Theoretical Economics 12 (2017), 53–78 1555-7561/20170053

Auction design without quasilinear preferences

Brian Baisa
Department of Economics, Amherst College

I study the canonical private value auction model for a single good without the quasilinearity restriction. I assume only that bidders are risk averse and the indivisible good for sale is a normal good. I show that removing quasilinearity leads to qualitatively different solutions to the auction design problem. Expected revenue is no longer maximized using standard auctions that allocate the good to the highest bidder. Instead, the auctioneer better exploits bidder preferences by using a mechanism that allocates the good to one of many different bidders, each with strictly positive probability. I introduce a probability demand mechanism that treats probabilities of winning the indivisible good like a divisible good in net supply 1. With enough bidders, it has greater expected revenues than any standard auction, and under complete information, it implements a Pareto efficient allocation.

Keywords. Auctions, multidimensional mechanism design, risk aversion, wealth effects.

JEL classification. C70, D44, D82.

1. Introduction

1.1 *Motivation*

Most of the auction design literature considers bidders with quasilinear preferences. In this paper, I revisit the canonical private value auction design problem and I remove the quasilinearity restriction. Instead, I assume only that bidders are risk averse and have positive wealth effects. This relaxation allows for a more complete description of bidder preferences but also complicates standard economic analysis. With quasilinearity, a bidder's incentives are described by her valuation. Without quasilinearity, bidder's incentives are also affected by her risk preferences, financial constraints, and wealth

Brian Baisa: bbaisa@amherst.edu
I gratefully acknowledge financial support from the Whitebox Advisors Fellowship. I received excellent comments from two anonymous referees, as well as seminars at Yale, Notre Dame, UNC, Johns Hopkins, Haverford College, Amherst College, Cambridge, the University of Michigan, and conferences hosted at Tel Aviv University, USC, Stony Brook, Lund, and the Informs annual meetings. I am especially grateful to Dirk Bergemann, Benjamin Polak, Larry Samuelson, and Johannes Hörner for numerous conversations and encouragement while advising me on this project. I would also like to thank Cihan Artunç, Lint Barrage, Tilman Börgers, Yeon-Koo Che, Eduardo Faingold, Amanda Gregg, Yingni Guo, Jun Ishii, Christopher Kingston, Vitor Farinha Luz, Drew Fudenberg, Adam Kapor, Phil Haile, Sofia Moroni, Steven Matthews, David Rappoport, Peter Troyan, and Kieran Walsh for helpful comments and conversations.

DOI: 10.3982/TE1951

effects. Thus, bidders' types are multidimensional and characterizing the optimal auction through a Myerson-like approach proves intractable. For this reason, I take a new approach to studying the design problem and obtain qualitatively different solutions relative to the quasilinear benchmark. I propose an alternative to standard auctions called the probability demand mechanism. The mechanism uses randomization to better exploit features of bidder preferences. While the multidimensionality of bidder types inhibits explicit characterizations of equilibrium behavior, I provide a partial characterization of bid behavior, and use this to obtain revenue comparisons. Specifically, I eliminate dominated strategies to bound a bidder's report. I use this bound on bid behavior to show that the probability demand mechanism has greater revenues than any standard auction when there are enough bidders.

There are many well studied settings where the quasilinearity restriction is violated. As an example, consider firms bidding on spectrum rights or oil tracts. The corporate finance literature shows that many firms have an internal spending hierarchy (Fazzari et al. 1988). Firms prefer to use internal versus external financing, because they pay higher interest on money borrowed from third parties. A firm may be able to place a relatively low bid in an auction without needing external financing, but to place a relatively high bid, the firm may need to obtain external financing and pay a higher interest rate on this debt. Consequently, even if firms are risk neutral, such a financing constraint makes them behave as though they have declining marginal utility of money.

I show that the auctioneer can increase revenue by using randomization. This result may seem counterintuitive with risk averse bidders, but the intuition follows directly from the assumption that the good is normal. When a good is normal, a bidder's willingness to pay for it increases with her wealth. Similarly, her willingness to pay for any given probability of winning the good increases with her wealth. Therefore, the bidder is willing to pay the most for her first marginal "unit" of probability of winning, before she has spent any of her wealth. Thus, the bidder is willing to buy a small probability of winning the good at a price per unit of probability that exceeds her willingness to pay for the entire good. Standard auctions that allocate the good to the highest "bidder" do not make use of this property of bidder preferences.

I construct a *probability demand mechanism* that uses lotteries to better exploit bidder preferences. The mechanism sells probabilities of winning the good like a divisible good that is in net supply 1. Bidders report a demand curve over probabilities of winning. The curve reports the probability of winning the bidder demands (Q) for a given price per unit of probability (P). The auctioneer uses an algorithm similar to that of the Vickrey auction for a divisible good to determine each bidder's probability of winning and expected payment.

I study the revenue properties of the probability demand mechanism in a setting that nests the benchmark independent private types case, but also allows for correlated types. Removing quasilinearity makes it difficult to explicitly solve for equilibria. Instead, I form a lower bound on bid behavior by using the normal good assumption. I use this partial characterization of bid behavior to construct a lower bound on expected revenues in the probability demand mechanism. With enough bidders, this lower bound on revenues strictly exceeds an analogously constructed upper bound on revenues from any standard auction. That is, with enough bidders the probability demand mechanism has higher expected revenues than any standard auction (Propositions 3 and 4). This class of standard auctions includes the first price, second price, and all pay auctions, as well as modifications of these formats to allow for entry fees and/or reserve prices. When there are relatively few bidders, I use a numerical example to show that the probability demand mechanism can have (nonnegligibly) greater revenues than standard auctions.

I also show that under complete information, the probability demand mechanism is an efficient mechanism. My motivation for studying the complete information case is driven by recent impossibility results regarding the dominant strategy implementation of Pareto efficient allocations for cases where bidders have non-quasilinear preferences. While implementing efficient allocations under incomplete information may prove to be impossible, under complete information any undominated Nash equilibrium of the probability demand mechanism is Pareto efficient. This is not true of standard auctions that assign the good to the highest bidder.

The rest of the paper proceeds as follows. The remainder of the Introduction relates my work to the current literature on auction design. Section 2 describes the model and specifies the assumptions I place on bidders' preferences. Section 3 motivates the use of probabilistic allocations. Section 4 outlines the construction of the probability demand mechanism. Section 5 focuses on revenue comparisons between the probability demand mechanism and standard auction formats. Section 6 provides a numerical example illustrating the practical applicability of my results. Section 7 discusses efficiency. The Appendix concludes.

1.2 *Related literature*

Most research that studies auctions with risk averse bidders fits into one of two categories: (i) comparing the performance of standard auction formats and (ii) studying the design of optimal auctions. In the first category, Matthews (1983, 1987) and Che and Gale (2006) compare first and second price auctions. These papers show that first price auctions yield higher revenue than second price auctions. The payoff environment considered by Che and Gale (2006) is closest to the one studied here. Their setting allows for risk aversion, wealth affects, and multidimensional heterogeneity.

This paper fits into the second category of papers that study the auction design problem. Maskin and Riley (1984) are the first to study the properties of revenue maximizing auctions when bidders are not quasilinear. Their paper studies the case where bidders' types are single dimensional and independent and identically distributed (i.i.d.). They show that the exact construction of the optimal mechanism depends on the distribution of types and the functional form of the bidders' utility. Their setting is general enough to include bidders with wealth effects and/or risk aversion, but is limited to cases with single-dimensional heterogeneity and i.i.d. types. This paper expands on their analysis by considering the auction design problem when bidders can have multidimensional and correlated types.

In a related line of research, Laffont and Robert (1996) and Pai and Vohra (2014) study the revenue maximizing auction design problem of Myerson (1981), but when bidders have budgets. They consider settings with i.i.d. types and show that the auctioneer can increase revenues by using randomization. In this paper, I show that with many bidders, randomization can increase revenue in a more general setting that allows for budgets as a limiting case, but also includes any case where bidders have positive wealth effects.

In my model, the multidimensionality of bidder types complicates the Myersonian approach of characterizing bidders' interim incentive constraints. Thus, I take a new approach to studying the auction design problem. Instead of explicitly characterizing equilibria, I show that by placing bounds on bid behavior, we can construct a mechanism that obtains higher revenues than standard auctions when there are many bidders. Armstrong (1999) uses a similar approach to study the problem of a multiproduct monopolist selling to a representative consumer. In Armstrong's model, consumers have multidimensional types. He describes qualitative features of an almost optimal solution for the monopolists when there are many products.

In addition, I show that my probability demand mechanism is Pareto efficient in a complete information setting. I focus on the complete information case because the recent work of Dobzinski et al. (2012) shows that when bidders have private budgets, there is no mechanism that is dominant strategy implementable, is Pareto efficient, and satisfies a no budget deficit condition. Thus, while it is impossible to obtain an efficient and detail-free mechanism with incomplete information, this paper is able to provide a prescription for efficient implementation under complete information.

