

Dynamics of bid optimization in online advertisement auctions

Christian Borgs
Microsoft Research
Redmond, WA

borgs@microsoft.com

Jennifer Chayes
Microsoft Research
Redmond, WA

jchayes@microsoft.com

Omid Etesami*
U.C. Berkeley
Berkeley, CA

etesami@cs.berkeley.edu

Nicole Immorlica
Microsoft Research
Redmond, WA

nickle@microsoft.com

Kamal Jain
Microsoft Research
Redmond, WA

kamalj@microsoft.com

Mohammad Mahdian†
Yahoo! Research
Santa Clara, CA

mahdian@yahoo-inc.com

Categories and Subject Descriptors

F.2 [Theory of Computation]: ANALYSIS OF ALGORITHMS AND PROBLEM COMPLEXITY

Keywords

Advertisement Auctions, Bidding Algorithms

ABSTRACT

We consider the problem of online keyword advertising auctions among multiple bidders with limited budgets, and study a natural bidding heuristic in which advertisers attempt to optimize their utility by equalizing their return-on-investment across all keywords. We show that existing auction mechanisms combined with this heuristic can experience cycling (as has been observed in many current systems), and therefore propose a modified class of mechanisms with small random perturbations. This perturbation is reminiscent of the small time-dependent perturbations employed in the dynamical systems literature to convert many types of chaos into attracting motions. We show that the perturbed mechanism provably converges in the case of first-price auctions and experimentally converges in the case of second-price auctions. Moreover, the point of convergence has a natural economic interpretation as the unique market equilibrium in the case of first-price mechanisms. In the case of second-price auctions, we conjecture that it converges to the “supply-aware” market equilibrium. Thus, our results can be alternatively described as a tâtonnement process for convergence to market equilibrium in which prices are adjusted on the side of the buyers rather than

*Work performed while author was an intern at Microsoft Research.

†Work performed while author was a postdoc at Microsoft Research.

the sellers. We also observe that perturbation in mechanism design is useful in a broader context: In general, it can allow bidders to “share” a particular item, leading to stable allocations and pricing for the bidders, and improved revenue for the auctioneer.

1. INTRODUCTION

Online search engine advertising is becoming an increasingly important and costly component of the marketing and sales strategies of many businesses. The corresponding auctions are the main source of revenue for many search engines and other Internet-related businesses. It is therefore of tremendous interest to understand and analyze the behavior of these auction systems, and to try and ensure that the system functions smoothly. In this paper, we consider the advertisement auction system as a whole and from a dynamic perspective. We first define a simple and natural bidding heuristic for budget-limited advertisers based on equalizing the “return-on-investment” (ROI) across keywords. We then observe that, when used by a set of advertisers, multiple copies of this heuristic may induce cycling behavior into the system. We propose circumventing this undesirable effect by introducing random perturbations, and see that this modified system converges to the market equilibrium (provably for first-price auctions and experimentally for second-price auctions). Thus our results may alternatively be interpreted as providing a tâtonnement process for convergence to market equilibrium in which prices are adjusted on the side of the buyers rather than the sellers.

Online search engine advertising is typically sold via keyword auctions (see, for example, Google’s AdWords, Yahoo’s Search Marketing, and MSN’s AdCenter). Each prospective advertiser chooses a set of keywords relevant to his products, and for each keyword submits a bid representing an estimate of his utility for a click when that word is displayed. He also submits a maximum budget which must be respected for the chosen time period. When each keyword appears, it is auctioned among all

interested advertisers with remaining budget, typically using a first-price or second-price auction mechanism (see [5, 1, 11, 16] for a comparison of these approaches).

As bidders have limited budgets, the bid optimization problem they face is essentially a discrete separable resource allocation problem [7]. One of the most popular metrics to assess the efficiency of various investment strategies is marginal “return-on-investment,” which in this context can be taken as the derivative of the utility with respect to the price. (See Section 3 for precise definitions.) Here we use an easily computable approximation to this quantity, namely the ratio rather than the derivative. For a particular advertiser, we define the ROI of a keyword at a given bid to be the ratio of the utility of this word to the price of the word, both at the given bid. In the bidding heuristic we consider, each budget-constrained advertiser bids an amount such that his ROI is equal across all keywords. Such heuristics are common in practice and have been proposed in other theoretical contexts as well [15].

Assume that the above bidding heuristic is employed by a set of advertisers. Two questions immediately arise. First, does there exist an underlying mechanism which causes these algorithms to *converge*? Second, if a convergent mechanism does exist, to what does it converge? In particular, how does this system impact revenue for the search engine provider? It is important to note that we consider these questions in light of the bidding dynamics defined by the specified heuristics, assuming all bidders adhere to these heuristics and use them truthfully regardless of the optimality of such a strategy. In particular, we do not study the properties of these systems in a strategic equilibrium.¹

The first question, namely the existence of a *convergent* mechanism, is more than just a theoretical question. Indeed, what appears to be chaotic cycling behavior has been observed in actual search engine auctions [11].² Moreover, for straightforward mechanisms used in conjunction with the ROI bidding heuristic, we can easily construct two-bidder examples which exhibit cycling, with the allocation oscillating between the bidders. These observations and examples are not surprising in light of the general phenomenon of heteroclinic cycles that can occur in both continuous [6] and discrete [14] dynamic systems with symmetry, sometimes leading to cycling chaos [3, 13].

In order to overcome this, we introduce an *online random bid perturbation* into our algorithm. In some sense, this perturbation is reminiscent of the small time-dependent perturbations employed in the dynamical sys-

tems literature to convert many types of chaos into attracting motions [12]. In mechanism design, perturbation has been proposed previously as a solution to spiteful bidding (bidding strategies which attempt to drive out competition by exhausting their budgets) [10]. Our results further motivate the introduction of perturbations to mechanism design as a technique for smoothing the dynamics of the system and permitting bidders to “share” items in arbitrary ratios.

