

Auctions with Unique Equilibria

SHUCHI CHAWLA, University of Wisconsin - Madison, shuchi@cs.wisc.edu.

JASON D. HARTLINE, Northwestern University, hartline@eecs.northwestern.edu.

We study Bayes-Nash equilibria in a large class of anonymous order-based auctions. These include the generalized first-price auction for allocating positions to bidders, e.g., for sponsored search. We show that when bidders' values are independent and identically distributed there is a unique Bayes Nash equilibrium; This equilibrium is symmetric and efficient. Importantly, our proof is simple and structurally revealing. This uniqueness result for the generalized first-price auction is in stark contrast to the generalized second-price auction where there may be no efficient equilibrium. This result suggests, e.g., that first-price payment semantics may have advantages over second-price payment semantics. Our results extend also to certain models of risk aversion.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]: Economics

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Bayes-Nash equilibrium, first-price auction, all-pay auction, GFP, GSP, price of anarchy

1. INTRODUCTION

The potential for multiplicity of equilibria, especially when some are undesirable, introduces considerable challenges in game theory. How do the players coordinate on one of these equilibria? How can the bad equilibria be removed? How much worse are the bad from the good equilibria? These issues are also important in mechanism design as a designed mechanism may have multiple equilibria, some of which do not implement the designer's objective. Finally, multiplicity of equilibria can make ex-post econometric analysis challenging.

In this paper we seek to understand properties of mechanisms that guarantee the existence of a unique equilibrium.

The standard equilibrium concept for games of incomplete information like auctions is Bayes-Nash equilibrium (BNE). It is assumed that the bidders' values are drawn from a known distribution and the bidders' strategies, which map their values to bids, are in best response to each other for the given distribution. Even the simplest of auctions can admit multiple Bayes-Nash equilibria. Consider, for example, the second-price (a.k.a., Vickrey) auction where the winner is the highest bidder and she is charged the second highest bid. This auction is popular in academic research because it is efficient (the bidder with the highest value receives the item) and truthful bidding is a dominant strategy. Unfortunately, the second-price auction has other equilibria besides this dominant strategy equilibrium. If the value distributions are supported on $[0, 1]$ then any one agent bidding one and all others bidding zero is a Bayes-Nash equilibrium. In contrast, the non-truthful first-price auction, where the highest bidder is the winner and she is charged her bid, and the all-pay auction, where the highest bidder is the winner and all bidders are charged their bids, often have a unique equilibrium [Lebrun 2006; Maskin and Riley 2003; Amann and Leininger 1996]. For

Shuchi Chawla is supported in part by NSF awards CCF-1101429 and CNS-0905134. Jason Hartline is supported in part by NSF awards CCF-1101717 and CCF-0830773.

Permission to make digital or hardcopies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credits permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

EC'13, June 16–20, 2013, Philadelphia, USA. Copyright © 2013 ACM 978-1-4503-1962-1/13/06...\$15.00

instance, when the bidders' values are i.i.d. from a distribution with bounded support, it is easy to find symmetric equilibrium that is unique and efficient [Maskin and Riley 1984].

The uniqueness of equilibria has been studied extensively in economic theory and yet is poorly understood. To understand the difficulty in characterizing equilibria, consider a two-bidder auction. In any equilibrium, each bidder plays a best response to the other bidder's strategy. This best response can be expressed in the form of a differential equation with both the bidders' strategies as variables. This leads to a pair of ordinary differential equations (ODEs), the solution to which determines an equilibrium of the game. Analyzing the properties of these ODEs can lead to characterizations of auctions that admit unique equilibria. However, this approach is highly sensitive to small changes in initial conditions, rules of the mechanism, etc.; furthermore, obtaining general solutions for settings beyond simple single-item first-price and all-pay auctions does not seem to be analytically tractable via this approach.

We focus on symmetric settings where agents are a priori identical and the mechanism is anonymous, that is, it treats the agents identically and a permutation of agents' bids permutes the outcome (allocations and payments). Even in symmetric settings, mechanisms can admit many different asymmetric equilibria in which some agents bid aggressively and the others respond by bidding conservatively.

Our main contribution is a new approach for ruling out asymmetric equilibria in a large class of symmetric auction games where the outcome for each agent is determined only by her own bid and her rank among the other bidders. Importantly, this restriction explicitly excludes "second-price" payment rules. This approach is structurally revealing in that our assumptions are explicitly used and it is intuitively clear why the assumptions are necessary for the approach.

Let us illustrate our approach through an example. Consider a two agent setting and a pair of pure strategies in BNE that are strictly increasing continuous functions of value. We rule out the possibility that the strategies cross at two distinct points without being equal in the middle. Induction would then imply that the strategies are equal everywhere. Central to the approach are two formulas for the utility of an agent. The first formula comes directly from the auction rules in terms of the agent's value, allocation probability, and expected payment. The second formula comes from the characterization of Bayes-Nash equilibria [Myerson 1981] which requires that an agent's utility is equal to the integral of her equilibrium probability of winning as a function of her value. The following argument proves the claim:

- (1) If the bids of the two agents are equal at some value (i.e. their strategies "cross"), then by the first formula for utility they obtain the same utility at this value.
- (2) If one of the agents bids strictly higher than the other agent at some value, then she has a strictly higher probability of winning; therefore, the second formula for utility implies that the higher-bidding agent obtains strictly higher utility than the lower-bidding agent.
- (3) For an interval with end points at which the strategies cross and within which one strictly exceeds the other, the difference in utilities at the end points must be equal by (1) and distinct by (2). Thus, such an interval cannot occur.

This argument can be extended with some effort to mixed equilibria and strategies that are only weakly increasing or discontinuous, as well as to some models of risk aversion in the agent preferences.

The *position auctions* popularized by Google and Yahoo! give an illustrative example of auctions for which our theorem does and does not apply. An ordered list of advertisements is shown along side search results on Internet search engines. The probability that an ad is clicked on is a decreasing function of the position of the ad in the list. Auctions are used to determine the assignment of advertisements to positions. In the generalized second price

auction (GSP) introduced by Google, the advertisers are assigned to slots in decreasing order of bids and each advertiser is charged the bid of the next lower advertiser if her ad is clicked on. In the generalized first price auction (GFP) introduced by Overture (acquired by Yahoo! but no longer in use [Jansen and Mullen 2008]) the same ranking is used, however, an advertiser is charged her bid if she is clicked on.