2. The model

2.1 *The payoff environment*

Consider a private value auction setting for an indivisible good with a single risk neutral seller and $N \geq 2$ buyers, indexed by $i \in \{1, \dots, N\}$. Bidder i's preferences are described by utility function u_i, where I let $u_i(1, w_i)$ denote bidder i's utility with wealth w_i when she owns the object. Similarly, $u_i(0, w_i)$ denotes bidder i's utility with wealth w_i when she does not own the object. Thus,

$$u_i : \{0, 1\} \times [\underline{w}, \infty) \to \mathbb{R},$$

where $\underline{w} \in \mathbb{R} \cup \{-\infty\}$ is a lower bound on a bidder's wealth and $w_i > \underline{w}$.[1]

The object is a "good;" thus,

$$u_i(1, w) > u_i(0, w) \quad \forall w \in [\underline{w}, \infty).$$

Additionally, I assume that bidders' preferences are strictly increasing and twice continuously differentiable in wealth.

Let $k(u_i, w_i)$ be bidder i's willingness to pay for the good when she has an initial wealth w_i and a utility function u_i. That is, k is implicitly defined as

$$u_i(1, w_i - k) = u_i(0, w_i). \tag{2.1}$$

[1] If $\underline{w} \in \mathbb{R}$, then assume that $u_i(x, w) = -\infty$ if $w < \underline{w}$.

I place two additional restrictions on bidder preferences. First, I assume that the good being sold is a normal good (i.e., positive wealth effects). My notion of positive wealth effects is analogous to the notion in the divisible goods case, where a bidder's demand for the good increases as her wealth increases for a constant price level.

Assumption 1 (Positive wealth effects). *Bidder i has positive wealth effects:*

$$\frac{\partial k(u_i, w)}{\partial w} > 0, \quad \forall w \in [\underline{w}, \infty).$$

Note that Assumption 1 can also be written in terms of primitives.[2] Second, I assume that bidders have strictly declining marginal utility from money.

Assumption 2 (Risk aversion). *Bidder i has declining marginal utility of money:*

$$\frac{\partial^2 u_i(x, w)}{\partial w^2} < 0 \quad \textit{for } x = 0, 1.$$

Let $\mathcal{U}$ be the set of all utility functions that satisfy Assumptions 1 and 2. Quasilinear preferences are not included in $\mathcal{U}$ as $\partial k(w, u_i)/\partial w = 0$ and $\partial^2 u_i(x, w)/\partial w^2 = 0$. However, quasilinear preferences are a limiting case of the environment.

2.2 *Incomplete information setting*

I describe bidder i's preferences by her type $t_i \in T \subset \mathbb{R}^m$, where $m \in \mathbb{N}$ and T is compact. I let the first element represent bidder i's initial wealth level w_i. A bidder with type t_i has preferences described by the utility function $u(\cdot, \cdot, t_i)$ when her type is t_i. I assume that bidders' preferences have declining marginal utility of money and positive wealth effects. That is, for any $t_i \in \mathbb{R}^m$, $u(\cdot, \cdot, t_i) \in \mathcal{U}$. In addition, I assume u is continuous in t_i. At the same time, this setup allows for heterogeneity across risk preferences, initial wealth, and financing constraints.

Note that by assuming that a bidder's type t_i is finite dimensional, we are considering a type space T that is a subspace of the infinite-dimensional space of utility functions that satisfy Assumptions 1 and 2. We could alternatively allow bidder types to be infinite dimensional, but then we would need to define a topology on the (infinite-dimensional) type space.[3]

[2]Specifically it requires that $(\partial/\partial w)u_i(1, w-k) > (\partial/\partial w)u_i(0, w)$ when $k = k(u_i, w)$.

[3]For example, suppose instead that bidders have types t where $t \in T \subset L^p$. The difference is important only when I make revenue comparisons between the probability demand mechanism and standard auctions. I will show that there is an upper bound on revenue for any mechanism that satisfies interim individual rationality, and I then show that with a large N, the probability demand mechanism attains this bound. I establish this bound by finding the highest price a bidder is willing to pay for a unit of probability (see Section 5.1). This will be called $\overline{p}(t)$. By using the compactness of T, I show that there exists a $x \in \mathbb{R}$ such that $\overline{p}(t) < x$ for all $t \in T$. Thus, expected revenues are bounded by x for any N. If the type space was infinite dimensional, then compactness does not imply $\exists x \in \mathbb{R}$ such that $\overline{p}(t) < x$ $\forall t \in T$. Thus, if the type space was infinite dimensional, I would need an assumption that states that there exists an $x > 0$ such that $x > \overline{p}(t)$ $\forall t \in T$.

I assume that the profile of bidder types $(t_1, \ldots, t_n)$ is conditionally independent. I introduce aggregate demand uncertainty by allowing for different states of the world. There is a finite number of states of the world, $s_1, \ldots, s_J \in S$. Conditional on state s, bidder types are i.i.d. draws of a random variable with distribution function $F(t|s)$, where $F : T \times S \to [0, 1]$. There is a $g(s_j)$ probability of state s_j occurring, where $\sum_{j=1}^{J} g(s_j) = 1$. Note that if $J = 1$, this is the benchmark i.i.d. case. A bidder observes her type, but not the state of the world.

2.3 *Allocations and mechanisms*

By the revelation principle, I can limit attention to direct revelation mechanisms. A mechanism describes how the good is allocated and how transfers are made. Let $\mathbb{A}$ be the set of all feasible assignments, where

$$\mathbb{A} := \left\{ a \middle| a \in \{0, 1\}^N \text{ and } \sum_{i=1}^{N} a_i \leq 1 \right\},$$

where $a_i = 1$ if bidder i is given the object. A feasible outcome ϕ specifies both transfers and a feasible assignment: $\phi \in \mathbb{A} \times \mathbb{R}^N$. I define $\Phi := \mathbb{A} \times \mathbb{R}^N$ as the set of feasible outcomes. A (probabilistic) allocation is a distribution over feasible outcomes. Thus, an allocation α is an element of $\Delta(\Phi)$.

Let $\mathbb{E}_\alpha[t_i]$ denote the expected utility of bidder i under allocation $\alpha \in \Delta(\Phi)$ when she has type t_i. Similarly, let $\mathbb{E}_\alpha[u_0]$ be the expected revenue for the auctioneer under allocation α. A direct revelation mechanism M maps reported types to an allocation. That is,

$$M : T^N \to \Delta(\Phi).$$

3. Probabilistic allocations

I propose a probability demand mechanism that uses randomization to increase revenue over standard auctions. The value of randomization stems from the normal good assumption. With a normal good, a bidder is willing to pay the most for her first unit of probability of winning the good, when her wealth is relatively high. To formalize this intuition, I show that a bidder's demand for probabilities of winning the indivisible good is similar to a consumer's demand for a divisible normal good. If we imagine that probabilities of winning are sold at a constant per unit price p, then bidder i has a demand curve for probability units q_i that is (i) decreasing in p and (ii) positive for some values of p that exceed her willingness to pay for the indivisible good.

Consider a gamble where bidder i wins the good with probability q and pays x in expectation. I call a payment scheme *efficient* if, given a bidder's expected payment and probability of winning the good, her payments maximize her expected utility. In the efficient payment scheme, bidder i pays p_w^* and p_l^* contingent on winning or losing, respectively, where

$$(p_w^*, p_l^*) = \arg\max_{p_w, p_l} qu(1, w_i - p_w, t_i) + (1-q)u_i(0, w_i - p_l, t_i)$$
$$\text{s.t. } x = qp_w + (1-q)p_l.$$

In the probability demand mechanism, bidders' payments are structured efficiently. Thus, given bidder i's probability of winning and expected payment, the auctioneer constructs a payment scheme to maximize her expected utility. This is useful, because with many bidders, the auctioneer is able to extract the additional surplus generated by efficient payments.

Assuming efficient payments, a bidder's indirect utility function V is a function of her expected payments and the probability she wins the good. It is defined as

$$V(q, -x, t_i) := \max_{p_w, p_l} qu(1, w_i - p_w, t_i) + (1-q)u_i(0, w_i - p_l, t_i)$$
$$\text{s.t. } x = qp_w + (1-q)p_l.$$

The indirect utility function gives the maximal expected utility for bidder i conditional on winning the object with probability q and paying x in expectation.

Bidder i's indirect utility function defines her probability demand curve q, where

$$q(p, t_i) := \arg\max_{q\in[0,1]} V(q, -qp, t_i). \tag{3.3}$$

I economize notation by writing bidder i's indirect utility function as $V_i(q, -x) = V(q, -x, t_i)$ and her probability demand curve as $q_i(p) = q(p, t_i)$. A bidder's probability demand curve has similar properties to demand curves for divisible normal goods.

Proposition 1. *If bidder i has type $t_i \in T$ and willingness to pay k_i, then the follwing statements hold:*

(i) Curve $q_i(p)$ is continuous and weakly decreasing.

(ii) Curve $q_i(k_i + \epsilon) > 0$ for some $\epsilon > 0$.

The proof is given in the Appendix. The first point is similar to that made by Garratt (2012), who shows that consumers' demands of probability units of an indivisible good satisfy a law of demand. In addition, the second point will be useful for my analysis, as it shows that any bidder with positive wealth effects is willing to accept a gamble where she pays a price per unit of probability that exceeds her willingness to pay for the indivisible good.

As an example, consider a bidder i, with initial wealth 100 and preferences described by

$$u(x, w) = 4\mathbb{I}_{x=1} + \sqrt{w}.$$

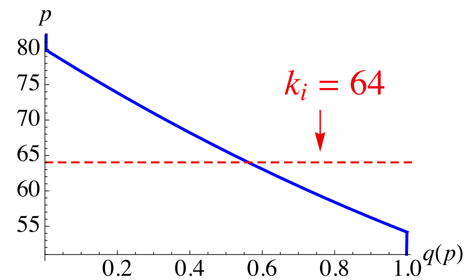

FIGURE 1. A probability demand curve.

