

Discrete Strategies in Keyword Auctions and their Inefficiency for Locally Aware Bidders*

Evangelos Markakis¹ and Orestis Telelis^{2**}

¹ Department of Informatics, Athens University of Economics and Business, Greece
markakis@gmail.com

² Department of Computer Science, University of Liverpool, U.K.
telelis@gmail.com

Abstract. We study formally discrete bidding strategies for the game induced by the Generalized Second Price keyword auction mechanism. Such strategies have seen experimental evaluation in the recent literature as parts of iterative best response procedures, which have been shown not to converge. We give a detailed definition of iterative best response under these strategies and, under appropriate discretization of the players' strategy spaces we find that the discretized configurations space *contains* socially optimal pure Nash equilibria. We cast the strategies under a new light, by studying their performance for bidders that act based on local information; we prove bounds for the worst-case ratio of the social cost of locally stable configurations, relative to the socially optimum cost.

1 Introduction

We study discrete bidding strategies for repeated keyword auction games, induced by the Generalized Second Price (GSP) mechanism. Sponsored search auctions have received considerable attention in the recent literature, as the premiere source of income for search engines that allocate advertisement slots. The GSP mechanism is implemented in different forms by **Google**, **Yahoo!**, and **Bing**. Other online enterprises also use flavors of GSP; e.g. **Google** exports its slot allocation and pricing system as a service. In the simplest form of the mechanism, advertisers are ranked in order of non-increasing bids and each of the first k is matched to one of k available slots, paying the next highest bid to his. In the current version bids are weighted by relevance parameters of advertisers. For one slot the GSP mechanism coincides with the VCG mechanism. For at least two slots however, the GSP auction *does not* retain the features of VCG, e.g., truthful reporting of valuations, and encourages strategic behavior.

Strategic behavior in GSP auctions raises the question of how should an advertiser decide on his bidding. A best response of a player i under a current

* This work was partly funded by the EU FP7 Network of Excellence Euro-NF, through the specific joint research project AMESA. A full version with all omitted proofs is available at the authors' homepages.

** This work was carried out while the author was at the Center for Mathematics and Computer Science (CWI) Amsterdam, The Netherlands.

bidding configuration is any bid value within the interval defined by the bids of at most two other players, that grants i his desired slot; but how should the exact value be decided? In practice bidders may hire consultants to design bidding strategies for them. Phenomena of competition have been observed in the adopted strategies, ranging from modest budget investment to aggressive bidding, inducing large prices for competitors. These issues have received attention in the recent literature [3, 11]. Most of the existing works concern *iterative best response* procedures, viewing a GSP auction as a repeated game. Cary *et al.* [3] studied strategies where players adjust their bid iteratively, synchronously or asynchronously – in a randomly chosen order – always targeting the slot that maximizes their profits. They introduced 3 bidding strategies and proved convergence for one of them to a single fixed point, the equilibrium described in [4].

We focus on the other two simple strategies introduced in [3], that have seen less theoretical treatment, but have been used in experimental comparisons [3, 9, 11]. The first is *Altruistic Bidding* (AB), where every player takes a slot by minimally outbidding the player who currently owns it. The second is *Competitor Busting* (CB), where a player minimally underbids the player who owns the slot above the one aimed for. Both require discretization of the players' strategy spaces by a *bidding unit* ϵ . This may change the original game entirely. Iterative AB and CB procedures have been observed not to converge for fixed ϵ [3, 9]; can we expect the best response state space to even have pure Nash equilibria (PNE)? How should ϵ be tuned so that the game in discrete strategies retains properties of the original game? The relevance of AB and CB is amplified for bidders that, due to lack of complete information, perform local best responses.

Contribution. We study iterative AB and CB best response procedures that differ from previous work [3] in that bidders only update their bid when they have incentive to target a *different* slot. We provide a detailed description of AB and redefine CB differently than it has appeared previously, to ensure its consistency with developments to follow (Section 3). We decide an upper bound on the discretization parameter ϵ to ensure that the induced discretized configurations space has a socially optimum locally envy-free PNE, analogous to the one identified in [4], that is also a PNE for the game in continuous strategies. We ensure that if iterative AB or CB converge to a socially optimum configuration, then this is a PNE even in continuous strategies. Subsequently we examine the case of bidders that take only local steps upwards or downwards due to incompleteness of available information (Section 4). We study the social inefficiency of locally stable configurations of the GSP auction and produce upper bounds on the inefficiency of configurations reached by *local* iterative AB/CB (L-AB/L-CB).

