Equilibrium Analysis of Dynamic Bidding in Sponsored Search Auctions*

Yevgeniy Vorobeychik
Computer Science & Engineering
University of Michigan
yvorobey@umich.edu

Daniel M. Reeves
Microeconomics & Social Systems
Yahoo! Research
dreeves@yahoo-inc.com

Abstract. We analyze symmetric pure strategy equilibria in dynamic sponsored search auction games using simulations, restricting the strategies to several in a class of greedy bidding strategies introduced by Cary et al. We show that a particular convergent strategy, "balanced bidding", also exhibits high stability to deviations in the dynamic setting. On the other hand, a cooperative strategy which yields high payoffs to all players is not sustainable in equilibrium play. Additionally, we analyze a repeated game in which each stage is a static complete-information sponsored search game. In this setting, we demonstrate a collusion strategy which yields high payoffs to all players and empirically show it to be sustainable over a range of settings. Finally, we show how a collusive strategy profile can arise even in the case of incomplete information.

1 Motivation

Sponsored search—the placement of advertisements along with search engine results—is currently a multi-billion dollar industry, with Google and Yahoo! the key players [Lahaie, 2006]. In the academic literature, much progress has been made in modeling sponsored search auctions as one-shot games of complete information, in which the players' values per click and click-through-rates are common knowledge. A typical justification for such an approach is the abundance of information in the system, since the advertisers have ample opportunity to explore, submitting and resubmitting bids at will. As the complexity of modeling the full dynamic game between advertisers that is actually taking place is quite intractable, static models provide a good first approximation. However, it ultimately pays to understand how relevant the dynamics really are to strategic choices of players.

* This is a revised and extended version of a paper appearing in the Third International Workshop on Internet and Network Economics (WINE'07). The research was conducted while the first author was at Yahoo! Research.

One question which has been addressed in the dynamic setting is whether it is reasonable to expect simple dynamic strategies to converge to Nash equilibria. Cary et al. [2007] explored several *greedy bidding strategies*, that is, strategies under which players submit bids with the goal of obtaining the most profitable slot given that other players bids are fixed. One of these strategies, *balanced bidding*, provably converges to a minimum revenue symmetric Nash equilibrium of the static game of complete information. This happens to be analytically tractable and has therefore received special attention in the literature [Varian, to appear, Lahaie and Pennock, 2007, Edelman et al., 2007]. Similar questions, particularly in the context of pools of vindictive agents, have been studied by Liang and Qi [2007].

Convergence of dynamic bidding strategies is only one of many relevant questions that arise if we try to account for the dynamic nature of the sponsored search game. Another is whether we can identify Nash equilibrium strategies in the dynamic game. This problem in general is quite hard as there are many possible actions and ways to account for the changing information structure. One approach, taken by Feng and Zhang [2007], is to model the dynamic process using a Markovian framework. Our own approach focuses on the set of greedy bidding strategies from Cary et al. [2007]. In motivating greedy bidding strategies, Cary et al. have argued that advertisers are unlikely to engage in highly fine-grained strategic reasoning and will rather prefer to follow relatively straightforward strategies. This motivation, however, only restricts attention to a set of plausible candidates. To identify which are likely to be selected by advertisers, we need to assess their relative stability to profitable deviations. For example, while we would perhaps like advertisers to follow a convergent strategy like balanced bidding, it is unclear whether players will find it more profitable to follow a non-convergent strategy.

Our goal is to provide some initial information about stability of a small set of greedy bidding strategies under incomplete information. Specifically, we use simulations to estimate the gain any advertiser can accrue by deviating from pure strategy symmetric equilibria in greedy bidding strategies. The results are promising: the convergent balanced bidding strategy is typically the most stable of the strategies we study.

To complement the analysis above, we examine the incentives when joint valuations are common knowledge, but the game is repeated indefinitely. Folk theorems [Mas-Colell et al., 1995] suggest that players may be able to increase individual profits (and decrease search engine revenue) by colluding. We demonstrate one such collusion strategy and show it to be effective over a range of sponsored search auction environments. Our analysis complements other approaches to studying collusion in auctions, both in the dynamic sponsored search context [Feng and Zhang, 2007] and in a general one-shot context [Krishna, 2002]. Finally, we extend the collusion result to the case of incomplete information.

2 Game Theoretic Preliminaries

We first review terminology, definitions, and core concepts from game theory that we employ throughout this paper. Our key solution concept is the Nash equilibrium and approximations thereof.

2.1 One-Shot Games of Incomplete Information

In much of this work we analyze the *strategic form*[1] of dynamic games, that is, *one-shot games of incomplete information* [Mas-Colell et al., 1995], denoted by $[I, \{R_i\}, \{T_i\}, F(\cdot), \{u_i(r,t)\}]$. where I refers to the set of players and $m = |I|$ is the number of players. R_i is the set of actions available to player $i \in I$, and $R_1 \times \cdots \times R_m$ is the joint action space. T_i is the set of types (private information) of player i, with $T = T_1 \times \cdots \times T_m$ representing the joint type space. Since we presume that a player knows its type prior to taking an action, but does not know types of others, we allow it to condition its action on own type. Thus, we define a strategy of a player i to be a function $s_i : T_i \rightarrow \mathbb{R}$, and use $s(t)$ to denote the vector $(s_1(t_1), \ldots, s_m(t_m))$. $F(\cdot)$ is the distribution over the joint type space.

We use s_{-i} to denote the joint strategy of all players other than player i. Similarly, t_{-i} designates the joint type of all players other than i. We define the payoff (utility) function of each player i by $u_i : R \times T \rightarrow \mathbb{R}$, where $u_i(r_i, r_{-i}, t_i, t_{-i})$ indicates the payoff to player i with type t_i for playing action $r_i \in R_i$ when the remaining players with joint types t_{-i} play r_{-i}. Given a strategy profile $s \in S$, the expected payoff of player i is $\tilde{u}_i(s) = E_t[u_i(s(t), t)]$.

Given a known strategy profile of players other than i, we define the best response of player i to s_{-i} to be the strategy s_i^* that maximizes expected utility $\tilde{u}_i(s_i, s_{-i})$. If we know the best response of every player to a strategy profile s, we can evaluate the maximum amount that any player can gain by deviating from s. Such an amount, which we also call *regret*, we denote by $\epsilon(s) = \max_{i \in I}[\tilde{u}_i(s_i^*, s_{-i}) - \tilde{u}_i(s_i, s_{-i})]$. When we use the term "stability" it is in the sense of low *regret*. Faced with a one-shot game of incomplete information, an agent would ideally play a strategy that is a best response to strategies of others. A joint strategy s where all agents play best responses to each other constitutes a *Nash equilibrium* ($\epsilon(s) = 0$); when applied to games of incomplete information, it is called a *Bayes-Nash equilibrium*.