Note that the utility function satisfies Assumptions 1 and 2.[4] If bidder i faces a gamble where she wins the good with probability q and pays x in expectation, then the efficient payment is such that she makes an equal payment in the win state and the lose state. Thus,

$$V_i(q, -x) = 4q + \sqrt{100 - x}.$$

The corresponding probability demand curve is illustrated in Figure 1.

Bidder i is only willing to pay 64 for the good, as $u(1, 100-64) = u(0, 100) = 10$. However, if probability units are sold at a constant per unit price p, then bidder i demands a positive probability of winning the good at any p that is below 80.

Standard auctions do not make use of this feature of bidder preferences. For example, in the first or second price auction, bids are bounded by a bidder's willingness to pay for the good. The auctioneer can increase her revenue by selling lotteries instead. Prior work has shown that using lotteries can increase seller revenue when bidders have budgets. I use Proposition 1 to show that we can similarly use randomization to increase the revenue for selling any normal good when there are enough bidders.

4. The probability demand mechanism

Proposition 1 shows that if a bidder is indifferent between accepting or rejecting a take-it-or-leave-it offer for the good at a price of k, there are gambles she strictly prefers where she wins the good with positive probability q and pays strictly greater than qk in expectation. A natural way to exploit the above bidder preference is to introduce randomization to the auction design. Perhaps the simplest such design is selling raffle tickets at a fixed price per ticket. For a straightforward example, imagine a case where all bidders are identical. The auctioneer sells each bidder a ticket that gives a $1/N$ probability of winning. With enough bidders, the auctioneer can set the (expected) price of the ticket to be greater than k/N (where k is a bidder's willingness to pay) and still sell all N tickets. Thus, this simple mechanism could already raise more money than a first or second price auction, where revenues are bounded by bidders' willingness to pay for the good. However, it is easy to find cases where raffles perform poorly. Determining the appropriate ticket price requires the auctioneer to have precise information on the distribution

[4]This utility function is not defined over negative wealth levels. This will not be relevant in the analysis shown below as the bidder's wealth will never approach zero. This is meant to be used as an illustration.

of bidder types. Even if the auctioneer were to know the underlying distribution of bidder types, there may be correlation across bidder types. If bidders have correlated types, in a relatively high demand state the auctioneer would want to sell more expensive raffle tickets. In a lower demand state, bidders may not want to buy tickets at this relatively high price. At the same time, the auctioneer is unable to extract information on the aggregate demand state using Cremer and McLean-style gambles because bidders are risk averse.

The probability demand mechanism sells the indivisible good as though it were a divisible good in net supply 1 sold through a Vickrey auction. It is approximately revenue maximizing with many bidders, while standard auctions and raffles are not.

4.1 *The probability demand mechanism*

In the probability demand mechanism, a bidder reports her probability demand curve $q_i(\cdot)$ and her type t_i. The reported q_i must be such that q_i is continuous, $q_i(0) = 1$, and $\lim_{p\to\infty} q_i(p) = 0$. The auctioneer uses the reported demand curves to calculate each bidder's probability of winning and expected payment. Given a bidder's probability of winning and expected payment, her payments are then structured efficiently. [5]

The probability that bidder i wins the object is calculated using the reported probability demand curves. Given the reported demand curves, the auctioneer calculates the (lowest) price for probabilities of winning the good that "clears the market." That is, she finds the (lowest) price p^* where the total reported demand for probabilities of winning the good equals 1:

$$p^* := \inf_p \left\{ p \,\middle|\, 1 = \sum_{i=1}^{N} q_i(p) \right\}. \tag{4.1}$$

The price p^* determines each bidder's probability of winning and it turns out that bidder i wins with probability $q_i(p^*)$.[6]

The price p^* is not the per unit price bidders pay for probabilities of winning the good. Instead each bidder faces a probability supply curve that represents her marginal price curve for probabilities of winning the good. The supply curve is the residual demand for probabilities of winning. It is analogous to the residual demand curve in a Vickrey auction for a divisible good. Thus, the price a bidder pays for a unit of probability depends on her rivals' actions.

[5] We could define the probability demand mechanism as a direct revelation mechanism where a bidder only reports her type to the auctioneer, and a proxy bidder then reports a demand curve for the bidder. The advantage of the current setup over this approach is that the current setup allows us to study cases where a bidder misreports her probability demand curve, without necessarily misreporting her risk preferences. In addition, the current setup illustrates the connection between the probability demand mechanism and the Vickrey auction for a divisible good.

[6] Since we assume q_i is continuous, we have that $\exists p$ such that $1 = \sum_{i=1}^{N} q_i(p)$.

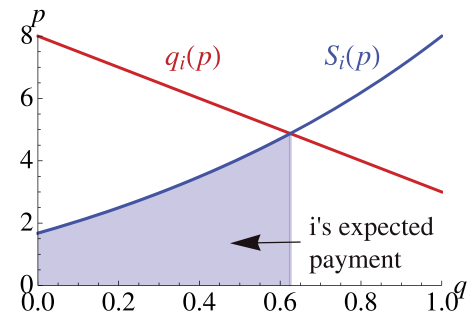

Figure 2. Expected payment in the probability demand mechanism.

Given a price p, a bidder's probability supply curve $S_i(p)$ states the amount of probability of winning the good that is not demanded by the $N-1$ other bidders:

$$S_i(p) = \begin{cases} 1 - \sum_{j \neq i} q_j(p) & \text{if } 1 - \sum_{j \neq i} q_j(p) > 0 \\ 0 & \text{otherwise.} \end{cases} \tag{4.2}$$

Bidder i's (reported) probability demand curve equals her probability supply curve at the price p^*. Thus, the market clearing price sets each bidder's probability demand curve equal to her probability supply curve.

Bidder i's expected payments are determined by treating her probability supply curve as her (expected) marginal price curve. Her expected payment to the auctioneer is X_i, where

$$X_i = \int_0^{p*} \alpha \, dS_i(\alpha). \tag{4.3}$$

I suppress notation in writing X_i; it is a function of the complete profile of reported demand curves $(q_1, \dots, q_N)$. Figure 2 illustrates this graphically.

Thus, the reported probability demand curves determine each bidder's expected transfers and probability of winning the good. If the profile of reported demand curves $(q_1, \dots, q_N)$ is such that bidder i wins with probability $q_i(p^*)$ and pays the auctioneer X_i in expectation, she makes an efficient payment. If bidder i reports her type t_i, she pays $p^*_{i,w}$ when she wins and $p^*_{i,l}$ when she loses, where

$$(p^*_{i,w}, p^*_{i,l}) = \arg\max_{p_w, p_l} q_i(p^*) u(1, w_i - p_w, t_i) + (1 - q_i(p^*)) u(0, w_i - p_l, t_i)$$
$$\text{s.t. } X_i = q_i(p^*) p_w + (1 - q_i(p^*)) p_l.$$

Definition 1 (The probability demand mechanism). The probability demand mechanism maps reported demand curves $(q_1, \dots, q_N)$ and reported types $(t_1, \dots, t_N)$ to a probabilistic allocation described by (4.1)–(4.4).

4.2 *Behavior in the probability demand mechanism*

I derive a bound on bidders' reports by showing that it is a dominated strategy for a bidder to underreport her demand for winning probabilities. The intuition for this result

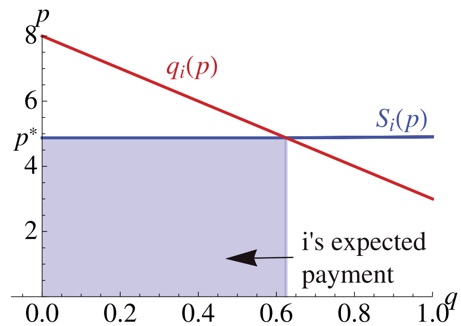

Figure 3. Expected payment when facing a perfectly elastic supply curve.

can be understood graphically. Consider the hypothetical case where bidder i faces a perfectly elastic probability supply curve and, thus, pays a constant marginal price for units of probability of winning. Let this price be p_E. The bidder seeks to maximize her expected utility given this price p_E. The solution to this maximization problem is $q_i(p_E)$, because $q_i(p_E)$ is defined as the probability of winning the good that she desires when she pays a price of p_E per unit of probability. Thus, truthful reporting is a best response. By truthfully reporting, she wins the good with probability $q_i(p_E)$ and pays $p_E q_i(p_E)$ in expectation. This case is illustrated in Figure 3.

Suppose, instead, bidder i faces a more inelastic (relative to perfectly elastic) probability supply curve. Suppose her residual probability demand curve still passes through the (arbitrary) point $(p_E, q_i(p_E))$. If bidder i truthfully reports her probability demand curve, she wins the good with the probability $q_i(p_E)$. Thus, her probability of winning the good is the same as it was when she faced the perfectly elastic supply curve. However, she pays less when she faces the more inelastic supply curve. The marginal price she pays for all but the final unit of probability she acquires is less than p_E. Thus, she pays X_i, which is less than $p_E q_i(p_E)$ for a $q_i(p_E)$ probability of winning. It is as though she faced a perfectly elastic supply curve with constant price p_E, and then is given a refund of $p_E q_i(p_E) - X_i \geq 0$. Positive wealth effects imply that this refund increases her demand of the good relative to the case where she simply pays the price of p_E per unit of probability. This case is illustrated in Figure 4.

