Do Externalities Degrade GSP's Efficiency?

Tim Roughgarden, Stanford
Éva Tardos, Cornell

We consider variants of the cascade model of externalities in sponsored search auctions introduced independently by Aggrawal et al. and Kempe and Mahdian in 2008, where the click-through rate of a slot depends also on the ads assigned to earlier slots. Aggrawal et al. and Kempe and Mahdian give a dynamic programming algorithm for finding the efficient allocation in this model. We give worst-case efficiency bounds for a variant of the classical Generalized Second Price (GSP) auction in this model.

Our technical approach is to first consider an idealized version of the model where an unlimited number of ads can be displayed on the same page; here, Aggrawal et al. and Kempe and Mahdian show that a greedy algorithm finds the optimal allocation. The game theoretic analog of this greedy algorithm can be thought of as a variant of the classical GSP auction. We give the first non-trivial worst-case efficiency bounds for GSP in this model.

In the more general model with limited slots, greedy algorithms like GSP can compute extremely bad allocations. Nonetheless, we show that an appropriate extension of the greedy algorithm is approximately optimal, and that the worst-case equilibrium inefficiency in the corresponding analog of GSP also remains bounded. In the context of these models, the GSP mechanisms suffer from two forms of suboptimality: that from using a simple allocation rule (the greedy algorithm) rather than an optimal one (based on dynamic programming), and that from the strategic behavior of the bidders (caused by using the GSP's critical bid pricing rule rather than one leading to a dominant-strategy implementation). Our results show that for this class of problems, the two causes of efficiency loss can be analyzed separately.

Categories and Subject Descriptors: [**Theory and Foundations**]: Auction Theory

General Terms: Auction, Generalized Second Price, AdAuction Externalities, price of anarchy

ACM Reference Format:
Roughgarden, T., Tardos, E. 2012. Do Externalities Degrade GSP's Efficiency? ACM X, X, Article X (February 2012), 16 pages.
DOI = 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

1. INTRODUCTION

The now standard method of selling advertisements on search pages is the Generalized Second Price (GSP) auction. Advertisers are charged only when their ad is clicked on. In the auction, each advertiser places a bid expressing their willingness to pay for a click, and these bids are resolved in a simple automated auction resulting in a sequence of ads displayed next to the organic search results. The classical model of GSP was introduced in [Edelman et al. 2007; Varian 2007]. The ads are ordered on the page

Tim Roughgarden was supported in part by NSF grant CCF-1016885, an ONR PECASE Award, and an AFOSR MURI grant. Email: tim@cs.stanford.edu. Éva Tardos was supported in part by NSF grants CCF-0910940, ONR grant N00014-08-1-0031, a Yahoo! Research Alliance Grant, and a Google Research Grant.
Author's addresses: Tim Roughgarden, Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305, tim@cs.stanford.edu; Éva Tardos, Computer Science Department, Cornell University, Ithaca NY 14853, eva@cs.cornell.edu

DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

from top to bottom, and the probability of an ad i getting a click when displayed on slot j is modeled as $\gamma_i \alpha_j$, the product of two separable quantities, where α_j is the click-through rate (CTR) of the assigned slot, and γ_i depends on the quality (or relevance) of the ad. Here the CTR α_i models that the value of showing an ad decreases as its position on the page falls (as higher slots have higher CTRs), and the quality factor γ_i models that more relevant ads may receive more clicks.

A limiting aspect of the standard model is the assumption that the click-through rate of a slot does not depend on the other ads shown on the page. Modeling such externalities among the ads displayed is one of the key issues in better understanding sponsored search. In this paper we consider variants of the model of externalities introduced independently in [Aggarwal et al. 2008] and [Kempe and Mahdian 2008], called the cascade model, which provides an endogenous model of click-through rates, where the click-through rate of a slot is a function of the players assigned to earlier slots. The proposed endogenous model of click-through rates was recently shown by [**?**] to be more realistic. [Aggarwal et al. 2008] and [Kempe and Mahdian 2008] give a dynamic programming algorithm for finding the efficient allocation in the model, and show that a greedy algorithm can find the optimal allocation in the idealized version of the model where an unlimited number of ads can be displayed on the same page. The game theoretic analog of the greedy algorithm can be thought of as a variant of the classical GSP auction. We give the first non-trivial worst-case efficiency bounds for GSP in these models.

We see our results for GSP in the cascade model as a positive example of the broader question: for which problems are there simple mechanisms that always have near-optimal performance at equilibrium? Dominant-strategy mechanism design has an elegant and powerful theory (see, e.g., Nisan [Nisan 2007]) but is not always popular in practice [Ausubel and Milgrom 2006; Rothkopf 2007]. On many occasions, practitioners have willingly sacrificed strong incentive properties in favor of simpler implementations. For example, in the standard single-shot pay-per-click model of [Edelman et al. 2007; Varian 2007] for sponsored search, the Vickrey-Clarke-Groves (VCG) mechanism (e.g., [Nisan 2007]) provides a dominant-strategy implementation of the welfare-maximizing outcome. Current search auctions, however, are based instead on the GSP auction, which shares its rank-by-bid allocation rule with the VCG mechanism but uses a simpler payment rule. The GSP mechanism has a welfare-maximizing equilibrium, but participants do not generally have dominant strategies [Edelman et al. 2007; Varian 2007].

Motivated by the prevalence of relatively simple mechanisms in practice, we are interested in developing a theory of how to design such simple mechanisms:

(Q1) For which problems are there simple mechanisms that always have near-optimal performance at equilibrium?
(Q2) Which simple mechanisms have the best equilibrium guarantees, especially for equilibrium concepts more general, and easier to attain, than Nash equilibria?

In this paper, we answer these questions for GSP-style auctions in several variants of the cascade model. Our analysis of the GSP game can be view as having two separate parts, corresponding to the two causes of inefficiency of GSP-style auctions: that from using a simple allocation rule (the greedy algorithm) rather than an optimal one (based on dynamic programming), and that from the strategic behavior of the bidders (made possible by using the GSP's critical bid pricing rather than one leading to a dominant-strategy implementation). The modularity of our analysis leads us to a third general question.

(Q3) Are there well-motivated mechanism design problems and a class of simple allocation rules for which analyzing the ”price of anarchy” of an induced game under the critical bid payment rule reduces to analyzing the approximation factor of the allocation rule with truthful bids?

1.1. Our Results

We consider the cascade model from a game theoretic perspective. One can think of the greedy algorithm of [Aggarwal et al. 2008] and [Kempe and Mahdian 2008] for the case with unbounded slots as a variant of the classical GSP auction. In more detail, recall that in the standard setup, advertisers are asked for bids b_i (expressing their willingness to pay for a click), and ads are sorted by the product of bid and quality factor $\gamma_i b_i$. If players would bid their true values, then this would result in the maximally efficient sorting of ads; thus, this mechanism fails to be fully efficient only because of strategic behavior by the players. The greedy algorithm of [Aggarwal et al. 2008] and [Kempe and Mahdian 2008] can also be thought of as sorting by a product of the form $\gamma_i v_i$; the difference is that the quality factor γ_i now depends also on the effect of the ad on the user viewing later ads on the page.