Indeed, in the case of a first-price auction, we prove that the introduction of random perturbations causes the mechanism to converge. This is by far the most technically complex part of the paper. We conjecture that the random perturbations will also eliminate cycling behavior and lead to convergence of an analogous second-price auction, a conjecture which is supported by simulations in Section 5. Furthermore, we can prove that, in the case of the perturbed first-price auction, the prices (and hence revenue) of our system converges to the unique market equilibrium. As a side note, this also gives an algorithm for computing the market equilibrium in our setting (incidentally, the algorithm is quite similar to that of Devanur et al. [4] for computing market equilibria), as well as a tâtonnement process for convergence to market equilibrium in which prices are adjusted on the side of the buyers rather than the sellers.

All of our results are supported by simulations, which we discuss in Section 5.

2. MODEL

Search engines often display advertisements alongside search results when a user performs a search. These advertisements appear in a dedicated area of the search results page, each one in a particular fixed subarea, or *slot*. An online advertisement auction is a mechanism for selling these slots based on the keyword which the user provided to the search engine.

We consider a setting in which m advertisers bid for the advertising slots of n keywords. Each keyword j has l slots and appears $q_j(t)$ times on day t (by “day” we mean some fixed unit of time; it does not necessarily have to be 24 hours). Advertiser i has a value v_{ij} for each click received when his advertisement is displayed on keyword j . Note that while advertisers value clicks, our auction is actually selling *impressions*, or the chance to appear in a keyword slot. We can convert the values per click to an expected *value per impression* u_{ijk} by taking the product of v_{ij} with the probability c_{ijk} that advertiser i receives a click when displayed in slot k of keyword j . This probability is called the *click-through-rate*. We assume these click-through-rates factor, that is, there exist β_{ij} for each bidder i and keyword j , and α_k for each slot k (independent of the advertiser and keyword)³ such that $c_{ijk} = \beta_{ij}\alpha_k$. Thus the per impression bid u_{ijk} for the k ’th slot can be written as $\alpha_k u_{ij}$ for some u_{ij} . We number slots in order of decreasing click-through-rate so $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_l$ and without loss of generality assume $\alpha_1 = 1$.

³The assumption that the click-through-rate can be decomposed in this way is a reasonable assumption and is used in practice.

¹A major difficulty in studying this setting as a strategic game is the repeated nature of the game. Folklore theorems show that repeated games (such as this one) have a plethora of equilibria, thereby making equilibrium analysis (without any restriction on the set of available strategies) unsuitable for predicting the behavior of the system. In this work, we are taking a different route: we fix a particular bidding strategy (whose variants are used in practice) and analyze the equilibrium this strategy.

²For an alternative justification of observed cycling patterns see [17].

Each advertiser submits a bid b_{ij} for each keyword representing the amount he is willing to pay for one impression in slot 1 of keyword j (i.e., u_{ij} above). By extension, we assume he is willing to pay $\alpha_k b_{ij}$ for an impression in slot k of keyword j .⁴ Advertisers additionally submit a daily budget B_i indicating the maximum amount they are willing to spend in a given day. Although in general these parameters may be adjusted at arbitrary times, for simplicity we assume they are updated at most daily and in the beginning of the day.

Upon a search for a particular keyword j , the advertisement auction then selects up to l advertisers i_1, \dots, i_l and assigns them to slots $1, \dots, l$, respectively. It then computes a price p_{jk} for each advertiser $i_k \in \{i_1, \dots, i_l\}$. The auction guarantees no bidder is charged more than his bid nor exceeds his budget. Furthermore, no bidder is awarded more than one slot per search query. We focus our attention on two particular auction mechanisms quite common in practice. The first is a *first-price mechanism* in which advertisers are awarded slots in a priority order determined by their bids. Advertisers are then charged a price equal to the minimum of their bid and remaining budget. The second mechanism is a generalization of the *second-price mechanism*. The allocation rule of this mechanism is identical to that of the first-price mechanism, but the pricing scheme is different. Each advertiser is now charged a price equal to the minimum of his remaining budget and the bid of the advertiser in the next slot. The pseudocode of these two mechanisms appears in Figure 1.

For our theoretical results, we simplify the model in the following ways. First, we study a setting in which there is only one slot per keyword. The single-slot setting is rich enough to capture the chaotic behavior our results circumvent and thus suffices to illustrate our main points.⁵ Second, we consider a continuous-time version of the auction: For each keyword j , there are a constant number q_j of searches each day, and these searches are evenly spaced throughout the day. We assume q_j 's are large and therefore we can model this process as one in which all keywords arrive continuously at a uniform rate throughout the day. The daily budget of advertiser i is B_i , and the total utility of advertiser i for showing his ad on keyword j throughout the entire day is u_{ij} (thus, his utility for being shown during an α fraction of the day is αu_{ij}). Without loss of generality we will assume $B_i \leq \sum_j u_{ij}$.

3. BID OPTIMIZATION HEURISTICS

In this section we describe a natural bidding heuristic for optimizing the utility of the advertisers. We consider the following abstraction of the bid optimization prob-

⁴Note we could just have easily described our results for a setting where advertisers submit a bid per click if we assume the click-through-rates of advertisers and slots are known or estimated.

