

Best-Response Auctions

Noam Nisan*

Michael Schapira†

Gregory Valiant‡

Aviv Zohar§

ABSTRACT

We present a new framework for auction design and analysis that we term “best-response auctions”. We use this framework to show that the simple and myopic best-response dynamics converge to the VCG outcome and are incentive compatible in several well-studied auction environments (Generalized Second Price auctions, and auctions with unit-demand bidders). Thus, we establish that in these environments, given that all other bidders are repeatedly best-responding, the best course of action for a bidder is to also repeatedly best-respond. Our results generalize classical results in economics regarding convergence to equilibrium and incentive compatibility of ascending-price English auctions. In addition, our findings provide new game-theoretic justifications for some well-studied auction rules. Best-response auctions provide a way to bridge the gap between the full-information equilibrium concept and the usual private-information auction theory.

1. INTRODUCTION

1.1 Why Best-Respond?

Arguably, the most striking application of “electronic markets and auctions” are the billions of penny “ad-auctions” that bankroll much of the Internet as we know it. It is quite an embarrassment that the pricing rule used in these auctions, almost universally, is the “Generalized Second Price” (GSP), rather than the more theoretically-motivated “Vickrey-Clarke-Groves” (VCG) rule which basic mechanism design theory would suggest¹.

*School of Eng. and Computer Science, The Hebrew University of Jerusalem. noam@cs.huji.ac.il

†Computer Science Dept. Princeton. ms7@cs.princeton.edu

‡Dept. of Computer Science, UC Berkeley. gvaliant@eecs.berkeley.edu

§Microsoft Research, Silicon Valley. avivz@microsoft.com

¹Of course there are various non-game-theoretic reasons why GSP was chosen, such as its cognitive simplicity rel-

In the Generalized Second-Price (GSP) auction, there are k slots that are to be sold. Each slot j has a click-through-rate (CTR) α_j such that $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_k$ (“higher” slots are clicked more often than “lower” ones). There are n bidders (advertisers) and each bidder i has a private value $v_i \in \mathbf{R}^+$ per click (and thus his value for the j 'th slot is $\alpha_j v_i$). GSP assigns the k slots to the k highest bidders (the highest bidder gets the highest slot, *etc.*) and charges each winning bidder a cost per-click that equals the bid of the bidder that was assigned the slot below his.

Two early landmark papers on ad auctions [3, 6] provide some explanation; they show that the theoretically motivated VCG pricing emerges as an equilibrium of the GSP rule. While at first sight this might lessen the embarrassment, at second sight it makes it even worse: the type of analysis under which this result is obtained is that of full-information, *i.e.*, this analysis is based on the implicit assumption that all bidders know each other's private values. That is, this explanation essentially reverts to classical economic theory without taking any information-economics considerations into account, disregarding the point of view that is at the foundation of mechanism design.

In particular, showing that VCG prices are an equilibrium of the GSP auction does not address the question of *how* the bidders may reach this equilibrium without having the required information. This issue was partially addressed by Cary et. al. [1], who show that if the bidders participate in a GSP auction repeatedly, and at each step best-respond to the others' previous bids, then the VCG equilibrium is indeed reached. The actual theorem is a bit more delicate as this only applies with high probability, when bidders' turns are chosen at random. This answer to the question “How?”, begs the question “Why?": Why would we expect the bidders to repeatedly best-respond to the others?

While there is much intuitive appeal to the simple and myopic repeated best-response strategy, we seek a game-theoretic justification. Perhaps some other long-sighted strategy could improve a bidder's utility? Is the “repeated-best-response” strategy really a best response overall, in the extended (multiple round) game? We present a positive answer to this question. We show that if all bidders in the GSP auction are repeatedly best-responding, then so should you. This result allows us to resurrect some of the partial information assumptions: there is no need to know the exact values of bidders, only that they are best-responding according to *some* values.

ative to VCG, but still one would expect some game-theoretic/economic justification.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC '11 San Jose, CA USA

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

We put this result in a wider context of auctions where repeated-best-response converge to the full-information equilibrium and do so in an incentive-compatible manner. A companion paper [4] considers more general (non-auction) mechanisms with this property. Our model bridges the gap between the full-information equilibrium concept and the usual private-information auction theory.

1.2 Illustration: Iterative Single-Item Auction

It is instructive to informally demonstrate the types of results we get, via the first-price auction of a single item. It is well known, though still quite intriguing, we believe, that the full-information equilibrium of such a first-price auction is achieved at the second price, that is, at the VCG outcome.

Consider a discretized setting where bids must belong to some pre-defined discrete set $\{0, \delta, 2\delta, \dots\}$ (that is, δ is the “minimal increment” for bids). In the full-information equilibrium of the first-price auction in this discretized setting, the highest bidder bids the lowest bid that exceeds the second highest valuation for the item v_2 and the second highest bidder bids the highest bid that does not exceed v_2 . The ascending English auction, where bidders repeatedly increase their bids by the minimal increment, and drop from the auction once their values are exceeded, can be regarded as an effective mechanism for achieving exactly this outcome.

We suggest the following iterative auction mechanism. The auction allows bidders to repeatedly announce bids, based on the most recent bid announcements of the other bidders. The prescribed strategy for each bidder is to repeatedly best-respond to the most recent bid announcements of the others. Under the first-price auction base rule, this means that when a bidder’s turn to act comes, if the bidder can outbid the current winning bid without exceeding his value for the item then he shall submit the minimal such bid (so as to minimize his payment). In the event that a bidder cannot outbid the current winning bid without exceeding his bid, we prescribe that the bidder break ties between the multiple possible best-responses in favor of the highest bid that does not exceed his value.

The auction terminates after all bidders consecutively “pass”, *i.e.*, repeat their last bid, in which case the outcome is according to the base auction rule on the final bids². If the auction does not terminate within a pre-determined number of rounds then the item remains unallocated.

Observe that if the initial bids are all 0, then this best-response dynamics proceeds with the sequence of minimally-increasing prices, and thus mimicks the behavior of the ascending-price English auction. We show that it is also true in general that, for any starting vector of bids, and every order of bidders’ best-responses, the second-price equilibrium of the full-information first-price auction is reached, provided that the tie-breaking rule is chosen carefully.

Our main result for this restricted environment is that repeated-best-response is indeed each bidder’s global best-response to the others’ repeated-best-response strategies. That is, given that all other bidders are following the prescribed behavior and repeatedly best-responding, a bidder cannot gain from not doing the same. Thus, the incentive compatibility of the English auction is but a special case of a more general result.

²The fixed termination rule is important in our settings, and our results do not hold without it. This is in contrast to the results in our companion paper [4] that hold more generally.