— We show that all Nash equilibria of the above analog of the GSP auction for the cascade model with unbounded slots have efficiency at least $1/4$th of the maximum possible,[1] i.e., the corresponding game has a price of anarchy of at most 4. Our approximation bounds also extend well beyond Nash equilibria, both to more general and easy-to-learn full-information equilibrium concepts (like coarse correlated equilibria), as well as to Bayes-Nash equilibria in the incomplete information setting, even when players' private valuations are correlated.
— Rather than focusing on the exact greedy algorithm of [Aggarwal et al. 2008] and [Kempe and Mahdian 2008], we show the above result by considering a class of algorithms: GSP with arbitrary quality factors γ_i, and prove that the price of anarchy in the corresponding GSP mechanism is at most 4 times worse than the approximation factor of the greedy algorithm for the underlying welfare maximization problem. Considering a class of algorithms makes our analysis flexible, effectively reducing the price of anarchy analysis to designing an (approximately) efficient greedy allocation rule.
— We use the above flexibility to extend the analysis for the GSP mechanism for variants of the cascade model, including the problem with only k slots, where the greedy allocation rule is no longer optimal. [Aggarwal et al. 2008] and [Kempe and Mahdian 2008] give a dynamic programming algorithm for this problem, which is much more complex than the ranking-based algorithms prevalent in sponsored search. We give approximately optimal greedy algorithms for these problems, by modifying the quality factors γ_i to take into account also the limited number of slots available, and show a price of anarchy analysis for the resulting variant of GSP.
— Finally, in the full-information cascade model with unbounded number of slots, we show that the game has a Nash equilibrium that is fully efficient. This is analogous to the result of [Edelman et al. 2007; Varian 2007] for GSP in the classical sponsored search model (without externalities).

[1]We assume throughout the paper that bidders are conservative in that they do not bid above their value. When players can ”bluff” by bidding more than their value, even the full-information version of the Vickrey auction has arbitrarily poor Nash equilibria. Further, bidding above their valuation is a dominated strategy in all our games.

1.2. Related Work

The classical model of sponsored search auctions has been introduced by [Edelman et al. 2007; Varian 2007]. In this model the Vickrey-Clarke-Groves (VCG) mechanism (e.g., [Nisan 2007]) provides a dominant-strategy implementation of the welfare-maximizing outcome. Current search auctions, however, are based instead on the Generalized Second Price (GSP) mechanism, which shares its rank-by-bid allocation rule with the VCG mechanism but uses a simpler payment rule. [Edelman et al. 2007; Varian 2007] show that the GSP mechanism has a welfare-maximizing equilibrium, but participants do not generally have dominant strategies. In addition to the welfare-maximizing Nash equilibrium, there can also be suboptimal Nash equilibria, even when every player bids at most its value [Aggarwal et al. 2006; Leme and Tardos 2010].

A sequence of recent papers consider the price of anarchy in the standard GSP auction of [Leme and Tardos 2010; Caragiannis et al. 2011; Lucier and Paes Leme 2011; Caragiannis et al. 2012]. There are no previous non-trivial results on the worst-case inefficiency of equilibria in other models of sponsored search, such as the models with externalities studied here.[2] The cascade model that we consider was proposed independently in [Aggarwal et al. 2008] and [Kempe and Mahdian 2008]; several other researchers have also studied models to accommodate externalities in sponsored search [Abrams et al. 2007; Giotis and Karlin 2008] or offer endogenous models of click-through rates [Athey and Nekipelov 2010; Athey and Ellison 2011; Ghosh and Mahdian 2008].

Several recent papers have studied the price of anarchy in simple mechanisms without dominant strategies [Bhawalkar and Roughgarden 2011; Borodin and Lucier 2010; Christodoulou et al. 2008; Hassidim et al. 2011; Johari and Tsitsiklis 2004; Lucier 2010; Lucier and Borodin 2010]. Among these, the closest to our paper in spirit is by [Lucier and Borodin 2010; Lucier 2010], who consider greedy allocation rules and study that additional inefficiency at equilibrium compared to the inefficiency of the allocation rule as an approximation algorithm. The setting in [Lucier and Borodin 2010; Lucier 2010] is greedy allocation rules for welfare maximization in combinatorial auctions. One of the main results in [Lucier and Borodin 2010] is that a c-approximate greedy algorithm, coupled with the natural definition of a critical bid payment rule, induces a mechanism in which the price of anarchy is at most $c + 1$. One drawback of the setting in [Lucier and Borodin 2010; Lucier 2010] is that welfare maximization in combinatorial auctions is an extremely difficult problem; in particular, for most special cases of interest, no greedy algorithm can obtain a reasonable approximation ratio. Another difference that is more technical: in the direct-revelation combinatorial auctions in [Lucier and Borodin 2010; Lucier 2010], a bidder can always "target" a given bundle (i.e., submit a non-zero bid only for that bundle), a fact that facilitates the price of anarchy analysis. In a single-parameter problem with multiple different quantities (like sponsored search), a bidder cannot bid directly on a given quantity; this restricted bidding vocabulary makes the analysis harder.

2. PRELIMINARIES

The classical model of Sponsored Search Auctions introduced by [Edelman et al. 2007; Varian 2007] assumes that a set of n advertisers want to display ads on a Web page of search results. Each advertiser i has a private valuation v_i for a web-surfer clicking on the ad. Advertisements can be displayed in multiple slots on a Web page. In the basic model for sponsored search auctions the bidder receiving the jth slot from the top

[2]Of course, we are not counting applications of the welfare-maximizing but relatively complicated VCG mechanism to such models [Aggarwal et al. 2008; Kempe and Mahdian 2008].

in a ranked list of ads gets a click with probability α_j, where α_j is the click-through rate (CTR) of the slot. We assume, as is natural and standard, that higher ads are better: $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_k$. We consider the simple allocation rule that assigns the ith highest bidder to the ith slot for each $i = 1, 2, \ldots, k$. An easy exchange argument shows that this (monotone) allocation rule maximizes the welfare when players bid their true values. The classical work of [Myerson 1981] yields a payment rule that extends this allocation to a truthful, dominant-strategy mechanism: given bids $b_1 \geq b_2 \geq \cdots \geq b_n$, charge the ith highest bidder a payment of

$$\sum_{j=i+1}^{k+1} (\alpha_{j-1} - \alpha_j) b_j \tag{1}$$

for each $i = 1, 2, \ldots, k$, where α_{k+1} should be interpreted as 0.

