



Dütting, P., Fischer, F., and Parkes, D. C. (2016) Truthful Outcomes from Non-Truthful Position Auctions. 17th ACM Conference on Economics and Computation, Maastricht, The Netherlands, 24-28 Jul 2016. p. 813. ISBN 9781450339360 (doi:[10.1145/2940716.2940731](https://doi.org/10.1145/2940716.2940731))

This is the author's final accepted version.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/134756/>

Deposited on: 20 January 2017

Truthful Outcomes from Non-Truthful Position Auctions*

Paul Dütting[†]

Felix Fischer[‡]

David C. Parkes[§]

Abstract

We exhibit a property of the VCG mechanism that can help explain the surprising rarity with which it is used even in settings with unit demand: a relative lack of robustness to inaccuracies in the choice of its parameters. For a standard position auction environment in which the auctioneer may not know the precise relative values of the positions, we show that under both complete and incomplete information a non-truthful mechanism supports the truthful outcome of the VCG mechanism for a wider range of these values than the VCG mechanism itself. The result for complete information concerns the generalized second-price mechanism and lends additional theoretical support to the use of this mechanism in practice. Particularly interesting from a technical perspective is the case of incomplete information, where a surprising combinatorial equivalence helps us to avoid confrontation with an unwieldy differential equation.

1 Introduction

The Vickrey-Clarke-Groves (VCG) mechanism stands as one of the pillars of mechanism design theory, but in the real world is used with surprising rarity. In past work, this mismatch between theory and practice has been attributed to a number of properties that affect the mechanism in certain settings, like susceptibility to collusion or prohibitive computational costs [2, 30]. Here we identify another property of the mechanism that may be problematic in practice, a relative lack of robustness to inaccuracies in the choice of its parameters. Unlike most of the known deficiencies it applies already in settings with unit demand.

1.1 The Model

We start from the standard position auction model of Edelman et al. [16] and Varian [33], where n bidders compete for the assignment of k positions. Bidders have unit demand and the value of bidder i for position j is given by $\beta_j \cdot v_i$, where v_i is a bidder-specific value and a non-increasing vector

*Valuable feedback from Dirk Bergemann, Allan Borodin, the anonymous referees, and seminar participants at Google, Facebook, Universität Zürich, and the University of Glasgow is gratefully acknowledged. The second author is supported by the Einstein Foundation Berlin. Part of the work was done while the first author was at the Department of Mathematics of the London School of Economics and Political Science and the second author at the Statistical Laboratory of the University of Cambridge.

[†]Departement Informatik, Eidgenössische Technische Hochschule Zürich, Universitätstrasse 6, 8092 Zürich, Switzerland. Email: paul.duetting@inf.ethz.ch

[‡]Institut für Mathematik, Technische Universität Berlin, Straße des 17. Juni 136, 10623 Berlin, Germany. Email: fischerf@math.tu-berlin.de.

[§]John A. Paulson School of Engineering and Applied Sciences, Harvard University, 33 Oxford Street, Cambridge, MA 02138, USA. Email: parkes@eecs.harvard.edu

$\beta = (\beta_1, \dots, \beta_k)$ describes the relative values of the positions. The most prominent application of this model is to sponsored search, which contributed over 90% to Google’s revenue of \$66 billion in 2014. What is rather curious is that all successful sponsored search services have used non-truthful auction mechanisms. Overture, the first company to provide such a service, used a generalized first-price (GFP) mechanism. Google and Microsoft use a generalized second-price (GSP) mechanism. Facebook does use the VCG mechanism to place ads, but not in the context of sponsored search and not in a position auction. It is hard to say in retrospect what led to the choice of non-truthful mechanisms over a truthful one, and changing the mechanism at this point would clearly come with huge financial risks. We will see, however, that under certain assumptions choosing the non-truthful mechanisms may have been wise even if it was not entirely deliberate.

A mechanism for the above setting accepts a bid b_i for each bidder i and from these bids determines a one-to-one assignment of bidders to positions and a monetary payment for each bidder. Bidder behavior can then be analyzed using game-theoretic reasoning, where it is commonly assumed that each bidder tries to maximize the value of its assigned position minus its payment. In the VCG mechanism in particular it is an equilibrium for each bidder to bid its true valuation, and we refer to the resulting assignment and payments in a given situation as the truthful VCG outcome.

In addition to the standard model of auction theory, where bidders have only probabilistic information about one another’s valuations, it has become common to analyze position auctions also in a complete-information model where valuations are common knowledge among the bidders. This is motivated by practical auctions that provide bidders with aggregate statistics of others’ bids and thus enable best-response bidding, by empirical support that has been given for a family of Nash equilibria [33, 16], and by cyclic bidding behavior observed in the absence of pure Nash equilibria [16].¹

The bid b_i can be interpreted alternatively as a vector of bids $\beta_j \cdot b_i$, one for each position j . Here we vary the standard model by assuming that mechanisms instead work with bids $\alpha_j \cdot b_i$, where $\alpha = (\alpha_1, \dots, \alpha_k)$ is another non-increasing vector and generally $\alpha \neq \beta$. We then ask for which values of α and β different mechanisms obtain the truthful VCG outcome in equilibrium. Intuitively this question concerns the robustness of mechanisms to uncertainty regarding the true value of β . It draws its motivation from current practice to infer β from data using machine learning techniques, which will hopefully produce a good estimate $\alpha \approx \beta$ but will usually mean that $\alpha \neq \beta$.² The truthful VCG outcome is a natural point of reference for at least two reasons. Under complete information it coincides with the bidder-optimal envy-free outcome, which happens to be unique [24]. Under incomplete information it is the unique outcome compatible with the maximization of social welfare [29].³

¹In sponsored search an auction is carried out for a specific search term. Intuitively a complete-information model seems more appropriate for common search terms, an incomplete-information model for those that are rare.

²Search engines do observe when an ad is clicked, and if the relative values of the positions were proportional to the relative number of clicks all auctions we consider could be implemented without knowledge of β [34]. This is, however, not currently being done [e.g., 6, 19]. Indeed there are good reasons why the value of a position may depend on other factors besides the probability of a click. Milgrom [27] for example considers a situation with two types of users of a search engine, one of them genuinely interested in the products being advertised and one merely curious. The two types come with different rates at which clicks on an advertisement result in a purchase, but are indistinguishable from the point of view of the search engine.

³Indeed, work on position auctions has traditionally focused on the maximization of welfare, which in practice can be seen as a maximization of customer satisfaction to ensure long-term success. On a more technical level, the focus on welfare allows for a unified treatment of complete and incomplete information environments. Extending our results to the maximization of revenue nevertheless is an interesting direction for future work.

1.2 Results

We show that under both complete and incomplete information a non-truthful mechanism obtains the truthful VCG outcome in an equilibrium for a strictly larger set of values of α and β than the VCG mechanism itself. Failure of the VCG mechanism to support the truthful VCG outcome in fact occurs already when α is very close to β .

The result for complete information concerns the GSP mechanism and is developed in Section 3 in two steps. Through direct arguments that combine the equilibrium conditions and the structure of allocation and payments in the truthful VCG outcome we obtain precise characterizations, for both the VCG and the GSP mechanism, of those values of α and β that enable this outcome in equilibrium. We then exploit the recursive nature of VCG payments to show that any violation of the characterization for the GSP mechanism necessarily leads to a violation of the characterization for the VCG mechanism.

In Section 4 the same type of result is shown for incomplete information, but here the VCG mechanism is compared to the GFP mechanism and a more elaborate argument is required to establish superiority of the latter. We begin with a standard technique for characterizing equilibria, by equating the expected payments in a welfare-maximizing equilibrium as given by Myerson's Lemma with the respective payments in the two mechanisms. This gives us a candidate equilibrium bidding function for each the two mechanisms, and each of these functions constitutes an equilibrium if and only if it is strictly increasing almost everywhere. In the case of the VCG mechanism we encounter an ordinary differential equation, but avoid the use of heavy machinery by appealing to a surprising combinatorial equivalence. Even with the bidding functions for the VCG and GFP mechanisms at hand it is not trivial to show that the latter is increasing for a larger set of values of α and β , and we use a surprising connection between the two functions to show that this is indeed the case.

That the VCG mechanism fails already when α is very close to β is shown by means of two examples in Sections 3.1 and 4.1. In these examples equilibria supporting the truthful VCG outcome cease to exist when mechanisms underestimate the value of less valuable positions. This makes the less valuable positions more attractive by reducing their associated payments, incentivizing bidders to shade their bid and do so more strongly as their value increases. The relatively lower payments in the VCG mechanism only magnify this effect and cause it to fail for smaller discrepancies between α and β . In settings with sufficiently many positions the failure also occurs when mechanisms overestimate the quality of some positions but underestimate the quality of others.

More generally, the GSP and GFP mechanisms seem to benefit from the relative simplicity of their payments, which for a given position depend only on one bid. In the VCG mechanism a particular bid may simultaneously affect the payments of many bidders, setting the correct equilibrium payments thus becomes impossible more quickly as α and β move out of alignment. An orthogonal requirement for equilibrium existence that favors different non-truthful mechanisms under complete and incomplete information is that a bidder's ability to control its own payment must match the degree of knowledge it has of other bidders' valuations. It is well known, for example, that the GFP mechanism may not possess a pure Nash equilibrium under complete information [16] and that the GSP mechanisms may not possess a welfare-maximizing Bayes-Nash equilibrium under incomplete information [18], even when $\alpha = \beta$.

The main focus of our analysis is on mechanisms and parameters currently in use. The investigation of additional mechanisms and parameters, and of the interaction between the auction mechanisms and the learning algorithm used to infer β from data, provide ample scope for future work.

1.3 Related Work

An increased robustness of non-truthful mechanisms for position auctions in the sense we discuss here was first suggested by Milgrom [27], who compared the GSP and VCG mechanisms in a complete-information setting. This result was then strengthened by Dütting et al., who identified a single value of α with which the GSP and GFP mechanisms respectively obtain the truthful VCG outcome under complete and incomplete information for all values of β [12], as well as a single GFP mechanism with more expressive bids that achieves the same under both complete and incomplete information [13]. What distinguishes our results from this past work is that they characterize, for any given value of β , which values of α enable the truthful VCG outcome. They thus apply in particular to mechanisms currently in use.

The performance of the VCG mechanism and that of alternative, non-truthful mechanisms have also been compared in the standard position auction model, where $\alpha = \beta$, and a number of authors have noted certain advantages or lack of disadvantages of the alternative mechanisms. Under complete information, the GSP mechanism obtains the truthful VCG outcome in a locally envy-free equilibrium, and payments that at least match those of the truthful VCG outcome in any locally envy-free equilibrium [16, 33]. Under incomplete information the GFP mechanism admits a unique Bayes-Nash equilibrium, which yields the truthful VCG outcome [9]. Each of the two mechanisms has severe disadvantages in the respective other setting, such as non-existence of a pure Nash equilibrium or of an efficient Bayes-Nash equilibrium [16, 18]. In cases where equilibria exist, however, the worst-case welfare loss is bounded in the sense of a small price of anarchy [23, 8, 32].⁴ In the standard model, and other things being equal, the VCG mechanism of course has the advantage of truthfulness. Our results concern a more general and arguably more realistic model and apply in a very similar way to both complete- and incomplete-information environments.

A concurrent line of research has sought to emphasize the advantages of non-revelation mechanisms for position auctions under incomplete information [9, 10, 21]. In addition to uniqueness of equilibria these include computational tractability, amenability to statistical inference, and the ability to approximately maximize welfare in general one-dimensional environments. Our results finally fit more generally into an increasing body of work that emphasizes robustness and simplicity in economic and algorithmic design. Relevant examples of this type of work in economics include the literature on robust mechanism design [4] and the design of the upcoming FCC Incentive Auctions [28, 14]. The former seeks to obtain mechanisms with more robust common knowledge assumptions and has in fact identified robustness properties that directly lead to simpler mechanisms as an important direction. The latter employs a greedy mechanism to achieve computational and strategic simplicity. Additional examples come from algorithmic game theory, where recent work has obtained simple mechanisms with near-optimal revenue [20, 1, 3, 36] or welfare [11, 25, 5, 17, 26, 15], but has also pointed out computational barriers to near-optimal equilibria [7, 31].