2.2 Complete Information Infinitely Repeated Games

The second model we use is an *infinitely repeated game* [Mas-Colell et al., 1995]. The model divides time into an infinite number of discrete stages and presumes that at each stage players interact strategically in a one-shot fashion (that is,

[1] Although strategies are dynamic in that players choose their *actions* as a function of history, our model of the meta-level strategic interaction is one-shot in that players choose the dynamic strategies (which dictate actions in specific states) once and follow these throughout.

no one agent can observe actions of others until the next stage). Naturally, all players care not just about the payoffs they receive in one stage, but all the payoffs in past and subsequent stages of the dynamic interaction. We assume that their total utility from playing the repeated game is a discounted sum of stage utilities. Formally, it can be described by the tuple $[I, \{R_i\}, u_i(r), \gamma_i]$, where I, R_i and $u_i(r)$ are as before, and γ_i is the amount by which each player discounts utility at each stage. That is, if we let $\bar{r} = \{r_1, r_2, \dots, r_i, \dots\}, r_j \in R$ be a sequence of choices by players indexed by the chronological sequence of stages, $U_i(\bar{r}) = \sum_{t=1}^{\infty} \gamma_i^{t-1} u_i(r_t)$.

Define a stage-k *subgame* of a repeated game as a restricted repeated game which begins at stage k rather than at stage 1. The solution concept that we will use for infinitely repeated games is the *subgame perfect Nash equilibrium* [Mas-Colell et al., 1995], which obtains when the players have no incentive to deviate from their sequence of strategic choices in any stage of the interaction.

3 Modeling Sponsored Search Auctions

A traditional model of sponsored search auctions specifies a ranking rule, which ranks advertisers based on their bid and some information about their relevance to the user query, click-through-rates for each player and slot, and players' valuations or distributions of valuations per click. Let a player i's click-through-rate in slot s be denoted by c_s^i and its value per click by v_i. Like many models in the literature (e.g., [Lahaie, 2006, Lahaie and Pennock, 2007]) we assume that click-through-rate can be factored into $e_i c_s$ for every player i and every slot s. If player i pays p_i^s in slot s, then its utility is $u_i = e_i c_s (v_i - p_i^s)$. The parameter e_i is often referred to as *relevance* of the advertiser i, and c_s is the slot-specific click-through-rate. We assume that the search engine has K slots with slot-specific click-through-rates $c_1 > c_2 > \dots > c_K$.

Lahaie and Pennock [2007] discuss a family of ranking strategies which rank bidders in order of the product of their bids b_i and some weight function w_i. They study in some depth a particular weight function $w(e_i) = e_i^q$, where q is a real number. In the analysis below, we consider two settings of q: 0 and 1. The former corresponds to rank-by-bid, b_i, whereas the latter is typically called rank-by-revenue, $e_i b_i$.

When players are ranked by their bids, two alternative pricing schemes have been studied: first-price (set price equal to player's bid) and generalized second-price (set price equal to next highest bid). For more than one slot, neither is incentive compatible. However, stability issues have induced the major search engines to use generalized second-price auctions. These have been generalized further to ranking by weighted bid schemes by using the price rule $p_i^s = \frac{w_{s+1} b_{s+1}}{w_i}$. The interpretation is that the bidder i pays the amount of the lowest bid sufficient to win slot s.

4 Dynamic Bidding Strategies

In much of this work we restrict the strategy space of players to four dynamic strategies. While this is a dramatic restriction, it allows us to gain some insight into the stability properties of the dynamic game and to identify particularly interesting candidates for further analysis in the future. Additionally, it has been argued as unlikely that players will engage in full-fledged strategic reasoning and will rather follow relatively straightforward dynamic strategies [Cary et al., 2007] such as the ones we consider. We now define a simple class of dynamic bidding strategies.

Definition 1 (Greedy Bidding Strategies). *A* greedy bidding strategy [Cary et al., 2007] *for a player i is to choose a bid for the next round of a repeated keyword auction that obtains a slot which maximizes its utility u_i assuming the bids of all other players remain fixed.*

If the player bids so as to win slot s which it is selecting according to a greedy bidding strategy, any bid in the interval (p_i^s, p_i^{s-1}) will win that slot at the same price. The particular rule which chooses a bid in this interval defines a member of the class of greedy bidding strategies. We analyze strategic behavior of agents who can select from four greedy bidding strategies specified below. For all of these, let s^* designate the slot which myopically maximizes player i's utility as long as other players' bids are fixed.

Definition 2 (Balanced Bidding). *The* Balanced Bidding [Cary et al., 2007] *strategy* BB *chooses the bid b which solves $c_{s^*}(v_i - p_i^{s^*}) = c_{s^*-1}(v_i - b)$. If s^* is the top slot, choose $b = (v_i + p_i^1)/2$.*

The Balanced Bidding strategy bids what the next higher slot would have to be priced to make the player indifferent about switching to it.

Definition 3 (Competitor Busting). *The* Competitor Busting [Cary et al., 2007] *strategy* CB *selects the bid $b = \min\{v_i, p_i^{s^*-1} - \epsilon\}$.*

Thus the CB strategy tries to cause the player that receives the slot immediately above s^* to pay as much as possible.

Definition 4 (Altruistic Bidding). *The* Altruistic Bidding [Cary et al., 2007] *strategy* AB *chooses the bid $b = \min\{v_i, p_i^{s^*} + \epsilon\}$.*

This strategy ensures the highest payoff (lowest price) of the player receiving the slot immediately above s^*.

Definition 5 (Random Bidding). *The* Random *strategy* RAND *selects the bid b uniformly randomly in the interval (p_i^s, p_i^{s-1}).*

5 Empirical Bayesian Meta-Game Analysis

In this section we construct and analyze a Bayesian meta-game played between advertisers (alternatively, bidders or players) who may choose one of four greedy bidding strategies described above. Being a one-shot game of incomplete information, the bidders can condition their strategic choices only on their own valuations. We do not allow conditioning based on relevances, as these are assumed to be a priori unknown both to the search engine and to the bidders. The reason we refer to the model as a meta-game is that we abstract away the dynamic nature of the game by enforcing a one-shot choice of a dynamic strategy, that is, once the strategy is chosen, the player must follow it forever after. While this is a strong assumption given the restriction of the strategy space, it is without loss of generality when no such restriction is imposed, since an optimal dynamic strategy is optimal in any subgame along the played path.

In order to construct the meta-game, we need to define player payoffs for every joint realization of values and relevances, as well as the corresponding choice of dynamic strategies. As is common for dynamic interactions, we define the payoff in the meta-game as the discounted sum of stage payoffs. In each stage, exactly one bidder, selected uniformly randomly, is allowed to modify its bid according to its choice of dynamic bidding strategy.[2] The corresponding stage payoff is an expected payoff given the ranking and payments of players as a function of joint bids, as defined in Section 3. We model the entire dynamic process—once relevances, values, and strategies are determined—using a simulator, which outputs a sample payoff at the end of a run of 100 stages. The discount factor is set to 0.95.[3] With this discount factor, the total contribution from stage 101 to infinity is 0.006, and we thus presume that the history thereafter is negligible.

