

Convergence of Position Auctions under Myopic Best-Response Dynamics*

Matthew Cary Aparna Das Benjamin Edelman
Ioannis Giotis Kurtis Heimerl Anna R. Karlin
Scott Duke Kominers Claire Mathieu Michael Schwarz

January 10, 2013

Abstract

We study the dynamics of multi-round position auctions, considering both the case of exogenous click-through rates and the case in which click-through rates are determined by an endogenous consumer search process. In both contexts, we demonstrate that the dynamic auctions converge to their associated static, envy-free equilibria. Furthermore, convergence is efficient, and the entry of low-quality advertisers does not slow convergence. Because our approach predominantly relies on assumptions common in the sponsored search literature, our results suggest that dynamic position auctions converge more generally.

Keywords. Position auctions, dynamic auctions, consumer search, best-response bidding.

1 Introduction

The *position auctions* used to allocate sponsored search links are often modeled as games in which advertisers submit bids and are assigned positions in descending bid order. The utility of each advertiser then depends upon its per-click valuation, the click-through rate of the position it receives, and the profile of bids.

*This work combines and extends previous papers of Cary *et al.* [6] and Kominers [13], which respectively appeared in the *Proceedings of the 8th ACM Conference on Electronic Commerce* and the *Proceedings of the 5th International Conference on Algorithmic Aspects in Information and Management*.

This research was supported in part by NSF grant CCF-0635147, a Yahoo! Research Alliance gift to Karlin, and Harvard College Program for Research in Science and Engineering and AMS-Simons grants to Kominers. Kominers also appreciates the hospitality of the Harvard Program for Evolutionary Dynamics, which hosted him during part of this research. The authors thank Zachary Abel, Susan Athey, Yiling Chen, Peter Coles, Noam D. Elkies, Drew Fudenberg, Andrea J. Hawksley, Kathleen L. McGinn, David Parkes, Michael Ostrovsky, Alvin E. Roth, Sergei Severinov, Dmitry Taubinsky, several referees, the associate editor, and participants in the Harvard EconCS Research Workshop for helpful comments and suggestions.

Early position auction models, such as [2], [9], and [17] assumed positions’ click-through rates to be entirely exogenous. New approaches have introduced the effects of consumer behavior, as in [8] and [3]. However, both approaches typically rely on static modeling frameworks—leaving open the question of whether and how outcomes converge to the static equilibrium.¹

In this paper, we study dynamic position auctions, considering both the case of exogenous click-through rates (using the static model of [9]) and the case in which click-through rates are determined by an endogenous consumer search process (using the static model of [3]). In each case, we demonstrate that the dynamic auctions converge to their associated static, envy-free equilibria. This convergence is efficient. Moreover, the entry of low-quality advertisers does not slow convergence.

Our contributions are threefold: First, we illustrate that the popular static position auction equilibrium of [9] arises under a natural best-response bidding dynamic. We also show that the addition of consumer behavior does not prevent the eventual stability of the dynamic position auction. Finally, we demonstrate the robustness of our approach: similar analysis yields the convergence of position auctions in significantly different settings.

We proceed as follows: In Section 2, we present a basic dynamic position auction framework. In Section 3, we present convergence results for both synchronous and asynchronous bidding in position auctions with fixed click-through rates. In Section 4, we obtain convergence results for the case of endogenous click-through rates arising from consumer search behavior. Finally, in Section 5, we conclude with comments on the generality of our approach.

2 Framework

We consider an auction in which N advertisers bid on $M < N$ sponsored link positions. Each advertiser π has a per-click valuation q_π ; we assume that these valuations are drawn independently from a public, atomless distribution with support on $[0, 1]$. For convenience, we label the advertisers π by $\{1, \dots, N\}$ so that the valuations q_π satisfy $q_1 > \dots > q_N$.

An assignment of advertisers to positions is an injection

$$\mathcal{P} : \{1, \dots, M\} \hookrightarrow \{1, \dots, N\}$$

such that advertiser $\mathcal{P}(j)$ is assigned position j .

We assume a dynamic setting with sequential rounds of play $t = 1, 2, \dots$. In each round $t > 0$, the search engine allocates positions through a generalized second-price auction:² the advertiser submitting the j -th highest bid in round t

¹There are a few notable exceptions. For example, [10] uses a dynamic position auction model to derive equilibrium refinements for the static framework of [9].

²The exact details of the implementation of such an auction without and with consumer search are specified by [9] and [3], respectively. The mechanism is analogous to a second-price ascending bid auction for the M positions.

is assigned position j and is charged a per-click price p_j^t equal to the $(j + 1)$ -st highest bid submitted in round t .³

3 Bidding Dynamics in the Presence of Fixed Click-Through Rates

For now, we assume that each position j has an exogenous *click-through rate* θ_j , interpreted as the probability that a consumer clicks on the advertisement in position j . (In Section 4.1, we instead derive positions’ click-through rates endogenously as a result of a consumer search process.) We assume that higher positions have higher click-through rates, i.e. $\theta_1 > \dots > \theta_M$

When modeled as a static game of complete information, the generalized second-price auction has a continuum of Nash equilibria (see [9, 17, 14]). One of these equilibria results in advertiser payments identical to those which would arise under the VCG mechanism; this equilibrium is also attractive to advertisers because it is the “cheapest” equilibrium that is (*locally*) *envy-free* in the sense that an advertiser assigned position j does not want to exchange positions with the advertiser assigned position $(j - 1)$ at the price the advertiser occupying position $(j - 1)$ is paying. In that equilibrium, bids follow a recursive formula:

$$b_{\pi_j} = \begin{cases} q_{\pi_j} - \frac{\theta_j}{\theta_{j-1}}(q_{\pi_j} - b_{\pi_{j+1}}) & 1 \leq j \leq k, \\ q_{\pi_j} & k < j \leq N, \end{cases} \quad (1)$$

where the j -th highest bid b_{π_j} is submitted by advertiser π_j . There are multiple Nash equilibria in the static game. However, [17] and [9] suggest that the equilibrium given by (1) is the most plausible because its bids are the *highest* advertisers can choose while ensuring that no advertiser will lose money if she accidentally exceeds the bid of the advertiser directly above her and thus moves up in the allocation of positions.⁴ Here and throughout subsequent analysis, we define $\theta_0 := 2\theta_1$, so that the bidding strategy is well-defined in the first position.

3.1 Balanced Bidding

We assume that advertisers π play myopic best-response strategies, submitting bids under the assumption that all other advertisers will repeat their previous bids. So that advertisers’ strategies are well-defined in the first round, we assume a random assignment of advertisers to initial positions, with all initial bids equal to 0. Under this assumption, we define the *restricted balanced bidding* strategy.

³This is a *rank-by-bid* mechanism: payment is proportional to the number of clicks that occur, but advertisers are ranked in strict bid order. In contrast, *rank-by-revenue* mechanisms sort advertisers after their bids are weighted by their ads’ expected click-through rates. We focus on rank-by-bid mechanisms in part because modeling a dynamic rank-by-revenue mechanism requires modeling fluctuations in advertisers’ expected click-through rates [16]. Furthermore, as shown in [3], rank-by-revenue mechanisms render equilibrium bidding behavior unclear in the presence of consumer search, even in the static context.

⁴[9] proposes a model where this equilibrium is a unique outcome.

