

Greedy Bidding Strategies for Keyword Auctions

Matthew Cary
University of Washington

Ioannis Giotis
University of Washington

Claire Mathieu
Brown University

Aparna Das
Brown University

Kurtis Heimerl
University of Washington

Michael Schwarz
Yahoo! Research

Ben Edelman
Harvard University

Anna R. Karlin
University of Washington

ABSTRACT

How should players bid in keyword auctions such as those used by Google, Yahoo! and MSN? We consider greedy bidding strategies for a repeated auction on a single keyword, where in each round, each player chooses some optimal bid for the next round, assuming that the other players merely repeat their previous bid. We study the revenue, convergence and robustness properties of such strategies. Most interesting among these is a strategy we call the *balanced bidding strategy* (BB): it is known that BB has a unique fixed point with payments identical to those of the VCG mechanism. We show that if all players use the BB strategy and update each round, BB converges when the number of slots is at most 2, but does not always converge for 3 or more slots. On the other hand, we present a simple variant which is guaranteed to converge to the same fixed point for any number of slots. In a model in which only one randomly chosen player updates each round according to the BB strategy, we prove that convergence occurs with probability 1. We complement our theoretical results with empirical studies.

Categories and Subject Descriptors

H.4.0 [Information Systems Applications]: General; G.0 [Mathematics of Computing]: General

General Terms

Algorithms, economics, theory

Keywords

Keyword auctions, bidding strategies

1. INTRODUCTION

Online search engine advertising is an appealing approach to highly targeted advertising, and is the major source of revenue for modern web search engines such as Google and

Yahoo!. The basic setup is as follows: when a person performs a query at a search engine, he receives a page of results that contains the links the search engine has deemed relevant to the search, together with *sponsored links*, i.e., paid advertisements. From the advertiser's perspective, this approach offers a great benefit: they can precisely target their ads based on users' search terms. For example, if a travel agent "buys" the search term "Tahiti," then when searching on the word "Tahiti", a user might be shown a link to the web page for that travel agent offering, say, plane tickets to Tahiti. If the user actually clicks on this link, he will be transferred to the aforementioned web page. For each such click, in which the advertiser receives a potential customer, the advertiser pays the search engine.

The process of determining which ads get assigned to which keywords (search terms) and how much each advertiser pays is resolved via *keyword auctions*. Advertisers choose which keywords they want to bid on and participate in auctions for those keywords. For a keyword of interest, each advertiser submits a bid stating the maximum amount it is willing to pay for a click. When a user submits a query on that keyword, an instantaneous auction is run to determine which of the advertisers currently bidding on that keyword are allocated advertising slots.

In the absence of budget constraints, there is one and only one *truthful* auction that can in principle be used: the Vickrey-Clarke-Groves or VCG mechanism [12, 3, 7]. This mechanism has the property that it is in the best interest of the participating advertisers to bid their true valuation of a click. Despite this appealing property of the VCG mechanism, no search engine uses the VCG mechanism. Rather, the most widely used auction is the non-truthful *Generalized Second Price* or GSP auction (described in Section 2).

The fact that the GSP mechanism is not truthful means that the participating advertisers are forced to undertake the complicated task of choosing a bidding strategy. Asdemir [1] and Edelman et al.[4] observe that instability and bidding wars can result from the use of the GSP mechanism. To make matters worse, on a typical search page, there is room for multiple sponsored links. The positioning of these sponsored links affects the chances that a sponsored link will be clicked on and thus these advertising slots have varying desirability from the perspective of advertisers. This makes the advertisers' utilities a discontinuous function of their bids. Overall, the resulting bidding is sufficiently complex that many advertisers hire consultants or intermediaries to do their bidding for them, often at significant cost.

In a typical day, an advertiser will choose one or more

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'07, June 11–15, 2007, San Diego, California, USA.

Copyright 2007 ACM 978-1-59593-653-0/07/0006 ...\$5.00.

search terms to target and compete in a potentially large number of keyword auctions for those terms. At the same time, advertisers have the opportunity to update their bids. This is typically done automatically via software robots [9].

In this paper, we undertake a systematic exploration of a natural class of greedy bidding strategies a software robot might use in a repeated keyword auction for a particular search term(s). The main theme is the following: If a particular advertiser A knew how the other bidders were going to bid in the next round, A would bid so as to maximize his utility. The greedy bidding strategies we study assume that the recent past is the best prediction for the future: A assumes that the other bidders will bid in the next round exactly what they bid in the current round. Given this assumption, A chooses his bid for the next round to maximize his utility relative to this postulated set of bids by the other bidders.

There is still substantial flexibility within this definition. Suppose that, given the presumed bids by the other bidders, the optimal advertising slot for A to target is slot s at price p_s , and that the price of the (higher click-through rate) slot $s - 1$ is p_{s-1} . The GSP mechanism allows a range of bid values that will result in the same outcome from A 's perspective (namely, any bid between p_s and p_{s-1}). We focus the bulk of our study on one particular choice in this range: *Balanced Bidding* (BB), in which the advertiser chooses his next bid b so as to be indifferent between successfully winning the targeted slot s at price p_s , or winning a slightly more desirable slot at price b . BB is a particularly interesting strategy, since a system in which all players bid according to BB has a unique fixed point in which players are bidding according to a Nash equilibrium of GSP and the payments to the search engine are identical to those of the VCG mechanism. (See Theorem 5.)

Our main results about BB are the following:

- For two slots, BB always converges to its unique fixed point (Theorem 6 part 1).
- For three or more slots, BB need not ever converge, assuming that in each round, *all* players simultaneously update their bids according to the BB greedy strategy (Theorem 6 part 2).
- In an asynchronous model, where exactly one randomly chosen player updates his bid each round according to the BB strategy, BB bidding always converges eventually (Theorem 6 part 4). This is not true if the bidders are not chosen randomly.
- We present a simple variation on BB with the same unique fixed point and prove that it converges to its fixed point even in the synchronous model in which all players update their bids according to BB at each step (Theorem 10).

