Expressiveness and Robustness of First-Price Position Auctions

PAUL DÜTTING, Stanford University
FELIX FISCHER, University of Cambridge
DAVID C. PARKES, Harvard University

It is desirable for an economic mechanism that its properties hold in a robust way across multiple equilibria and under varying assumptions regarding the information available to the participants. In this paper we focus on the design of position auctions and seek mechanisms that guarantee high revenue in every efficient equilibrium under both complete and incomplete information. Our main result identifies a generalized first-price auction with multi-dimensional bids as the only standard design capable of achieving this goal, even though valuations are one-dimensional. The fact that expressiveness beyond the valuation space is necessary for robustness provides an interesting counterpoint to previous work, which has highlighted the benefits of simple bid spaces. From a technical perspective, our results are interesting because they establish equilibrium existence for a multi-dimensional bid space, where standard techniques for establishing equilibrium existence break down.

Categories and Subject Descriptors: F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Simplicity-Expressiveness Tradeoffs, Generalized First-Price Auction, Target-Profit Strategies

1. INTRODUCTION

We consider a standard position auction setting in which k positions are to be assigned to n agents in a one-to-one fashion and agents agree on the relative values of the positions. Such a setting can be described by two vectors $v = (v_1, \ldots, v_n)$ and $\beta = (\beta_1, \ldots, \beta_k)$, where value v_i is private to agent i and β is publicly known. The valuation of agent i for position j is then given by $\beta_j \cdot v_i$, and we will assume for convenience that $\beta_1 \geq \beta_2 \geq \cdots \geq \beta_k$. This is a one-dimensional setting, as the private information of each agent consists of a single number. A prime example of position auctions can be found in the context of sponsored search, where agents correspond to advertisers and positions correspond to slots in which advertisements can be displayed. Each slot comes with a click-through rate and each advertiser with a value per click.

In position auction settings, the same auction format is typically applied across very different problem instances. Search engines for example use the same format to auction off both valuable, high frequency keywords and a vast number of keywords corresponding to less frequent user queries. For high volume keywords it may be reasonable

to consider a complete information model, where agents know one another's value per click. For less frequent keywords, on the other hand, an incomplete information model, where an agent only has probabilistic beliefs about the values of others, seems more appropriate. Naturally, an auction used for both high frequency and low frequency keywords should possess appropriate strategic equilibria under both complete and incomplete information.

In addition to equilibrium existence, auction designs are subject to a tradeoff among various performance criteria. A goal that is shared by both the agents and the auctioneer is the maximization of social welfare, i.e., of the sum of the agents' valuations for the positions they are assigned. For the agents this goal corresponds to an efficient use of the available resources, for the auctioneer to the provision of a satisfactory service to its customers. At the same time, the auctioneer also seeks to maximize its revenue, i.e., the sum of the payments it receives from the agents.

We thus arrive at the following question:

> Does there exist a position auction that possesses an efficient equilibrium and achieves good revenue in every efficient equilibrium, under both complete and incomplete information?

As a benchmark in both cases we adopt the revenue achieved in the truthful equilibrium of the Vickrey-Clarke-Groves auction, henceforth termed the truthful VCG revenue. This benchmark has been used previously in settings with complete information [e.g., Lucier et al. 2012]. An alternative benchmark for settings with incomplete information sacrifices efficiency [Myerson 1981], which seems inappropriate given our focus on efficient equilibria. The truthful VCG revenue, on the other hand, is the maximum revenue achievable by any efficient mechanism [Myerson 1981].

Before we proceed any further, it is worth noting that revenue equivalence, which states that equilibria resulting in the same allocation must also yield the same revenue, is not enough to resolve the above question for either complete or incomplete information. Under complete information, revenue equivalence does not normally hold. Under incomplete information, where it does apply, it does not guarantee existence of an efficient equilibrium.

1.1. Candidate Auctions

We address the question by considering variants of the three designs commonly used for position auctions: the Vickrey-Clarke-Groves (VCG) auction, the generalized second-price (GSP) auction, and the generalized first-price (GFP) auction.[1] For each of these designs we distinguish between a simplified variant with one-dimensional bids and an expressive variant with multi-dimensional bids:

— In the simplified variant each agent i specifies a single bid b_i, and this bid is multiplied by $\beta_1 \geq \beta_2 \geq \cdots \geq \beta_k$ to obtain bids for the different positions.
— In the expressive variant each agent i submits a separate bid $b_{i,j}$ for each position j, where we require bids to be non-increasing from position 1 to position k.

We further distinguish between two allocation rules: the first assigns positions to agents so as to maximize social welfare with regard to the reported valuations; the second considers each position in turn, from first to last, and greedily assigns it to an agent with maximum bid for the position among those not previously assigned a position. For simplified bids the greedy allocation rule also maximizes reported social welfare, and hence the two allocation rules are equivalent. This equivalence does not

---

[1]Google and Microsoft use the GSP auction, Facebook the VCG auction. The GFP auction was used by Overture, the first company to provide a successful sponsored search service.

Table I. Negative results for standard auction designs and focus of this work

bid space	allocation rule	VCG-based payments	GSP-based payments	GFP-based payments
simplified	efficient / greedy	efficient low revenue Nash equilibrium [Dütting et al. 2011a]	efficient low revenue Nash equilibrium [Dütting et al. 2011a], no Bayes-Nash equilibrium [Gomes and Sweeney 2009]	no Nash equilibrium [Edelman et al. 2007]
expressive	efficient	efficient low revenue Nash equilibrium [Milgrom 2010]	efficient low revenue Nash equilibrium [Milgrom 2010]	
	greedy	efficient low revenue Nash equilibrium [Milgrom 2010]	efficient low revenue Nash equilibrium [Milgrom 2010]	this work

generally hold in the expressive case. Our main positive result is enabled by the greedy allocation rule, while all negative results hold for both allocation rules.

Finally, for a given allocation rule, payments are defined as follows: the VCG auction charges each agent the difference in social welfare of the other agents when the agent is absent and when it is present, both with regard to reported valuations; the GSP auction charges each agent the next-highest bid on the position it is assigned coming from an agent that is not assigned a higher position; the GFP auction charges each agent its bid on the position it is assigned.

1.2. Results

It turns out that most candidate auctions are disqualified by prior work, see Table I for an overview. All variants of the VCG and GSP auctions—with simplified or expressive bids, efficient or greedy allocation rule—have an efficient complete information equilibrium with arbitrarily small revenue compared to the truthful VCG revenue [Milgrom 2010; Dütting et al. 2011a]. This result is rather robust, and continues to hold for example when multipliers $\alpha \neq \beta$ are used [Dütting et al. 2011a]. An additional disadvantage of the simplified GSP auction is that it may not have any (Bayes-Nash) equilibrium when information is incomplete [Gomes and Sweeney 2009]. The simplified GFP auction, on the other hand, has a unique (Bayes-Nash) equilibrium under incomplete information [Chawla and Hartline 2013], but may not have an equilibrium under complete information [Edelman et al. 2007].

The only remaining candidate is the expressive GFP auction, and we show that the variant with greedy allocation rule indeed possesses the desired robustness property:

— Under complete information it has at least one equilibrium, and all of its equilibria are efficient and yield at least the truthful VCG revenue.
— Under incomplete information it has an efficient equilibrium, and every efficient equilibrium yields the truthful VCG revenue.