Definition 1. The *restricted balanced bidding* (RBB) strategy is the strategy for advertiser π which, given the bids of the other advertisers in round t ,

- targets the position s_π^* which maximizes the utility of advertiser π among the positions with no higher expected click-through rate than her current position s_π ,
- chooses the bid b_π^* for round $t + 1$ so as to satisfy the equation

$$\theta_{s_\pi^*}(q_\pi - p_{s_\pi^*}^t) = \theta_{s_\pi^* - 1}(q_\pi - b'). \quad (2)$$

The static bid profile (1) is the unique fixed point of the RBB strategy. Thus, we see that the condition (2) is not ad-hoc—it arises as the local bidding condition in round $t + 1$ of our model when all advertisers play according to an envy-free symmetric strictly monotone equilibrium strategy. Indeed, if advertiser π expects other advertisers to repeat their bids from round t , then she must bid as in (2) if she is to be indifferent between receiving position $s_\pi^* - 1$ at price b_π^* and receiving position s_π^* at price $p_{s_\pi^*}^t$.⁵

The requirement that advertisers only target positions with click-through rates no higher than those of their current positions is less natural. This is a technical condition which is necessary in order to obtain convergence in the synchronous bidding model. As we show in Section 3.3, convergence results for synchronous bidding under RBB imply that convergence obtains in an asynchronous bidding model even when this technical condition is lifted. Thus, this condition does not appear overly confining.

Would an advertiser be well advised to choose the RBB strategy? We simulated a set of advertisers playing a variety of other strategies, and a focus advertiser choosing between playing RBB and matching the other advertisers' strategy. Table 1 presents the percent increases in payoff to the focus advertiser by switching to RBB in an environment with geometric click-through rates and exponential valuations. For every alternative strategy that we examined, the focus advertiser increased his payoff by choosing RBB. We found qualitatively similar results for other distributions of click-through rates (geometric, pareto, uniform) and valuations (exponential, normal, uniform); on the whole the RBB strategy yielded more favorable individual payoffs than alternatives.

⁵Recall that $p_{s_\pi^*}^t$ is the price of position s_π^* in round $t + 1$ if all advertisers other than π repeat their bids from round t .

All other bidders use this bidding strategy x	% payoff increase to focus bidder by switching from x to RBB	
		from x to BB
Jammer (bids at top of largest gap)	234%	234%
Bids at middle of biggest gap	204%	204%
ROI maximizer (seeks maximum return on advertising spend)	38%	39%
Mean bidder (bids midway between BB point and next-highest bid)	1.0%	0.97%
Kind bidder (bids midway between BB point and next-lowest bid)	2.6%	1.6%
Bid at midpoint of best response interval	2.3%	2.6%
Reinforcement learner		
Reinforcing over positions	5.0%	5.4%
Reinforcing over kind/balanced/mean within best response interval	0.92%	0.60%

5

Table 1: Payoff Improvements from Switching to the BB Strategy

Simulation details: Number of bidders and number of positions were held constant at 10. Valuations were drawn from an exponential distribution with mean 1. Bidders updated bids asynchronously in random order. Click-through rates were geometric with rate of decline 0.8. Each point reflects the averaged results of 1000 simulations. Comparison strategies included strategies used by practitioners (e.g., strategies offered by commercial bidding tools, as in [1]) as well as strategies suggested by the academic literature (e.g., reinforcement learning, as in [11] and the references therein).

Strategy details: The *jammer* finds the largest price gap (by subtracting each bid from the bid above it, and finding the gap of maximum of these differences), then bids at the top of this gap. In light of the GSP payment rule, this strategy provides the jammer with the largest possible discount (in cents per click). The *middle of biggest gap* strategy targets the same

position, but sets its bid as the arithmetic mean of the top and bottom of the gap. The *ROI maximizer* seeks to maximize its percentage return on advertising spending, maximizing the ratio of its value to its advertising expenditure. The *mean*, *kind*, and *midpoint of best response interval* bidders all find the best response interval and bid, respectively, halfway between the BB point and the next-highest bid, halfway between BB and next-lowest, and at the midpoint of the best response interval. *Reinforcement learning* (RL) bidders consider previous payoffs when selecting parameters for further bidding. Our first RL bidder-type reinforces over ranked positions, optimizing over the prospect of seeking each ordinal position. Our second RL bidder-type always places a bid in the best response interval and uses reinforcement learning to select one of nine positions within that interval (mean, three equally-spaced positions between mean and BB, BB, three equally-spaced positions between BB and kind, and kind).

3.2 Synchronous Bidding Dynamics

We now consider the behavior of dynamic position auctions under *synchronous bidding*, in which all advertisers see all round- $(t-1)$ bids and—simultaneously—submit bids for round t . This problem corresponds to an instance of the *synchronous, distributed assignment problem* considered by Bertsekas [4]. Applying an analog of Bertsekas’s approach yields the following convergence result.

Theorem 2. *In the synchronous model of dynamic bidding in which each advertiser bids every round, the RBB strategy always converges to its fixed point.*

Proof. We denote $\gamma_j := \frac{\theta_j}{\theta_{j-1}}$ and let $\gamma^{**} := \max_j \gamma_j$.

Lemma 3. *Advertiser π prefers to target position j over position $j-1$ in round $t+1$ if and only if $(1-\gamma_j)q_\pi + \gamma_j p_j^t < p_{j-1}^t$.*

Proof. This follows from the fact that advertiser π prefers to target position j over position $j-1$ if and only if

$$\theta_j(q_\pi - p_j^t) > \theta_{j-1}(q_\pi - p_{j-1}^t). \quad \square$$

Lemma 4. *At every round t such that*

$$t > t_1 := 2 + \log_{\gamma^{**}}((1-\gamma^{**})(q_M - q_{M+1})/q_{M+1}),$$

we have

$$\begin{cases} b_\pi > q_{M+1} & \pi < M+1, \\ b_\pi = q_\pi & \pi \geq M+1, \end{cases}$$

where b_π is the bid of advertiser $1 \leq \pi \leq N$.

Proof. If b is the $(M+1)$ -st highest bid, then $b \leq q_{M+1}$. If $b < q_{M+1}$ in some round t , then in the next round any advertiser $\pi \in \{1, 2, \dots, M+1\}$ will either bid $b'_\pi = q_\pi$ or target some position $j \in \{1, \dots, M\}$ with bid

$$\begin{aligned} b'_\pi &:= (1-\gamma_j)q_\pi + \gamma_j p_j^t \\ &\geq (1-\gamma^{**})q_{M+1} + \gamma^{**}b \\ &\geq b + (1-\gamma^{**})(q_{M+1} - b). \end{aligned}$$

In both of these cases, $q_{M+1} - b'_\pi \leq \gamma^{**}(q_{M+1} - b)$.

It follows that

$$q_{M+1} - b < (1-\gamma^{**})(q_M - q_{M+1})$$

within at most $t \leq \log_{\gamma^{**}}((1-\gamma^{**})(q_M - q_{M+1})/q_{M+1})$ rounds. Then, the bidders $\pi \in \{1, \dots, M\}$ will bid at least

$$\begin{aligned} (1-\gamma_j)q_\pi + \gamma_j p_j^t &\geq (1-\gamma_j)q_\pi + \gamma_j b \\ &\geq b + (1-\gamma_j)(q_\pi - b) \\ &> b + (1-\gamma^{**})(q_M - q_{M+1}) > q_{M+1} \end{aligned}$$

in round $t+1$. In round $t+2$, advertiser $M+1$ will bid q_{M+1} while advertisers $\pi \in \{1, \dots, M\}$ bid above q_{M+1} . \square

Lemma 4 proves that, within finitely many rounds, the set of advertisers competing for the M positions will stabilize and that this set will be the collection of advertisers of maximal quality, $\{1, \dots, M\}$. Furthermore, at this time, the $N - M$ advertisers $\{M + 1, \dots, N\}$ will bid their values in every subsequent round. Thus, we may assume that these rounds have already elapsed; all that remains is to show that the bids for the M actual positions eventually converge to the desired fixed point. Since the fixed point is unique, it suffices to prove convergence.

For any $j \in [0, M]$, we say that the advertisers assigned positions $[j + 1, M]$ are *stable* if their allocation is in order of decreasing quality and their prices satisfy equation (2). If all M positions are stable, then we have reached the fixed point of the RBB strategy.