Theorem: (Informal) *Under the first-price auction base rule, repeated best-response strategies (with a specific tie-breaking rule) converge to the VCG outcome in an incentive compatible manner.*

1.3 GSP and Unit-Demand Auctions

As a generalization of the single item auction, we consider a framework for auction design that we term “*best-response auctions*”. Given a base auction rule, a best-response auction allows bidders to repeatedly announce bids in response to the most recent bids of the others, and the prescribed strategy for each bidder is to repeatedly best-respond (under the base auction rule). A tie breaking rule specifies which bid to choose in cases where there are several different best-responses. This process goes on until (1) all bidders consecutively pass, in which case the outcome is according to the base auction rule on the final bids; or (2) the auction does not terminate within a pre-determined number of rounds, in which case all goods on sale remain unallocated.

Under this framework, we consider two well-studied auction environments: the GSP auction mentioned above and the auction of multiple heterogeneous items among unit-demand bidders, which has received much attention since the seminal paper of Demange, Gale, and Sotomayer (see [5]). In the unit-demand setting, each bidder i has a value $v_i(j)$ for each item j , and is interested in purchasing exactly one item. We consider a base auction rule where each item is sold individually in a first-price auction.

For both of these environments, we prove the general result that best-response dynamics converge to the optimal allocation of items, and are also in global equilibrium, *i.e.*, that each bidder is indeed best-off following the repeated-best-response strategy if the others are also doing so. In both cases, our best-response auction mechanism is *randomized*—in each round, a single bidder is selected uniformly at random to act. (Our results for GSP auctions actually also imply that deterministic repeated *better*-response dynamics studied in [1] also converge to the optimal allocation and are incentive compatible.)

Specifically, we show the following.

Theorem: (Informal) *Under the GSP auction base rule, repeated best-response strategies (with a specific tie-breaking rule) converge to the VCG outcome in an incentive compatible manner.*

Theorem: (Informal) *Under our auction base rule for unit-demand bidders, repeated best-response strategies (with a specific tie-breaking rule) converge to the VCG outcome in an incentive compatible manner.*

We point out that generalized ascending-price English auction mechanisms that implement the VCG outcome in an incentive-compatible manner exist in both GSP auctions [3] and in auctions with unit-demand bidders [2]. As in the single-item first-price auction, these mechanisms can be regarded as specific executions of best-response dynamics and thus, as special cases of best-response auctions.

When designing best-response auctions we face two main challenges: (1) establishing convergence, that is, that best-response dynamics reach a stable outcome (with probability 1); and (2) establishing incentive compatibility, that is, that the outcome reached is such that no bidder has incentive to deviate from the prescribed behavior. The focus of our main

effort in each of the two auction environments we consider is different: in the case of the GSP auction, convergence was known and our effort is concentrated on proving incentives; in the unit-demand auction most of our effort is focused on proving convergence.

1.4 Future Research

We view this work as a first step towards a more general research agenda. Our results establish that in several well-studied settings where the VCG outcome matches the full-information equilibrium, repeated best-response dynamics converge to the VCG outcome in an incentive compatible manner. We still lack a general understanding of when such results can be obtained (our results establish sufficient, but not necessary, conditions). What more general principle underlies our results, that is, under what conditions does repeated best-response lead to optimal outcomes *and* is in global equilibrium?

Our results in this work pertain to convergence to *optimal* outcomes (that is, to an allocation of resources for which the overall social welfare is maximized). An interesting direction for future research is understanding in what environments repeated best-response provides reasonable *approximation* guarantees (again, in an incentive compatible manner). The answer to these questions can serve as the basis for the design of natural and simple to implement auction mechanisms.

In addition, our work focuses on a specific natural and myopic update rule—repeated best-response. We believe the examination of other dynamics (*e.g.*, fictitious play, regret minimization) from the perspective introduced here is also of interest.

1.5 Organization

We present our game-theoretic framework and illustrate our approach via single-item first-price auctions in Section 2. We present our convergence and incentive compatibility results for single-item first-price auctions, GSP auctions, and auctions with unit-demand bidders, in Section 3, Section 4, and Section 5, respectively.

2. OUR FRAMEWORK: BEST-RESPONSE AUCTIONS

A best-response auction can be regarded as a (finite time) execution of best-response dynamics in a game determined by the auction base rule. From this perspective, the outcome of a best-response auction corresponds to a pure Nash equilibrium of the game. We now present basic game theoretic terminology and explain the connection to best-response auctions.

2.1 Background: Games

We use standard game-theoretic notation. Let Γ be a *normal-form game* with n players $1, \dots, n$. We denote by S_i be the *strategy space* of the i^{th} player. Let $S = S_1 \times \dots \times S_n$, and let $S_{-i} = S_1 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_n$ be the cartesian product of all strategy spaces but S_i . Each player i has a *utility function* u_i that specifies i 's *payoff* in every strategy-profile of the players. For each strategy $s_i \in S_i$, and every $(n-1)$ -tuple of strategies $s_{-i} \in S_{-i}$, we denote by $u_i(s_i, s_{-i})$ the utility of the strategy profile in which player i plays s_i and all other players play their strategies in s_{-i} . We will make use of the following definitions.

DEFINITION 2.1 (BEST-RESPONSE STRATEGIES). A *strategy* $s_i \in S_i$ is a best response of player i to a strategy profile $s_{-i} \in S_{-i}$ of the other players if

$$s_i \in \operatorname{argmax}_{s'_i \in S_i} u_i(s'_i, s_{-i}).$$

DEFINITION 2.2 (PURE NASH EQUILIBRIA). A *strategy profile* s is a pure Nash equilibrium if, for every player i , s_i is a best response of i to s_{-i} .

Under best-response dynamics, players take turns in best-responding to other players' current strategies. The process proceeds as follows: Start from some arbitrary strategy profiles and allow players to best-respond, one by one in some order, to the current strategy profile. This goes on until a pure Nash equilibrium is reached. We refer to best-response dynamics in which the order of players' activations is cyclic as "*round-robin best-response dynamics*" and to best-response dynamics in which players are chosen uniformly at random to best-respond as "*randomized best-response dynamics*".

2.2 Best-Response Auctions and Games

Under a best-response auction, bidders repeatedly best-respond (given the auction best-rule) to the most recent bids of others. This goes on until a "stable state" is reached, or "sufficiently long" time has passed. Therefore, a best-response auction can be regarded as a finite-time execution of best-response dynamics in a game determined by the auction base-rule.

To illustrate this point, let us revisit the example of a single-item first-price auction discussed in the Introduction. We consider a discretized setting where there is a single item for sale and n bidders $1, \dots, n$. Each bidder i has a private value $v_i \in [0, M\delta]$ (where $M > 0$ is fixed and δ is some small real number) for getting the item. The goal is to assign the object to the bidder that values it the most (breaking ties lexicographically if there are multiple such bidders).