Related Work. A considerable amount of work in sponsored search auctions concerns the strategic behavior of the bidders. As mentioned above, Cary *et al.* [3] defined and studied three bidding strategies, *Altruistic Bidding* (AB), *Competitor Busting* (CB) and *Balanced Bidding* (BB). CB has been observed often in practice [2, 13]. Using CB advertisers try to exhaust the budget of their competitors by placing the highest possible bid that will guarantee them the slot they decide to target. Altruistic bidding is the opposite of CB, whereas BB balances between

these two extremes. For BB the authors showed that, under some conditions, it converges to the efficient locally envy-free equilibrium characterized in [4]. For AB and CB it was shown that they do not generally converge. Experimental analysis of AB and CB revealed low and high revenue respectively.

The performance of these strategies is analyzed in Bayesian settings in [12, 10]. In [9], vindictive strategies are studied for games where bidders are either vindictive or cooperative. Regarding efficiency of equilibria, the first upper bounds on the Price of Anarchy with respect to the social welfare in GSP Auctions were obtained by Lahaie [7]. Tighter upper bounds were obtained for *conservative* bidders (that do not outbid their valuation) by Leme and Tardos in [8]. It was shown that the price of anarchy is at most equal to the golden ratio for the complete information game and at most 8 for the Bayesian setting.

2 Definitions & Preliminaries

An instance of the GSP Auction game has a set of n players (bidders), a set of k slots and a tuple $\langle \{\theta_j\}_{j=1}^k, \{\rho_i\}_{i=1}^n, \{v_i\}_{i=1}^n \rangle$. $\theta_j \in [0, 1]$ is the probability that a link displayed in slot j is clicked (*Click-Through Rate* - CTR), $\rho_i \in [0, 1]$ is the probability that an advertisement by player i is clicked (*relevance* of i) and v_i is the valuation of i . We use \hat{v}_i for $\rho_i v_i$, the *expected revenue* of i . Assume $\theta_1 \geq \dots \geq \theta_k > 0$, $\hat{v}_1 \geq \dots \geq \hat{v}_n$ and define $\gamma_j = \theta_j / \theta_{j-1}$, $\gamma = \max_j \gamma_j$ for $j \geq 2$.

The GSP Mechanism. Players issue collectively a bid vector $\mathbf{b} = (b_1, \dots, b_n)$; they are ranked in order of non-increasing *declared expected revenue* $\hat{b}_i = \rho_i b_i$ and matched to slots in order of non-increasing CTR. This is the *Rank-By-Revenue* (RBR) rule. When all bidders' relevances are equal, the players are ranked by non-increasing bid b_i (*Rank-By-Bid* rule - RBB). Under RBB, a player i receiving a slot j pays the bid of the $(j + 1)$ -th player. Under RBR, i pays the declared expected revenue of the bidder i' receiving slot $j + 1$ divided by ρ_i , i.e. $\rho_{i'} b_{i'} / \rho_i$.

Given a bid configuration \mathbf{b} , we denote by $b_{(j)}$, $\rho_{(j)}$, $v_{(j)}$, the bid, relevance and valuation of the player occupying slot j . \mathbf{b}_{-i} is the strategy profile \mathbf{b} without the bid of player i and $\mathbf{b}_{-(j)}$ denotes exclusion of the bid of the player occupying slot j . Define $\mathbf{b}(j) = b_{(j)}$, and $\mathbf{b}_{-i}(j)$, $\mathbf{b}_{-(i)}(j)$ will be the bid of the player occupying slot j in \mathbf{b}_{-i} and $\mathbf{b}_{-(i)}$ respectively. We use $\hat{\mathbf{b}}$ for the vector of declared expected revenues as above. The utility of a player occupying slot j under \mathbf{b} is:

$$u_{(j)}(\mathbf{b}) = \theta_j \rho_{(j)} \left(v_{(j)} - \frac{\rho_{(j+1)} \hat{b}_{(j+1)}}{\rho_{(j)}} \right) = \theta_j (\hat{v}_{(j)} - \hat{b}_{(j+1)}).$$