Both GSP and GFP have the desirable property that the exact click rates of the slots need not be known; this fact gives a significant motivation to prefer these auctions over the dominant strategy Vickrey-Clarke-Groves (VCG) auction which requires knowledge of the click rates to calculate payments [see Varian 2009]. Neither GSP nor GFP have truth-telling as a dominant strategy. Bayes-Nash analysis of GSP is quite challenging even with i.i.d. distributions [Gomes and Sweeney 2009]. First, the equilibria are not unique. Second, symmetric equilibria may not exist. Third, there may be no efficient equilibria (i.e., the so called *price of stability* is greater than one). Fourth, some equilibria may be worse than others (i.e., the so called *price of anarchy* is greater than the price of stability). Our theorem shows, in contrast, that (with i.i.d. and bounded support distributions) the simple symmetric equilibrium of the GFP auction is unique and efficient (i.e., the price of anarchy is one).¹

Many relevant auction environments are structurally equivalent to position auctions and for these environments calculating the corresponding weights of the positions (i.e., click rates) may not be easy. For instance, when the bidders are randomly assigned to roles in a matroid set system then the feasibility constraint induced is a position auction but there is no obvious way to calculate the weights [Hartline and Yan 2011]. In such settings, our result shows that GFP implements the same outcome as the VCG mechanism in a unique equilibrium without knowledge of the weights and, therefore, might be preferable.

Related work. Several works have noted that first-price and all-pay auctions generally admit unique equilibria. Lebrun [2006] and Bajari [2001] prove uniqueness for first-price auctions under certain assumptions on the value distribution. Amann and Leininger [1996] prove similar results for all-pay auctions with two agents. Maskin and Riley [2003] study two-agent auctions with general pricing rules that are bid-dependent. All of these works are limited in the allocation rules they consider – highest bid wins subject to a reserve price. On the other hand, they also consider asymmetric settings while we do not.

Our work is closely related to the canonical problem in Bayesian implementation theory [Jackson 2001]: given a social choice rule, is there a mechanism such that *all* of its BNE implement the social choice rule? While that literature presents near-complete characterizations for the class of social choice rules that are implementable, the actual mechanisms that implement them are of theoretical interest only. In contrast, we show that rank-based social choice rules are implementable by a practical and prevalent mechanism.

While most of the literature in algorithmic mechanism design has focused on the design of dominant strategy truthful mechanisms, several settings necessitate considering first-price and all-pay auctions. This includes, for example, settings where agents have budget constraints [e.g., Bhattacharya et al. 2010], and crowdsourcing contests where all players make a contribution [e.g., Chawla et al. 2012; DiPalantino and Vojnovic 2009]. All of these works study properties of one canonical equilibrium for the mechanism and do not address the possibility of existence of other equilibria. Our work shows that in symmetric settings the canonical equilibrium is in fact the only BNE.

The most general mechanisms that our work applies to are position auctions. Varian [2007] and Edelman et al. [2007] independently introduced the generalized second-price auction, a.k.a. position auction, as a simplified model for sponsored search. Several works [Cara-

¹GFP has been criticized in literature for leading to unstable bidding behavior in the sponsored search context. This can generally be attributed to informational asymmetries [Edelman et al. 2007]. Our paper does not address this aspect.

giannis et al. 2012; Paes Leme and Tardos 2010; Lucier and Paes Leme 2011] have studied the inefficiency of these mechanisms in the full information setting in the form of the price of anarchy (the efficiency of the worst Nash equilibrium) and the price of stability (the efficiency of the best Nash equilibrium). Gomes and Sweeney [2009] were the first to study these mechanisms in a Bayesian setting. They showed that GSP does not always admit a BNE. The best known bound on the Bayesian price of anarchy for GSP is 2.927 due to Caragiannis et al. [2012]. When bidders have non-identical distributions, Syrgkanis and Tardos [2013] bound the price of anarchy of GFP by $2e/(e-1) \approx 3.16$; in contrast, for identical continuous bounded-support distributions our results imply the price of anarchy is one (i.e., GFP is optimal).

Many other auction environments are equivalent to position auctions. Hartline and Yan [2011] show that matroid permutation environments are equivalent to position auctions. Moreover, Ha and Hartline [2012] show that within a general downwards-closed permutation environment a position auction can be found subject to which optimization gives a good approximation to the optimal mechanism for the original problem.

2. DEFINITIONS AND ASSUMPTIONS

2.1. Mechanisms, strategies, and Bayes-Nash equilibria

In a single-parameter mechanism design problem, we have n agents, each with a value v_i drawn independently from a known distribution F_i . A **mechanism** (\mathbf{x}, \mathbf{p}) for this problem maps a vector of bids, \mathbf{b} , to allocations $x_i(\mathbf{b})$ that specify the probability with which agent i gets service, and prices $p_i(\mathbf{b})$ that specify the expected amount that agent i needs to pay.

A **pure strategy** s_i for agent i maps the value of the agent to a bid. A simple property of pure strategies in BNE is that they are non-decreasing functions of value: $v_1 < v_2$ implies $s_i(v_1) \leq s_i(v_2)$. Lemma 3.9 in Section 3.2 extends this observation to mixed strategies.

The goal of any individual agent i is to bid so as to maximize her (linear) utility $u_i(v_i, \mathbf{b}) = v_i x_i(\mathbf{b}) - p_i(\mathbf{b})$. When the strategies of agents other than i are fixed and clear from context, we use $u_i(v_i, b)$, $x_i(b)$, and $p_i(b)$ to denote the agent's expected utility, allocation, and payment respectively when others' bids are drawn according their respective strategies, e.g., $u_i(v_i, b) = \mathbf{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}}[u_i(v_i, b, \mathbf{s}_{-i}(\mathbf{v}_{-i}))]$. When the value or the bid of the agent are clear from the context, we drop the respective argument to the functions.

A profile of strategies $\mathbf{s} = (s_1, \dots, s_n)$ forms a **Bayes-Nash equilibrium** (BNE) if for all i , v_i , and b' ,

$$u_i(v_i, s_i(v_i)) \geq u_i(v_i, b').$$

Let $u_i(v_i)$ denote the utility that agent i gets at value v_i when playing (implicit) pure strategy s_i . Then we have

$$u_i(v_i) = v_i x_i(s_i(v_i)) - p_i(s_i(v_i)). \quad (1)$$

If s_i is a best response to \mathbf{s}_{-i} , an alternate way of calculating utility that follows from Myerson's [1981] characterization of Bayes-Nash equilibria; here $v_i > v_{\min}$:

$$u_i(v_i) = u_i(v_{\min}) + \int_{z=v_{\min}}^{z=v_i} x_i(s_i(z)) dz. \quad (2)$$

Similar formulas hold for mixed strategies.

2.2. Assumptions

In this paper we study symmetric settings with identical agents and anonymous auctions. Our main result is to rule out the existence of asymmetric equilibria in a large class of auctions. This large class of auctions is defined by the assumption that the allocation and payment of each agent are based on the agent's bid and rank among the other agents, but not the other agents' specific bids.

Definition 2.1. The *rank* of bidder i is given by the number of bidders with strictly higher bid and the number of bidders with equal bid. I.e., for bids \mathbf{b} , the rank of i is $(|\{j : b_j > b_i\}|, |\{j : b_j = b_i\}|)$. Ranks are totally ordered as $(w, t) < (w', t')$ if $w < w'$ or $w = w'$ and $t < t'$.