Expected payoff to a particular player for a fixed value per click, relevance, and strategy is estimated using a sample average of payoffs based on 1000 draws from the distribution of valuations and relevances of other players. The metric for quality with which a particular strategy profile s approximates a Bayes-Nash equilibrium is the estimate of $\epsilon(s)$, which is the sample average gain from playing a best response to s over 100 draws from the player's value and relevance distributions. For each of these 100 draws, the gain from playing a best response to s is computed as the difference between the highest expected payoff for any strategy in the restricted set and the expected payoff from s_i, estimated as described above.

Since the meta-game is constructed numerically for every choice of values, relevances, and strategies of all players, an in-depth analysis of all strategies in the game is hopeless. Instead, we focus much of our attention on four pure

[2] This condition ensures the convergence of *balanced bidding* dynamics.

[3] While this is a very conservative discount for each bidding stage, our offline experiments suggest that our results are not particularly sensitive to it (for example, results which use average payoff per round as long-term utility seem to be qualitatively similar).

symmetric strategy profiles, in which each player chooses the same dynamic strategy for any valuation. While this seems an enormous restriction, it turns out to be sufficient for our purposes, as these happen to contain near-equilibria.

5.1 Equal Relevances

In this section we focus on the setting in which all players' relevances are equal and assume that values per click are distributed normally with mean 500 and standard deviation 200.[4] Three sponsored search auction games are considered: in one, 5 advertisers compete for 2 slots; in the others, 20 and 50 advertisers respectively compete for 8 slots.

Figure 1 presents average $\epsilon(s)$ and payoffs for all four pure symmetric profiles in strategies which are constant functions of player values per click. The first observation we can make is that BB has a very low $\epsilon(s)$ in every case, suggesting that it has considerable strategic stability in the restricted strategy space. This result can also be claimed with high statistical confidence, as 99% confidence intervals are so small that they are not visible in the figure. In contrast, AB manifests very high $\epsilon(s)$ in the plot and we can be reasonably certain that it is not sustainable as an equilibrium. The picture that emerges is most appealing to the search engine: AB, which yields the greatest payoffs to players (and least to the auctioneer), is unlikely to be played, whereas BB yields the lowest player payoffs in the restricted strategy space.

5.2 Independently Distributed Values and Relevances

We now consider the setting in which relevances of players are not identical, but are rather identically distributed—and independently from values per click—according to a uniform distribution on the interval [0,1]. Since now the particulars of the bid ranking scheme come into play, we present results for the two schemes that have received the most attention: rank-by-bid ($q = 0$) and rank-by-revenue ($q = 1$).

Figures 2a and b present the results on stability of each symmetric pure strategy profile to deviations for $q = 0$ and $q = 1$ respectively. We can see that there are really no qualitative differences between the two settings, and indeed between the setting of independently distributed values and relevances and the previous one in which relevances were set to a constant for all players. A possible slight difference is that RAND and CB strategies appear to have better stability properties when $q = 0$. However, this could be misleading since the payoffs to players are also generally lower when $q = 0$. The most notable quality we previously observed, however, remains unchanged: BB is an equilibrium (or nearly so) in all games for both advertiser ranking schemes, and AB is highly unstable, whereas BB yields a considerably lower payoff to advertisers than AB in all settings.

[4] For this and subsequent settings we repeated the experiments with an arguably more realistic lognormal distribution and found the results to be qualitatively unchanged.

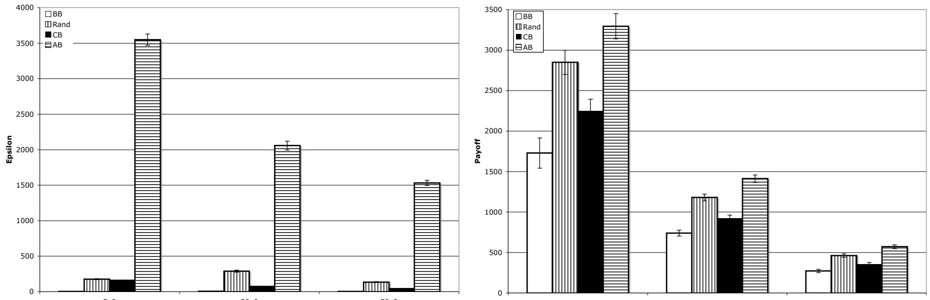

(a) (b)

Fig. 1. (a) Experimental $\epsilon(s)$ and (b) symmetric payoff for every pure symmetric profile in constant strategies with associated 99% confidence bounds.

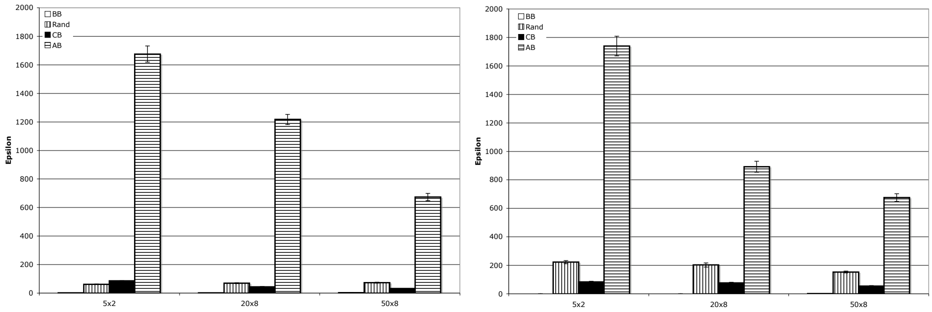

(a) (b)

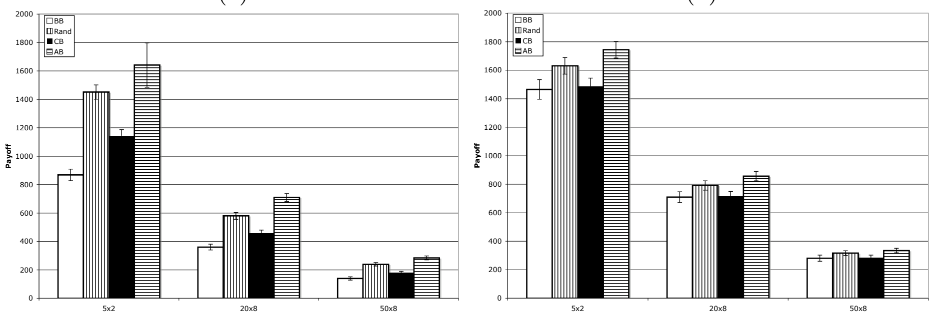

(c) (d)

Fig. 2. Experimental $\epsilon(s)$ (a) when $q = 0$ (b) when $q = 1$ for every pure symmetric profile; experimental payoff (c) when $q = 0$ (d) when $q = 1$ for every pure symmetric profile.