A few comments are in order regarding our focus on efficient equilibria. Since coordination among agents is relatively straightforward under complete information, the existence of efficient complete information equilibria with low revenue, which see agents coordinate to extract a larger share of the surplus, provides a rather strong argument against the use of a particular design. In this sense, a restriction to efficient equilibria only strengthens the negative results for the VCG and GSP auctions. At the same time, our results for the expressive GFP auction rule out the possibility of low revenue in any complete information equilibrium. Under incomplete information, coordination among agents seems more difficult and hence less likely, especially if there exists a

natural efficient equilibrium. The incomplete information equilibrium we identify for the expressive GFP auction can be reached via so-called target-profit strategies [e.g., Bernheim and Whinston 1986]. It has been argued that their conceptual simplicity makes these strategies a very natural candidate for adoption in the absence of dominant strategies [Milgrom 2004].

In summary, our results identify an expressive auction with first-price payments as the only standard design capable of providing the desired robustness property across complete and incomplete information environments. That expressiveness is in this sense necessary for robustness provides an interesting counterpoint to previous work on position auctions, which has highlighted the benefits of simplicity [Milgrom 2010; Dütting et al. 2011a].

1.3. Techniques

Our analysis of the complete information case is technically similar to the classic analysis of Bernheim and Whinston [1986] that links equilibria of first-price auctions to the core, and to more recent approaches that also make this connection [Day and Milgrom 2008; Hoy et al. 2013]. The common feature is the use of target-profit strategies. Specifically, we show that having each agent i bid its value $\beta_j \cdot v_i$ for position j minus its utility u_i in the truthful equilibrium of the VCG auction, or zero if this difference is negative, yields an efficient equilibrium. By construction, revenue in this equilibrium equals the truthful VCG revenue. We then exclude the possibility of inefficient equilibria by showing how inefficiencies in the allocation lead to opportunities for beneficial unilateral deviation. To establish the revenue guarantee for efficient equilibria, we consider certain unilateral deviations from the proposed efficient equilibrium, and use the fact that these deviations cannot be beneficial to derive lower bounds on payments in equilibrium.

In the incomplete information setting, the combination of one-dimensional valuations and multi-dimensional bids poses two fundamental challenges. First, Myerson's [1981] classic equilibrium characterization only provides a necessary, but not a sufficient condition: it tells us that payments in every efficient equilibrium must equal those in the truthful equilibrium of the VCG auction. Since bids are multi-dimensional, there are many different bids that satisfy the condition and thus many candidate equilibrium bids. Second, the standard technique to verify that a particular candidate is an equilibrium involves integrating the derivative of an agent's utility, as a function of both valuation and bid, along a path between two bids. This technique breaks down in our setting, where the bid space has higher dimension than the valuation space and the utility function may not be defined everywhere on the path.

In the candidate equilibrium we consider, agent i's bid on position j equals its expected truthful VCG payment conditioned on being allocated position j. Like their counterparts under complete information, the bids in the incomplete information setting thus have an interpretation as target-profit strategies. To prove that they form an equilibrium we carry out an induction from the last position to the first, and show that the conjectured equilibrium bid on position j is optimal for agent i, given that the other agents bid according to the conjectured equilibrium and agent i bids according to the conjectured equilibrium on positions $j + 1$ to k. We believe that similar techniques could be used to show equilibrium existence in more general settings, including settings with multi-dimensional valuations.

1.4. Related Work

The design of expressive auctions for specific applications is an important topic of contemporary mechanism design [e.g., Aggarwal et al. 2009; Ghosh and Sayedi 2010; Constantin et al. 2011; Dütting et al. 2011b; Dobzinski et al. 2012; Dütting et al. 2012; Goel

et al. 2012]. The intuition that expressiveness is generally desirable is supported by work of Benisch et al. [2008], who showed that the maximum social welfare a mechanism can achieve strictly increases with a measure of expressiveness based on a concept from computational learning theory.

The classic analysis of position auctions is due to Varian [2007] and Edelman et al. [2007]. Working under the assumption that agents have complete information, these authors showed that the GSP auction—although not truthful—has certain desirable equilibria, in which the revenue is at least as high as in the truthful equilibrium of the VCG auction. For incomplete information environments, Gomes and Sweeney [2009] showed that the GSP auction may not possess an efficient equilibrium. The GFP auction always has a unique, efficient (Bayes-Nash) equilibrium under incomplete information, which yields the truthful VCG revenue [Chawla and Hartline 2013], but may not have a (Nash) equilibrium under complete information [Edelman et al. 2007]. Our analysis differs from these earlier results in its use of an expressive bidding language, which overcomes the negative result for complete information settings. In a series of papers, Paes Leme and Tardos [2010], Lucier and Paes Leme [2011], Caragiannis et al. [2011], and Syrgkanis and Tardos [2013] showed that the GSP auction has a small constant price of anarchy under both complete and incomplete information.[2] Lucier et al. [2012] and Caragiannis et al. [2012] showed that the revenue of the GSP auction can be arbitrarily small compared to the truthful VCG revenue under complete information, but under incomplete information always is within a constant factor of the optimal revenue of Myerson [1981] when reserve prices are chosen appropriately.

The role of expressiveness in position auction environments was highlighted by Milgrom [2010] and Dütting et al. [2011a], who considered VCG and GSP auctions under complete information and argued that a restriction of the bidding space to a subspace of the valuation space can rule out zero revenue equilibria without introducing new and potentially bad ones. Also in a complete information setting and for VCG and GSP auctions, Blumrosen et al. [2008] and Abrams et al. [2009] bounded the reduction in equilibrium quality resulting from a restriction of the bidding space to a subspace of the valuation space.

Independently from our work, Hoy et al. [2013] argued in favor of an expressive design for first-price auctions in complete information environments. The authors used target-profit strategies to establish the existence of an efficient complete information equilibrium that yields the truthful VCG revenue. They also showed that every cooperatively envy-free outcome—a strengthening of the envy-freeness concept from individual bidders to groups of bidders—is efficient and yields the truthful VCG revenue. Like the original envy-freeness concept, cooperative envy-freeness leaves the standard equilibrium framework. Our results regarding efficiency of equilibria and revenue in efficient complete information equilibria strengthen these results for position auctions from cooperatively envy-free equilibria to Nash equilibria. In addition we also consider incomplete information settings, requiring that the same design has robust performance under both complete and incomplete information. In performing our analysis, we successfully apply the notion of target-profit strategies to an incomplete information setting. Lifting this concept from complete to incomplete information had been an open problem since its introduction by Bernheim and Whinston [1986].

To the best of our knowledge, the study of position auctions that admit efficient equilibria and yield high revenue in every efficient equilibrium under both complete

---

[2]The price of anarchy compares the minimum social welfare in any equilibrium to the maximum social welfare of any outcome. That greedy algorithms can achieve a small price of anarchy, potentially smaller than that of an efficiently computable outcome, was previously highlighted by Gairing [2009] in the context of covering games.

and incomplete information, and the use of additional expressiveness to achieve this goal, are also novel to the present work.

2. PRELIMINARIES

We study a setting with a set $\{1, \ldots, k\}$ of *positions* ordered by quality and a set $N = \{1, \ldots, n\}$ of *agents* with *unit demand* and *one-dimensional valuations* for the positions. More formally, write $\mathbb{R}^k_{\geq} = \{x \in \mathbb{R}^k : x_j > 0, x_j \geq x_{j'} \text{ if } j < j'\}$ for the set of k-dimensional vectors whose entries are positive and non-increasing. For $\beta \in \mathbb{R}^k_{\geq}$, let $\mathbb{R}^k_{\beta} = \{x \in \mathbb{R}^k : x = \beta v, v \in \mathbb{R}_+\}$ be the one-dimensional subspace of $\mathbb{R}^k$ spanned by β. Agent i's valuation can then be represented by a vector $\beta v_i \in \mathbb{R}^k_{\beta}$ in this subspace, such that $\beta_j v_i \geq 0$ is the agent's value for position j. We assume that vector β is common knowledge among the agents. Our goal is to assign the positions to agents in order to maximize total value or *social welfare*. An assignment of agents to positions satisfying this property is commonly called *efficient*.