It was a best response for bidder i to truthfully report her demand curve when she faced a constant marginal price curve, but with an upward sloping marginal price curve she has an incentive to overreport her demand curve. Since bidders report downward sloping demand curves, bidder i will always face an upward sloping marginal price curve. The precise amount that bidder i wants to overreport her demand curve depends on the elasticity of the supply curve she expects to face. Her incentive to overreport is greater when facing a more inelastic supply curve (larger "refund"). What is clear is that it is never a best reply for bidder i to underreport her demand curve. This observation allows us to use bidder i's truthful report as a lower bound on her actual report.

Proposition 2. *Assume bidder i has probability demand curve $q_i(p)$. Reporting a probability demand curve $\tilde{q}_i$, where $\tilde{q}_i(p) < q_i(p)$ for some $p \in \mathbb{R}_+$, is weakly dominated by*

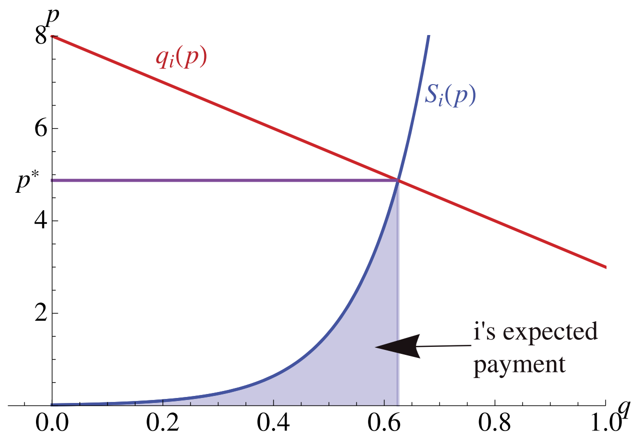

FIGURE 4. Expected payment when facing a relative inelastic supply curve.

reporting $\overline{q}$*, where*

$$\overline{q_i}(p) = \max\{q_i(p), \tilde{q}_i(p)\}.$$

Proposition 2 shows that truthful reporting can serve as a lower bound on a bidder's possible report. This lower bound on a bidder's report enables revenue comparisons between the probability demand mechanism and other auctions.

5. Revenue comparisons

With many bidders, a lower bound on expected revenues from the probability demand mechanism exceeds the expected revenues of a large class of standard auctions.

5.1 *A revenue upper bound for all mechanisms*

Consider the highest per unit price where a bidder still demands a positive probability of winning. Let $\overline{p}_i$ be such a per unit price for bidder i—her choke price for probabilities of winning. This is a function of a bidder's type,

$$\overline{p}(t_i) := \sup_p\{p|q(p, t_i) > 0\}.$$

Since preferences u are continuous in t, it follows that $\overline{p}(t_i)$ is continuous in t_i. I let f_s be the density of bidder types conditional on the state of the world being s. Similarly, let $\overline{P}(s) = \max\{\overline{p}(t)|t \in \text{supp}(f_s)\}$ be the maximal choke price given the state, and let $\overline{P}$ be the expectation of $\overline{P}(s)$,

$$\overline{P} := \mathbb{E}(\overline{P}(s)).$$

Since T is compact, $\overline{P}(s) < \infty, \forall s$ and hence, $\overline{P} < \infty$. The price $\overline{P}$ is an upper bound on the expected revenues from any interim individually rational mechanism. A mechanism with expected revenues that exceed $\overline{P}$ necessarily violates some bidder's interim individual rationality constraint because bidder i is never willing to pay more than $\overline{p}(t_i)$ for a unit of probability.

5.2 *Revenues from the probability demand mechanism*

Proposition 2 shows that with many bidders, a lower bound on expected revenue from the probability demand mechanism approaches the expected revenues upper bound $\overline{P}$. For ease of notation I write $R(N)$ as the expected revenue from truthful reporting in the probability demand mechanism when there are N bidders.

PROPOSITION 3. *For any $\epsilon > 0$, $\exists N^*$ such that for all $N > N^*$,*

$$R(N) > \overline{P} - \epsilon.$$

To see the intuition for this result, suppose that the bidder with the highest willingness to pay for a unit of probability is bidder 1. Bidder 1 is willing to pay $\overline{p}(t_1)$ for her first marginal unit of probability. Thus bidder 1 demands a strictly positive probability of winning for any price that is under $\overline{p}(t_1)$. With many bidders, there are many other bidders with types similar to bidder 1. Thus, there are many bidders who demand a positive probability of winning at a price slightly lower than $\overline{p}(t_1)$. With enough bidders, this means the residual demand for units of probability at price $\overline{p}(t_1) - \epsilon$ is zero. If a bidder does win the good with a positive probability, she pays a per unit price that exceeds $\overline{p}(t_1) - \epsilon$. Thus, expected revenues exceed $\overline{p}(t_1) - \epsilon$. In expectation, this means revenues exceed $\overline{P} - \epsilon$, which gives the above result. By combining this with Proposition 2, we see that when bidders play undominated strategies, expected revenues exceed $\overline{P} - \epsilon$ when there are many bidders.

Figure 5 illustrates this for a special case where bidders all have initial wealth of 100 and preferences u given by

$$u(x, w) = 4\mathbb{I}_{x=1} + \sqrt{w}.$$

By assuming that bidders truthfully report their probability demand curves, I obtain a lower bound on a bidder's residual probability demand. As N increases, the lower bound on the residual probability demand curves approaches the expected revenue upper bound of $\overline{P} = 80$.

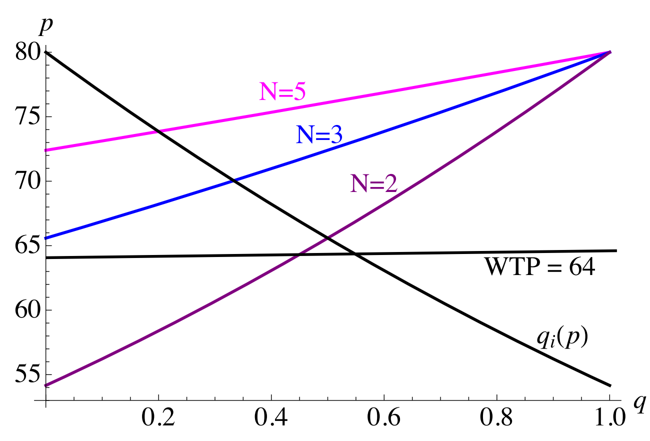

FIGURE 5. Bounds on residual probability demand curves.

5.3 *Revenue comparisons with standard auction formats*

While the probability demand mechanism approaches the expected revenue upper bound of any individually rational mechanism, the expected revenues from standard auction formats do not approach this upper bound, even with many bidders.

To see this, first consider the first and second price auctions. In each format, it is a dominated strategy for a bidder to submit a bid that exceeds her willingness to pay for the good. Thus, expected revenues are bounded by the highest willingness to pay of any bidder, which I call $\overline{K}$, where

$$\overline{K} = \mathbb{E}\Big(\max_{i=1,\dots,N} k_i\Big).$$

Yet, Proposition 1 shows that each bidder is willing to purchase a positive amount of probability of winning at a price per unit that exceeds her willingness to pay for the (entire) good. That is, if bidder i has preferences such that, $u_i(1, w_i - k_i) = u_i(0, w_i)$, then $\overline{p}(t_i) > k_i$. In the probability demand mechanism, with many bidders, the expected revenue approaches the expected highest price any agent is willing to pay for positive probability of winning the good, $\overline{P}$. Thus, the expected revenues from the probability demand mechanism $\overline{P}$ strictly exceed the expected revenues upper bound first price or second price auction $\overline{K}$ when there are sufficiently many bidders.

Corollary 1. *Assume that bidders play undominated strategies. When N is sufficiently large, expected revenues from the probability demand mechanism exceed expected revenues from the first or second price auctions.*

This result generalizes to a broad class of indirect mechanisms where bidders submit single-dimensional bids. In particular, I focus on “highest bid wins” mechanisms, where a bidder receives the object only if she submits the highest bid and leaves the auction at no cost by bidding 0. Thus, I study mechanisms where each bidder reports a bid $b_i \in \mathbb{R}_+$. The indirect mechanism M maps the N bids to a distribution over feasible outcomes:

$$M : \mathbb{R}_+^N \to \Delta(\Phi).$$

If $b_i = 0$, bidder i makes no transfers and wins the good with zero probability. This is equivalent to allowing bidders free exit from the auction. I allow for a minimum bid, which I call $b_{\min}$.

Definition 2 (Highest bid mechanism). The indirect mechanism M is a highest bid mechanism if bidder i is given the object if and only if she submits the highest bid and it is at least the minimum bid $b_{\min}$:

$$\max_{j\neq i} b_j > b_i \quad\Longrightarrow\quad a_i = 0 \quad\text{and}\quad b_i > \max_{j\neq i} b_j, \qquad b_{\min} \quad\Longrightarrow\quad a_i = 1.$$

If $b_i = 0$, then bidder i pays 0 and $a_i = 0$.

This class of mechanisms includes the first price, second price, and all pay auctions. It also includes each of these formats with entry fees or reserve prices.

Most commonly studied auction formats have the property that along the equilibrium path, the probability of a tie is zero. With a sufficient amount of heterogeneity in preferences, this is to be expected. This property is always true of any equilibrium of a first price or all pay auction. I say that an equilibrium of a highest bid mechanism is a "no-tie" equilibrium if in equilibrium there is a zero probability of a tie along the equilibrium path.