5.3 Correlated Values and Relevances

In the final set of experiments we draw values and relevances from a joint distribution with a correlation coefficient of 0.5. As may by now be expected, BB remains a near-equilibrium both when we set $q = 0$ and $q = 1$ (Figures 3a and b). However, when $q = 0$, RAND and CB are now also near-equilibria when the number of players and slots is relatively large—and, indeed, more so as the number of players grows from 20 to 50. As a designer, this fact may be somewhat disconcerting, as BB remains the strategy with the lowest payoffs to players (and, consequently, will likely yield the highest search engine payoffs) when $q = 0$; by comparison, payoffs to players when RAND is played are considerably higher than BB (Figure 3c). In all the cases, however, altruistic bidding remains highly unstable, to the bidders' great chagrin, as it is uniformly more advantageous in terms of payoffs (Figures 3c and d).

5.4 Truthfulness in Balanced Bidding

One question that arises is whether there is any incentive for a bidders to misrepresent their valuations. In the context of balanced bidding, this would mean that a bidder submits bids as if it had a valuation other than its own. Myopically this should not be the case since each bidder submits bids that greedily maximize their utility. In the long term, however, perhaps such deception would pay off. First, during the dynamics of balanced bidding a bidder could obtain a higher slot at a considerable premium. That untruthfulness in this sense is unprofitable at a fixed point seems a foregone conclusion, since balanced bidding is guaranteed to converge in our settings to a symmetric equilibrium. However, there may well be an asymmetric equilibrium with the resulting allocation and bids. Below, we show that this is impossible if the fixed point is a minimum symmetric equilibrium. Particularly, we now show that at a fixed point of balanced bidding the utility of an untruthful bidder is no higher than if it were truthful (and converged to the truthful fixed point). For convenience, we restrict the remainder of our analysis in this section to a setting in which all relevances are constant although the analysis would remain qualitatively unchanged if we had not.

Lemma 1. *Consider a situation in which all but one player is bidding according to balanced bidding dynamics, and one, the deviant, is considering whether to bid truthfully. The utility of the deviant in a truthful fixed point is no lower than in a fixed point reached when it is playing balanced bidding as if it had another value per click. Furthermore, bidding in order to get a higher slot yields a strictly lower utility in a fixed point for generic values per click.*

[The proof is in the appendix.] Considering Lemma 1 and the fact that balanced bidding is greedy and, consequently, a player cannot obtain immediate gain by deviating suggests that the only gain from an untruthful variant of balanced bidding is through transient payoffs—that is, as a side-effect of responses by agents who follow balanced bidding honestly. Such effects are difficult to

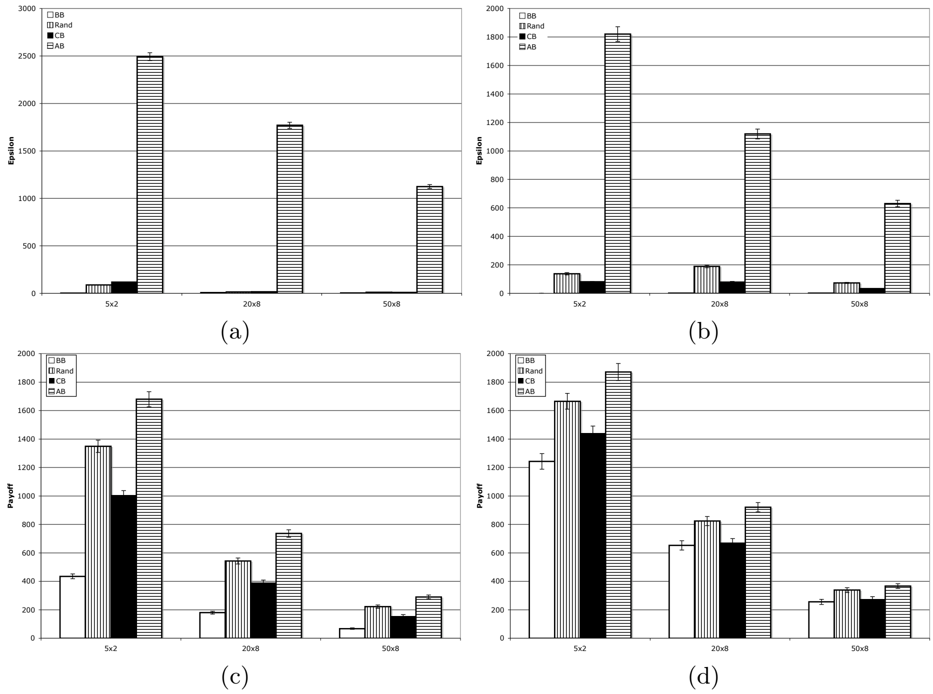

Fig. 3. Experimental $\epsilon(s)$ (a) when $q = 0$ (b) when $q = 1$ for every pure symmetric profile; experimental payoff (c) when $q = 0$ (d) when $q = 1$ for every pure symmetric profile.

study analytically; thus, we use simulations to gauge regret from truthful balanced bidding. In Figure 4 we present results for constant relevances, as well as settings when relevances and values are independent and correlated. The sponsored search settings considered are also as above. In the experiments, we allowed deviations to be factors of $k = \{0, 0.5, 1.5, 2\}$ of the player's valuation. Thus, for example, a player with $k = 0.5$ and value per click of 200 would play balanced bidding as if its value were 100. As the figure suggests, when there are few players, gains from bidding as if you are someone else appear relatively high: when values and relevances are correlated our results suggest that regret can be as high as 90 (roughly 22% of total payoff). However, regret drops off considerably as the number of players increases, down to about 10% of payoff with 50 players. Regret is lower when relevances are drawn independently of values or when they are constant.

Overall, we can observe that balanced bidding seems to be somewhat less stable when we consider the possibility that bidders may play untruthfully, that is, play as if they their value per click was different from what it actually is. Of course, we have already shown that if any agent did play untruthfully, the corresponding fixed point would not then be a fixed point, as the untruthful player would then want to deviate. Note, however, that our discount factor of 0.95 is actually extremely conservative: while reasonable as an annual rate, it is unlikely to be so low per bidding round. As such, the importance of the result in Lemma 1 is likely to be considerably greater than our empirical results suggest. In any case, whether the observed incentives to deviate are strong enough or not would remain on other factors, such as the cost of determining a deviation that carries substantial gain.

As a final piece of evidence for the efficacy of truthful balanced bidding, we compare its regret to that from playing several variants of untruthful bidding. Figure 5 displays the regret for several symmetric profiles with varying degrees of untruthfulness exhibited by the players in a game with 5 players and 2 slots. As we can see, the regret from truthful balanced bidding ($k = 1$) is far overshadowed by that from untruthful profiles.

6 Repeated Game

6.1 Common Knowledge of Values

It is common in the sponsored search auction literature to assume that the player valuations and click-through-rates are common knowledge, suggesting that the resulting equilibria are rest points of natural bidder adjustment dynamics. The justification offered alludes to the repeated nature of the agent interactions. Yet, the equilibrium concept used is a static one. If a game is infinitely repeated, the space of Nash equilibrium strategies expands considerably [Mas-Colell et al., 1995]. Thus, if we take the dynamic story seriously, it pays to seek subgame perfect equilibria in the repeated game, particularly if they may offer considerably better payoffs to players than the corresponding static Nash equilibria.

