=\underline{b}$ has $\phi_f^{-1}(v)=\beta_f(v)=\{\underline{b}\}$. Hence,

$$F_f(\underline{b})=\Pr\{\phi_f(B_f(\theta))\leq\underline{b}\}=\Pr\{B_f(\theta)\leq\underline{b}\}. \tag{A.5}$$

Combining the results, we have

$$X_f(v)=(\Pr\{B_f(\theta)\leq\underline{b}\})^{n-1}=F_f(\underline{b})^{n-1}=F_f(v)^{n-1},$$

as was to be shown.

Next consider $(\hat{v},\hat{v}')$. By (c), $\beta_f(v)\setminus\{\underline{b}\}\neq\emptyset$ for each $v\in(\hat{v},\hat{v}')$. Since $\beta_f(\cdot)$ is nondecreasing, by Lemma 1, and is bounded above by K (see (a)), $\beta_f(v)$ collapses to a singleton, say $\{b_f(v)\}$, for almost every $v\in(\hat{v},\hat{v}')$. Hence, for such v,

$$X_f(v)=x_f(b_f(v))=(\Pr\{B_f(\theta)\leq b_f(v)\})^{n-1}=(\Pr\{\phi_f(B_f(\theta))\leq v\})^{n-1}=F_f(v)^{n-1}. \tag{A.6}$$

The first and last equalities follow from the respective definitions, the second follows from the fact that there is no mass at $b_f(v)\in\mathscr{B}_f\setminus\{\underline{b}\}$, and the third equality is immediate if $b_f(\cdot)$ is strictly increasing at v, or else it follows from the fact that there is no mass at $b_f(v)$.

Finally, consider any $v>\hat{v}'$. Since there is no mass point at $\overline{b}_f$ (by (b)), we have $BR_f(v)=\{\overline{b}_f\}$. It follows that $X_f(v)=x_f(\overline{b}_f)=1$. Meanwhile, $\phi_f(B_f)\leq\hat{v}'$, $\forall B_f\in\mathscr{B}_f$, so $F_f(v)=1$ for $v>\hat{v}'$. Consequently, $X_f(v)=F_f(v)^{n-1}$ for all $v>\hat{v}'$. Since we have established $X_f(v)=F_f(v)^{n-1}$ for almost every $v\in\mathscr{V}$, the proof is complete. □

References

Asker, John and Estelle Cantillon (2003), "Equilibrium in scoring auctions." Unpublished paper, Harvard Business School. [96]

[37] Either region may not exist if $\hat{v}=\underline{b}$ or $\hat{v}=\hat{v}'$.

Benoît, Jean-Pierre and Vijay Krishna (2001), "Multiple-object auctions with budget constrained bidders." *Review of Economic Studies*, 68, 155–179. [106]

Che, Yeon-Koo (1993), "Design competition through multidimensional auctions." *RAND Journal of Economics*, 24, 668–680. [96]

Che, Yeon-Koo and Ian Gale (1996), "Financial constraints in auctions: effects and antidotes." *Advances in Applied Microeconomics*, 6, 97–120. [106]

Che, Yeon-Koo and Ian Gale (1998), "Standard auctions with financially constrained bidders." *Review of Economic Studies*, 65, 1–21. [97, 106]

Chew, Soo Hong and Mei Hui Mao (1995), "A Schur concave characterization of risk aversion for non-expected utility preferences." *Journal of Economic Theory*, 67, 402–435. [111]

Dekel, Eddie (1986), "An axiomatic characterization of preferences under uncertainty: weakening the independence axiom." *Journal of Economic Theory*, 40, 304–318. [111]

Fang, Hanming and Sérgio O. Perreiras (2001), "Standard auctions with financially constrained bidders: comment." Unpublished paper, University of Pennsylvania. [97]

Fang, Hanming and Sérgio O. Perreiras (2003), "On the failure of the linkage principle with financially constrained bidders." *Journal of Economic Theory*, 110, 374–392. [106]

Holt, Jr., Charles A. (1980), "Competitive bidding for contracts under alternative auction procedures." *Journal of Political Economy*, 88, 433–445. [96, 106]

Kagel, John H. (1995), "Auctions: a survey of experimental research." In *The Handbook of Experimental Economics* (John H. Kagel and Alvin E. Roth, eds.), 501–585, Princeton University Press, Princeton, New Jersey. [106]

Karni, Edi and Zvi Safra (1989), "Dynamic consistency, revelations in auctions and the structure of preferences." *Review of Economic Studies*, 56, 421–434. [110]

Laffont, Jean-Jacques and Jacques Robert (1996), "Optimal auction with financially constrained buyers." *Economics Letters*, 52, 181–186. [106]

Machina, Mark J. (1982), " 'Expected utility' analysis without the independence axiom." *Econometrica*, 50, 277–323. [110]

Maskin, Eric S. (2000), "Auctions, development and privatization: efficient auctions with liquidity-constrained buyers." *European Economic Review*, 44, 667–681. [106]

Maskin, Eric S. and John G. Riley (1984), "Optimal auctions with risk averse buyers." *Econometrica*, 52, 1473–1518. [106]

Maskin, Eric S. and John G. Riley (1985), "Auction theory with private values." *American Economic Review Papers and Proceedings*, 75, 150–155. [96, 101]

Matthews, Steven A. (1983), "Selling to risk averse buyers with unobservable tastes." *Journal of Economic Theory*, 30, 370–400. [96, 106]

Matthews, Steven A. (1987), "Comparing auctions for risk averse buyers: a buyer's point of view." *Econometrica*, 55, 633–646. [96, 106, 110]

Milgrom, Paul R. and Ilya Segal (2002), "Envelope theorems for arbitrary choice sets." *Econometrica*, 70, 583–601. [100]

Myerson, Roger B. (1981), "Optimal auction design." *Mathematics of Operations Research*, 6, 58–73. [96, 101]

Neilson, William S. (1994), "Second price auctions without expected utility." *Journal of Economic Theory*, 62, 136–151. [110]

Quiggin, John (1982), "A theory of anticipated utility." *Journal of Economic Behavior and Organization*, 3, 323–343. [111]

Rhodes-Kropf, Matthew and S. Viswanathan (forthcoming), "Financing auction bids." *RAND Journal of Economics*. [106]

Riley, John G. (1989), "Expected revenue from open and sealed bid auctions." *Journal of Economic Perspectives*, 3 (3), 41–50. [96, 101]

Riley, John G. and William F. Samuelson (1981), "Optimal auctions." *American Economic Review*, 71, 381–392. [96, 101, 106]

Volij, Oscar (2002), "Payoff equivalence in sealed bid auctions and the dual theory of choice under risk." *Economics Letters*, 76, 231–237. [110]

Wakker, Peter (1994), "Separating marginal utility and probabilistic risk aversion." *Theory and Decision*, 36, 1–44. [111]

Zheng, Charles Z. (2001), "High bids and broke winners." *Journal of Economic Theory*, 100, 129–171. [106]

Submitted 2005-6-24. Final version accepted 2005-9-22. Available online 2005-12-2.