Definition 3 (No-tie equilibrium). A Bayesian Nash equilibrium of a highest bid mechanism is a *no-tie equilibrium* if in equilibrium,

$$P\Big(b_i = \max_{j\neq i} b_j | b_i > b_{\min}\Big) = 0 \quad \forall i.$$

Whether a mechanism has a no-tie equilibrium depends on the underlying distribution of preferences and states and on how the mechanism M structures payments.

For a given distribution of preferences, there is an upper bound on the expected revenues in any no-tie equilibrium of a highest bid mechanism. The upper bound is independent of the number of bidders and is strictly less than the revenue upper bound derived for any interim individually rational mechanism.

Proposition 4. *There exists an $\alpha > 0$ such that for any N, the expected revenues from any no-tie Bayesian Nash equilibrium of a highest bid mechanism are less than $\overline{P} - \alpha$.*

This shows that when there are many bidders, any no-tie equilibrium of a highest bid mechanism gives strictly lower expected revenues than the probability demand mechanism.

6. A numerical example

The results from the previous section show that the expected revenues from the probability demand mechanism exceed the expected revenues from standard auction formats when there are sufficiently many bidders. This leads to other questions. First, how many bidders are needed for the probability demand mechanism to generate greater revenues than standard auction formats? Second, how much greater are the revenues from the probability demand mechanism than other auction formats?

The answers to both questions depend on the assumed distribution of preferences. To further study these questions, I consider a particular setting that is embedded in my model: financially constrained bidders. Each bidder must borrow money to finance her payments to the auctioneer. The interest rate rises in the amount that she borrows. This is a similar setting to that studied by Che and Gale (1998). I find that revenues from the probability demand mechanism exceed those from standard auction formats, even with a small number of bidders. The differences in revenues are nonnegligible.

For the example, assume each bidder has a valuation of the good v_i, where $v_i \sim$ uniform$[5, 15]$. I depart from the quasilinear environment by assuming that bidders are financially constrained. So as to make payments, bidders borrow money from the bank. The interest rate paid on a loan of m dollars is $r(m)$, where

$$r(m) = \frac{m}{100}.$$

Thus, the bidder's payoff is given by

$$u_i(x, -m) = v_i \mathbb{I}_{x=1} - m(1 + r(m)).$$

The financing constraint is the only departure from the quasilinear environment.

Even with few bidders, the probability demand mechanism has expected revenues that exceed the revenues of standard auction formats. Using the methodology developed in Section 3, a bidder's probability demand curve can be expressed as

$$q_i(p) = \begin{cases} \max\{1, \frac{50}{p}(\frac{v_i - p}{p})\} & \text{if } \frac{50}{p}(\frac{v_i - p}{p}) > 0 \\ 0 & \text{if } p > v_i. \end{cases}$$

I compare the lower bound on expected revenues from the probability demand mechanism to the expected revenues of the first and second price auctions. I assume bidders truthfully report their demand curves to obtain the lower bound on revenues in the probability demand mechanism. Applying results of Che and Gale (1998) gives the equilibrium bidding function for the first price auction. In the second price auction, it is a dominant strategy for a bidder to bid her willingness to pay for the good.[7] Figure 6 illustrates the revenue comparisons between the three formats using Monte Carlo simulations.

The dark line marked with circles is the revenue lower bound for the probability demand mechanism. The gray line with squares represents the revenues of both the first and second price auction. The results of Che and Gale (1998) show the first price auction has greater expected revenues than the second price auction when bidders face financing constraints. In this environment, the difference in expected revenues between first and second price auctions is relatively small when compared to the expected revenue difference between either format and the probability demand mechanism. When there are four or more bidders, the lower bound on expected revenues from the probability demand mechanism exceeds the expected revenues of the first and second price auctions. Also, the difference in expected revenues between the two formats grows as the number of bidders increases. As the number of bidders increases, expected revenue from the probability demand mechanism approaches 15. Yet in the first and second price auctions, as the number of bidders increases, the expected revenues approach the highest possible willingness to pay of any bidder. Here, this is 13.24. When there are 21 or more bidders, the lower bound on revenues from the probability demand mechanism will actually exceed any bidder's willingness to pay for the (entire) good in expectation.

[7] The equilibrium bid functions are $b_f(v_i) = 10\sqrt{25 + 5/N + (N-1)/N v_i} - 50$ and $b_s(v_i) = 10\sqrt{25 + v_i} - 50$.

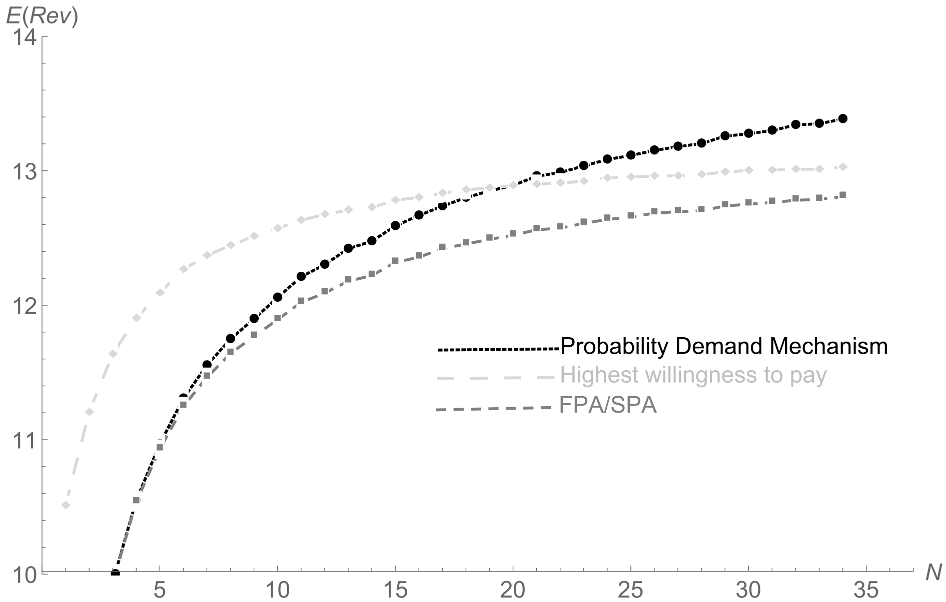

Figure 6. Revenue comparisons between formats.

7. Pareto efficiency

Now consider the question of efficient auction design. When bidders have quasilinear preferences, a second price auction implements a Pareto efficient allocation. This is not the case when we remove quasilinearity. As an example, consider a case with two bidders who each have initial wealth of 100 and preferences $u_i(x, w) = 4x + \sqrt{w}$. If the goods are sold by a second price auction, both bidders bid their willingness to pay, which is 64. Thus, the object is randomly allocated to one of the two bidders and is sold for 64. Since the bidders pay their willingness to pay, conditional on winning, both bidders have an expected utility of 10. As an alternative, suppose that each bidder buys a lottery ticket that gives a $\frac{1}{2}$ probability of winning. Assume that each lottery ticket is sold for a price of 32. In this case, a bidder gets expected utility $\frac{1}{2}(4) + \sqrt{100 - 32} \approx 10.25$, and the auctioneer's revenue is 64. Thus, the outcome of the second price auction is Pareto dominated.

Without quasilinearity, the impossibility result of Dobzinski et al. (2012) shows that under incomplete information, it is impossible to construct a mechanism that implements a Pareto efficient allocation in dominant strategies. While it is impossible to implement an efficient allocation under incomplete information, I show that with complete information, the probability demand mechanism implements a Pareto efficient allocation. Specifically, I show that for any profile of bidder types, there is a Nash equilibrium that implements an efficient allocation, and that any Nash equilibrium in undominated strategies is Pareto efficient. This is not true of standard auctions, as we see in the above example.

My notion of Pareto efficiency under complete information is equivalent to most notions of ex post Pareto efficiency used in games of incomplete information (see, for example, Section 3 of Holmström and Myerson 1983). Suppose that bidder types are $(t_1, \dots, t_N)$.[8] I say an allocation $\alpha \in \Delta(\Phi)$ is Pareto efficient if there is no other allocation $\alpha' \in \Delta(\Phi)$ that gives greater or equal expected revenues and increases at least one bidder's expected utility without decreasing any other bidder's expected utility.

Definition 4 (Pareto efficiency). An allocation $\alpha \in \Delta(\Phi)$ is Pareto efficient if $\nexists\, \alpha' \in \Delta(\Phi)$ such that

$$\mathbb{E}_{\alpha}[t_i] \geq \mathbb{E}_{\alpha'}[t_i] \quad \forall i = 1, \dots, N$$

and

$$\mathbb{E}_{\alpha}[u_0] \geq \mathbb{E}_{\alpha'}[u_0],$$

where at least one of the above statements holds with a strict inequality.