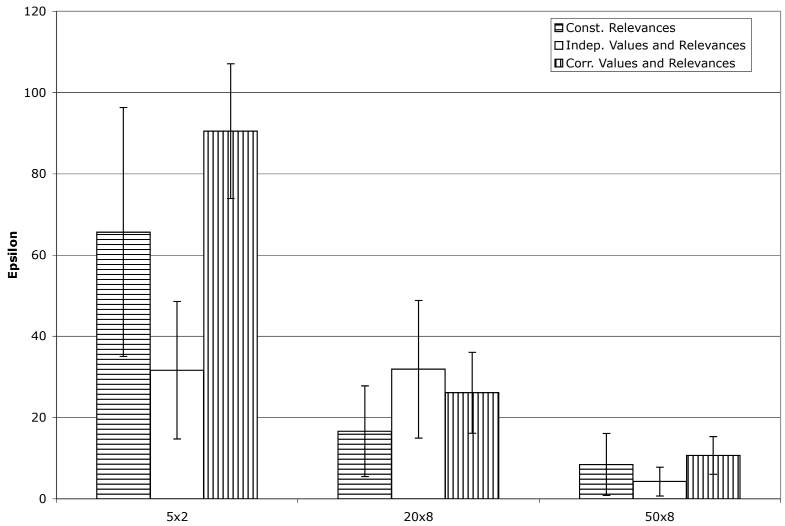

Fig. 4. Experimental $\hat{\epsilon}(s)$ and the associated 95% confidence intervals for balanced bidding based on allowed deviations to untruthful play.

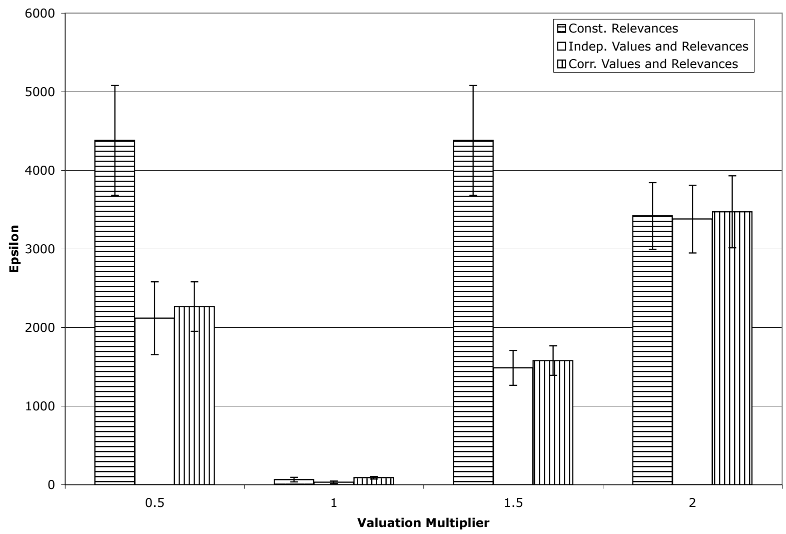

Fig. 5. Experimental $\hat{\epsilon}(s)$ and the associated 95% confidence intervals for balanced bidding profiles with $k = 0.5, 1, 1.5, 2$, where k is the multiplier of player's value; note that $k = 1$ is truthful balanced bidding. Deviations are allowed to $k = 0, 0.5, 1, 1.5, 2$. The setting is a game with 5 players and 2 slots.

As typical analysis of repeated interactions goes, our subgame perfect equilibrium consists of two parts: the main path, and the deviation-punishment path. The main path has players jointly follow an agreed-upon profitable strategy profile, whereas the deviation path punishes any deviant. The trick, of course, is that for the equilibrium to be subgame perfect, the punishment subgame must itself be in equilibrium, yet must be sufficiently bad to discourage deviation.

A natural candidate for punishment is the worst (in terms of player payoffs) Nash equilibrium in the static game. Clearly, such a path would be in equilibrium, and is likely to offer considerable discouragement to deviants. A desirable main path would have players pay as little as possible, but needs to nevertheless discourage bidders who do not receive slots from outbidding those who do. Furthermore, all "slotless" bidders should remain slotless in the deviation subgame, since it is then clear that no incentives to deviate exist among such bidders, and we need only consider bidders who occupy some slot.

For the remainder of this section, we assume that the valuations are generic and bidders are indexed by the number of the slot they obtain in a symmetric Nash equilibrium.[5] Define the dynamic strategy profile *COLLUSION* as follows:

- *Main path*: $\forall s > K$, $b_s = v_s$. For all others, $b_s = \frac{w_{K+1}}{w_s} v_{K+1} + (K - s + 1)\epsilon$, where ϵ is some very small (negligible) number. Note that this yields the same ordering of bidders who receive slots as any symmetric Nash equilibrium of the game.
- *Deviation path*: play the maximum revenue symmetric Nash equilibrium strategies in every stage game. This yields the maximum revenue to the auctioneer and the lowest utilities to the players of any Nash equilibrium in the stage game [Varian, to appear].

Whether the delineated strategy constitutes a subgame perfect Nash equilibrium depends on the magnitude of the discount factor, γ_i, of every player i. The relevant question is then how large does γ need to be to enable enforcement of *COLLUSION*. For example, $\gamma_i = 0$ will deter nothing, since there are no consequences (the game is effectively a one-stage game). Below, we give the general result to this effect.

Theorem 1. *The* COLLUSION *strategy profile is a subgame perfect Nash equilibrium if, for all players i,*

$$\gamma_i \geq \max_{s \leq K, t \leq s} \frac{(c_t - c_s)(v_s - \frac{w_{K+1} v_{K+1}}{w_s}) - (c_t \frac{w_t}{w_s}(K - t + 1) - c_s(K - s))\epsilon}{c_t(v_s - \frac{w_{K+1} v_{K+1}}{w_s}) - c_s v_s - c_t \frac{w_t}{w_s}(K - t + 1)\epsilon + V_{sum}}, \quad (1)$$

where $V_{sum} = \sum_{t=s+1}^{K} w_{t-1} v_{t-1}(c_{t-1} - c_t) + w_K v_K c_K$.

[The proof is in the appendix.] The lower bound on the discount factor in Equation 1 depends on the particular valuation vector, the relative merits of slots, and the total number of slots, and it is not immediately clear whether there

[5] Via a simple extension of the results by Varian [to appear] we can show that in a symmetric Nash equilibrium, bidders are ranked by $w_s b_s$.

actually are reasonable discount factors for which deviations can be discouraged. To get a sense of how sustainable such an equilibrium could be, we study the effect of these parameters on the lower bound of the discount factor. To do this, we let the relevances of all players be constant, fix the number of players at 20 and take 100 draws of their valuations from the normal distribution with mean 500 and standard deviation 200. We vary the number of slots between 2 and 15, recording the average, minimum, and maximum values of the lower bound. Furthermore, we normalize c_1 to 1 and let $\frac{c_s}{c_{s+1}} = \delta$ for all $s \leq K-1$. The results are displayed in Figure 6 for different values of δ.