Suppose that, at some round $t > t_1$, the set $S = [s + 1, M]$ of stable positions is not the full set $[1, M]$. Let P denote the set of advertisers in positions $[1, s]$, and let b denote the minimum bid of these advertisers. Define a partial order \sqsupseteq on stable sets: $S' \sqsupseteq S$ if either $S \subsetneq S'$ or if the advertiser of minimum quality in $(S \cup S') \setminus (S' \cap S)$ belongs to S' .

In round $t + 1$, all advertisers in S repeat their bids. We let the new lowest bid of advertisers in P be b'_π , bid by advertiser π . We must consider three cases:

Case 1: $b'_\pi < p_s^t$. We let j be the position targeted by π . By Lemma 3 and the definition of RBB, we have $p_j^t < (1 - \gamma_j)q_\pi + \gamma_j p_j^t = b'_\pi < p_{j-1}^t$.

We denote by $\pi_j \in S$ the advertiser who assigned position j in round t . By the stability of S , we have $p_{j-1}^t = (1 - \gamma_j)q_{\pi_j} + \gamma_j p_j^t$. Then, we have

$$p_{j-1}^t = (1 - \gamma_j^t(q_{\pi_j}))q_{\pi_j} + \gamma_j^t(q_{\pi_j})p_j^t > (1 - \gamma_j)q_\pi + \gamma_j p_j^t > p_j^t,$$

as advertiser π_j is assigned position j in round t (whence $q_{\pi_j} \geq p_j^t$). Furthermore, as S is stable, we have

$$p_{j-1}^t = (1 - \gamma_j)q_{\pi_j} + \gamma_j p_j^t \tag{3}$$

by the definition of the RBB strategy. As (3) is larger than the bid of π , we find $q_{\pi_j} > q_\pi$. Likewise, we find that $q_{\pi_{j-1}} < q_\pi$. Thus, $S' := \{\pi' \in S : q_{\pi'} < q_\pi\} \cup \{\pi\}$ is stable and $S' \sqsupseteq S$.

Case 2: π targets position s . Then π is allocated position s and $S \cup \{\pi\} \sqsupseteq S$ is stable.

Case 3: π targets some position $j \leq s - 1$. Then, S remains stable and the minimum bid of advertisers in P has increased. We will show that this case may occur only finitely many times between occurrences of Cases 1 and 2.

As in Section 4.1, we respectively denote the qualities of the advertisers in positions $1, \dots, M$ by $q_{\pi_1}, \dots, q_{\pi_M}$. We then let

$$\epsilon := \frac{\theta_M}{\theta_1} (1 - \gamma^{**}) \min_{\pi \neq \pi'} |q_\pi - q_{\pi'}|$$

and let $x := \log_{1/\gamma^{**}}((q_1 - q_{M+1})/\epsilon)$. We will see that at most x instances of Case 3 may occur between instances of Cases 1 and 2.

Lemma 5. *If $p_{s-1}^t > q_\pi - \epsilon$, then advertiser π prefers position s to any position $j < s$.*

Proof. We have

$$\begin{aligned} q_\pi - p_s^t &= (1 - \gamma_{s+1})(q_\pi - q_{\pi_{s+1}}) + \gamma_{s+1}p_{s+1}^t \\ &\geq (1 - \gamma^{**}) \min_{\pi \neq \pi'} |q_\pi - q_{\pi'}|. \end{aligned} \quad (4)$$

The ratio of the expected utility of position $k < s$ to that of position s is less than

$$\begin{aligned} \frac{\theta_k(q_\pi - p_{s-1}^t)}{\theta_s(q_\pi - p_s^t)} &\leq \epsilon \frac{\theta_k}{\theta_s(q_\pi - p_s^t)} \\ &\leq \epsilon \frac{\theta_1}{\theta_s(q_\pi - p_s^t)} \leq 1. \end{aligned} \quad \square$$

Now suppose that Case 3 occurs for x consecutive rounds. We let π be the advertiser in P of minimal quality q_π and denote by $b^{(t')}$ the minimal bid of advertisers in P after t' consecutive rounds of Case 3. If $\pi' \in P$ submits the minimal bid $b^{(t'+1)}$ in the next round, then

$$\begin{aligned} b^{(t'+1)} &\geq (1 - \gamma^{**})q_{\pi'} + \gamma^{**}b^{(t')} \\ &\geq (1 - \gamma^{**})q_\pi + \gamma^{**}b^{(t')} \\ &\geq q_\pi - \gamma^{**}(q_\pi - b^{(t')}). \end{aligned}$$

After x consecutive rounds of Case 3, we have

$$b^{(x)} \geq q_\pi - (\gamma^{**})^x (q_\pi - b^{(0)}).$$

Hence, $b^{(x)} \geq q_\pi - \epsilon$. It follows from Lemma 5 that π will target position s in the next round, so the next round is an instance of Case 2. Thus, we have shown that Case 3 may occur only finitely many times between instances of Cases 1 and 2. □

Theorem 2 shows that under the RBB strategy, advertisers' bids converge to the static bid profile (1). In general, the proof of Theorem 2 yields a rather weak bound on the convergence time: tracing the argument shows convergence within

$$O\left(2^M \left(\frac{\log(1 - \gamma^*)}{\log \gamma^*} + \log_{(1/\gamma^*)} \frac{q_1}{\min_{1 \leq \pi \leq M} (q_\pi - q_{\pi+1})} + \log_{(1/\gamma^*)} \frac{\theta_1}{\theta_M}\right)\right)$$

rounds, where $\gamma^* := \max_j \frac{\theta_j}{\theta_{j-1}}$ determines an advertiser's maximum marginal loss from being lowered one position.

A tighter bound can be obtained in the case when the click-through rates are geometrically decreasing, i.e. when $\theta_j = \delta^{j-1}$, for $0 < \delta < 1$.⁶

⁶According to [12, 14], real click-through rates are roughly geometric with $\delta = 0.7$.

Theorem 6. *In the synchronous model of dynamic bidding with geometrically decreasing click-through rates, $\theta_j = \delta^{j-1}$ (with $0 < \delta < 1$), the RBB strategy converges to bid profile (1) within*

$$O\left(M^3 + \log_\delta\left(\left(1 - \delta\right)\frac{q_M - q_{M+1}}{q_{M+1}}\right)\right)$$

rounds.

In the proof of Theorems 2 and 6, the number of rounds until convergence is constant in N , holding $\max_{1 \leq \pi \leq N} q_\pi$ fixed. That implies the following result:

Corollary 7. *The entry of low-quality advertisers does not slow the auction's convergence.*

3.3 Asynchronous Bidding Dynamics

We now examine convergence of position auctions under *asynchronous bidding* dynamics, in which an advertiser chosen (uniformly) at random updates its bid in each round, according to the following envy-free bidding strategy.

Definition 8. The *balanced bidding* (BB) strategy is the strategy for advertiser π which, given the bids of the other advertisers in round t ,

- targets the position s_π^* which maximizes the utility of advertiser π ,
- chooses the bid b_π^* for round $t + 1$ so as to satisfy (2).

BB performs well against a variety of other strategies. The final column of Table 1 presents payoff increases for BB in an environment with geometric click-through rates and exponential valuations. Similar to our results for RBB, an advertiser switching to BB obtained increased payoffs. Results were qualitatively similar for other distributions of click-through rates (geometric, pareto, uniform) and valuations (exponential, normal, uniform).

Unlike RBB, the BB strategy does not restrict the positions that advertisers may target. However, like RBB, BB has a unique fixed point—the static bid profile (1). Our next result shows that asynchronous bidding under BB converges, and that this convergence is again mediated by the constant γ^* .

Theorem 9. *In the asynchronous model of dynamic bidding in which advertisers bid in a (uniformly) random order and follow the balanced bidding strategy, bids converge to the bid profile (1) with probability 1 and expected convergence time*

$$O\left(t_1(N \log M) + N \log N + M^{2^M(1+x)}\right),$$

where (as in the proof of Theorem 2)

$$t_1 = 2 + \log_{\gamma^*}\left(\frac{(1 - \gamma^*)(q_M - q_{M+1})}{q_{M+1}}\right),$$

$$x = \log_{1/\gamma^*}\left(\frac{q_1 - q_{M+1}}{\frac{\theta_M}{2\theta_1}(1 - \gamma^*) \min_{\pi \neq \pi'} |q_\pi - q_{\pi'}|}\right).$$

This result follows from the asynchronous-case extension of the arguments of [4]. For the full proof, see Appendix B.