The sponsored search auctions used by modern search engines are derived from this "rank-by-bid" allocation rule above, but they use the critical bid payment rule in lieu of the relatively complicated one in (1). The critical bid payment rule corresponding to this allocation charges the ith highest bidder b_{i+1} for each click (and $\alpha_i b_{i+1}$ in total), which is a higher price than in (1). The resulting mechanism does *not* have dominant strategies. Intuitively, depending on others' bids, a player might have an incentive to underbid, with the goal of getting a clicks with slightly smaller probability at a much cheaper price. We will assume throughout the paper that "conservative" bidding strategies are used in that every player bids at most its valuation. We will see that bidding above the valuation is a dominated strategy in our games. Further, when players can "bluff" by bidding more than their value, even the full-information version of the Vickrey auction has arbitrarily poor Nash equilibria. (Consider two bidders with value 1 and 0 who bid 0 and 1, respectively.) With conservative bidding, every Nash equilibrium under payment rule (1) is fully efficient.

Modern search engines distinguish ads also by relevance by assigning each ad i a "quality score" γ_i, and rank ads by sorting bid – quality factor products $\gamma_i b_i$. The results of [Edelman et al. 2007; Varian 2007] also extend to a more general model including quality factors, assuming that the quality factors are publicly known and effect the probability of ad i getting a click as follows: when the ad is displayed in slot j resulting in a click by probability $\alpha_j \gamma_i$. The natural extension of GSP for the problem with quality scores is the critical value payment scheme, charging each player the smallest bid that would allow it to keep its slot in the assignment. More formally, if the bid in slot i was b_i and has quality factor γ_i, then the per-click payment p_i of bidder in slot i is set by the equality

$$p_i \gamma_i = b_{i+1} \gamma_{i+1} \tag{2}$$

for each $i = 1, 2, \ldots, k$, where b_{n+1} should be interpreted as 0.

Our main result concerns the Cascade Model of Sponsored Search Auction introduced independently by [Aggarwal et al. 2008] and [Kempe and Mahdian 2008]. As before, assume a set of n advertisers want to display ads on a Web page of search results. Each advertiser i has a private valuation v_i per click and two public parameters p_i, q_i. Given an assignment of advertisers to k slots, a potential customer scans the ads, beginning at the top. If the ith ad is looked at, then it is clicked on with probability p_i, and the next ad is also looked at with probability q_i. (These two events can be assumed independent, or not; it doesn't matter to any of the computations in this paper.)

[Aggarwal et al. 2008] and [Kempe and Mahdian 2008] show that when $k = n$ and simple greedy algorithm finds the efficient allocation. We can view this greedy algorithm as the sort by bid $\times$ quality factor allocation rule using quality factors

$\gamma_i = p_i/(1 - q_i$. However, with a limited number of slots, maximizing welfare in this model involves dynamic programming. Thus neither the allocation nor payment rule of the VCG mechanism for this problem is as simple as the GSP mechanism. Nonetheless, we will prove that coupling GSP with a suitable definition of quality factors, the corresponding critical bid payment rule (2) yields a mechanism in which every Nash equilibrium has welfare at least a constant times that of the maximum possible.

2.1. Solution Concepts

In all the models we consider, we will think of the allocation resulting from the variant of the GSP mechanism with arbitrary quality factors γ_i as a game. Each player is asked for a bid b_i, then ads get sorted and allocated along the page in the order of $b_i\gamma_i$, and charge per-click is computed using the formula (2). When advertisement with value v_i gets displayed in a slot so that the probability of getting a click is x_i and has to pay p_i for each click, than we view the resulting utility to be $u_i = x_i(v_i - p_i)$. In the full information setting, a pure Nash equilibrium of this game is a bid vector b, such that for each player i and any other bid z for this player $u_i(z, b_{-i}) \leq u_i(b)$. As mentioned in the introduction, we will assume that players are conservative and use bids $b_i \leq v_i$, i.e., they never bid above their valuation. We note that for the class of allocation and payments rules used by GSP, overbidding is dominated strategy.

LEMMA 2.1. *For any quality factors γ_i, bidding above the player's valuation $b_i > v_i$ in the GSP mechanism is dominated by bidding v_i.*

To evaluate the quality of Nash equilibria we will use the Price of Anarchy: the smallest factor μ such that all Nash equilibria (assuming conservative bidding) have welfare at least a $1/\mu$ fraction of the welfare in the efficient allocation.

Finding equilibrium in games may be computationally hard, and even when the equilibrium can be found algorithmically, it can be challenging for the players to coordinate on an equilibrium. In contrast, simple and natural strategies can guarantee that players have no-regret in a sequence of repeated play of the game in the following sense. Suppose the sequence of bid vectors in a repeated play is $b^1, b^2, \ldots$. To evaluate the regret of player i we consider the player's best strategy with hindsight, keeping all other bids b^t_{-i} fixed: $\max_z \sum(u_i(z, b^t_{-i})$. The regret of a player is $\max_z(\sum_t(u_i(z, b_{-i}) - \sum_t u_i(b^t)))$. Let $r_i(T)$ denote the regret of player i after T rounds of play. We say that a play of the game has the no-regret (or vanishing regret) property for player i if $r_i(T)/T$ goes to 0 as T goes to infinity. There are many simple and natural strategies (see, e.g., [Arora et al. 2012; Hart and Mas-Colell 2000] and the citations within these papers) that guarantee a player no-regret for any play by the other players. When all players use such no-regret strategies, the resulting play converges to a coarse correlated equilibrium. Formally, a coarse correlated equilibrium is a probability distribution of bid vectors, such that for all players i $\mathbb{E}(u_i(b)) \geq \mathbb{E}(u_i(z, b_{-i}))$ for all possible bids z, where the expectation is taken over the distribution of bids.

Our bounds for the price of anarchy also extend bounding the price of total anarchy [Blum et al. 2008], which is the smallest factor μ such that the expected value of the social welfare in any coarse correlated equilibrium is at least a $1/\mu$ fraction of the welfare in the efficient allocation.

Finally, we will consider also the Bayesian model when valuations of users are not full information, but rather are drawn from a distribution $\mathcal{V}$, where valuations of different bidders can be arbitrarily correlated. In this case a player's strategy is a (possibly randomized) bidding function $b_i(z)$ that assigns bids to each valuation z. A set of bidding functions is a Bayes-Nash equilibrium if $\mathbb{E}(u_i(b)|v_i = \nu) \geq \mathbb{E}(u_i(z, b_{-i})|v_i = \nu)$ for all values ν of player i and all possible bids z, where the expectation is taken with respect to the valuations of all players other than i, and the randomness in the bidding functions. We define the price of anarchy for this solution concept via expected welfare, the smallest μ such that the expected efficiency at any equilibrium is at least a $1/\mu$ fraction of the expected welfare of the efficient allocation.