First, focus on Figure 6c which shows the results for $\delta = 1.428$, an empirically observed click-through-rate ratio [Lahaie and Pennock, 2007]. As the figure suggests, when the number of slots is between 0 and 5, it seems likely that *COLLUSION* can obtain as a subgame perfect equilibrium, as the requirements on the discount factor are not too strong. When the number of slots grows, however, the incentives to deviate increase, and when the number of slots is above 10, such a collusive equilibrium no longer seems likely.

Figures 6a, b, and d display similar plots for other settings of δ. These suggest that as δ rises, incentives to deviate rise, since when there is a greater dropoff in slot quality for lower slots, players have more to gain by moving to a higher slot even for a one-shot payoff.

Above, our punishment path involved bidders playing a maximum symmetric equilibrium. Such a punishment path is quite strong, since, as we mentioned, it yields the lowest utilities to the players of any Nash equilibrium in the stage game. An alternative and considerably weaker punishment would be to play a *minimum* symmetric equilibrium. Define the dynamic strategy profile *COLLUSION2* as follows:

- *Main path*: $\forall s > K$, $b_s = v_s$. For all others, $b_s = v_{K+1} + (K-s+1)\epsilon$, where ϵ is some very small (negligible) number. Again, this yields the same ordering of bidders who receive slots as any symmetric Nash equilibrium of the game.
- *Deviation path*: play the *minimum* revenue symmetric Nash equilibrium strategies in every stage game.

One may wonder why we would ever stipulate a weaker punishment. The answer is that a maximum revenue symmetric equilibrium deviation path actually requires the common knowledge of values, which in the generalized price auction will be different from bids for all players who receive slots. As we will see below, there is considerable value in being able to relax this assumption. Particularly, it is convenient that balanced bidding converges to a minimum symmetric Nash equilibrium as its unique fixed point. Thus, if we assume that all bidders follow this strategy (and do so truthfully), we can simply punish by stipulating that bidders revert to their fixed point bids, making *COLLUSION2* a much more plausible strategy.

We can easily extend Theorem 1 to *COLLUSION2*:

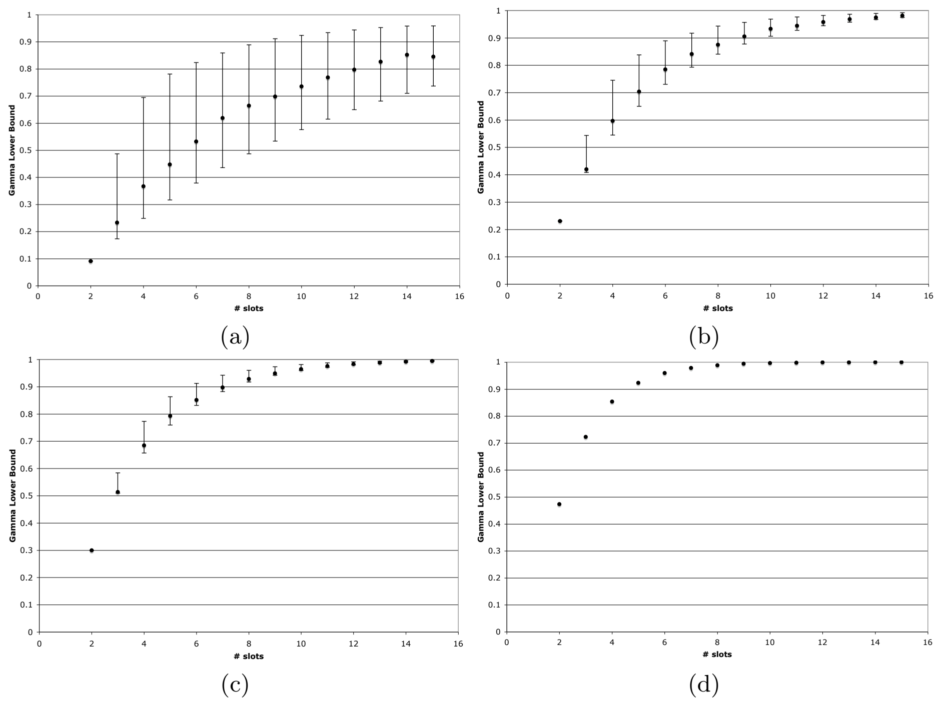

Fig. 6. Lower bounds on the discount factor as the number of available slots varies when (a) $\delta = 1.1$, (b) $\delta = 1.3$, (c) $\delta = 1.428$, and (d) $\delta = 1.9$.

Theorem 2. *The* COLLUSION2 *strategy profile is a subgame perfect Nash equilibrium if, for all players* i,

$$\gamma_i \geq \max_{s \leq K, t \leq s} \frac{(c_t - c_s)(v_s - \frac{w_{K+1}v_{K+1}}{w_s}) - (c_t \frac{w_t}{w_s}(K-t+1) - c_s(K-s))\epsilon}{c_t(v_s - \frac{w_{K+1}v_{K+1}}{w_s}) - c_s v_s - c_t \frac{w_t}{w_s}(K-t+1)\epsilon + V_{sum}}, \quad (2)$$

where $V_{sum} = \sum_{t=s+1}^{K} w_t v_t (c_{t-1} - c_t) + w_{K+1} v_{K+1} c_K$.

Naturally, we would like to know how much weaker *COLLUSION2* is than *COLLUSION*. The answer can be observed in Figure 7: it does appear substantially weaker. While collusion still seems likely when the number of slots is below about 8, even with few slots there are instances when collusion is not feasible, as the higher range of bounds on γ seem very near 1.

A final question we would like to raise here is whether the choice of ranking function affects feasibility of collusion and, if so, how much. We study this in the cases when values and relevances are distributed independently and when they are correlated. Otherwise, the setup is as above. The results are presented in Figure 8. As the plots suggest, increasing q from 0 to 1 (that is, increasingly emphasizing relevances in the ranking) does reduce incentives to collude somewhat. However, the reduction is not all that significant—the bound on γ increases by at most 5%. It does appear, however, that with higher settings of q there are more instances in which collusion is not feasible at all, with higher ranges of γ near 1 for most slot configurations.

6.2 Incomplete Information

While in the complete information setting we enforced collusion using a maximum symmetric equilibrium as punishment, we could have enforced it also with a minimum symmetric equilibrium, accepting that it would not obtain as an equilibrium in a smaller range of settings. The advantage of the latter approach, however, is that we no longer need to know valuations of all players at all—rather, since balanced bidding converges to a set of minimum symmetric equilibrium bids, we can simply punish by reverting to that same set of bids.