The *social welfare* $SW(\mathbf{b})$ of \mathbf{b} is $SW(\mathbf{b}) = \sum_{j=1}^k \theta_j \hat{v}_{(j)} = \sum_{j=1}^k \theta_j \rho_{(j)} v_{(j)}$. We assume a deterministic tie-breaking rule in case there are ties in the ranking. Edelman *et al.* [4] identified a PNE configuration \mathbf{b}^* for the GSP auction game with optimum social welfare $SW(\mathbf{b}^*) = \sum_j \theta_j \rho_j v_j$ and payments equal to the ones in the efficient dominant strategy equilibrium of the VCG mechanism. This equilibrium is also *locally envy-free*, i.e. every bidder i under \mathbf{b}^* is indifferent of receiving at price $\rho_i b_i^*$ the slot right above the one he occupies under \mathbf{b}^* .

Local Stability. In Section 4, motivated by the costs incurred to players for learning the competitors' bids, we assume that a player only learns the price of the slots right above/below the slot he currently occupies and only considers these local deviations. In case of ties, i.e., other players above/below him bidding the same, we assume that he learns the price of the first slot below the ties. This inspires a definition of *local stability*, which is a relaxation of Nash equilibrium.

Definition 1. Let \mathbf{b} be a bid configuration of the Generalized Second Price Auction game with k slots and $n \geq k$ players. Fix any slot $j_0 \in \{1, \dots, k\}$ and let $j_1 = j_0 + 1$, $j_2 = j_0 - 1$. Define $j'_1 = \min(\{n\} \cup \{j | \hat{b}_{(j)} < \hat{b}_{(j_1)}\})$ and $j'_2 = \max(\{1\} \cup \{j | \hat{b}_{(j)} > \hat{b}_{(j_2)}\})$. The bid configuration \mathbf{b} is locally stable if:

1. For any slot j_0

$$\text{if } j_0 \neq k \text{ and } j'_1 \leq k + 1, \theta_{j_0}(\hat{v}_{(j_0)} - \hat{b}_{(j_0+1)}) \geq \theta_{j'_1-1}(\hat{v}_{(j_0)} - \hat{b}_{(j'_1)}), \quad (1)$$

$$\text{if } j_0 \neq 1, \theta_{j_0}(\hat{v}_{(j_0)} - \hat{b}_{(j_0+1)}) \geq \theta_{j'_2+1}(\hat{v}_{(j_0)} - \hat{b}_{(j'_2+2)}), \quad (2)$$

2. For any player i who does not win a slot under \mathbf{b} , $\hat{v}_i \leq \hat{b}_{(k)}$.

The definition states that no player has an incentive to move to the next feasible slot upwards or downwards under \mathbf{b} . j'_1 and j'_2 determine the slot that the bidder at slot j_0 can target, in case that due to ties he cannot aim for the one right above/below him. The condition $j'_1 \leq k + 1$ in (1) states that a bidder may not be able to deviate downwards if all the remaining bidders have equal score. For non-winning players, we assume they know the *bidding entry level* to competition, $\hat{b}_{(k)} = \rho_{(k)} b_{(k)}$. The last constraint prescribes that no such bidder has incentive to target slot k . In analogy to the *Price of Anarchy* [6], we quantify the inefficiency of locally stable configurations by the following worst-case ratio:

Definition 2. The Local Stability Ratio of a GSP Auction game is defined as $\Lambda = \sup_{\mathbf{b}} \frac{\sum_j \theta_j \hat{v}_j}{SW(\mathbf{b})}$, where the supremum is over all locally stable configurations.

We note that the notion of a locally stable configuration and hence the notion of the Local Stability Ratio can be defined for a much wider context. They are applicable to any game where the outcome is a ranking, and for every action profile b any player is allowed, in a well defined manner, to deviate upwards or downwards in the ranking and determine his new payoff. *Ranking Games* [1] constitute one such interesting class of games. (GSP Auctions differ from games studied in [1] in that a player's payoff does not depend only on his rank).

3 Discrete Bidding Strategies

We focus on *conservative* bidders [8] that never outbid their valuation v_i in fear of receiving a negative payoff. Our discussion throughout the paper is in terms

of equal relevances and the RBB ranking rule. All results extend for RBR. We assume a discretization of the continuous strategy space $[0, v_i]$ of player i , in multiples of *bidding step* $\epsilon > 0$; i.e., the strategy space of i is $\Sigma_i = \{0, \epsilon, 2\epsilon, \dots, \lfloor v_i \rfloor_\epsilon\}$, where $\lfloor x \rfloor_\epsilon$ will henceforth denote *the maximum multiple of ϵ that is at most x* .