⁴In the language of the price-of-anarchy literature we essentially seek to characterize those values of α and β for which the price of stability is one. Arguments similar to the ones used to establish the price-of-anarchy guarantees for the standard model also apply to the more general setting we study here, providing welfare guarantees that degrade gracefully in α and β .

2 Preliminaries

We study the standard setting of position auctions with k *positions* ordered by quality and n *agents* with *unit demand* and *one-dimensional valuations* for the positions. Denote by $\mathbb{R}_{\geq}^k = \{x \in \mathbb{R}^k : x_j > 0, x_j \geq x_{j'} \text{ if } j < j'\}$ the set of k -dimensional vectors whose entries are positive and non-increasing. Given $\beta \in \mathbb{R}_{\geq}^k$, which we assume to be common knowledge among the agents, the valuation of a particular agent i can then be represented by a scalar $v_i \in \mathbb{R}$, such that $\beta_j v_i \geq 0$ is the agent's value for position j .

A *mechanism* in this setting receives a profile $b \in \mathbb{R}^n$ of bids from the agents, assigns positions to agents in a one-to-one fashion, and charges each agent a non-negative payment. It can be represented by a pair (g, p) of an *allocation rule* $g : \mathbb{R}^n \rightarrow S_n$ and a *payment rule* $p : \mathbb{R}^n \rightarrow \mathbb{R}^n$, such that for each $i \in \{1, \dots, n\}$, $g_i(b) = j$ for $j \in \{1, \dots, k\}$ means that agent i is assigned position j , and $p_i(b)$ is the payment charged to agent i . Allocation rules and mechanisms are called *efficient* if they maximize $\sum_{i=1}^n \beta_{g_i(b)} v_i$, where we use the convention that $\beta_j = 0$ if $j > k$. An efficient allocation can be obtained by assigning positions in non-increasing order of agents' valuations. We will be concerned exclusively with mechanisms that maximize welfare with respect to the bids, and henceforth denote by g an efficient allocation rule that breaks ties in an arbitrary but consistent manner. For a given vector $v \in \mathbb{R}^n$ of values or bids and $i \in \{1, \dots, n\}$, we write $v_{(i)}$ for the $(n - i + 1)$ st order statistic of v , such that $v_{(1)} \geq \dots \geq v_{(n)}$, and use the convention that $v_{(i)} = 0$ if $i > n$.

In reasoning about strategic behavior we make the usual assumption of quasi-linear preferences and consider two different models of information regarding the preferences of other agents. Under quasi-linear preferences, the *utility* $u_i(b, v_i)$ of agent i with value v_i , in a given mechanism and for a given bid profile b , is equal to its valuation for the position it is assigned minus its payment, i.e., $u_i(b, v_i) = \beta_{g_i(b)} v_i - p_i(b)$. In the *complete information* model the values v_i are common knowledge among the agents. A bid profile b is a *Nash equilibrium* of a given mechanism 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 $i \in N$,

$$u_i(b, v_i) = \max_{x \in \mathbb{R}} u_i((b_{-i}, x), v_i),$$

where $(b_{-i}, x) = (b_1, \dots, b_{i-1}, x, b_{i+1}, \dots, b_n)$.

In the *incomplete information* model values v_i are drawn independently from a continuous distribution with density function f , cumulative distribution function F , and support $[0, \bar{v}]$ for some finite $\bar{v} \in \mathbb{R}_+$ we assume to be common knowledge among the agents.⁵ Our results in addition require existence and boundedness of the first three derivatives of F . Since valuations are independent and identically distributed, an efficient allocation for all value profiles can only be obtained from a symmetric profile b of bidding functions, i.e., one where $b = (b, \dots, b)$ for some bidding function $b : \mathbb{R} \rightarrow \mathbb{R}$. The quantity of interest for strategic considerations under incomplete information is the *expected utility* $u_i^b(x, v_i)$ of agent i with value v_i given that the agent bids $x \in \mathbb{R}$ and all other agents use bidding function b , which is given by

$$u_i^b(x, v_i) = \mathbb{E}_{v_j \sim F, j \neq i} \left[u_i \left(v_i, (b(v_1), \dots, b(v_{i-1}), x, b(v_{i+1}), \dots, b(v_n)) \right) \right].$$

⁵An analytical characterization of equilibria in the case of non-identical distributions is, unfortunately, well beyond the state of the art even for very simple settings. For a single item and two bidders with values drawn uniformly from distinct intervals, for example, this question was posed by Vickrey [35] and answered only recently, almost half a century later, by Kaplan and Zamir [22].

A symmetric profile of bidding functions b then is a *Bayes-Nash equilibrium* if no agent has an incentive to change its bid, i.e., if for all $i \in \{1, \dots, n\}$ and $v_i \in [0, \bar{v}]$,

$$u_i^b(b(v_i), v_i) = \max_{x \in \mathbb{R}} u_i^b(x, v_i).$$

A mechanism that obtains an efficient allocation in both Nash and Bayes-Nash equilibrium is the so-called Vickrey-Clarke-Groves (VCG) mechanism. It uses the efficient allocation rule g and a payment rule p^β that charges each agent its externality on the other agents, which is equal to the additional utility agents assigned lower slots would obtain by moving up one slot, i.e.,

$$p_i^\beta(b) = \sum_{j=g_i(b)}^k (\beta_j - \beta_{j+1}) b_{(j+1)},$$

It is well known that the VCG mechanism makes it optimal for each agent to bid its true valuation, and we refer to the resulting allocation and payments for a given valuation profile as the truthful VCG outcome for that profile.

While computation of payments in the VCG mechanism requires knowledge of β , we will be interested instead in the ability of mechanisms to support the truthful VCG outcome in equilibrium when only an inaccurate estimate α of the vector β of relative values is available to the designer. To this end we consider parameterized variants of three mechanisms that have been used in practice and therefore studied quite extensively: the α -VCG mechanism, the α -GFP mechanism, and the α -GSP mechanism. The three mechanisms all use allocation rule g , and their payment rules p^V , p^F , and p^S respectively charge an agent its externality, its bid on the position it is assigned, and the next-lower bid on that position, i.e.,

$$\begin{aligned} p_i^V(b) &= \sum_{j=g_i(b)}^k (\alpha_j - \alpha_{j+1}) b_{(j+1)}, \\ p_i^F(b) &= \alpha_{g_i(b)} b_i, \quad \text{and} \\ p_i^S(b) &= \alpha_{g_i(b)} b_{(g_i(b)+1)}. \end{aligned}$$

We will sometimes drop superscripts when the mechanism we are referring to is clear from the context.

3 Complete Information

We begin our analysis with the complete-information case. Here, when $\alpha = \beta$, the α -VCG mechanism has a truthful equilibrium, the α -GSP mechanism has an equilibrium that yields the truthful VCG outcome [16, 33], and the α -GFP mechanism may not have any equilibrium [9]. When $\alpha \neq \beta$ the α -VCG mechanism loses its truthfulness, and it makes sense to ask under what conditions the α -VCG mechanism and the α -GSP mechanism preserve the truthful VCG outcome in equilibrium. To build intuition we first look at the special case with three positions and three agents, before moving on to the general case.

3.1 Three Positions and Three Agents

In the special case, valuations are given by vectors $v \in \mathbb{R}^3$ and $\beta \in \mathbb{R}_{\geq}^3$ while mechanisms use a vector $\alpha \in \mathbb{R}_{\geq}^3$ that may differ from β . Without loss of generality we can assume that $v_1 \geq v_2 \geq v_3$ and that $\alpha_1 = \beta_1 = 1$. Our goal will be to understand which combinations of $\alpha = (1, \alpha_2, \alpha_3)$ and $\beta = (1, \beta_2, \beta_3)$ allow for the existence of a bid vector $b = (b_1, b_2, b_3)$ that leads to an efficient assignment and truthful VCG payments for all positions. Efficiency requires that $b_1 \geq b_2 \geq b_3$ and truthful VCG payments are given by

$$\begin{aligned} p_1^\beta(v) &= (\beta_1 - \beta_2)v_2 + (\beta_2 - \beta_3)v_3, \\ p_2^\beta(v) &= (\beta_2 - \beta_3)v_3, \quad \text{and} \\ p_3^\beta(v) &= 0. \end{aligned}$$

In the α -GSP mechanism, the agent assigned position i pays $\alpha_i b_{(i+1)}$ when $i \in \{1, 2\}$ and zero when $i = 3$. We thus obtain the truthful VCG payments if $\alpha_1 b_{(2)} = p_1^\beta(v)$ and $\alpha_2 b_{(3)} = p_2^\beta(v)$, i.e., if $b_{(2)} = p_1^\beta(v)/\alpha_1$ and $b_{(3)} = p_2^\beta(v)/\alpha_2$. Together with efficiency this yields the following necessary condition, which in fact is sufficient as well:

$$b_1 \geq b_2 = \frac{p_1^\beta(v)}{\alpha_1} \geq b_3 = \frac{p_2^\beta(v)}{\alpha_2}.$$

In the α -VCG mechanism the payment of the agent assigned position i is $(\alpha_1 - \alpha_2)b_{(2)} + (\alpha_2 - \alpha_3)b_{(3)}$ if $i = 1$, $(\alpha_2 - \alpha_3)b_{(3)}$ if $i = 2$, and zero if $i = 3$, which is equal to the truthful VCG payments if $b_{(2)} = (p_1^\beta(v) - p_2^\beta(v))/(\alpha_1 - \alpha_2)$ and $b_{(3)} = p_2^\beta(v)/(\alpha_2 - \alpha_3)$. Together with efficiency we obtain the following necessary condition, which again is also sufficient:

$$b_1 \geq b_2 = \frac{p_1^\beta(v) - p_2^\beta(v)}{\alpha_1 - \alpha_2} \geq b_3 = \frac{p_2^\beta(v)}{\alpha_2 - \alpha_3}.$$

Crucial for both mechanisms is the second inequality, concerning the relationship between b_2 and b_3 , and by solving it for α_2 we obtain

$$\alpha_2 \geq \alpha_1 \frac{p_2^\beta(v)}{p_1^\beta(v)} \quad \text{and} \quad \alpha_2 \geq (\alpha_1 - \alpha_3) \frac{p_2^\beta(v)}{p_1^\beta(v)} + \alpha_3$$

for the α -GSP and the α -VCG mechanism, respectively. On the other hand $p_1^\beta(v) \geq p_2^\beta(v)$ and thus $p_2^\beta(v)/p_1^\beta(v) \leq 1$. This means that

$$\alpha_1 \frac{p_2^\beta(v)}{p_1^\beta(v)} \leq (\alpha_1 - \alpha_3) \frac{p_2^\beta(v)}{p_1^\beta(v)} + \alpha_3,$$

making the condition for the α -GSP mechanism easier to satisfy. If $p_2^\beta(v)/p_1^\beta(v) < 1$ the inequality will in fact be strict, and will produce examples where the α -GSP mechanism possesses an equilibrium of the desired type but the α -VCG mechanism does not. For concreteness let $\alpha = (1, 0.6, 0.3)$, $\beta = (1, 0.7, 0.3)$, and $v = (20, 10, 10)$. Then $p_1^\beta(v) = 7$, $p_2^\beta(v) = 4$, and $p_3^\beta(v) = 0$, and thus $\alpha_2 = 0.6 \geq \alpha_1 \cdot p_2^\beta(v)/p_1^\beta(v) = 4/7$ while $\alpha_2 = 0.6 < (\alpha_1 - \alpha_3) \cdot p_2^\beta(v)/p_1^\beta(v) + \alpha_3 = 0.7$. More generally, and holding everything else fixed, any value of $\alpha_2 \geq 4/7 \approx 0.57$ would suffice for the α -GSP