An important perspective on these results is the following: The Nash equilibria of GSP are fairly straightforward to characterize and have been understood for some time [5, 11, 10]. What was not known was if there was a natural bidding strategy that would lead to these Nash equilibria. Here we show how players can get to the most natural of the GSP equilibria using a greedy strategy.

In Section 5, we explore a number of relevant issues. We begin by experimentally studying the BB strategy, including comparing the convergence of BB in synchronous and

asynchronous settings. We next compare the VCG revenue to what would be obtained if the bidders used other natural greedy strategies: *competitor busting* (CB), in which the advertiser chooses the highest bid value consistent with the target slot, and *altruistic bidding* (AB), in which the advertiser chooses the lowest bid value consistent with the target slot. Finally, we explore the revenue obtained by the search engines under GSP compared with VCG, assuming that bidders eventually end up in some Nash equilibrium of GSP (which of course is not necessarily the case). We complement these experiments with some theoretical bounds (Theorem 16).

2. MODEL AND DEFINITIONS

Definition 1. A keyword auction is defined by:

- A set of k slots with click-through rates (CTRs) $\theta_1 > \dots > \theta_k$, where θ_i is the probability that the user will click on the advertisement in slot i .
- A set of n players (advertisers) participating in the auction, each one having a private valuation v_i for a click, $v_1 > \dots > v_n$.
- Based on knowledge of the mechanism and their own private valuations, each player submits a bid to the auction. We denote by b_i the bid submitted by player i .
- The mechanism:
 - computes an allocation π of the slots to k different players: π_s is the identity of the player that is allocated slot s .
 - charges a price p_s to the player π_s for each click on his advertisement.
- If player i is allocated slot s at price p_s , player i 's expected utility is $\theta_s(v_i - p_s)$.

The generalized second price auction is the most common mechanism in use.

Definition 2. The generalized second price mechanism (GSP) for keyword auctions uses the following allocation and payment rules:

- Players are allocated slots in decreasing order of bids.¹
- For each slot s , the payment p_s of player π_s is b_{π_s+1} .

Players who do not win a slot make no payment and gain no utility.

The GSP mechanism is not truthful, but has a continuum of Nash equilibria that are well understood [5, 11, 10]. One of these equilibria results in player payments identical to those that would be made if the mechanism being employed was VCG. We call this Nash equilibrium of GSP the *VCG*

¹Technically, there is also a “relevance” or per-advertiser click-through rate r_i associated with each advertiser i , and bidders are actually ranked (assigned slots) in decreasing order of ranking score, where the ranking score of bidder i is $b_i r_i$. In this paper we assume $r_i = 1$ for all i . All results in the paper extend easily to the case where the r_i values are arbitrary.

equilibrium of GSP. GSP also has Nash equilibria in which the revenue to the search engine is either higher or lower than that in the VCG equilibrium (see, for example, Varian [11] or Theorem 16 in this paper). It is not clear whether any of these equilibria are actually reached in real keyword auctions. It is also not clear what bidding strategy players might employ to reach a particular equilibrium. One of the goals of this paper is to gain some insight into these issues.

We consider a *repeated keyword auction*, with a fixed set of n players and k slots. The participants in the repeated auction have the opportunity to update their bids in between successive rounds. How should players bid? Naturally, a player's main objective is to maximize his own expected utility over multiple rounds of the auction. However, without any real insight into the bidding strategies followed by the rest of the players, it is difficult for one player to make predictions about the future bids of other players and hence choose an optimal bidding strategy. Thus, a natural approach is to assume that the immediate past is the best predictor of the future. This leads to a natural greedy-like bidding scheme where a player assumes that all the other bids will remain fixed in the next round. Under this assumption, the rational choice for a player j is to bid so as to win a slot s that maximizes his utility $u_j = \theta_s(v_j - p_s)$. This leads to the following definition.

Definition 3. Greedy Bidding Strategies

A greedy bidding strategy for a player j is to choose a bid for the next round of a repeated keyword auction round so as to maximize his utility u_j , assuming the bids of all other players b_{-j} in the next round will remain fixed to their values in the previous round.

Given b_{-j} , denote by $p_{-j}(s)$ the payment player j would make if he bids so as to win slot s . Let s^* be the slot the greedy bidder j will target. Then if player j is greedy, he will bid $b' \in (p_{-j}(s^* - 1), p_{-j}(s^*))$. As b' is not fully specified, this defines a class of strategies that are distinguished by the choice of b' within the allowed range.

Advertisers participating in a keyword auction are often business competitors. One of the most commonly observed secondary objectives besides gaining the optimal slot is the desire to “push” the prices paid by other advertisers higher. This is naturally done by bidding in the high end of the range mentioned earlier. However, there is the risk that a change in other players' bids could result in paying a higher price than expected. Balancing these objectives leads to the following bidding strategy.

Definition 4. Balanced Bidding

The Balanced greedy bidding strategy BB is the strategy for a player j that, given b_{-j}

- next targets the slot s_j^* which maximizes his utility (greedy bidding choice), that is,

$$s_j^* = \operatorname{argmax}_s \{\theta_s(v_j - p_{-j}(s))\}.$$

- chooses its bid b' for the next round so as to satisfy the following equation:

$$\theta_{s_j^*}(v_j - p_{-j}(s_j^*)) = \theta_{s_j^*-1}(v_j - b').$$

If s_j^* is the top slot, we (arbitrarily) choose $b' = (v_j + p_{-j}(1))/2$. We can thus deal with all slots uniformly by defining $\theta_0 = 2\theta_1$.

The intuition behind the bid selection is that the player bids high enough to force the prices paid by his competitors to rise, but not so high that if one of his competitors were to just undercut him, he would mind getting a higher slot at a price just below his own bid of b' . The BB strategy is an appealing strategy for the following reason.

Theorem 5. [5] *If all the players are following the BB strategy in an auction with all distinct θ s, then the system has a unique fixed point. In this fixed point, the revenue of the auctioneer (and payments of each player) is equal to that of the VCG equilibrium. The equilibrium bids b_j^* of the players in the fixed point of BB satisfy the following equations:*

$$b_j^* = \begin{cases} v_j & \text{if } j \geq k+1 \text{ and} \\ \gamma_j b_{j+1}^* + (1 - \gamma_j)v_j & \text{if } 2 \leq j \leq k. \end{cases} \quad (1)$$

where $\gamma_j = \theta_j/\theta_{j-1}$.