We view sponsored search auctions as repeated games, and we study the bidding strategies **AB** and **CB** in the context of iterative best response procedures. In each iteration, given a current configuration $\mathbf{b} = (b_1, \dots, b_n)$, a player i is chosen at random to respond to \mathbf{b}_{-i} by choosing a bid b'_i , so as to optimize his utility $u_i(\mathbf{b}_{-i}, b'_i)$. To do so, player i aims for the most profitable slot, $j^*(i)$, which he may win by a bid $b'_i \in (\mathbf{b}_{-i}(j^*(i)), \mathbf{b}_{-i}(j^*(i) - 1)]$; i.e., b'_i *strictly beats* $\mathbf{b}_{-i}(j^*(i))$ and equals at most $\mathbf{b}_{-i}(j^*(i) - 1)$, the bid issued by a player occupying slot $j^*(i) - 1$. Due to discretization and possible ties, it may occur that no $b'_i \in \Sigma_i$ grants the desired slot to i . Hence we define $j^*(i) = \arg \max_j [\theta_j(v_i - \mathbf{b}_{-i}(j))]$, where the max is taken over slots j for which $\Sigma_i \cap (\mathbf{b}_{-i}(j(i)), \mathbf{b}_{-i}(j(i) - 1)] \neq \emptyset$. If there is no such slot, then the bidder does not alter his bid. If bidder i is not occupying any slot under the current configuration \mathbf{b} , it may be the case that there is no slot giving him positive utility, in which case the bidder does not alter his bid either. Finally, if $j^*(i)$ equals the currently occupied slot by i , then i does not alter his bid. We consider two simple ways of selecting an extremal bid in this range, namely *Altruistic Bidding (AB)* and *Competitor Busting (CB)*.

Altruistic Bidding. **AB** [3] dictates that player i first computes his optimal slot $j^*(i)$ and then submits the most altruistic bid that is feasible and beats $\mathbf{b}_{-i}(j^*(i))$. Hence if $j^*(i) = 1$, he issues the bid $\mathbf{b}_{-i}(j^*(i)) + \epsilon$, otherwise he bids:

$$b'_i = \min[(\Sigma_i \cap \{\mathbf{b}_{-i}(j^*(i)) + \epsilon, \dots, \mathbf{b}_{-i}(j^*(i) - 1)\}) \setminus \{b_i\}]$$

Competitor Busting. **CB** expresses competitive behavior of player i , in that i incurs the highest possible payment to the player receiving the slot right above $j^*(i)$. We define the bid b'_i issued by i to be the maximum feasible bid that grants i slot $j^*(i)$, except if $j^*(i) = 1$. In this case set $b'_i = \mathbf{b}_{-i}(1) + \epsilon$, otherwise:

$$b'_i = \max[(\Sigma_i \cap \{\mathbf{b}_{-i}(j^*(i)) + \epsilon, \dots, \mathbf{b}_{-i}(j^*(i) - 1)\}) \setminus \{b_i\}]$$

Generally, b'_i *equals* (if feasible) $\mathbf{b}_{-i}(j^*(i) - 1)$, except for when $\mathbf{b}_{-i}(j^*(i) - 1) = b_i$, in which case $b'_i = \mathbf{b}_{-i}(j^*(i) - 1) - \epsilon$. This definition of **CB** differs from the one in [3], where $b'_i = \mathbf{b}_{-i}(j^*(i) - 1) - \epsilon$ *always*. Note that, assuming that $j^*(i)$ differs from currently occupied slot by i under \mathbf{b} , we forbid $b'_i = b_i$.

We need a *tie-breaking* rule, for when a newly submitted bid ties with an existing bid of another player. If bidder i best-responds by $b'_i = \mathbf{b}_{-i}(j')$ for slot j' then bidding b'_i grants i slot $j' + 1$ (or lower if there are more ties). For iterative best response this rule facilitates *dynamic temporal tie-breaking*, i.e. bidding the same bid as some player i' , but later than i' , may only grant a lower slot than i' .