⁵In fact, it is straightforward to generalize our convergence result (Theorem 1) to the multi-slot setting (essentially the only thing that needs to be changed is Equation 1). However, the point to which the system converges can no longer be characterized as a market equilibrium.

lem for advertiser i . We want to specify a bid b_{ij} on each keyword j . We assume that if advertiser i bids b_{ij} on keyword j then his day-long charge and net utility (i.e., total value minus total charge) on that keyword is given by $P_j(b_{ij})$ and $U_j(b_{ij})$ respectively.⁶ The optimization problem is now to choose $\{b_{ij}\}$ such that $\sum_j U_j(b_{ij})$ is maximized subject to $\sum_j P_j(b_{ij}) \leq B_i$. Through the use of Lagrangian relaxation, we see that a necessary condition for the optimality of bids b_{ij}^* is the existence of a constant λ (the Lagrangian multiplier) such that for all j with $U_j(b_{ij}^*) > 0$,

$$dU_j/dP_j |_{b_{ij}=b_{ij}^*} = \lambda$$

if such derivatives exist. This derivative is known as the *marginal return-on-investment* (marginal ROI) and measures how the net utility of an advertiser changes as he modifies his investment. Thus, for an optimal set of bids $\{b_{ij}^*\}$, we know advertiser i has the same marginal ROI at b_{ij}^* across all keywords. This marginal ROI is exactly the Lagrangian multiplier λ above.

The marginal ROI is usually difficult to estimate, and is even undefined when P_j or U_j are discontinuous. Thus, it is useful to approximate the marginal ROI of keyword j at bid b by the ROI of keyword j at that bid, where ROI is defined as $\text{ROI}_j(b) = U_j(b)/P_j(b)$. This suggests one method for optimizing the bids of the advertiser: set the bids b_{ij} such that $\text{ROI}_j(b_{ij})$ approximately equals some constant ROI for all j .

If the prices were fixed and known to the advertiser, determining an optimal bidding vector would be a simple calculation. Suppose the price of the k^{th} slot for keyword j is p_{jk} . We further introduce an artificial slot $l+1$ with price zero and utility zero indicating that the advertiser does not appear in any slot on that keyword. A bidding strategy is now a selection of affordable slots $s_j \in \{1, \dots, l+1\}$ for each keyword j , where a selection is affordable if the sum of prices is at most the budget of the advertiser. This problem is a natural extension of the knapsack problem [8] and has a similar FPTAS.

In fact, the idea of the ROI heuristic is similar to the well-known 2-approximation for knapsack. It tries to maintain the invariant that for some constant $R = \text{ROI}$, $R \in (u_{js_j}/p_{js_j}, u_{j(s_j+1)}/p_{j(s_j+1)})$ for all keywords j , and searches for the maximum possible R subject to the budget constraint. Thus, if the advertiser has budget left over at the end of the day, he finds the keyword j with minimum u_{js_j}/p_{js_j} and chooses slot s_j+1 for keyword j on the following day. Otherwise, if he ran out of budget early, he finds the keyword j with maximum u_{js_j}/p_{js_j} and chooses slot s_j-1 for that keyword on the following day.

An alternative way to implement the ROI heuristic is through a tâtonnement-like process, where the advertiser iteratively incrementing bids on keywords with relatively large ROI and decrementing bids on keywords

⁶Note that we assume the charge and net utility of advertiser i for keyword j is a function of his bid for keyword j alone and does not depend on the bids of i for other keywords. Although this is not strictly true, it is a reasonable approximation and serves to develop our intuition for our heuristic.

<p>First-Price Mechanism Let S be the set of bidders $\{i : s_i \leq B_i\}$. For $k = 1$ to l do Let $i = \operatorname{argmax}_{i \in S}(b_{ij})$, Set $S = S - \{i\}$, Assign i to slot k, Charge i price $\min(\alpha_k b_{ij}, B_i - s_i)$.</p>	<p>Second-Price Mechanism Let S be the set of bidders $\{i : s_i \leq B_i\}$. For $k = 1$ to l do Let $i = \operatorname{argmax}_{i \in S}(b_{ij})$, Set $S = S - \{i\}$, Assign i to slot k, Charge i price $\min(\alpha_k \max_{i' \in S} b_{i'j}, B_i - s_i)$.</p>
--	--

Figure 1: Pseudocode for the first and second-price auctions, respectively. The parameter s_i is the current total daily charge of advertiser i .

with relatively small ROI by small increments. The advantage of this method is that it requires the minimal amount of information. In particular, it does not even need to know the price of the slots above and below the current slot. It is easy for an advertiser to calculate the ROI for each keyword in hindsight at the end of the day. Based on this idea, we consider the following ROI-based heuristic bidding algorithm for advertiser i .

ALGORITHM 1. *On each day t , all bids of advertiser i are determined by a single parameter $R_i(t) \in (0, 1]$.⁷ The parameter $R_i(t)$ is adjusted based on the performance of advertiser i 's bids on the previous day. Starting from an arbitrary $R_i(0) \in (0, 1]$ for day $t = 0$, advertiser i sets*

$$R_i(t+1) = \begin{cases} R_i(t)e^{-\epsilon} & \text{if } i \text{ runs out of money} \\ & \text{before the end of day } t \\ \min(R_i(t)e^{\epsilon}, 1) & \text{otherwise} \end{cases}$$

where $\epsilon > 0$ is a small constant. Finally, he sets the bid $b_{ij}(t)$ of keyword j to

$$b_{ij}(t) = R_i(t)u_{ij}.$$

Note since $R_i(t) \in (0, 1]$, $b_{ij}(t) \leq u_{ij}$.

Before discussing the dynamics of this algorithm, let us note that an added advantage of the above bidding heuristic is that it can be adapted to cases where an advertiser only knows her budget and the relative utility of various keywords (i.e., the ratio of u_{ij} 's), and not the exact value of the utilities. In this case, the bidding algorithm for advertiser i can initially set her largest utility to B_i and the other utilities according to the specified ratios, and then adjust these values by changing $R_i(t)$. This is useful in practice since for an advertiser estimating the ratio of the values of various keywords is a considerably simpler task than estimating the exact utilities.