The *VCG auction* solicits a single-dimensional bid $b_i \in \mathbb{R}_+$ from each agent $i \in N$. This bid is then extended to a k-dimensional bid by multiplying it with vector β. Agent i's bid on position j is thus $\beta_j v_i$. The allocation is chosen so as to maximize the reported social welfare, i.e., the sum of the bids. The payment of agent i is the amount by which its presence reduces the social welfare of the other agents with regard to their bids.

The *expressive GFP auction* solicits a vector $b_i \in \mathbb{R}^k_{\geq}$ of bids from each agent $i \in N$, where $b_{i,j}$ is interpreted as agent i's bid on position j. The allocation is computed in a greedy manner by going through positions 1 to k, and assigning the current position to an agent with maximum bid on that position among the agents not yet assigned a higher position. The payment of an agent i assigned position j is its bid $b_{i,j}$ on this position.

We assume *quasi-linear utilities*, such that the utility $u_i(b, v_i)$ of agent i with value v_i, in a given auction and for a given bid profile $b = (b_1, \ldots, b_n)$, is equal to its valuation for the position it is assigned minus its payment for that position. To be able to reason about strategic behavior we further need to specify what agents know about one another's valuations.

In the *complete information* setting, the values v_i are common knowledge among the agents. A bid profile $(b_1, \ldots, b_n)$ is a *Nash equilibrium* of an auction if no agent has an incentive to change its bid assuming that the other agents don't change their bids, i.e., if for every $i \in N$ and every b'_i,

$$u_i((b_1, \ldots, b_i, \ldots, b_n), v_i) \geq u_i((b_1, \ldots, b'_i, \ldots, b_n), v_i).$$

In the *incomplete information* setting, values v_i are drawn independently from a distribution F with density f and support $[0, \bar{v}]$ for some finite $\bar{v} \in \mathbb{R}_+$. Distribution F is assumed to be common knowledge among the agents. A profile $(b_1, \ldots, b_n)$ of *bidding functions* then is a *Bayes-Nash equilibrium* if no agent has an incentive to change its bidding function assuming that the values of the other agents are drawn from F, i.e., if for every $i \in N$, every $v_i \in [0, \bar{v}]$, and every bidding function b'_i,

$$\mathbb{E}_{v_j \sim F, j \neq i} \Big[u_i \big((b_1(v_1), \ldots, b_{i-1}(v_{i-1}), b_i(v_i), b_{i+1}(v_{i+1}), \ldots, b_n(v_n)), v_i \big) \Big] \geq$$

$$\mathbb{E}_{v_j \sim F, j \neq i} \Big[u_i \big((b_1(v_1), \ldots, b_{i-1}(v_{i-1}), b'_i(v_i), b_{i+1}(v_{i+1}), \ldots, b_n(v_n)), v_i \big) \Big].$$

Because our environment is one-dimensional we can appeal to Myerson's characterization of the expected payments in a Bayes-Nash equilibrium.

THEOREM 2.1 (MYERSON [1981]). *Consider a position auction, and assume that agents use bidding functions such that agent i with valuation v_i is assigned position*

s with probability $P_s^i(v_i)$. *Then the bidding functions are a Bayes-Nash equilibrium of the auction only if, for every agent i,*

(1) *the expected allocation $\sum_{s=1}^{k} P_s^i(v_i)\beta_s$ is non-decreasing in v_i and*
(2) *the expected payment is*

$$p_i(v_i) = \sum_{s=1}^{k} P_s^i(v_i)\beta_s v_i - \int_{z=0}^{v_i} \sum_{s=1}^{k} P_s^i(z)\beta_s \, dz,$$

where $p_i(0) = 0$.

Since an efficient allocation satisfies Condition 1, monotonicity, we have the following corollary.

COROLLARY 2.2. *In an efficient Bayes-Nash equilibrium of any position auction, the expected payment of every agent i is equal, for every value v_i, to the expected payment of the agent in the truthful equilibrium of the VCG auction.*

3. COMPLETE INFORMATION

We begin our analysis of the expressive GFP auction for settings with complete information. We show that it always has a Nash equilibrium, that all its Nash equilibria are efficient, and that payments in every Nash equilibrium are at least the truthful VCG payments.

The first result is proved by showing that the bid profile where the bid of agent i on position j equals its value $\beta_j \cdot v_i$ minus its truthful VCG utility u_i, or zero if this difference is negative, constitutes a Nash equilibrium. To establish the second result we argue that any inefficiency in the allocation leads to an opportunity for unilateral deviation. For the third result we use the fact that no beneficial unilateral deviations exist in a Nash equilibrium to derive lower bounds on payments.[3]

THEOREM 3.1. *Assume that agent valuations in a position auction on k positions are in subspace $\mathbb{R}_\beta^k$ of $\mathbb{R}^k$, for $\beta \in \mathbb{R}_\geq^k$. Then,*

(1) *the expressive GFP auction has an efficient Nash equilibrium with payments equal to the truthful VCG payments,*
(2) *every Nash equilibrium of the expressive GFP auction is efficient, and*
(3) *the payments in every Nash equilibrium of the expressive GFP auction are at least the truthful VCG payments.*

PROOF. We prove the claims one by one.

For the first claim, consider without loss of generality the case where agents are ordered by decreasing value, such that $v_1 \geq v_2 \geq \cdots \geq v_n$, and the efficient allocation where agent i is assigned position i, for $1 \leq i \leq k$. Denote by u_i the truthful VCG utility of agent i, for $1 \leq i \leq n$, and by p_i the truthful VCG payment for position i, for $1 \leq i \leq k$. Then $u_i = \beta_i v_i - p_i$ for $1 \leq i \leq k$ and $u_i = 0$ for $i > k$. We further claim that the bid profile $b \in (\mathbb{R}_\geq^k)^n$ with

$$b_{i,j} = \max(\beta_j v_i - u_i, 0)$$

for $i = 1, \ldots, n$ and $j = 1, \ldots, k$ is an equilibrium of the GFP auction that is efficient and yields the truthful VCG payments.

---

[3] An alternative proof could use the connections between the VCG outcome and Walrasian equilibria [Gul and Stacchetti 1999] and between Walrasian equilibria and equilibria of expressive first-price auctions [Hassidim et al. 2011]. Our proof from first principles has the advantage that it makes the role of target-profit strategies explicit.

Under this bid profile, an efficient allocation assigns position i to agent i at price p_i. With the greedy allocation rule, this outcome can be obtained by letting agent i point to position i and breaking ties in favor of the agent that points to a given position. To see that b is indeed an equilibrium first observe that agent i cannot decrease its bid for position i without being assigned a position other than i. Now assume for contradiction that agent i has a beneficial deviation to a position $j \neq i$, such that

$$\beta_i v_i - p_i < \beta_j v_i - p_j - \epsilon,$$

for every $\epsilon > 0$. Here we use that agent i can bid $p_j + \epsilon$ on positions j and above to win one of these positions, and that it values each of them at least as highly as position j. The left-hand side of this inequality equals the utility of agent i in the truthful equilibrium of the VCG auction, whereas the right-hand side equals the utility agent i would get if it was instead assigned position j at price $p_j + \epsilon$. The inequality contradicts the fact that the truthful VCG equilibrium is envy-free.