In general, the bound in Theorem 9 is doubly exponential in the number of positions, M . However if M is constant (as is the case in practice), then the bound is nearly linear in the number of bidders, N .

We ran simulations to better evaluate convergence speed in practice. Figure 1 presents convergence time as the number of bidders and number of positions increase identically. For bidders drawn from exponential distributions, convergence speed is roughly linear in the number of bidders and positions. Results were similar for bidders drawn from other distributions. A best-fit line (computed using OLS regression) shows the number of bids required per bidder; the fit is tight, and the slope of the line is not statistically significantly different from 1. We also simulated other distributions of click-through rates (geometric, pareto, uniform) and valuations (exponential, normal, uniform) in all combinations. Results were qualitatively similar, and a linear best-fit line remained a good fit of the number of bids required per bidder, although best-fit slope ranged from 1.02 to 2.81 (with standard error in every case below 0.1).

4 Bidding Dynamics in the Presence of Consumer Search

4.1 Consumer Search Process

We now consider an auction in which click-through rates θ_k arise endogenously as the result of a consumer search process.

We interpret advertisers' per-click valuations in terms of advertiser quality: q_π is the probability of meeting an individual consumer's need; an advertiser receives a payoff of 1 every time it meets a consumer's need. With this structure, advertisers' per-click valuations are again given by the values q_π , as the expected per-click revenue of advertiser π is $q_\pi \cdot 1 = q_\pi$.

At all times, a continuum of consumers seek to meet their needs by searching through the sponsored link list. Consumers are assumed to be ignorant of the positions' dynamics, so their beliefs are independent of the dynamics of advertisers' bid updating. (A typical consumer will only use a search engine to seek a given product once, hence no consumer has an opportunity to learn about the positions' dynamics.) Additionally, consumers are assumed to believe that the advertisers' bidding strategies are strictly monotone in their qualities. This is a reasonable assumption for our purposes. Indeed, Proposition 5 of [3] shows that the static position auction game in this framework has a symmetric pure strategy monotone equilibrium; this equilibrium is the unique equilibrium of both the static game and the dynamic game in our model.

Each consumer i must pay a search cost of c_i for each click on a sponsored search link. Search costs c_i are assumed to be distributed according to a public, atomless distribution with support on $[0, 1]$ and CDF G . Since search is costly and consumers believe the links to be sorted in order of descending quality, they

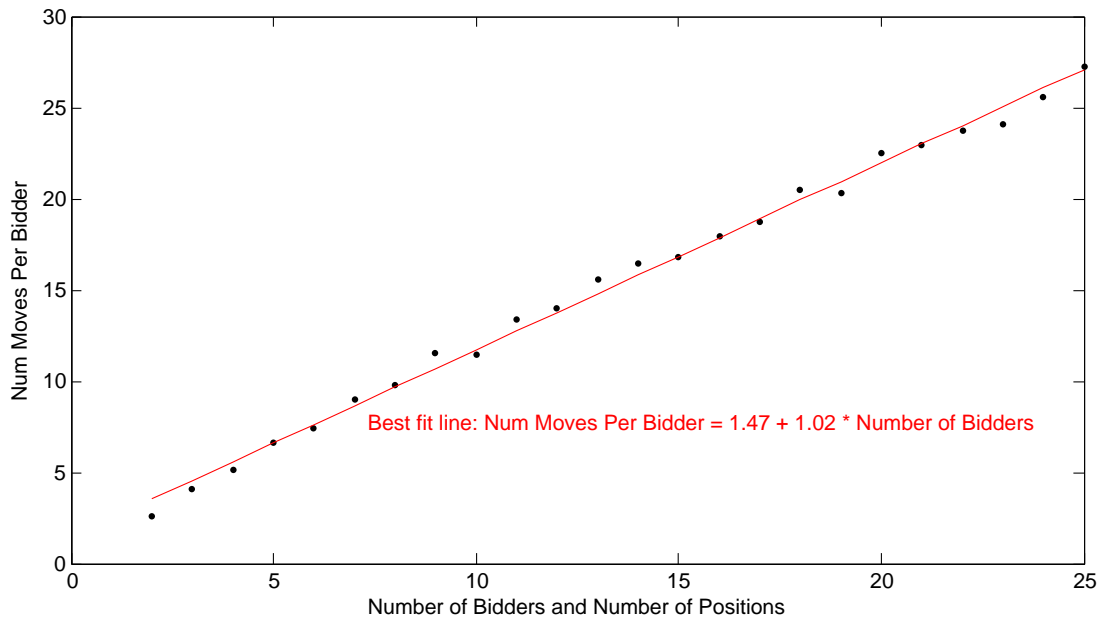


Figure 1: Convergence speed as number of bidders changes.

Simulation details: Number of bidders and number of positions were increased identically as indicated. Bidders updated bids asynchronously in random order. Valuations were drawn from an exponential distribution with mean of 1. All bidders used the BB bidding strategy. Click-through rates were geometric with rate of decline 0.8. Convergence was declared when all bids were within 10^{-6} of the [9] equilibrium bids indicated in (1). Each plotted point reflects the averaged results of 100 simulations.

examine links from top to bottom. Consumers update their predictions about advertisers' qualities in a Bayesian manner: when a website does not meet consumer i 's need, she reduces her estimate of the lower websites' qualities and continues searching if and only if the expected value of clicking on the next link exceeds c_i . Since consumers are unable to learn about the positions' dynamics, we assume that consumers are slightly myopic: an individual consumer will maintain her belief that the advertisers below position j have lower qualities than do the advertisers in positions $1, \dots, j$, even if she discovers while searching that the advertisers are not sorted in order of quality.

Suppose that the advertisers with quality scores $q_{\pi_1}^{\mathcal{P}}, \dots, q_{\pi_M}^{\mathcal{P}}$ are respectively assigned positions $1, \dots, M$ in some assignment \mathcal{P} mapping advertisers to positions. Let $z_{\pi_1}^{\mathcal{P}}, \dots, z_{\pi_M}^{\mathcal{P}}$ be Bernoulli variables taking the value 1 with probabilities $q_{\pi_1}^{\mathcal{P}}, \dots, q_{\pi_M}^{\mathcal{P}}$. Then a consumer whose need has not been met by the advertisers in the first $j \geq 1$ positions expects that the quality of the firm in position $j + 1$ is

$$\bar{q}_{j+1}^{\mathcal{P}} := E(q_{\pi_{j+1}} | z_{\pi_1}^{\mathcal{P}} = \dots = z_{\pi_j}^{\mathcal{P}} = 0).$$

The expected probability of the advertiser in the first position meeting a consumer's need is always $\bar{q}_1^{\mathcal{P}} := E(q_{\pi_1})$. All consumers can compute this value, as the distribution of advertiser quality is assumed to be public. From these definitions, it is apparent that $\bar{q}_j^{\mathcal{P}} > \bar{q}_{j+1}^{\mathcal{P}}$ for any $1 \leq j \leq M$. Then the advertiser assigned to position j will receive $(1 - q_{\pi_1}^{\mathcal{P}}) \cdots (1 - q_{\pi_{j-1}}^{\mathcal{P}}) \cdot G(\bar{q}_j^{\mathcal{P}})$ clicks. (See Proposition 2 of [3].)

For convenience, we denote

$$\bar{q}_j^t := \bar{q}_j^{\mathcal{P}^t}, \quad \bar{q}_j := \bar{q}_j^{\mathcal{P}^*},$$

where \mathcal{P}^t is the assignment of positions in round t and \mathcal{P}^* is the assignment of advertisers to positions in order of descending valuation. With our assumption that $q_1 \geq \dots \geq q_N$, it follows that $q_{\pi_j}^{\mathcal{P}^*} = q_j$. Then, by construction, $\bar{q}_j \geq \bar{q}_j^{\mathcal{P}}$ for any assignment \mathcal{P} and $1 \leq j \leq M$.