In this section, I place one additional restriction on the strategy space of the mechanism. I assume that bidders must report probability demand curves where q_i^{-1} is continuous. This ensures that a bidder faces a continuous marginal price curve, and in equilibrium all winning bidders have the same marginal willingness to pay for a unit of probability. Without this additional restriction on the strategy space, there exists an efficient Nash equilibrium of the probability demand mechanism, but I cannot show that all Nash equilibria in undominated strategies are Pareto efficient.[9]

There exists a Nash equilibrium of the probability demand mechanism that implements a Pareto efficient allocation. Let p^* be the market clearing price if bidders report their types truthfully. Suppose that bidder i plays the strategy $\tilde{q}_i$, where

$$\tilde{q}_i(p) = \begin{cases} q_i(p) & \text{if } p \geq p^* \\ 1 & \text{if } p < p^*. \end{cases}$$

Remark 1. The pure strategy profile $(\tilde{q}_1, \dots, \tilde{q}_N)$ is a Nash equilibrium of the probability demand mechanism.

Thus, there exists a Nash equilibrium in undominated strategies. Next, I show that any Nash equilibrium in undominated strategies is Pareto efficient. I use the notation U_i to describe the set of undominated strategies for bidder i.

Proposition 5. *Suppose that $(\hat{q}_1, \dots, \hat{q}_N)$ is a pure strategy Nash equilibrium of the probability demand mechanism. In addition, suppose that $\hat{q}_i \in U_i\ \forall i$. Then $\Gamma(\hat{q}_1, \dots, \hat{q}_N)$ is Pareto efficient.*

[8] It is without loss of generality to assume $w_i = 0\ \forall i$. To see this, note that a bidder with initial wealth w_i and preferences $\hat{u}_i$ behaves the same as a bidder with initial wealth 0 and utility $u_i(x, d) = \hat{u}_i(x, w_i + d)$.

[9] Assumptions 1 and 2 imply bidders' actual probability demand curves have continuous inverses. This holds because Assumptions 1 and 2 imply that a bidder's marginal willingness to pay for additional units of probability strictly decreases as her expected payment increases.

Appendix

Proof of Proposition 1. Recall that $q_i(p)$ is defined as

$$q_i(p) = \arg \max_{q \in [0,1]} V(q, -qp).$$

Since V is continuous and increasing in both arguments, then q_i is weakly decreasing in p.

Next, I show that $q_i(p)$ is continuous in p. Suppose that q_i is discontinuous at some $\hat{p} > 0$ and let $q_l := \lim_{p \to \hat{p}^+} q_i(p)$ and $q_h = \lim_{p \to \hat{p}^-} q_i(p)$. Thus, $q_h > q_\ell$ by assumption. Since V_i is continuous in both arguments, $V_i(q_l, -\hat{p}q_l) = V_i(q_h, -\hat{p}q_h) = \frac{1}{2}(V_i(q_l, -\hat{p}q_l) + V_i(q_h, -\hat{p}q_h))$.

Let x_w and x_l be the efficient payments when i wins with probability q_l and pays $\hat{p}q_l$ in expectation. Similarly, let y_w and y_l be the efficient payments when i wins with probability q_h and pays $\hat{p}q_h$ in expectation. For simplicity, I use the notation $G(w) := u_i(1, w)$ and $B(w) := u_i(0, w)$. Rewriting $(V_i(q_l, -\hat{p}q_l) + V_i(q_h, -\hat{p}q_h))/2$ gives

$$\frac{q_l G(w_i - x_w) + (1 - q_l) B(w_i - x_l) + q_h G(w_i - y_w) + (1 - q_h) B(w_i - y_l)}{2}.$$

If $x_w \neq y_w$ and/or $x_l \neq y_l$, then Jensen's inequality implies

$$\frac{q_l G(w_i - x_w) + q_h G(w_i - y_w)}{2} < \frac{q_l + q_h}{2} G\left(w_i - \frac{q_l x_w + q_h y_w}{q_l + q_h}\right)$$

and/or

$$\frac{(1 - q_l) B(w_i - x_l) + (1 - q_h) B(w_i - y_l)}{2} < \left(\frac{1 - q_l + 1 - q_h}{2}\right) B\left(w_i - \frac{(1 - q_l) x_l + (1 - q_h) y_l}{1 - q_l + 1 - q_h}\right).$$

Let $q_m = (q_l + q_h)/2$, $z_w = (q_l x_w + q_h y_w)/(q_l + q_h)$, and $z_l = ((1 - q_l) x_l + (1 - q_h) y_l)/(1 - q_l + 1 - q_h)$. Note that $\hat{p}q_m = q_m z_m + (1 - q_m) z_l$ and

$$V_i(q_m, -q_m \hat{p}) > \tfrac{1}{2}(V_i(q_l, -\hat{p}q_l) + V_i(q_h, -\hat{p}q_h)) = V_i(q_l, -\hat{p}q_l).$$

This contradicts that $q_l = \arg\max_{q \in [0,1]} V(q, -q\hat{p})$.

Finally I show that $q_i(k_i + \epsilon) > 0$ for some $\epsilon > 0$. Given that q_i is continuous and weakly decreasing, it suffices to show that $q_i(k_i) > 0$. Suppose that $q_i(k_i) = 0$. Then

$$B(w_i) = V_i(1, -k_i) = G(w_i) = V_i(0, 0) \geq V_i\left(\tfrac{1}{2}, -\tfrac{1}{2}k_i\right).$$

Since B and G are strictly increasing, continuous, and differentiable, then B^{-1} and G^{-1} are strictly increasing, continuous, and differentiable. We can then rewrite a bidder's willingness to pay as a function of her initial wealth, $k(u_i, w) = w - G^{-1}(B(w))$. By the inverse function theorem,

$$\frac{\partial k(u_i, w)}{\partial w} = 1 - \frac{1}{G'(G^{-1}(B(w)))} B'(w).$$

Positive wealth effects imply $(\partial k(u_i, w))/(\partial w) > 0$. Thus,

$$\frac{\partial k(u_i, w)}{\partial w} > 0 \quad \Longrightarrow \quad G'(w - k(u_i, w)) > B'(w).$$

At $w = w_i$ this implies that

$$G'(w_i - k_i) > B'(w_i).$$

Thus, for a sufficiently small $\epsilon > 0$,

$$\tfrac{1}{2}(G(w_i - k_i + \epsilon) + B(-\epsilon)) > \tfrac{1}{2}(G(w_i - k_i) + B(w_i)) = u_i(0, w_i) = V_i(0, 0).$$

The definition of V_i implies

$$V_i\left(\tfrac{1}{2}, -\tfrac{1}{2}k_i\right) \geq \tfrac{1}{2}(G(w_i - k_i + \epsilon) + B(-\epsilon)) > V_i(0, 0).$$

This contradicts that $V_i(0, 0) \geq V_i(\frac{1}{2}, -\frac{1}{2}k_i)$. □

Proof of Proposition 2. I prove this by showing that if stating the demand curve $\tilde{q}_i$ instead of $\overline{q}_i$ does change bidder i's payoff, it must lower the payoff.

Consider a case where i faces a perfectly elastic residual probability demand curve (i.e., a constant marginal price per unit of probability). It is a best response for her to truthfully reveal her demand curve. Let $p_E > 0$ be the constant marginal price for units of probability. If her payoff is changed by reporting $\tilde{q}_i$, then $\tilde{q}_i(p_E) < \overline{q}_i(p_E)$. By the definition of $\overline{q}_i$, then $\overline{q}_i(p_E) = q_i(p_E)$. That is, if reporting $\tilde{q}_i$ does change her payoff, it is the case that, $\tilde{q}_i$ is strictly below her demand for probability at p_E. By the construction of the probability demand curve, truthful reporting is a best response to a perfectly elastic supply curve. Thus, she decreases her payoff by reporting type $\tilde{q}_i$. Recalling that V_i is her indirect utility function under efficient payments, it follows that

$$V_i(q_i(p_E), -p_E q_i(p_E)) \geq V_i(q, -p_E q) \quad \text{for any } q \in (0, 1).$$

Now consider instead that bidder i faces a more inelastic (relatively to perfectly elastic) residual probability demand curve. Assume that the supply curve is such that $S_i(p_E) = q_i(p_E)$. Once again if her payoff is changed by reporting $\tilde{q}_i$, then $\tilde{q}_i(p_E) < \overline{q}_i(p_E) = q_i(p_E)$, using the same argument as before. Thus, she wins with a lower probability by reporting $\tilde{q}$. Assume that if she reports $\tilde{q}_i$, she pays $\tilde{X}$ in expectation and wins with probability $\tilde{q}(\tilde{p}^*)$, where $\tilde{q}_i(\tilde{p}^*) = S_i(\tilde{p}^*)$. Since her marginal price is strictly below p_E, $\tilde{X} < p_E \tilde{q}_i(\tilde{p}^*)$. If she instead reports $\overline{q}_i$, she wins with probability $q_i(p_E)$ and pays a marginal price below p_E for the incremental probability of winning gained by reporting $\overline{q}_i$. Thus, she pays $X \leq \tilde{X} + p_E(q_i(p_E) - \tilde{q}_i(\tilde{p}^*))$. Let $Y = p_E \tilde{q}_i(\tilde{p}^*) - \tilde{X} > 0$. Thus, it is sufficient to show that

$$V_i(q_i(p_E), -p_E q_i(p_E) + Y) \geq V_i(\tilde{q}_i(\tilde{p}^*), -p_E \tilde{q}_i(\tilde{p}^*) + Y).$$

Since we have already shown that the above expression holds true at $Y = 0$, it is then sufficient to show that when $Y > 0$ and $q \leq q_i(p_E)$,

$$\frac{d}{dq} V_i(q, -p_E q + Y) \geq 0.$$

Since the above function is concave in q (see Proof of Proposition 1), it suffices to show this at $q = q_i(p_E)$. Note that

$$\frac{d}{dq}V_i(q, -p_E q + Y) = \frac{\partial}{\partial q}V_i(q, -p_E q + Y) - p_E \frac{\partial}{\partial Y}V_i(q, -p_E q + Y).$$