For the second claim, consider a Nash equilibrium $b = (b_1, \ldots, b_n)$ and assume for contradiction that it leads to an inefficient assignment. Then there exist agents i, j with $v_i > v_j$ such that agent i is assigned position s and agent j is assigned position $t < s$.

First assume that agent i bids $b_{j,t} + \epsilon$ on positions t and above, which means that it wins one of these positions. Since b is an equilibrium this deviation is not beneficial, i.e., for every $\epsilon > 0$,

$$\beta_s v_i - b_{i,s} \geq \beta_t v_i - b_{j,t} - \epsilon. \tag{1}$$

Now consider the situation where agent j bids according to bid vector b'_j with

$$b'_{j,\ell} = \begin{cases} b_{i,s} + \epsilon & \text{if } 1 \leq \ell \leq s \\ 0 & \text{otherwise} \end{cases}$$

for some $\epsilon > 0$. We claim that with these bids agent j will either be assigned a position above s, or will compete for position s with bids that are at most $b_{i,s}$ and will therefore be assigned position s. For the latter observe that agents other than j who are assigned a position above s when agent j bids according to b_j can only be assigned a higher position when agent j bids according to b'_j. This suffices because agents other than j who were assigned position s or below bid at most $b_{i,s}$ on position s.

Since b is an equilibrium, agent j does not benefit from bidding according to b'_j, and thus for every $\epsilon > 0$,

$$\beta_t v_j - b_{j,t} \geq \beta_s v_j - b_{i,s} - \epsilon. \tag{2}$$

By adding (1) and (2) and rearranging,

$$\beta_s v_i + \beta_t v_j \geq \beta_s v_j + \beta_t v_i - 2\epsilon$$

and thus

$$v_j \geq v_i - \frac{2\epsilon}{\beta_t - \beta_s}$$

for every $\epsilon > 0$. This contradicts the assumption that $v_i > v_j$.

For the third claim, consider a Nash equilibrium $b = (b_1, \ldots, b_n)$ and assume without loss of generality that it leads to an assignment where agent i is assigned position i for $i = 1, \ldots, k$. Further assume that the assignment is efficient, i.e., that $v_1 \geq v_2 \geq \cdots \geq v_k$. For $1 \leq i \leq k$, agent $i + 1$ does not benefit from bidding $b_{i,i} + \epsilon = p_i + \epsilon$ on position i and above, so

$$\beta_{i+1} v_{i+1} - p_{i+1} \geq \beta_i v_{i+1} - p_i - \epsilon$$

for every $\epsilon > 0$. Thus, for every $\epsilon > 0$,

$$p_k \geq \beta_k v_{k+1} - \epsilon \quad \text{and}$$
$$p_i \geq (\beta_i - \beta_{i+1})v_{i+1} + p_{i+1} - \epsilon \quad \text{for } 1 \leq i < k,$$

which proves the claim. $\square$

4. INCOMPLETE INFORMATION

We now consider environments with incomplete information and show our main result, that the expressive GFP auction always has an efficient equilibrium that yields the truthful VCG revenue.

THEOREM 4.1. *Assume that agent valuations in a position auction on k positions are drawn independently from a continuous distribution on $\mathbb{R}_\beta^k$ with bounded support. Then the expressive GFP auction has an efficient Bayes-Nash equilibrium with the same payments as the truthful equilibrium of the VCG auction.*

We prove this result by constructing a bidding function $b^* : \mathbb{R} \to \mathbb{R}_\geq^k$, and showing by induction that an agent maximizes its utility by bidding according to b^*, assuming that all other agents bid according to b^* as well. To this end, we define in Section 4.1 a function $b_j^* : \mathbb{R} \to \mathbb{R}$ for each position j that maps a valuation v to the expected truthful VCG payment $b_j^*(v)$ an agent with valuation v would face if it was allocated position j. The equilibrium bidding function b^* will then be given by $b^*(v) = (b_1^*(v), \ldots, b_k^*(v))$. We will say that an agent with valuation v bids truthfully on position j (according to b^*) if it bids $b_j^*(v)$, and that it bids truthfully if it bids truthfully on all positions. The property we show by induction is that independently of the agent's bids on positions $1, \ldots, j-1$ and assuming truthful bids on positions $j+1, \ldots, k$, it is optimal to bid truthfully on position j. For this we apply the standard technique, and integrate the derivative of the utility function from the truthful bid on position j to a conjectured beneficial deviation on position j to derive a contradiction.

Denote by $u^*((x_1, \ldots, x_k), v)$ the expected utility of an agent with valuation v who bids $b_j^*(x_j)$ on position j, for $1 \leq j \leq k$, while all other agents bid truthfully. The proof of Theorem 4.1 uses two lemmata, which we respectively prove in Sections 4.2 and 4.3.

LEMMA 4.2. *Fix a particular agent and a position j. Assume that all other agents bid truthfully and that the agent bids truthfully on positions $j+1, \ldots, k$. Then the derivative of the agent's expected utility with respect to the bid on position j vanishes at the truthful bid, i.e.,*

$$\frac{d}{dx_j} u^*((x_1, \ldots, x_j, v, \ldots, v), v)\Big|_{x_j = v} = 0.$$

LEMMA 4.3. *Fix a particular agent and a position j. Assume that all other agents bid truthfully and that the agent bids truthfully on positions $j+1, \ldots, k$. Then, the derivative in the valuation of the derivative in the bid on position j of the agent's expected utility is non-negative, i.e.,*

$$\frac{d}{dv}\frac{d}{dx_j} u^*((x_1, \ldots, x_j, v, \ldots, v), v) \geq 0.$$

Using the lemmata we prove our main theorem.

PROOF OF THEOREM 4.1. Fix a particular agent and assume that all other agents bid truthfully. Suppose that we have established the claim for positions $j+1, \ldots, k$, and that we want to establish it for position j. The claim trivially holds for $j = k$, so

we know from the induction hypothesis that

$$u^*(x_1, \ldots, x_j, v, \ldots, v) = \max\{x_{j+1}, \ldots, x_k : u^*(x_1, \ldots, x_j, x_{j+1}, \ldots, x_k)\}.$$

To show that the claim holds for position j, assume for contradiction that there exists $v' \in \mathbb{R}_{\geq 0}$ such that

$$u^*(x_1, \ldots, x_{j-1}, v', v, \ldots, v) > u^*(x_1, \ldots, x_{j-1}, v, v, \ldots, v).$$

First consider the case where $v' < v$. Then,

$$u^*((x_1, \ldots, x_{j-1}, v, \ldots, v), v) - u^*((x_1, \ldots, x_{j-1}, v', v, \ldots, v), v) =$$

$$\int_{v'}^{v} \frac{d}{dx_j} u^*((x_1, \ldots, x_{j-1}, x_j, v, \ldots, v), v) \, dx_j \geq$$

$$\int_{v'}^{v} \frac{d}{dy} u^*((x_1, \ldots, x_{j-1}, y, x_j, \ldots, x_j), x_j)\Big|_{y=x_j} dx_j = 0,$$

where the inequality and the last equality respectively hold by Lemma 4.3 and Lemma 4.2. This is a contradiction.

