

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

Evangelos Markakis*

Orestis Telelis†

Abstract

We study formally two simple discrete bidding strategies in the context of *iterative best response procedures*, for the game induced by the Generalized Second Price keyword auction mechanism. These strategies have seen experimental evaluation in the recent literature as parts of best response procedures, which have been shown not to converge. Here 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 best response state space *contains* socially optimum pure Nash equilibrium states of the game. We cast the strategies under a new light, arguing that they constitute natural choices for conservative myopic bidders that need to act based only on local information. For this case we provide bounds for the worst-case ratio of the social welfare of the reached locally stable states relative to the socially optimum welfare. Our work aims at providing a reference description for the strategies under study and their properties, thus filling a gap in the recent literature. We make several interesting observations experimentally, that give rise to challenging open problems.

*Department of Informatics, Athens University of Economics & Business, Greece (markakis@gmail.com)

†Department of Computer Science, University of Liverpool, United Kingdom (telelis@gmail.com)

1 Introduction

We study best response 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 they constitute the premiere source of income for search engines that allocate web page advertisement slots. The GSP mechanism is nowadays implemented in different forms by Google (through the AdWord system), Yahoo! and Bing (formerly MSN Live). Other online enterprises also use variants of this mechanism, for allocating advertisement slots in their websites; Google in particular, offers management (slot allocation and pricing) of advertisement slots in websites as a service, through its AdSense system.

The GSP mechanism is quite simple to implement. In its simplest form advertisers are ranked according to their bids from the highest to the lowest. If there are k available slots, each of the first k advertisers in the ranking wins the corresponding slot and pays a price equal to the bid of the next advertiser; bidders who do not win a slot do not pay anything. In the form that is mostly used now, every bid is also weighted by a relevance parameter of the advertiser. In case of a single slot the GSP mechanism is identical to the celebrated VCG mechanism [11]. For games with at least two slots however (in practice they can be 8 – 9), the GSP auction is different and *does not* retain some of the desirable characteristics of the VCG mechanism. In particular, truthful reporting of valuations is no longer a dominant strategy. The induced game also has a plethora of Nash equilibria, including equilibria in which overbidding occurs.

Encouragement of strategic behavior in GSP Auctions, which has also been observed in real data ever since the early days of keyword auctions [7], gives rise to the question of how should an advertiser decide on his bidding behavior. The strategy space of any advertiser is in principle infinite; a best response of a player i under a current bidding configuration is any bid value within the interval defined by the bids of at most two other players, that grants i his most desired slot; but how should the exact value be decided? Is there a reasonable characterization of the players' strategies that prescribes equilibria of the game reached under strategic behavior? In practice it is known that some bidders receive service by enterprises that undertake the task for designing bidding strategies for them, also referred to as auto-bidders [7]. Several phenomena of competition have been observed in the strategies adopted by the bidders ranging from planned conservatism in investing a modest budget, to harming competition by aggressive bidding so as to induce large prices for their competitors (and revenue to the auctioneer). These issues have received a lot of attention in the recent literature [6, 19]. Most of the existing works study bidding strategies as *iterative best response* procedures, viewing a GSP auction as a repeated game. Cary *et al.* [6] studied such strategies, where players adjust their bid iteratively, either synchronously or asynchronously in a randomly chosen order, always targeting the slot that maximizes their profits. They introduced three types of bidding strategies and proved convergence for the strategy referred to as *Balanced Bidding* (BB), which has a single fixed point, the locally envy-free equilibrium described by Edelman *et al.* [8].

We focus on the other two simple and intuitive strategies introduced in [6], that have received significantly less theoretical treatment, although they have been used in experimental comparisons [6, 15, 19]. 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 strategies require inevitable discretization of the players' strategy spaces by a parameter ϵ , which can be thought of as the monetary bidding unit. Such a discretization raises several issues however, as it may change the original game of continuous strategies entirely. Iterative best response procedures incorporating AB and CB have been observed not to converge for fixed values of ϵ [6, 15]; can we expect the

Λ	PNE	BNE	Locally Stable		
			Geom. CTR	Zipf CTR	General
Upper	$\phi < 1.62$ [14]	$\frac{2e}{(e-1)}$ [16]	$\frac{\alpha}{\alpha-1}, \alpha > 1$	k	$(1 - \gamma)^{-1}$
Lower	1.25913 [16]		1.95	4.446	$\Omega(\sqrt{c^k}), c > 1$

Table 1: **Bounds on Price of Anarchy and Local Stability Ratio.**

best response state space to even have one pure Nash equilibrium? How should this parameter be tuned so that the best response state space retains some of the properties of the original (continuous) game? The relevance of AB and CB is amplified in the context of bidders that, due to lack of complete information, restrain themselves to local conservative best responses, i.e. local moves targeting a slot above or below the currently allocated slot. In this context an interesting question that arises is whether such locally stable profiles suffer significant efficiency loss. The theme of our work is a formal treatment of AB and CB in relation to the original game of continuous strategies and under the light of locally aware bidders.

Contribution. We study iterative best response procedures embodying AB and CB and differ from previous work [6] in that bidders only update their bid when they find it more profitable to target a different slot than the one they currently occupy. We provide a detailed description of both bidding strategies and, in fact, redefine CB slightly differently than it has appeared in the literature, to ensure its consistency with theoretical developments to follow (Section 4). We decide an upper bound on the discretization parameter ϵ to ensure that the induced best response state space does have a socially optimum pure Nash equilibrium analogous to the one identified by Edelman et al. [8], that is a pure Nash equilibrium for the original game (i.e., in continuous strategies). Also, we ensure that if iterative AB or CB converge to a socially optimum configuration, then this is a pure Nash equilibrium even in continuous strategies. Subsequently we turn our attention to the case of bidders that take local steps due to incompleteness of available information to them (Section 5). We investigate the social inefficiency of configurations where players have no incentive to move to the slot directly above or below their current slot, and introduce the notion of Local Stability Ratio to quantify the efficiency loss, in analogy to the Price of Anarchy. We provide upper and lower bounds on the Local Stability Ratio for general instances as well as for certain families of CTR distributions (in particular, *Geometrically* decreasing and *Zipf*), and we then use our analysis to produce upper bounds for stable configurations reached by locally aware adaptations of iterative AB and CB (termed L-AB and L-CB). Our results on the Local Stability Ratio are summarized in Table 1, where we also compare them with the known results on the Price of Anarchy (discussed in detail in section 2). Finally, in Section 6 we provide experimental evidence on the convergence of CB L-CB, and L-AB and conclude.

2 Related Work

A considerable amount of work in sponsored search auctions concerns the strategic behavior of the bidders and the analysis of appropriate bidding strategies. Asdemir in [1] considered an auction game with two bidders that play interchangeably by applying strategies based on their payoff history. Markov Perfect Equilibria were identified for this extensive form game. Cary *et al.* [6] investigated theoretically iterative bid adjustment strategies, viewing GSP auctions as a repeated game. The authors formally defined and studied three strategies, namely *Altruistic Bidding* (AB), *Competitor Busting* (CB) and *Balanced Bidding* (BB). CB, also referred to as vindictive bidding, has

been observed quite often in practice [4, 21]. It is an aggressive bidding behavior where 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 balanced bidding is an attempt to achieve a balance between these two extremes. For BB the authors showed that, under some conditions, it converges to the efficient locally envy-free equilibrium characterized in [8, 18]. For AB and CB it was shown that they do not generally converge. The authors also analyzed experimentally AB and CB to conclude that AB results generally in low revenue, while CB results in high revenue.

The properties of AB BB and CB have been studied further in other works as well. The performance of these strategies are analyzed in Bayesian settings in [20, 17]. In [10] a mechanism is proposed so as to discourage bidders from exhibiting busting behavior. In [15], vindictive strategies are studied for games where bidders can choose to either be vindictive or cooperative. Convergence to different notions of equilibrium have also been considered in [5]. For bidders participating in multiple keyword auctions at the same time, a bidding protocol was designed in [2], where in each round each bidder has to decide on splitting his budget among the various keywords. The convergence of their protocol was analyzed both empirically and theoretically. We do not consider here such settings but focus only on auctions for a single keyword.

Regarding the efficiency of equilibria in GSP auctions, the first upper bounds on the Price of Anarchy with respect to the social welfare were obtained by Lahaie [13]. Recently, tighter upper bounds were obtained for *conservative* bidders by Paes Leme and Tardos in [14]. It was shown that for bidders that do not outbid their valuation, the price of anarchy is at most equal to the golden ratio $\phi < 1.62$ for the complete information game and at most 8 for the Bayesian setting. For the Bayesian Price of Anarchy, Lucier and Paes Leme proved recently an improved upper bound of $\frac{2e}{e-1}$ and provided a lower bound of 1.25913 for the Price of Anarchy of pure equilibria, improving on a simpler lower bound of 1.25 from [14]. These results are summarized in table 1.