The following assumptions, which are made throughout the paper, are sufficient for the non-existence of asymmetric equilibria.

Definition 2.2. Our standard assumptions on the auction are

- *continuous bounded-support distribution:* The distributions of bidders' values are independent, identical, continuous (i.e. contain no point masses), and have bounded support. We use F to denote the cumulative distribution function for the value distributions, and let $[v_{\min}, v_{\max}]$ denote its support.
- *rank-and-bid-based allocation rule:* The allocation rule $x_i(b_i, \mathbf{b}_{-i})$ is symmetric across agents, and is given by a non-decreasing function of the agent's bid b_i and the rank of her bid among the other bids.
- *bid-based payment rule:* The payment rule $p_i(b_i, \mathbf{b}_{-i})$ is symmetric across agents, and is a continuous increasing function of the agent's bid b_i and whether or not she wins. We write the payment rule as $p_i(b_i) = p^{\text{part}}(b_i) + x_i(b_i)p^{\text{win}}(b_i)$ where p^{part} and p^{win} are continuous functions representing the agent's payments for participation and winning respectively.

2.3. Two agent reduced form

Our approach will be to reduce the question of existence of asymmetric equilibria in n agent auctions to that in two-agent auctions. For instance, for any symmetric auction, symmetric distribution on values, and fixed strategies of $n - 2$ agents, the induced game on the remaining two players is symmetric. If all possible induced two-agent games only have symmetric equilibria, then the general n -agent game only has symmetric equilibria as the same argument can be applied to all pairs of agents simultaneously. As an example, an n -agent single-item auction has a rank-based allocation rule, i.e., only the top ranked agent is served. The induced two-agent auction for a given distribution of the other $n - 2$ agent bids is both rank- and bid-based: the top ranked agent (of the two) is served with a probability that is increasing in her bid, in particular, the probability that she is also top ranked amongst the other bids.

Definition 2.3. Any n -agent rank-and-bid-based auction, Bayes-Nash equilibrium, and pair of agents i and j induces a *two-agent reduced form*. Its allocation rule is composed of three components specifying the probability that i wins conditioned on her rank relative to j (win, lose, or tie).

$$\begin{aligned} \alpha_W(b_i) &= \mathbf{E}_{\mathbf{v}_{-i}}[x_i(b_i, \mathbf{s}_{-i}(\mathbf{v}_{-i})) \mid b_i > s_j(b_j)], & (\text{win}) \\ \alpha_L(b_i) &= \mathbf{E}_{\mathbf{v}_{-i}}[x_i(b_i, \mathbf{s}_{-i}(\mathbf{v}_{-i})) \mid b_i < s_j(b_j)], \text{ and} & (\text{lose}) \\ \alpha_T(b_i) &= \mathbf{E}_{\mathbf{v}_{-i}}[x_i(b_i, \mathbf{s}_{-i}(\mathbf{v}_{-i})) \mid b_i = s_j(b_j)]. & (\text{tie}) \end{aligned}$$

Importantly, if we consider the induced auction on any subset of agents, then the rank-and-bid-based property of the allocation rule is inherited by the reduced form. This implies that in the above definitions of α_W and α_L the conditional distribution of b_j is irrelevant for the calculation as all bids that satisfy the conditioning render i with the same distribution on rank.

PROPOSITION 2.4. *The two-agent reduced form of an n -agent rank-and-bid-based allocation rule in any Bayes-Nash equilibrium is rank-and-bid based.*

3. NON-EXISTENCE OF ASYMMETRIC EQUILIBRIA

The main technical ingredient of our approach is to rule out the existence of asymmetric equilibria for two-agent settings under our standard assumptions (Definition 2.2). Our main theorem is as follows.

THEOREM 3.1. *In the Bayes-Nash equilibrium of any two-agent symmetric auction with rank-and-bid-based allocation, bid-based payment, and values drawn independently and identically from a continuous bounded-support distribution, the agents play identical strategies everywhere except on a measure zero set of values.*

Proposition 2.4 immediately extends Theorem 3.1 to n -agent auctions with rank-and-bid-based allocations where the following two corollaries forbid Bayes-Nash equilibria wherein symmetric agents follow asymmetric strategies.

COROLLARY 3.2. *In the Bayes-Nash equilibrium of any symmetric n -agent auction with rank-and-bid-based allocation and bid-based payment, any subset of agents with values drawn independent and identically from a continuous bounded-support distribution play identical strategies everywhere except a measure zero set of values.*

COROLLARY 3.3. *In the Bayes-Nash equilibrium of any symmetric n -agent auction with rank-and-bid-based allocation, bid-based payment, and agents with values drawn independently and identically from a continuous bounded-support distribution, all agents play identical strategies except on a measure zero set of values.*

The remainder of this section is devoted to proving Theorem 3.1. We break the proof into two parts: Lemma 3.8 in the following subsection proves the non-existence of pure asymmetric BNE, and Lemma 3.10 in Section 3.2 extends it to mixed BNE.

3.1. Ruling out asymmetric pure equilibria

We begin by considering the possibility of asymmetric *pure* Bayes-Nash equilibria in two-agent auctions. Recall that the rank-and-bid-based allocation rule of a two-agent auction is defined by functions α_W , α_T , and α_L (Definition 2.3). Given strategies s_1 and s_2 for the two agents in BNE, we can write the allocation rule x_1 (and likewise x_2) as follows, where the probabilities are taken over b_2 drawn from the second agent's equilibrium bid distribution.

$$x_1(b) = \alpha_W(b)\Pr[b_2 < b] + \alpha_T(b)\Pr[b_2 = b] + \alpha_L(b)\Pr[b_2 > b] \quad (3)$$

Note that when s_1 and s_2 are not identical, $x_1(b)$ and $x_2(b)$ are not necessarily equal because the probabilities in the above expression may be different for the two agents (cf. Figure 1).

The high level argument proceeds as follows. Suppose there is some interval of values wherein with the same value agent 2 would strictly outbid agent 1. We will be able to identify a (possibly) larger interval where for any value agent 2 weakly outbids agent 1 (and therefore has weakly higher allocation probability) and at the lower end agent 2 has at least the utility of agent 1 and vice versa at the higher end. This *utility crossing* will be observed using equation (1). We will then use utility crossing, the assumption that agent 2 strictly outbids agent 1 on the initial interval (and therefore has strictly higher allocation probability), and equation (2) to obtain a contradiction.

Continuous and Strictly Increasing Strategies. As an illustration we first walk through this approach for the special case where strategies are continuous and strictly increasing. This case will be easier to deal with because ties in bids are a measure zero event. We start with a claim that is important for both parts of the results (see Figure 1).