If $v' > v$, we can proceed analogously to show that the deviation is not beneficial. Note that when all other agents bid according to b^*, it is enough to consider bids $b_j^*(v)$ with v in the support of F. Any other bid will be dominated by a bid of this type. $\square$

4.1. Truthful VCG Payments and Allocation Probabilities

We begin by formally defining the position-specific bidding functions b_j^*, and computing their derivative with respect to the valuation. We then derive a recursive formulation of the allocation probabilities, which will be used in the proofs of Lemma 4.2 and Lemma 4.3.

Bid $b_j^*(v)$ equals the truthful VCG payment for position j given valuation v and conditioned on allocation of position j. This quantity is equal to the sum of the differences $\beta_s - \beta_{s+1}$ multiplied by the expected value of the $s+1$-highest valuation among all agents conditioned on v being the j-highest valuation and assuming that valuations are drawn independently from distribution F. Formulaically,

$$b_j^*(v) = \sum_{s=j}^{k} (\beta_s - \beta_{s+1}) \int_0^v \frac{(n-j)!}{(n-s-1)!(s-j)!} \left(\frac{F(u)}{F(v)}\right)^{n-s-1} \left(1 - \frac{F(u)}{F(v)}\right)^{s-j} \frac{f(u)}{F(v)} u \, du$$

Using the fact that $(1 - \frac{F(u)}{F(v)})^{s-j} = \sum_{t=0}^{s-j}(-1)^t(\frac{F(u)}{F(v)})^t$, and letting $Z_{n-s+t}(v) = (\frac{1}{F(v)})^{n-s+t}\int_0^v F(u)^{n-s+t} \, du$, we have that

$$b_j^*(v) = \sum_{s=j}^{k} (\beta_s - \beta_{s+1}) \sum_{t=0}^{s-j}(-1)^t \binom{s-j}{t} \frac{(n-j)!}{(n-s-1)!(s-j)!} \frac{1}{(n-s+t)} \left(v - Z_{n-s+t}(v)\right).$$

Using that $\frac{d}{dv}(v - Z_{n-s+t}(v)) = (n-s+t)\frac{f(v)}{F(v)} Z_{n-s+t}(v)$, we obtain

$$\frac{d}{dv} p_j(v) = \sum_{s=j}^{k} (\beta_s - \beta_{s+1}) \sum_{t=0}^{s-j}(-1)^t \binom{s-j}{t} \frac{(n-j)!}{(n-s-1)!(s-j)!} \frac{f(v)}{F(v)} Z_{n-s+t}(v). \quad (3)$$

Denote by $P_{s,m}(x)$ the probability that an agent is assigned position s against m opposing agents if it reports a valuation vector $x \in \mathbb{R}_{\geq}^k$. Then $P_{s,m}(x)$ can be written

recursively as

$$P_{1,m}(x) = F(x_1)^m, \quad \text{and}$$

$$P_{s,m}(x) = \binom{m}{m-s+1} F(x_s)^{m-s+1}\left(1 - \sum_{t=1}^{s-1} P_{t,s-1}(x)\right) \tag{4}$$

The intuition behind this formulation is that the agent is assigned position s if $m - s + 1$ of the opposing agents have valuations smaller than x_s, and if the agent is not assigned one of the positions $1, .., s-1$ against the remaining $s-1$ agents. An important observation is that $P_{s,m}(x)$ does not depend on x_ℓ for $\ell > s$.

4.2. Proof of Lemma 4.2

To prove Lemma 4.2, we write the expected utility that agent i with value v can achieve with a report $x \in \mathbb{R}_{\geq}^k$ as a sum of the contributions $T_s(x, v) = P_{s,n-1}(x)(\beta_s v - b_s^*(x_s))$ of position s. We then group these contributions into those of positions $s < j$, those of positions j and $j + 1$, and those of positions $s > j + 1$, and argue for each group separately that the derivative in x_j vanishes at $x_j = v$.

For the contribution $\sum_{s=1}^{j-1} T(x, v)$ of positions $s < j$ this is rather straightforward, as neither the allocation probability $P_{s,n-1}(x)$, nor the utility $\beta_s v - b_s^*(x_s)$ subject to allocation, depends on x_j. Hence the derivative in x_j is zero everywhere, and in particular at $x_j = v$.

To prove the claim for $T_j(x, v) + T_{j+1}(x, v)$, we first apply the recursive formulation of the allocation probabilities to compute the derivatives in x_j of $T_j(x, v)$ and $T_{j+1}(x, v)$. We then observe that the derivative of $T_j(x, v) + T_{j+1}(x, v)$ vanishes at $x_j = v$ if and only if a certain differential equation involving the bids $b_j^*(v)$ and $b_{j+1}^*(v)$ is satisfied. Finally, we use the formulas for the truthful VCG payments conditioned on allocation and their derivatives to show that this differential equation is satisfied.

LEMMA 4.4. *Fix a particular agent and a position j. Assume that all other agents bid truthfully and that the agent bids truthfully on positions $j + 1, \ldots, k$. Then,*

$$\frac{d}{dx_j}\left(T_j(x, v) + T_{j+1}(x, v)\right)\Big|_{x_j = v} = 0.$$

PROOF. First consider the contribution $T_j(x, v) = P_{j,n-1}(x)(\beta_j v - b_j^*(x_j))$ of position j. By applying (4) to $P_{j,n-1}(x)$,

$$T_j(x, v) = \binom{n-1}{n-j} F(x_j)^{n-j}\left(1 - \sum_{t=1}^{j-1} P_{t,j-1}(x)\right)(\beta_j v - b_j^*(x_j)),$$

and thus

$$\frac{d}{dx_j} T_j(x, v) = \binom{n-1}{n-j}\left(1 - \sum_{t=1}^{j-1} P_{t,j-1}(x)\right)\left((n-j)F(x_j)^{n-j-1}f(x_j)\beta_j v - \right.$$

$$\left. (n-j)F(x_j)^{n-j-1}f(x_j)b_j^*(x_j) - F(x_j)^{n-j}\frac{d}{dx_j}b_j^*(x_j)\right).$$

Now consider the contribution $T_{j+1}(x, v) = P_{j+1,n-1}(x)(\beta_{j+1} v - b_{j+1}^*(x_{j+1}))$ of position $j + 1$. By applying (4) to $P_{j+1,n-1}(x)$,

$$T_{j+1}(x, v) = \binom{n-1}{n-j-1} F(v)^{n-j-1}\left(1 - \sum_{t=1}^{j} P_{t,j}(x)\right)(\beta_{j+1} v - b_{j+1}^*(v)).$$

By pulling $P_{j,j}(x)$ out of the sum and applying (4) to it, we obtain

$$T_{j+1}(x,v) = \binom{n-1}{n-j-1}F(v)^{n-j-1} \cdot \left(1 - \sum_{t=1}^{j-1} P_{t,j}(x) - \right.$$

$$\left. \binom{j}{1}F(x_j)\left(1 - \sum_{t=1}^{j-1} P_{t,j-1}(x)\right)\right)(\beta_{j+1}v - b_{j+1}^*(v)),$$

and thus

$$\frac{d}{dx_j}T_{j+1}(x,v) = -\binom{n-1}{n-j-1}F(v)^{n-j-1}\binom{j}{1}f(x_j)\left(1 - \sum_{t=1}^{j-1} P_{t,j-1}(x)\right)(\beta_{j+1}v - b_{j+1}^*(v)).$$