Roughgarden [Roughgarden 2009] developed a framework that identifies price of anarchy bounds that naturally extend to some of the above solution concepts. Using the slightly relaxed version from [Lucier and Paes Leme 2011; Caragiannis et al. 2012] we say that a utility game (λ, μ) semi-smooth if for any valuation vector v there is a bid vector b^* so that b^* results in the efficient allocation, and for all other bid vectors b the following inequality holds:

$$\sum_i u_i(b_i^*, b_i) \geq \lambda \sum_i u_i(b^*) - \mu \sum_i u_i(b).$$

In [Roughgarden 2009] it was shown that whenever this inequality holds in a full-information game, the price of total anarchy (i.e., for easy-to-learn outcomes) is bounded by $\frac{\lambda}{1+\mu}$. In [Lucier and Paes Leme 2011; Caragiannis et al. 2012] it was proved, essentially, that if b_i^* is a function only of v_i (and not of v_{-i}), then the price of anarchy of Bayes-Nash equilibria is bounded by $\frac{\lambda}{1+\mu}$ for every distribution over player valuations (even correlated ones).

2.2. The class of games

Our results concern variants of the Cascade Model of Sponsored Search Auctions. Our main result concerns the model with k slots introduced independently by [Kempe and Mahdian 2008] and [Aggarwal et al. 2008].

Example 2.2 (*Cascade Model of Sponsored Search Auctions*). A set of n advertisers want to display ads on a Web page of search results. Each advertiser i has a private valuation v_i per click and two public parameters p_i, q_i. Given an assignment of advertisers to k slots, a potential customer scans the ads, beginning at the top. If the ith ad is looked at, then it is clicked on with probability p_i, and the next ad is also looked at with probability q_i. (These two events can be assumed independent, or not; it doesn't matter to any of the computations in this paper.) Maximizing welfare in this model involves dynamic programming. Thus neither the allocation nor payment rule of the VCG mechanism for this problem is simple in our sense. Nonetheless, we prove that coupling a suitable nonadaptive greedy allocation rule with the corresponding critical bid payment rule yields a mechanism in which every Nash equilibrium has welfare at least a constant times that of the maximum possible.

We start the technical developments by considering a simpler model also introduced by [Kempe and Mahdian 2008] and [Aggarwal et al. 2008] that has unlimited number of slots.

Example 2.3 (*Unbounded Cascade Model of Sponsored Search Auctions*). Each advertiser i has a private valuation v_i per click and two public parameters p_i, q_i. Given an assignment of advertisers to slots, a potential customer scans the ads, beginning at the top. If the ith ad is looked at, then it is clicked on with probability p_i, and, with probability q_i, the next ad is also looked at (where the two probabilities do not have to be independent). [Kempe and Mahdian 2008] and [Aggarwal et al. 2008] show that a simple greedy algorithm can be used to maximize the welfare of the assignment: sort advertisers in nondecreasing order of $\frac{v_i p_i}{1-q_i}$.

We can think of the mechanism based on this greedy algorithm as the classical GSP using a quality factor $\gamma_i = \frac{p_i}{1-q_i}$. In the Appendix we further extend this analogy by showing that the full information game with quality factors γ_i has a Nash equilibrium that implements the VCG outcome (like in the standard GSP model [Edelman et al. 2007; Varian 2007]).

Our first technical result will be to bound the price of anarchy of the above greedy mechanism. We will do this in the more general context of allowing the greedy algorithm to use quality factors γ_i smaller than what is needed in the optimal greedy algorithm. Such a modified greedy algorithm may not result in (even approximately) efficient allocation. We compare the welfare of the Nash equilibrium with the welfare of the greedy outcome with the true valuations in the spirit of the high-level outline suggested by (Q3) in the introduction. This flexibility allows us to extend the result to variants of the model where the welfare-maximizing allocation rule is not simple, like in the Cascade Model with limited number of slots, as well as a few other models defined below where algorithms that find a welfare-maximizing allocation are not known.

Example 2.4 (*Cascade Model with Periodic Gaps*). Here is a model that combines features of the unbounded cascade model, and the limited slot model. The model has two extra parameters k and $\delta \in (0,1)$. The integer k models the number of ads on a single screen, and δ models the decreased probability of looking at ads below the current screen. Formally, this model is identical to the unbounded cascade model with one exception. For slots i that are integer multiples of k, if an ad is placed in this slot, the continuation probability is only $q_i\delta$.

Example 2.5 (*Unbounded Cascade Mechanisms with Reserve Price*). Instead of limiting the ads shown by having only limited number of slots, we can add a reserve price r. The mechanism discards all bids below the reserve price, runs on the remaining ads, and charges all players at least the reserve price, charging the critical bid of the modified mechanism. Reserve price limits the participants to only those with value $v_i \geq r$, but is also limiting the strategy space of the players, as bids below r are discarded. Our bounds also extend to the Unbounded Cascade Mechanisms with Reserve Price.

For simplicity of presentation we will assume throughout the paper that $p_j = 1$ for all j. It is easy to adopt all results to the case with general p_j by using the "impression value" $v'_j = p_j v_j$ everywhere in place of the value v_j.

3. GREEDY ALGORITHM AND THE PRICE OF ANARCHY

The first goal of this section is to analyze the price of anarchy of a class of greedy mechanism in the setting of Example 2.3 with unlimited slots. As mentioned earlier, the greedy algorithm with "quality factors" $\gamma_i = (1-q_i)^{-1}$, sorting the ads in decreasing order of $v_i\gamma_i$ results in the efficient allocation. We will consider a class of mechanisms that greedily sorts bids in decreasing order of $b_i\gamma_i$ for a "quality factor" $\gamma_i \leq (1-q_i)^{-1}$. Following the high-level outline of the introduction, we compare the resulting welfare of an equilibrium to the welfare of the greedy outcome with ads modified to have $q'_i \leq q_i$ such that $\gamma_i = (1-q'_i)^{-1}$. We show that if the bidders are conservative in the sense that $b_i \leq v_i$ for all i, then the welfare of the resulting equilibrium of the mechanism is at most a constant factor worse than the welfare of the greedy outcome with the modified q'_i values. This flexibility in choosing quality factors makes our analysis more flexible, and allows us to use it in models where such greedy algorithm isn't optimal.