CLAIM 1. *Let s_1 and s_2 be continuous strictly-increasing strategies in a Bayes-Nash equilibrium. Then for any value v , $s_1(v) < s_2(v)$ implies $x_1(s_1(v)) < x_2(s_2(v))$ and $s_1(v) = s_2(v)$ implies $x_1(s_1(v)) = x_2(s_2(v))$.*




Fig. 1. Strategies as described in Claim 1 with $s_2(v) > s_1(v)$.

PROOF. Let $b_1 = s_1(v)$, $b_2 = s_2(v)$. Since strategies are continuous and strictly increasing they have a well defined inverse. These inverses specify the probability agent 1 beats agent 2 as $F(s_2^{-1}(b_1))$ and the probability that agent 2 beats agent 1 as $F(s_1^{-1}(b_2))$. Assume $b_1 \leq b_2$ which implies that $s_2^{-1}(b_1) \leq v \leq s_1^{-1}(b_2)$ and write $x_1(b_1)$ from (3) as:

$$\begin{aligned} x_1(b_1) &= \alpha_W(b_1)F(s_2^{-1}(b_1)) + \alpha_L(b_1)(1 - F(s_2^{-1}(b_1))) \\ &\leq \alpha_W(b_1)F(s_1^{-1}(b_2)) + \alpha_L(b_1)(1 - F(s_1^{-1}(b_2))) \\ &\leq \alpha_W(b_2)F(s_1^{-1}(b_2)) + \alpha_L(b_2)(1 - F(s_1^{-1}(b_2))) \\ &= x_2(b_2). \end{aligned}$$

Here we used the fact ties are a measure zero event (and therefore the α_T component of (3) is zero) and that α_W and α_L are non-decreasing functions.

In order to prove the strict inequality for $b_1 < b_2$, we note that strictness follows if $\alpha_W(y) > \alpha_L(y)$ at any $y \in [b_1, b_2]$, or if at least one of the functions α_W or α_L takes on a larger value at b_2 than at b_1 . Moreover, the absence of any of those conditions is a violation of Bayes-Nash equilibrium as follows. The allocation probability for agent 2 for bids in $[b_1, b_2]$ is constant; however, the payment at b_2 is strictly larger than the payment at b_1 . This contradicts the assumption that bidding b_2 in this situation is a best response for agent 2. ■

We now formalize the notion of utility crossing and prove that utility crossing on an interval implies equal strategies on that interval. We will state the definition and results in terms of agent 2 potentially outbidding agent 1, though the analogous definitions and results apply with the roles of the agents reversed.

Definition 3.4. Interval $I = (\underline{z}, \bar{z})$ satisfies *utility crossing* if $s_2(v) \geq s_1(v)$ on $v \in I$ and $u_2(\underline{z}) \geq u_1(\underline{z})$ and $u_1(\bar{z}) \geq u_2(\bar{z})$.

LEMMA 3.5. *For a continuous strictly increasing strategy profile \mathbf{s} in Bayes-Nash equilibrium and any interval $I = (\underline{z}, \bar{z})$ that satisfies utility crossing, it must be that $s_2(v) = s_1(v)$ on all $v \in I$ except for a measure zero set of values.*

PROOF. Utility crossing requires that agent 2 weakly outbids agent 1 on the entire interval I , i.e., $s_2(v) > s_1(v)$ for $v \in I$. Therefore, by Claim 1, the allocation probability of agent 2 on these values is at least that of agent 1, i.e. $x_2(v) \geq x_1(v)$. Suppose for a contradiction that agent 2 strictly outbids agent 1, i.e. $s_2(v) > s_1(v)$, on a measurable subset of values in I . Again by Claim 1, on these values the allocation probability of agent 2 strictly exceeds that of agent 1, i.e. $x_2(v) > x_1(v)$. Equation (2) then implies that the difference in utility between \underline{z} and \bar{z} for agent 2 strictly exceeds that for agent 1, i.e. $u_2(\bar{z}) - u_2(\underline{z}) > u_1(\bar{z}) - u_1(\underline{z})$. However, this contradicts utility crossing which requires that agent 1 starts at \underline{z} with at most the utility of agent 2 and ends at \bar{z} with at least the utility of agent 2. ■

We now prove the utility crossing lemma.

LEMMA 3.6 (UTILITY CROSSING). *For a continuous strictly increasing strategy profile \mathbf{s} in Bayes-Nash equilibrium, any interval $I' = (\underline{z}', \bar{z}')$ for which $s_2(v) > s_1(v)$ is contained in an interval $I = (\underline{z}, \bar{z})$ that satisfies utility crossing.*

PROOF. Choose \bar{z} as the smallest value that is at least \bar{z}' where the strategies are identical or v_{\max} if none exists. To prove utility crossing at \bar{z} we must show that $u_1(\bar{z}) \geq u_2(\bar{z})$. If the strategies are equal at \bar{z} , this is implied directly by Claim 1. If the strategies are not equal (when $\bar{z} = v_{\max}$) then we give a deviation of agent 1 which guarantees the same utility as agent 2 (of course the utility of agent 1 for her chosen strategy exceeds that of any deviation). With value v_{\max} the higher-bidding agent 2 (who is outbidding agent 1 at all values) always beats agent 1 and therefore obtains allocation probability $\alpha_W(s_2(v_{\max}))$. If lower-bidding agent 1 with value v_{\max} deviates to mimic higher-bidding agent 2 (bidding $s_2(v_{\max})$) she always beats agent 2 (except for a measure zero probability of tie) and therefore obtains allocation probability $\alpha_W(s_2(v_{\max}))$. The two agents' payments are also the same, so (1) implies that their utilities are the same.

The exact same argument can be applied to \underline{z} defined as the largest value that is at least \underline{z}' where the strategies are identical or v_{\min} if non exists. Notice that in the case $\underline{z} = v_{\min}$ the higher-bidding agent 2 can deviate to mimic the lower-bidding agent 1's strategy which always loses and has allocation probability $\alpha_L(s_2(v_{\min}))$. ■

Lemmas 3.5 and 3.6 combine to show that pure continuous strictly-increasing strategies cannot be in Bayes-Nash equilibrium if they are asymmetric.

LEMMA 3.7. *In any pure continuous strictly-increasing-strategy Bayes-Nash equilibrium of a two-agent symmetric auction with rank-and-bid-based allocation, bid-based payment, and values drawn independently and identically from a continuous bounded-support distribution, the strategies must be identical.*

General Strategies. Strategies in Bayes-Nash equilibrium are always weakly monotone. While the preceding discussion assumed continuous and strictly-increasing strategies, the following discussion allows general strategies which may have jumps and may be constant on some intervals of values. These constant intervals result in point masses in the bid distribution and ties in bid are possible.