We conclude that the derivative in x_j of the contribution $T_j(x,v) + T_{j+1}(x,v)$ of positions j and $j+1$ vanishes at $x_j = v$ if and only if

$$\binom{n-1}{n-j}\left((n-j)F(v)^{n-j-1}f(v)\beta_j v - \right.$$

$$\left. (n-j)F(v)^{n-j-1}f(v)b_j^*(v) - F(v)^{n-j}\frac{d}{dx_j}b_j^*(x_j)\Big|_{x_j=v}\right)$$

$$- \binom{n-1}{n-j-1}F(v)^{n-j-1}\binom{j}{1}f(v)(\beta_{j+1}v - b_{j+1}^*(v)) = 0.$$

By using $\binom{n-1}{n-j-1}\binom{j}{1} = \binom{n-1}{n-j}(n-j)$ to simplify and by rearranging, we obtain the following differential equation:

$$\frac{d}{dx_j}b_j^*(x_j)\Big|_{x_j=v} = (n-j)\frac{f(v)}{F(v)}\left[(\beta_j v - b_j^*(v)) - (\beta_{j+1}v - b_{j+1}^*(v))\right].$$

We first observe that the v parts of $b_j^*(v)$ and $b_{j+1}^*(v)$ cancel $\beta_j v$ and $\beta_{j+1}v$. This is the case because for $z \in \{j, j+1\}$, the v part of $b_z^*(v)$ is equal to

$$\sum_{s=z}^{k}(\beta_s - \beta_{s+1})\underbrace{\sum_{t=0}^{s-z}(-1)^t\binom{s-z}{t}\frac{(n-j)!}{(n-s-1)!(s-z)!}\frac{v}{(n-s+t)}}_{=1} = \sum_{s=z}^{k}(\beta_s - \beta_{s+1})v = \beta_z v.$$

It remains to show that $(n-j)\frac{f(v)}{F(v)}$ times the Z part of $b_{j+1}^*(v)$ minus the Z part of $b_j^*(v)$ is equal to the derivative in x_j of $b_j^*(x_j)$ at $x_j = v$. Formulaically, the former can be expressed as

$$(n-j)\frac{f(v)}{F(v)}\left[\sum_{s=j}^{k}(\beta_s - \beta_{s+1})\sum_{t=0}^{s-j}(-1)^t\binom{s-j}{t}\frac{(n-j)!}{(n-s-1)!(s-j)!} \cdot \right.$$

$$\frac{1}{(n-s+t)}Z_{n-s+t}(v) - \sum_{s=j+1}^{k}(\beta_s - \beta_{s+1})\sum_{t=0}^{s-j-1}(-1)^t \cdot \tag{5}$$

$$\left.\binom{s-j-1}{t}\frac{(n-j-1)!}{(n-s-1)!(s-j-1)!}\frac{1}{(n-s+t)}Z_{n-s+t}(v)\right].$$

We prove the identity by showing that for all s and t, the terms in (5) are identical to the corresponding terms in (3).

For $s = j$, the only possible value for t is $t = 0$, so the term in (5) is

$$(n-j)\frac{f(v)}{F(v)}(\beta_s - \beta_{s+1})(-1)^t \binom{s-j}{t}\frac{(n-j)!}{(n-s-1)!(s-j)!}\frac{1}{(n-s+t)}Z_{n-s+t}(v).$$

Since $(n-j) = (n-s+t)$, we obtain the corresponding term in (3).

For $s > j$ and any t in the correct range, the term in (5) is

$$(n-j)\frac{f(v)}{F(v)}\left[(\beta_s - \beta_{s+1})(-1)^t \binom{s-j}{t}\frac{(n-j)!}{(n-s-1)!(s-j)!}\cdot\right.$$

$$\frac{1}{(n-s+t)}Z_{n-s+t}(v) - (\beta_s - \beta_{s+1})(-1)^t \binom{s-j-1}{t}$$

$$\left.\frac{(n-j-1)!}{(n-s-1)!(s-j-1)!}\frac{1}{(n-s+t)}Z_{n-s+t}(v)\right],$$

which using $\binom{s-j-1}{t} = \binom{s-j}{t}\frac{s-j-t}{s-j}$ can be rewritten as

$$(\beta_s - \beta_{s+1})(-1)^t \binom{s-j}{t}\frac{(n-j)!}{(n-s-1)!(s-j)!}\frac{f(v)}{F(v)}\left[\frac{n-j}{n-s+t} - \frac{s-j-t}{n-s+t}\right]Z_{n-s+t}(v).$$

Since $\frac{n-j}{n-s+t} - \frac{s-j-t}{n-s+t} = 1$, we obtain the corresponding term in (3). $\qquad\square$

Next we consider the contribution $\sum_{s=j+2}^{k} T_s(x, v)$ of positions $s > j + 1$.

LEMMA 4.5. *Fix a particular agent. Assume that all other agents bid truthfully and that the agent bids truthfully on positions $j + 1, \ldots, k$. Then,*

$$\frac{d}{dx_j}\left(\sum_{s=j+2}^{k} T_s(x, v)\right)\bigg|_{x_j=v} = 0.$$

Note that for position $s > j + 1$, the contribution $T_s(x, v) = P_{s,n-1}(x, v)(\beta_s v - b_s^*(x_s))$ only depends on x_j through the allocation probability $P_{s,n-1}(x, v)$. It therefore suffices to show that the derivative in x_j of $P_{s,n-1}(x, v)$ vanishes at $x_j = v$. We establish this claim by means of two auxiliary lemmata, which are proved in the appendix and again exploit the recursive formulation of the allocation probabilities.

LEMMA 4.6. *Fix a particular agent and a position j. Assume that all other agents bid truthfully and that the agent bids truthfully on positions $j + 1, \ldots, k$. Then, for all $m \geq j + 1$,*

$$\frac{d}{dx_j}\left(P_{j,m}(x) + P_{j+1,m}(x)\right)\bigg|_{x_j=v} = 0.$$

LEMMA 4.7. *Fix a particular agent and a position j. Assume that all other agents bid truthfully and that the agent bids truthfully on positions $j + 1, \ldots, k$. Then, for all m and ℓ such that $m \geq \ell \geq j + 2$,*

$$\frac{d}{dx_j}P_{\ell,m}(x)\bigg|_{x_j=v} = 0.$$

PROOF OF LEMMA 4.5. For position $s > j+1$ we first apply (4) to $P_{s,n-1}(x)$ to obtain

$$T_s(x, v) = \binom{n-1}{n-s}F(x_s)^{n-s}\left(1 - \sum_{t=1}^{s-1} P_{t,s-1}(x)\right)(\beta_s v - b_s^*(x_s)),$$

and then split $\sum_{t=1}^{s-1} P_{t,s-1}(x)$ into two parts to obtain

$$T_s(x,v) = \binom{n-1}{n-s} F(x_s)^{n-s} \left(1 - \sum_{t=1}^{j-1} P_{t,s-1}(x) - \sum_{t=j}^{s-1} P_{t,s-1}(x) \right) (\beta_s v - b_s^*(x_s)).$$

The derivative is thus

$$\frac{d}{dx_j} T_s(x,v) = \binom{n-1}{n-s} F(x_s)^{n-s} \left(-\frac{d}{dx_j} \sum_{t=j}^{s-1} P_{t,s-1}(x) \right) (\beta_s v - b_s^*(x_s)),$$

and we use Lemma 4.6 and Lemma 4.7 to conclude that it vanishes at $x_j = v$. $\qquad\square$

4.3. Proof of Lemma 4.3

We now turn to Lemma 4.3, and begin by recalling the results for the one-dimensional case. In this case the expected utility for report x given value v is equal to

$$\sum_{s=1}^{k} \beta_s P_{s,n-1}(x)(v-x) + \sum_{s=1}^{k} \beta_s \int_0^x P_{s,n-1}(t)\, dt,$$

which for truthful report $x = v$ simplifies to

$$\sum_{s=1}^{k} \beta_s \int_0^v P_{s,n-1}(t)\, dt. \tag{6}$$

We will use this formula to express the expected utility from positions $j+1,\ldots,k$ for which both agent i and the other agents report their valuations truthfully.