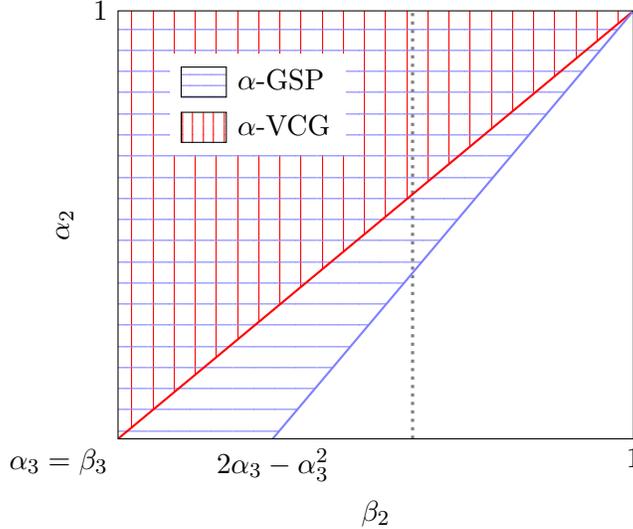


Figure 1: Comparison of the α -GSP and α -VCG mechanisms under complete information, for a setting with three positions, three agents, and valuations $v_1 \geq v_2 = v_3$. The hatched areas indicate the combinations of α_2 and β_2 for which the mechanisms respectively obtain the truthful VCG outcome in equilibrium, when $\alpha_1 = \beta_1 = 1$ and $\alpha_3 = \beta_3$. The dotted line illustrates the performance of the mechanisms for a particular value of β_2 . When $\alpha_3 = \beta_3 = 0.3$ it would lie at $\beta_2 = 0.7$ and would intersect the curve for the α -GSP mechanism at $\alpha_2 = 4/7$ and that for the α -VCG mechanism at $\alpha_2 = 0.7$. Any point on the dotted line between these two intersection points corresponds to a value of α_2 for which the α -GSP mechanism has an equilibrium of the desired type and the α -VCG mechanism does not.

mechanism, whereas the α -VCG mechanism would require that $\alpha_2 \geq 0.7$. Figure 1 illustrates the relationship between the inequalities in a parametrized manner, holding $\alpha_1 = \beta_1 = 1$ and $\alpha_3 = \beta_3$ fixed and letting α_2 and β_2 range over $[\alpha_3, 1]$. It shows that the α -VCG mechanism fails to preserve the truthful VCG outcome even if it only slightly underestimates β_2 , while the α -GSP mechanism is relatively robust in this regard.

3.2 The General Case

Perhaps surprisingly, superiority of the α -GSP mechanism over the α -VCG mechanism in preserving the truthful VCG outcome also holds in general. The following result establishes a weak superiority for arbitrary numbers of agents and positions and arbitrary valuations. Examples in which only the α -GSP mechanism preserves the truthful VCG outcome are straightforward to construct, and indeed we have already done so for a specific setting.

Theorem 1. *Let $\alpha, \beta \in \mathbb{R}_{\geq}^k$, $v \in \mathbb{R}^n$. Then the α -GSP mechanism obtains the truthful VCG outcome in a Nash equilibrium for valuations given by β and v whenever the α -VCG mechanism does.*

We show this result by characterizing the combinations of α and β for which the α -GSP and α -VCG mechanisms respectively obtain the truthful VCG outcome in equilibrium. To this end

consider a situation with k positions and n agents, and assume without loss of generality that $n \geq k$ and that agents are indexed such that $v_1 \geq \dots \geq v_n$. For $j = 1, \dots, k$ let $p_j^\beta(v) = \sum_{i=j}^k (\beta_i - \beta_{i+1}) \cdot v_{i+1}$ denote the truthful VCG price of position j , and for notational convenience set $p_i^\beta(v) = 0$ when $i \geq k + 1$.

The characterization is provided in terms of two lemmas. Under both mechanisms the bid b_1 of the agent with the highest value has to be a maximum bid but is otherwise unconstrained. The bids b_2, \dots, b_k of the agents with the 2nd- to k th-highest values are completely determined, as is the highest bid $b_{(k+1)}$ of any agent not assigned a position, if such an agent exists. They are given by functions depending on α and on the truthful VCG payments $p^\beta(v)$ and have to form a non-increasing sequence. The formal statements of the lemmas also take care of some corner cases, like those where $\alpha_j = 0$ in the case of the α -GSP mechanism or $\alpha_{j-1} - \alpha_j = 0$ in the case of the α -VCG mechanism, and are therefore slightly more complicated. Detailed proofs of these and all other lemmas are given in the appendix.

Lemma 1. *Let $\alpha, \beta \in \mathbb{R}_{\geq}^k$, $v \in \mathbb{R}^n$ such that $v_1 \geq v_2 \geq \dots \geq v_n$. Then the following are equivalent:*

- (a) *Bid vector $b \in \mathbb{R}^n$ yields an efficient equilibrium of the α -GSP mechanism in which agent i pays $p_i^\beta(v)$.*
- (b) *The bids satisfy $b_1 \geq b_2 \geq \dots \geq b_k \geq b_{(k+1)}$, and for $j = 1, \dots, k$ either $\alpha_j = p_j^\beta(v) = 0$ or $b_{(j+1)} = p_j^\beta(v)/\alpha_j$.*

Lemma 2. *Let $\alpha, \beta \in \mathbb{R}_{\geq}^k$, $v \in \mathbb{R}^n$ such that $v_1 \geq v_2 \geq \dots \geq v_n$. Then the following are equivalent:*

- (a) *Bid vector $b \in \mathbb{R}^n$ yields an efficient equilibrium of the α -VCG mechanism in which agent i pays $p_i^\beta(v)$.*
- (b) *The bids satisfy $b_1 \geq b_2 \geq \dots \geq b_k \geq b_{(k+1)}$, and for $j = 1, \dots, k$ either $\alpha_j = \alpha_{j+1}$ and $p_j^\beta(v) = p_{j+1}^\beta(v)$ or $b_{(j+1)} = \frac{p_j^\beta(v) - p_{j+1}^\beta(v)}{\alpha_j - \alpha_{j+1}}$.*

The proof of Theorem 1 now exploits the recursive structure of VCG payments to show that the characterization for the α -VCG mechanism is more demanding.

Proof of Theorem 1. Assume for contradiction that there exist $\alpha, \beta \in \mathbb{R}_{\geq}^k$ such that the α -GSP mechanism does not preserve the truthful VCG outcome on β while the α -VCG mechanism does. By Lemma 1, there are only two possible reasons for the failure of the α -GSP mechanism to support the truthful VCG outcome. Either there exists a position $i \leq k$ such that $\alpha_i = 0$ and $p_i^\beta(v) > 0$, or a position $i < k$ with $\alpha_i \geq \alpha_{i+1} > 0$ such that

$$\frac{p_i^\beta(v)}{\alpha_i} < \frac{p_{i+1}^\beta(v)}{\alpha_{i+1}}.$$

In both cases we can assume without loss of generality that i is the largest such position.

In the former case, for any bid vector $b \in \mathbb{R}^n$, $p_i^\alpha(b) = (\alpha_i - \alpha_{i+1})b_{(i+1)} + p_{i+1}^\alpha(b) = 0$, which clearly prevents the α -VCG mechanism from supporting the truthful VCG outcome.

In the latter case, we can make two simple observations. First $\alpha_i > \alpha_{i+1}$, as $\alpha_i = \alpha_{i+1}$ would imply $p_i^\beta(v) < p_{i+1}^\beta(v)$, which is clearly impossible. Second there must exist a position j with

$i < j \leq k$ such that $\alpha_j \neq \alpha_{j+1}$, because $i < k$, $\alpha_{i+1} > 0$, and $\alpha_{k+1} = 0$. Without loss of generality we can assume that j is the smallest such position. Now, by writing $p_i^\beta(v)$ as $p_i^\beta(v) - p_{i+1}^\beta(v) + p_{i+1}^\beta(v)$,

$$\frac{p_i^\beta(v) - p_{i+1}^\beta(v) + p_{i+1}^\beta(v)}{\alpha_i} < \frac{p_{i+1}^\beta(v)}{\alpha_{i+1}},$$

and by rearranging,

$$\frac{p_i^\beta(v) - p_{i+1}^\beta(v)}{\alpha_i} < \frac{p_{i+1}^\beta(v)}{\alpha_{i+1}} - \frac{p_{i+1}^\beta(v)}{\alpha_i}.$$

Since $\alpha_i > \alpha_{i+1}$, we can multiply both sides by $\alpha_i/(\alpha_i - \alpha_{i+1})$ to obtain

$$\frac{p_i^\beta(v) - p_{i+1}^\beta(v)}{\alpha_i - \alpha_{i+1}} < \frac{p_{i+1}^\beta(v)}{\alpha_{i+1}}. \quad (1)$$

By writing $p_j^\beta(v)$ as $p_j^\beta(v) - p_{j+1}^\beta(v) + p_{j+1}^\beta(v)$, and using that $p_{j+1}^\beta(v) \leq \frac{\alpha_{j+1}}{\alpha_j} p_j^\beta(v)$ by choice of i ,

$$p_j^\beta(v) \leq p_j^\beta(v) - p_{j+1}^\beta(v) + \frac{\alpha_{j+1}}{\alpha_j} p_j^\beta(v),$$

and by rearranging,

$$p_j^\beta(v) - \frac{\alpha_{j+1}}{\alpha_j} p_j^\beta(v) \leq p_j^\beta(v) - p_{j+1}^\beta(v).$$

Since $\alpha_j > \alpha_{j+1}$, we can divide both sides by $\alpha_j - \alpha_{j+1}$ to obtain

$$\frac{p_j^\beta(v)}{\alpha_j} \leq \frac{p_j^\beta(v) - p_{j+1}^\beta(v)}{\alpha_j - \alpha_{j+1}}. \quad (2)$$

By choice of j , for any position m with $i < m < j$ we have that $\alpha_s = \alpha_{s+1}$. By the assumption that the α -VCG mechanism preserves the truthful VCG outcome and by Lemma 2, this implies that $p_s^\beta(v) = p_{s+1}^\beta(v)$. If no such m exists then $j = i + 1$, and in both cases,

$$\frac{p_{i+1}^\beta(v)}{\alpha_{i+1}} = \frac{p_j^\beta(v)}{\alpha_j} \quad (3)$$

Combining (1), (2), and (3),

$$\frac{p_i^\beta(v) - p_{i+1}^\beta(v)}{\alpha_i - \alpha_{i+1}} < \frac{p_{i+1}^\beta(v)}{\alpha_{i+1}} = \frac{p_j^\beta(v)}{\alpha_j} \leq \frac{p_j^\beta(v) - p_{j+1}^\beta(v)}{\alpha_j - \alpha_{j+1}}. \quad (4)$$

On the other hand, since $\alpha_i > \alpha_{i+1}$ and $\alpha_j > \alpha_{j+1}$, we know from Lemma 2 that any bid vector b' that yields the truthful VCG outcome must set b'_{i+1} to the left-hand side of (4) and b'_{j+1} to the right-hand side of (4). Since $i < j$ it must also hold that $b'_{i+1} \geq b'_{j+1}$, which is a contradiction. \square

4 Incomplete Information

We now turn to incomplete-information environments, where agents only possess probabilistic information regarding one another's valuations. Here the α -GSP mechanism may fail to possess an efficient equilibrium even when $\alpha = \beta$.⁶ When $\alpha = \beta$, the α -VCG mechanism of course maintains its truthful dominant-strategy equilibrium. Another good mechanism in this case is the α -GFP mechanism, which differs from the α -GSP mechanism in its use of first-price rather than second-price payments. While sharing the latter's non-truthfulness it possesses a unique Bayes-Nash equilibrium for any value of α , and this equilibrium yields the truthful VCG outcome [9].

Given these results it is quite natural to ask how successful the α -VCG and α -GFP mechanisms are in maintaining the truthful VCG outcome when $\alpha \neq \beta$. The answer to this question is strikingly similar to the complete-information case in that the non-truthful mechanism is again more robust, for arbitrary values of α and β and independent and identically distributed valuations according to any distribution satisfying mild technical conditions. Our analysis uses Myerson's classical characterization of possible equilibrium bids to identify, for either of the two mechanisms, conditions on α and β that are necessary and sufficient for equilibrium existence. The conditions for the α -VCG mechanism turn out to be more demanding. Just as we did for complete-information environments, we begin by considering a special case, this time with two positions, three agents, and valuations drawn uniformly at random from the unit interval. The special case is used to build intuition, and introduce the necessary machinery, for the general result.