Our best-response single-item auction can be regarded as round-robin best-response dynamics in the following game.

The 1st-price auction game: The bidders are the players. Each bidder's strategy space S_i is a predetermined discrete set $\{0, \delta, 2\delta, \dots, M\delta\}$. The utility of player i from a strategy (bid) tuple $b = (b_1, \dots, b_n)$ is $u_i(b) = v_i - b_i$ if i is the highest bidder (breaking ties lexicographically) and 0 otherwise.

Best-response single-item auction: Under our best-response auction, players take part in round-robin best-response dynamics until a pure Nash equilibrium is reached, in which case the item is allocated to the highest bidder who is charged his bid, or until $n^2 M$ rounds have passed, in which case the item remains unallocated.

To fully specify the best-response single-item auction we must determine how bidders break ties between multiple best responses. Importantly, this tie-breaking rule must be "uncoupled", *i.e.*, depend solely on the player's private information (utility function) and not on information that is unavailable to him. Our tie-breaking rule in the 1st-price auction game context has the following simple form: fix, for each player i , an a-priori full order \prec_i on S_i (that can depend on u_i), and instruct player i to break ties between multiple best-responses according to \prec_i .

Tie-breaking rules for the 1st-price auction game: Each player i 's tie breaking rule \prec_i is as follows: $\forall s, t \in S_i$,

if $s > v_i \vee t > v_i$ then $s \prec_i t$ iff $s > t$. Otherwise, $s \prec_i t$ iff $s < t$.

In Section 3, we prove that the best-response single-item auction converges to the (discretized) outcome of the second-price Vickrey auction in an incentive compatible manner.

3. WARMUP: BEST-RESPONSE SINGLE-ITEM AUCTION

We now present our results for the single item first-price auction rule. Specifically, we show that for any starting vector of bids, and every round-robin order of bidders' best-responses, the second-price equilibrium of the full-information first-price auction is reached, given our tie-breaking rule. Moreover, repeated-best-response is indeed each bidder's global best-response to the others' repeated-best-response strategies.

THEOREM 3.1. *The best-response single-item auction converges to the (discretized) outcome of the second-price Vickrey auction in an incentive compatible manner.*

PROOF. To prove the theorem we first establish convergence to the VCG outcome. We then prove the incentive properties of the auction.

Convergence: Assume all players are following the repeated best-response strategy with the prescribed tie breaking rules, and that the auction evolves from some arbitrary profile of bids. We will observe the dynamics in iterations of n rounds, during which every player is activated once. Notice that after the first such iteration, no bidder bids above his private value. This is because a bid above one's value is never a best response: Any winning bid above the value is strictly worse than bidding below it, and all losing bids above the value have the same utility, but our tie breaking rules prescribe a preference for bidding below the value in this case. For the same reasons, no player will ever bid above his value again. We say that a player is "out of the race" if from some point on, he bids his value exactly, and never changes his bid.

The proof continues by showing that as time progresses, players eliminate certain bids and will never bid them again: Let us look first at the weakest player, that has all ties broken against him. We argue that he will never bid 0 after the first iteration of activations, unless his value happens to be 0, in which case he will be out of the race. A bid of 0 by that player will always be a losing bid, and so, if it is under the player's value it will never be a best response to the bids of others (tie breaking rules prefer higher losing bids as long as they are under the valuation). If the weakest player is out of the race, then the second weakest player will, for similar reasons, never bid 0 after the following iteration or he will drop out of the race. This goes on either until we remain with only one player that is still in the race, or until the weakest player bids at least δ (and bid at least that during the rest of the process). In the latter case, after the next iteration, all other players will bid at least δ , and will do so during the rest of the process (as the weakest player in the race is never going to bid below that). Again, the weakest player still in the race will either drop out of the race or from this point on or never bid anything below 2δ . This line of reasoning goes on until we are left with only a single player in the race. At that point all other players bid

their valuation and never change their bids. The remaining player will therefore bid the second price (or δ above that to break ties) and will never change his bid – this is his best response. All players will then pass, and we have converged.

Incentive Compatibility: Let us assume that some player i deviates from the repeated best-response strategy and gains by doing so. We argue that the player can only gain from the deviation if he wins the item because of it (otherwise his utility is 0, and before the deviation he did no worse). In particular, all players must have passed in the last iteration, as the mechanism would not allocate the item otherwise. If the i gains by deviating, then he must be getting the item at a price p that is lower price than the price without the deviation. Since the price is lower, both the highest and second highest bidder (only one of whom might be player i) still wish to purchase the item at price p and so the final bidding profile cannot be stable, and at least one player would not pass. We therefore reach a contradiction. \square

4. BEST-RESPONSE GSP AUCTIONS

The following "GSP game" was presented in [3, 6]. Our best-response GSP auction can be regarded as randomized best-response dynamics in this game.

The Game. The bidders are the players. Each bidder's strategy space S_i is some predetermined discrete set $\{0, \delta, 2\delta, \dots, M\delta\}$ for some arbitrarily small δ . For every bid-vector $b = (b_1, \dots, b_n) \in S$, $\pi_j(b)$ denotes the bidder that is awarded the j 'th slot. The utility of the bidder that gets slot j when the bid-vector is $b = (b_1, \dots, b_n)$ is thus $u_{\pi_j(b)} = \alpha_j(v_{\pi_j(b)} - b_{\pi_{j+1}(b)})$.

Tie-breaking rules: Consider the case that bidder i 's best-response results in bidder i getting slot j . Then, i should select the bid b_i such that $\alpha_j \cdot (v_i - b_{\pi_{j+1}(b)}) = \alpha_{j-1} \cdot (v_i - b_i)$. Intuitively, i should choose a bid b_i such that i is indifferent between getting the j 'th slot and paying the next highest bid, or getting the $(j-1)$ 'th slot and paying b_i . If none of bidder i 's best-responses result in i getting a slot, then i should bid his highest feasible bid, *i.e.*, the maximal bid b_i that does not exceed v_i .

Convergence: Cary et al. [1] study convergence of deterministic and randomized best- and better-response dynamics to the VCG outcome in GSP games and prove two positive results for convergence to the VCG outcome: (1) probabilistic convergence for best-response dynamics (with the above tie-breaking rules); and (2) deterministic convergence for specific better-response dynamics.

Incentive compatibility: We now show that the positive results in [1] are actually achieved in an incentive-compatible manner (no bidder is motivated to deviate from repeated best- and better-response). To establish this, we prove that there does not exist a "bad state" in GSP games, in the sense that there is no strategy profile from which all players but (possibly) a single player i do not wish to deviate, and that i strictly prefers to the pure Nash equilibrium that would be reached if all players follow the repeated best-response strategies. Observe that the nonexistence of such a bad state is indeed sufficient (and actually, also necessary) to establish the incentive compatibility of best-response auctions for this environment.