When $Y = 0$, the necessary first order condition defining $q_i(p_E)$ implies that

$$\frac{d}{dq}V_i(q, -p_E q) = 0 \quad \text{if } q = q_i(p_E).$$

Thus, it suffices to show that increasing Y does not decrease the above derivative:

$$\frac{d}{dY}\frac{d}{dq}V_i(q, -p_E q + Y) \geq 0 \quad \text{if } q = q_i(p_E), Y > 0.$$

Rewriting the above expression, we find

$$\frac{d}{dY}\frac{d}{dq}V_i(q, -p_E q + Y) = \frac{\partial}{\partial Y}\frac{\partial}{\partial q}V_i(q, -p_E q + Y) - p_E \frac{\partial^2}{\partial Y^2}V_i(q, -p_E q + Y).$$

Note that by the envelope theorem, $(\partial/\partial q)V_i(q, -p_E q + Y) = u_i(1, w - x_w) - u_i(0, w - x_\ell)$, where x_w (or x_ℓ) is the efficient payment conditional on winning (or losing) with probability q and paying $-p_E q + Y$ in expectation. As Y increases, the efficient payments made when winning and losing both decrease; thus, bidder i finishes with a greater wealth conditional on winning. Since this increases her utility conditional on winning, then $(\partial/\partial Y)(\partial/\partial q)V_i(q, -p_E q + Y) \geq 0$. The second term is negative since bidders have declining marginal utility of money. Thus,

$$\frac{d}{dY}\frac{d}{dq}V_i(q, -p_E q + Y) \geq 0 \quad \text{if } q = q_i(p_E), Y > 0.$$

This implies,

$$V_i(q_i(p_E), -p_E q_i(p_E) + Y) \geq V_i(\tilde{q}_i(\tilde{p}^*), -p_E \tilde{q}_i(\tilde{p}^*) + Y),$$

which is what we wanted to show. □

Proof of Proposition 3. I use the notation, $q(p, t_i)$ to represent the probability demand curve for a bidder with type t_i. Fix $\epsilon > 0$. Let $\tau(\epsilon, \delta, s) := \{t | q(\overline{P}(s) - \epsilon, t) > \delta, t \in \text{supp}(f)\}$. This is the set of all types that demand at least a δ probability of winning in state s, at price $\overline{P}(s) - \epsilon$. By the definition of $\overline{P}(s)$, the set $\tau(\epsilon, \delta, s)$ is nonempty when $\delta > 0$ is sufficiently small.

Note that

$$\sum_{i=1}^{N} q_i(\overline{P}(s) - \epsilon, t_i) \geq \delta \sum_{1=i}^{N} \mathbb{I}_{t_i \in \tau(\epsilon, \delta, s)}.$$

This states that the total demand at price $\overline{P}(s) - \epsilon$ is greater than δ times the number of bidders whose demand strictly exceeds δ at $\overline{P}(s)$.

Suppose that $\delta\sum_{1=i}^{N}\mathbb{I}_{t_i\in\tau(\epsilon,\delta,s)} > 2$. It follows that $S_i(\overline{P}(s)-\epsilon)=0$ for all i. That is, if at price $\overline{P}(s)-\epsilon$, all bidders demands for probabilities of winning exceed 2, then there is zero residual demand for each bidder at price $\overline{P}(s)-\epsilon$. Thus, if i wins with positive probability, she pays a marginal price per unit that exceeds $\overline{P}(s)-\epsilon$. Thus, $X_i\geq q_i(p^*)(\overline{P}(s)-\epsilon)$ for all i. Since $\sum q_i(p^*)=1$, this implies that total expected transfers exceed $\overline{P}(s)-\epsilon$:

$$\sum_{i=1}^{N} X_i > \overline{P}(s)-\epsilon.$$

When $\delta>0$ is sufficiently small, there is a strictly positive probability that a randomly drawn bidder has type $t\in\tau(\epsilon,\delta,s)$. That is, there is a positive probability that a randomly drawn bidder demands at least a δ probability of winning the good when she pays a price of $\overline{P}-\epsilon$. Represent this probability as

$$\nu(\epsilon,\delta,s)=\int_{t\in\tau(\epsilon,\delta,s)} f_s(t)\,dt.$$

Suppose $\delta>0$ is such that $\nu(\epsilon,\delta,s)>0$. Then, by the law of large numbers, as $N\to\infty$ the probability that the sum $\delta\sum_{1=i}^{N}\mathbb{I}_{\tau\in\upsilon(\epsilon,\delta,s)}$ exceeds 2 approaches 1. Note that we can apply the law of large numbers as $\mathbb{I}_{t_i\in\tau(\epsilon,\delta,s)}$ is a Bernoulli random variable that is independent of i and equals 1 with probability $\upsilon(\epsilon,\delta,s)$.

Thus, for any $\alpha\in(0,1)$, there is a finite $N(\alpha)$ such that there is a $1-\alpha$ probability that $\delta\sum_{1=i}^{N}\mathbb{I}_{t_i\in\upsilon(\epsilon,\delta,s)} > 2$. This holds for all s. Since α and ϵ are arbitrary, let them be arbitrarily close to 0 when N is sufficiently large. Thus with a sufficiently large N, total payments exceed $\overline{P}(s)-\epsilon$ with probability $1-\alpha$, where both α and ϵ are arbitrarily small. Since there are only finitely many states of the world, we can say that for any $\epsilon>0$, $\alpha\in(0,1)$, there is an N that is sufficiently large such that there is at least a $1-\alpha$ probability that revenues exceed $\overline{P}(s)-\epsilon$ $\forall s\in S$. This yields the desired result. □

Proof of Proposition 4. Let $V(q,-x,t_i)$ be the indirect utility function of type $t_i\in T$, under efficient payments. Before I begin the proof, it will be useful to show that $\forall t_i$, $\exists\beta>0$ such that

$$V(0,0,t_i) > V(q,\beta-q\overline{p}(t_i),t_i)$$

for any $q\geq\frac{1}{2}$. That is, any gamble where type t_i wins with probability $q\geq\frac{1}{2}$ and pays $q\overline{p}(t_i)-\beta$ necessarily violate her individual rationality constraint. This is useful, because then we know that if bidder i wins with a high interim probability, she must pay a price per unit of probability that is strictly below $\overline{p}(t_i)-\beta$.

Recall that we have already shown that $V(q,-qp,t_i)$ is strictly concave in q for any t_i (this is shown in Proof of Proposition 1). In addition, since $\overline{p}(t_i)$ is bidder i's maximal willingness to pay for a unit of probability, then

$$\frac{d}{dq}V(q,-q\overline{p}(t_i))<0\quad\forall q\in(0,1).$$

If this did not hold, then bidder i would demand a positive probability when the price of units is $\overline{p}(t_i)$. Thus,

$$V(0,0,t_i) > V\left(\tfrac{1}{2}, -\tfrac{1}{2}\overline{p}(t_i), t_i\right) > V(q, -q\overline{p}(t_i), t_i) \quad \forall q > \tfrac{1}{2}.$$

Since V is continuous, then $\forall q \in [\frac{1}{2}, 1]$, $\exists R(q, t_i) > 0$ such that

$$V(0,0,t_i) = V(q, R(q,t_i) - q\overline{p}(t_i), t_i).$$

Let $\underline{R}(t_i) = \min_{q\in[\frac{1}{2},1]} R(q,t_i)$ and $\alpha = \min_{t_i\in T}(\underline{R}(t_i))/2$. Thus, we have that $\forall t_i$, $\exists \alpha > 0$ such that

$$V(0,0,t_i) > V(q, \alpha - q\overline{p}(t_i), t_i).$$

Let $\mathcal{Q}_i(t_i)$ be i's interim probability of winning. I analogously define $\mathcal{X}_i(t_i)$ as bidder i's interim expected payment. Let Q be the probability that, from the ex ante perspective, the good is won:

$$\sum_{i=1}^{N} \int_{t\in T} \mathcal{Q}_i(t) f(t)\, dt = Q.$$

Feasibility requires that $Q \leq 1$. I assume $Q > 0$, or else individual rationality requires that expected revenue is 0. Let $\underline{p} = \min_{t\in T} \overline{p}(t)$. This is the lowest amount that any bidder is willing to pay for her first marginal unit of probability. Note that

$$(1-Q)\underline{p} + \sum_{i=1}^{N} \int_{t\in T} \mathcal{Q}_i(t)\overline{p}(t) f(t)\, dt \leq \overline{P}.$$

The right hand side is the highest expected willingness to pay (for a unit of probability) of any bidder. The left hand side is the expected willingness to pay of the winner if we instead gave the object to a bidder with the lowest possible willingness to pay whenever the direct mechanism states the object is not sold. In other words, the left hand side is the expected value of the winner's $\overline{p}(t)$ term if, instead, we always assign the good, but do not necessarily give the good to the bidder with the highest $\overline{p}(t)$ term.