3 Definitions & Preliminaries

An instance of the Generalized Second Price (GSP) Auction game consists of a set of players $N = \{1, \dots, n\}$ players, a set of k slots and a tuple $\langle \{\theta_j\}_{j=1}^k, \{\rho_i\}_{i=1}^n, \{v_i\}_{i=1}^n \rangle$, where $\theta_j \in [0, 1]$ denotes the probability that a link displayed in slot j is clicked by the viewer (*Click-Through Rate* - CTR), $\rho_i \in [0, 1]$ is the *relevance* of advertiser i – the probability that an advertisement by i is clicked – and v_i is the valuation of advertiser i . We refer to advertisers also as players or bidders. We use \hat{v}_i for the product $\rho_i v_i$, the *expected revenue* of player i . We assume indexing of slots and players so that: $\theta_1 \geq \theta_2 \geq \dots \geq \theta_k > 0$ and $\hat{v}_1 \geq \hat{v}_2 \geq \dots \geq \hat{v}_n$. For every slot $j \geq 2$ we define $\gamma_j = \theta_j / \theta_{j-1}$ and let $\gamma = \max_j \gamma_j$.

The GSP Mechanism. In a GSP auction, the players issue collectively a bid vector $\mathbf{b} = (b_1, b_2, \dots, b_n)$. The mechanism ranks players by order of non-increasing *declared expected revenue* $\hat{b}_i = \rho_i b_i$ and matches them with the slots in order of non-increasing CTR respectively. This is the *Rank-By-Revenue* (RBR) rule. When all bidders' relevances are equal, the players are ranked by non-increasing bid (*Rank-By-Bid* rule - RBB) b_i . Under RBB every player i receiving a slot j pays the bid of the player i' receiving the slot $j + 1$. Under RBR, the receiver of slot j pays the declared expected revenue of the receiver of slot $j + 1$ divided by his relevance ρ_i , i.e. $\hat{b}_{i'} / \rho_i = \rho_i b_{i'} / \rho_i$.

Given a bid configuration \mathbf{b} , we reserve notation $b_{(j)}$, $\rho_{(j)}$, $v_{(j)}$, to refer to the bid, relevance and valuation of the player occupying slot j under the used ranking rule. $\beta_{(j)} = \theta_j \rho_{(j)}$. By \mathbf{b}_{-i} we refer to the strategy profile \mathbf{b} without the bid of player i and $\mathbf{b}_{-(j)}$ will denote exclusion of the bid

of the player occupying slot j . Finally, we 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, i.e., in the absence of i or of the person who occupied slot 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:

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

The *social welfare* of \mathbf{b} is denoted by $SW(\mathbf{b})$ and defined as: $SW(\mathbf{b}) = \sum_{j=1}^k \theta_j \hat{v}_{(j)} = \sum_{j=1}^k \theta_j \rho_{(j)} v_{(j)}$.

We assume here that the search engine has a deterministic tie-breaking rule in case there are ties in the ranking. All our results are independent of the tie-breaking rule that is used. Edelman *et al.* [8] identified a pure Nash equilibrium 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 by players equal to their payments 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 the slot right beyond the one he is assigned under \mathbf{b}^* , at price $\rho_i b_i^*$. Formally, \mathbf{b} satisfies: $\theta_j (\hat{v}_{(j)} - \hat{b}_{(j+1)}) \geq \theta_{j-1} (\hat{v}_{(j)} - \hat{b}_{(j)})$. The total revenue of the auctioneer in any locally envy free equilibrium is at least as large as in \mathbf{b}^* .

Local Stability. In most of the literature on analysis of bidding strategies, it is assumed that the bids of the other players as the game evolves is common knowledge. This assumption is often justified by the fact that the bidders can potentially apply learning techniques to estimate their opponents' bids. However, learning the remaining bids with such techniques incurs a cost (in time, budget, effort, etc) and given the dynamic nature of these games, it is conceivable that by the time one estimates all the remaining bids, the game may have already switched to a different bid vector. With this in mind, we are interested in studying games where players only make local moves and are not aware of the entire bid vector but only of a small part of it. In particular, in Section 5, we will assume that a player, when deciding to update his strategy, only makes the effort to learn the price of the slots above and below the slot he is currently occupying and is only considering these local moves. In the case of ties, i.e., many people below him bidding the same, we will make the assumption that he actually learns the price of the first slot below the ties. This gives rise to the following definition, which is a weakening of the notion of Nash equilibrium¹.

Definition 1 (Local Stability) *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

(a) If $j_0 \neq k$ and $j'_1 \leq k + 1$, then: $\theta_{j_0} (\hat{v}_{(j_0)} - \hat{b}_{(j_0+1)}) \geq \theta_{j'_1-1} (\hat{v}_{(j_0)} - \hat{b}_{(j'_1)})$

(b) If $j_0 \neq 1$, then: $\theta_{j_0} (\hat{v}_{(j_0)} - \hat{b}_{(j_0+1)}) \geq \theta_{j'_2+1} (\hat{v}_{(j_0)} - \hat{b}_{(j'_2+2)})$

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

¹Note also that a locally envy-free equilibrium is locally stable but the reverse is not true. A locally stable configuration is not necessarily locally envy-free.

A few comments are in order. The intuition behind the definition is simply that no player has an incentive to move upwards or downwards to the next feasible slot given the configuration \mathbf{b} . The indices j'_1 and j'_2 in the definition above are used in order to determine the slot that the bidder at slot j_0 can actually target in case that due to ties he cannot target the slot right above and/or right below him. The condition $j'_1 \leq k + 1$ in (1a) expresses the fact that a bidder may not be able to make a downward deviation if all the remaining bidders have equal score. For players who do not win a slot, we assume that they know the *bidding entry level* to competition, i.e. the quantity $\hat{b}_{(k)} = \rho_{(k)} b_{(k)}$ that wins the last slot. The last constraint above prescribes that no such bidder has an incentive to win the last slot. Any pure Nash equilibrium is of course locally stable. In analogy to the *Price of Anarchy* [12], 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* [3] constitute one such interesting class of games. (GSP Auctions differ from games studied in [3] in that a player's payoff does not depend only on his rank).

4 Discrete Bidding Strategies

We will consider a discretization of the strategy space of the bidders as we feel that this is a more accurate model of real auctions. Upon such a discretization, we define and study two natural bidding strategies that have been proposed in previous works, in order to capture altruistic and vindictive behavior. We depart from the existing literature on these heuristics by introducing variants of them, so as to eliminate certain sources of instability as we explain below. For simplicity, our description in this section and the proofs throughout the paper are in terms of equal relevances of bidders and the RBB ranking rule. All the results straightforwardly extend for the RBR rule.

We assume a discretization of the continuous strategy space of a player i into bids that are multiples of a *bidding step* $\epsilon > 0$. Thus, given $\epsilon > 0$ we assume that the strategy space of every player i is $\Sigma_i = \{r \times \epsilon | r \in \mathbb{Z}^+\}$. When players are assumed to be *conservative*, i.e. never outbid their valuation, the discrete strategy space of each player i will be restricted to $\Sigma_i = \{r \times \epsilon | r \in \mathbb{Z}^+ : r \times \epsilon \leq v_i\}$ or, equivalently, $\Sigma_i = \{0, \epsilon, 2\epsilon, \dots, \lfloor v_i \rfloor_\epsilon\}$, where $\lfloor v_i \rfloor_\epsilon$ will henceforth denote *the maximum multiple of ϵ that is at most v_i* ².

We view sponsored search auctions as repeated games, and study bidding strategies in the context of iterative best response procedures. We consider an asynchronous model³ of bid updating, where 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, denoted by $j^*(i)$, which can be granted to i by a bid $b'_i \in (\mathbf{b}_{-i}(j^*(i)), \mathbf{b}_{-i}(j^*(i) - 1)] \setminus \{b_i\}$, i.e., a bid other than b_i that *strictly beats* $\mathbf{b}_{-i}(j^*(i))$ and at most matches $\mathbf{b}_{-i}(j^*(i) - 1)$; the latter is issued by a player occupying the slot immediately above

² $\epsilon \leq \min_i v_i$ is assumed.

³We do not consider synchronous models of updating bids as the asynchronous model captures better how these games are played in practice.

$j^*(i)$. Due to discretization and possible presence of ties, it may occur that there is no $b'_i \in \Sigma_i$ that grants the desired slot to i (e.g. if the bid of the player right above equals the price of $j^*(i)$). 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 $(\mathbf{b}_{-i}(j(i)), \mathbf{b}_{-i}(j(i) - 1)) \setminus \{b_i\} \neq \emptyset$. If there is no such slot, then the bidder may not alter his current 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. Furthermore, we impose one more rule on the strategy of player i that differs from the existing definitions of such best response bidding strategies: if the optimal slot he computes, $j^*(i)$, is the same as the slot he is currently occupying, then the bidder does not alter his bid.