4.1 Two Positions and Three Agents

Let v_1, v_2, v_3 be drawn independently from the uniform distribution on $[0, 1]$. Let $\alpha, \beta \in \mathbb{R}_{\geq}^2$ with $\alpha_2, \beta_2 > 0$, and assume without loss of generality that $\alpha_1 = \beta_1 = 1$. Our goal will again be to characterize the values of α and β for which given mechanisms of interest, in this case the α -GFP and α -VCG mechanisms, admit an efficient equilibrium. Behavior under incomplete information can be described by a vector of bidding functions, one for each agent, that map the agent's value to its bid. It is clear that in a symmetric setting like ours efficient outcomes can only result from symmetric bidding functions, so we will be interested in functions $b^F : \mathbb{R} \rightarrow \mathbb{R}$ that yield an efficient equilibrium in the α -GFP mechanism and functions $b^V : \mathbb{R} \rightarrow \mathbb{R}$ that achieve the same in the α -VCG mechanism.

The standard technique for equilibrium analysis under incomplete information uses a seminal result of Myerson that characterizes the expected allocation and payments in equilibrium in terms of the allocation probabilities induced by a mechanism and agents' bidding functions. The result was originally formulated for truthful mechanisms, but equivalent conditions exist for arbitrary bidding functions that instead of being in equilibrium provide a best response among values in their range. The latter is obviously a necessary condition for equilibrium, and can be turned into a sufficient condition by arguing that no better response exists outside the range. For our setting and notation we have the following result.

Lemma 3 (Myerson [29]). *Consider a position auction for an environment with n agents, k positions, and $\beta \in \mathbb{R}_{\geq}^k$. Assume that agents use a bidding function b with range X , and that an*

⁶Gomes and Sweeney [18] gave a characterization of those values of α that enable equilibrium existence in this case. The result can be strengthened in our setting to show that for some values of β no choice of α leads to an efficient equilibrium.

agent with value v is consequently assigned position $s \in \{1, \dots, k\}$ with probability $P_s(v)$. Then $u(b(v), v) = \max_{x \in X} u(x, v)$ for all $v \in [0, \bar{v}]$ if and only if the following holds:

- (a) the expected allocation $\sum_{s=1}^k P_s(v)\beta_s$ is non-decreasing in v , and
- (b) the payment function p satisfies

$$\mathbb{E}[p(v)] = p(0) + \sum_{s=1}^k \beta_s \int_0^v \frac{dP_s(z)}{dz} z dz. \quad (5)$$

All mechanisms we consider set $p(0) = 0$ and use an efficient allocation rule, for which

$$P_s(v) = \binom{n-1}{s-1} (1-F(v))^{s-1} (F(v))^{n-s}$$

and (a) is satisfied. Together with our assumptions on F , efficiency mandates further that b must increase almost everywhere.

In the special case with two positions and three agents with values distributed uniformly on the unit interval we have that $P_1(v) = F^2(v) = v^2$ and $P_2(v) = \binom{2}{1} F(v)(1-F(v)) = 2v(1-v)$, payments in any efficient equilibrium can thus be described by a function $p^E : \mathbb{R} \rightarrow \mathbb{R}$ satisfying

$$\begin{aligned} \mathbb{E}[p^E(v)] &= \beta_1 \int_0^v \frac{dP_1(z)}{dz} z dz + \beta_2 \int_0^v \frac{dP_2(z)}{dz} z dz \\ &= \frac{2}{3} \beta_1 v^3 + \beta_2 v^2 - \frac{4}{3} \beta_2 v^3. \end{aligned} \quad (6)$$

A candidate equilibrium bidding function for the α -GFP mechanism can now be obtained by writing the expected payment in terms of bidding function b^F , equating the resulting expression with (6), and solving for b^F . In the α -GFP mechanism an agent with value v that is allocated position s pays $\alpha_s b^F(v)$, its expected payment therefore satisfies

$$\begin{aligned} \mathbb{E}[p^F(v)] &= P_1(v)\alpha_1 b^F(v) + P_2(v)\alpha_2 b^F(v) \\ &= (\alpha_1 v^2 + 2\alpha_2 v - 2\alpha_2 v^2) b^F(v). \end{aligned} \quad (7)$$

By Lemma 3 the expressions in (6) and (7) must be the same. Equating them yields

$$b^F(v) = \frac{2/3 \cdot v^3 - 4/3 \cdot \beta_2 v^2 + \beta_2 v^2}{v^2 - 2\alpha_2 v^2 + 2\alpha_2 v}$$

when $v > 0$, and we can set $b^F(0) = 0$ for convenience.⁷ Bidding below $b^F(0) = 0$ is impossible, bidding above $b^F(\bar{v})$ is dominated,⁸ and b^F satisfies the second condition of Lemma 3 by construction. The α -GFP mechanism thus has an efficient equilibrium if and only if b^F is increasing almost everywhere. Taking the derivative we obtain

$$\frac{db^F(v)}{dv} = \frac{(\frac{4}{3}v - \frac{8}{3}\beta_2 v + \beta_2)(v - 2\alpha_2 v + 2\alpha_2)}{(v - 2\alpha_2 v + 2\alpha_2)^2} - \frac{(1 - 2\alpha_2)(\frac{2}{3}v^2 - \frac{4}{3}\beta_2 v^2 + \beta_2 v)}{(v - 2\alpha_2 v + 2\alpha_2)^2}.$$

⁷Application of l'Hospital's rule shows that $\lim_{v \rightarrow 0} b^F(v) = 0$, so this choice makes b^F increasing.

⁸Since equilibrium bidding functions must be increasing almost everywhere, bidding above $b^F(\bar{v})$ would not increase the probability of winning, and it would also not lead to a lower payment.

The sign of this expression is determined by the sign of its numerator, and it turns out that the numerator is positive at 0 and, depending on the value of β_2 , either non-decreasing everywhere on $[0, 1]$ or decreasing everywhere on $[0, 1]$. Indeed, $db^F(v)/dv|_{v=0} = \beta_2/(2\alpha_2) > 0$, and the derivative of the numerator, $(4/3 - 8/3\beta_2)(v - 2\alpha_2v + 2\alpha_2)$, is non-negative when $\beta_2 \leq 1/2$ and negative when $\beta_2 > 1/2$. In the case where $\beta_2 > 1/2$ we need that

$$\left. \frac{db^F(v)}{dv} \right|_{v=1} = \left(\frac{4}{3} - \frac{5}{3}\beta_2 \right) - (1 - 2\alpha_2) \left(\frac{2}{3} - \frac{1}{3}\beta_2 \right) \geq 0,$$

which holds when

$$\alpha_2 \geq \frac{2\beta_2 - 1}{2 - \beta_2}.$$

We conclude that the α -GFP mechanism possesses an efficient equilibrium if and only if $\beta_2 \leq 1/2$ or $\alpha_2 \geq (2\beta_2 - 1)/(2 - \beta_2)$.

Analogously, in the α -VCG mechanism, the payment of an agent with value v satisfies

$$\begin{aligned} \mathbb{E}[p^V(v)] &= P_1(v) \left[(\alpha_1 - \alpha_2) \int_0^v \frac{2t}{v^2} b^V(t) dt + \alpha_2 \int_0^v \frac{2(v-t)}{v^2} b^V(t) dt \right] + P_2(v) \alpha_2 \int_0^v \frac{1}{v} b^V(t) dt \\ &= (2\alpha_1 - 4\alpha_2) \int_0^v t b^V(t) dt + 2\alpha_2 \int_0^v b^V(t) dt, \end{aligned} \quad (8)$$

where $2t/v^2 = 2F(t)f(t)/F(v)^2$ and $2(v-t)/v^2 = 2F(v-t)f(t)/F(v)^2$ are the densities of the second and third highest values given that the agent's value v is the highest, and $1/v = f(t)/F(v)$ is the density of the third highest value given that v is the second highest. By Lemma 3 the expressions in (6) and (8) must again be the same. Taking the derivatives of both and solving for $b^V(v)$ yields

$$b^V(v) = \frac{2v^2 - 4\beta_2v^2 + 2\beta_2v}{2v - 4\alpha_2v + 2\alpha_2}$$

when $v < 1$, and we can extend b^V appropriately when $v = 1$.⁹ By the same argument as before, the α -VCG mechanism has an efficient equilibrium if and only if b^V is increasing almost everywhere. Taking the derivative we obtain

$$\frac{db^V(v)}{dv} = \frac{(4v - 8\beta_2v + 2\beta_2)(2v - 4\alpha_2v + 2\alpha_2)}{(2v - 4\alpha_2v + 2\alpha_2)^2} - \frac{(2 - 4\alpha_2)(2v^2 - 4\beta_2v^2 + 2\beta_2v)}{(2v - 4\alpha_2v + 2\alpha_2)^2}.$$

When $\alpha_2 < 1$ the sign of this expression is determined by its numerator, which is positive at 0 and, depending on the value of β_2 , either non-decreasing everywhere on $[0, 1]$ or decreasing everywhere on $[0, 1]$. Indeed, $db^F(v)/dv|_{v=0} = \beta_2/\alpha_2 > 0$, and the derivative of the numerator, $(4 - 8\beta_2)(2v - 4\alpha_2v + 2\alpha_2)$, is non-negative when $\beta_2 \leq 1/2$ and negative when $\beta_2 > 1/2$. When $\beta_2 > 1/2$ we need that

$$\left. \frac{db^V(v)}{dv} \right|_{v=1} = \frac{(4 - 6\beta_2)(2 - 2\alpha_2) - (2 - 4\alpha_2)(2 - 2\beta_2)}{(2 - 2\alpha_2)^2} \geq 0,$$

⁹We have assumed that $\alpha_2 > 0$, so the denominator vanishes only when $v = \alpha_2 = 1$. If $\beta_2 < 1$, then $\lim_{v \rightarrow 1} b^V(v) = \infty$. If $\beta_2 = 1$, application of l'Hospital's rule shows that $\lim_{v \rightarrow 1} b^V(v) = 1$.

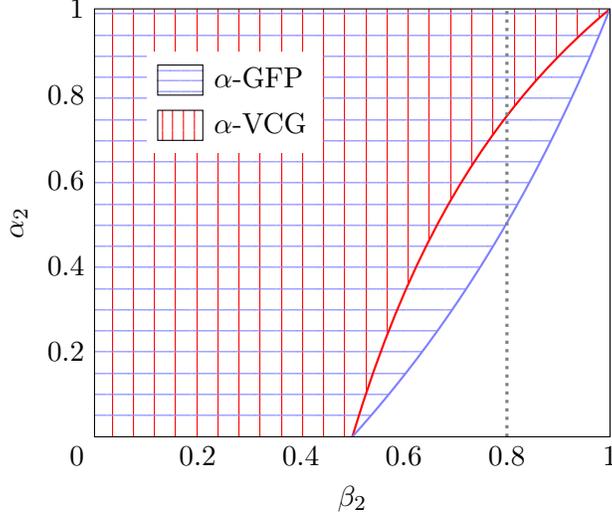


Figure 2: Comparison of the α -GFP and α -VCG mechanisms under incomplete information, for a setting with two positions, three agents, and valuations drawn independently and uniformly from $[0, 1]$. The hatched areas indicate the combinations of α_2 and β_2 for which the mechanisms obtain the truthful VCG outcome in equilibrium, when $\alpha_1 = \alpha_2 = 1$. The dotted line at $\beta_2 = 0.8$ intersects the curve for the α -GFP mechanism at $\alpha_2 = 0.5$ and that for the α -VCG mechanism at $\alpha_2 = 0.75$. For all points between the intersection points the α -GFP mechanism has an equilibrium of the desired type and the α -VCG mechanism does not.

which for $\alpha_2 < 1$ holds when

$$\alpha_2 \geq 2 - \frac{1}{\beta_2}.$$

When $\alpha_2 = 1$ the above reasoning still applies as long as $v < 1$, so $b^V(v)$ is increasing almost everywhere when

$$\lim_{v \rightarrow 1} \frac{db^V(v)}{dv} \geq 0.$$

This is indeed the case, as $\lim_{v \rightarrow 1} db^V(v)/dv = \infty$ when $\beta_2 < 1$, and $\lim_{v \rightarrow 1} db^V(v)/dv = 1$ when $\beta_2 = 1$ by applying l'Hospital's rule twice. We conclude that the α -VCG mechanism possesses an efficient equilibrium if and only if $\beta_2 \leq 1/2$ or $\alpha_2 \geq 2 - 1/\beta_2$.