PROPOSITION 4.1. *Let N be the unique pure Nash equilibrium (under tie-breaking) in a GSP game (this is a VCG outcome). Then, there does not exist a state $b = (s_1, \dots, s_n) \in S$, and player $i \in [n]$, such that $\forall j \neq i, s_j$ is a best-response (under tie-breaking) to s and $u_i(s) > u_i(N)$.*

For ease of notation in our proof of the above proposition, let $b \in S$ denote the bad state; σ_j is the bidder that is allocated the j 'th slot when the bid-vector is b ; π_j is the bidder that gets slot j in the VCG outcome, and VCG_j is that bidder's bid in the VCG outcome. Observe that if the unstable bidder does not win a slot then his utility is 0 and so he is certainly not better off than in the unique pure Nash equilibrium outcome. Thus, the unstable bidder is awarded some slot \bar{j} . Our proof of Proposition 4.1 will follow easily from the following lemma.

LEMMA 4.2. *For each bidder i that wins a slot lower than \bar{j} (that is, a slot $r > \bar{j}$) v_i is at least as high as that of the bidder that got the same slot in the VCG outcome, and b_i is at least as high as the bid of the bidder that got the same slot in the VCG outcome. That is,*

$$\forall r > \bar{j} \quad VCG_r \leq b_{\sigma_r} \quad \text{and} \quad v_{\pi_r} \leq v_{\sigma_r}.$$

PROOF. In the VCG outcome the $n - k$ losing bidders are the bidders with the lowest values and they bid these values. Observe that in b the losing bidders are the $n - k$ bidders with the lowest bids, and these bids are equal to their values, and thus for all $i \geq k + 1$, we have that $v_{\sigma_i} \geq v_{\pi_i}$.

The lemma now follows easily by induction: assuming that for all $i \geq r \geq \bar{j} + 2$, it holds that $v_{\sigma_i} \geq v_{\pi_i}$, and $b_{\sigma_i} \geq VCG_i$. Consider the bidder σ_{r-1} , since this bidder is following the prescribed best-responses, his value must be at least v_{σ_r} , and thus his value must be in the top $r - 1$ valuations, and thus $v_{\sigma_{r-1}} \geq v_{\pi_{r-1}}$. Additionally, he will bid the prescribed "indifference price";

$$\begin{aligned} b_{\sigma_{r-1}} &= v_{\sigma_{r-1}} - \frac{\alpha_{r-1}}{\alpha_{r-2}} (v_{\sigma_{r-1}} - b_{\sigma_r}) \\ &\geq v_{\pi_{r-1}} - \frac{\alpha_{r-1}}{\alpha_{r-2}} (v_{\pi_{r-1}} - b_{\sigma_r}) \\ &\geq v_{\pi_{r-1}} - \frac{\alpha_{r-1}}{\alpha_{r-2}} (v_{\pi_{r-1}} - b_{\pi_r}) = VCG_{r-1}, \end{aligned}$$

where the second and third lines above follow from noting that the indifference price in the first line is an increasing function of both $v_{\sigma_{r-1}}$ and b_{σ_r} . \square

Proof of Proposition 4.1: The above lemma shows that the "unstable" bidder $\sigma_{\bar{j}}$ is forced to pay for the \bar{j} 'th slot a price that is at least as high as the payment of the bidder who gets that slot in the VCG outcome, and thus the unstable bidder would prefer to be allocated the \bar{j} 'th slot in the VCG allocation. To conclude, since the VCG outcome ensures that each bidder gets at least as much utility from his allocation as he would get were he to receive the allocation allotted to another player, the unstable bidder gets at most the utility he would get were to get item \bar{j} in the VCG outcome, which gives at most as much utility as he would get in the actual VCG outcome. \square

5. BEST-RESPONSE UNIT-DEMAND AUCTIONS

Our best-response auction for unit-demand bidders can be regarded as best-response dynamics in the following class of games.

The 1st-price unit-demand auction game: The bidders are the players. Each bidder's strategy is a vector of non-negative bids $b = (b_1, \dots, b_m)$, where each b_j represents the bid for item j and belongs to some predetermined discrete set $\{0, \delta, 2\delta, \dots, M\delta\}$ for some arbitrarily small δ , which we refer to as the *discretization parameter*. Given an n -tuple of bid vectors (i.e., players' strategies) (b^1, \dots, b^n) , the induced price-vector $\vec{p} = (\mathbf{p}_1, \dots, \mathbf{p}_m)$ is such that $\forall j \in [m]$, $p_j = \max_{i \in [n]} b_j^i$, where b_j^i is bidder i 's bid for item j . To complete the specification of the game we need the following:

DEFINITION 5.1 (OVERDEMANDED SETS). *Given an n -tuple of bid vectors $\mathbf{b} = (b^1, \dots, b^n)$, and the induced price vector \mathbf{p} , $\forall j \in [m]$, $B(j) := \{i \in [n] \mid b_j^i = p_j\}$ (that is, the set of highest bidders for j); $\forall i \in [n]$, $T(i) := \{j \in [m] \mid b_j^i = p_j\}$ (that is, the set of items bidder i is a highest bidder for). $X \subseteq [m]$ is overdemanded if $|\{i \mid T(i) \subset X\}| > |X|$. We say that such a set of items X is a minimal overdemanded set if no subset of X is overdemanded.*

Given an n -tuple of bid vectors $\mathbf{b} = (b^1, \dots, b^n)$, the allocation of items $\Gamma(\mathbf{b})$ is defined as follows: (1) all items that are in minimal overdemanded sets remain unallocated; (2) allocate bidder i the item with the lowest index for which i is the single highest bidder, if such an item exists; and (3) allocate all remaining items so that item j be given to a bidder i such that $i \in B(j)$, breaking ties among such allocations lexicographically, and allocating at most one item to each bidder. We are now ready to describe the utility function u_i of each bidder, i . Given an n -tuple of bid vectors $\mathbf{b} = (b^1, \dots, b^n)$, let $\Gamma_i(\mathbf{b})$ be the set of items allocated to i in $\Gamma(\mathbf{b})$, and note that $|\Gamma_i(\mathbf{b})| \in \{0, 1\}$. Then if $|\Gamma_i(\mathbf{b})| = 1$, $u_i(\mathbf{b}) = v_{i\Gamma_i(\mathbf{b})} - p_{\Gamma_i(\mathbf{b})}$, and $u_i = 0$ otherwise, where $\vec{p} = (p_1, \dots, p_m)$ is the price vector associated with bids \mathbf{b} .