As a first step, we will consider the social welfare generated by any ordering of the ads, and express the resulting welfare in two different forms, as a convex combination of values, and as an integral. For expressing values as a convex combination, we will associate a value of an ad with the event that this ad is looked at and then the process stops. Consider ads displayed in order $i_1, i_2, \ldots$. Ad i_k will now get clicked on with probability $\prod_{j<k} q_{i_j}$, the probability of continuing after every ad in earlier slots. The ad contributes $v_{i_k} \prod_{j<k} q_{i_j}$ to social welfare. We want to associate this welfare with the event that i_k is the last ad looked at. This event has probability $s_k = (1 - q_{i_k}) \prod_{j<k} q_{i_j}$, the probability of the ad getting a click and then not continuing. To make the treatment uniform for all slots, including the last one, we add $s_{n+1} = \prod_{j \le n} q_{i_j}$. With this definition, $(s_1, s_2, \ldots, s_{n+1})$ is a probability distribution and if we define a modified value of an ad i to be $V_i = \frac{v_i}{1-q_i}$, we can express the social value of the resulting allocation as a convex combination of the modified values as follows.

LEMMA 3.1. *Given a sorted order of ads $i_1, i_2, \ldots$, the expected welfare $\sum_k v_{i_k} \prod_{j<k} q_{i_j}$ is the same as the expected modified value of the last ad looked at, which is expressed as $\sum_{k \le n} s_k V_{i_k}$.*

Next we will express the above social welfare also as an integral over the interval $[0, 1]$. This intergal will be helpful in comparting the social value generated by two different orderings. Assume as before that ads are ordered as i_1, i_2, etc, and recall the values s_k defined above, denoting the probability of the ad in slot k is the last to get clicked on. Intuitively, we want to think of the integral of an initial segment $[0, x]$ as expressing the value of an initial segment of the ads: the segment $[0, s_1]$ the value generated by ad i_1, the next segment $[s_1, s_1 + s_2]$ the value generated by ad i_2, etc. To do this we define a function $f(x)$ as a piecewise constant function setting $f(x) = V_{i_k}$ for x in $(\sum_{j<k} s_j, \sum_{j \le k} s_j]$, an interval of length s_k, and setting $f(x) = 0$ for $x \in (1 - s_{n+1}, 1]$. We note that $\sum_{j<k} s_j$ is the probability that we stop the process before getting to the kth slot, so $x_k = 1 - \sum_{j<k} s_j$ is the probability of the ad in the kth slot getting clicked on. The integral of the segment of f with value V_{i_k} is exactly ad i_kth contribution to social welfare. Summing over ads we get the following

LEMMA 3.2. *Given a sorted order of ads $i_1, i_2, \ldots$, the expected welfare resulting from displaying ads in this order is expressed as $\int_0^1 f(x)dx$.*

To facilitate evaluating the resulting welfare with different quality factors we will also consider the welfare with the decreased continuation probabilities q'_j. Let f' denote the analogous function defined for the same ordering using the probabilities q' to define modified V'_i, and use the fact that $\int_0^1 f'(x)dx = \sum_k v_{i_k} \prod_{j<k} q'_{i_j}$.

Now consider the Nash equilibrium of the game defined via the quality scores and critical value pricing. Assume that f and f' are the functions defined above associated with a pure Nash equilibrium of the above game. We are going to compare the quality of this equilibrium to the welfare resulting from a greedy algorithm that has access to the true values, but sorts by $\gamma_i v_i$, and is using modified continuation probabilities q'_i such that $\gamma_i = (1 - q'_i)^{-1}$ to evaluate welfare. With the modified continuation probabilities and the true values this greedy algorithm results in the efficient allocation, but due to using smaller continuation probabilities, the allocation may not be efficient with the real continuation probabilities.

It will be useful to also express the value resulting from the above greedy order via an integral. Let $g'(x)$ be the analogous function defined using the ads in decreasing order of $v_i\gamma_i$ values, and continuation probabilities q'_j. Note that the modified value used for the interval associated with ad j in function g' is $v_i\gamma_i$, and hence the function g' is monotone decreasing throughout the interval $[0, 1]$. The integral $\int_0^1 g'(x)dx$ expresses the welfare of the outcome of the greedy algorithm with continuation probabilities q'.

THEOREM 3.3. *For functions f, f' and g' defined above $\int_0^1 g'(x)dx \leq 2(\int_0^1 f(x)dx + \int_0^1 f'(x)dx)$. The left hand side is the social value of the greedy order with continuation probabilities q' while the two integrals on the right hand side each express the social value of the Nash equilibrium with continuation probabilities q and q' respectively.*

Before proceeding to the proof of the theorem, we state the corollary summarizing the result. The same sort has higher value with higher continuation probabilities, and hence $\int_0^1 f(x)dx \geq \int_0^1 f'(x)dx$.

COROLLARY 3.4. *The social welfare of the mechanism using $\gamma_i \leq (1-q_i)^{-1}$ at Nash equilibrium is at least a $1/4th$ fraction of the welfare of the greedy outcome.*

PROOF. Assume ads are numbered $1, 2, \ldots$ in decreasing order of $v_i\gamma_i$. Consider an ad i in the greedy order, and let x_i^g denote the probability that ad i gets clicked on in the greedy order with continuation probability q_i', and $s_i^g = x_i^g(1-q_i')$ the probability that this is the last ad clicked on. The interval associated with ad i in the function $g'(.)$ is $[1-x_i^g, 1-x_i^g+s_i^g]$ and $g'(x) = v_i\gamma_i$ in this interval.

The proof is analogous the the price of anarchy proof of [Lucier and Paes Leme 2011; Caragiannis et al. 2011, 2012] considering for each ad i the value it can get by a possible deviation from Nash by bidding $b_i = v_i/2$. If this alternate bid b_i' gets i to a slot where the ad would get viewed with probability y with the modified probabilities (which is the product of the continuation probabilities q' above the new position of the ad), then the deviation would result in a value of at least $\frac{1}{2}v_iy$ for advertiser i, as the real probability of getting clicked on is even higher, and the price charged cannot exceed the bid. This implies that i must derive this much utility also at the equilibrium resulting in $u_i(N) \geq \frac{1}{2}v_iy$. We consider a few cases.

(a) If $y \geq x_i^g$, then the ad must get a value $u_i(N)$ of at least $\frac{1}{2}v_iy \geq \frac{1}{2}v_ix_i^g$ also in its position in Nash, which is $1/2$ the value derived by advertiser i in the greedy outcome.

Next consider the integral expressing the social value of Nash. For any $x < 1-y$ the function value $f'(x)$ is defined by an ad j that is in an earlier slot than where bid $\frac{1}{2}v_i$ takes i, and hence $b_j\gamma_j \geq \frac{1}{2}v_i\gamma_i$. Due to the conservative assumption it must also be the case that $v_j\gamma_j \geq b_j\gamma_j \geq \frac{1}{2}v_i\gamma_i$. The function value $f'(x) = v_j\gamma_j$, so we get $f'(x) \geq \frac{1}{2}v_i\gamma_i$.