We consider two simple ways of selecting an extremal within the range $(\mathbf{b}_{-i}(j^*(i)), \mathbf{b}_{-i}(j^*(i) - 1)) \setminus \{b_i\}$, referred to as *Altruistic Bidding* (AB) and *Competitor Busting* (CB) respectively.

Altruistic Bidding. AB was first studied in [6] and dictates that player i first computes his optimal slot $j^*(i)$ and then submits the most altruistic bid that is feasible and beats $\mathbf{b}_{-i}(j^*(i))$. If $j^*(i) = 1$, he issues the bid $\mathbf{b}_{-i}(j^*(i)) + \epsilon$. Otherwise:

$$b'_i = \begin{cases} \min[(\Sigma_i \cap \{\mathbf{b}_{-i}(j^*(i)) + \epsilon, \dots, \mathbf{b}_{-i}(j^*(i) - 1)\}) \setminus \{b_i\}] & \text{if } i \text{ is conservative} \\ \min(\{\mathbf{b}_{-i}(j^*(i)) + \epsilon, \dots, \mathbf{b}_{-i}(j^*(i) - 1\}) \setminus \{b_i\}) & \text{otherwise} \end{cases}$$

Competitor Busting. CB (also referred to as *Vindictive Bidding*) has been studied more extensively than AB since observations of real data has shown that it occurs in practice quite frequently [7]. CB expresses competitive behavior of player i , in that he makes the payment incurred to the player receiving the slot right above $j^*(i)$, as high as possible. We define the bid b'_i issued by i in this case to be the maximum feasible bid he can issue to win his desired slot, except when he targets the first slot. Hence if $j^*(i) = 1$, he issues the bid $\mathbf{b}_{-i}(j^*(i)) + \epsilon$, otherwise he bids:

$$b'_i = \begin{cases} \max[(\Sigma_i \cap \{\mathbf{b}_{-i}(j^*(i)) + \epsilon, \dots, \mathbf{b}_{-i}(j^*(i) - 1)\}) \setminus \{b_i\}] & \text{if } i \text{ is conservative} \\ \max(\{\mathbf{b}_{-i}(j^*(i)) + \epsilon, \dots, \mathbf{b}_{-i}(j^*(i) - 1\}) \setminus \{b_i\}) & \text{otherwise} \end{cases}$$

We point out some technical but important details regarding the definitions: notice that since we assume that $j^*(i)$ is different than the slot that bidder i is currently occupying under \mathbf{b} , we forbid i to play $b'_i = b_i$ and gain $j^*(i)$. This is done in order to avoid certain pathological situations that may arise in the presence of ties and lead to cycling behavior. Furthermore, we allow (if it is feasible) b'_i to equal the bid $\mathbf{b}_{-i}(j^*(i) - 1)$, except for the case where $\mathbf{b}_{-i}(j^*(i) - 1) = b_i$. In this case one has to play $\mathbf{b}_{-i}(j^*(i) - 1) - \epsilon$. This differentiates the definition from the ones used previously in [6], where a player plays always a bid smaller by a term ϵ than $\mathbf{b}_{-i}(j^*(i) - 1)$.

For AB and CB, we also need a *tie-breaking* rule for the cases where a newly submitted bid ties with an existing bid of another player. Given a configuration \mathbf{b} , if bidder i best-responds by $b'_i = \mathbf{b}_{-i}(i')$ for some player $i' \neq i$ who receives slot j' under \mathbf{b}_{-i} , then bidding b'_i grants i slot $j + 1$ (or lower if there are more ties). In the context of iterative best response procedures this rule facilitates *dynamic temporal tie-breaking*, i.e. a player bidding the same bid as some player i' but later than i' may only be granted a lower slot than i' .

AB and CB as defined both here and in [6] are iterative best response procedures. The most crucial difference between our definitions and those of [6] is the fact that in [6] a bidder may update his bid even when the optimal slot for him has not changed, unlike our definitions. Thus our version, especially for CB, gives rise to a less aggressive bidding behavior that is likely to overcome certain sources of instability and cycling. This is demonstrated further with experimental evidence in

Section 6. Naturally, discretization of the players' strategy spaces in multiples of the bidding step ϵ , may introduce certain configurations that are stable, but not necessarily pure Nash equilibria in continuous strategies. In the next section we identify an upper bound on the bidding step that guarantees existence of a socially optimum pure Nash equilibrium state in the discretized space. We refer to the non-conservative versions of the strategies as *unrestricted*.

4.1 Socially Optimal Pure Nash Equilibria

Although AB and CB have seen experimental study in the recent literature [6], it is not known whether the induced configuration spaces of these strategies do actually maintain any pure Nash equilibrium of the original game in continuous strategies. Here, by conditioning on the bidding step ϵ we establish for conservative players the existence of a socially optimum locally envy-free pure Nash equilibrium, which is a discretized version of the equilibrium identified by Edelman *et al.* in [8]. Our result is additionally strengthened by the fact that, if our (unrestricted) iterative best response procedures converge to a socially optimum configuration \mathbf{b} , then \mathbf{b} is a pure Nash equilibrium of the game even with continuous strategies. For ease of exposition, assume⁴ that $v_1 > v_2 > \dots > v_n > 0$ and let Δv denote the minimum among the distances between two valuations or the distance of a valuation from 0: $\Delta v = \min\{\{|v_i - v_j| : i, j \in N\} \cup \{|v_i| : i \in N\}\}$. Assume also that $\theta_1 > \theta_2 > \dots > \theta_k > 0$. Define a critical bidding step $\epsilon^* = \min\{\gamma, 1 - \gamma\} \times \Delta v$. We show first the following:

Theorem 1 *For any bidding step $\epsilon \leq \epsilon^*$ and for conservative bidders, the configuration space of the GSP Auction game with discrete strategies contains at least one configuration \mathbf{b} , that is socially optimum pure Nash equilibrium for the GSP Auction game even with continuous strategies, where:*

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

Proof. First we verify that the bids are indeed conservative. For $j \geq k + 1$, this is obvious. For $2 \leq j \leq k$, we can prove it inductively starting from k . It is easy to verify that $b_k < v_k$. Then for any other $j < k$, we have that $b_j \leq (1 - \gamma_j)v_j + \gamma_j b_{j+1}$. But by the induction hypothesis $b_{j+1} \leq (1 - \gamma_{j+1})v_{j+1} + \gamma_{j+1} b_{j+2} < v_{j+1}$. Finally for $j = 1$, we have that:

$$\begin{aligned} b_1 = b_2 + \epsilon &\leq (1 - \gamma_2)v_2 + \gamma_2 b_3 + \epsilon \\ &\leq (1 - \gamma_2)v_2 + \gamma_2 b_3 + (1 - \gamma_2)(v_1 - v_2) \\ &= \gamma_2 b_3 + (1 - \gamma_2)(v_1 - v_2) < v_1 \end{aligned}$$

The second inequality above follows from the assumption we have made on the bidding step.

Then we verify that the bids are properly ordered so that the emerging ranking of players is the socially optimum one. For $j_1 > j_2 > k$, $b_{j_1} = \lfloor v_{j_1} \rfloor_\epsilon \leq \lfloor v_{j_2} \rfloor_\epsilon = b_{j_2}$, because $v_{j_1} < v_{j_2}$. For the j -th slot, with $2 \leq j \leq k$, we have:

$$\begin{aligned} b_j &= \lfloor (1 - \gamma_j)v_j + \gamma_j b_{j+1} \rfloor_\epsilon \geq (1 - \gamma_j)v_j + \gamma_j b_{j+1} - \epsilon \\ &\geq (1 - \gamma_j)v_j + \gamma_j b_{j+1} - (1 - \gamma_j)(v_j - v_{j+1}) \\ &= b_{j+1} + (1 - \gamma_j)(v_j - b_{j+1}) - (1 - \gamma_j)(v_j - v_{j+1}) > b_{j+1} \end{aligned}$$

⁴If there are identical valuations or valuations equal to 0, the definition of Δv would be adjusted to consider only differences among non-identical valuations and also absolute values of non-zero valuations.