Our main theorem is that there exist tie-breaking rules under which the randomized best-response auction converges and is approximately incentive-compatible, in that by unilaterally deviating from the prescribed best-responses, one can increase one's utility by at most a factor proportional to the discretization parameter δ .

THEOREM 5.2. *there exist tie-breaking rules under which the randomized best-response auction converges and is approximately incentive-compatible. Specifically, in an instance of the unit-demand auction game with n players, m items, with player valuations v_{ij} , for $i \in [n], j \in [m]$, and discretization parameter δ , given any initial configuration of bids,*

- *After at least $k \leq n + m + (m^2 + mn) \cdot \max_{i \in [n], j \in [m]} v_{ij} / \delta$ turns, with probability at least $1/n^k$, the unit-demand auction game is at a strategy profile for which no player has a prescribed best-response that deviates from his current bid.*
- *Given any strategy profile of the unit-demand auction game for which no player has a prescribed best-response that deviates from his current bid, the utility that each*

player receives is within $\delta(m+2)$ of the utility they receive in the VCG allocation of the corresponding unit-demand auction.

Our proof of the above theorem consists of three main components. In Proposition 5.3, we show that all the pure Nash equilibrium of the unit-demand auction game (under the prescribed tie-breaking rules) are similar to the VCG outcome, in that the final prices are close to the VCG payments from which it follows that the utility each player receives is similar to the utility that the player would receive under the VCG outcome. In Proposition 5.10, we show that given any configuration, there exists a prescribed sequence of at most $k \leq n+m+(m^2+mn) \cdot \max_{i \in [n], j \in [m]} v_{ij}/\delta$ players such that after such a sequence of players best-responding, the game will be at a pure equilibrium, and no player will henceforth wish to deviate from his current strategy. The final component of our proof of Theorem 5.2 is Proposition 5.11, which shows that the above convergence occurs in an incentive-compatible manner, in that by unilaterally deviating from the prescribed best responses, no player can benefit his utility by more than $\delta(m+2)$.

Before proceeding further, we now describe the prescribed tie-breaking rules. Before rigorously stating these rules, we first informally describe and motivate them. In short, a player first chooses the/an item that he desires most at the lowest price he would need to bid for that item in order to receive it, given everyone else's bids. His bid for that most-desired item, say item i , is then the lowest bid he would need to make for that item so as to receive that item; his bid for all other items is such that he would receive δ more utility for receiving any of these other items at the prices he bids, than he would receive for his "most-desired", namely his bids b_j for all items $j \neq i$ satisfy $v_j - b_j = v_i - b_i + \delta$, where v_j denotes that player's valuation for item j .

Motivated by the Hungarian method, and similar approaches to computing efficient allocations (optimal weighted matchings), it is intuitively clear that players should roughly bid "level prices"—prices such that the player is indifferent to which item he receives at the prices he bids. Additionally, there must be some mechanism that allows the prices to fall in the case that, for example, an efficient matching has been found, yet the prices are all elevated (say the VCG allocation and the VCG prices increased by some constant); as we will see, by having each player bid δ less than his "level prices" for all but the most-desired item, we allow the prices to both rise, and fall, ultimately eventually finding a near-efficient allocation, with prices roughly the VCG prices. We now formally define the tie-breaking rules.

Tie-breaking rules. Given an n -tuple of bid vectors $\mathbf{b} = (b^1, \dots, b^n)$, let $p_{\Gamma_i(\mathbf{b})}$ denote the price i pays for $\Gamma_i(\mathbf{b})$ (and $\Gamma_i(\mathbf{b}) = \emptyset$ if i is not allocated any items); $\forall j \in [m]$, p_j^* denotes the minimum bid i must place on item j (while bidding 0 on all other items) so as to receive j in the resulting allocation. Bidder i should break ties as follows.

- If $\max_{r \in [m]}(v_{ir} - p_r^*) < 0$, then bid $\max(0, v_{ir} - \delta)$ for all items.
- If $\max_{r \in [m]}(v_{ir} - p_r^*) \geq 0$, and either $\Gamma_i(\mathbf{b}) = \emptyset$ or $v_{i\Gamma_i(\mathbf{b})} - p_{\Gamma_i(\mathbf{b})} < \max_{r \in [m]}(v_{ir} - p_r^*)$, then do the following. Let $s := \operatorname{argmax}_{r \in [m]}(v_{ir} - p_r^*)$. Bid p_s^* for item s , and $\max(0, v_{is} - v_{ir} + p_r^* - \delta)$ for all items

$r \neq s$. Throughout, we shall call item s (the most desired item at the time of bidding) the "target item" for bidder i .

- If $v_{i\Gamma_i(\mathbf{b})} - p_{\Gamma_i(\mathbf{b})} = \max_{r \in [m]}(v_{ir} - p_r^*)$, bid $p_{\Gamma_i(\mathbf{b})}$ for item $\Gamma_i(\mathbf{b})$, and $\max(0, v_{ir} - v_{i\Gamma_i(\mathbf{b})} + p_{\Gamma_i(\mathbf{b})} - \delta)$ for all items $r \neq \Gamma_i(\mathbf{b})$.

5.1 Equilibria of Unit-Demand Auctions

In this section we prove that all pure Nash equilibria of the unit-demand auction game are "close" to the VCG outcome. The proof relies fundamentally on the well-known fact that the VCG outcome and prices are the minimal competitive equilibrium prices. In particular, to show that the prices at an equilibrium in the unit-demand auction game cannot be much lower than the VCG prices, we show that no δ -approximate competitive equilibrium can have prices too much lower than the VCG prices. Our proof that the equilibrium prices can not be too much higher than the VCG prices is slightly more involved, though at its core, leverages the fact that given any competitive equilibrium for which the prices are higher than the VCG prices, there exists some way of reallocating the items and decreasing some subset of the bids that only increases each player's utility.

PROPOSITION 5.3. *Let G be a unit-demand auction game, with discretization parameter δ . Then, all pure Nash equilibrium in G under the above tie-breaking rules are close to the VCG outcome.*

Specifically, letting $\vec{q} = (q_1, \dots, q_m)$ denote the VCG prices, and $u_1^{VCG}, \dots, u_n^{VCG}$ the associated utilities of each bidder in the VCG outcome, then if bid vectors $\mathbf{b} = b^1, \dots, b^n$ are at equilibrium in the unit-demand auction game, then the associated allocation $\Gamma(\mathbf{b})$ and price vector $\vec{p} = (p_1, \dots, p_m)$ satisfy:

- $\forall i \in [m], |p_i - q_i| \leq (m+1)\delta$.
- $\forall i \in [n], |u_i(\mathbf{b}) - u_i^{VCG}| \leq (m+2)\delta$.