(b) If $y \leq x_{i+1}^g$, smaller than the probability of the next ad getting clicked on in the greedy outcome, than we get $f(x) \geq \frac{1}{2}v_i\gamma_i$ throughout the interval $[1-x_i^g, 1-x_{i+1}^g]$ that corresponds to ad i in the integral of g', and so $\int_{1-x_i^g}^{1-x_{i+1}^g} f'(x)dx \geq \frac{1}{2}v_i\prod_{j<i}q_j' = \frac{1}{2}v_ix_i^g$.

Combining these two bounds we get the following in both cases.

$$u_i(N) + \int_{1-x_i^g}^{1-x_{i+1}^g} f'(x)dx \geq \frac{1}{2}v_ix_i^g.$$

Finally, we claim that this bound also holds when $x_{i+1}^g < y < x_i^g$. The ad i derives value $u_i(N) \geq \frac{1}{2}v_iy$, while the contribution of the integral is at least $\frac{1}{2}(x_i^g - y)v_i\gamma_i$. Summing these two we get at least $\frac{1}{2}v_ix_i^g$, as claimed.

Summing over all advertisers, and using that social welfare at Nash can be expressed at $\int_0^1 f(x)dx \geq \int_0^1 f'(x)dx$, and is at least the sum of advertisers utilities we get the claimed theorem.

□

Note that the proof used only one property of the outcome that bidder i has utility at least as much as his deviation to the bid $b'_i = \frac{1}{2}v_i$. Using the terminology from [Lucier and Paes Leme 2011; Caragiannis et al. 2012] we showed that for the Unbounded Cascade model the GSP game using $\gamma_i = (1-q_i)^{-1}$ is $(1/2, 1)$-semi-smooth. This allows up to bound the outcome quality even in the Bayesian setting when types are drawn from a possibly correlated distribution, and they are valid for learning outcomes. We assume that the continuation probabilities q_i are known and are used in the mechanism, only valuations are private.

THEOREM 3.5. *Consider the Unbounded Cascade Model using GSP mechanism with $\gamma_i = (1-q_i)^{-1}$ in the Bayesian setting when player types are drawn from a distribution $\mathcal{F}$ that is possibly arbitrarily correlated. The expected social welfare in outcomes when players have no regret about one particular alternate bid $b'_i = v_i/2$ is at least $1/4$th of the maximum possible social welfare. Implying that the price of total anarchy for this GSP game in the Bayesian Unbounded Cascade Model setting with correlated bids is bounded by 4.*

Next we derive similar bounds for the other models introduced. A simple mechanism the Example 2.4 in section 2 modeling page breaks follows immediately.

To deal with the occasional δ decrease of the probability, we will use $q'_i = q_i\delta^{1/k}$, and quality factors $\gamma_i = (1-q'_i)^{-1}$. The difference in the click-through rates with continuation probabilities q' and q differs by at most a factor of δ at any slot, and hence the greedy algorithm that is optimal for q' is at most a factor of δ off from the true optimum. Theorem 3.5 immediately gives us the following.

COROLLARY 3.6. *Consider the Unbounded Cascade Model with periodic gaps using GSP mechanism with $\gamma_i = (1-\delta^{1/k}q_i)^{-1}$ in the Bayesian setting when player types are drawn from a distribution $\mathcal{F}$ that is possibly arbitrarily correlated. The expected social welfare in outcomes when players have no regret about one particular alternate bid $b'_i = v_i/2$ is at least $\delta/4$th of the maximum possible social welfare. Implying that the price of total anarchy for this GSP game in the Bayesian Unbounded Cascade Model with periodic gaps in setting with correlated bids is bounded by $4/\delta$.*

Next consider the mechanism of Example 2.5 Unbounded Cascade Model with reserve price r. In presence of the reserve price, the mechanism allocates only bidders with $b_i \geq r$, and changes participant at least the reserve price r. Slightly modifying the proof to consider alternate bids $(v_i + r)/2$ we get the following.

THEOREM 3.7. *Consider the Unbounded Cascade Model using GSP mechanism with reserve price r and with $\gamma_i = (1-q_i)^{-1}$ in the Bayesian setting when player types are drawn from a distribution $\mathcal{F}$ that is possibly arbitrarily correlated. The expected social welfare in outcomes when players with values $v_i \geq r$ have no regret about one particular alternate bid $b'_i = (v_i + r)/2$ is at least $1/4$th of the maximum possible social welfare. Implying that the price of total anarchy for this GSP game in the Bayesian Unbounded Cascade Model with reserve price r in setting with correlated bids is bounded by 4.*

PROOF. The proof follows closely the the proof on Theorem 3.3, but we use $\gamma_i = (1-q_i)^{-1}$, so won't distinguish f and f', or g and g'. For an ad with value v_i we consider deviating bid $b'_i = (v_i + r)/2$. Let probability y the click probability resulting from this deviation from Nash, as in the proof on Theorem 3.3. We get that $u_i(N) \geq (v_i - b'_i)y = \frac{1}{2}(v_i - r)y$. For $x < (1-y)$, the value $f(x)$ corresponds to an ad j with $v_j\gamma_j > b'_i\gamma_i$, which implies that $f(x) \geq b'_i\gamma_i = \frac{v_i+r}{2}\gamma_i$ and hence $f'(x) - r\gamma_i \geq \frac{1}{2}(v_i - r)\gamma_i$. We combine these, as was done in the proof on Theorem 3.3 using the notation x_i^g we get

$$u_i(N) + \int_{1-x_i^g}^{1-x_{i+1}^g} f(x)dx \geq \frac{1}{2}(v_i - r)\prod_j q_j.$$

Adding these bounds for all players gives us a bound on total utility of players at Nash. To get the social welfare we can add the revenue, which is at least r for every click, and as we have unlimited slots the number of clicks is $1-x$ for $x = \prod_j q_j$ independent of the order the ads. We get the following:

$$\sum_i u_i(N) + r(1-x) + \int_0^{1-x} f(x)dx \geq \frac{1}{2}\sum_i (v_i - r)\prod_j q'_j + (1-x)r \geq \frac{1}{2}\int_0^1 g(x)dx.$$

The first two terms together are at most social welfare at Nash, so the left hand side is at most twice the social welfare, while the right hand side is the welfare of the efficient outcome, proving the bound of 4 on the price of anarchy. □

Finally, we prove the first result mentioned in the introduction, a greedy mechanism for the Example 2.2 of the Cascade Model with k slots. The case $k = 1$ is the classic Vickrey auction, so we will assume here $k > 1$. To guarantee good approximation quality with the true valuations, we will use quality factors $\gamma_i = \min((1-q'_i)^{-1}$ for $q'_i = \min(q_i, 1-1/k)$.