3. CONVERGENCE PROPERTIES OF THE BB STRATEGY

We study the convergence properties of the BB strategy in a repeated keyword auction under two models. We refer to our standard model, where all players simultaneously update their bids on each round as the *synchronous* model. In the *asynchronous* model, in each round, exactly one player updates her bid, while the other players merely repeat their previous bids. We consider both the case in which the player performing the update is arbitrary and the case where the player performing the update is chosen at random. This has been studied, for example in Zhang [14]. The asynchronous model is closer to actual practice in sponsored search, although the synchronous model applies in on-line settings where bids are updated in batches.

Theorem 6. *Consider a repeated keyword auction where all players are using the BB strategy starting with arbitrary initial bids. We have:*

1. A 2-slot auction system always converges to its fixed point in both the synchronous and asynchronous models. The number of rounds until convergence in the synchronous model is $O(\log((v_2 - v_3)/v_3))$, where the constant depends on the click-through rates θ_1 and θ_2 .
2. There exists a 3-slot auction system and set of initial bids which does not converge in the synchronous model.
3. There exists a 3-slot auction system, a set of initial bids, and an order in which the players update which does not converge in the asynchronous model.
4. In the asynchronous model where players bid in random order, no matter how many slots there are, the system always converges to its fixed point.

Proof of Part 1:

Lemma 7. *At every round t such that $t > t_1 = 2 + \log_{\gamma^*}((1 - \gamma^*)(v_2 - v_3)/v_3)$, where $\gamma^* = \max\{\theta_1/\theta_0, \theta_2/\theta_1\}$, we have:*

$$\begin{aligned} b_1, b_2 &> v_3, \\ b_3 &= v_3. \end{aligned}$$

PROOF. Let b denote the third highest bid. By the second round, b can never be more than v_3 . Suppose for some round that b is less than v_3 . Take a player i in $\{1, 2, 3\}$. In the next round, i will bid her value or target some slot $j \in \{1, 2\}$ and bid $b'_i = (1 - \gamma_j)v_i + \gamma_j p_j \geq (1 - \gamma^*)v_3 + \gamma^*b = b + (1 - \gamma^*)(v_3 - b)$. In either case,

$$(v_3 - b'_i) \leq \gamma^*(v_3 - b).$$

Initially $v_3 - b \leq v_3$.

It takes at at most $r = \log_{\gamma^*}((1 - \gamma^*)(v_2 - v_3)/v_3)$ before $v_3 - b < (1 - \gamma^*)(v_2 - v_3)$. (Recall that $\gamma^* < 1$ as the θ_i 's are decreasing.) In round $r + 1$, bidders $i \in \{1, 2\}$ will each bid either $v_i > v_3$ or target a slot $j \in \{1, 2\}$ and bid

$$\begin{aligned} b'_i &= (1 - \gamma_j)v_i + \gamma_j p_j \geq (1 - \gamma^*)v_2 + \gamma^*b \\ &= b + (1 - \gamma^*)(v_2 - b) > b + (1 - \gamma^*)(v_2 - v_3) \\ &> v_3, \end{aligned}$$

hence in either case their bids are both above v_3 . In round $r + 2$, player 3 will then bid v_3 while players 1 and 2 keep on bidding above v_3 ; this concludes the proof of the lemma. \square

Once the conditions of Lemma 7 hold, the price of slot 2 is fixed at $p_2 = v_3$. Let $T_2 = b_2^* = (1 - \theta_2/\theta_1)v_2 + (\theta_2/\theta_1)p_2$ and $T_1 = (1 - \theta_2/\theta_1)v_1 + (\theta_2/\theta_1)p_2$. If the last bids of players 1 and 2 were b_1 and b_2 then, in the next round, their bids are:

$$\begin{aligned} b'_1 &= \begin{cases} T_1 & \text{if } b_2 > T_1, \\ (v_1 + b_2)/2 & \text{otherwise.} \end{cases} \\ b'_2 &= \begin{cases} T_2 & \text{if } b_1 > T_2, \\ (v_2 + b_1)/2 & \text{otherwise.} \end{cases} \end{aligned} \quad (2)$$

Let $b_{\min} = \min(b_1, b_2)$ be the minimum of the two bids.

Lemma 8. *After at most t_2 rounds, we have $b_{\min} \geq T_2$, where*

$$t_2 \leq t_1 + 2 \frac{\theta_1 - \theta_2}{\theta_2}.$$

PROOF. Assume that at time $t > t_1$, we have $b_{\min} < T_2$. Then it is easy to check that at the next round we have $b'_{\min} \geq (v_2 + b_{\min})/2$. This implies

$$b'_{\min} - b_{\min} \geq \frac{v_2 - b_{\min}}{2} \geq \frac{v_2 - T_2}{2} = \frac{\theta_2}{\theta_1} \frac{v_2 - p_2}{2} = \delta.$$

Since at time t_1 we have $b_{\min} \geq v_3$, we will reach $b'_{\min} \geq T_2$ after an additional number of rounds bounded by $(T_2 - v_3)/\delta$. \square

Finally, since at round t_2 we have $b_{\min} \geq T_2$, at time $t_2 + 1$ we will have $b'_2 = T_2 < T_1$, and therefore at time $t_2 + 2$ we will have $b''_2 = T_2 = b_2^*$ and $b'_1 = (v_1 + T_2)/2 = b_1^*$; we have reached equilibrium. This proves Part 1 of the Theorem in the synchronous model.

The proof of convergence in the asynchronous model is similar, but complicated by the random order of bidding. Lemma 7 holds with a constant factor increase in t_1 to allow for all bidders to make a bid. Lemma 8 is easier to check as in this case only the top two players alter their bids, therefore after t_2 bid alternations we reach the same outcome. \blacksquare

Proof of Part 2: Unfortunately, an auction system where all the players follow the BB strategy is not guaranteed to converge when there are more than 2 slots. We will show

this with the following counterexample. Let there be three slots with $\theta_1 = 1$, $\theta_2 = 2/3$, $\theta_3 = 1/3$ and four players with values 161, 160, 159 and 100. Let the initial bids of 130.5, 130, 129.5 and 100, respectively. Then the bidding evolves as below; in particular it is not convergent.