We assume that $\forall i, j, v_{ij} \geq 0$. We begin by making two basic observations:

OBSERVATION 5.4. *At equilibrium \mathbf{b} , with associated prices $\vec{p} = (p_1, \dots, p_m)$, for all bidders $i \in [n]$,*

$$u_i(\mathbf{b}) \geq \max_{j \in [m]}(v_{ij} - p_j) - \delta$$

OBSERVATION 5.5. *If bidder i 's target item was item j , and in the current allocation $\Gamma_i(\mathbf{b}) \neq \{j\}$, then i has a best-response that deviates from his bid b^i .*

We now begin our proof of Proposition 5.3.

LEMMA 5.6. *If bid vectors $\mathbf{b} = b^1, \dots, b^n$ with associated prices $\vec{p} = (p_1, \dots, p_m)$ are at equilibrium in the unit-demand auction game, then every item i such that $p_i \geq \delta$ is allocated in $\Gamma(\mathbf{b})$.*

PROOF. Assume for the sake of contradiction that this is not the case, and some item $i \in [m]$ has $p_i \geq \delta$ but is unallocated. Let $B(i)$ be the set of bidders with the highest bid for item i . Since i is unallocated, every bidder in $B(i)$ that is not allocated an item must be bidding δ less than their value, but if there is such a player, then by bidding his value he

would get the item; thus all bidders in $B(i)$ must already be allocated an item, which, because \mathbf{b} is at equilibrium, must be each player's 'target' item (by Observation 5.5). But this is also a contradiction, since a player in $B(i)$ would value item i at price p_i δ more than the item he receives, and thus the best-response dynamics would prescribe him to change his target item to item i , and bid p_i for item i and drop his other bids accordingly. \square

LEMMA 5.7. *Let $\vec{q} = (q_1, \dots, q_m)$ denote the VCG prices. If bid vectors $\mathbf{b} = b^1, \dots, b^n$ with associated prices $\vec{p} = (p_1, \dots, p_m)$ are at equilibrium in the unit-demand auction game, then $p_i - q_i \leq (m+1)\delta$.*

PROOF. Assume for the point of contradiction that there is some item i such that $p_i - q_i \geq (m+2)\delta$. We split the argument into two cases: the case where there is some item j for which $p_j \leq q_j + \delta$, and the case where $p_j \geq q_j + 2\delta$ in every coordinate.

In the first case, by the pigeon-hole principle, there must exist some set $I \subset [m]$ such that $\forall i \in I, p_i - q_i \geq 2\delta$, and additionally, for all $j \notin I, p_i - q_i \geq 2\delta + p_j - q_j$. For example, if we sort the items in decreasing order according to $p_i - q_i$, there must be two items in consecutive order for which their respective values of $p - q$ differ by at least 2δ .

Now, consider the set of bidders $B(I)$ who have the highest bid for some item of I , and let $Q(I)$ denote the bidders who receive items in I in the VCG allocation. We claim that $B(I) = Q(I)$. First, note that $|B(I)| \geq |I|$, since all items in I have price at least 2δ , and thus by Lemma 5.6 are allocated. Since $|B(I)| \geq |I| \geq |Q(I)|$, it suffices to show that for each $i \in B(I)$, bidder i receives some item in I in allocation Q . First, note that bidder i must receive some item in allocation Q , since $\exists j \in [m]$ such that $v_{ij} - p_j \geq 0$, and thus $v_{ij} - q_j \geq p_j - q_j \geq 2\delta$, so bidder i would not be content without an item given prices \vec{q} . Let $j \in [m]$ denote an item such that $i \in B(j)$. From our definition of the set I , it follows that bidder i must receive some item that is also in set I , since for any item $k \notin I, p_j - q_j - 2\delta \geq p_k - q_k$, but because the procedure terminated, we know that $v_{ij} - p_j \geq v_{ik} - p_k - \delta$. Adding these two equations yields $v_{ij} - q_j \geq v_{ik} - q_k + \delta$, and thus bidder i must receive an item in I in allocation Q , and thus $B(I) = Q(I)$; that is, the highest bidders for items in the set I must be the set of bidders who are allocated items in I (in both allocation $\Gamma(\mathbf{b})$ and the VCG allocation Q).

To conclude this case, consider the directed graph (with no self-loops) whose vertex set is $B(I)$, and where there is a directed edge from i to i' if bidder i' has a highest bid for the item that i is allocated in $\Gamma(\mathbf{b})$. Because the protocol terminated, and the conclusion of the previous paragraph (that the highest bidders for items in the set I must be the set of bidders who are allocated items in I), each vertex must have out-degree at least one, and thus there must be some directed cycle C in the graph. The allocation Γ' induced by modifying the allocation $\Gamma(\mathbf{b})$ according to the cycle C would be a valid matching with fewer bidders receiving their target items than in $\Gamma(\mathbf{b})$, and thus at least one bidder would have a best response (which would get him the utility he would get in Γ').

We now consider the case that $\forall j \in [m], q_j \leq p_j - 2\delta$. The set of bidders having a highest bid in \vec{p} must be the set of bidders receiving items in the allocation Q , because for such a bidder i , there is an item j such that $v_{ij} \geq p_j \geq q_j + 2\delta$. Now, the same argument as in the previous case (in which a

directed graph as above with a cycle is constructed) yields the desired contradiction. \square

We now prove the other direction—that the equilibrium prices can be at most $(m+1)\delta$ less than the corresponding VCG prices.

LEMMA 5.8. *Let $\vec{q} = (q_1, \dots, q_m)$ denote the VCG prices. If bid vectors $\mathbf{b} = b^1, \dots, b^n$ with associated prices $\vec{p} = (p_1, \dots, p_m)$ are at equilibrium in the unit-demand auction game, then $q_i - p_i \leq (m+1)\delta$.*

PROOF. We first argue that there is some item $j \in [m]$ such that $p_j \geq q_j - \delta$. Assume otherwise, and note that any bidder $i \in [n]$ who does not receive an item in allocation $\Gamma(\mathbf{b})$ must have valuations $v_{ij} \leq p_j + \delta \leq q_j - \delta$, for all items j , and thus will not receive or desire an item in the VCG allocation Q . Additionally, since all bidders receiving items in allocation Q must also receive items in $\Gamma(\mathbf{b})$ (because otherwise, they would desire an item, and have a best-response), and all m items must be allocated in the VCG outcome since $q_i > 2\delta$ for all i , we conclude that the set of bidders receiving an item in allocation $\Gamma(\mathbf{b})$ must be the same as the set receiving an item in allocation Q . Decreasing every component of \vec{q} by δ and keeping allocation Q would yield a smaller competitive equilibrium, since every bidder receiving an item would still be at equilibrium given prices $\vec{q} - \delta$, and any bidder i who does not receive an item will still not desire an item, since, as noted above, it must be the case that $v_{ij} \leq q_j - \delta$. This contradicts the minimality of the competitive equilibrium prices \vec{q} .