4.2 Balanced Bidding

We again assume that each advertiser π plays a myopic best-response strategy in each round, submitting bids under the assumption that other advertisers repeat their bids. As before, we assume a random assignment of advertisers to positions at the beginning of the game (round 0), with all initial bids set to 0. Under this assumption, we define a *restricted balanced bidding* strategy, which is analogous to that of Definition 1.

Definition 10. The *restricted balanced bidding* (RBB) strategy is the strategy for advertiser π which, given the bids of the other advertisers in round t ,

- targets the position s_{π}^* which maximizes the utility of advertiser π among the positions with no higher expected click-through rate than the advertiser's current position s_{π} ,

- chooses the bid b_π^* for round $t + 1$ so as to satisfy the equation

$$(1 - q_\pi)G(\bar{q}_{s_\pi^*}^t)(q_\pi - p_{s_\pi^*}^t) = G(\bar{q}_{s_\pi^* - 1}^t)(q_\pi - b_\pi^*). \quad (5)$$

(As above, we define $G(\bar{q}_0^t) := 2G(\bar{q}_1^t)$ so that the strategy is well-defined in the first position.)

Instead, we could require that each advertiser π update its bid with attention to long-run equilibrium click-through rates, choosing b_π^* to satisfy

$$(1 - q_\pi)G(\bar{q}_{s_\pi^*})(q_\pi - p_{s_\pi^*}^t) = G(\bar{q}_{s_\pi^* - 1})(q_\pi - b_\pi^*). \quad (6)$$

With this bidding assumption, the proof of Theorem 11 applies directly. In this case, the assumption that consumers always search from the top down is unnecessary, as this search behavior arises endogenously.

As in the model without consumer search, condition (5) arises as the local bidding condition in round $t + 1$ when all advertisers play according to an envy-free, symmetric, strictly monotone equilibrium strategy. Furthermore, just as in the model without consumer search, convergence results for synchronous bidding under RBB imply convergence in an asynchronous bidding model in which advertisers are not required to target positions with expected click-through rates (weakly) lower than those of their current positions.

The key difference between Definition 10 and Definition 1 is (5), providing for round-dependent click-through rates as well as adding a factor of $1 - q_\pi$ which arises from the consumer search process.⁷

4.3 Synchronous Bidding Dynamics

At the unique fixed point of the RBB strategy, the advertisers choose their bids according to the recursive strategy presented in Proposition 5 of [3]. (This is a direct consequence of the proof of Proposition 5 of [3], which shows that these bids satisfy the condition (5).) That is, the bids at the fixed point are given by

$$b_{\pi_j} = \begin{cases} q_{\pi_j} - \frac{G(\bar{q}_j)}{G(\bar{q}_{j-1})}(1 - q_{\pi_j})(q_{\pi_j} - b_{\pi_{j+1}}) & 1 < j \leq M, \\ q_{\pi_j} & M < j \leq N, \end{cases} \quad (7)$$

where the j -th highest bid b_{π_j} is submitted by advertiser π_j . Our main result is the convergence of the dynamic position auction to this fixed point.

Theorem 11. *In the synchronous model of dynamic bidding in which click-through rates are determined by consumer search and each advertiser bids every round, the RBB strategy always converges to its fixed point within finitely many rounds.*

⁷Since consumers search from the top down, advertiser π expects to lose a fraction of clicks equal to $1 - q_\pi$ when switching from position $s_\pi^* - 1$ to position s_π^* , as $1 - q_\pi$ is the expected fraction of consumers meeting their needs at position $s_\pi^* - 1$. Including consumer search also changes the computation of the position s_π^* , but this does not materially affect our arguments.

The proof of Theorem 11, which we present in Appendix C, is analogous to the proof of Theorem 2. However, the addition of consumer search to the model produces variation in click-through rates across rounds. This increases the technical difficulty of the argument; in particular, it renders the results of [4] inapplicable.

The additional complexity introduced by consumer search weakens the bound on convergence time. However, even in the presence of consumer search, the entry of low-quality advertisers does not affect the speed of convergence.

4.4 Asynchronous Bidding Dynamics

As in the case without consumer search, it is possible to extend Theorem 11 to obtain a convergence result for an asynchronous auction in a less restricted strategy space.

Definition 12. The *balanced bidding* (BB) strategy is the strategy for advertiser π which, given the bids of the other advertisers in round t ,

- targets the position s_π^* which maximizes the utility of advertiser π ,
- chooses the bid b_π^* for round $t + 1$ so as to satisfy the equation

$$(1 - q_\pi)G(\bar{q}_{s_\pi^*}^t)(q_\pi - p_{s_\pi^*}^t) = G(\bar{q}_{s_\pi^* - 1}^t)(q_\pi - b_\pi^*). \quad (8)$$

(As in RBB, we define $G(\bar{q}_0^t) := 2G(\bar{q}_1^t)$, so that this strategy is well-defined in the first position.)

The BB strategy is a natural envy-free bidding strategy, analogous to Definition 8. Unlike RBB, each advertiser π playing BB may target any position. However, as in RBB, the bid condition (8) arises as the advertisers' envy-free condition in a symmetric, strictly monotone equilibrium. The bid profile (7) is the unique fixed point of the BB strategy. (This follows directly from Proposition 5 of [3].)

Our next result, which follows from the proof of Theorem 11, shows that the dynamic position auction converges under BB in an *asynchronous bidding model* in which advertisers update their bids asynchronously, bidding in a uniformly random order.

Theorem 13. *Consider the asynchronous model of dynamic bidding in which click-through rates are determined by consumer search, and in which advertisers*

- *bid in a (uniformly) random order and*
- *follow the balanced bidding strategy.*

Then bids converge to the bid profile (7) with probability 1 and expected convergence time

$$O\left(t_1(N \log M) + N \log N + M^{2^M(1+x)}\right).$$

The proof follows directly from the analysis in the proof of Theorem 9. As before, this result continues to hold if each advertiser π updates with attention to long-run equilibrium click-through rates, choosing bids b_π^* to satisfy (6) instead of (8).

5 Discussion and Conclusion

Standard models of position auctions largely ignore the question of how bids converge to equilibrium. Some models assume that advertisers already know each others' values. (For example, players could learn others' values through previous interactions that are outside the model.) Other models offer an approach to information revelation that is intended to capture the essence of actual practice despite important differences from reality. (For example, the generalized English Auction of [9] is a metaphor for GSP but not a realistic description of the true sponsored search bidding game.) In contrast, the preceding results establish that if players follow a simple bidding rule in the spirit of myopic best response, bids eventually converge to the equilibrium proposed in [9] and [17]. Notably, the convergence process we present arises in strikingly different environments, including the fixed click-through rate framework of Edelman *et al.* [9] and the consumer search model of Athey and Ellison [3].

To identify the conditions supporting our convergence results, we review the key steps of our proofs of convergence:

1. We restrict the strategy space to a class of plausible bidding behaviors (as in Definition 10),
2. We establish that the advertisers with the lowest valuations must eventually bid their values (as in Lemma 17), and
3. We demonstrate that the advertisers who win positions in equilibrium eventually sort monotonically (as in Cases 1–3 within the proof of Theorem 11).

These steps are consistent with standard assumptions in the sponsored search literature. For example, some restriction of the strategy space is required to ensure that the bidding equilibrium is unique. Such a restriction is usually included even in static models (e.g., envy-freeness in [9] and [3]). Step 1 combines this selection restriction with the requirement that bidders adjust their bids to maximize individual short-term profits. Since it is always weakly dominant for advertisers to bid their values, the advertisers who actually receive positions in any position auction should be those with the highest valuations. Hence, Step 2 seems particularly natural. Finally, Step 3 uses monotonicity of the equilibrium strategy and best-response dynamics to establish convergence of the auction among the advertisers with the largest valuations.⁸ Despite these modest assumptions about bidder behavior, we are able to demonstrate

⁸Without monotonicity, “bid cycling” may occur. See [7] and [5].

convergence—suggesting that the “well-behaved” properties of sponsored search include convergence under reasonable assumptions.