Bidder Value	Round		
	1	2	3
161	130.5	145.5	130.5
160	130	145.25	130
159	129.5	144.75	129.5
100	100	100	100

There is one inactive bidder, with the lowest value, who will never be able to bid for a slot as the prices are all above his utility. In the first round, the remaining bidders all target the lowest slot. The bid from the highest-valued player is low enough so that for the next round, all three top bidders target the first slot. With these high bids, the price for any player for the top slot is too high when compared with the price of the last slot. Thus for the third round all three top players target the last slot, and the cycle continues.

In this example the click-through rates follow a simple geometric sequence, similar to those observed in practice [6]. Note that in this example the bids are well-behaved, in the sense that the bids are in the same order as the players' values. Hence such regularity is not sufficient for convergence. Finally, even though we used our convention of $\theta_0 = 2$, a similar example can be constructed where the players are cycling while targeting intermediate slots. This concludes the proof of part 2. \blacksquare

Proof of Part 3: Consider the following three slot example with $\theta_1 = 1$, $\theta_2 = 0.1$ and $\theta_3 = 0.09$ and the top three bidders alternating in sequence.

Value	Round						
	1	2	3	4	5	6	7
102	19.2	80.8	80.8	80.8	19.2	19.2	19.2
101	19.1	19.1	90.9	90.9	90.9	19.1	19.1
100	59.6	59.6	59.6	95.45	95.45	95.45	59.6
10	10	10	10	10	10	10	10

Proof of Part 4: A proof sketch is given the Appendix. \blacksquare

Notes

- The example above showing that for 3 slots, BB doesn't converge in the synchronous model is not anomalous. In Section 5, we show experimentally that in a significant fraction of instances, BB does not converge in the synchronous model.
- The bound we give on the time to convergence of BB in the random asynchronous model is extremely loose. We have run simulations to study the speed of convergence in this setting and have found experimentally that convergence is actually quite fast, certainly no more than a polynomial in n , the number of bidders.

4. CONVERGENCE PROPERTIES OF THE RESTRICTED-BB STRATEGY

We will now examine a variant of the BB strategy, called RBB, where the players can only aim for their current slot or

a slot of lower click-through rate than the one they currently have. The RBB strategy is designed so that it has the same unique fixed point as BB. However, by restricting the degree to which a player can be greedy, we will be able to show that even in the synchronous model, RBB always converges to the VCG equilibrium.

Definition 9. *The Restricted Balanced Bidding (RBB) bidding strategy is the strategy where given b_{-j} from the previous round, player j*

- *targets the slot s_j^* which maximizes his utility among the slots with no higher click-through rate than his current slot s_j , that is,*

$$s_j^* = \operatorname{argmax}\{\theta_s(v_j - p_{-j}(s)) : s \geq s_j\}.$$

- *chooses her bid b' for the next round so as to satisfy the following equation: $\theta_{s_j^*}(v_j - p_{-j}(s_j^*)) = \theta_{s_{j-1}^*}(v_j - b')$.*

To make sure this is well-defined for the first slot, we define $\theta_0 = 2\theta_1$.

Theorem 10.

1. *The system defined by a repeated keyword auction in which all the players are following the RBB strategy has a unique fixed point at which players are bidding according to the VCG equilibrium, i.e. the equilibrium bids b_j^* are given by Equation (1).*
2. *In the synchronous model the RBB strategy always converges to its fixed point. The number of steps until convergence is 2^k times*

$$O\left(k + \frac{\log(1 - \gamma^*)}{\log \gamma^*} + \log_{(1/\gamma^*)} \frac{v_1 - v_{k+1}}{\min_{1 \leq i \leq k} (v_i - v_{i+1})}\right),$$

where $\gamma^* = \max_i \theta_i / \theta_{i-1}$.

PROOF. The proof of Part 1 of the Theorem is very similar to that of Theorem 5. (See Edelman et al.[5].) To prove Part 2, we first bound the number of steps until the price of slot k and the set of players who will be allocated slots have converged. As before, we define $\gamma_i = \theta_i / \theta_{i-1}$.

Lemma 11. *Player p prefers to target slot j rather than slot $j - 1$ if and only if*

$$(1 - \gamma_j)v_p + \gamma_j p_j < p_{j-1}.$$

The proof of this Lemma is simple algebra and is omitted.

Lemma 12. *At every round t such that $t > t_1 = 2 + \log_{\gamma^*}((1 - \gamma^*)(v_k - v_{k+1})/v_{k+1})$, where $\gamma^* = \max_{i>0} \theta_i / \theta_{i-1}$, we have:*

$$\begin{aligned} b_i &> v_{k+1} & \forall i \leq k, \\ b_i &= v_i & \forall i \geq k+1. \end{aligned}$$

PROOF. Let b denote the $(k+1)$ st highest bid. By definition, b can never be more than v_{k+1} . Suppose for some round that b is less than v_{k+1} . Take any player i in $\{1, 2, \dots, k+1\}$. In the next round, i will either bid her value or target some slot $j \in \{1, \dots, k\}$ and bid $b'_i = (1 - \gamma_j)v_i + \gamma_j p_j \geq (1 - \gamma^*)v_{k+1} + \gamma^*b \geq b + (1 - \gamma^*)(v_{k+1} - b)$. In either case,

$$(v_{k+1} - b'_i) \leq \gamma^*(v_{k+1} - b).$$

Initially $v_{k+1} - b \leq v_{k+1}$. How many rounds does it take before $v_{k+1} - b < (1 - \gamma^*)(v_k - v_{k+1})$? At most $r \leq \log_{\gamma^*}((1 - \gamma^*)(v_k - v_{k+1})/v_{k+1})$. In round $r + 1$, bidders $i \in \{1, \dots, k\}$ will each bid either $v_i > v_{k+1}$ or bid at least

$$\begin{aligned} b'_i &= (1 - \gamma_j)v_i + \gamma_j p_j \geq (1 - \gamma^*)v_k + \gamma^*b \\ &\geq b + (1 - \gamma^*)(v_k - b) > b + (1 - \gamma^*)(v_k - v_{k+1}) \\ &> v_{k+1}, \end{aligned}$$

hence in either case their bids are above v_{k+1} . In round $r + 2$, player $k + 1$ will then bid v_{k+1} while players $1, 2, \dots, k$ continue to bid above v_{k+1} ; the other players don't get a slot, so they bid their value, and this concludes the proof. \square

We now need to prove that the allocation of the k slots players to these k players converges to a fixed point.