It is now not hard to see that the equilibrium condition for the α -GFP mechanism is easier to satisfy than that for the α -VCG mechanism. In fact, the α -VCG mechanism may fail to preserve the truthful VCG outcome even when α_2 is very close to β_2 . When $\beta_2 = 0.8$, for example, any value of $\alpha_2 \geq 0.5$ would suffice for the α -GFP mechanism, while the α -VCG mechanism would require that $\alpha_2 \geq 0.75$. An illustration is provided in Figure 2. Figure 3 compares the derivatives of the respective bidding functions for $\beta_2 = 0.8$ and varying values of α_2 .

4.2 The General Case

We proceed to show that the α -GFP mechanism is superior in general to the α -VCG mechanism in preserving the truthful VCG outcome. The following result establishes a weak superiority for

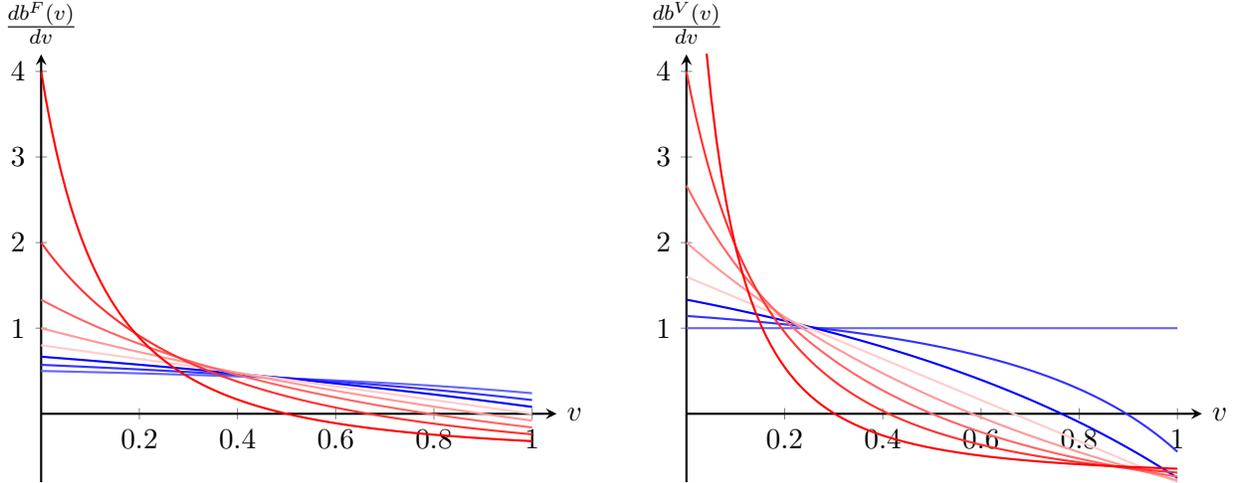


Figure 3: Derivatives of the candidate bidding functions for the α -GFP and α -VCG mechanisms in a setting with three agents with values distributed uniformly on $[0, 1]$ and two positions with $\beta = (1, 0.8)$, when $\alpha_1 = 1$ and α_2 ranges from 0.8 to 0.1.

any number of positions and agents, and arbitrary symmetric valuation distributions. Examples in which only the α -GFP mechanism preserves the truthful VCG outcome are straightforward to construct, and indeed we have already done so for a specific setting.

Theorem 2. *Let $\alpha, \beta \in \mathbb{R}_{\geq}^n$. Let $v \in \mathbb{R}^n$, with components drawn independently from a continuous distribution with bounded support. Assume that $n > k$. Then the α -GFP mechanism obtains the truthful VCG outcome in a Bayes-Nash equilibrium for valuations given by β and v whenever the α -VCG mechanism does.*

To obtain this general result we will follow the same basic strategy as in the special case, but will have to overcome two major difficulties on the way.

The first difficulty concerns the equilibrium bidding function for the α -VCG mechanism. Whereas deriving a bidding function for the α -GFP mechanism remains relatively straightforward even for an arbitrary number of positions and arbitrary valuation distributions, the situation becomes significantly more complex for the α -VCG mechanism due to the dependence of its payment rule on the bids for all lower positions. Specifically, when equating the two expressions for the expected payment in equilibrium, (6) and (8) in the special case, and taking derivatives on both sides, the integrand in the latter no longer depends only on t , the variable of integration. Instead, the conditional densities of the values of agents assigned lower positions introduce a dependence on v . When taking the derivative one would expect to obtain a differential equation, and a closed form solution to this differential equation would be required to continue with the rest of the argument. We take a different route and use a rather surprising combinatorial equivalence to obtain an alternative expression for the expected payment that only depends on t .

A second difficulty arises when trying to show that b^F is increasing for a wider range of values of α and β than b^V . In the special case we could argue directly about the derivatives of the bidding functions, but this type of argument becomes infeasible rather quickly when increasing the number of positions or moving to general value distributions. The key insight that will allow us to generalize

the result is that there exist functions $A : \mathbb{R} \rightarrow \mathbb{R}$ and $B : \mathbb{R} \rightarrow \mathbb{R}$ such that $b^F(v) = A(v)/B(v)$ and $b^V(v) = A'(v)/B'(v)$, where A' and B' respectively denote the derivatives of A and B . This relationship is easily verified for the bidding functions in (6) and (8) but continues to hold in general. We use it to show that at the minimum value of v for which $db^F(v)/dv$ is non-positive, should such a value exist, $db^V(v)/dv$ is non-positive as well.

We begin by deriving candidate equilibrium bidding functions for the two mechanisms. Due to the more complicated structure of the payments, the case of the α -VCG mechanism is significantly more challenging.

Lemma 4. *Let $\alpha, \beta \in \mathbb{R}_{\geq}^n$, where $\alpha_j > 0$ and $\beta_j > 0$ for $j \in \{1, \dots, k\}$. Suppose valuations are drawn from a distribution with support $[0, \bar{v}]$, probability density function f , and cumulative distribution function F . Then, an efficient equilibrium of the α -GFP mechanism must use a bidding function b^F with*

$$b^F(v) = \frac{\sum_{s=1}^k \beta_s \int_0^v \frac{dP_s(t)}{dt} t dt}{\sum_{s=1}^k \alpha_s P_s(v)}.$$

If b^F is increasing almost everywhere, it constitutes the unique efficient equilibrium. Otherwise no efficient equilibrium exists.

Lemma 5. *Let $\alpha, \beta \in \mathbb{R}_{\geq}^n$, where $\alpha_j > 0$ and $\beta_j > 0$ for $j \in \{1, \dots, k\}$. Suppose valuations are drawn from a distribution with support $[0, \bar{v}]$, probability density function f , and cumulative distribution function F . Then, an efficient equilibrium of the α -VCG mechanism must use a bidding function b^V with*

$$b^V(v) = \frac{\sum_{s=1}^k \beta_s \frac{dP_s(v)}{dv} v}{\sum_{s=1}^k \alpha_s \frac{dP_s(v)}{dv}}.$$

If b^V is increasing almost everywhere, it constitutes the unique efficient equilibrium. Otherwise no efficient equilibrium exists.

Even with the candidate bidding functions b^F and b^V in hand, the cases where the α -GFP and α -VCG mechanisms respectively admit an efficient equilibrium seem difficult to compare. What will ultimately drive the proof of Theorem 2 is a rather curious relationship between the two bidding functions that is straightforward to verify given Lemma 4 and Lemma 5: the numerator of b^V is equal to the derivative of the numerator of b^F , and the denominator of b^V is equal to the derivative of the denominator of b^F .

Corollary 1. *Let $b^F : \mathbb{R} \rightarrow \mathbb{R}$ and $b^V : \mathbb{R} \rightarrow \mathbb{R}$ be the candidate equilibrium bidding functions for the α -GFP and α -VCG mechanisms as defined in Lemma 4 and Lemma 5. Then*

$$b^F(v) = \frac{A(v)}{B(v)} \quad \text{and} \quad b^V(v) = \frac{A'(v)}{B'(v)},$$

where $A(v) = \sum_{s=1}^k \beta_s \int_0^v \frac{dP_s(t)}{dt} t dt$ and $B(v) = \sum_{s=1}^k \alpha_s P_s(v)$.

To show that the α -GFP mechanism possesses an efficient equilibrium whenever the α -VCG mechanism does we recall that equilibrium existence is equivalent to a bidding function that is increasing almost everywhere. We first consider the candidate bidding function for the α -GFP mechanism and show that at $v = 0$, either its derivative is positive or both its derivative and

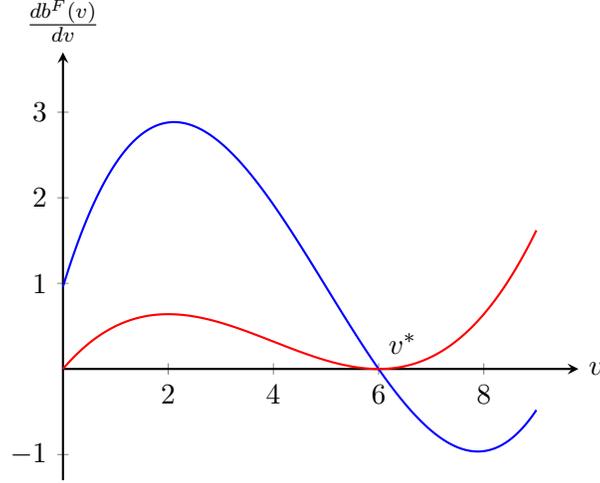


Figure 4: Possible forms of the derivative of the candidate equilibrium bidding function b^F when the α -GFP mechanism does not possess an equilibrium. Both the derivative and the second derivative are non-negative at zero, so if the former is non-positive anywhere on $(0, \bar{v}]$ there must be a value $v^* > 0$ where it either touches or cuts the x -axis from above.

second derivative are non-negative. Failure to possess an equilibrium thus implies existence of a value $v^* > 0$ where the derivative cuts the x -axis from above, or of a set of such values with positive measure where it touches the x -axis. In a second step we then show that the candidate bidding function for the α -VCG mechanism behaves roughly in the same way at these values. The situation is illustrated in Figure 4.

Lemma 6. *Let $b^F : \mathbb{R} \rightarrow \mathbb{R}$ be the candidate equilibrium bidding function for the α -GFP mechanisms as defined in Lemma 4. Then,*

$$\left. \frac{db^F(v)}{dv} \right|_{v=0} = \frac{n-k}{n-k+1} \cdot \frac{\beta_k}{\alpha_k}.$$

Lemma 7. *Let $b^F : \mathbb{R} \rightarrow \mathbb{R}$ be the candidate equilibrium bidding function for the α -GFP mechanisms as defined in Lemma 4. Then, for $n = k$,*

$$\left. \frac{d^2b^F(v)}{dv^2} \right|_{v=0} \geq 0.$$

Proof of Theorem 2. Assume that the α -GFP mechanism does not possess an equilibrium, and recall that this implies the existence of a set of value with positive measure where the derivative of b^F is not strictly increasing. By Lemmas 6 and 7, there must thus exist a set of values $v^* > 0$ with positive measure where

$$\left. \frac{db^F(v)}{dv} \right|_{v=v^*} = 0 \quad \text{and} \quad \left. \frac{d^2b^F(v)}{dv^2} \right|_{v=v^*} \leq 0,$$

or one such value where the equality holds and the inequality is strict.