The latter strict inequality follows by the fact that $b_{j+1} < v_{j+1}$. We show next that \mathbf{b} constitutes a pure Nash equilibrium even in continuous strategies. For any two distinct slots $j < i$ the following should hold at a Nash equilibrium:

$$\left. \begin{array}{l} \theta_j(v_j - b_{j+1}) \geq \theta_i(v_j - b_{i+1}) \\ \theta_i(v_i - b_{i+1}) \geq \theta_j(v_i - b_j) \end{array} \right\} \Rightarrow \left\{ \begin{array}{l} b_{j+1} \leq [1 - (\theta_i/\theta_j)]v_j + (\theta_i/\theta_j)b_{i+1} \\ (1 - (\theta_i/\theta_j))v_i + (\theta_i/\theta_j)b_{i+1} \leq b_j \end{array} \right. \quad (2)$$

We prove validity of inequalities (2) inductively. In fact, we prove these inequalities for any $j \leq i$. For the first one of (2), we fix $i > 1$ and consider j ranging from i down to 1. For $j = i$ we obtain $b_{j+1} = b_{i+1} \leq b_{i+1}$. Consider now a slot j , such that $j + 1 \leq i$, $j \geq 1$. By the induction hypothesis, we have: $b_{j+2} \leq [1 - (\theta_i/\theta_{j+1})]v_{j+1} + (\theta_i/\theta_{j+1})b_{i+1}$. Then:

$$\begin{aligned} b_{j+1} &\leq (1 - \gamma_{j+1})v_{j+1} + \gamma_{j+1}b_{j+2} \\ &\leq (1 - \gamma_{j+1})v_{j+1} + \gamma_{j+1} [(1 - (\theta_i/\theta_{j+1}))v_{j+1} + (\theta_i/\theta_{j+1})b_{i+1}] \\ &= [1 - (\theta_i/\theta_j)]v_{j+1} + (\theta_i/\theta_j)b_{i+1} \leq [1 - (\theta_i/\theta_j)]v_j + (\theta_i/\theta_j)b_{i+1} \end{aligned}$$

For the second inequality from (2), fix $j < k$ and let i range from j to k . We obtain trivially $b_j \geq b_{j+1}$ for $i = j$. We also show that for $i = j + 1$, the desired inequality holds. We have already shown that $b_j > b_{j+1}$. Hence, when $i = j + 1$:

$$b_j \geq b_{j+1} + \epsilon \geq (1 - \gamma_{j+1})v_{j+1} + (\gamma_{j+1})b_{j+2} = [1 - (\theta_i/\theta_j)]v_i + (\theta_i/\theta_j)b_{i+1}$$

Consider a slot i with $i - 1 > j$, $i \leq k$. By induction hypothesis, we have: $[1 - (\theta_{i-1}/\theta_j)]v_{i-1} + (\theta_{i-1}/\theta_j)b_i \leq b_j$. Then:

$$\begin{aligned} b_j &\geq [1 - (\theta_{i-1}/\theta_j)]v_{i-1} + (\theta_{i-1}/\theta_j)b_i \geq \\ &\geq [1 - (\theta_{i-1}/\theta_j)]v_{i-1} + (\theta_{i-1}/\theta_j)[(1 - \gamma_i)v_i + \gamma_i b_{i+1} - \epsilon] \\ &= [1 - (\theta_i/\theta_j)]v_i + (\theta_i/\theta_j)b_{i+1} + [1 - (\theta_{i-1}/\theta_j)](v_{i-1} - v_i) - (\theta_{i-1}/\theta_j)\epsilon \\ &\geq [1 - (\theta_i/\theta_j)]v_i + (\theta_i/\theta_j)b_{i+1} + (1 - \gamma)\Delta v - \gamma(1 - \gamma)\Delta v \\ &\geq [1 - (\theta_i/\theta_j)]v_i + (\theta_i/\theta_j)b_{i+1} + (1 - \gamma)^2\Delta v \geq [1 - (\theta_i/\theta_j)]v_i + (\theta_i/\theta_j)b_{i+1} \end{aligned}$$

Stability of players $i > k$ is ensured, since they already bid the maximum they can. \square

Proposition 1 *For $\epsilon \leq \epsilon^*$, if unrestricted iterative AB or CB converge to a socially optimum and conservative configuration, then this is a pure Nash equilibrium of the GSP Auction game in continuous strategies.*

Proof. Let \mathbf{b} denote a stable configuration for (unrestricted) iterative AB or CB, with respect to a GSP Auction game instance. We call a bidding profile \mathbf{b} *upwards stable* if no player has an incentive to move to a higher slot than the one he is allocated. In the same manner, we call a profile *downwards stable* if no player has an incentive to move to a lower slot. The proof consists of the following two Lemmas:

Lemma 1 *Let \mathbf{b} denote a stable configuration for (unrestricted) iterative AB or CB for any $\epsilon > 0$. Then \mathbf{b} is upwards stable even in continuous strategies.*

Proof. Since \mathbf{b} is a stable configuration for AB or CB, it is upwards stable in discrete strategies. Suppose that there was a feasible deviation in the continuous strategy space of a player for a higher slot, say slot i . This means that $b_{(i-1)} \neq b_{(i)}$ (if $i = 1$, the remaining argument still holds). But since bidders play according to AB or CB, the gap between these two bids has to be at least ϵ . But then such a deviation would always be possible also by usage of discrete strategies, because players are unrestricted. In particular the bid $b_{(i)} + \epsilon$ would grant slot i to the player. \square

Lemma 2 *Let \mathbf{b} denote a stable configuration for (unrestricted) iterative AB or CB, with respect to a GSP Auction game instance and suppose $\epsilon \leq \epsilon^*$. If \mathbf{b} is a conservative and socially optimum profile, then it is downwards stable even in continuous strategies.*

Proof. We give the proof for CB; an analogous proof for AB appears in **Appendix A**. Suppose that the claim of the Lemma is false. Let $i_0 = (j_0)$ be a player occupying slot j_0 under \mathbf{b} and assume that it is feasible to target a slot $j > j_0$ with a continuous strategy that would incur i_0 utility larger than $u_{i_0}(\mathbf{b})$.

Suppose first that $j \leq k - 1$. Then since targeting slot j is feasible by i_0 , we have that $b_{(j)} \neq b_{(j+1)}$ and in particular $b_{(j+1)} \leq b_{(j)} - \epsilon$, as all bids evolve as multiples of ϵ . If $b_{i_0} \neq b_{(j)}$, then player i_0 could target slot j with the CB strategy by bidding $b_{(j)}$ and improve his utility, a contradiction. Hence suppose that $b_{i_0} = b_{(j)}$. In this case, the CB strategy would prescribe player i_0 to bid $b_{(j)} - \epsilon$ (recall that we defined AB and CB so that one cannot win a different slot from the one currently allocated to him by using the same bid). If $b_{(j+1)} < b_{(j)} - \epsilon$, this implies that again, player i_0 could win slot j and improve his utility by following CB, a contradiction. Hence the only remaining case is when both $b_{i_0} = b_{(j)}$ and $b_{(j+1)} = b_{(j)} - \epsilon$. Then however, if there was a continuous strategy b' that would grant slot j to i_0 , we would have:

$$\begin{aligned}
u_{i_0}(b', \mathbf{b}_{-i}) &= \theta_j(v_{i_0} - b_{(j+1)}) &= \theta_j(v_{i_0} - b_{(j)} + \epsilon) \\
&= \theta_{j_0}(v_{i_0} - b_{(j)}) + (\theta_j - \theta_{j_0})(v_{i_0} - b_{(j)}) + \theta_j \epsilon \\
&\leq \theta_{j_0}(v_{i_0} - b_{(j)}) + (\theta_j - \theta_{j_0})(v_{i_0} - b_{(j)}) + (\theta_{j_0} - \theta_j)(v_{i_0} - v_{(j)}) \\
&\leq \theta_{j_0}(v_{i_0} - b_{(j)}) = u_{i_0}(\mathbf{b})
\end{aligned}$$

which means that i_0 is not better off by aiming for slot j . In the analysis above, the social optimality of \mathbf{b} guarantees that $v_{i_0} \geq v_{(j)}$. Hence we can substitute ϵ by $(1 - \gamma)(v_{i_0} - v_{(j)})$, which in turn is at most $(\gamma^{-1} - 1)(v_{i_0} - v_{(j)})$. Finally the last inequality is derived by the assumption that \mathbf{b} was a conservative bidding profile. \square

Lemmas 1 and 2 complete the proof of the Proposition. \square

Regarding convergence of AB and CB, we defer the discussion to Section 6.