At any time, for any $i \in [1, k]$, consider the players allocated slots $[i + 1, k]$. They are called *stable* if their bids and prices satisfy Equation (1), that is, the allocation is in order of decreasing values, and if $\pi(j)$ is the player currently allocated slot j , then the last bids of those players satisfied:

$$b_{\pi(j)} = \gamma_j b_{\pi(j+1)} + (1 - \gamma_j)v_{\pi(j)},$$

for every $j \in [i + 1, k]$.

If the players allocated all slots $[1, k]$ form a stable set, then we have reached the fixed point of the RBB strategy.

Assume that the current setting is not (yet) a fixed point of the RBB strategy. Let A be the maximum stable set, with associated $i \geq 2$, and let B be the set of players in slots $[1, i]$. Let b_{\min} denote the minimum bid from players of B .

We define a partial order over sets of players. We say that $A' \succ A$ if either $A \subset A'$ and $A \neq A'$, or if the player of minimum value in the symmetric difference of A and A' belongs to A' ². This will be our measure of progress.

In the next round, observe that, following the RBB strategy, players in A still bid in the same way as before. Let b'_{\min} be the new minimum bid from players of B , and p be the player of B whose bid is b'_{\min} . There are three cases to consider.

1. p bids below p_i . Let $j \in [i + 1, k]$ be the slot which was targeted by p . By definition, p prefers slot j to slot $j - 1$, and so, by Lemma 11, the bid of p is less than p_{j-1} . By definition, the bid is $(1 - \gamma_j)v_p + \gamma_j p_j > p_j$, thus it falls in the interval (p_{j-1}, p_j) and p will be allocated slot j . Recall that $\pi(j) \in A$ denotes the player who was in slot j in the previous round. Since set A is stable, by definition we have $p_{j-1} = (1 - \gamma_j)v_{\pi(j)} + \gamma_j p_j$. Since this is greater than the bid of p , it follows that $v_{\pi(j)} > v_p$. Moreover, since p prefers slot j to slot $j + 1$, by Lemma 11 again, we must have $(1 - \gamma_{j+1})v_p + \gamma_{j+1}p_{j+1} > p_j = (1 - \gamma_{j+1})v_{\pi(j+1)} + \gamma_{j+1}p_{j+1}$, and so $v_p > v_{\pi(j+1)}$. Thus $A' = \{p' \in A : v_{p'} < v_p\} \cup \{p\}$ is a stable set, and $A' \succ A$.
2. p targeted slot i . Then p is allocated slot i , $A' = A \cup \{p\}$ is a stable set, and $A' \succ A$.
3. p targeted some slot $j \leq i - 1$. Then A is still a stable set, and $b_{\min} = p_{i-1}$ has increased: $b'_{\min} = (1 - \gamma_j)v_p + \gamma_j p_j \geq b_{\min} + (1 - \gamma^*)(v_p - b_{\min})$.

²This corresponds to a lexicographic ordering of the sets.

We will prove that Case 3 can only happen a bounded number x of times, (where x depends on the θ_j 's and the v_j 's but not on the bids) before Case 1 or 2 must occur. Thus, the maximum stable set must change at least once every x rounds, and when that happens, it is replaced by a set which is larger in the \succ ordering. This implies that the system converges to a fixed point and that the number of rounds until convergence is bounded by $2^k(x+1)$, hence the Theorem.

First, a useful technical lemma.

Lemma 13. *Let $\varepsilon = (1/2)\theta_k(1-\gamma^*) \min_{q \neq q'} |v_q - v_{q'}|/\theta_1$. If $p_{i-1} > v_p - \varepsilon$ and $v_p > p_i$, then player p prefers slot i to any slot $j < i$.*

PROOF. From player p 's viewpoint, the utility of slot i is $\theta_i(v_p - p_i)$, the utility of slot $j < i$ is $\theta_j(v_p - p_j) < \theta_j(v_p - p_{i-1})$, and the ratio is

$$\frac{\theta_j(v_p - p_{i-1})}{\theta_i(v_p - p_i)} \leq \varepsilon \frac{\theta_j}{\theta_i(v_p - p_i)} \leq \varepsilon \frac{\theta_1}{\theta_k(v_p - p_i)}.$$

Now,

$$\begin{aligned} v_p - p_i &= v_p - ((1 - \gamma_{i+1})v_{\pi(i+1)} + \gamma_{i+1}p_{i+1}) \\ &= (1 - \gamma_{i+1})(v_p - v_{\pi(i+1)}) + \gamma_{i+1}(v_p - p_{i+1}), \end{aligned}$$

which is at least $(1 - \gamma^*) \min_{q \neq q'} |v_q - v_{q'}|$. Plugging this into the previous expression proves the Lemma. \square

Now, let $x = \log_{1/\gamma^*}((v_1 - v_{k+1})/\varepsilon)$. Assume that Case 3 happens for x consecutive rounds. Let p_{\min} be the player in B whose value is minimum and v_{\min} be its value. Let $b_{\min}^{(t)}$ be the minimum bid of players in B after t rounds, $0 \leq t \leq x$. If $p \in B$ is the player defining the minimum bid in round $t+1$, we have:

$$b_{\min}^{(t+1)} \geq (1 - \gamma^*)v_p + \gamma^*b_{\min}^{(t)} \geq (1 - \gamma^*)v_{\min} + \gamma^*b_{\min}^{(t)}.$$