Let $\underline{g} := \min_{s=1,\dots,S} g(s)$. That is, each state occurs with a probability of at least $\underline{g}$. Given the state of the world s, let $G(b|s)$ be the distribution of highest submitted bids. Note that $G(b|s)$ is continuous and weakly increasing over $(b_{\min}, \infty)$ by the no-tie assumption. Let $b^*(s) = \inf_{b\geq b_{\min}}\{b | G(b|s) \geq \frac{1}{2}\}$. That is, given state s, there is at least a $\frac{1}{2}$ probability that the highest bid is below $b^*(s)$. Let $\overline{b}^* = \max_{s\in S} b^*(s)$. Thus, if $b_i(t) \geq \overline{b}^*$, then $\mathcal{Q}_i(t) \geq \frac{1}{2}\forall i$, as there is at least a $\frac{1}{2}$ probability that $\overline{b}^*$ is a winning bid in any state of the world.

Suppose that $\overline{b}^* > b_{\min}$. Let A_i be the set of all types for bidder i such that $A_i := \{t | \mathcal{Q}_i(t) \geq \frac{1}{2}\}$. We know that if t is such that $b_i(t) \geq \overline{b}^*$, then $t \in A_i$. Thus, the probability the good is won by a bidder with type $t \in A_i$ is

$$\sum_{i=1}^{N} \int_{t\in A_i} \mathcal{Q}_i(t) f(t)\, dt \geq \frac{1}{2}\underline{g},$$

because there is a $\frac{1}{2}$ probability that the good is won by a bidder who bids $b > \overline{b}^*$ in at least one state.

Ex ante expected revenues are

$$\sum_{i=1}^{N} \int_{t\in T} \mathcal{X}_i(t) f(t)\, dt = \sum_{i=1}^{N} \int_{t\in A_i} \mathcal{Q}_i(t) \frac{\mathcal{X}_i(t)}{\mathcal{Q}_i(t)} f(t)\, dt + \sum_{i=1}^{N} \int_{t\in A_i^C} \mathcal{Q}_i(t) \frac{\mathcal{X}_i(t)}{\mathcal{Q}_i(t)} f(t)\, dt.$$

Note that if $t \in A_i$, then individual rationality requires that $(\mathcal{X}_i(t))/(\mathcal{Q}_i(t)) \leq \overline{p}(t) - \beta$. In addition, individual rationality requires that $\mathcal{X}_i(t)/\mathcal{Q}_i(t) \leq \overline{p}(t)\ \forall t \in T$. Thus, we can rewrite the right hand side as

$$\sum_{i=1}^{N} \int_{t\in T} \mathcal{Q}_i(t)\overline{p}(t) f(t)\, dt - \beta \sum_{i=1}^{N} \int_{t\in A_i} \mathcal{Q}_i(t) f(t)\, dt \leq \sum_{i=1}^{N} \int_{t\in T} \mathcal{Q}_i(t)\overline{p}(t) f(t)\, dt - \frac{1}{2}\beta \underline{g},$$

where the inequality follows because $\sum_{i=1}^{N} \int_{t\in A_i} \mathcal{Q}_i(t) f(t)\, dt \geq \frac{1}{2}\underline{g}$. Recalling that

$$\sum_{i=1}^{N} \int_{t\in T} \mathcal{Q}_i(t)\overline{p}(t) f(t)\, dt \leq \overline{P},$$

we can rewrite the right hand side of the above expression and say that expected revenue is bounded by

$$\overline{P} - \frac{1}{2}\beta\underline{g}.$$

Now suppose instead that $\overline{b}^* \leq b_{\min}$. Then there is at least a 0.5 probability that the highest bid does not exceed $b_{\min}$. Thus, $Q \leq \frac{1}{2}$. Recall that earlier we showed that

$$(1 - Q)\underline{p} + \sum_{i=1}^{N} \int_{t\in T} \mathcal{Q}_i(t)\overline{p}(t) f(t)\, dt \leq \overline{P}$$

and that revenue is bounded by $\sum_{i=1}^{N} \int_{t\in T} \mathcal{Q}_i(t)\overline{p}(t) f(t)\, dt$. Thus, expected revenue is bounded by

$$\sum_{i=1}^{N} \int_{t\in T} \mathcal{Q}_i(t)\overline{p}(t) f(t)\, dt \leq \overline{P} - (1 - Q)\underline{p}.$$

Because $Q \leq \frac{1}{2}$, we have that expected revenue is bounded by $\overline{P} - \frac{1}{2}\underline{p}$. Thus, revenues are bounded by $\min\{\overline{P} - \frac{1}{2}\underline{p}, \overline{P} - \frac{1}{2}\beta\underline{g}\}$. Letting $\alpha = \max\{\frac{1}{2}\underline{p}, \frac{1}{2}\underline{dg}\}$ yields our desired result. □

Proof of Remark 1. Suppose that bidder i reports demand curve $\hat{q}_i$. If bidder i wins the same number of probability units by reporting $\hat{q}_i$, her payoff is unchanged. Suppose that bidder i wins a strictly greater number of probability units q_h by bidding $\hat{q}_i$. She pays X_h in expectation for these q_h probability units. Since $S_i(p) = 0$ if $p < p^*$, then $X_h \geq p^* q_h$. Recall that by definition of q_i,

$$V_i(q_i(p^*), -p^* q_i(p^*)) \geq V_i(q_h, -p^* q_h) \geq V_i(q_h, -X_h).$$

Thus, reporting $\tilde{q}_i$ is a best reply to $\tilde{q}_{-i}$. □

Proof of Proposition 5. Suppose that market clearing price bidders report $(\hat{q}_1, \dots, \hat{q}_N)$ is $\hat{p}^*$. Thus, the marginal price of an additional unit of probability is $\hat{p}^*$. Suppose that bidder i wins $\hat{q}_{i,\mathrm{NE}} := \hat{q}_i(\hat{p}^*)$ units of probability and pays $\hat{x}_i$ in expectation in the Nash equilibrium (NE). Hence, bidder i best responds by reporting $\hat{q}_i$ when her rivals report $\hat{q}_{-i}$. Her indirect utility is then

$$V_i(\hat{q}_{\mathrm{NE}}, -\hat{x}_i).$$

If $\hat{q}_{i,\mathrm{NE}} = 0$, then $\hat{p}^*$ is such that $q_i(\hat{p}^*) = 0$. If not, then $q_i(\hat{p}^*) > 0$ and bidder i can increase her payoff by reporting her demand truthfully. Recall that bidder i faces a (continuous) marginal price curve for units of probability. Note that continuity is guaranteed by the fact that all bidders report probability demand curves that have continuous inverses. Thus, the necessary first order condition implies that if $\hat{q}_{i,\mathrm{NE}} > 0$, then

$$\frac{\partial V_i(\hat{q}_{\mathrm{NE}}, -\hat{x}_i)}{\partial q} = \hat{p}^* \frac{\partial V_i(\hat{q}_{\mathrm{NE}}, -\hat{x}_i)}{\partial x}.$$

Let $\omega_i = \hat{q}_{i,\mathrm{NE}} \hat{p}^* - \hat{x}_i$. Suppose that instead of starting with initial wealth 0, each bidder starts with an initial wealth of ω_i. Then $\hat{p}^*$ is the Walrasian equilibrium for units of probability. This holds as the function $V_i(q, x - qp)$ is concave in q (shown in Proof of Proposition 1) and the necessary first order condition stated above holds as $\omega_i - \hat{q}_{i,\mathrm{NE}} \hat{p}^* = \hat{x}_i$. The first welfare theorem then implies that this is Pareto efficient. □

References

Armstrong, Mark (1999), "Price discrimination by a many-product firm." *Review of Economic Studies*, 66, 151–168. [56]

Che, Yeon-Koo and Ian Gale (1998), "Standard auctions with financially constrained standard auctions with financially constrained bidders." *Review of Economic Studies*, 65, 1–21. [67, 68]

Che, Yeon-Koo and Ian Gale (2006), "Revenue comparisons for auctions when bidders have arbitrary types." *Theoretical Economics*, 1, 95–118. [55]

Dobzinski, Shahar, Ron Lavi, and Noam Nisan (2012), "Multi-unit auctions with budget limits." *Games and Economic Behavior*, 74, 486–503. [56, 69]

Fazzari, Steven M., R. Glenn Hubbard, and Bruce C. Petersen (1988), "Financing constraints and corporate investment." Brookings Papers on Economic Activity, 141–195. [54]

Garratt, Rodney (2012), "Lotteries and the law of demand." *New Insights into the Theory of Giffen Goods, Lecture Notes in Economics and Mathematical Systems*, 655, 161–171. [59]

Holmström, Bengt and Roger B. Myerson (1983), "Efficient and durable decision rules with incomplete information." *Econometrica*, 51, 1799–1819. [70]

Laffont, Jean-Jacques and Jacques Robert (1996), "Optimal auction with financially constrained buyers." *Economic Letters*, 52, 181–186. [55]

Maskin, Eric and John Riley (1984), "Optimal auctions with risk averse buyers." *Econometrica*, 52, 1473–1518. [55]

Matthews, Steven A. (1983), "Selling to risk averse buyers with unobservable tastes." *Journal of Economic Theory*, 30, 370–400. [55]

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

Myerson, Roger B. (1981), "Optimal auction design." *Mathematics of Operations Research*, 6, 58–73. [55]

Pai, Mallesh M. and Rakesh Vohra (2014), "Optimal auctions with financially constrained bidders." *Journal of Economic Theory*, 150, 383–425. [55]

Co-editor George J. Mailath handled this manuscript.

Manuscript received 2 September, 2014; final version accepted 21 December, 2015; available online 13 January, 2016.