For an arbitrary value v ,

$$\frac{db^F(v)}{dv} = \frac{A'(v)B(v) - B'(v)A(v)}{(B(v))^2} = 0$$

requires that

$$A'(v)B(v) - B'(v)A(v) = 0. \quad (9)$$

Assuming (9),

$$\frac{d^2b^F(v)}{dv^2} = \frac{A''(v)B(v) - B''(v)A(v)}{(B(v))^2} \leq 0$$

requires that

$$A''(v)B(v) - B''(v)A(v) \leq 0, \quad (10)$$

Consider any $v^* > 0$, and observe that $A(v^*) > 0$ and $A'(v^*) > 0$. For $v = v^*$ we can thus rewrite (9) as $B(v^*) = \frac{B'(v^*)A(v^*)}{A'(v^*)}$, and substitute this into (10) to obtain

$$A''(v^*) \frac{B'(v^*)A(v^*)}{A'(v^*)} - B''(v^*)A(v^*) \leq 0.$$

Dividing by $A(v^*) > 0$ and multiplying by $A'(v^*) > 0$ yields

$$A''(v^*)B'(v^*) - A'(v^*)B''(v^*) \leq 0,$$

and thus

$$\left. \frac{b^V(v)}{dv} \right|_{v=v^*} = \frac{A''(v^*)B'(v^*) - A'(v^*)B''(v^*)}{(B'(v^*))^2} \leq 0.$$

It is, moreover, easily verified that the inequality holds strictly when $d^2b^F(v)/dv^2|_{v=v^*} < 0$. There thus exists a set of values v^* with positive measure where $\frac{b^V(v)}{dv} \leq 0$, and the claim follows. \square

A Proof of Lemma 1

First assume that (a) is satisfied. Efficiency requires that $b_1 \geq b_2 \geq \dots \geq b_k \geq b_{(k+1)}$, and implies that the payment of agent $j \in \{1, \dots, n\}$ in the α -GSP mechanism is equal to $\alpha_j b_{(j+1)}$. It follows that $\alpha_j b_{(j+1)} = p_j^\beta(v)$ and thus either $\alpha_j = 0$ and $p_j^\beta(v) = 0$ or $b_{(j+1)} = p_j^\beta(v)/\alpha_j$.

Now assume that (b) is satisfied, and observe that $b_1 \geq b_2 \geq \dots \geq b_k \geq b_{(k+1)}$ leads to an efficient assignment. Moreover, either $\alpha_j = p_j^\beta(v) = 0$ or $b_{(j+1)} = p_j^\beta(v)/\alpha_j$, so in both cases $\alpha_j b_{(j+1)} = p_j^\beta(v)$. It remains to be shown that b yields an equilibrium. To this end, assume for contradiction that some agent $i \in \{1, \dots, n\}$ would benefit strictly from changing its bid and being assigned position $i' \neq i$ as a consequence. The agent would then pay $\alpha_{i'} b_{(i')} \geq \alpha_{i'} b_{(i'+1)}$ if $i' < i$, and $\alpha_{i'} b_{(i'+1)}$ if $i' > i$. This is at least the payment of the agent assigned position i' in the efficient allocation, contradicting envy-freeness of the VCG payments [e.g., 24].

B Proof of Lemma 2

First assume that (a) is satisfied. Efficiency requires that $b_1 \geq b_2 \geq \dots \geq b_k \geq b_{(k+1)}$, and implies that the payment of agent k in the α -VCG mechanism is equal to $\alpha_k b_{(k+1)}$. If $\alpha_k = \alpha_{k+1} = 0$, this means that $p_k^\beta(v) = p_{k+1}^\beta(v) = 0$. Otherwise $\alpha_k - \alpha_{k+1} > 0$, and rearranging yields that $b_{(k+1)} = p_k^\beta(\theta)/\alpha_k$. For the remaining bids we can now argue inductively. Consider $s \in \{1, \dots, n\}$, and suppose that any agent $j \in \{s+1, \dots, k\}$ pays $p_j^\beta(v)$. If $\alpha_s \neq \alpha_{s+1}$, then agent s pays $(\alpha_s - \alpha_{s+1})b_{s+1} + p_{s+1}^\alpha(b) = (\alpha_s - \alpha_{s+1})b_{s+1} + p_{s+1}^\beta(v) = p_s^\beta(v)$, and thus $b_{s+1} = \frac{p_s^\beta(v) - p_{s+1}^\beta(v)}{\alpha_s - \alpha_{s+1}}$ as claimed. If instead $\alpha_s = \alpha_{s+1}$, then agents s and $s+1$ must pay the same and thus $p_s^\beta(v) = p_{s+1}^\beta(v)$.

Now assume that (b) is satisfied, and observe that $b_1 \geq b_2 \geq \dots \geq b_k \geq b_{(k+1)}$ leads to an efficient assignment. That payments are as required again follows by induction. For the base case observe that either $\alpha_k = \alpha_{k+1} = 0$ and agent k pays $0 = p_k^\beta(v)$, or $\alpha_k - \alpha_{k+1} > 0$ and agent k pays $\alpha_k b_{(k+1)} = (\alpha_k - \alpha_{k+1}) \frac{p_k^\beta(v)}{\alpha_k - \alpha_{k+1}} = p_k^\beta(v)$. For the inductive step assume that payments are as required up to position $s+1$. Then either $\alpha_s \neq \alpha_{s+1}$ and $b_{s+1} = \frac{p_s^\beta(v) - p_{s+1}^\beta(v)}{\alpha_s - \alpha_{s+1}}$ so that agent s pays

$$(\alpha_s - \alpha_{s+1})b_{s+1} + p_{s+1}^\beta(v) = (\alpha_s - \alpha_{s+1}) \frac{p_s^\beta(v) - p_{s+1}^\beta(v)}{\alpha_s - \alpha_{s+1}} + p_{s+1}^\beta(v) = p_s^\beta(v),$$

or $\alpha_s = \alpha_{s+1}$ and $p_s^\beta(v) = p_{s+1}^\beta(v)$ so that agent s pays

$$(\alpha_s - \alpha_{s+1})b_{s+1} + p_{s+1}^\beta(v) = p_{s+1}^\beta(v) = p_s^\beta(v).$$

It remains to be shown that b yields an equilibrium. To this end, assume for contradiction that some agent $i \in \{1, \dots, n\}$ would benefit strictly from changing its bid and being assigned position $i' \neq i$ as a consequence. The agent would then pay $\sum_{s=i'}^{i-1} (\alpha_s - \alpha_{s+1})b_{(s)} + p_i^\beta(v) \geq \sum_{s=i'}^{i-1} (\alpha_s - \alpha_{s+1})b_{(s+1)} + p_i^\beta(v)$ if $i' < i$, and $p_{i'}^\beta(v)$ if $i' > i$. This is at least the payment of the agent assigned position i' in the efficient allocation, contradicting envy-freeness of the VCG payments [e.g., 24].

C Proof of Lemma 4

Since efficient equilibria must be symmetric, we can write an efficient equilibrium of the α -GFP mechanism in terms of a bidding function $b^F : [0, \bar{v}] \rightarrow \mathbb{R}_{\geq 0}$. An agent with value v who is allocated position s then pays $\alpha_s b^F(v)$, and we have that

$$\mathbb{E} [p^F(v)] = \sum_{s=1}^k \alpha_s P_s(v) b^F(v). \quad (11)$$

The expected payment in an efficient equilibrium is given by Lemma 3, and by equating (11) with (5) and solving for $b^F(v)$ we obtain

$$b^F(v) = \frac{\sum_{s=1}^k \beta_s \int_0^v \frac{dP_s(t)}{dt} t dt}{\sum_{s=1}^k \alpha_s P_s(v)}.$$

Bidding below $b^F(0) = 0$ is impossible and bidding above $b^F(\bar{v})$ is dominated, so the claim follows from Lemma 3.

D Proof of Lemma 5

Efficiency again requires symmetry, so any efficient equilibrium of the α -VCG mechanism can be described by a bidding function $b^V : [0, \bar{v}] \rightarrow \mathbb{R}_{\geq 0}$.

Denote by $p^V(v)$ the payment in the α -VCG mechanism of an agent with value v , and by $p_s^V(v)$ the same payment under the condition that the agent has the s -highest value overall. These quantities are random variables that depend on the values of $n - 1$ other agents, and we have that

$$\mathbb{E}[p^V(v)] = \sum_{s=1}^k P_s(v) \cdot \mathbb{E}[p_s^V(v)], \quad (12)$$

where, as before, $P_s(v)$ is the probability that v is the s -highest of n values drawn independently from F . The conditional payment $p_s^V(v)$ depends on the conditional densities of the valuations of agents assigned lower positions, and on their resulting bids. For $s \in \{1, \dots, k\}$ and $\ell \in \{s, \dots, k\}$, denote by

$$I_{\ell,s}(v, t) = \frac{(n-s)f(t) \binom{n-s-1}{n-\ell-1} F(t)^{n-\ell-1} (F(v) - F(t))^{\ell-s}}{F(v)^{n-s}}$$

the density at t of the $(\ell + 1)$ -highest of n values drawn independently from F , under the condition that the s -highest value is equal to v . Then

$$\mathbb{E}[p_s^V(v)] = \sum_{\ell=s}^k (\alpha_\ell - \alpha_{\ell+1}) \cdot \int_0^v I_{\ell,s}(v, t) b^V(t) dt,$$

and by substituting into (12), exchanging the order of summation and integration, and grouping by coefficients of α_s , we obtain

$$\begin{aligned} \mathbb{E}[p^V(v)] &= \sum_{s=1}^k P_s(v) \sum_{\ell=s}^k (\alpha_\ell - \alpha_{\ell+1}) \int_0^v I_{\ell,s}(v, t) b^V(t) dt \\ &= \int_0^v \sum_{s=1}^k \alpha_s \left[\sum_{\ell=1}^s P_\ell(v) \cdot I_{s,\ell}(v, t) - \sum_{\ell=1}^{s-1} P_\ell(v) \cdot I_{s-1,\ell}(v, t) \right] b^V(t) dt. \end{aligned} \quad (13)$$

Note that the roles of s and ℓ have been reversed, such that $s \geq \ell$ henceforth. We now recall that

$$P_\ell(v) = \binom{n-1}{\ell-1} (1 - F(v))^{\ell-1} F(v)^{n-\ell}$$

and consider each of the two sums inside the parentheses in turn.