4. DYNAMICS OF THE SYSTEM

In Section 3 we defined a heuristic for bidding in an advertisement auction. In order to better understand the properties of a system where bidders are using such

⁷This parameter is related to the target return-on-investment by $R_i(t) = 1/(\text{ROI} + 1)$ where ROI is the target return on investment of advertiser i .

a heuristic, we need to analyze the interplay of bidding algorithms of various bidders. One might wonder if such a system could ever stabilize, and whether the resulting prices would be logical in some sense (i.e., be simultaneously “reasonable” for the advertisers and generate sufficient revenue for the search engine). In fact, the following example shows that the combination of the first-price auction with the ROI heuristic may result in an unstable situation with low prices.

EXAMPLE 1. *Suppose there is just one keyword with one slot and 1000 impressions. There are two advertisers a and b , each advertiser with a budget of \$500 and a utility of \$1 for each impression of the keyword. Consider the first-price auction mechanism. Assume a bids \$0.5e ϵ , and b bids \$0.5. Bidder a is going to win all the impressions until he runs out of the budget around the end of the day, but he is going to decrease his bid for tomorrow to \$0.5, since he ran out of budget today. On the other hand, b is going to bid \$0.5e ϵ on the following day. Thus, a and b will interchange roles. This way the allocation of the impressions alternates between a and b daily.*

It is easy to see that the above example works for the second price mechanism as well. The results of Section 5 confirm that such examples arise in a variety of plausible scenarios, resulting in oscillating allocations and dampened revenue. We avoid such situations by applying a random perturbation to the bids of the advertisers in determining the allocation, as defined below. In this section we study variants of the first and second-price auctions with perturbations. We prove that the perturbed first-price auction, coupled with multiple copies of the bid optimization algorithm presented in Section 3, converges to a fixed allocation and set of prices corresponding to the *market equilibrium*. We conjecture a similar result for the perturbed second-price auction, supporting our conjecture with simulation results in Section 5.

4.1 Perturbations

In order to get rid of situations like the one explained in Example 1, we modify the auction mechanism to slightly perturb the bids before running the auction, thereby giving the bidder with a smaller bid some chance of winning if his bid is close to the largest bid. The perturbations are defined as follows. On each day t , advertiser i bids a value $b_{ij}(t)$ for the day-long possession

of keyword j . When a search on keyword j occurs, we perturb the bids as follows:

$$b'_{ij} = b_{ij}(t)e^{-\eta_i},$$

where η_i is a uniformly random number in $[0, \delta]^8$, independently generated for each bidder/query pair, and $\delta > 0$ is a constant. The auction mechanisms are run exactly as described in Section 2, but the allocation is determined according to the perturbed bids $b'_{ij}(t)$.

Perturbations essentially allow advertisers to bid such that they share the keyword in *any portion* they please. That is, fixing the bids of other advertisers on a particular keyword, a given advertiser can choose to receive in expectation any fraction α of the day-long procession of the keyword by adjusting his bid appropriately. Note that such a sharing property can not be achieved by introducing a randomized tie-breaking rule; applying the perturbation to the bids themselves is significantly more powerful. Notice how this affects the advertisers in Example 1.

EXAMPLE 2. *Again, consider the scenario from the previous example. However, now suppose the bids are perturbed as described above and notice the instability we observed before won't happen. Indeed a and b share the impressions almost equally in expectation, and so neither bidder runs out of budget. Therefore, they will increase their bids until their bids get close to \$1 at which time both the price and allocations remain stable. In this case the perturbation both removed the cycling and improved auctioneer's revenue by a factor of two.*

4.2 Convergence to Equilibria

We now discuss our main theoretical results, namely the convergence properties of our perturbed mechanisms with multiple bid optimization algorithms. Throughout the remainder of this section, we assume there is just one slot per keyword.⁹

We consider both perturbed first-price and perturbed second-price auctions. In each of these auctions, the allocation rule awards the keyword slot to the bidder with the highest perturbed bid b'_{ij} . The winning advertiser is then charged a price equal to the minimum of his remaining budget and unperturbed bid b_{ij} in the case of the first-price auction¹⁰, or the minimum of his remaining budget and the perturbed bid of the closest competitor in the case of the second-price auction. Once the spending of an advertiser during a day reaches his daily budget, he is withdrawn from all further auctions during that day.

We now state our principal result. Namely, we prove that in a perturbed first-price auction where bidders bid according to the ROI heuristic, Algorithm 1 of Section 3, both the prices and the daily utilities of the advertisers,

⁸The choice of the distribution for perturbation is essentially arbitrary, and our results hold for other reasonable perturbation models (e.g., Gaussian perturbations) as well.

⁹It is not hard to see that Theorem 1 holds for the multi-slot case with essentially the same proof.

¹⁰Note that our results hold if the pricing rule charges the winning bidder his perturbed bid b'_{ij} as well.

and hence the revenue of the auctioneer, converge to that of the market equilibrium in the sense of Arrow and Debreu [2] when goods correspond to the ad spaces and the money (see Appendix A).

More formally, let $s_i(t) \in [0, B_i]$ denote the spending of advertiser i on day t . Let $\tau_i(t) \in [0, 1]$ denote the moment during day t when advertiser i spends all his budget (or 1 if he does not spend all his budget). Finally let $r_i(t)$ denote the spending rate of advertiser i in the beginning of the day *before* anyone runs out of budget. In other words,

$$r_i(t) = \sum_{j=1}^n \frac{b_{ij}(t)}{\delta} \int_0^\delta \prod_{i' \neq i} \Pr[b_{ij}(t)e^{-x} > b_{i'j}(t)e^{-\eta_{i'}}] dx \quad (1)$$

Note that the rate of spending only increases as other advertisers run out of budget, and therefore we have $s_i(t) \geq r_i(t)\tau_i(t)$. We first show these parameters converge, namely, that after some time no advertiser runs out of budget early and each advertiser either spends most of his budget or is bidding nearly his utility on all keywords. The proof of the following theorem appears at the end of this section.