Discretization of the players' strategy spaces in multiples of ϵ may introduce stable configurations that are not PNE in continuous strategies. Although **AB** and **CB** have seen experimental study in the recent literature [3], it is not known whether their induced state spaces maintain any PNE of the original game in

continuous strategies. By conditioning on ϵ , we establish existence of a socially optimum locally envy-free PNE, which is a discretized version of the PNE identified by Edelman *et al.* in [4]. Our result is additionally strengthened by the fact that, if our iterative best response procedures converge to a socially optimum configuration \mathbf{b} , then \mathbf{b} is a PNE of the game even with continuous strategies³ Let Δv denote the minimum among the distances between two valuations or the distance of a valuation from 0: $\Delta v = \min\{|v_i - v_j| : i, j \in N\} \cup \{|v_i| : i \in N\}$.

Theorem 1. *For any bidding step $\epsilon \leq \epsilon^* = (\gamma^{-1} - 1)\Delta v$, the configuration space of the GSP Auction game with discrete strategies contains at least one configuration \mathbf{b} , that is socially optimum and locally envy-free pure Nash equilibrium for the GSP Auction game even with continuous strategies, given by:*

$$b_j = \begin{cases} b_2 + \epsilon, & \text{if } j = 1 \\ \lfloor (1 - \gamma_j)v_j + \gamma_j b_{j+1} \rfloor_\epsilon, & \text{if } 2 \leq j \leq k \\ \lfloor v_j \rfloor_\epsilon, & \text{if } j \geq k + 1 \end{cases}$$

Also, if iterative AB or CB converges to a socially optimum configuration, then this is a pure Nash equilibrium of the GSP Auction game in continuous strategies.

Regarding the convergence of iterative AB/CB, we found examples showing that AB does not always converge, even for bidding step $\epsilon \leq \epsilon^*$ and geometrically decreasing (well separated) CTRs. We were not able to prove or disprove convergence of CB, despite extensive experimentation (reported in the full version). Resolving convergence for CB is therefore an interesting open problem. Convergence of local versions of these strategies – discussed next – also remains open.

4 Locally Aware Bidders & Local Stability

It is commonly assumed in the literature that bids of other players are observable. In principle one could apply learning techniques to estimate all the other bids as shown in [2]. Such a practice incurs however costs in time and money and, given the dynamic nature of these games, the game may have switched to a different bid vector by the time one estimates all remaining bids. Modeling the uncertainty about other bidders' offers is one approach to this issue [11]. Here we take a different approach and assume that bidders have only local knowledge about the bid vector and make only local moves, adhering to the following rules:

- 1.** They estimate the prices only for the slots right above or below their current slot and – in the absence of ties – will only move one slot upwards or downwards. In case of ties, a bidder learns the price of the first slot above or below him that he can actually target. If none of these moves are beneficial, no deviation occurs.
- 2.** Bidders not receiving a slot only learn the price of the last slot or – in case of ties – the price of the first slot from the end that they can target.

³ More accurately, *there is a tie-breaking rule* for the one-shot game in continuous strategies that renders \mathbf{b} a PNE. However, the socially optimum locally envy-free PNE described in Theorem 1 is independent of choice of tie-breaking rules.

The restrictions of AB/CB for such *locally aware* bidders (L-AB/L-CB) are natural strategies in this setting. If iterative L-AB or L-CB converge, they will converge to a *locally stable* configuration (in ϵ -discrete strategies), as in Definition 1. We analyze first the inefficiency of locally stable configurations *in continuous strategies*. Subsequently, we consider the performance of iterative L-AB and L-CB.

Theorem 2. *The GSP Auction game in continuous strategies with conservative bidders has Local Stability Ratio at least $\Omega(\sqrt{\alpha^k})$, for any constant $\alpha > 1$.*

In the proof of this result we used a game instance with $\gamma = 1$. However, fitting of real data in previous works [5] has shown that CTRs are well separated ($\gamma < 1$), by following a power law distribution. Geometrically decreasing CTRs $\theta_j \propto \alpha^{1-j}$ for $\alpha = 1.428$, were observed in [5]. Other authors [10] have used a Zipf distribution, where $\theta_j = j^{-\alpha}$, for $\alpha \geq 1$. For such cases with $\gamma < 1$ we obtain:

Theorem 3. *The GSP Auction game in continuous strategies with conservative bidders has Local Stability Ratio at most $(1 - \gamma)^{-1}$, assuming $\gamma < 1$.*