Given that there is some item $j \in [m]$ such that $p_j \geq q_j - \delta$, we now argue that $\max_{j \in [m]} (q_j - p_j) \leq (m+1)\delta$. Assume otherwise, and consider the set of items I such that $\forall j \in I, k \notin I, q_j - p_j \geq 2\delta + q_k - p_k$, and $q_j - p_j \geq 2\delta$. Such a set exists and is nonempty by the pigeonhole principle, the fact that some item j satisfies $p_j \geq q_j - \delta$, and our assumption that some item k satisfies $q_k - p_k \geq (m+2)\delta$.

Note that $Q(I) = \Gamma_I(\mathbf{b})$, since bidders who get an item of I in allocation Q would not be at equilibrium at prices \vec{p} unless they got an item of I . Next, for a bidder i who does not receive an item in allocation $\Gamma(\mathbf{b})$, it must be the case that $i \notin Q(I)$, and for item $j \in I, v_{ij} \leq q_j - \delta$. For a bidder i who receives an item $j \notin I$ in allocation $\Gamma(\mathbf{b})$, for any item $k \in I$, we must have $(v_{ij} - p_j) - (v_{ik} - p_k) \geq -\delta$, which from the definition of the set I implies that $(v_{ij} - q_j) - (v_{ik} - q_k) \geq \delta$. Thus the price vector \vec{q}' defined by decreasing every coordinate of \vec{q} by δ will also be a competitive equilibrium, which is a contradiction, as desired. \square

Proposition 5.3 now follows from the above two lemmas, and the observation that if the equilibrium price vector satisfies $|p_i - q_i| \leq (m+1)\delta$, and for all bidders $i \in [n], v_{i\Gamma_i(\mathbf{b})} - p_{\Gamma_i(\mathbf{b})} \geq \max_{j \in [m]} (v_{ij} - p_j) - \delta$, then the utility that bidder i derives from the outcome associated with \mathbf{b} is at most $(m+2)\delta$ different than the utility he would receive in the VCG outcome.

5.2 Convergence of Repeated Best-Response

We first show that round-robin best-response dynamics can oscillate indefinitely in unit-demand auction games. We then consider randomized best-response dynamics; a single player is chosen to play uniformly at random at each turn. In this randomized setting, to show convergence, it suffices

to show that given any initial conditions, there exists some finite *prescribed* sequence of players such that after the prescribed sequence of players best-response, an equilibrium will be reached. At a high-level, we prescribe such a list of players so that the dynamics of the prices have two general phases. In the first phase, the prices of all items only increase. This first phase resembles an “ascending-price auction” amongst some subset of the players. The second phase then closely resembles a “descending-price auction,” in which the price of every item can only decrease.

PROPOSITION 5.9. *There exists an initial configuration of the unit-demand auction game, such that if bidders perform best-responses in round-robin order, the dynamics cycle indefinitely.*

PROOF. Consider the auction with bidder set $\{A, B, C\}$ and two items x, y . Let each bidder value each item for 2 units. Consider the initial values with bidder A bidding $1, 1 + \epsilon$ for items x, y respectively, bidder B bidding $1 + \epsilon, 1 + 2\epsilon$, for the two items, and bidder C bidding $1 + \epsilon, 1$. Suppose the round-robin order has A responding first, then B , then C . We now consider the dynamics: let $(\cdot, \cdot)_A$ denote the vector of bids for player A response, and correspondingly for $(\cdot, \cdot)_B$ and $(\cdot, \cdot)_C$. It is easy to verify that the prescribed best-responses will be as follows:

$$(1 + 2\epsilon, 1 + \epsilon)_A, (1, 1 + \epsilon)_B, (1 + \epsilon, 1 + 2\epsilon)_C, (1 + \epsilon, 1)_A \\ (1 + 2\epsilon, 1 + \epsilon)_B, (1, 1 + \epsilon)_C, \dots$$

We now see that after two full rounds of bidding, the configuration of bids is exactly as it was prior to the first best-response, except with the bids for items x switched with those for item y . Thus, by the symmetry of the bidders’ valuations, and the fact that at every step, they have a unique most-desired item, after four complete rounds of bidding we will be in the exact same configuration as prior to the first best-response. \square

PROPOSITION 5.10. *Let n be the number bidders, and m the number of items. Given any initial configuration, after $k = n + m + (m^2 + mn) \cdot \max_{i \in [n], j \in [m]} v_{ij} / \delta$ best-responses, in which the player best-responding at each turn is chosen uniformly at random from $[n]$, the probability that an equilibrium has been reached is at least $1/n^k$.*

PROOF. It suffices to show that that given any initial configuration, there exists at least one sequence of players of length at most k such that if the bidders are chosen according to that sequence, an equilibrium is reached. For any initial configuration, we will describe such a sequence of players; roughly, we will select the players such that dynamics will have two phases: the ‘ascending phase’, in which the vector of prices \vec{p} associated with the bids \mathbf{b} will only increase (as in an ascending auction with reserve prices), and the ‘descending phase’ that will resemble a descending-price auction. Let $\vec{p}^t = (p_1^t, \dots, p_m^t)$ denote the vector of maximum prices at time t , and say $t = 0$ corresponds to the initial configuration.

We first let each player make a single best-response bid. Observe that after this initial set of n turns, each bid is at most $\max_{i \in [n], j \in [m]} v_{ij}$. We now begin the ‘ascending phase’ of the dynamics at time $t = n$. At each time step, we choose a player $i \in [n]$ to best-respond who has the property that

prior to his turn, he was not allocated an item, but after his turn, he will be allocated an item (and we say that the ‘ascending phase’ has ended when no such bidder can be chosen). Since no price can ever decrease, after at most m steps of the ‘ascending phase’, either the price for some item was increased by at least δ , or all items are allocated, in which case the next bidder must increase the price of some item. Thus by time $t = n + m^2 \cdot \max_{i \in [n], j \in [m]} v_{ij} / \delta$ turns the ascending phase will have terminated (since otherwise at least one of the m items would have a price above the maximum value of any item to any bidder, which is impossible). Finally, we let each player who is not allocated an item have a turn. Let time $t = t_1$ be at the completion of this extra set of turns. At this time, every bidder who could derive at least δ utility from receiving some item $j \in [m]$ and paying price $p_j^{t_1}$ is allocated an item. Additionally, for every item j , all players $i \in [n]$ who are not allocated an item at time t_1 are bidding $v_{ij} - \delta$.