After x rounds, we get $b_{\min}^{(x)} \geq v_{\min} - (\gamma^*)^x(v_{\min} - b_{\min}^{(0)})$, hence $b_{\min}^{(x)} \geq v_{\min} - (\gamma^*)^x(v_i - v_{k+1})$. Plugging in the value of x yields $b_{\min}^{(x)} \geq v_{\min} - \varepsilon$. From Lemma 13, we know that p_{\min} prefers slot i to any slot $j < i$. In the next round p_{\min} targets slot i and has to be the minimum bidder from B , therefore we are now in Case 2. Thus there are at most x occurrences of Case 3 between any two occurrences of Case 1 or Case 2, and the Theorem is proved. \blacksquare

5. EMPIRICAL EVALUATION

In the previous section, we showed how simple greedy bidding strategies can lead to the VCG equilibrium. Are the theoretical convergence bounds tight? How desirable is this outcome for bidders and for the search engine? In this section, we study these questions empirically. We first study the convergence and bidder utility of the balanced bidding strategies, before turning to other greedy bidding strategies. We finish by comparing search engine revenue in equilibrium with the VCG revenue.

5.1 Experimental Setup

In all experiments, we define the keyword auction as follows: we use three slots and four players. Following the study of Feng et al.[6], we choose the click-through rates as a geometrically decreasing sequence by $\theta_i = \delta^{i-1}$ for some

value of δ between 0 and 1, and plot our results as a function of δ . We take the average of 150 instances where for each instance, the values of the four players are each independently chosen from a normal distribution with mean 500 and deviation 200.

To analyze a bidding strategy, for each of the 150 instances, the simulated strategy is run from starting bids of 1 for all players, except for AB (defined later), which is run from starting bids equal to the minimum player value. Some experiments, such as those comparing equilibrium revenue with VCG revenue, are not affected by the choice of asynchronous or synchronous settings. Others, such as BB convergence, do depend on the setting, and in such cases we will discuss the differences.

5.2 Balanced Bidding Strategies: BB and RBB

We first demonstrate the frequency of synchronous non-convergence of BB shown in Theorem 6 part 2. Figure 1a plots the probability that BB converges. (When $\delta > .95$, it takes much longer for the system to reach a fixed point or a cycle, and our simulations did not always reach that point, so we plot two curves representing lower and upper bounds to the convergence probability.) This figure indicates that the non-convergent example in Theorem 6 is not pathological: *The non-convergence of BB is a practically observed phenomenon which occurs for a significant fraction of instances.*

While we have proved that the BB strategy always converges in an asynchronous setting, our theoretical bounds on the rate of convergence are quite weak. Figure 1b plots the average number of bids per player until convergence, as function of the number of players. The average number of bids per player increases linearly in the number of players. This suggests a $O(n^2)$ overall rate of convergence.

Our analysis suggests that bidders should use RBB rather than BB in the synchronous setting, since it gets to the same fixed point and has the advantage of always converging. But, bidders considering this recommendation might wonder about its effect on their utility. To study this concern, we take one special player p and fix his value at 500, while the values of the other players continue to be drawn from a normal distribution with mean 500 and standard deviation 200. We then plot, as a function of δ , the ratio of player p 's utility under RBB relative to his utility under BB.

We see in Figure 2 that when $\delta < .6$, player p 's utility in RBB is only about 90% of the utility he would get in BB. Interestingly, as δ approaches 1, this is reversed: when δ is close to 1, RBB yields better utility than BB. We think this results from frequent non-convergence of BB when δ is small. It is claimed [6] that in practice the click-through rates are fit well by a geometric sequence with δ about .7. In that range RBB and BB yield about the same revenue.

Note that our result appears to be robust: we have observed that changing the value of the player under focus from 500 to 400 or to 600 (so as to make him either a relatively lower valued player or a relatively higher valued player) does not change the graph significantly.

5.3 Other Greedy Bidding Strategies: CB and AB

So far we have focused on BB and RBB, and have found that these strategies will result in search engine revenue equal to the VCG revenue. Thus we can use the VCG revenue as




Figure 1: Convergence of BB




Figure 2: RBB vs. BB Player Utility (BB=1)

a benchmark for comparing search engine revenue. In this section we examine two other greedy bidding strategies and study their impact on search engine revenue, compared to VCG revenue. For the following we let $\varepsilon_{\text{price}}$ be a suitably small bid increment, for example 1¢.

5.3.1 Competitor Busting

A popular bidding strategy used in practice is known as *competitor busting* [13]. This strategy is also referred to as anti-social or vindictive bidding [2, 15], and may be used by as many as 40% of the bidders on Yahoo! [15]. The idea is that a player bids as high as possible while retaining her desired slot, in order to make competitors pay as much as possible and exhaust their advertising resources.

Definition 14. *The Competitor-Busting greedy bidding strategy (CB) is the strategy for a player j that, given b_{-j}*

- *next targets the slot s_j^* which maximizes her utility (greedy bidding choice), that is,*

$$s_j^* = \operatorname{argmax}\{\theta_s(v_j - p_{-j}(s)) : s \geq s_j\}.$$

- *chooses her next bid as $b' = \min\{v_j, p_{-j}(s_{j-1}^*) - \varepsilon_{\text{price}}\}$.*

Convergence issues for CB are much more serious than for BB: in general, the CB strategy does not have a fixed point! Indeed, a player p in slot i will only have non-negligible utility if the player in slot $i+1$ cannot raise her bid any further, because it is equal to her value. In this case, unless the situation is degenerate, p will prefer to move to a different slot, which will in turn lead other players to want different slots.

Thus, the only fixed point for CB is when all players are bidding their values and those bids happen to be a Nash

equilibrium for the GSP strategy. Bids equal to the values form a Nash equilibrium if and only if $\theta_i(v_i - v_{i+1}) \geq \theta_j(v_i - v_{j+1})$ for each i and $j > i$; for a random instance, these constraints are satisfied with significant probability when δ is close to 0.