5 Locally Aware Bidders & Local Stability

In this Section we turn our attention to the social inefficiency of locally stable configurations with respect to local versions of AB and CB for conservative players and an analysis of the local stability ratio, defined in Section 3. Most of the analysis of bidding strategies in the literature is based on the quite strong assumption that the bids of the other players as the game evolves are observable. In principle one could apply learning techniques to estimate all the other bids as shown in [4]. Such a practice incurs however costs in time and budget and, given the dynamic nature of these

games, it is conceivable that by the time one estimates all the remaining bids, the game may have already switched to a different bid vector. One way to approach this issue is to introduce models of uncertainty about other bidders' offers as is considered e.g. by Vorobeychik in [19]. Here we take a different approach and focus on bidders who have only local knowledge about the bid vector and make only local moves. In particular, we will assume they abide by the following rules, when they consider updating their bid:

1. Bidders will only try to estimate the prices for the slots right above or below their current slot and they will only consider moving one slot upwards or downwards from a given bid configuration. In the case of ties, we assume that the bidder learns the price of the first slot above or below him that he can actually target (e.g., if many people below a certain bidder bid the same then it is not possible to target the slot right below him, hence he may only target the first slot below the ties).
2. For bidders not receiving a slot, we assume they only learn the *entry level of competition*, i.e., the price of the last slot or, in case of ties, the price of the first slot from the end that they may target.

We refer to such bidders as *locally aware*. The restrictions of AB and CB to locally aware bidders (referred to as L-AB and L-CB) are very natural choices in this setting. And if they converge, they will converge to a *locally stable* configuration (in discrete strategies defined by the used bidding step ϵ), as given by Definition 1. We start by showing that in general the social inefficiency $\mathbf{\Lambda}$ of locally stable configurations can be very high (even in continuous strategies).

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

Proof. We construct an instance of the game with $\rho_i = 1$ for every bidder i , i.e., the bound holds even for the rank-by-bid rule. Let the number of slots be $k = 2\lambda$ for some integer $\lambda \geq 1$ and suppose there are $k + 1$ players. For $j = 1, \dots, \lambda$, define $\theta_j = c^{1-j}$ for some constant $c > 1$. For $j = \lambda + 1, \dots, k$ let $\theta_j = c^{1-\lambda} = \theta_\lambda$. For the valuations of players, define $v_j = c^{1-j}$, for $j = 1, \dots, k$ and $v_{k+1} = \delta$ for some small constant $0 < \delta < v_{2\lambda}$. First we describe the assignment and compute its inefficiency and then we show existence of a bid configuration that results in this assignment and is in fact locally stable. We will consider the assignment of each player $j = 1, \dots, \lambda$ to slot $\lambda + j$ and assignment of players $j = \lambda + 1, \dots, k$ to slots $j - \lambda$. The social welfare for a bid configuration, \mathbf{b} , that results in this assignment is:

$$\begin{aligned}
SW(\mathbf{b}) &= \sum_{j=1}^{\lambda} c^{1-j} c^{1-j-\lambda} + \sum_{j=1+\lambda}^{2\lambda} c^{1-\lambda} c^{1-j+\lambda} \\
&= c^{2-\lambda} \sum_{j=1}^{\lambda} c^{-2j} + c^2 \sum_{j=1+\lambda}^{2\lambda} c^{-j} \\
&= c^{2-\lambda} \sum_{j=1}^{\lambda} c^{-2j} + c^{2-\lambda} \sum_{j=1}^{\lambda} c^{-j} \leq 2c^{2-\lambda} \sum_{j=1}^{\lambda} c^{-j} \\
&\leq \frac{2c^{3-\lambda}}{c-1}
\end{aligned}$$

For the socially optimum assignment we have: $SW(\mathbf{b}^*) = \sum_{j=1}^{\lambda} c^{2-2j} + \sum_{j=1+\lambda}^{2\lambda} c^{1-\lambda} c^{1-j} \geq 1$. Thus we obtain the lower bound: $\mathbf{\Lambda} \geq SW(\mathbf{b}^*)/SW(\mathbf{b}) \geq \frac{c-1}{2} c^{\lambda-3} = \Omega(c^{k/2})$. We now define a bid vector \mathbf{b} that produces the desired assignment:

$$b_i = \begin{cases} \delta & \text{if } i \in \{1, \dots, \lambda\} \cup \{k+1\} \\ (1 - \gamma_{j+1})v_{(j+1)} + \gamma_{j+1}b_{(j+2)} & \text{if } i \in \{\lambda+1, \dots, 2\lambda\} \end{cases} \quad (3)$$

It is easy to see that for $j = 1, \dots, \lambda$, the bids $b_{(j)}$ are distinct, corresponding respectively to the players $\lambda+1, \dots, 2\lambda$. As a tie-breaking rule to determine the assignment, we will use the rule that favors players with higher index. Hence with vector \mathbf{b} , for $j = 1, \dots, \lambda$, player j will occupy respectively slot $j + \lambda$ and player $k + 1$ will not get a slot. It is easy to see though that our result is independent of the tie-breaking rule. Even if we had a different assignment for the last λ slots this would not affect neither the local stability, nor the asymptotic behavior of $\mathbf{\Lambda}$ (since CTRs are identical for the last λ slots). For adequately small $\delta < v_{2\lambda}$ the bids are conservative for the players who occupy slots $j = \lambda, \dots, 2\lambda$. For $j = 1, \dots, \lambda - 1$ the bids are $b_{(j)} = \frac{c-1}{c}v_{(j+1)} + \frac{1}{c}b_{(j+2)}$ and they can also be shown to be conservative by induction, since $v_{(j)} \geq v_{(j+1)}$.

It remains to show that \mathbf{b} is locally stable. We first show local stability with respect to *upward* deviations. For a player occupying slot j , for $j = 1, \dots, \lambda$, he may target slot $j - 1$ at price $b_{(j-1)}$ and then his payoff would be: $\theta_{j-1}(v_{(j)} - b_{(j-1)}) = \theta_{j-1}(v_{(j)} - (1 - \gamma_j)v_{(j)} - \gamma_j b_{(j+1)}) = \theta_{j-1}(\gamma_j v_{(j)} - \gamma_j b_{(j+1)}) = \theta_j(v_{(j)} - b_{(j+1)})$. But this is precisely his current payoff, hence he does not have incentive to aim for slot $j - 1$. For the players who occupy slots $\lambda + 1, \dots, 2\lambda$, due to the ties they can only aim at slot λ for the price of δ . But since the current price they pay is δ and since the CTRs in all these slots are identical with the CTR of slot λ , there is no incentive for moving upwards. Now consider deviations *downwards*. Due to ties, the players at slots $j = \lambda - 1, \dots, 2\lambda$, could target the last slot. But then they would pay the same price δ , while $\theta_{j+1} \leq \theta_j$. For players occupying the slots $j = 1, \dots, \lambda - 2$ we need to show that $\theta_j(v_{(j)} - b_{(j+1)}) \geq \theta_{j+1}(v_{(j)} - b_{(j+2)})$, which is equivalent to $c \cdot v_{(j)} - (c-1)v_{(j+2)} - b_{(j+3)} \geq v_{(j)} - b_{(j+2)}$. This in turn becomes $(c-1)v_{(j)} \geq (c-1)v_{(j+2)} + b_{(j+3)} - b_{(j+2)}$, which is true for all $j \leq \lambda - 2$. \square

In the proof we use an instance with $\gamma = 1$. In practice however, the CTRs have been observed to be well separated. Previous works have fitted actual data from sponsored search auctions by power law distributions. It has been reported in [9] that geometrically decreasing CTRs $\theta_j \propto \alpha^{1-j}$ fit real data extremely well for $\alpha = 1.428$. Other authors [17] have used a Zipf distribution, where $\theta_j = j^{-\alpha}$, for $\alpha \geq 1$. We will prove an upper bound on $\mathbf{\Lambda}$ for conservative bidders and game instances with $\gamma < 1$. For power law distributions of CTRs, including Zipf and geometrically decreasing distributions, it immediately yields upper bounds showing that $\mathbf{\Lambda}$ is significantly lower than for general instances.

We will show an upper bound simultaneously for continuous and discrete strategies. To this end we will use the notion of the *marginal bid* β_i of each player i . This will be the largest bid that player i may issue. For the case of continuous strategies it will be $\beta_i = v_i$, while for discrete strategies with bidding step ϵ we have $\beta_i = \lfloor v_i \rfloor_\epsilon$. In using the marginal bids, we assume such an ordering of the players that $\beta_1 \geq \dots \geq \beta_n$. First we prove the following lemma, which states essentially that, at a locally stable configuration, every bidder with index $i \leq k$ is assigned a slot (rearranging indices appropriately in case of ties).