There are clearly limitations to our study: First, the assumption that the advertisers are able to implement myopic best-responses may not always be realistic, as it requires knowledge of both current prices and click-through rates. However, both of these data can be learned as long as they are not changing too rapidly. Second, it is not *a priori* clear whether advertisers should update myopically. However, a recent paper of Nisan *et al.* [15] partially assuages this concern, by showing that for GSP auctions, an individual advertiser’s best course of action is to best respond whenever all other advertisers best-respond repeatedly.⁹ Finally, our work is not immune to the standard criticisms about the practical relevance of position auction models, such as the fact that they do not incorporate budgets.

References

- [1] Atlas search - campaign tools. Available at <http://advertising.microsoft.com/atlas-search-bid-management>. Historic version, listing specific bidding strategies, is preserved by Archive.org at <http://web.archive.org/web/20070312022600/http://www.atlasonepoint.com/products/bidmanager/rulesbased?>.
- [2] AGGARWAL, G., GOEL, A., AND MOTWANI, R. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce (2006)*, pp. 1–7.
- [3] ATHEY, S., AND ELLISON, G. Position auctions with consumer search. *Quarterly Journal of Economics* 126 (2011), 1213–1270.
- [4] BERTSEKAS, D. P. The auction algorithm: A distributed relaxation method for the assignment problem. *Annals of Operations Research* 14 (1988), 105–123.
- [5] BU, T.-M., DENG, X., AND QI, Q. Dynamics of strategic manipulation in ad-words auction. In *3rd Workshop on Sponsored Search Auctions (SSA2007), in conjunction with WWW2007 (2007)*.
- [6] CARY, M., DAS, A., EDELMAN, B., GIOTIS, I., HEIMERL, K., KARLIN, A. R., MATHIEU, C., AND SCHWARZ, M. Greedy bidding strategies for keyword auctions. In *Proceedings of the 8th ACM Conference on Electronic Commerce (2007)*, pp. 262–271.
- [7] CARY, M., DAS, A., EDELMAN, B., GIOTIS, I., HEIMERL, K., KARLIN, A. R., MATHIEU, C., AND SCHWARZ, M. On best-response bidding in GSP auctions. NBER Working Paper W13788, 2008.

⁹This provides a game-theoretic justification to our model in the case that click-through rates are determined exogenously.

- [8] CHEN, Y., AND HE, C. Paid placement: Advertising and search on the internet. *Economic Journal* 121 (2011), F309–F328.
- [9] EDELMAN, B., OSTROVSKY, M., AND SCHWARZ, M. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review* 97 (2007), 242–259.
- [10] EDELMAN, B., AND SCHWARZ, M. Optimal auction design and equilibrium selection in sponsored search auctions. *American Economic Review* 100 (2010), 597–602.
- [11] EREV, I., AND ROTH, A. E. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review* 88 (1998), 848–881.
- [12] FENG, J., BHARGAVA, H. K., AND PENNOCK, D. M. Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *INFORMS Journal on Computing* 19 (2007), 137–148.
- [13] KOMINERS, S. D. Dynamic position auctions with consumer search. In *Proceedings of the 5th International Conference on Algorithmic Aspects in Information and Management* (2009), vol. 5564 of *Lecture Notes in Computer Science*, Springer, pp. 240–250.
- [14] LAHAIE, S. An analysis of alternative slot auction designs for sponsored search. In *Proceedings of the 7th ACM Conference on Electronic Commerce* (2006), pp. 218–227.
- [15] NISAN, N., SCHAPIRA, M., VALIANT, G., AND ZOHAR, A. Best-response auctions. In *Proceedings of the 12th ACM Conference on Electronic commerce* (2011), pp. 351–360.
- [16] PIN, F., AND KEY, P. Stochastic variability in sponsored search auctions: Observations and models. In *Proceedings of the 12th ACM Conference on Electronic commerce* (2011), pp. 61–70.
- [17] VARIAN, H. R. Position auctions. *International Journal of Industrial Organization* 25 (2006), 1163–1178.

A Proof of Theorem 6

We begin the analysis after the M highest-value advertisers are all bidding above q_{M+1} and the other $N - M$ advertisers are all bidding their values. By Lemma 4, it will take at most $\log_\delta((1 - \delta)(q_M - q_{M+1})/q_{M+1})$ rounds to reach this point. Then we bound the number of rounds it takes for the M advertisers with values q_1, \dots, q_M to obtain positions $1, \dots, M$ respectively, with bids as in (1).

For a pair of advertisers π and π' , we say that an *inversion* occurs in round $t + 1$ if π is ranked above π' in round t , and then ranked below π' in some future

round. This is a *bad inversion* if $q_\pi > q_{\pi'}$; otherwise, it is a *good inversion*. First we prove that bad inversions do not occur under RBB when the click-through rates are geometrically decreasing; hence, the correct ordering is reached after at most $O(M^2)$ inversions. Second we prove there can be at most $M - 1$ rounds between each good inversion. These two facts together imply that convergence of the top M positions in $O(M^3)$ rounds.

Lemma 14. *Let π and π' be advertisers such that π has a higher value than π' and suppose π holds a position above π' in round t . Then in all future rounds $t' > t$, π holds a position above π' .*

Proof. Let j be the position occupied by π' in round t and let j' be the position that π' targets in round $t + 1$. Let k and k' denote the corresponding positions for advertiser π .

First, we observe that $k' \leq j'$: Since π' targets j' , π' has higher utility at position j' than at any position s such that $j \leq s \leq j'$. Thus, we have

$$\theta_{j'}(q_{\pi'} - p_{j'}^t) \geq \theta_s(q_{\pi'} - p_s^t). \quad (9)$$

Rearranging (9), we find

$$q_{\pi'} \geq (\theta_{j'} p_{j'}^t - \theta_s p_s^t) / (\theta_{j'} - \theta_s). \quad (10)$$

Since $q_\pi \geq q_{\pi'}$, (10) implies that

$$q_\pi > (\theta_{j'} p_{j'}^t - \theta_s p_s^t) / (\theta_{j'} - \theta_s). \quad (11)$$

Hence in round $t + 1$, π targets either position j' or a position above j . In either case $k' \leq j'$.

Given that $k' \leq j'$, π will bid $b_\pi = (1 - \delta)q_\pi + \delta p_{k'}^t$ in round $t + 1$, and π' will bid $b_{\pi'} = (1 - \delta)q_{\pi'} + \delta p_{j'}^t$.¹⁰ This implies the desired result. \square

Lemma 14 implies that bad inversions do not occur under RBB when click-through rates are geometrically decreasing; hence, every inversion is a good inversion. We show next that there can be at most $M - 1$ rounds between good inversions.

Now, we let $S = [s + 1, M]$ denote the maximum stable set in round t (as defined in the proof of Theorem 2), and recall that any advertiser assigned a position in S repeats his bid in round $t + 1$.

Let π_1, \dots, π_M represent the current assignment of advertisers to positions $1, \dots, M$. Our next lemma shows that in every round where there are no inversions, an additional position is added to the stable set. As we have reached the RBB fixed point once all M positions are stable, this implies that there can be at most $M - 1$ rounds between two good inversions.

Lemma 15. *If there is no inversion in round $t + 1$, then π_s targets position s and bids $b_{\pi_s} = \delta b_{\pi_{s+1}} + (1 - \delta)q_{\pi_s}$ (as in (2)) in round $t + 1$.*

¹⁰Recall that as click-through rates are geometrically decreasing, we have $\gamma_s = \delta$ for all positions s .