We will make the same basic argument but require a more general version of Claim 1 (which we give as Claim 2, below). Importantly, when there are constant bid intervals the potential for ties falsifies the claim that (at the same value) equal bids result in equal

allocation probabilities. This could happen if both agents' strategies are equal at a given value and the constant intervals are of different lengths. This will complicate our search for an interval that satisfies utility crossing as in addition to the one strategy being at least the other on the interval we also need the first allocation probability to be at least the other. Moreover, it also makes it more challenging finding a mimicking deviation to show the utility crossing property.

The following definitions are with respect to fixed and implicit strategies s_1 and s_2 for the agents. For a given bid b and $i \in \{1, 2\}$, let $\underline{v}_i(b)$ and $\bar{v}_i(b)$ denote the end-points of the interval of values over which $s_i(v) = b$; if these are distinct then there is a point-mass in the bid distribution. We generalize Claim 1 as follows (with proof at the end of the section).

CLAIM 2. *Let s_1 and s_2 be strategies in a Bayes-Nash equilibrium. Then for any value v , $s_1(v) < s_2(v)$ implies $x_1(s_1(v)) < x_2(s_2(v))$. Furthermore, when $s_1(v) = s_2(v) = b$, $\underline{v}_2(b) \leq \underline{v}_1(b)$ and $\bar{v}_2(b) \leq \bar{v}_1(b)$ implies $x_1(s_1(v)) \leq x_2(s_2(v))$.*

We now generalize the search for an interval that satisfies utility crossing (Lemma 3.6) and the argument that such an interval contradicts asymmetric strategies (Lemma 3.5) to general strategies.

LEMMA 3.8. *In any pure Bayes-Nash equilibrium of a two-agent symmetric auction with rank-and-bid-based allocation, bid-based payment, and values drawn independently and identically from a continuous bounded-support distribution, the agents play identical strategies everywhere except on a measure zero set of values.*

PROOF. Consider strategies s_1 and s_2 in BNE. Suppose that the utilities of the two agents are equal at all values. Then, equation (2) implies that their allocation functions can differ at only finitely many values, and we are done.

Otherwise, let \underline{z} denote the smallest value at which the two agents receive different utilities. The agents' strategies start to diverge at \underline{z} ; without loss of generality assume that agent 1's bids are lower than agent 2's bids around this value. Let \bar{z} be the first value at which the two strategies cross again after diverging at \underline{z} , that is, s_1 first becomes larger than s_2 . Formally, \bar{z} is defined to be the infimum of all values larger than \underline{z} with $s_1(v) > s_2(v)$ with two exceptions: (1) if the infimum does not exist, we define $\bar{z} = v_{\max}$; (2) if $s_1(\bar{z}) = s_2(\bar{z}) = b$, we redefine $\bar{z} = \underline{v}_1(b)$.

Note that for all values $v \in (\underline{z}, \bar{z})$, we have $s_1(v) \leq s_2(v)$, with strictly inequality over a part of this interval. Furthermore, at all such values v with $s_1(v) = s_2(v) = y$, we must have $\underline{v}_2(y) \leq \underline{v}_1(y)$ and $\bar{v}_2(y) \leq \bar{v}_1(y)$. Therefore, we can use Claim 2 to conclude $x_1(s_1(v)) \leq x_2(s_2(v))$, with strict inequality over a part of this interval. Therefore, applying equation (2) to the interval $[\underline{z}, \bar{z}]$ we get $u_1(\bar{z}) - u_1(\underline{z}) < u_2(\bar{z}) - u_2(\underline{z})$.

We now show that the interval (\underline{z}, \bar{z}) satisfies utility crossing. By the selection of \underline{z} , $u_2(\underline{z}) \geq u_1(\underline{z})$. We now show that $u_1(\bar{z}) \geq u_2(\bar{z})$ by exhibiting a deviation for every $\epsilon > 0$ that that gives agent 1 with value \bar{z} utility at least $u_2(\bar{z}) - \epsilon$. Let $b = s_2(\bar{z})$. Consider a deviation of agent 1 at \bar{z} to a bid b' picked to satisfy the following conditions: (1) $b' > b$, and, (2) each component of the expected payment at b' is no more than $\epsilon/2$ larger than the corresponding payment at b . It is easy to see by the continuity of the payment function that such a bid always exists.

We now have the following series of observations:

- We claim $x_1(b') \geq x_2(b)$.
- By our choice of b' , $p_1(b') \leq p_2(b) + \epsilon$.
- Consequently, agent 1's utility from deviating at \bar{z} is at least $u_2(\bar{z}) - \epsilon$.

It remains to prove that $x_1(b') \geq x_2(b)$. This claim is very similar to Claim 2. However, note that after agent 1 deviates, the two agents are not necessarily playing a best response to each other, so we cannot apply Claim 2. Instead we will give a direct proof.

We claim that $\underline{v}_1(b) \leq \bar{v}_1(b) \leq \underline{v}_2(b') \leq \bar{v}_2(b')$. Note that $b' > b$ implies that $\bar{v}_2(b') \geq \underline{v}_2(b') \geq \bar{v}_2(b) \geq \bar{z}$. Now, if $s_1(\bar{z}) \neq b$, then by the definition of \bar{z} , at all values $v \geq \bar{z}$, $s_1(v) > b$. So, $\bar{v}_1(b) \leq \bar{z}$, and our claim follows. If on the other hand, if $s_1(\bar{z}) = b$, then, because the two strategies cross each other at bid b , we have $\bar{v}_2(b) \geq \bar{v}_1(b)$, and again our claim follows.

Now, we can write $x_1(b')$ as:

$$\begin{aligned} x_1(b') &= \alpha_W(b')F(\underline{v}_2(b')) + \alpha_T(b')(F(\bar{v}_2(b')) - F(\underline{v}_2(b'))) + \alpha_L(b')(1 - F(\bar{v}_2(b'))) \\ &\geq \alpha_W(b')F(\underline{v}_1(b)) + \alpha_T(b')(F(\bar{v}_1(b)) - F(\underline{v}_1(b))) + \alpha_L(b')(1 - F(\bar{v}_1(b))) \\ &\geq \alpha_W(b)F(\underline{v}_1(b)) + \alpha_T(b)(F(\bar{v}_1(b)) - F(\underline{v}_1(b))) + \alpha_L(b)(1 - F(\bar{v}_1(b))) \\ &= x_2(b) \end{aligned}$$

Here the first inequality follows from noting $\alpha_W(b) \geq \alpha_T(b) \geq \alpha_L(b)$ as well as $\underline{v}_1(b) \leq \bar{v}_1(b) \leq \underline{v}_2(b') \leq \bar{v}_2(b')$. The second follows by noting that α_W , α_T , and α_L are non-decreasing functions. This completes the proof. ■

The proof of Claim 2 follows the same basic structure as that of Claim 1; however, ties in bids are properly accounted for.