Lemma 3 *Let M be the set of players i with marginal bid equal to the k -th highest marginal bid $\hat{\beta}$ over all players. Under any conservative locally stable configuration \mathbf{b} the slots are allocated to players i with $\beta_i > \hat{\beta}$ and to a subset of players from M .*

Proof. Define $M = \{i | \beta_i = \hat{\beta}\}$ where $\hat{\beta}$ is the k -th highest marginal bid. Define $U = \{i | \beta_i > \hat{\beta}\}$ and $L = \{i | \beta_i < \hat{\beta}\}$. Note that $|U| \leq k - 1$. Assume there is a locally stable configuration \mathbf{b} under which a player $i \in U$ is not allocated a slot. Then there is certainly a player $j \in M \cup L$ that wins

a slot. This means that $b_{(k)} \leq b_j \leq \beta_j \leq v_j$. Clearly $u_i(\mathbf{b}) = 0$. If player i considers targeting the last slot (similar argument holds if due to ties he targets a higher slot) his payoff would be:

$$\theta_k(v_i - b_{(k)}) \geq \theta_k(\beta_i - \beta_j) > 0,$$

which contradicts local stability of \mathbf{b} . Hence assume now that all bidders of U are winners and it remains to show that the remaining winners are from M . Suppose there is a bidder $j \in L$ who wins a slot under \mathbf{b} . Then there is a bidder $i \in M$ who does not win a slot. By definition $\beta_i > \beta_j$. Again it holds that $b_{(k)} \leq b_j \leq \beta_j \leq v_j$. By repeating the argument above we obtain a contradiction to local stability of \mathbf{b} . \square

Our upper bound proof hereafter is built upon the concept of a *serial partition* of the k slots, with reference to a locally stable configuration \mathbf{b} , defined as follows:

Definition 3 A **serial partition** $\mathcal{P}(\mathbf{b})$ of slots with reference to a locally stable bid vector \mathbf{b} is a collection of $\ell \geq 1$ slot indices $\{j_r | r = 1, \dots, \ell\}$ such that:

$$\begin{aligned} j_r &= \arg \max_{j=1+j_{r-1}, \dots, k} \beta_{(j)}, \quad \text{for } r = 1, \dots, \ell \\ \{1, \dots, k\} &= \bigcup_{r=1}^{\ell} \{1 + j_{r-1}, \dots, j_r\}, \quad \text{where } j_0 = 0, j_\ell = k \end{aligned}$$

Hence a serial partition is simply a partition of the slots into groups such that each group contains consecutive slots and the last slot in each group is allocated to the person with the highest marginal bid (within this group and the remaining groups up to slot k). For a serial partition of slots according to a locally stable configuration \mathbf{b} we can prove the following lower bound on the social welfare:

Lemma 4 Let $\{j_r | r = 1, \dots, \ell\}$ denote a serial partition of slots with reference to a locally stable configuration \mathbf{b} . The social welfare $SW(\mathbf{b})$ under both continuous strategies and discrete strategies with $\epsilon \leq \gamma \Delta v$ is:

$$SW(\mathbf{b}) \geq (1 - \gamma) \sum_{r=1}^{\ell} \beta_{(j_r)} \sum_{j=1+j_{r-1}}^{j_r} \theta_j$$

Proof. By definition, $j_r \leq k$. Let $i_r = (j_r)$, the player occupying slot j_r and let j denote the slot that player i_r would target when attempting to deviate upwards with a local move. If there are at least two players, say i, i' , who receive the consecutive slots right above j_r , with $b_i = b_{i'}$, let j denote the smallest slot index that such a player receives. Otherwise, let $j = j_r - 1$. Now, given the local stability of \mathbf{b} there are two cases. Either player i_r does not have incentive to aim for slot j or *he cannot*, due to discretization of his strategy space. In the first case we have $\theta_{j_r}(v_{i_r} - b_{(j_r+1)}) \geq \theta_j(v_{i_r} - b_{(j)})$, which yields:

$$b_{(j)} \geq [1 - (\theta_{j_r}/\theta_j)]v_{i_r} + (\theta_{j_r}/\theta_j)b_{(j_r+1)} \geq [1 - (\theta_{j_r}/\theta_{j_r-1})]v_{i_r} \geq (1 - \gamma)v_{i_r} \geq (1 - \gamma)\beta_{i_r} \quad (4)$$

In the second case, which may happen under discretized strategies, it may occur that $\beta_{i_r} < b_{(j)} + \epsilon$ (which means $\beta_{i_r} \leq b_{(j)}$), even though $v_{i_r} > b_{(j)}$. Then, it must be $b_{(j)} + \epsilon > v_{i_r}$, hence $b_{(j)} > v_{i_r} - \epsilon$. Because $\epsilon \leq \gamma \cdot \Delta v \leq \gamma v_{i_r}$, we obtain again $b_{(j)} > (1 - \gamma)v_{i_r} \geq (1 - \gamma)\beta_{i_r}$. Notice that, by the way we defined j , if there is a slot j' , with $j < j' < j_r$, it should be $b_{(j')} = b_{(j)}$. Because \mathbf{b} is *conservative*, we obtain that for any slot index $l = 1, \dots, j_r - 1$, we have $v_{(l)} \geq b_{(l)} \geq b_{(j)} \geq (1 - \gamma)\beta_{i_r}$. Using (4) and summing the social welfare over the groups $r = 1, \dots, \ell$, we get:

$$\begin{aligned}
SW(\mathbf{b}) &= \sum_{r=1}^{\ell} \sum_{j=1+j_{r-1}}^{j_r} \theta_j v_{(j)} \\
&\geq (1-\gamma) \sum_{r=1}^{\ell} \sum_{j=1+j_{r-1}}^{j_r} \theta_j \beta_{(j_r)} \\
&= (1-\gamma) \sum_{r=1}^{\ell} \beta_{(j_r)} \sum_{j=1+j_{r-1}}^{j_r} \theta_j
\end{aligned}$$

□

Remark 1 Note that in the above lemma we only used upwards local stability.

We now provide the following upper bound on the optimal social welfare with respect to a serial partition of the slots.

Lemma 5 Consider a serial partition of the slots with reference to a locally stable bid vector \mathbf{b} . For discrete strategies and for any $\epsilon > 0$, the social optimum is upper bounded as follows:

$$SW(\mathbf{b}^*) \leq \sum_{r=1}^{\ell} \beta_{(j_r)} \sum_{j=1+j_{r-1}}^{j_r} \theta_j + \epsilon \sum_{j=1}^k \theta_j$$

For continuous strategies, the same bound applies after setting $\epsilon = 0$.

Proof. First notice that $SW(\mathbf{b}^*) \leq \sum_{j=1}^k \theta_j (\beta_j + \epsilon)$, which follows by $\beta_i = \lfloor v_i \rfloor_{\epsilon}$. Thus at most $\epsilon \sum_j \theta_j$ may be lost in expressing the optimum social welfare in terms of the k highest marginal bids. It remains to upper bound the term $\sum_{j=1}^k \theta_j \beta_j$. Now consider a serial partition $\mathcal{P}(\mathbf{b})$ of slots with reference to a locally stable bidding vector \mathbf{b} . We claim that for $r = 1, \dots, \ell$, $\beta_{(j_r)} \geq \beta_j$ for $j = 1 + j_{r-1}, \dots, j_r$. This is clearly true for $r = 1$, since by definition $\beta_{(j_1)} = \max_{j=1, \dots, k} \beta_{(j)} = \beta_1$ and, by lemma 3, only players with the k highest marginal bids receive a slot. For $r > 1$ suppose that $\beta_{(j_r)} < \beta_j$, for some $j \in \{1 + j_{r-1}, \dots, j_r\}$, in particular it must then be that $\beta_{(j_r)} < \beta_{1+j_{r-1}}$. By the definition of j_r and the serial partition with respect to \mathbf{b} , it cannot be the case that player $1 + j_{r-1}$ is assigned under \mathbf{b} a slot among $1 + j_{r-1}, \dots, k$. Hence by lemma 3, he must have been assigned a slot j in the range $1, \dots, j_{r-1}$. This means that –again by lemma 3– a player with marginal bid $\beta_i \geq \beta_{1+j_{r-1}} > \beta_{(j_r)}$ has to be assigned a slot within $\{1 + j_{r-1}, \dots, k\}$. But by the definition of j_r , this is impossible, hence we obtain the desired lower bound that $\beta_{(j_r)} \geq \beta_j$ for $j = 1 + j_{r-1}, \dots, j_r$. In total we have:

$$\sum_{j=1}^k \theta_j \beta_j = \sum_{r=1}^{\ell} \sum_{j=1+j_{r-1}}^{j_r} \theta_j \beta_j \leq \sum_{r=1}^{\ell} \beta_{(j_r)} \sum_{j=1+j_{r-1}}^{j_r} \theta_j$$

□

Theorem 3 The GSP Auction game with conservative bidders has Local Stability Ratio at most $(1-\gamma)^{-1}$ in continuous strategies and at most $2(1-\gamma)^{-1}$ in discrete strategies, assuming $\gamma < 1$ and taking any bidding step $\epsilon \leq \epsilon^* = \min\{\gamma, 1-\gamma\} \times \Delta v$. Assuming rational valuations, there is appropriately small $\epsilon \leq \epsilon^*$, so that the Local Stability Ratio in discrete strategies is at most $(1-\gamma)^{-1} + \gamma^{-1}$.