THEOREM 1. *Given utilities u_{ij} , budgets B_i , and constants $\delta > 0$ and $\gamma > 0$, there exist constants $\epsilon > 0$ and $t_0 < \infty$, such that for all $t \geq t_0$ and all i , we have*

1. $\tau_i(t) \geq 1 - \gamma$, and
2. $s_i(t) \geq (1 - \gamma)B_i$ or $R_i(t) \geq 1 - \gamma$.

Here ϵ and t_0 can be chosen as $\epsilon = \Theta(\gamma \min\{1, \delta/C^2\})$ and $t_0 = (2 \log C)/\epsilon - \log(\min_i R_i(0))$ with $C = \max_i (\sum_j u_{ij}/B_i)$.

The above theorem allows us to characterize the equilibrium of our system. Let $L_i(t) = B_i - s_i(t)$ be the unused portion of advertiser i 's budget at the end of day t . Then the following theorem holds.

THEOREM 2. *Given $\delta > 0$ and $\gamma > 0$, let $t = t(\delta, \gamma) \geq t_0$, where t_0 is defined as in Theorem 1. Let $p_j(t)$ be the maximum price at which keyword j is sold in day t , and let $x_{ij}(t)$ be the fractional daily allocation of word j to advertiser i on day t . As δ, γ go to zero, the price vector $p_j(t)$ converges to that of the market equilibrium, and the total utilities of the advertisers including their unused budgets, $L_i + \sum_j u_{ij}x_{ij}(t)$, converge to the utilities of an equilibrium allocation.*

Notice that convergence of the price vector implies also convergence of the total revenue $\sum_i p_i$ for the auctioneer. The proof of Theorem 2, which makes substantial use of the stability results in Theorem 1, is deferred to Appendix A.

PROOF OF THEOREM 1. We first show Statement 1, i.e. that after some finite time nobody runs out of budget early. More precisely, we will show that for every $0 < \lambda < 1$, ϵ small enough and $t \geq T_\lambda$ (where T_λ is a constant depending on λ), we have $\tau_i(t) \geq 1 - \lambda$ for all $1 \leq i \leq n$. Let $k(t)$ be the first advertiser who finishes his budget on day t . The proof of Statement 1 follows from the following two claims.

CLAIM 1. If $\tau_{k(t)}(t-1) < 1$, then

$$\tau_k(t)(t) \geq \min(e^\epsilon \tau_{k(t)}(t-1), 1).$$

CLAIM 2. If $\tau_{k(t)}(t-1) = 1$, then $\tau_{k(t)}(t) \geq 1 - \lambda$, provided ϵ is chosen in such a way that $2C\epsilon e^\epsilon \leq \lambda\delta$.

To see that these two claims imply Statement 1 of the theorem, set $\tau_{min}(t) = \min_i \tau_i(t)$. Claims 1 and 2 together imply $\tau_{min}(t) \geq \min(1 - \lambda, e^\epsilon \tau_{min}(t-1))$. We know that $\tau_{min}(1) \geq \min_i B_i / (\sum_j u_{ij}) = 1/C$. Therefore for $t \geq T_\lambda = \epsilon^{-1} \log(C(1 - \lambda))$, we have $\tau_{min}(t) \geq 1 - \lambda$, as required.

PROOF OF CLAIM 1. Throughout this proof, let $k = k(t)$. If $\tau_k(t) = 1$, then the claim is true. Assume $\tau_k(t) < 1$. Note that since $\tau_k(t-1) < 1$, $R_k(t) = R_k(t-1)e^{-\epsilon}$ and for $i \neq k$, $R_i(t) \geq R_i(t-1)e^{-\epsilon}$. Consider an imaginary scenario in which on day t , $\hat{R}_i(t) = R_i(t-1)e^{-\epsilon}$ for all bidders i . By (1), the spending rate $\hat{r}_k(t)$ of bidder k in the imaginary scenario is at least that of the real scenario ($\hat{r}_k(t) \geq r_k(t)$). Furthermore, $\hat{r}_k(t) = r_k(t-1)e^{-\epsilon}$ since advertisements in the imaginary scenario are sold to advertisers with the same probabilities as day $t-1$ and at a price $e^{-\epsilon}$ times the price of day $t-1$. Therefore, we have

$$r_k(t-1)\tau_k(t-1) \leq B_k = \tau_k(t)r_k(t) \leq \tau_k(t)r_k(t-1)e^{-\epsilon}$$

which implies Claim 1. \square

In order to prove Claim 2, we first prove the following lemma.

LEMMA 1. For all t and all i , we have $|r_i(t) - r_i(t-1)| \leq (2C\epsilon e^\epsilon / \delta)B_i$.

PROOF. Note that $R_i(t) \leq R_i(t-1)e^\epsilon$ and $R_{i'}(t) \geq R_{i'}(t-1)e^{-\epsilon}$ for $i' \neq i$. Consider an imaginary scenario in which on day t , $\hat{R}_i(t) = R_i(t-1)e^{2\epsilon}$ and $\hat{R}_{i'}(t) = R_{i'}(t-1)$ for $i' \neq i$. Then $\hat{R}_i(t) \geq e^\epsilon R_i(t)$ and $\hat{R}_i(t)/\hat{R}_{i'}(t) \geq R_i(t)/R_{i'}(t)$, which implies that now $\hat{r}_i(t) \geq r_i(t)e^\epsilon$. We couple the perturbed bids $\hat{b}'_{i'j}(t)$ of the imaginary scenario with the perturbed bids $b'_{i'j}(t-1)$ of day $t-1$ in such a way that $\hat{b}'_{i'j}(t) = b'_{i'j}(t-1)$ if $i' \neq i$ and $\Pr[\hat{b}'_{ij}(t) \neq b'_{ij}(t-1)] = 2\epsilon/\delta$. Namely, we set

$$\hat{b}'_{ij}(t) = \begin{cases} b'_{ij}(t-1) & \text{if } b'_{ij}(t-1) \geq \hat{b}_{ij}(t) \exp(-\delta) \\ b'_{ij}(t-1)e^\delta & \text{otherwise} \end{cases}$$