PROOF OF CLAIM 2. Let $b_1 = s_1(v)$ and $b_2 = s_2(v)$. If $b_1 < b_2$, then we have $\underline{v}_2(b_1) \leq \bar{v}_2(b_1) < v < \underline{v}_1(b_2) \leq \bar{v}_1(b_2)$. Otherwise, by assumption, $\underline{v}_2(b_1) \leq \underline{v}_1(b_2)$ and $\bar{v}_2(b_1) \leq \bar{v}_1(b_2)$. Now we can proceed as in the proof of Lemma 3.8 and write $x_1(b_1)$ as:

$$\begin{aligned} x_1(b_1) &= \alpha_W(b_1)F(\underline{v}_2(b_1)) + \alpha_T(b_1)(F(\bar{v}_2(b_1)) - F(\underline{v}_2(b_1))) + \alpha_L(b_1)(1 - F(\bar{v}_2(b_1))) \\ &\leq \alpha_W(b_1)F(\underline{v}_1(b_2)) + \alpha_T(b_1)(F(\bar{v}_1(b_2)) - F(\underline{v}_1(b_2))) + \alpha_L(b_1)(1 - F(\bar{v}_1(b_2))) \\ &\leq \alpha_W(b_2)F(\underline{v}_1(b_2)) + \alpha_T(b_2)(F(\bar{v}_1(b_2)) - F(\underline{v}_1(b_2))) + \alpha_L(b_2)(1 - F(\bar{v}_1(b_2))) \\ &= x_2(b_2) \end{aligned}$$

Here we used the fact that α_W , α_T , and α_L are non-decreasing functions.

In order to prove the strict inequality for $b_1 < b_2$, we note that strictness follows if $\alpha_W(y) > \alpha_L(y)$ at any $y \in [b_1, b_2]$, or if at least one of the functions α_W , α_T , and α_L takes on a larger value at b_2 than at b_1 . In the absence of any of those conditions, it follows that the allocation probability for the agents is independent of their own bid as well as their opponent's bid as long as their bid is in the range $[b_1, b_2]$. However, the payment at b_2 is strictly larger than the payment at b_1 . This contradicts the fact that agent 2 is playing a best response to agent 1 and bids b_2 . ■

3.2. Ruling out asymmetric mixed equilibria

We now turn to mixed equilibria and show that when the distribution F has no point masses, mixed equilibria are essentially equivalent to pure equilibria. Therefore, Lemma 3.8 extends to mixed equilibria as well. In particular, we show that every mixed Bayes-Nash equilibrium is equivalent to some pure equilibrium in that the allocation rules they induce on the agents' values are identical (almost) everywhere.

A mixed strategy s_i for agent i maps the value of the agent to a distribution over bids. Let $s_i(v_i)$ denote the support of the distribution over bids given by this strategy when the agent's value is v_i . We begin by noting a simple property of (mixed) strategies forming a Bayes-Nash equilibrium, namely that they are non-decreasing functions of value. A strategy s_i is non-decreasing if for all $v_1 < v_2$, $b_1 \in s_i(v_1)$, and $b_2 \in s_i(v_2)$, $b_1 \leq b_2$. The proof of this lemma uses the assumption that the payment function is a strictly increasing function of bids; It is deferred to the appendix.

LEMMA 3.9. *Every (mixed) strategy in a Bayes-Nash equilibrium of an auction with bid-based payments is non-decreasing.*

Armed with this lemma, we note that for every agent i we can “round” the distribution of bids at every value v to any one of the bids in the support $s_i(v)$. Because the value distribution has no point masses, this transformation preserves the distribution of bids almost everywhere, and therefore is a pure equilibrium. We now present this argument formally.

LEMMA 3.10. *Let \mathbf{s} be a mixed Bayes-Nash equilibrium for an auction with bid-based payments and values drawn independently from continuous bounded-support distributions. Then there exists a pure Bayes-Nash equilibrium \mathbf{s}' such that $x_i(s_i(v_i))$ is equal to $x_i(s'_i(v_i))$ at all values except for a set of measure zero.*

PROOF. We define \mathbf{s}' as follows. For every value v_i , set $s'_i(v_i) = \min\{b \in s_i(v_i)\}$, the minimum bid in the support of $s_i(v_i)$. We claim that \mathbf{s}' is a BNE. Note first that s'_i is a best response to \mathbf{s}_{-i} . Next we claim that for every i , the cumulative distribution of \mathbf{b}_{-i} drawn from \mathbf{s}_{-i} is identical to the cumulative distribution of \mathbf{b}'_{-i} drawn from \mathbf{s}'_{-i} . This implies that s'_i is a best response to \mathbf{s}'_{-i} as well.

So consider some agent i and a bid b . We will prove that $\Pr_{\mathbf{s}_i}[b_i \leq b] = \Pr_{\mathbf{s}'_i}[b_i \leq b]$. Let $v = \max\{z : \exists b' \in s_i(z) \text{ with } b' \leq b\}$. Note that by Lemma 3.9, for all $v' < v$ and $b' \in s(v')$, $b' \leq b$. So, $\Pr_{\mathbf{s}_i}[b_i \leq b]$ is at least $\Pr_{v' \sim F}[v' < v]$ and at most $\Pr_{v' \sim F}[v' \leq v]$. Because there is no point mass at v , the two probabilities are equal to $F(v)$. On the other hand, for all $v' > v$ and $b' \in s(v')$, $b' > b$. Therefore, $\Pr_{\mathbf{s}'_i}[b_i \leq b] = \Pr_{v' \sim F}[v' \leq v] = F(v)$.

To complete the proof, we note that Lemma 3.9 implies that s_i and s'_i differ only at a countable number of values v_i , namely those where the functions “jump”. Therefore, the agent’s allocation stays the same at all but a countable number of values. ■

Theorem 3.1 follows directly from Lemmas 3.8 and 3.10.

4. NON-MULTIPLICITY OF SYMMETRIC EQUILIBRIA IN RANK-BASED AUCTIONS

To rule out multiple symmetric equilibria, we need to make additional assumptions on the allocation rule.

Definition 4.1. An auction has a *win-vs-tie-strict* rank-and-bid-based allocation rule if for every bid b an agent has strictly higher probability of allocation when she is ranked first versus when she is tied with everyone (unless both are zero).

Notice that any position auction in which the first position has strictly higher weight than the n th position is win-vs-tie-strict. Tying gives an allocation probability equal to the average of the n position weights and winning gives allocation probability equal to the first position weight. The latter strictly exceeds the former.