Denoting

$$J_{\ell,s} = \binom{n-1}{\ell-1} \binom{n-\ell-1}{n-s-1} (n-\ell),$$

we have that

$$\begin{aligned} \sum_{\ell=1}^s P_\ell(v) \cdot I_{s,\ell}(v, t) &= \sum_{\ell=1}^s J_{\ell,s} \cdot (1 - F(v))^{\ell-1} f(t) F(t)^{n-s-1} (F(v) - F(t))^{s-\ell} \\ &= \sum_{\substack{1 \leq \ell \leq s \\ 0 \leq x \leq \ell-1 \\ 0 \leq y \leq s-\ell}} J_{\ell,s} \binom{\ell-1}{x} \binom{s-\ell}{y} (-1)^{\ell+y-x-1} f(t) F(v)^{s-x-y-1} F(t)^{n+y-s-1}, \end{aligned}$$

where the second equality holds because by the binomial theorem

$$(1 - F(v))^{\ell-1} = \sum_{x=0}^{\ell-1} \binom{\ell-1}{x} (-F(v))^{\ell-x-1} \quad \text{and}$$

$$(F(v) - F(t))^{s-\ell} = \sum_{y=0}^{s-\ell} \binom{s-\ell}{y} F(v)^{s-\ell-y} (-F(t))^y.$$

We claim that the terms with $x + y < s - 1$ cancel out, i.e., that

$$\sum_{\substack{1 \leq \ell \leq s \\ 0 \leq x \leq \ell-1 \\ 0 \leq y \leq s-\ell \\ x+y < s-1}} J_{\ell,s} \binom{\ell-1}{x} \binom{s-\ell}{y} (-1)^{\ell+y-x-1} f(t) F(v)^{s-x-y-1} F(t)^{n+y-s-1}$$

$$= \sum_{\substack{0 \leq z \leq s-2 \\ 0 \leq y \leq z \\ z-y+1 \leq \ell \leq s-y}} J_{\ell,s} \binom{\ell-1}{z-y} \binom{s-\ell}{y} (-1)^{\ell+2y-z-1} F(v)^{s-z-1} f(t) F(t)^{n+y-s-1} = 0.$$

Indeed, the first equality follows by setting $z = x + y$ and observing that in both sums ℓ takes exactly the values between $x + 1 = z - y + 1$ and $s - y$. The second equality holds because for any z and y with $0 \leq z \leq s - 2$ and $0 \leq y \leq z$,

$$\sum_{\ell=z-y+1}^{s-y} J_{\ell,s} \binom{\ell-1}{z-y} \binom{s-\ell}{y} (-1)^{\ell+2y-z-1}$$

$$= \sum_{\ell=z-y+1}^{s-y} \binom{n-1}{\ell-1} \binom{n-\ell-1}{n-s-1} (n-\ell) \binom{\ell-1}{z-y} \binom{s-\ell}{y} (-1)^{\ell+2y-z-1}$$

$$= \frac{(n-1)!}{(n-s-1)!(z-y)!y!} \sum_{\ell=z-y+1}^{s-y} \frac{(-1)^{\ell+2y-z-1}}{(\ell-z+y-1)!(s-\ell-y)!}$$

$$= \frac{(n-1)!}{(n-s-1)!(z-y)!y!} \sum_{j=0}^m \frac{(-1)^{j+y}}{j!(m-j)!}$$

$$= \frac{(n-1)!(-1)^y}{(n-s-1)!(z-y)!y!m!} \sum_{j=0}^m (-1)^j \binom{m}{j}$$

$$= \frac{(n-1)!(-1)^y}{(n-s-1)!(z-y)!y!m!} (1 + (-1))^m = 0, \tag{14}$$

where the third equality follows by setting $j = \ell - z + y - 1$ and $m = s - z - 1$ and the fifth equality holds by the binomial theorem. Thus, actually,

$$\sum_{\ell=1}^s P_{\ell}(v) \cdot I_{s,\ell}(v, t) = \sum_{\substack{1 \leq \ell \leq s \\ 0 \leq x \leq \ell-1 \\ 0 \leq y \leq s-\ell \\ x+y=s-1}} J_{\ell,s} \binom{\ell-1}{x} \binom{s-\ell}{y} (-1)^{\ell+y-x-1} f(t) F(v)^{s-x-y-1} F(t)^{n+y-s-1}$$

$$= \sum_{\ell=1}^s J_{\ell,s} \binom{\ell-1}{\ell-1} \binom{s-\ell}{s-\ell} (-1)^{s-\ell} f(t) F(v)^0 F(t)^{n-\ell-1}$$

$$\begin{aligned}
&= \sum_{\ell=1}^s J_{\ell,s} \cdot (-1)^{s-\ell} F(t)^{n-\ell-1} f(t) \\
&= \sum_{\ell=1}^s \binom{n-1}{s-1} (n-s) \binom{s-1}{\ell-1} (-1)^{s-\ell} F(t)^{n-\ell-1} f(t) \\
&= \sum_{\ell=0}^{s-1} \binom{n-1}{s-1} (n-s) \binom{s-1}{\ell} (-1)^{s-\ell-1} F(t)^{n-\ell-2} f(t) \\
&= \binom{n-1}{s-1} (1-F(t))^{s-1} (n-s) F(t)^{n-s-1} f(t), \tag{15}
\end{aligned}$$

where the third equality holds because

$$\begin{aligned}
J_{\ell,s} &= \binom{n-1}{\ell-1} \binom{n-\ell-1}{n-s-1} (n-\ell) = \frac{(n-1)!}{(n-\ell)!(\ell-1)!} \frac{(n-\ell-1)!}{(s-\ell)!(n-s-1)!} (n-\ell) \\
&= \frac{(n-1)!}{(\ell-1)!(s-\ell)!(n-s-1)!} = \frac{(n-1)!}{(n-s)!(s-1)!} \frac{(s-1)!}{(s-\ell)!(\ell-1)!} (n-s) \\
&= \binom{n-1}{s-1} \binom{s-1}{\ell-1} (n-s)
\end{aligned}$$

and the fifth equality because by the binomial theorem

$$\sum_{\ell=0}^{s-1} \binom{s-1}{\ell} (-1)^{s-\ell-1} F(t)^{s-\ell-1} = (1-F(t))^{s-1}.$$

Analogously, for the second term in the parentheses of (13),

$$\begin{aligned}
\sum_{\ell=1}^{s-1} P_{\ell}(v) \cdot I_{s-1,\ell}(v,t) &= \sum_{\ell=1}^{s-1} J_{\ell,s-1} \cdot (1-F(v))^{\ell-1} f(t) F(t)^{n-s} (F(v) - F(t))^{s-\ell-1} \\
&= \sum_{\substack{1 \leq \ell \leq s-1 \\ 0 \leq x \leq \ell-1 \\ 0 \leq y \leq s-\ell-1}} J_{\ell,s-1} \cdot \binom{\ell-1}{x} \binom{s-\ell-1}{y} (-1)^{\ell+y-x-1} f(t) F(v)^{s-x-y-2} F(t)^{n+y-s},
\end{aligned}$$

where the second equality holds because by the binomial theorem

$$\begin{aligned}
(1-F(v))^{\ell-1} &= \sum_{x=0}^{\ell-1} \binom{\ell-1}{x} (-F(v))^{\ell-x-1} \quad \text{and} \\
(F(v) - F(t))^{s-\ell-1} &= \sum_{y=0}^{s-\ell-1} \binom{s-\ell-1}{y} F(v)^{s-\ell-y-1} (-F(t))^y.
\end{aligned}$$

We claim that the terms with $x + y < s - 2$ cancel out, i.e., that

$$\begin{aligned} & \sum_{\substack{1 \leq \ell \leq s-1 \\ 0 \leq x \leq \ell-1 \\ 0 \leq y \leq s-\ell-1 \\ x+y < s-2}} J_{\ell, s-1} \binom{\ell-1}{x} \binom{s-\ell-1}{y} (-1)^{\ell+y-x-1} f(t) F(v)^{s-x-y-2} F(t)^{n+y-s} \\ &= \sum_{\substack{0 \leq z \leq s-3 \\ 0 \leq y \leq z \\ z-y+1 \leq \ell \leq s-y-1}} J_{\ell, s-1} \binom{\ell-1}{z-y} \binom{s-\ell-1}{y} (-1)^{\ell+2y-z-1} f(t) F(v)^{s-z-2} F(t)^{n+y-s} = 0. \end{aligned}$$

Indeed, the first equality follows by setting $z = x + y$ and observing that in both sums ℓ takes exactly the values between $x + 1 = z - y + 1$ and $s - y - 1$. The second equality holds because for any z and y with $0 \leq z \leq s - 3$ and $0 \leq y \leq z$,

$$\sum_{\ell=z-y+1}^{s-y-1} J_{\ell, s-1} \binom{\ell-1}{z-y} \binom{s-\ell-1}{y} (-1)^{\ell+2y-z-1} = \sum_{\ell=z-y+1}^{r-y} J_{\ell, r} \binom{\ell-1}{z-y} \binom{r-\ell}{y} (-1)^{\ell+2y-z-1} = 0,$$

where the first equality follows by setting $r = s - 1$ and the second equality holds by (14). Thus, actually,

$$\begin{aligned} \sum_{\ell=1}^{s-1} P_{\ell}(v) \cdot I_{s-1, \ell}(v, t) &= \sum_{\substack{1 \leq \ell \leq s-1 \\ 0 \leq x \leq \ell-1 \\ 0 \leq y \leq s-\ell-1 \\ x+y=s-2}} J_{\ell, s-1} \binom{\ell-1}{x} \binom{s-\ell-1}{y} (-1)^{\ell+y-x-1} f(t) F(v)^{s-x-y-2} F(t)^{n+y-s} \\ &= \sum_{\ell=1}^{s-1} J_{\ell, s-1} \binom{\ell-1}{\ell-1} \binom{s-\ell-1}{s-\ell-1} (-1)^{s-\ell-1} f(t) F(v)^0 F(t)^{n-\ell-1} \\ &= \sum_{\ell=1}^{s-1} J_{\ell, s-1} \cdot (-1)^{s-\ell-1} F(t)^{n-\ell-1} f(t) \\ &= \sum_{\ell=1}^{s-1} \binom{n-1}{s-1} (s-1) \binom{s-2}{\ell-1} (-1)^{s-\ell-1} F(t)^{n-\ell-1} f(t) \\ &= \sum_{\ell=0}^{s-2} \binom{n-1}{s-1} (s-1) \binom{s-2}{\ell} (-1)^{s-\ell-2} F(t)^{n-\ell-2} f(t) \\ &= \binom{n-1}{s-1} (1 - F(t))^{s-2} (s-1) F(t)^{n-s} f(t), \end{aligned} \tag{16}$$

where the third equality holds because

$$\begin{aligned} J_{\ell, s-1} &= \binom{n-1}{\ell-1} \binom{n-\ell-1}{n-s} (n-\ell) = \frac{(n-1)!}{(n-\ell)!(\ell-1)!} \frac{(n-\ell-1)!}{(s-\ell-1)!(n-s)!} (n-\ell) \\ &= \frac{(n-1)!}{(\ell-1)!(s-\ell-1)!(n-s)!} = \frac{(n-1)!}{(n-s)!(s-1)!} \frac{(s-2)!}{(s-\ell-1)!(\ell-1)!} (s-1) \\ &= \binom{n-1}{s-1} \binom{s-2}{\ell-1} (s-1) \end{aligned}$$

and the fifth equality because by the binomial theorem

$$\sum_{\ell=0}^{s-2} \binom{s-2}{\ell} (-1)^{s-\ell-2} F(t)^{s-\ell-2} = (1 - F(t))^{s-2}.$$

By substituting (15) and (16) into (13), we conclude that

$$\begin{aligned} \mathbb{E} [p^V(v)] &= \int_0^v \sum_{s=1}^k \alpha_s \left(\binom{n-1}{s-1} (1 - F(t))^{s-1} (n-s) F(t)^{n-s-1} f(t) - \right. \\ &\quad \left. \binom{n-1}{s-1} (1 - F(t))^{s-2} (s-1) F(t)^{n-s} f(t) \right) b^V(t) dt \\ &= \sum_{s=1}^k \alpha_s \int_0^v \frac{dP_s(t)}{dt} b^V(t) dt. \end{aligned} \tag{17}$$

The expected payment in an efficient equilibrium is again given by Lemma 3. We can thus equate (17) with (5), take derivatives on both sides, and solve for $b^V(v)$ to obtain

$$b^V(v) = \frac{\sum_{s=1}^k \beta_s \frac{dP_s(v)}{dv} v}{\sum_{s=1}^k \alpha_s \frac{dP_s(v)}{dv}}.$$

Bidding below $b^V(0) = 0$ is impossible and bidding above $b^V(\bar{v})$ is dominated, so the claim follows from Lemma 3.