Proof. Under RBB advertiser π_s can only target position s or a lower position $j > s$. If π_s targets and wins position s in round $t + 1$, then by RBB he must bid as in (2), as desired.

Now, suppose π_s targets a position $j > s$. Then its bid b_{π_s} must be less than the current bid b of advertiser π_j . Since $j \in S$, π_j submits bid $b_j = b$ in round $t + 1$; hence, $b_{\pi_s} < b_{\pi_j} = b$, and an inversion (between π_s and π_j) occurs in round $t + 1$.

Finally, suppose π_s targets but does not win position s . Since π_s is assigned position s in round t , this implies there is an inversion (between π_s and the advertiser who wins position s) in round $t + 1$. \square

By Lemma 15, we see that if there is no inversion in round t , then $[s, M]$ is stable in round $t + 1$. Combining this fact with our prior observations completes the proof.

B Proof of Theorem 9

At round t , a player π is *activated* if π is the advertiser who updates his bid (while others repeat their previous bids). To prove convergence in the asynchronous case, we revisit each step of the proof of Theorem 2 (showing convergence in the synchronous case). We analyze the expected number of asynchronous activations required to achieve each round of the synchronous proof, given that the active advertiser is selected with uniform probability in each round.

Step 1. We must first assure that $b_\pi > q_{M+1}$ for $\pi \leq M$, and $b_\pi = q_\pi$ for $\pi \geq M + 1$. Lemma 3 showed this round could be obtained in t_1 rounds in the synchronous setting. To use Lemma 3 in the asynchronous case, we must activate each of the top $k + 1$ advertisers once per t_1 rounds. Because the active advertiser is picked uniformly from the N advertisers, in expectation $O(N \log(M))$ activations are required per round to ensure that all M highest-value advertisers update their bids. Thus, after an average of $O(t_1 N \log(M))$ activations, the top $M + 1$ advertisers will all be bidding as required. Then each losing advertiser must be activated once so that each losing advertiser bids its value. To ensure that all the $N - (M + 1)$ losing advertisers are activated at least once requires an additional $O(N \log(N - M))$ activations in expectation. Thus in total the expected number of activations to complete Step 1 in the asynchronous setting is $O(t_1 N (\log(M) + \log(N)))$.

Step 2. Continuing the argument of Theorem 2, we must assure that the M highest-value advertisers are sorted correctly. Progress entails increasingly better stable sets (where “better” is defined by the lexicographic ordering \square of Theorem 2).

After Step 1, no losing advertiser can afford a slot. Thus, if a losing advertiser is activated, he bids his value, which must be less than b_M . Thus the bids from losing advertisers will not interfere with the convergence of the M highest-value

advertisers' bids. However, losing advertisers can still consume activations. We therefore later add a budget of rounds to cover the expected number of losing advertisers' activations.

We now define a sequence \mathcal{T} of activations of the top M bidders. The sequence \mathcal{T} will be partitioned into phases, corresponding to stable sets (defined in the proof of Theorem 2). Let S be the current stable set, let $[j + 1, M]$ be the slots occupied by the advertisers of S , and let P be the set of advertisers occupying slots $[1, j]$. Let π_{\min} be the advertiser in P whose value is minimal. Consider the three cases enumerated in the proof of Theorem 2. The occurrence of Case 1 or Case 2 always results in a new stable set. For each stable set, Case 3 can occur at most x rounds, after which a bid from π_{\min} is guaranteed to result in a new stable set.

To achieve these results in the asynchronous case, we repeatedly activate the advertiser in slot j until either Case 1 or Case 2 occurs. If neither case occurs after x activations, then we activate advertiser π_{\min} . At this point a new stable set is formed, the current phase ends, and we move on to the sequence of activations required by the next stable set. This completes the definition of sequence \mathcal{T} . As there are 2^M stable sets and each stable set requires at most $x + 1$ activations, the length \mathcal{T} is bounded by $t_2 := 2^M(x + 1)$.

How long will it take to activate advertisers in the order given by sequence \mathcal{T} , if losers are excluded? With M advertisers, the probability of activating the correct advertiser is at least $1/M$ per round. Thus the correct sequence of activations will occur after M^{t_2} rounds, in expectation.

Finally, we add back rounds to account for losing advertisers. We expect $(1 - \frac{M}{N})M^{t_2}$ losing advertiser activations to occur during Step 2. The second stage therefore requires $M^{t_2} + (1 - \frac{M}{N})M^{t_2} = O(M^{t_2})$ total activations in expectation.

Summation. Adding the expected rounds required in Step 1 to those from Step 2, we see that convergence is reached in $O(t_1(N \log M) + N \log N + M^{t_2})$ rounds, in expectation. It follows that convergence occurs within finitely many rounds with probability 1.

C Proof of Theorem 11

We denote $\gamma_j^t(q) := (1 - q) \frac{G(\bar{q}_j^t)}{G(\bar{q}_{j-1}^t)}$ and let

$$\gamma^*(q) := (1 - q) \max_{\mathcal{P}} \left[\max_{j > 0} \left(\frac{G(\bar{q}_j^{\mathcal{P}})}{G(\bar{q}_{j-1}^{\mathcal{P}})} \right) \right], \quad \gamma^{**} := \max_{1 \leq \pi \leq N} \gamma^*(q_\pi).$$

We observe that, by construction,

$$\gamma_j^t(q_\pi) \leq \gamma^*(q_\pi) \leq \gamma^{**} < 1$$

for any $t > 0$, $1 \leq j \leq M$, and $1 \leq \pi \leq N$. The last of these inequalities follows from the fact that $\bar{q}_j^{\mathcal{P}} > \bar{q}_{j+1}^{\mathcal{P}}$ for any \mathcal{P} and $1 \leq j \leq M$, since then $\frac{G(\bar{q}_j^{\mathcal{P}})}{G(\bar{q}_{j-1}^{\mathcal{P}})} < 1$.

We begin with two lemmata.

Lemma 16. *Advertiser π prefers to target position j over position $j - 1$ in round $t + 1$ if and only if $(1 - \gamma_j^t(q_\pi))q_\pi + \gamma_j^t(q_\pi)p_j^t < p_{j-1}^t$.*

Proof. Upon algebraic manipulation, this follows from the fact that advertiser π prefers to target position j over position $j - 1$ if and only if

$$(1 - q_\pi)G(\bar{q}_j^t)(q_\pi - p_j^t) > G(\bar{q}_{j-1}^t)(q_\pi - p_{j-1}^t). \quad \square$$

Lemma 17. *At every round t such that*

$$t > t_1 := 2 + \log_{\gamma^{**}}((1 - \gamma^{**})(q_M - q_{M+1})/q_{M+1}),$$

we have

$$\begin{cases} b_\pi > q_{M+1} & \pi < M + 1, \\ b_\pi = q_\pi & \pi \geq M + 1, \end{cases}$$

where b_π is the bid of advertiser $1 \leq \pi \leq N$.

Proof. If b is the $(M + 1)$ -st highest bid, then $b \leq q_{M+1}$. If $b < q_{M+1}$ in some round t , then in the next round any advertiser $\pi \in \{1, 2, \dots, M + 1\}$ will either bid $b'_\pi = q_\pi$ or target some position $j \in \{1, \dots, M\}$ with bid

$$\begin{aligned} b'_\pi &:= (1 - \gamma_j^t(q_\pi))q_\pi + \gamma_j^t(q_\pi)p_j^t \\ &\geq (1 - \gamma_j^t(q_\pi))q_{M+1} + \gamma_j^t(q_\pi)b \\ &= b + (1 - \gamma_j^t(q_\pi))(q_{M+1} - b) \\ &\geq b + (1 - \gamma^{**})(q_{M+1} - b). \end{aligned}$$

In both of these cases, $q_{M+1} - b'_\pi \leq \gamma^{**}(q_{M+1} - b)$.