The restriction to win-vs-tie-strict rank-and-bid-based allocation rules implies in the two-agent reduced form of the auction that agents benefit from breaking ties with other agents; In other words, α_W is strictly larger than α_T . A simple example illustrates why this strictness property is necessary for the uniqueness of symmetric BNE. Consider a first-price or all-pay auction over two agents that allocates with probability 1 to the higher bidder, with probability 1 to both agents if they have equal bids, and with zero probability otherwise. This setting satisfies all of our standard assumptions. It admits a canonical symmetric BNE, \mathbf{s} , in which agents’ strategies are strictly increasing and the higher value agent always wins. However, there is another BNE in which both agents always bid 0 — then both agents always get allocated with probability 1 and pay nothing, and therefore, neither has an incentive to deviate. In fact, for every $b \in [0, s(v_{\max})]$, there exists a symmetric BNE in which each agent follows the canonical strategy s up to some value and then switches to bidding b .

From the win-vs-tie-strictness property we can deduce that the bid functions are continuous and therefore agents' ranking by bids is identical to their ranking by value in any symmetric equilibrium. The payment identity then implies that the equilibrium is unique. This standard approach is formalized in the lemmas below.

LEMMA 4.2. *In any symmetric Bayes-Nash equilibrium of an n -agent auction with win-vs-tie-strict rank-and-bid-based allocation rule, bid-based payment rule, and i.i.d. distribution on values, there are no point masses in the bid distribution, except at bids with zero probability of allocation.*

PROOF. As the agents' value distributions as well as strategies are identical, each agent bids according to the same bid distribution. Suppose that there is a point mass in the bid distribution at some bid b . Consider an agent i with value v_i for which $b \in s_i(v_i)$. Then we claim that for any $\epsilon > 0$, the agent's allocation probability from bidding $b + \epsilon$ is strictly larger than from bidding b , say by an amount δ . This is because on the one hand, for every setting of other agents' bids \mathbf{b}_{-i} , agent i 's allocation at $b + \epsilon$ is at least as large as that at b , and on the other hand, when the other agents all bid b (an event that happens with non-zero probability), agent i gets a strictly larger allocation by bidding $b + \epsilon$ and winning as opposed to tying. Given that payments are continuous functions of bids, we can now pick ϵ small enough so that deviating from b to $b + \epsilon$ increases agent i 's payment by less than δ . Then, the deviation improves agent i 's utility at value v_i contradicting the fact that the agent is playing a best response. ■

We can conclude that in auctions where the allocation is a function of rank alone and not bids, each agent's allocation rule is uniquely pinned down by the fact that in equilibrium agent bids are ordered strictly as values are ordered.

Definition 4.3. An auction is *rank-based* if the allocation rule depends on the agent's rank alone. An auction is *rank-based with a reserve price* if the allocation rule does not depend on the bids except for a reserve price below which lower bids are rejected.

LEMMA 4.4. *In any symmetric Bayes-Nash Equilibrium of an n -agent auction with win-vs-tie-strict that is rank-based with a reserve price, has a bid-based payment rule, and i.i.d. continuous distribution on values, the equilibrium orders those bidders that bid above the reserve price strictly by value.*

PROOF. By Lemma 4.2 there are no point masses in the bid distribution except at bids below the reserve price, so bid functions must be strictly increasing above a certain value. Symmetry then implies that higher valued agents are ranked strictly above lower valued agents. ■

THEOREM 4.5. *There is only one symmetric Bayes-Nash Equilibrium of an n -agent auction with win-vs-tie-strict rank-based allocation rule, bid-based payment rule, and i.i.d. continuous distribution on values.*

PROOF. For rank-based allocation rules, each agent allocation rule is determined by the order of bids which is must be the same in all symmetric BNE. The payment identity (that follows from equation (2)) then implies that the payment rule as a function of value is fixed. By continuity of the payment rule there is a unique inverse which gives the unique symmetric equilibrium strategies. ■

We now extend this theorem to rank-based auctions with a reserve price. We show, in particular, that in any two symmetric BNE, agents switch from bidding 0 to bidding above

the reserve at the same value. Then, we can again argue that all symmetric equilibria induce the same payment rule.

THEOREM 4.6. *There is only one symmetric Bayes-Nash Equilibrium of an n -agent auction with win-vs-tie-strict rank-based allocation rule with a reserve price, bid-based payment rule, and i.i.d. continuous distribution on values.*

PROOF. Let r be the reserve price of the auction. We can assume without loss of generality that agents' strategies map values to bids in the range $\{0\} \cup [r, \infty)$. For a symmetric BNE \mathbf{s} , let v^* denote the smallest value v at which $s(v) \geq r$. We will argue that v^* is the same for all BNE \mathbf{s} . Then the result follows by noting as in Lemma 4.4 that the agent's allocation at values above v^* is uniquely defined at all BNE.

We first observe that $s(v^*) = r$: if instead $s(v^*) > r$, an agent can reduce his bid at v^* , obtaining the same probability of allocation and a smaller payment, thereby improving his utility. Second, the agent is indifferent between bidding 0 and bidding r at v^* . This is because if the agent receives strictly positive utility from bidding r at v^* , then she would receive strictly positive utility from bidding r at a value smaller than but close enough to v^* . This contradicts the fact that the agent is playing a best response and bidding 0 at values below v^* .

Finally, the utility of an agent at v^* from bidding r is given by $u(v^*) = x(v^*)(v^* - p^{\text{win}}(r)) - p^{\text{part}}(r)$, where $x(z)$ is defined to be the expected allocation of the agent when all other agents with values above z bid more than r and all other agents with values below z bid 0. Note that x , and therefore also u , is a strictly increasing function of v^* . So, there is a unique value v^* at which the utility becomes zero, proving our claim. ■

5. RISK AVERSE AGENTS

So far we have focused on agents with quasilinear preferences. We now consider settings with risk averse agents. Following a standard approach (see, e.g., [Maskin and Riley 1984; Matthews 1983; Fu et al. 2013]), we model risk aversion via a concave utility function. Formally, we define the *wealth* of an agent as the total value he gets from the mechanism's outcome minus the payment that he makes. The agent's utility is a concave function U of his wealth (for deterministic two-part bid-based payments):

$$u_i(v_i, \mathbf{b}) = x_i(\mathbf{b})U(v_i - p^{\text{win}}(\mathbf{b}) - p^{\text{part}}(\mathbf{b})) + (1 - x_i(\mathbf{b}))U(-p^{\text{part}}(\mathbf{b})) \quad (4)$$

We will now argue that our uniqueness of equilibria result extends to settings with risk aversion. In keeping with our symmetry assumption we assume that all agents have identical risk profiles, that is, their utility functions are identical. We use U to denote this common function.

There are two ways in which our argument for the non-existence of asymmetric equilibria in the two player risk neutral setting depended on the agents' utility functions. First, we considered local deviations for agents that improved their utilities; We used equation (1) to relate the utility to the agent's allocation probability and payment. The deviating agent either copied exactly the opponent's strategy, obtaining the same allocation and payment as the opponent, or employed a bid slightly larger than that of her opponent. In the latter case, we argued that the deviating agent receives a strictly larger probability of allocation than her opponent, and in choosing the deviation carefully we ensured that the increase in payment is small enough to produce an overall increase in the agent's utility. Second, we used Myerson's revenue equivalence theorem (2) to relate the utility that an agent accrues over an interval of values to her allocation probability over that interval. Using this observation we argued that if one agent has consistently higher allocation probability over a value interval than the other, then she accrues a larger utility.