LEMMA 3.8. *The quality of the greedy algorithm that assigns ads to k slots in sorted order of values $v_i\gamma_i$ is an $(1-1/k)^k(1-(1-1/k)^k)$ at least an 8th fraction, and $\approx e^{-1}(1-e^{-1}) \approx 4.3$ approximation algorithm for large k.*

PROOF. With only k slots, the difference in click probabilities using continuation probabilities q or q' is at most a factor $(1-\frac{1}{k})^k$. Recall the integral $\int_0^1 g'(x)dx$ expressing the welfare of the greedy sort with continuation probabilities q'. Let x_{k+1}^g denote product of the first k continuation probabilities q' in this order. The welfare with continuation probabilities q' generated by the first k ads is expressed by $\int_0^{1-x_{k+1}^g} g'(x)dx$. Since $q'_i \leq 1-1/k$ we get $x_{k+1}^g < (1-1/k)^k$. The function g' is monotone decreasing, so the integral is at least a $1-x_{k+1}^g$ fraction of the total value, and $1-x_{k+1}^g \geq 1-(1-1/k)^k$.
Combining these two bounds the social welfare of the first k ads in the sorted order of $v_i\gamma_i$ is at least a $(1-1/k)^k(1-(1-1/k)^k)$ fraction of the maximum possible welfare.

□

Combining the approximation result with our mechanism we get the main theorem.

THEOREM 3.9. *Consider the Cascade Model using GSP mechanism with k slots and with $\gamma_i = (1-q'_i)^{-1}$ as defined above in the Bayesian setting when player types are drawn from a distribution $\mathcal{F}$ that is possibly arbitrarily correlated. The expected social welfare in outcomes when players with values v_i have no regret about one particular alternate bid $b'_i = v_i/2$ is at least a $\frac{1}{4}(1-1/k)^k(1-(1-\frac{1}{k}))$ fraction of the maximum possible social welfare. Implying that the price of total anarchy for this GSP game in the Bayesian Unbounded Cascade Model with reserve price r in setting with correlated bids is bounded by $4(1-1/k)^{-k}(1-(1-\frac{1}{k}))^{-1}$.*

PROOF. Following the high-level outline of the introduction, we use the approximation bound and the proof of Theorem 3.3 to bound the price of anarchy. As usual, assume ads are numbered in decreasing order of $v_i\gamma_i$. Our goal is to compare the greedy outcome with the Nash equilibrium.

As in the previous proof consider the deviating bid $b'_i = v_i/2$ for an ad i as before. If this deviating bid gives the bidder one of the first k slots, we can use the argument in the proof above to show that

$$u_i(N) + \int_{1-x_i^g}^{1-x_{i+1}^g} f'(x)dx \geq \frac{1}{2} v_i \prod_{j<i} q'_j.$$

Let i be the first ad in the greedy order, where such a deviating bid doesn't result in ad i getting one of the top k slots. Adding the above bounds for ads before i, we get that

$$\sum_{j<i} u_i(N) + \int_0^{1-x_i^g} f'(x)dx \geq \frac{1}{2} \int_0^{1-x_i^g} g'(x)dx.$$

Now consider ad i, and let y be the product of the q' probabilities of the k ads placed in Nash. If the deviation $b'_i = v_i/2$ doesn't get the ad i one of the top k slots, it must be the case that the product $b_j\gamma_j$ for all ads placed in the Nash outcome have $v_j\gamma_j \geq b_j\gamma_j \geq v_i\gamma_i/2$, which implies that for all $x \leq y$ associated with ads in the top k slots have $f'(x) \geq \frac{1}{2}v_i\gamma_i$.

$$\sum_{j<i} u_i(N) + \int_0^{1-y} f'(x)dx \geq \frac{1}{2} \int_0^{1-y} g'(x)dx.$$

The left hand side is at at most twice the social welfare at Nash, while the right hand side is closely related to the value of the greedy outcome. The integral $\int_0^1 g'(x)dx$ is at least a $(1-1/k)^k$ fraction of the optimal social welfare, and taking the integral only till $1-y$ possibly looses an additional factor of $(1-(1-1/k)^k)$ as $y \leq (1-1/k)^k$. This shows that the value of the Nash equilibrium is at least a $\frac{1}{4}(1-1/k)^k(1-(1-\frac{1}{k}))$ fraction of the maximum possible.

□

REFERENCES

ABRAMS, Z., GHOSH, A., AND VEE, E. 2007. Cost of conciseness in sponsored search auctions. In *The 3rd International Workshop on Internet and Network Economies (WINE)*. 326–334.

AGGARWAL, G., FELDMAN, J., MUTHUKRISHNAN, S., AND PÁL, M. 2008. Sponsored search auctions with Markovian users. In *The 4th Workshop on Ad Auctions*.

AGGARWAL, G., GOEL, A., AND MOTWANI, R. 2006. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce (EC)*. 1–7.

ARORA, S., HAZAN, E., AND KALE, S. 2012. The multiplicative weights update method: a meta-algorithm and some applications. *Theory of Computing*. To appear.

ATHEY, S. AND ELLISON, G. 2011. Position auctions with consumer search. *Quarterly Journal of Economics 126,* 3, 1319–1374.

ATHEY, S. AND NEKIPELOV, D. 2010. A structural model of sponsored search advertising auctions. In *The 6th Workshop on Ad Auctions*.

AUSUBEL, L. M. AND MILGROM, P. 2006. The lovely but lonely vickrey auction. In *Combinatorial Auctions. P. Cramton, Y. Shoham, R. Steinberg (eds.), Chapter 1*. MIT Press.

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

BLUM, A., HAJIAGHAYI, M., LIGETT, K., AND ROTH, A. 2008. Regret minimization and the price of total anarchy. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*. 373–382.

BORODIN, A. AND LUCIER, B. 2010. On the limitations of greedy mechanism design for truthful combinatorial auctions. In *37th International Colloqium on Automata, Languages and Programming (ICALP)*. 90–101.

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

CARAGIANNIS, I., KAKLAMANIS, C., KANELLOPOULOS, P., KYROPOULOU, M., LUCIER, B., LEME, R. P., AND ÉVA TARDOS. 2012. On the efficiency of equilibria in generalized second price auctions. Submitted.

CHRISTODOULOU, G., KOVÁCS, A., AND SCHAPIRA, M. 2008. Bayesian combinatorial auctions. In *35th International Colloqium on Automata, Languages and Programming (ICALP)*. 820–832.

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

EDELMAN, B., OSTROVSKY, M., AND SCHWARZ, M. 2007. Internet advertising and the Generalized Second-Price Auction: Selling billions of dollars worth of keywords. *American Economic Review 97,* 1, 242–259.

GHOSH, A. AND MAHDIAN, M. 2008. Externalities in online advertising. In *Proceedings of the 17th International World Wide Web Conference (WWW)*. 161–168.