Proof. The case of continuous strategies follows immediately by lemmas 4 and 5, in the limit as $\epsilon \rightarrow 0$; then $\beta_i \rightarrow v_i$. For the case of discrete strategies we obtain:

$$\mathbf{\Lambda} \leq \frac{1}{1-\gamma} + \left(\epsilon \sum_j \theta_j \right) \times \left[(1-\gamma) \sum_{r=1}^l \beta_{(j_r)} \sum_{j=1+j_{r-1}}^{j_r} \theta_j \right]^{-1}$$

Because $\beta_i \geq \epsilon$, the latter yields $\mathbf{\Lambda} \leq 2(1-\gamma)^{-1}$. Using rationality of valuations, we can pick bidding step ϵ so that $v_n = \beta_n = r \cdot \epsilon$, for integer $r \geq 1$, where v_n is the smallest valuation. Then for every bidder i we have that $v_i \geq v_n = \beta_n \geq \Delta v$, hence $\beta_i \geq \Delta v$. By using this and substituting $\epsilon \leq \gamma \Delta v$ in the lower bound above, we obtain $\mathbf{\Lambda} \leq (1-\gamma)^{-1} + \gamma^{-1}$. \square

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

Proof. For geometrically decreasing CTRs with decay factor $\alpha > 1$, we have $\gamma_j = \alpha^{-1}$ for $j = 2, \dots, k$. For Zipf CTRs we have $\theta_j = j^{-\alpha}$, for some $\alpha \geq 1$. Hence $\gamma_j = \left(\frac{j-1}{j}\right)^\alpha$ and $\gamma \leq \left(\frac{k-1}{k}\right)^\alpha \leq \frac{k-1}{k}$. \square


This result along with the observations in [9] implies a constant upper bound on the performance of locally stable configurations under a model of geometrically decreasing CTRs. Furthermore, for this case an upper bound of $\alpha/(\alpha-1)$ has been given for the performance of Nash equilibria by Lahaie [13]. Surprisingly, we have been able to provide the same upper bound on the social inefficiency of the relaxed notion of locally stable configurations. Finally, our analysis for theorem 3 yields an upper bound on the efficiency of stable configurations of AB and CB as well as local iterative AB and CB, since such configurations are locally stable.

Corollary 2 *For $\gamma < 1$, the social inefficiency of stable configurations with respect to 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.*

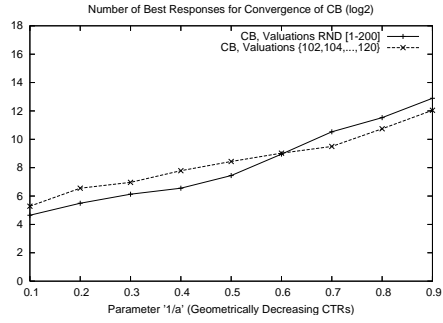
6 Experimental Investigation

We performed experiments to obtain indications regarding two issues; the tightness of the bounds for the inefficiency $\mathbf{\Lambda}$ of locally stable configurations (in continuous strategies) of the GSP Auction game, and for the convergence properties of the discussed procedures.


To examine tightness of bounds given in Corollary 1, we constructed numerically $n \times n$ game instances for which a bid configuration exists, that renders locally stable a *reverse* assignment of players to slots, i.e. the j -th slot is assigned to player $n - j + 1$. We solved numerically two non-linear programs, that express bounds on $\mathbf{\Lambda}$ (described in **Appendix B**), one for geometrically decreasing CTRs (for up to 20 players) and one for Zipf distributed ones (for up to 17 players). In each case we solved for (i) unrestricted $\alpha \geq 1$ and (ii) $\alpha = 2$ (geom. decreasing) and $\alpha = 1$ (Zipf) respectively. Fig. 1a depicts the evolution of the inefficiency $\mathbf{\Lambda}$ of these reversed locally stable configurations with number of players. It suggests a linear evolution in the case of Zipf distributed CTRs with the number of slots, but a modest 1.5-bounded evolution for geometrically decreasing ones with $\alpha = 2$ and 2-bounded for unrestricted α .



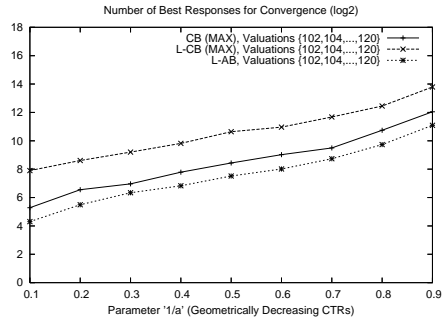
(a) Lower Bounds for of Reversed Assignments.



(b) Convergence Time of CB.



(c) Randomly Chosen Valuations.



(d) Densely Chosen Valuations.

Figure 1: Lower Bounds on Λ of Locally Stable Configurations (continuous) and Convergence Times for CB, L-CB, L-AB

Regarding convergence of AB, CB, we have found examples showing that iterative AB does not converge even in the iterative best response described here for $\epsilon \leq \epsilon^*$ and even for geometrically decreasing CTRs that are well separated. We provide a numerical example in **Appendix B**. For the procedures CB, L-CB and L-AB we were to our surprise unable to identify non-convergent instances during extensive experimentations. Most importantly, we found that, for sufficiently small bidding step ϵ below the bound ϵ^* utilized throughout the paper, the social welfare and the revenue in the reached stable states matched always the ones obtained by the equilibrium defined in Theorem 1. We believe that it is interesting to resolve convergence, particularly for CB.

Figures 1b, 1c and 1d show (in a \log_2 scale) how many best responses it takes for the procedures to converge on two kinds of 10×8 instances. We made 9 instances per kind with geometrically decreasing CTRs, parameterized by $\alpha^{-1} \in \{0.1, 0.2, \dots, 0.9\}$. One kind of instances contained valuations chosen uniformly at random in $[1, 200]$ while the other contained the exact dense set of valuations $\{102, 104, 106, \dots, 120\}$. The depicted data points are the maximum number of best responses per instance, encountered over 1000 executions of iterative best responses, initialized from a random initial configuration. We notice that in all cases there appears to be an exponential growth of the number of steps as α^{-1} approaches 1.

References

- [1] K. Asdemir. Bidding patterns in search engine auctions. In *Proceedings of the 2nd Workshop on Sponsored Search Auctions*, 2006.

- [2] C. Borgs, J. T. Chayes, N. Immorlica, K. Jain, O. Etesami, and M. Mahdian. Dynamics of bid optimization in online advertisement auctions. In *Proceedings of the International Conference on World Wide Web (WWW)*, pages 531–540, 2007.
- [3] F. Brandt, F. A. Fischer, P. Harrenstein, and Y. Shoham. Ranking games. *Artificial Intelligence*, 173(2):221–239, 2009.
- [4] F. Brandt and G. Weiß. Antisocial agents and vickrey auctions. In *Intelligent Agents, 8th International Workshop (ATAL)*, pages 335–347, 2001.
- [5] T. Bu, X. Deng, and Q. Qi. Forward looking nash equilibrium for keyword auction. *Information Processing Letters*, 105(2):41–46, 2008.
- [6] M. Cary, A. Das, B. Edelman, I. Giotis, K. Heimerl, A. R. Karlin, C. Mathieu, and M. Schwarz. Greedy bidding strategies for keyword auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 262–271, 2007.
- [7] B. Edelman and M. Ostrovsky. Strategic bidder behavior in sponsored search auctions. *Decision Support Systems*, 43(1):192–198, 2007.
- [8] B. Edelman, M. Ostrovsky, and M. Schwartz. Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review*, 97(1):242–259, 2007.
- [9] J. Feng, H. K. Bhargava, and D. M. Pennock. Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *INFORMS Journal on Computing*, 19(1):137–148, 2007.
- [10] A. Grillo, A. Lentini, M. Naldi, and G.F. Italiano. Penalized second price: A new pricing algorithm for advertising in search engines. In *Accepted for presentation at the 8th Conference on Communications Networks and Services Research (CNSR2010)*, Montreal, 11-14 May 2010.
- [11] T. Groves. Incentives in Teams. *Econometrica*, 41(4):617–631, 1973.
- [12] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In *STACS*, pages 404–413, 1999.
- [13] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 218–227, 2006.
- [14] R. Paes Leme and E. Tardos. Pure and bayes-nash price of anarchy for generalized second price auction. In *Foundations of Computer Science (FOCS)*, to appear, 2010.
- [15] L. Liang and Q. Qi. Cooperative or vindictive: Bidding strategies in sponsored search auction. In *Workshop on Internet and Network Economics (WINE)*, pages 167–178, 2007.
- [16] B. Lucier and R. P. Leme. Improved Social Welfare Bounds for GSP at Equilibrium. *CoRR*, abs/1011.3268, 2010.
- [17] M. Naldi, G. D’Acquisto, and G. F. Italiano. The value of location in keyword auctions. *Electronic Commerce Research and Applications (ECRA)*, 9(2):160–170, 2010.
- [18] H. R. Varian. Position Auctions. *International Journal of Industrial Organization*, 25:1163–1178, 2007.