We now describe how to modify these arguments to the risk averse setting. In the first case we can replace equation (1) by equation (4). It follows that whenever two agents receive the same allocation and payment, they also obtain the same utility. Further we observe that a discontinuous increase in an agent’s allocation probability leads to a discontinuous increase in her utility as well, and therefore, we can once again pick a deviation carefully so as to obtain a net increase in utility.

The second case requires a (weak) utility equivalence in risk averse settings similar to Myerson’s theorem for risk neutral agents. We say that a utility function satisfies *weak utility equivalence* if for any mechanism with a bid-based payment rule, over any interval $[v_1, v_2]$, for an equilibrium strategy s , we can write the difference of utilities of the agent at v_2 and v_1 as a function of the agent’s allocation probability over that interval:

$$u(v_1) - u(v_2) = \int_{z=v_1}^{z=v_2} Q(x(s(z)), z) dz.$$

Here Q is strictly increasing in its first argument.

Matthews [1983] shows that utility functions exhibiting *constant absolute risk aversion* (CARA) satisfy weak utility equivalence. Fu et al. [2013] show that this condition also holds in first-price auctions for capacitated utility functions, namely those that are linear up to a certain wealth and then constant thereafter.

We therefore obtain the following theorem:

THEOREM 5.1. *In the Bayes-Nash equilibrium of any symmetric n -agent auction with rank-and-bid-based allocation, bid-based payment, values drawn from i.i.d. from a continuous bounded-support distribution, and identical utility-equivalent risk-averse preferences, all agents play identical strategies except on a measure zero set of values.*

Similar extensions follow for Corollary 3.2 and Theorem 4.6.

REFERENCES

- AMANN, E. AND LEININGER, W. 1996. Asymmetric all-pay auctions with incomplete information: The two-player case. *Games and Economic Behavior* 14, 1, 1 – 18.
- BAJARI, P. 2001. Comparing competition and collusion: a numerical approach. *Economic Theory* 18, 187–205.
- BHATTACHARYA, S., GOEL, G., GOLLAPUDI, S., AND MUNAGALA, K. 2010. Budget constrained auctions with heterogeneous items. In *Proceedings of the 42nd ACM symposium on Theory of computing*. ACM, New York, NY, USA, 379–388.
- CARAGIANNIS, I., KAKLAMANIS, C., KANELLOPOULOS, P., KYROPOULOU, M., LUCIER, B., LEME, R. P., AND É.TARDOS. 2012. On the efficiency of equilibria in generalized second price auctions. <http://arxiv.org/abs/1201.6429>.
- CHAWLA, S., HARTLINE, J., AND SIVAN, B. 2012. Optimal crowdsourcing contests. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 856–868.
- DIPALANTINO, D. AND VOJNOVIC, M. 2009. Crowdsourcing and all-pay auctions. In *Proceedings of the 10th ACM conference on Electronic commerce*. ACM, New York, NY, USA, 119–128.
- EDELMAN, B., OSTROVSKY, M., AND SCHWARZ, M. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review* 97, 1, 242–259.
- FU, H., HARTLINE, J., AND HOY, D. 2013. Prior-independent auctions for risk-averse agents. In *Proceedings of the 13th ACM conference on Electronic commerce*. ACM.

- GOMES, R. AND SWEENEY, K. 2009. Bayes-nash equilibria of the generalized second price auction. In *Proceedings of the 10th ACM conference on Electronic commerce*. ACM, 107–108.
- HA, B. AND HARTLINE, J. 2012. Mechanism design via consensus estimates, cross checking, and profit extraction. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 887–895.
- HARTLINE, J. AND YAN, Q. 2011. Envy, truth, and profit. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, 243–252.
- JACKSON, M. 2001. A crash course in implementation theory. *Social Choice and Welfare* 18, 655–708.
- JANSEN, B. AND MULLEN, T. 2008. Sponsored search: an overview of the concept, history, and technology. *International Journal of Electronic Business* 6, 114–131.
- LEBRUN, B. 2006. Uniqueness of the equilibrium in first-price auctions. *Games and Economic Behavior* 55, 1, 131 – 151.
- LUCIER, B. AND PAES LEME, R. 2011. Gsp auctions with correlated types. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, New York, NY, USA, 71–80.
- MASKIN, E. AND RILEY, J. 1984. Optimal auctions with risk averse buyers. *Econometrica* 52, 1473–1518.
- MASKIN, E. AND RILEY, J. 2003. Uniqueness of equilibrium in sealed high-bid auctions. *Games and Economic Behavior* 45, 2, 395 – 409.
- MATTHEWS, S. A. 1983. Selling to risk averse buyers with unobservable tastes. *Journal of Economic Theory* 30, 2, 370 – 400.
- MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 58–73.
- PAES LEME, R. AND TARDOS, E. 2010. Pure and bayes-nash price of anarchy for generalized second price auction. In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, 735–744.
- SYRGKANIS, V. AND TARDOS, E. 2013. Composable and efficient mechanisms. In *Symposium on Theory of Computation*. ACM.
- VARIAN, H. 2007. Position auctions. *International Journal of Industrial Organization* 25, 6, 1163–1178.
- VARIAN, H. R. 2009. Online ad auctions. *The American Economic Review* 99, 2, 430–434.

A. DEFERRED PROOFS

PROOF OF LEMMA 3.9. Consider an agent i , and fix the strategies of all other agents. (We drop the subscript i in the rest of this proof to improve readability.) Suppose there are values $v_1 < v_2$ and bids $b_1 \in s(v_1)$ and $b_2 \in s(v_2)$ with $b_1 > b_2$. We can relate the utilities of the agent at v_1 and v_2 as follows:

$$u(v_2, b_1) = u(v_1, b_1) + (v_2 - v_1)x(b_1)$$

$$u(v_2, b_2) = u(v_1, b_2) + (v_2 - v_1)x(b_2)$$

Now, $b_2 \in s(v_2)$ implies $u(v_2, b_1) \leq u(v_2, b_2)$. Likewise, $b_1 \in s(v_1)$ implies $u(v_1, b_1) \geq u(v_1, b_2)$. Finally, $b_1 > b_2$ implies $x(b_1) \geq x(b_2)$. Therefore, it must be the case that $u(v_2, b_1) = u(v_2, b_2)$, $u(v_1, b_1) = u(v_1, b_2)$, and $x(b_1) = x(b_2)$. However, these in turn imply that the payments that the agent makes at b_1 and b_2 are equal, which contradicts the fact that the payment function is strictly increasing. ■