It follows that

$$q_{M+1} - b < (1 - \gamma^{**})(q_M - q_{M+1})$$

within at most $t \leq \log_{\gamma^{**}}((1 - \gamma^{**})(q_M - q_{M+1})/q_{M+1})$ rounds. Then, the bidders $\pi \in \{1, \dots, M\}$ will bid at least

$$\begin{aligned} (1 - \gamma_j^t(q_\pi))q_\pi + \gamma_j^t(q_\pi)p_j^t &\geq (1 - \gamma_j^t(q_\pi))q_\pi + \gamma_j^t(q_\pi)b \\ &\geq b + (1 - \gamma_j^t(q_\pi))(q_\pi - b) \\ &> b + (1 - \gamma^{**})(q_M - q_{M+1}) > q_{M+1} \end{aligned}$$

in round $t + 1$. In round $t + 2$, advertiser $M + 1$ will bid q_{M+1} while advertisers $\pi \in \{1, \dots, M\}$ bid above q_{M+1} . \square

Lemma 17 proves that, within finitely many rounds, the set of advertisers competing for the M positions will stabilize and that this set will be the collection of advertisers of maximal quality, $\{1, \dots, M\}$. Furthermore, at this time, the $N - M$ advertisers $\{M + 1, \dots, N\}$ will bid their values in every subsequent

round. Thus, we may assume that these rounds have already elapsed; all that remains is to show that the bids for the M actual positions eventually converge to the desired fixed point. Since the fixed point is unique, it suffices to prove convergence.

For any $j \in [0, M]$, we say that the advertisers assigned positions $[j + 1, M]$ are *stable* if their allocation is in order of decreasing quality and their prices satisfy equation (5). If all M positions are stable, then we have reached the fixed point of the RBB strategy.

Suppose that, at some round $t > t_1$, the set $S = [s + 1, M]$ of stable positions is not the full set $[1, M]$. Let P denote the set of advertisers in positions $[1, s]$, and let b denote the minimum bid of these advertisers. Define a partial order \sqsupseteq on stable sets: $S' \sqsupseteq S$ if either $S \subsetneq S'$ or if the advertiser of minimum quality in $(S \cup S') \setminus (S' \cap S)$ belongs to S' .

In round $t + 1$, all advertisers in S repeat their bids. We let the new lowest bid of advertisers in P be b'_π , bid by advertiser π . We must consider three cases:

Case 1: $b'_\pi < p_s^t$. We let j be the position targeted by π . By Lemma 16 and the definition of RBB, we have $p_j^t < (1 - \gamma_j^t(q_\pi))q_\pi + \gamma_j^t(q_\pi)p_j^t = b'_\pi < p_{j-1}^t$.

We denote by $\pi_j \in S$ the advertiser who assigned position j in round t . By the stability of S , we have $p_{j-1}^t = (1 - \gamma_j^t(q_{\pi_j}))q_{\pi_j} + \gamma_j^t(q_{\pi_j})p_j^t$. Then, we have

$$p_{j-1}^t = (1 - \gamma_j^t(q_{\pi_j}))q_{\pi_j} + \gamma_j^t(q_{\pi_j})p_j^t > (1 - \gamma_j^t(q_\pi))q_\pi + \gamma_j^t(q_\pi)p_j^t,$$

from which it follows that

$$(q_{\pi_j} - q_\pi) \left(1 + \frac{G(\bar{q}_j^t)}{G(\bar{q}_{j-1}^t)} ((q_{\pi_j} - p_j^t) + q_\pi - 1) \right) > 0. \quad (12)$$

Since advertiser π_j is assigned position j in round t , we know that $q_{\pi_j} \geq p_j^t$. Furthermore, $0 < \frac{G(\bar{q}_j^t)}{G(\bar{q}_{j-1}^t)} \leq 1$, so

$$\frac{G(\bar{q}_j^t)}{G(\bar{q}_{j-1}^t)} ((q_{\pi_j} - p_j^t) + q_\pi - 1) > -1.$$

It follows that (12) holds if and only if $q_{\pi_j} > q_\pi$. Likewise, we find that $q_{\pi_{j-1}} < q_\pi$. Thus, $S' := \{\pi' \in S : q_{\pi'} < q_\pi\} \cup \{\pi\}$ is stable and $S' \sqsupseteq S$.

Case 2: π targets position s . Then π is allocated position s and $S \cup \{\pi\} \sqsupseteq S$ is stable.

Case 3: π targets some position $j \leq s - 1$. Then, S remains stable and the minimum bid of advertisers in P has increased. We will show that this case may occur only finitely many times between occurrences of Cases 1 and 2.

As in Section 4.1, we respectively denote the qualities of the advertisers in positions $1, \dots, M$ by $q_{\pi_1}, \dots, q_{\pi_M}$. We then let

$$\epsilon := \frac{G(\bar{q}_M)}{2G(\bar{q}_1)}(1 - \gamma^{**}) \min_{\pi \neq \pi'} |q_\pi - q_{\pi'}| \left(\prod_{j=1}^M (1 - q_j) \right)$$

and let $x := \log_{1/\gamma^{**}}((q_1 - q_{M+1})/\epsilon)$. We will see that at most x instances of Case 3 may occur between instances of Cases 1 and 2.

Lemma 18. *If $p_{s-1}^t > q_\pi - \epsilon$, then advertiser π prefers position s to any position $j < s$.*

Proof. We have

$$\begin{aligned} q_\pi - p_s^t &= (1 - \gamma_{s+1}^t(q_{\pi_{s+1}}))(q_\pi - q_{\pi_{s+1}}) + \gamma_{s+1}^t(q_{\pi_{s+1}})p_{s+1}^t \\ &\geq (1 - \gamma^{**}) \min_{\pi \neq \pi'} |q_\pi - q_{\pi'}|. \end{aligned} \quad (13)$$

The ratio of the expected utility of position $k < s$ to that of position s is less than

$$\begin{aligned} \frac{G(\bar{q}_k^t)(q_\pi - p_{s-1}^t)}{\left(\prod_{j=k}^s (1 - q_{\pi_j}) \right) G(\bar{q}_s^t)(q_\pi - p_s^t)} &\leq \epsilon \frac{G(\bar{q}_k^t)}{\left(\prod_{j=k}^s (1 - q_{\pi_j}) \right) G(\bar{q}_s^t)(q_\pi - p_s^t)} \\ &\leq \epsilon \frac{G(\bar{q}_1)}{\left(\prod_{j=k}^s (1 - q_{\pi_j}) \right) G(\bar{q}_M^t)(q_\pi - p_s^t)} \leq 1, \end{aligned}$$

where the last inequality follows from (13), the fact that $G(\bar{q}_s^t) \geq G(\bar{q}_M)$ (since $t > t_1$) and from the definition of ϵ . \square

Now suppose that Case 3 occurs for x consecutive rounds. We let π be the advertiser in P of minimal quality q_π and denote by $b^{(t')}$ the minimal bid of advertisers in P after t' consecutive rounds of Case 3. If $\pi' \in P$ submits the minimal bid $b^{(t'+1)}$ in the next round, then

$$\begin{aligned} b^{(t'+1)} &\geq (1 - \gamma^*(q_{\pi'}))q_{\pi'} + \gamma^*(q_{\pi'})b^{(t')} \\ &\geq (1 - \gamma^*(q_{\pi'}))q_\pi + \gamma^*(q_{\pi'})b^{(t')} \\ &= q_\pi - \gamma^*(q_{\pi'})(q_\pi - b^{(t')}) \\ &\geq q_\pi - \gamma^{**}(q_\pi - b^{(t')}). \end{aligned}$$

After x consecutive rounds of Case 3, we have

$$b^{(x)} \geq q_\pi - (\gamma^{**})^x (q_\pi - b^{(0)}).$$

Hence, $b^{(x)} \geq q_\pi - \epsilon$. It follows from Lemma 18 that π will target position s in the next round, so the next round is an instance of Case 2. Thus, we have shown that Case 3 may occur only finitely many times between instances of Cases 1 and 2.