- [19] Y. Vorobeychik. Simulation-based game theoretic analysis of keyword auctions with low dimensional bidding strategies. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2009.
- [20] Y. Vorobeychik and D. M. Reeves. Equilibrium analysis of dynamic bidding in sponsored search auctions. In *Workshop on Internet and Network Economics (WINE)*, pages 155–166, 2007.
- [21] Y. Zhou and R. M. Lukose. Vindictive bidding in keyword auctions. In *ICEC*, pages 141–146, 2007.

Appendix A

Proof of Proposition 1 for AB

Let $i_0 = (j_0)$ be a player occupying slot j_0 under \mathbf{b} and assume that a slot $j > j_0$ would incur i_0 utility larger than $u_{i_0}(\mathbf{b})$, but there is no available strategy for player i_0 that grants him slot j . The deviation in this case would occur as follows: i_0 leaves slot j_0 , all players occupying slots $j_0 + 1, \dots, j$ would shift one slot index downwards to occupy slots $j_0, \dots, j - 1$, and player i_0 would have to beat the bid of the player $(j + 1)$ (under \mathbf{b}) so as to gain slot j . The new utility of i_0 would be $u_{i_0}(\mathbf{b}') = \theta_j(v_{i_0} - b_{(j+1)})$. Thus an appropriate bid according to AB would be $b'_{i_0} = b_{(j+1)} + \epsilon$. Such a bid may not work due to one of the two following reasons:

(I) There is a slot j' such that $j' > j_0$ and $j' < j$, which is granted to i_0 by bid b'_{i_0} . Let b' be the bid of the player that occupies j' under \mathbf{b}_{-i_0} . If $b' \neq b_{(j+1)}$, then it must be $b' - b_{(j+1)} < \epsilon$ a contradiction. Otherwise $b' = b_{(j+1)}$. Then aiming for slot j is not a best response for i_0 , because $\theta_j(v_{i_0} - b_{(j+1)}) < \theta_{j'}(v_{i_0} - b_{(j+1)}) = \theta_{j'}(v_{i_0} - b')$, assuming all CTRs differ.

(II) $b'_{i_0} = b_{i_0}$, in which case the new bid grants i_0 the slot that he already occupies under \mathbf{b} , by definition of AB. Assuming $j \neq j_0$, there must exist at least one slot j' with $j' > j_0$ and $j' < j$ such that $b_{(j')} = b_{i_0} = b'_{i_0}$ for, otherwise, it must be $j = j_0$. Then notice that, under \mathbf{b} , player i_0 pays $b_{(j')}$. And for $\epsilon \leq \epsilon^* = \min\{\gamma, 1 - \gamma\} \times \Delta v$, we have:

$$\begin{aligned}
 u_{i_0}(\mathbf{b}') = \theta_j(v_{i_0} - b_{(j+1)}) &= \theta_j(v_{i_0} - b_{(j+1)} - \epsilon) + \theta_j\epsilon \\
 &= \theta_{j_0}(v_{i_0} - b_{(j')}) + (\theta_j - \theta_{j_0})(v_{i_0} - b_{(j')}) + \theta_j\epsilon \\
 &\leq \theta_{j_0}(v_{i_0} - b_{(j')}) + (\theta_j - \theta_{j_0})(v_{i_0} - b_{(j')}) + \theta_j(1 - \gamma)(v_{i_0} - v_{(j')}) \\
 &\leq \theta_{j_0}(v_{i_0} - b_{(j')}) + (\theta_j - \theta_{j_0})(v_{i_0} - b_{(j')}) + \theta_j(\gamma^{-1} - 1)(v_{i_0} - v_{(j')}) \\
 &\leq \theta_{j_0}(v_{i_0} - b_{(j')}) + (\theta_j - \theta_{j_0})(v_{i_0} - b_{(j')}) + (\theta_{j_0} - \theta_j)(v_{i_0} - b_{(j')}) \\
 &= \theta_{j_0}(v_{i_0} - b_{(j')}) = u_{i_0}(\mathbf{b})
 \end{aligned}$$

which states that aiming for j cannot be more profitable for i_0 . The previous to last inequality is derived by the assumption that \mathbf{b} was a conservative configuration, hence $v_{(j')} \geq b_{(j')}$. \square

Appendix B

θ_j	i	v_i	b_i	u_i
2^0	0	60	58.5	2.3
2^{-1}	1	59	57.7	0.7
2^{-2}	2	58	57.6	0.425
2^{-3}	3	57	56.3	0.1375
2^{-4}	4	56	55.9	0.09375
2^{-5}	5	55	54.5	0.034375

(a)

θ_j	i	v_i	b_i	u_i
2^0	0	60	58.5	2.3
2^{-1}	1	59	57.7	0.7
2^{-2}	2	58	57.6	0.525
2^{-3}	4	56	55.9	0.175
2^{-4}	3	57	54.6	0.15625
2^{-5}	5	55	54.5	0.034375

(b)

θ_j	i	v_i	b_i	u_i
2^0	0	60	58.5	2.4
2^{-1}	2	58	57.6	1
2^{-2}	1	59	56	0.775
2^{-3}	4	56	55.9	0.175
2^{-4}	3	57	54.6	0.15625
2^{-5}	5	55	54.5	0.034375

(c)

θ_j	i	v_i	b_i	u_i
2^0	0	60	58.5	2.4
2^{-1}	2	58	57.6	0.95
2^{-2}	3	57	56.1	0.25
2^{-3}	1	59	56	0.3875
2^{-4}	4	56	55.9	0.09375
2^{-5}	5	55	54.5	0.034375

(d)

θ_j	i	v_i	b_i	u_i
2^0	0	60	58.5	2.4
2^{-1}	2	58	57.6	0.9
2^{-2}	1	59	56.2	0.725
2^{-3}	3	57	56.1	0.1375
2^{-4}	4	56	55.9	0.09375
2^{-5}	5	55	54.5	0.034375

(e)

θ_j	i	v_i	b_i	u_i
2^0	0	60	58.5	2.4
2^{-1}	2	58	57.6	0.85
2^{-2}	3	57	56.3	0.2
2^{-3}	1	59	56.2	0.3875
2^{-4}	4	56	55.9	0.09375
2^{-5}	5	55	54.5	0.034375

(f)

Figure 2: A cycle of the AB iterative best response, with geometrically decreasing CTRs. Bold fonts mark the agent that deviates; its new position is enclosed by horizontal lines in each state. The bidding step is 0.1, while $\gamma = 0.5$ and $\Delta v = 1$.

Non-Linear Programs for Finding Lower Bounds

$$\text{Zipf} \left\{ \begin{array}{l} \Lambda \cdot \left(\sum_{j=1}^n j^{-\alpha} v_{n-j+1} \right) = \sum_{j=1}^n j^{-\alpha} v_j \\ b_j \leq \left(1 - \left(\frac{n-j}{n-j+1} \right)^\alpha \right) v_{j+1} + \left(\frac{n-j}{n-j+1} \right)^\alpha b_{j-1} \\ b_j \geq \left(1 - \left(\frac{n+1-j}{n+2-j} \right)^\alpha \right) v_{j-1} + \left(\frac{n+1-j}{n+2-j} \right)^\alpha b_{j-2} \\ b_j \leq v_j, b_j + t \leq b_{j+1}, v_j \leq v_{j-1} \end{array} \right. \quad \text{Geom.} \left\{ \begin{array}{l} \Lambda \cdot \left(\sum_{j=1}^n \alpha^{1-j} v_{n-j+1} \right) = \sum_{j=1}^n \alpha^{1-j} v_j \\ b_j \leq (1 - \alpha^{-1}) v_{j+1} + \alpha^{-1} b_{j-1} \\ b_j \geq (1 - \alpha^{-1}) v_{j-1} + \alpha^{-1} b_{j-2} \\ b_j \leq v_j, b_j + t \leq b_{j+1}, v_j \leq v_{j-1} \end{array} \right.$$

We solved these NLPs by maximizing Λ , for Zipf and Geometrically decreasing distributions of CTRs respectively. An important detail is the t parameter which must be set greater equal to a positive value, to separate the bids.