Indeed, Figure 4 confirms this finding. Like in Figure 1, we plot two curves, which are lower and upper bounds on the probability of convergence. (The space between the two curves corresponds to instances that have neither converged nor begun to cycle; as cycle detection is inaccurate in the asynchronous setting the true answer is probably nearer the lower curve). We see that CB frequently converges when δ is small, and rarely converges when δ is large. Figure 4 was plotted in the asynchronous setting; the results are essentially identical in the synchronous setting, showing that CB does not exhibit different cycling behavior depending on the format (in contrast to BB).

Given that CB is a well-known strategy, the auctioneer might wonder how much revenue he obtains if every player followed the CB strategy. Figure 3 compares the revenue obtained from running GSP with the players following CB in an asynchronous setting, with the revenue which would have been obtained from running our benchmark, VCG (the graph has the VCG revenue normalized to 1). This is good news for the auctioneer: *If the players follow the CB strategy, then the auctioneer's revenue is higher than the VCG revenue, by up to 20%.* Accordingly, player utility goes down. The results are very similar in the synchronous setting, showing the revenue properties of CB are robust.

5.3.2 Altruistic Bidding

Since competitor busting primarily benefits the auction-




Figure 3: Comparing revenue from CB with the VCG revenue (VCG=1)




Figure 4: Convergence of CB

eer, the players might consider trying a completely different approach: altruism. A natural complement to the competitor busting strategy, where players try to hurt other bidders as much as possible, the altruistic strategy has players bidding as low as possible to win their desired slot.

Definition 15. *The Altruistic greedy bidding strategy (AB), is the strategy for a player j that, given b_{-j}*

- *next targets the slot s_j^* which maximizes her utility (greedy bidding choice), that is,*

$$s_j^* = \operatorname{argmax}\{\theta_s(v_j - p_{-j}(s)) : s \geq s_j\}.$$

- *chooses her next bid as $b' = \min\{v_j, p_{-j}(s_j^*) + \varepsilon_{\text{price}}\}$. If there is no slot giving positive utility, bid v_j .*

Convergence issues for AB are even more serious than for CB: the AB strategy has *no* fixed points when the θ_i 's are separated. The auctioneer might worry about how much revenue he would obtain if every player followed the AB strategy. As can be seen in Figure 5, this fear is only somewhat justified: AB produces low revenue, but more than the minimum Nash equilibrium of GSP. Like with the CB revenue, these results from the asynchronous setting are very similar to those in the synchronous setting. It would be interesting to see if there are any auctions in practice where players bid in this way, or if a bidding strategy converging to the minimum Nash revenue exists.

5.4 Revenue in GSP Equilibria

How much revenue does GSP bring to search engines, compared to the revenue generated by the VCG mechanism? In order to evaluate this we assume that bidders eventually end up in some Nash equilibrium of GSP (which of course is not necessarily the case). Figure 6 plots the average, over instances, of the ratio between the revenue in some Nash equilibrium of GSP and the VCG revenue. Since GSP has many different Nash equilibria, we consider three extreme points: the maximum revenue Nash equilibrium, the minimum revenue Nash equilibrium, and the maximum revenue Nash equilibrium when the bidders are *debt-averse*, i.e. never bid above their value. Kitts et al.[8] show evidence that bidders usually follow a debt-averse bidding strategy.

From these experiments, we conclude the following: *the equilibrium revenue of GSP is within $\pm 40\%$ of the VCG revenue, unless δ is close to 1 and bidders are willing to bid above their value.* This should be contrasted with Part 1 of Theorem 16, where we exhibit an example of a GSP equilibrium with revenue much smaller than the VCG revenue: our simulations indicate that such an example is an anomaly.

We also see the following interesting behavior: *If bidders are debt-averse, then in equilibrium GSP generates at most 15% more revenue than VCG.* This indicates that the impact of choosing GSP over VCG is not as significant when rational behavior leads to debt-averse bidding. The influence of




Figure 5: AB total revenue (VCG=1)




Figure 6: Comparing revenue from GSP equilibria with the VCG revenue (VCG=1)

debt-averse behavior on the maximum revenue is rigorously justified in Part 3 of Theorem 16.

5.5 Theoretical results

The following theorem compares the auctioneer revenue obtainable by the GSP mechanism to the revenue which could have been obtained by using the VCG mechanism instead.

- Theorem 16.** 1. For every $K > 0$, there exists a keyword auction and a Nash equilibrium of the GSP mechanism whose revenue is at most $1/K$ times the revenue of the VCG mechanism, and moreover, every bidder i bids $b_i \leq v_i$.
2. For every $K > 0$, there exists a keyword auction and a Nash equilibrium of the GSP mechanism whose revenue is at least K times the revenue of the VCG mechanism.
3. If every bidder i bids $b_i \leq v_i$ (i.e. is debt-averse), then for every keyword auction and Nash equilibrium of the GSP mechanism, the revenue is at most α^* times the VCG revenue, where $\alpha^* = \max_i \theta_i / (\theta_i - \theta_{i+1})$.

PROOF. *Part 1:* Here is an example showing that the revenue to the auctioneer from a GSP equilibrium may be arbitrarily smaller than that the revenue from the VCG equilibrium. The example is for a two-slot auction, where $\theta_1 = 1$ and $\theta_2 = \frac{1}{2}$. The GSP revenue is $2\theta_1 + (1/2)\theta_2 = 2.25$ whereas the VCG revenue is $(x+1/2)\theta_1 + (1/2)\theta_2 = x + (3/4)$.

A GSP equilibrium much smaller than VCG

player	value	bid	slot won	GSP pr.	VCG pr.
1	$x+1$	$x+1$	1	2	$x + \frac{1}{2}$
2	x	2	2	$\frac{1}{2}$	$\frac{1}{2}$
3	1	1	—	0	0

Part 2: Here is an example showing that the revenue to the auctioneer from a GSP equilibrium may be arbitrarily larger than that the revenue from the VCG equilibrium. The example is for a two-slot auction, where $\theta_1 = 1$ and $\theta_2 = \frac{1}{2}$.