To compute the derivative in x_j of the expected utility we first observe that the contribution $T_s(x,v)$ is independent from x_j for $s < j$, and thus

$$\frac{d}{dx_j} \left(\sum_{s=1}^{k} T_s(x,v) \right) = \frac{d}{dx_j} \left(\sum_{s=j}^{k} T_s(x,v) \right) = \frac{d}{dx_j} \left(T_j(x,v) + \sum_{s=j+1}^{k} T_s(x,v) \right).$$

For the contribution $T_j(x,v)$ of position j,

$$\frac{d}{dx_j} T_j(x,v) = \frac{d}{dx_j} \left(P_{j,n-1}(x)(\beta_j v - b_j^*(x_j)) \right)$$

$$= \beta_j v \frac{d}{dx_j} P_{j,n-1}(x) - b_j^*(x_j) \frac{d}{dx_j} P_{j,n-1}(x) - P_{j,n-1}(x) \frac{d}{dx_j} b_j^*(x_j).$$

For the contributions $T_s(x,v)$ of positions $s > j$ we use (6) to obtain

$$\frac{d}{dx_j} \sum_{s=j+1}^{k} T_s(x,v) = \frac{d}{dx_j} \left(\sum_{s=j+1}^{k} \beta_s \int_0^v P_{s,n-1}(x_1,\ldots,x_j,t,\ldots,t) dt \right)$$

$$= \sum_{s=j+1}^{k} \beta_s \int_0^v \frac{d}{dx_j} P_{s,n-1}(x_1,\ldots,x_j,t,\ldots,t) dt.$$

Taking the derivative in v yields

$$\frac{d}{dv} \left(\frac{d}{dx_j} \sum_{s=1}^{k} T_s(x,v) \right) = \beta_j \frac{d}{dx_j} P_{j,n-1}(x) + \sum_{s=j+1}^{k} \beta_s \frac{d}{dx_j} P_{s,n-1}(x) = \frac{d}{dx_j} \sum_{s=j}^{k} \beta_s P_{s,n-1}(x).$$

The final step is to argue that this expression is non-negative.

That the β fraction won increases in the report x_j on position j holding everything else fixed follows by an ex-post argument. If the agent was allocated a position $s < j$ then changing its reported valuation x_j for position j has no effect and it will still be allocated position s. If the agent was allocated position $s = j$, it will still be allocated this position for higher x_j. If the agent was allocated a position $s > j$ or no position at all, then by increasing the reported valuation x_j for position j it will either be allocated the same position as before or position j, which means that the β fraction won will increase weakly.

5. CONCLUSIONS

Analyzing position auctions through the lens of robustness, we have asked whether there exists a single design that performs well under complete and incomplete information, in the sense that it obtains the truthful VCG revenue in every efficient equilibrium of both settings. We have shown that an expressive GFP auction achieves the desired robustness property, and is in fact the only standard design that does.

Our results send a clear, and surprising, message: if the goal is robustness against uncertainty about the information agents have about one another, then it is necessary to allow for expressiveness beyond the valuation space. This also provides a counterpoint to recent work on position auctions that has highlighted the benefits of simplicity.

An interesting question for future work is whether expressiveness is necessary for robustness in other contexts as well. Natural candidates include settings with multi-dimensional valuations such as combinatorial auctions, where simplified designs have recently received a lot of attention [Christodoulou et al. 2008; Bhawalkar and Roughgarden 2011; Feldman et al. 2013; Dütting et al. 2013; Babaioff et al. 2014], and two-sided settings with strategic buyers and sellers, such as the assortative matching problem considered by Hoppe et al. [2009].

REFERENCES

ABRAMS, Z., GHOSH, A., AND VEE, E. 2009. Cost of conciseness in sponsored search auctions. In *Proceedings of the 5th Workshop on Internet and Network Economics*. 326–334.

AGGARWAL, G., MUTHUKRISHNAN, S., PÁL, D., AND PÁL, M. 2009. General auction mechanism for search advertising. In *Proceedings of the 18th International Conference on World Wide Web*. 241–250.

BABAIOFF, M., LUCIER, B., PAES LEME, R., AND NISAN, N. 2014. On the efficiency of the Walrasian mechanism. In *Proceedings of the 15th ACM Conference on Economics and Computation*. Forthcoming.

BENISCH, M., SADEH, N., AND SANDHOLM, T. 2008. A theory of expressiveness in mechanisms. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence*. 17–23.

BERNHEIM, B. D. AND WHINSTON, M. D. 1986. Menu auctions, resource allocation and economic influence. *Quarterly Journal of Economics 101*, 1–31.

BHAWALKAR, K. AND ROUGHGARDEN, T. 2011. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*. 700–709.

BLUMROSEN, L., HARTLINE, J., AND NONG, S. 2008. Position auctions and non-uniform conversion rates. In *Proceedings of the 8th Ad Auctions Workshop*. 1–11.

CARAGIANNIS, I., KAKLAMANIS, C., KANELLOPOULOS, P., AND KYROPOULOU, M. 2011. On the efficiency of equilibria in generalized second price auctions. In *Proceedings of the 12th ACM Conference on Electronic Commerce*. 81–90.

CARAGIANNIS, I., KAKLAMANIS, C., KANELLOPOULOS, P., AND KYROPOULOU, M. 2012. Revenue guarantees in sponsored search auctions. In *Proceedings of the 20th European Symposium on Algorithms*. 253–264.

CHAWLA, S. AND HARTLINE, J. 2013. Auctions with unique equilibria. In *Proceedings of the 14th ACM Conference on Electronic Commerce*. 181–196.

CHRISTODOULOU, G., KOVÁCS, A., AND SCHAPIRA, M. 2008. Bayesian combinatorial auctions. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*. 820–832.

CONSTANTIN, F., RAO, M., HUANG, C.-C., AND PARKES, D. C. 2011. On expressing value externalities in position auctions. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence*. 644–649.

DAY, R. AND MILGROM, P. 2008. Core-selecting package auctions. *International Journal of Game Theory 36,* 3, 393–407.

DOBZINSKI, S., LAVI, R., AND NISAN, N. 2012. Multi-unit auctions with budget limits. *Games and Economic Behavior 74,* 2, 486–503.

DÜTTING, P., FISCHER, F., AND PARKES, D. C. 2011a. Simplicity-expressiveness tradeoffs in mechanism design. In *Proceedings of the 12th ACM Conference on Electronic Commerce*. 341–350.

DÜTTING, P., HENZINGER, M., AND STARNBERGER, M. 2012. Auctions for heterogeneous items and budget limits. In *Proceedings of the 8th Workshop on Internet and Network Economics*. 44–57.

DÜTTING, P., HENZINGER, M., AND STARNBERGER, M. 2013. Valuation compressions in VCG-based combinatorial auctions. In *Proceedings of the 9th Conference on Web and Internet Economics*. 146–159.

DÜTTING, P., HENZINGER, M., AND WEBER, I. 2011b. An expressive mechanism for auctions on the web. In *Proceedings of the 20th International Conference on World Wide Web*. 127–136.