As the ratio of $\hat{b}_{ij}(t)$ to $b_{ij}(t-1)$ is $e^{2\epsilon}$, it is easy to see that this coupling results in the desired probability. Thus, even if advertiser i wins all auctions in which $\hat{b}'_{ij}(t) \neq b'_{ij}(t-1)$ (which happens at most a $2\epsilon/\delta$ fraction of the times), we have

$$\hat{r}_i(t) \leq r_i(t-1) + \frac{2\epsilon}{\delta} \sum_j u_{ij} e^{2\epsilon} \leq r_i(t-1) + \frac{2\epsilon}{\delta} CB_i e^{2\epsilon}$$

Using that $\hat{r}_i(t) \geq r_i(t)e^\epsilon$, this implies $r_i(t) \leq r_i(t-1) + (2C\epsilon e^\epsilon / \delta)B_i$. The matching upper bound on $r_i(t-1)$ in terms of $r_i(t)$ is proved by exchanging the roles of t and $t-1$. \square

PROOF OF CLAIM 2. Let $k = k(t)$. By the previous lemma and our condition on ϵ , we have

$$r_k(t) \leq r_k(t-1) + \lambda B_k \leq B_k(1 + \lambda) = r_k(t)\tau_k(t)(1 + \lambda)$$

where we used the assumption $\tau_k(t-1) = 1$ to conclude that $r_k(t-1) \leq B_k$. This gives $\tau_k(t) \geq 1/(1 + \lambda) \geq 1 - \lambda$, proving the claim. \square

Now we will prove Statement 2. Note that $r_i(t) \geq B_i(1 - \gamma)$ implies $s_i(t) \geq B_i(1 - \gamma)$ (this is because either $s_i(t) = B_i$ or $\tau_i(t) = 1$ in which case $s_i(t) \geq r_i(t)$). Therefore, it is enough to show that for all $t \geq 2T_\lambda - \log(\min_i R_i(0))$ and all i , one of the following holds:

$$r_i(t) \geq (1 - \gamma)B_i, \quad (2)$$

$$R_i(t) \geq e^{-\epsilon} \quad (3)$$

so long as ϵ is less than γ . We first prove the following claim.

CLAIM 3. For $2C\lambda \leq \gamma$, $4C\epsilon e^\epsilon \leq \gamma\delta$, and $(t-1) \geq T_\lambda$, we have $s_i(t-1) - r_i(t) \leq \gamma B_i$.

PROOF. By Statement 1, $\tau_{min}(t-1) \geq (1 - \lambda)$, and therefore $s_i(t-1) \leq r_i(t-1)(1 - \lambda) + \lambda \sum_j u_{ij} \leq r_i(t-1) + \gamma B_i/2$ provided $2C\lambda \leq \gamma$. Moreover, by Lemma 1 and our condition on ϵ , we have $r_i(t-1) \leq r_i(t) + \gamma B_i/2$. Therefore $s_i(t-1) \leq r_i(t) + \gamma B_i$. \square

The proof of Statement 2 now follows by backwards induction. First suppose neither (2) nor (3) holds on day t and $t-1 \geq T_\lambda$. We will show neither inequalities holds on day $t-1$. Indeed, by the above claim, $s_i(t-1) \leq r_i(t) + \gamma B_i < B_i$ and hence $R_i(t) = \min(R_i(t-1)e^\epsilon, 1) \geq R_i(t-1)$. Therefore (3) did not hold on day $t-1$ as well, which implies that $R_i(t) = R_i(t-1)e^\epsilon$. Now using an argument similar to Claim 1, we can show that $r_i(t) \geq r_i(t-1)e^\epsilon$. It follows that (2) did not hold on day $t-1$ either.

For the base case, notice that as long as neither (2) nor (3) holds, we saw in the above paragraph that $R_i(t) = R_i(t-1)e^\epsilon$ and so for $t \geq 2T_\lambda - \log(\min_i R_i(0))$, inequality (3) will hold.

The above result shows that the prices in a perturbed first-price mechanism converge. We believe that a similar result holds for a perturbed second price auction (see next section for evidence of this in simulation results). However, our proof technique fails for the second price auction. Given the convergence result, in Theorem 2 (whose proof is presented in Appendix A) we show that for the first price auction, the prices converge to the market equilibrium prices. For the second-price auction, assuming our conjecture on the convergence of the system, we can similarly show that the prices tend to approximate equilibria for a new notion of market equilibrium, called the *self-competition-free* or *supply-aware* market equilibrium (see [9]). A supply-aware equilibrium for a market with additive utilities is a regular market equilibrium for a modified setting in which the utility of each buyer for each item is capped to the utility they derive by buying the entire supply of the item. The simulations in Section 5 support our intuitions.




Figure 2: Change in bids, Examples 1 and 2




Figure 4: Change in efficiency, random instance




Figure 3: Change in revenue, Examples 1 and 2




Figure 5: Change in revenue, random instance

5. SIMULATIONS

In this section, we present the results of simulating the bid optimization algorithm of Section 3 for various auction mechanisms. In particular, we compare the behavior of the bid optimization algorithm in the equilibrium for the first and second-price auctions with and without perturbation.

Parameters of the simulation: We have implemented the simulation program in Matlab. In all our simulations, we assume that $\alpha_k = 1/k$ (i.e., click-through rates of different slots follow a power law with exponent -1). We assume that throughout the day, each keyword is searched for 1000 times, and these searches occur in a random order. At the end of each day, the bid optimization algorithm is run to update the bids of each advertiser. For most simulations, the parameters ϵ (determining the aggressiveness of the bid optimization algorithm in changing bids) and δ (determining the extent of the perturbations for perturbed mechanisms) are set to 0.01 and 0.1, respectively.