We now consider the question of whether there exists a collusive strategy for agents who do not initially know each other's valuations. Since players will not have incentive to reveal these to each other honestly, they will need to be induced. The idea is to let players play a symmetric Nash equilibrium for as many stages as necessary so that, given the discount factor, it will not pay for players to misrepresent themselves initially in order to exploit their fellow colluders later. After these initial stages playing a symmetric equilibrium, the players will be able to infer each other's valuations and play *COLLUSION* from then on.

Lemma 2. *For generic valuations, a deviation from a minimum symmetric equilibrium bid to a bid that obtains a higher slot yields a strictly lower than equilibrium utility.*

Proof. For generic values, in a minimum symmetric equilibrium $w_{j+1}b_{j+1} < w_j b_j$. Thus $c_i(v_i - p_i) \geq c_j(v_i - \frac{w_{j+1}b_{j+1}}{w_i}) > c_j(v_i - \frac{w_j b_j}{w_i})$. □

Theorem 3. *Let b_{sym} be a vector of symmetric equilibrium bids and suppose that the* COLLUSION *subgame is in equilibrium for a particular vector of δ_i. Then there exists an $S > 0$ such that the following is an SPNE:*

1. *Play b_{sym} for S stages.*
2. *Play* COLLUSION *from stage $S + 1$ to ∞.*

Proof. Since payoffs are discounted (and, thus, the discounted sum converges), for any $\delta > 1$, there is S large enough such that $\sum_{i=S+1}^{\infty} \gamma^i u_s < \delta$, where u_s is the stage payoff during the *COLLUSION* subgame. This is true for any discount factor and any u_s. As the payments are arbitrarily close to being identical during *COLLUSION*, they can be made close enough to identical to eliminate any incentive to reduce the price during collusion, since it would yield a strictly lower click-through-rate. Finally, by Lemma 2, all players strictly prefer to stay in their slot than to switch to a lower-numbered slot (lower-numbered is better, i.e., higher on the page). Thus, there exists $\delta > 0$ such that they still strictly prefer their current slot even if they get δ more from switching to the other. Thus, in particular, if their maximum payoff from the *COLLUSION* subgame does not exceed δ, they will have no incentive to switch. □

Now, consider the following dynamic strategy:

1. Play balanced bidding until a fixed point is reached.
2. Play fixed point for S stages.
3. Play *COLLUSION* from stage $S + 1$ to ∞.

We just showed that if *COLLUSION* itself is in equilibrium, we can find S large enough such that the subgame comprised of steps 2 and 3 above is also. We have also experimentally indicated that balanced bidding may itself be nearly a Bayes-Nash equilibrium in an array of settings when deviations to several other greedy bidding strategies are allowed. However, we assumed that players will play these strategies truthfully—that is, they would not play them as if their valuations were different from what they actually are. We studied incentives to play like someone else in balanced bidding separately. There, our results are somewhat less clear: in some settings, there does appear to be considerable incentive to pretend to have a different valuation, while in others such incentives seem relatively small. However, the ambiguity seems primarily due to our conservative choice of a discount factor. With a more reasonable discount factor, we argued that truthfulness seems unlikely to have significant regret. When incentives to lie are small, we expect the entire dynamic strategy in steps 1-3 to have relatively low regret, as long as the *COLLUSION* subgame does.

7 Conclusion

We have started with a set of greedy bidding strategies from Cary et al. [2007] and analyzed them in a dynamic sponsored search auction setting. Many of the results are more favorable for search engines than advertisers: a high-payoff strategy profile is not sustainable in equilibrium, whereas a low-payoff profile is reasonably stable. On the other hand, when complete information about valuations and click-through-rates is available, there are possibilities for collusion that yield high payoffs to players, sustainable over a range of settings. In the case of incomplete information, a two-stage equilibrium profile may be employed in which for some number of rounds depending on the discount factors bidders play a minimum symmetric equilibrium and then switch to a collusive strategy once the valuations become known.

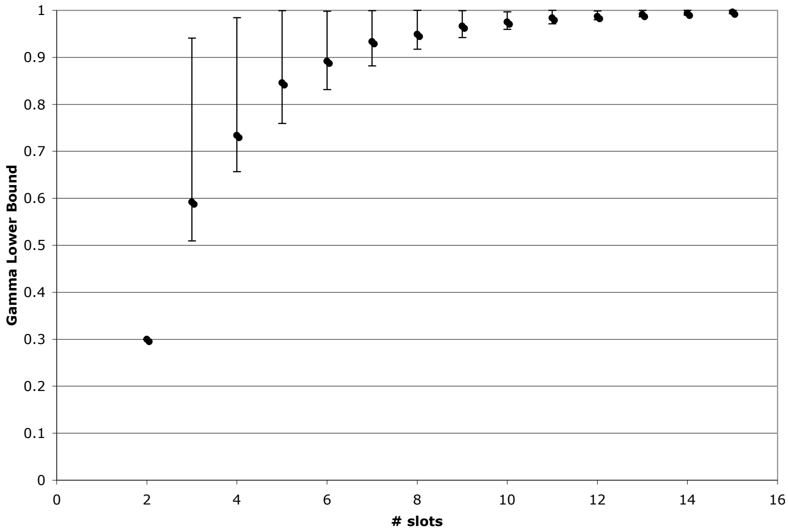

Fig. 7. Lower bounds on the discount factor as the number of available slots varies when $\delta = 1.428$.

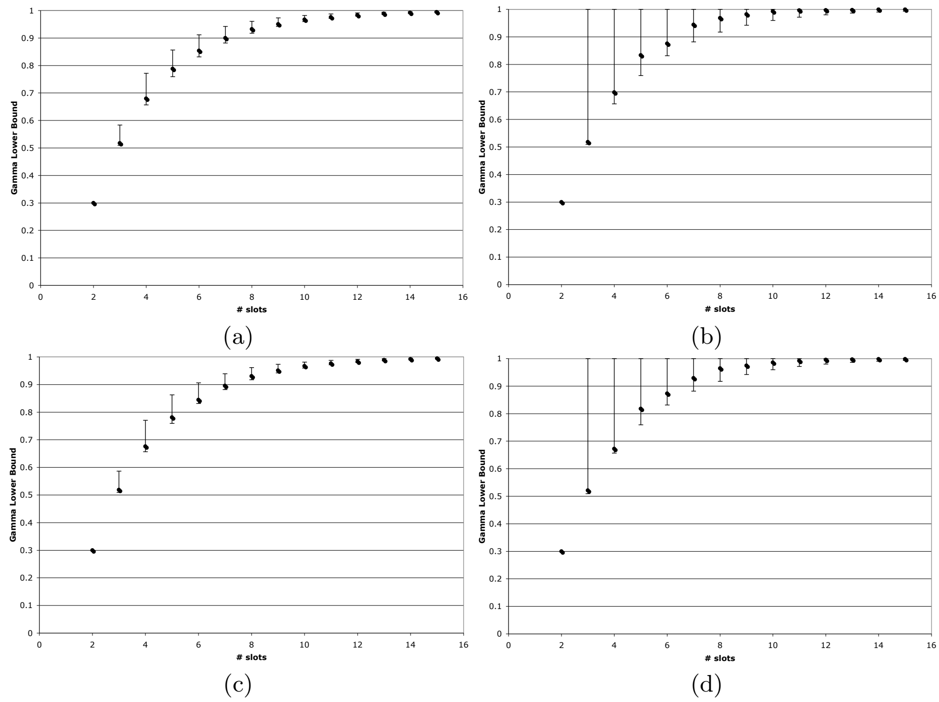

Fig. 8. Lower bounds on the discount factor as the number of available slots varies when (a) and (b) values and relevances are independent; and (c) and (d) when values and relevances are correlated. The figures on the left (a) and (c) correspond to $q = 0$, while those on the right (b) and (d) correspond to $q = 1$.