A GSP equilibrium much larger than VCG

player	value	bid	slot won	GSP pr.	VCG pr.
1	x	x	1	$(x+2)/2$	$\frac{3}{2}$
2	2	$\frac{(x+2)}{2}$	2	$\frac{1}{2}$	$\frac{1}{2}$
3	1	1	—	0	0

These examples rely on player values that are arbitrarily separated. The second example was constructed using the following expression from Varian [11], which gives the payments that achieve the maximum Nash revenue.

$$\theta_i p_i = \sum_{j>i} v_{j-1}(\theta_j - \theta_{j+1}) \quad (3)$$

Compare this with the VCG payments:

$$\theta_i p_i = \sum_{j>i} v_j(\theta_j - \theta_{j+1}) \quad (4)$$

Observe that if the values are close to each other, the maximum revenue from a Nash equilibrium is close to the VCG revenue: if $v_{i+1} \geq \alpha v_i$, then the payments from (3) are at most a factor of $1/\alpha$ from the payments of (4).

Notice that the second example, with Nash revenue larger the VCG revenue, entails the second player bidding much more than his value. This example therefore is not debt-averse (as we defined that concept in Section 5.4). Under the more realistic debt-averse assumption we obtain the third part of this theorem.

Part 3: Let R^M be the maximum debt-averse Nash revenue. Note that the payment of i , p_i^M , is at most v_{i+1} . If $p_i^M > v_{i+1}$, as all players are bidding at most their values, there would not be enough winners to fill the top i slots. Thus the maximum risk-free Nash equilibrium revenue is $R^M \geq \sum_{1 \leq i \leq k} \theta_i v_{i+1}$. On the other hand, the VCG payment of i is $\theta_i p_i^{\text{VCG}} \geq (\theta_i - \theta_{i+1})v_{i+1}$. Thus the total VCG revenue is at most $\sum \theta_i p_i^{\text{VCG}} \geq \sum (\theta_i - \theta_{i+1})v_{i+1} = \sum (\theta_i - \theta_{i+1})\theta_i v_{i+1}/\theta_i \geq R^M/\alpha^*$ as required. \square

6. REFERENCES

- [1] ASDEMIR, K. Bidding patterns in search engine auctions. In *Second Workshop on Sponsored Search Auctions* (2006), ACM Electronic Commerce.
- [2] BRANDT, F., AND WEISS, G. Antisocial agents and vickrey auctions. In *Proceedings Of The 8th Workshop on Agent Theories Architectures* (2001).
- [3] CLARKE, E. H. Multipart pricing of public goods. *Public Choice* 11 (1971), 17–33.
- [4] EDELMAN, B., AND OSTROVSKY, M. Strategic bidder behavior in sponsored search auctions. In *First Workshop on Sponsored Search Auctions* (2005), ACM Electronic Commerce.
- [5] EDELMAN, B., OSTROVSKY, M., AND SCHWARZ, M. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. To appear in *American Economic Review*, 2006.
- [6] FENG, J., BHARGAVA, H. K., AND D.M., P. Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *INFORMS Journal on Computing* (2006). Forthcoming.
- [7] GROVES, T. Incentives in teams. *Econometrica* 41 (1973), 617–631.
- [8] KITTS, B., LAXMINARAYAN, P., LEBLANC, B., AND MEECH, R. A formal analysis of search auctions including predictions on click fraud and bidding tactics. In *First Workshop on Sponsored Search Auctions* (2005), ACM Electronic Commerce.
- [9] KITTS, B., AND LEBLANC, B. Optimal bidding on keyword auctions. *Electronic Markets* (2004), 186–201.
- [10] LAHAIE, S. An analysis of alternative slot auction designs for sponsored search. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce* (New York, NY, USA, 2006), ACM Press, pp. 218–227.
- [11] VARIAN, H. Position auctions. To appear in *International Journal of Industrial Organization*, 2006.
- [12] VICKERY, W. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance* (1961), 8–37.
- [13] WEB-CITE. Shooting the moon and budget busting with overture’s auto-bidding, April 2003.
- [14] ZHANG, X. Finding edgeworth cycles in online advertising auctions. MIT Sloan School of Management, Working Paper, 2005.
- [15] ZHOU, Y., AND LUKOSE, R. Vindictive bidding in keyword auctions. In *Second Workshop on Sponsored Search Auctions* (2006), ACM Electronic Commerce.

APPENDIX

Proof Sketch of Theorem 6, Part 3.

At time t , we say that a player p is *activated* if p is the player who updates his bid while the other bidders repeat their previous bids.

Lemma 17. *Let T be a certain function of $n, k, (\theta_i), (v_j)$. For every starting configuration, there exists a sequence of player activations, of length at most T , such that the resulting configuration is a fixed point.*

PROOF. The idea will be to emulate the proof of Theorem 10. First, one can show that the result of Lemma 12 holds in the asynchronous setting, so that after t_1 rounds, $b_i > v_{k+1}$ for $i \leq k$, and $b_i = v_i$ for $i \geq k+1$, where t_1 is approximately n times the t defined in Lemma 12. The rest of the sequence is partitioned into phases, corresponding to stable sets, defined as in the proof of Theorem 10. Let A be the current stable set, $[i+1, k]$ be the slots occupied by the player of A , and B be the set of players occupying slots $[1, i]$. Let p_{\min} be the player in B whose value is minimum, v_{\min} be its value, and p_{i-1} be the price of slot $i-1$.

Consider the three cases enumerated in the proof of Theorem 10. We repeatedly activate the player currently in slot i until either Case 1 or Case 2 occurs, or $p_{i-1} > v_{\min} - \varepsilon$, where ε is defined as in Lemma 13. We then activate player p_{\min} . At this point, Case 1 or Case 2 must occur, a new stable set is defined, and the phase ends. This completes the definition of the sequence.

This argument bounds the length T of the constructed sequence of player activations by t_1 plus k times the bound in Theorem 10, Part 2, which is independent of the starting configuration and bids. \square

With this Lemma, it is easy to complete the proof of Theorem 6: In a random sequence, at every step we have probability at least $1/n$ to choose the next activation as in Lemma 17, hence the sequence will occur after about n^T steps on average, and in any case, it will occur with probability 1 after a finite time. This proves convergence.