A small example: We start by showing the outcome of the simulation for the instance explained in Examples 1 and 2 for 500 days. In this instance, there are two advertisers and one keyword with one ad slot. Each advertiser has a utility of \$1 and a daily budget of \$500. Both advertisers start by bidding \$0.20 on each

keyword. The graph of the bid of the first advertiser as a function of time for each of the four mechanisms is shown in Figure 2 (the second advertiser has similar bids). As we see in this figure, in unperturbed mechanisms, the bids of the advertisers grow only to \$0.50, and after that remain constant, whereas in perturbed mechanisms, the bids grow to \$1. The revenue of the mechanisms are compared in Figure 3.¹¹ Since the utilities in this example are equal, the efficiency of all mechanisms are constant over time.

A larger example: We have simulated the bid optimization algorithm with different mechanisms on larger instances generated at random. Figures 6, 4 and 5 show the changes in the bids on two keywords, and the efficiency and the revenue of the auctions (per day) as a function of the day for an instance with $n = 20$ bidders, $m = 10$ keywords, and one slot per keyword. In this instance, each advertiser bids on each keyword with

¹¹The decrease in the revenue of the perturbed second-price auction (compared to the first-price) is due to the fact that after a short while, the randomness in the system could cause the bid of one of the advertisers to be slightly more than the other, resulting in the advertiser running out of budget earlier than the other advertiser, and the other advertiser getting the remaining ad spaces in that day for free.

probability $1/3$, and the value of the bids are drawn uniformly at random from $[0, 1]$. The daily budgets of the advertisers are $3000, 3000/2, 3000/3, \dots, 3000/20$.¹² As Figure 6 shows, the mechanisms with perturbation avoid having bids that are almost equal and frequently change order, whereas in mechanisms without perturbation, such situations are common. This can be observed from the diagram of efficiency in Figure 4, where it can be observed that the efficiency of the allocation on odd-numbered days are significantly lower than the efficiency of the mechanism on even-numbered days.

Random instances: We have simulated the bid optimization algorithm with each of the four auction mechanisms on a set of 150 randomly generated instances to measure the average behavior of the algorithm in different auctions. The instances are generated similar to the way described in the previous example, with 10 bidders, 5 keywords, and 3 slots per keyword. We have simulated the auctions for 300 days, and measured the following parameters: the convergence of system, and the efficiency and the revenue of the auction.

Convergence. To measure the convergence, we check the properties required in the statement of Theorem 1, and compute the fraction of bidders for whom both of these properties are satisfied at the end of the simulation. We say we have *perfect convergence* if these conditions (for $\gamma = 0.1$) are satisfied for all bidders and *good convergence* if they are satisfied for 90% (i.e., all but at most one) of the bidders after 1000 steps. Figure 7 shows the distribution of the number of converged bidders, and Figure 8 compares the percentage of the times perfect or good convergence is achieved on the four mechanisms. In this figure, mechanisms 1, 2, 3, and 4 represent the first price, the second price, the perturbed first price, and the perturbed second price mechanisms, respectively. These figures confirm our result that perturbed mechanisms are significantly more likely to converge to an equilibrium.

Revenue and Efficiency. The comparison of the revenue and the efficiency of the mechanisms reveals that in this set of instances, the revenue and the efficiency of the perturbed mechanisms are consistently (between 79% and 97% of the times) more than the unperturbed mechanisms. However, the difference is small (between 1.5% and 5% on average).

Acknowledgments: We would like to thank Max Chickering, Uri Feige, and Chris Meek for many fruitful discussions regarding the systems proposed in this paper.

6. REFERENCES

- [1] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, 2006.
- [2] K. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22:265–290, 1954.
- [3] M. Dellnitz, M. Field, M. Golubitsky, A. Hohmann, and J. Ma. Cycling chaos. *Intern. J. Bifur. & Chaos*, 5:1243–1247, 1995.
- [4] N.R. Devanur, C.H. Papadimitriou, A. Saberi, and V.V. Vazirani. Market equilibrium via a primal-dual-type algorithm. In *Proceedings of the 43rd Symposium on Foundations of Computer Science*, 2002.
- [5] J. Feng, H.K. Bhargava, and D. Pennock. Comparison of allocation rules for paid placement advertising in search engines. In *Proceedings of the 5th International Conference on Electronic Commerce*, 2003.
- [6] M.J. Field. Equivariant dynamical systems. *Trans. Amer. Math. Soc.*, 259:185–205, 1980.
- [7] T. Ibaraki and N. Katoh. *Resource Allocation Problems: Algorithmic Approaches*. MIT Press, Cambridge, MA, 1988.
- [8] O.H. Ibarra and C.E. Kim. Fast approximation algorithms for the knapsack and sum of subset problems. *Journal of the ACM*, 22(4):463–468, 1975.
- [9] K. Jain and K. Talwar. Truth revealing market equilibria. Manuscript, 2004.
- [10] C. Meek, D. Chickering, and D. Wilson. Stochastic and contingent-payment auctions. In *1st Workshop on Sponsored Search Auctions*, 2005.
- [11] M. Ostrovsky, B. Edelman, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. In *2nd Workshop on Sponsored Search Auctions*, 2006.
- [12] E. Ott, C. Grebogi, and J.A. Yorke. Controlling chaos. *Physics Review Letters*, 64, 1990.
- [13] A. Palacios. Cycling chaos in one-dimensional couple iterated maps. *Intern. J. Bifur. & Chaos*, 12m:1859–1868, 2002.
- [14] A. Palacios. Heteroclinic cycles in coupled systems of difference equations. *J. Difference Eqs & Appl.*, 9:671–686, 2002.
- [15] P. Rusmevichientong and D. Williamson. An adaptive algorithm for selecting profitable keywords for search-based advertising services. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, 2006.
- [16] H. Varian. Position auctions. Manuscript, 2006.
- [17] X.M. Zhang and J. Feng. Price cycles in online advertising auctions. In *Proceedings of the 26th International Conference on Information Systems (ICIS)*, 2005.