Bibliography

Matthew Cary, Aparna Das, Ben Edelman, Ioannis Giotis, Kurtis Heimerl, Anna R. Karlin, Claire Mathieu, and Michael Schwarz. Greedy bidding strategies for keyword auctions. In *Eighth ACM Conference on Electronic Commerce*, 2007.

Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 9(1):242–259, March 2007.

Juan Feng and Xiaoquan Zhang. Dynamic price competition on the internet: advertising auctions. In *Eighth ACM Conference on Electronic Commerce*, pages 57–58, 2007.

Vijay Krishna. *Auction Theory*. Academic Press, 1st edition, 2002.

Sebastien Lahaie. An analysis of alternative slot auction designs for sponsored search. In *Seventh ACM Conference on Electronic Commerce*, 2006.

Sebastien Lahaie and David M. Pennock. Revenue analysis of a family of ranking rules for keyword auctions. In *Eighth ACM Conference on Electronic Commerce*, 2007.

Li Liang and Qi Qi. Cooperative or vindictive: bidding strategies in sponsored search auctions. In *Third International Workshop on Internet and Network Economics*, Lecture Notes in Computer Science. Springer Verlag, 2007.

Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

Hal Varian. Position auctions. *International Journal of Industrial Organization*, to appear.

Appendix

A Proof of Lemma 1

Let the deviant be bidder i with value v_i and suppose w.l.o.g. that in a truthful fixed point it gets slot i and in an untruthful fixed point it gets slot j. First, consider $j = i$, that is, its bidding did not affect the slot it received. Since its payment is b_{i+1} which would remain unchanged in a minimum symmetric equilibrium as it only depends on valuations of players ranked below bidder i, its utility will remain unchanged.

Suppose it bids as if its value were $v' < v_i$ and gets slot $j > i$. Since b_{j+1} is unaffected by its untruthfulness, as it only depends on values (and bidding behavior) of bidders ranked below bidder j, and since the bidder was ranked in slot i in a minimum symmetric equilibrium when it bid truthfully, deviation to slot j was unprofitable then and it remains so now.

Finally, suppose i bids as if its value were $v' > v_i$ and gets slot $j < i$. In a truthful minimum symmetric equilibrium,

$$b_{j+1} = v_i(x_{i-1} - x_i) + \sum_{t=j}^{i-1} v_{t+1}(x_{t+1} - x_t) + \sum_{t=i+1}^{K} v_{t+1}(x_{t+1} - x_t).$$

In a minimum symmetric equilibrium which obtains when i is untruthful in the above sense, its payment in slot j is

$$b'_{j+1} = v'_i(x_{i-1} - x_i) + \sum_{t=j}^{i-1} v'_{t+1}(x_{t+1} - x_t) + \sum_{t=i+1}^{K} v'_{t+1}(x_{t+1} - x_t),$$

where v'_t is the value of bidder that would get position t if i played untruthfully. Note that bidders ranked lower than i will retain the same ranking and, as we have already observed, will have the same bids as in a truthful fixed point. Thus,

$$\sum_{t=i+1}^{K} v'_{t+1}(x_{t+1} - x_t) = \sum_{t=i+1}^{K} v_{t+1}(x_{t+1} - x_t).$$

Furthermore, all the bidders in slots $j..i-1$ will now move down one slot. Thus,

$$\sum_{t=j}^{i-1} v'_{t+1}(x_{t+1} - x_t) = \sum_{t=j}^{i-1} v_t(x_{t+1} - x_t) \geq \sum_{t=j}^{i-1} v_{t+1}(x_{t+1} - x_t).$$

Finally, the bidder previously in slot $i-1$ will now be in slot i. Thus,

$$v'_i(x_{i-1} - x_i) = v_{i-1}(x_{i-1} - x_i) \geq v_i(x_{i-1} - x_i).$$

As a result, $b'_{j+1} \geq b_{j+1}$, and if bidder i had no incentive to switch to slot j in a truthful minimum symmetric equilibrium, it certainly will not now, since it would face at least as high a price (and probably higher). Note that these inequalities are strict when values are generic; thus, obtaining a higher slot than under a truthful fixed point yields strictly lower utility to i. □

B Proof of Theorem 1

Take a player s (recall that players are indexed according to the slots they occupy) and let the discount factor of that player be γ. First, note that if $s \geq K+1$, the player can only win a slot by paying more than v_s, and thus has no incentive to deviate.

Suppose that $s \leq K$. If the player s never deviates, it will accrue the payoff of $u_s = c_s(v_s - (K-s)\epsilon - \frac{w_{K+1}v_{K+1}}{w_s}$ at every stage. With γ as the discount factor, the resulting total payoff would be $\sum_{i=0}^{\infty} \gamma^i u_s = \frac{u_s}{1-\gamma}$. For ϵ sufficiently small, there will be no incentive to deviate to an inferior slot, since it offers a strictly lower click-through-rate with negligible difference in payment. The one-shot payoff for deviating to $t \le s$ is $u'_s = c_t(v_s - \frac{w_t}{w_s}(K-t+1)\epsilon - \frac{w_{K+1}v_{K+1}}{w_s})$. For all stages thereafter, the utility will be $u^p_s = c_s(v_s - \sum_{t=s+1}^{K+1} w_{t-1}v_{t-1}\frac{c_{t-1}-c_t}{c_s}) = c_s v_s - \sum_{t=s+1}^{K+1} w_{t-1}v_{t-1}(c_{t-1}-c_t) = c_s v_s - \sum_{t=s+1}^{K} w_{t-1}v_{t-1}(c_{t-1}-c_t) - w_K v_K c_K$. Since this utility will be played starting at the second stage, the total utility from deviating is $u'_s + \frac{\gamma u^p_s}{1-\gamma}$. For deviations to be unprofitable, it must be that for every $s \le K$ and every $t \le s$, $\frac{u_s}{1-\gamma} \ge u'_s + \frac{\gamma u^p_s}{1-\gamma}$, or, alternatively, $u_s \ge (1-\gamma)u'_s + \gamma u^p_s$. Plugging in the expressions for utilities and rearranging gives us the result. $\square$