E Proof of Lemma 6

By Corollary 1, $b^F(v) = A(v)/B(v)$ with $A(v) = \sum_{s=1}^k \beta_s \int_0^v \frac{dP_s(t)}{dt} t dt$ and $B(v) = \sum_{s=1}^k \alpha_s P_s(v)$. Writing the derivative as a limit of difference quotients, applying l'Hospital's rule to each of the two resulting terms, and respectively substituting x for 2δ and δ we obtain

$$\begin{aligned} \left. \frac{db^F(v)}{dv} \right|_{v=0} &= \lim_{\delta \rightarrow 0} \left(\frac{A(2\delta)/B(2\delta) - A(\delta)/B(\delta)}{\delta} \right) \\ &= \lim_{\delta \rightarrow 0} \frac{A(2\delta)}{\delta \cdot B(2\delta)} - \lim_{\delta \rightarrow 0} \frac{A(\delta)}{\delta \cdot B(\delta)} \\ &= \lim_{\delta \rightarrow 0} \frac{A'(2\delta) \cdot 2}{\delta \cdot B'(2\delta) \cdot 2 + B(2\delta)} - \lim_{\delta \rightarrow 0} \frac{A'(\delta)}{\delta \cdot B'(\delta) + B(\delta)} \\ &= \lim_{x \rightarrow 0} \frac{\left(\sum_{s=1}^k \beta_s \frac{dP_s(x)}{dx} \cdot x \right) \cdot 2}{\left(\sum_{s=1}^k \alpha_s \frac{dP_s(x)}{dx} \cdot x \right) + \left(\sum_{s=1}^k \alpha_s P_s(x) \right)} - \\ &\quad \lim_{x \rightarrow 0} \frac{\left(\sum_{s=1}^k \beta_s \frac{dP_s(x)}{dx} \cdot x \right)}{\left(\sum_{s=1}^k \alpha_s \frac{dP_s(x)}{dx} \cdot x \right) + \left(\sum_{s=1}^k \alpha_s P_s(x) \right)}. \end{aligned}$$

To analyze these limits we extend by $1 = (F(x)^{n-k-1} \cdot x)^{-1} / (F(x)^{n-k-1} \cdot x)^{-1}$, factor $(F(x)^{n-k-1} \cdot x)^{-1}$ into the numerator and denominator, and consider each of the terms in the numerator and denominator in turn.

Using γ as a placeholder for α or β and replacing $P_s(x)$ by its definition,

$$\begin{aligned} \frac{\sum_{s=1}^k \gamma_s \cdot \frac{dP_s(x)}{dx} \cdot x}{F^{n-k-1}(x) \cdot x} &= \sum_{s=1}^k \gamma_s \left[\binom{n-1}{s-1} (n-s) F^{k-s}(x) (1-F(x))^{s-1} f(x) \right. \\ &\quad \left. - \binom{n-1}{s-1} (s-1) F^{k-s+1}(x) (1-F(x))^{s-2} f(x) \right] \\ &= \sum_{s=1}^k \sum_{\ell=0}^{s-1} \gamma_s (-1)^\ell \binom{n-1}{s-1} (n-s) \binom{s-1}{\ell} F(x)^{k-s+\ell} f(x) \\ &\quad - \sum_{s=1}^k \sum_{\ell=0}^{s-2} \gamma_s (-1)^\ell \binom{n-1}{s-1} (s-1) \binom{s-2}{\ell} F(x)^{k-s+\ell+1} f(x). \end{aligned}$$

Similarly,

$$\begin{aligned} \frac{\sum_{s=1}^k \alpha_s P_s(x)}{F^{n-k-1}(x) \cdot x} &= \frac{1}{x} \cdot \sum_{s=1}^k \alpha_s \binom{n-1}{s-1} F(x)^{k-s+1} (1-F(x))^{s-1} \\ &= \frac{F(x)}{x} \cdot \sum_{s=1}^k \sum_{\ell=0}^{s-1} \alpha_s (-1)^\ell \binom{n-1}{s-1} \binom{s-1}{\ell} F(x)^{k-s+\ell}. \end{aligned}$$

Since $\lim_{v \rightarrow 0} F(x)^d = 0$ for $d > 0$, the only terms that survive in the limit are those where the exponent of $F(x)$ is zero. For $s \in \{1, \dots, k\}$ and $\ell \in \{0, \dots, s-1\}$, $k-s+\ell = 0$ only if $s = k$ and $\ell = 0$. For $s \in \{1, \dots, k\}$ and $\ell \in \{0, \dots, s-2\}$, $k-s+\ell-1 \neq 0$. Using that $\lim_{x \rightarrow 0} F(x)/x = f(0)$, we thus obtain

$$\begin{aligned} \left. \frac{db^F(v)}{dv} \right|_{v=0} &= \frac{\beta_k \binom{n-1}{k-1} (n-k) f(0) \cdot 2}{\alpha_k \binom{n-1}{k-1} (n-k) f(0) + \alpha_k \binom{n-1}{k-1} f(0)} - \frac{\beta_k \binom{n-1}{k-1} (n-k) f(0)}{\alpha_k \binom{n-1}{k-1} (n-k) f(0) + \alpha_k \binom{n-1}{k-1} f(0)} \\ &= \frac{2(n-k)\beta_k}{(n-k+1)\alpha_k} - \frac{(n-k)\beta_k}{(n-k+1)\alpha_k} = \frac{n-k}{n-k+1} \cdot \frac{\beta_k}{\alpha_k} \end{aligned}$$

as claimed.

F Proof of Lemma 7

By Corollary 1, $b^F(v) = A(v)/B(v)$ with $A(v) = \sum_{s=1}^k \beta_s \int_0^v \frac{dP_s(t)}{dt} t dt$ and $B(v) = \sum_{s=1}^k \alpha_s P_s(v)$. For $n = k$, by Lemma 6,

$$\left. \frac{db^F(v)}{dv} \right|_{v=0} = \left. \frac{A'(v)B(v) - A(v)B'(v)}{B(v)^2} \right|_{v=0} = 0.$$

Since

$$B(0) = \sum_{s=1}^k \alpha_s P_s(0) \geq \alpha_k P_k(0) = \alpha_k (1-F(0))^{n-1} = \alpha_k > 0,$$

this implies that

$$\left. (A'(v)B(v) - A(v)B'(v)) \right|_{v=0} = 0.$$

Thus

$$\begin{aligned} \left. \frac{d^2 b^F(v)}{dv^2} \right|_{v=0} &= \left. \frac{(A''(v)B(v) - A(v)B''(v))B(v)^2 - (A'(v)B(v))^2}{B(v)^4} \right|_{v=0} - \left. \frac{A(v)B'(v)2B(v)B'(v)}{B(v)^4} \right|_{v=0} \\ &= \left. \frac{A''(v)B(v) - A(v)B''(v)}{B(v)^2} \right|_{v=0}. \end{aligned}$$

We have already seen that $B(0) > 0$. Moreover, $A(0) = 0$ by the definition of A and $B''(0) < \infty$ by assumption on the value distributions, so it suffices to show that

$$A''(0) = \left(\sum_{s=1}^k \beta_s \frac{d^2 P_s(v)}{dv^2} \cdot v + \sum_{s=1}^k \beta_s \frac{dP_s(v)}{dv} \right) \Big|_{v=0} \geq 0.$$

Also by assumptions on the value distributions, $d^2 P_s(v)/dv^2 < \infty$ for all v , so the first term vanishes. The second term is

$$\begin{aligned} \sum_{s=1}^k \beta_s \frac{dP_s(v)}{dv} \Big|_{v=0} &= \sum_{s=1}^k \beta_s \left(\binom{n-1}{s-1} (n-s) F(v)^{n-s-1} (1-F(v))^{s-1} f(v) \right. \\ &\quad \left. - \binom{n-1}{s-1} (s-1) F(v)^{n-s} (1-F(v))^{s-2} f(v) \right) \Big|_{v=0} \\ &= \beta_{k-1} (k-1) f(0) - \beta_k (k-1) f(0) \geq 0, \end{aligned}$$

where we have used the definition of $P_s(v)$ and the fact that the only non-zero terms are those where the exponent of $F(v)$ is zero. Since $\beta_{k-1} \geq \beta_k$ and $f(0) > 0$, this shows the claim.

References

- [1] S. Alaei, H. Fu, N. Haghpanah, and J. D. Hartline. The simple economics of approximately optimal auctions. In *Proceedings of the 54th Symposium on Foundations of Computer Science*, pages 628–637, 2013.
- [2] L. M. Ausubel and P. Milgrom. The lovely but lonely Vickrey auction. In P. Cramton, Y. Shoham, and P. Steinberg, editors, *Combinatorial Auctions*, chapter 1, pages 17–40. MIT Press, 2006.
- [3] M. Babaioff, N. Immorlica, B. Lucier, and S. M. Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *Proceedings of the 55th Symposium on Foundations of Computer Science*, pages 21–30, 2014.
- [4] D. Bergemann and S. Morris. An introduction to robust mechanism design. *Foundations and Trends in Microeconomics*, 8(3):169–230, 2013.
- [5] K. Bhawalkar and T. Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 700–709, 2011.

- [6] M. Bisson. Bing Ads auction explained: How bid, cost-per-click and quality score work together, 2015. URL <http://advertise.bingads.microsoft.com/en-us/blog/27821/>. Accessed October 26, 2015.
- [7] Y. Cai and C. H. Papadimitriou. Simultaneous bayesian auctions and computational complexity. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 895–910, 2014.
- [8] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 81–90, 2011.
- [9] S. Chawla and J. Hartline. Auctions with unique equilibria. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, pages 181–196, 2013.
- [10] S. Chawla, J. D. Hartline, and D. Nekipelov. Mechanism design for data science. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 711–712, 2014.
- [11] G. Christodoulou, A. Kovács, and M. Schapira. Bayesian combinatorial auctions. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*, pages 820–832, 2008.
- [12] P. Dütting, F. Fischer, and D. C. Parkes. Simplicity-expressiveness tradeoffs in mechanism design. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 341–350, 2011.
- [13] P. Dütting, F. Fischer, and D. C. Parkes. Expressiveness and robustness of first-price position auctions. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 57–74, 2014.
- [14] P. Dütting, V. Gkatzelis, and T. Roughgarden. The performance of deferred-acceptance auctions. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 187–204, 2014.
- [15] P. Dütting, T. Kesselheim, and É. Tardos. Algorithms as mechanisms: The price of anarchy of relax-and-round. In *Proceedings of the 16th ACM Conference on Economics and Computation*, pages 187–201, 2015.
- [16] B. Edelman, M. Ostrovsky, and M. Schwartz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.
- [17] M. Feldman, H. Fu, N. Gravin, and B. Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, pages 201–210, 2013.
- [18] R. Gomes and K. S. Sweeney. Bayes-Nash equilibria of the generalized second-price auction. *Games and Economic Behavior*, 86:421–437, 2014.
- [19] Google, Inc. Actual cost-per-click (CPC), 2015. URL <https://support.google.com/adwords/answer/6297>. Accessed October 26, 2015.

- [20] J. D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, pages 225–234, 2009.
- [21] J. D. Hartline and S. Taggart. Non-revelation mechanism design via i.i.d. position auctions, 2015. Working paper.
- [22] T. Kaplan and S. Zamir. Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Economic Theory*, 50(2):269–302, 2012.
- [23] R. Paes Leme and É. Tardos. Pure and Bayes-Nash price of anarchy for generalized second price auction. In *Proceedings of the 51st Symposium on Foundations of Computer Science*, pages 735–744, 2010.
- [24] H. B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91(3):461–479, 1983.
- [25] B. Lucier and A. Borodin. Price of anarchy for greedy auctions. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 537–553, 2010.
- [26] B. Lucier and V. Syrgkanis. Greedy algorithms make efficient mechanisms. In *Proceedings of the 16th ACM Conference on Economics and Computation*, pages 221–238, 2015.
- [27] P. Milgrom. Simplified mechanisms with an application to sponsored-search auctions. *Games and Economic Behavior*, 70(1):62–70, 2010.
- [28] P. Milgrom and I. Segal. Deferred-acceptance auctions and radio spectrum reallocation. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 185–186, 2014.
- [29] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [30] M. H. Rothkopf. Thirteen reasons why the Vickrey-Clarke-Groves process is not practical. *Operations Research*, 55(2):191–197, 2007.
- [31] T. Roughgarden. Barriers to near-optimal equilibria. In *Proceedings of the 55th Symposium on Foundations of Computer Science*, pages 71–80, 2014.
- [32] V. Syrgkanis and É. Tardos. Composable and efficient mechanisms. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, pages 211–220, 2013.
- [33] H. Varian. Position auctions. *International Journal of Industrial Organization*, 25:1163–1178, 2007.
- [34] H. R. Varian and C. Harris. The VCG auction in theory and practice. *American Economic Review*, 104(5):442–445, 2014.
- [35] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.
- [36] A. C.-C. Yao. An n -to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 92–109, 2015.