APPENDIX

A. PROOF OF THEOREM 2

We start by recalling some standard definitions, as applied to our setting. Given the prices p_j for keywords, an optimal allocation x_{ij} for advertiser i is any solution

¹²The choice of budgets as a power law distribution with exponent -1 is motivated by the classical observation that income distribution often follows such a power law.




Figure 6: Change in the bids in a random instance

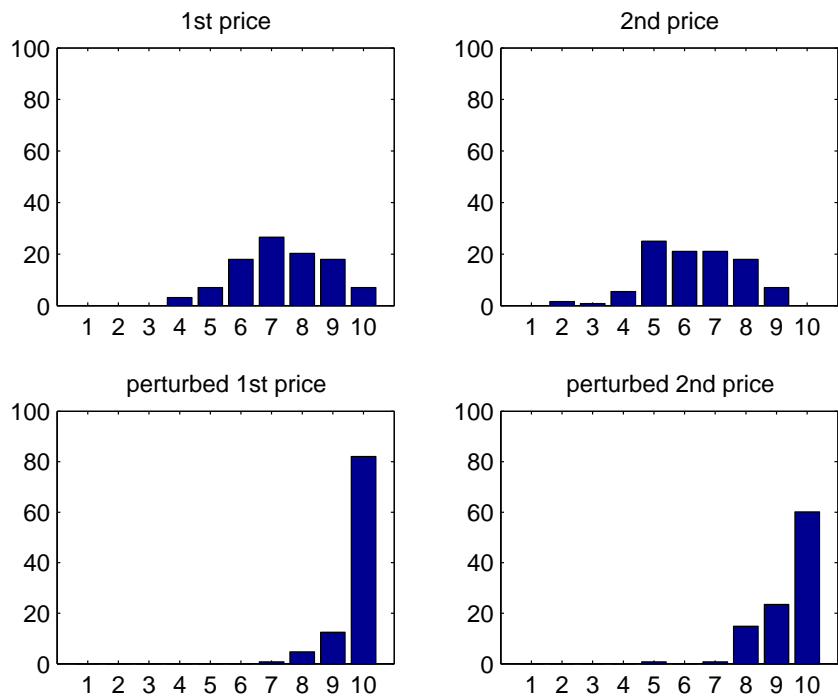


Figure 7: Distribution of the number of converged bidders




Figure 8: Distribution of the number of converged bidders

to the following linear program:

$$\begin{aligned}
& \text{maximize} && L_i + \sum_j u_{ij} x_{ij} \\
& \text{subject to} && L_i + \sum_j p_j x_{ij} = B_i \\
& && \forall j : x_{ij} \geq 0 \\
& && L_i \geq 0.
\end{aligned}$$

Here x_{ij} is the fractional amount of word j assigned to the advertiser i , and L_i is the amount of money unspent by i . A price vector is called a *market equilibrium* price vector if there exist allocations x_{ij} that satisfy the following two conditions:

- At the given price vector, x_{ij} is an optimal allocation for each advertiser i .
- For each keyword j , we have $\sum_i x_{ij} = 1$ (recall that the supply of each keyword was assumed to be 1).

The next theorem follows from the classical results in the economic literature (see, for example, Arrow and Debreu[2]) by considering the market with one commodity for each keyword and an additional commodity termed “money”. The proof of this theorem is deferred to the full version of the paper.

THEOREM 3. *There exists an equilibrium price vector. Moreover, the market equilibrium prices are unique, and can be characterized as the set specified by the following convex program.*

$$\forall i, j : \frac{L_i + \sum_{j'} u_{ij'} x_{ij'}}{B_i} \geq \frac{u_{ij}}{p_j} \quad (4)$$

$$\forall i : \frac{L_i + \sum_{j'} u_{ij'} x_{ij'}}{B_i} \geq 1 \quad (5)$$

$$\forall j : \sum_i x_{ij} \leq 1 \quad (6)$$

$$\sum_j p_j + \sum_i L_i \leq \sum_i B_i \quad (7)$$

$$\forall i, j : x_{ij} \geq 0 \quad (8)$$

$$\forall j : p_j \geq 0. \quad (9)$$

Now let us return to the proof of Theorem 2. We show that as δ and γ approach zero, the constraints in Theorem 3 becomes satisfied. In fact, constraints (5), (6), (8), and (9) are always satisfied: constraint (5) is satisfied because no advertiser buys any keyword at a higher price than his utility, and the other three constraints are satisfied because Algorithm 1 always computes a feasible allocation and non-negative prices. Therefore, the only constraints that we need to check are (4) and (7). But it is easy to use Theorem 1 to show that these constraints are satisfied approximately, i.e., there is a value $\rho(\delta, \gamma)$ that approaches zero as δ and γ approach zero so that:

$$\forall i, j : \frac{L_i + \sum_{j'} u_{ij'} x_{ij'}}{B_i} \geq (1 - \rho(\delta, \gamma)) \frac{u_{ij}}{p_j}$$

$$\sum_j p_j + \sum_i L_i \leq (1 + \rho(\delta, \gamma)) \sum_i B_i.$$

The prices and allocation of our algorithm must satisfy these constraints. Consider the convex region specified by these relaxed constraints. As δ and γ go to zero, the constraints approach those of Theorem 3, implying that the price and utility vectors converge to the unique equilibrium price and utility vectors, respectively. This completes the proof of Theorem 2.