We now begin the ‘descending phase’ of the dynamics. At each time step we choose a bidder $i \in [n]$ who has a best-response bid that differs from their current bid, provided that they are allocated an item in the timestep prior to their turn. Thus at each turn, no bid will ever be increased, and at least one bid will drop by at least δ , and thus this stage will terminate after at most $nm \cdot \max_{i \in [n], j \in [m]} v_{ij} / \delta$ turns. Next, we argue that a player i who is not allocated an item at time $t > t_1$ must be bidding $v_{ij} - \delta$ for all items. This is clearly true at time t_1 . Now, note that if bidder i' best-responds at time $t + 1$ and does not increase any of his bids, every bidder receiving an item at time t will also receive one at time $t + 1$, and thus at every time $t \geq t_1$, it must be the case that the players who are not allocated an item must not have been allocated an item at time t_1 .

Finally, any additional players whose best-responses differ from their current bids are given turns. We claim that the only players who have best-responses during this stage are players who are not allocated any item, but who share the highest bid for some item with at least one other such player. Consider the first player, i , to play in this stage; he must not be allocated an item prior to his turn (otherwise his turn would have been part of the ‘descending phase’), and thus he must bid $v_{ij} - \delta$ for all items prior to his turn, and his best-response bid must be to increase his bid for some item j to v_{ij} , and he must be allocated item j as a result of his turn. Assume for the sake of contradiction that there is some other bidder i' with a bid of v_{ij} for item j . i' must not be allocated item j prior to the best response of bidder i (otherwise i would not have been able to get item j unless i' also had a highest bid for a non-overdemanded item j' , in which case i' would have had a best-response that should have taken place as part of the ‘descending phase’). But if i' was not allocated item j prior to the turn of i , but i is allocated j after his turn, then either j was not the target item of i' , in which case i' would have had a best-response as part of the descending phase (in which he made j his target item and decreased his other bids), or item j was his target item, in which case he would also have a best-response as part of the descending phase. Thus after the best-response of i , he is the unique bidder to bid v_{ij} for item j , and all his other bids remain unchanged, and thus the above arguments hold for all additional players who best-respond in this final stage.

To conclude, note that the final stage does not affect any

bidder who was allocated an item prior to this final stage, and thus all bidders receiving an item after this stage still satisfy the property that they do not have a best-response deviation from their current bid. Additionally, all bidders i who do not receive an item are bidding $v_{ij} - \delta$ for all items, and additionally raising their bids to v_{ij} will not result in their getting an item. Thus an equilibrium has been reached. \square

5.3 Incentive Compatibility

In this section we show that if any player unilaterally deviates from the prescribed best-responses in the unit-demand auction game, he can not significantly increase his utility, thus completing our proof of Theorem 5.2 to establish that the randomized best-response mechanism is incentive compatible in Model A.

Our proof approach is simple: we show that if, by deviating from the prescribed best-responses, player i induces the dynamics to converge to some outcome \mathbf{b} , it is the case that the outcome is, in fact, a pure Nash equilibrium of a related instance of the unit-demand auction game. Thus by Proposition 5.3 it follows that the outcome is close to the VCG outcome of this modified game. We then argue that player i can not prefer the VCG outcome of the modified game to the VCG outcome of the original game, and thus there is no incentive to deviate from the prescribed best-response dynamics.

PROPOSITION 5.11. *Let N be a pure Nash equilibrium (under tie-breaking) in a unit-demand auction game with discretization parameter δ . There does not exist a state $\mathbf{b} = (b_1, \dots, b_n) \in S$, and player $i \in [n]$, such that $\forall j \neq i$, b_j is a best-response (under tie-breaking) to \mathbf{b} and $u_i(\mathbf{b}) > u_i(N) + 2(m+2)\delta$.*

PROOF. Let G denote the instance of the unit-demand auction game. Assume for the point of contradiction that there is some bidder $i \in [n]$, and some configuration \mathbf{b} such that all bidders $j \neq i$ have no best-response deviation from their bids b_j , but for which $u_i(\mathbf{b}) > u_i(N) + 2(m+1)\delta$. Let \vec{p} denote the corresponding vector of prices associated to \mathbf{b} . For ease of exposition, assume without loss of generality that $i = 1$.

First, note that $\Gamma_1(\mathbf{b}) \neq \emptyset$. We will now construct a related game G' whose set of bidders is $\{c^*, 2, \dots, n, c_1, \dots, c_m\}$, and whose set of items is $[m]$. The valuation are defined as follows: for all $i \in \{2, \dots, n\}$, the valuations are identical to the valuations of the corresponding players in game G . Let $v_{c^* \Gamma_1(\mathbf{b})} = p_{\Gamma_1(\mathbf{b})}$, and for all items $j \neq \Gamma_1(\mathbf{b})$, $v_{c^* j} = p_j + \delta$. For the remaining bidders c_i , let $v_{c_i i} = p_i + \delta$, and $v_{c_i j} = 0$ for all $j \neq i$. Define bids \mathbf{b}' for game G' as follows: each player $i \in \{2, \dots, m\}$ bids as in \mathbf{b} , bidder c_i bids p_i for item i and 0 for all other items, and bidder c^* bids \vec{p} . Note that bids \mathbf{b}' induce the same price vector as \mathbf{b} in game G . Additionally, bids \mathbf{b}' are at equilibrium, and induce allocation Γ as in game G (except with bidder c^* replacing bidder 1). Thus by Proposition 5.3, the utility that bidder c^* gets is within $(m+2)\delta$ of the VCG outcome in auction G' .

The valuations in G' for bidders $2, \dots, n$ are as in game G , and since adding additional players can only increase the VCG prices, the VCG prices for game G' are at most the VCG prices in game G'' which is defined to be identical to game G but with player 1 replaced by player c^* . To conclude, since the VCG mechanism is truthful, bidder 1 would not

prefer the VCG outcome of game G'' (item $\Gamma_1(\mathbf{b})$ at price $p_{\Gamma_1(\mathbf{b})}$) to the true VCG allocation of game G , which is at most $(m+2)\delta$ more profitable than the equilibrium that would be reached were he to best-respond. \square

6. REFERENCES

- [1] Matthew Cary, Aparna Das, Benjamin Edelman, Ioannis Giotis, Kurtis Heimerl, Anna R. Karlin, Claire Mathieu, and Michael Schwarz. Greedy bidding strategies for keyword auctions. In *EC*, pages 262–271, 2007.
- [2] Gabrielle Demange, David Gale, and Marilda Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94(4):863–72, August 1986.
- [3] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, March 2007.
- [4] Noam Nisan, Michael Schapira, Gregory Valiant, and Aviv Zohar. Best-response mechanisms. In *ICS*, 2011.
- [5] Alvin E. Roth and Marilda A. Oliveira Sotomayor. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press, 1990.
- [6] Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, December 2007.