GIOTIS, I. AND KARLIN, A. 2008. On the equilibria and efficiency of the GSP mechanism in keyword auctions with externalities. In *The 4th International Workshop on Internet and Network Economies (WINE)*. 629–638.

HART, S. AND MAS-COLELL, A. 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica 68,* 5, 1127–1150.

HARTLINE, J. D. AND ROUGHGARDEN, T. 2009. Simple versus optimal mechanisms. In *10th ACM Conference on Electronic Commerce (EC)*. 225–234.

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

JOHARI, R. AND TSITSIKLIS, J. N. 2004. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research 29,* 3, 407–435.

KEMPE, D. AND MAHDIAN, M. 2008. A cascade model for externalities in sponsored search. In *The 4th International Workshop on Internet and Network Economies (WINE)*. 585–596.

LEME, R. P. AND TARDOS, E. 2010. Pure and Bayes-Nash price of anarchy for generalized second price auction. In *51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 735–744.

LUCIER, B. 2010. Beyond equilibria: Mechanisms for repeated combinatorial auctions. In *Proceedings of the Innovations in Computer Science Conference (ICS)*. 166–177.

LUCIER, B. AND BORODIN, A. 2010. Price of anarchy for greedy auctions. In *21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 537–553.

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

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

MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research 6*, 58–73.
NISAN, N. 2007. Introduction to mechanism design (for computer scientists). In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, Eds. Cambridge University Press, Chapter 9, 209–241.
NISAN, N. AND RONEN, A. 2001. Algorithmic mechanism design. *Games and Economic Behavior 35,* 1/2, 166–196.
ROTHKOPF, M. H. 2007. Thirteen reasons why the Vickrey-Clarke-Groves process is not practical. *Mathematics of Operations Research 55,* 2, 191–197.
ROUGHGARDEN, T. Algorithmic game theory: Some greatest hits and future directions. In *TCS '08*. 21–42.
ROUGHGARDEN, T. 2009. Intrinsic robustness of the price of anarchy. In *41st ACM Symposium on Theory of Computing (STOC)*. 513–522.
VARIAN, H. R. 2007. Position auctions. *International Journal of Industrial Organization 25,* 6, 1163–1178.

Appendix: Price of Stability

Here we show that in the Unbounded Cascade model of Example 2.3, if the mechanism is using quality factors $\gamma = (1-q_i)^{-1}$ (so the sort is optimal with the true values), VCG gives rise to a Nash equilibrium, analogous to the results of [Edelman et al. 2007; Varian 2007] in the case of the classical GSP model, and hence the price of Stability is 1.

THEOREM 3.10. *In the Unbounded Cascade model in the the full information setting the GSP mechansim with quality factors $\gamma_i = (1-q_i)^{-1}$ has a Nash equilibrium that implements the VCG outcome, and hence is socially optimal.*

PROOF. Assume that ads are numbered in decreasing order of $v_i\gamma_i$. For user i the total VGG payment is

$$P_i = \sum_{j>i} v_j(\prod_{k<j} q_k - \prod_{k<j,k\neq i} q_k) = \prod_{k<i} q_k[\sum_{j>i} v_j(1-q_i) \prod_{i<k<j} q_k]$$

Advertiser i gets $\prod_{k<i} q_k$ clicks, so to pay the price P_i overall, the per click price needs to be

$$p_i = \sum_{j>i} v_j(1-q_i) \prod_{i<k<j} q_k.$$

To get this per-click price to be the critical price, the advertiser $i+1$ has to bid to make $p_i/(1-q_i) = b_{i+1}/(1-q_{i+1})$, which makes

$$b_{i+1} = (1-p_{i+1}) \sum_{j>i} v_j \prod_{i<k<j} q_k.$$

We first claim that $b_i \leq v_i$. This is true as by the greedy order of numbering, we know that $v_i(1-q_j) \geq v_j(1-q_i)$ for all $j \geq i$. Using this with $i+1 > j$ we get that $b_{i+1} \leq \sum_{j>i} v_{i+1}(1-p_j) \prod_{i<k<j} q_k = v_{i+1}(1-\prod_{i<k} q_k)$.

Next we first claim that these bids will cause the optimal sort to happen, that is, we claim that $m_i = b_i\gamma_i$ is monotone. We can write m_i as

$$m_i = \frac{b_i}{1-q_i} = \sum_{j\geq i} v_j \prod_{i\leq k<j} q_k = v_i + q_i \sum_{j>i} v_j \prod_{i<k<j} q_k$$
$$= \frac{v_i}{(1-q_i)}(1-q_i) + q_i m_{i+1}.$$

Recall that $v_i \geq b_i$, hence $v_i/(1-q_i) \geq m_i$, and so m_i is a convex combination of m_{i+1} and a larger term $v_i/(1-q_i)$.

Finally, we want to also claim that the bids b_i form a Nash equilibrium. Consider an advertiser i. Deviating to a smaller bid has the same effect as claiming a smaller value in VCG, as the per-click VCG payments for bidder i depend on the bidders under the slot and $(1-q_i)$. VCG is truthful, and hence bidding to get a slot with lower click probability cannot be beneficial.

We need to verify that it is not beneficial to bid to get a higher click probability slot, as the payment resulting from a deviation is now dependent on the current per-click VCG payment in a higher slot, which depended on bidder i also. To do this, we will show that any ad placed in a later slot prefers slot $i+1$ to slot i. Applying this to all ad j and all $i<j$ shows that j doesn't want to take any earlier slot. Recall that such an ad must have $v\gamma < v_i\gamma_i$. Compared to slot $i-1$, slot i has a q_i smaller click-through rate. To prove that slot i is preferred, we need to argue that the per-click value increases by at least this factor.

The per-click price that advertiser placed lower has to pay for the new slot i is c such that $c/(1-q) = b_i/(1-q_i) = m_i$. Then we get, $(v-c)/(1-q) = v/(1-q) - m_i$. In slot $i+1$ we get the same expression with m_{i+1} or m_{i+2} depending of the ad was originally in slot i or even lower. The move from $i+1$ to i changes the scaled per-click benefit from $v/(1-q) - m_{i+1}$ (or more if the ad was in slot $i+1$) to $v/(1-q) - m_i$. Recall the relation

$$m_i = \frac{v_i}{(1-q_i)}(1-q_i) + q_i m_{i+1}$$

The scaled per-click benefit $v/(1-q) - m_i$ is bounded as

$$v/(1-q) - m_i = v/(1-q) - \frac{v_i}{(1-q_i)}(1-q_i) - q_i m_{i+1} \leq q_i(v/(1-q) - m_{i+1})$$

showing that for such bidders slot $i+1$ is no worse (and typically better) then slot i, and hence deviation to an earlier slot is not beneficial. □