EDELMAN, B., OSTROVSKY, M., AND SCHWARTZ, M. 2007. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review 97,* 1, 242–259.

FELDMAN, M., FU, H., GRAVIN, N., AND LUCIER, B. 2013. Simultaneous auctions are (almost) efficient. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*. 201–210.

GAIRING, M. 2009. Covering games: Approximation through non-cooperation. In *Proceedings of the 5th Workshop on Internet and Network Economics*. 184–195.

GHOSH, A. AND SAYEDI, A. 2010. Expressive auctions for externalities in online advertising. In *Proceedings of the 19th International Conference on World Wide Web*. 371–380.

GOEL, G., MIRROKNI, V. S., AND PAES LEME, R. 2012. Polyhedral clinching auctions and the adwords polytope. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*. 107–122.

GOMES, R. D. AND SWEENEY, K. S. 2009. Bayes-Nash equilibria of the generalized second price auction. In *Proceedings of the 10th ACM Conference on Electronic Commerce*. 107–108.

GUL, F. AND STACCHETTI, E. 1999. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory 87,* 1, 95–124.

HASSIDIM, A., KAPLAN, H., MANSOUR, Y., AND NISAN, N. 2011. Non-price equilibria in markets of discrete goods. In *Proceedings of the 12th ACM Conference on Electronic Commerce*. 295–296.

HOPPE, H. C., MOLDOVANU, B., AND SELA, A. 2009. The theory of assortative matching based on costly signals. *Review of Economic Studies 76,* 1, 253–281.

HOY, D., JAIN, K., AND WILKENS, C. A. 2013. A dynamic axiomatic approach to first-price auctions. In *Proceedings of the 14th ACM Conference on Electronic Commerce*. 583–584.

LUCIER, B. AND PAES LEME, R. 2011. GSP auctions with correlated types. In *Proceedings of the 12th ACM Conference on Electronic Commerce*. 71–80.

LUCIER, B., PAES LEME, R., AND TARDOS, E. 2012. On revenue in the generalized second price auction. In *Proceedings of the 21st International Conference on World Wide Web*. 361–370.

MILGROM, P. 2004. *Putting Auction Theory to Work*. Cambridge University Press.

MILGROM, P. 2010. Simplified mechanisms with an application to sponsored-search auctions. *Games and Economic Behavior 70,* 1, 62–70.

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 51st Symposium on Foundations of Computer Science*. 735–744.

SYRGKANIS, V. AND TARDOS, E. 2013. Composable and efficient mechanisms. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*. 211–220.

VARIAN, H. 2007. Position auctions. *International Journal of Industrial Organization 25*, 1163–1178.

A. PROOF OF LEMMA 4.6

First consider the allocation probability $P_{j,m}(x)$ for position j. Applying (4) to $P_{j,m}(x)$ yields

$$P_{j,m}(x) = \binom{m}{m-j+1} F(x_j)^{m-j+1} \left(1 - \sum_{t=1}^{j-1} P_{t,j-1}(x) \right),$$

and thus

$$\frac{d}{dx_j} P_{j,m}(x) = \binom{m}{m-j+1}(m-j+1)F(x_j)^{m-j+1}(m-j)f(x_j)\left(1 - \sum_{t=1}^{j-1} P_{t,j-1}(x) \right).$$

Now consider the allocation probability $P_{j+1,m}(x)$ of position $j+1$. Applying (4) to $P_{j+1,m}(x)$ yields

$$P_{j+1,m}(x) = \binom{m}{m-j} F(v)^{m-j} \left(1 - \sum_{t=1}^{j} P_{t,j}(x) \right).$$

Pulling $P_{j,j}(x)$ out of the sum and applying (4) to it yields

$$P_{j+1,m}(x) = \binom{m}{m-j} F(v)^{m-j} \left(1 - \sum_{t=1}^{j-1} P_{t,j}(x) - \binom{j}{1} F(x_j) \left(1 - \sum_{t=1}^{j-1} P_{t,j-1}(x) \right) \right),$$

and thus

$$\frac{d}{dx_j}P_{j+1,m}(x) = \binom{m}{m-j}F(v)^{m-j}\left(-\binom{j}{1}f(x_j)\left(1-\sum_{t=1}^{j-1}P_{t,j-1}(x)\right)\right).$$

We conclude that the derivative in x_j of $P_{j,m}(x)+P_{j+1,m}(x)$ vanishes at $x_j = v$ if and only if

$$\binom{m}{m-j+1}(m-j+1)F(v)^{m-j}f(v) - \binom{m}{m-j}F(v)^{m-j}\binom{j}{1}f(v) = 0.$$

Since $\binom{m}{m-j}\binom{j}{1} = \binom{m}{m-j+1}(m-j+1)$, this is indeed the case.

B. PROOF OF LEMMA 4.7

We prove the claim by induction on m, starting with $m = j + 2$. In this case the only possible value of ℓ is $\ell = j + 2$, so it suffices to show that $\frac{d}{dx_j}P_{j+2,j+2}(x)\ |_{x_j=v}= 0$. Applying (4) to $P_{j+2,j+2}(x)$ shows that

$$P_{j+2,j+2}(x) = \binom{j+2}{1}F(v)\left(1-\sum_{t=1}^{j+1}P_{t,j+1}(x)\right).$$

By pulling $P_{j,j+1}(x)$ and $P_{j+1,j+1}(x)$ out of the sum this can be rewritten as

$$P_{j+2,j+2}(x) = \binom{j+2}{1}F(v)\left(1-\sum_{t=1}^{j-1}P_{t,j+1}(x) - \Big(P_{j,j+1}(x) + P_{j+1,j+1}(x)\Big)\right),$$

and thus

$$\frac{d}{dx_j}P_{j+2,j+2}(x) = \binom{j+2}{1}F(v)\frac{d}{dx_j}\left(-(P_{j,j+1}(x) + P_{j+1,j+1}(x))\right).$$

Using Lemma 4.6 we conclude that the derivative vanishes at $x_j = v$.

For the inductive step assume that the claim is true for all $m' < m$. We have to show that for any ℓ with $m \geq l \geq j + 2$ it holds that $\frac{d}{dx_j}P_{\ell,m}(x)\ |_{x_j=v}= 0$. Applying (4) to $P_{\ell,m}(x)$ yields

$$P_{\ell,m}(x) = \binom{m}{m-\ell+1}F(v)^{m-\ell+1}\left(1-\sum_{t=1}^{\ell-1}P_{t,\ell-1}(x)\right).$$

By splitting $\sum_{t=1}^{\ell-1}P_{t,\ell-1}(x)$ into three parts we obtain

$$P_{\ell,m}(x) = \binom{m}{m-\ell+1}F(v)^{m-\ell+1}\left(1-\sum_{t=1}^{j-1}P_{t,\ell-1}(x) - \sum_{t=j}^{j+1}P_{t,\ell-1}(x) - \sum_{t=j+2}^{\ell-1}P_{t,\ell-1}(x)\right),$$

and thus

$$\frac{d}{dx_j}P_{\ell,m}(x) = \binom{m}{m-\ell+1}F(v)^{m-\ell+1}\left(-\frac{d}{dx_j}\sum_{t=j}^{j+1}P_{t,\ell-1}(x) - \frac{d}{dx_j}\sum_{t=j+2}^{\ell-1}P_{t,\ell-1}(x)\right).$$

Using Lemma 4.6 and the induction hypothesis we conclude that the derivative again vanishes at $x_j = v$.