Corollary 1. *For geometrically decreasing click through rates with decay factor $\alpha > 1$ and conservative bidders, $\Lambda \leq \frac{\alpha}{\alpha-1}$. For click-through rates following the Zipf distribution with $\theta_j = j^{-\alpha}$, for $\alpha \geq 1$, $\Lambda \leq [1 - (1 - 1/k)^\alpha]^{-1} \leq k$.*

Corollary 1 and empirical observations [5] imply a constant upper bound on Λ for geometrically decreasing CTRs. We were not able to find matching lower bounds for Theorem 3 or Corollary 1. We give experimental results in figure 1, for the inefficiency of “reverse” assignments of players to slots, in games with $k = n$ slots, $n = 2, 3, \dots, 20$. The depicted results were found by solving non-linear programs (one for each curve), that express local stability of the reverse assignment and have Λ as objective function. Tightness of $\Lambda \leq k$ is evident for Zipf-distributed CTRs. Finally, our analysis for Theorem 3 can be used in bounding the inefficiency of stable configurations of iterative L-AB and L-CB:

Theorem 4. *For $\gamma < 1$ and $\epsilon \leq \epsilon^*$, the Local Stability Ratio of stable configurations with respect to iterative L-AB and L-CB is at most $(1 - \gamma)^{-1} + \gamma^{-1}$. Moreover, this bound applies to stable configurations with respect to AB and CB.*

References

1. Brandt, F., Fischer, F.A., Harrenstein, P., Shoham, Y.: Ranking Games. *Artificial Intelligence* 173(2), 221–239 (2009)
2. Brandt, F., Weiß, G.: Antisocial Agents and Vickrey Auctions. In: Meyer, J.J.C., Tambe, M. (eds.) *ATAL 2001*. LNCS, vol. 2333, pp. 335–347. Springer, Heidelberg (2002)
3. Cary, M., Das, A., Edelman, B., Giotis, I., Heimerl, K., Karlin, A.R., Mathieu, C., Schwarz, M.: Greedy bidding strategies for keyword auctions. In: MacKie-Mason, J.K., Parkes, D.C., Resnick, P. (eds.) *Proceedings of the 8th ACM Conference on Electronic Commerce (EC)*. pp. 262–271. ACM (2007)




Fig. 1. Local Stability Ratio Λ of Reversed Assignments.

4. Edelman, B., Ostrovsky, M., Schwartz, M.: Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review* 97(1), 242–259 (2007)
5. Feng, J., Bhargava, H.K., Pennock, D.M.: Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *INFORMS Journal on Computing* 19(1), 137–148 (2007)
6. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. *Computer Science Review* 3(2), 65–69 (2009)
7. Lahaie, S.: An analysis of alternative slot auction designs for sponsored search. In: Feigenbaum, J., Chuang, J.C.I., Pennock, D.M. (eds.) *Proceedings of the 7th ACM Conference on Electronic Commerce (EC)*. pp. 218–227. ACM (2006)
8. Leme, R.P., Tardos, E.: Pure and Bayes-Nash price of anarchy for generalized second price auction. In: *Proceedings of the 51st IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, to appear (2010)
9. Liang, L., Qi, Q.: Cooperative or Vindictive: Bidding strategies in Sponsored Search Auction. In: Deng, X., Graham, F.C. (eds.) *WINE 2007*. LNCS, vol. 4858, pp. 167–178. Springer, Heidelberg (2007)
10. Naldi, M., D’Acquisto, G., Italiano, G.F.: The value of location in keyword auctions. *Electronic Commerce Research and Applications* 9(2), 160–170 (2010)
11. Vorobeychik, Y.: Simulation-based Game Theoretic Analysis of Keyword Auctions with Low Dimensional Bidding Strategies. In: *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)* (2009), <http://www.cs.mcgill.ca/~uai2009/proceedings.html>
12. Vorobeychik, Y., Reeves, D.M.: Equilibrium analysis of dynamic bidding in sponsored search auctions. In: Deng, X., Graham, F.C. (eds.) *WINE 2007*. LNCS, vol. 4858, pp. 155–166. Springer, Heidelberg (2007)
13. Zhou, Y., Lukose, R.M.: Vindictive Bidding in Keyword Auctions. In: Gini, M.L., Kauffman, R.J., Sarppo, D., Dellarocas, C., Dignum, F. (eds.) *Proceedings of the International Conference on Electronic Commerce (ICEC)*. ACM International Conference Proceeding Series, vol. 258, pp. 141–146. ACM (2007)