Discrete Choice Models of Bidder Behavior in Sponsored Search

Quang Duong[1] and Sébastien Lahaie[2]

[1] University of Michigan
Ann Arbor, MI 48109
qduong@umich.edu

[2] Yahoo! Research
New York, NY 10018
lahaies@yahoo-inc.com

Abstract. There are two kinds of bidders in sponsored search: most keep their bids static for long periods of time, but some do actively manage their bids. In this work we develop a model of bidder behavior in sponsored search that applies to both active and inactive bidders. Our observations on real keyword auction data show that advertisers see substantial variation in rank, even if their bids are static. This motivates a discrete choice approach that bypasses bids and directly models an advertiser's (perhaps passive) choice of rank. Our model's value per click estimates are consistent with basic theory which states that bids should not exceed values, even though bids are not directly used to fit the model. An empirical evaluation confirms that our model performs well in terms of predicting realized ranks and clicks.

1 Introduction

A major portion of the revenue of search engines such as Google and Bing comes from advertising next to search results. Advertisers bid for placement on keywords relevant to their business, a practice known as sponsored search. A central problem in empirical modeling of sponsored search is to infer bidder values from their observed bidding behavior. Information on bidder values can inform virtually all aspects of the keyword auction design, including changes to the ranking rule to improve revenue [7]; reserve pricing policies [9]; and the impact of improved click-through rate models on efficiency [6]. With bidder values at hand, counterfactual experiments can be performed to evaluate the effect of changes to auction parameters before live testing, and even to compare the current design to more classical auctions such as VCG [1].

In this work we develop a model of bidder behavior in sponsored search that applies to both active and inactive bidders. We observe that most advertisers in Yahoo's sponsored search market keep their bids essentially constant for long periods of time (e.g., several weeks), as others have noticed independently in Bing data [10]. On the other hand, some advertisers on competitive, high-volume keywords do actively manage their bids [3]. We propose a discrete choice model of bidder decisions that can identify values under both kinds of behavior. Our key insight is that even though an advertiser's bids may show little variation, its *rank* typically varies considerably because of exogenous changes in the auction's parameters, such as the weights (related to click-through rates) placed on bids for ranking.

Our approach is to bypass bids and instead directly model an advertiser's (perhaps passive) choice of rank across auctions. Because there are only a small number of ad slots—no more than twelve—available on a search results page, an advertiser's choice of rank lends itself well to discrete choice modeling [11]. Besides value per click, our model also provides a useful estimate of the advertiser's *regret* variance, which captures how consistent its behavior is with a single value per click. We evaluate our model in terms of its ability to predict advertiser rank and realized clicks in future auctions, against both simple baselines assuming constant rank and click-through rates, and the recent stochastic variability model of Pin and Key [10].

Related Work. The earliest empirical estimates of advertiser values in sponsored search appear in the work of Varian [12]. His approach is to develop an equilibrium concept to model bidding, and jointly estimate bidder values on individual auction instances (i.e., on single queries) by minimizing deviation from equilibrium. This method, however, does not extend easily to several auction instances over time [1]. In another early work Borgers et al. [2] estimate values using a revealed preference approach: an advertiser's bid updates imply bounds on its value, assuming best-response, and with enough observations the bounds can pin down the value. However, this approach is ineffective if advertisers do not update their bids often, which is very common as previously mentioned.

More recently, Athey and Nekipelov [1] have developed an approach tailored to advertisers with static bids. By modeling the distribution over an advertiser's opponent bids, they derive a marginal cost, or "incremental cost per click" curve, and obtain a value based on where the advertiser bids along the curve. (A rational agent sets marginal cost equal to marginal value.) Pin and Key [10] have developed a simplified version of this approach that yields very similar predictions but is much more scalable. For advertisers that update their bids often, their method must estimate a separate value corresponding to each bid, which may be problematic unless there is good reason to believe their value per click indeed changes with each bid update.

2 The Model

In this section we provide the necessary background on sponsored search needed to understand advertisers' decision problems. We describe the basic model of sponsored search introduced in [4, 12]; for a survey of the literature see [8]. We then present our discrete choice logit model of bidder behavior; for a full treatment of logit and other discrete choice models see the monograph [11].

2.1 Sponsored Search

We first focus on a single search query for a given keyword. When the query is issued, an auction is run to allocate the ad slots on the search results page among advertisers bidding on the keyword. Let K be the number of slots and N be the number of agents, where $N > K$. The core of the current auction mechanism (ranking and pricing) used by major search engines is known as the *generalized second-price auction* (GSP) [4]. Each agent i places a bid b_i, and the search engine assigns weights w_i that depend on the ad's past click-through rates. The ads are then ranked in descending order of their *score* $w_i b_i$. Without loss of generality, we can re-index the agents so that $w_1 b_1 \geq w_2 b_2 \geq \ldots \geq w_N b_N$.

Agents are charged only when a click is received. In the GSP, payment follows a second-price rule: an agent is charged the lowest bid it could have placed while maintaining its position. In particular, to maintain its position, agent i must bid so that $w_i b_i \geq w_{i+1} b_{i+1}$, and so its price per click (PPC) is $w_{i+1} b_{i+1} / w_i$. In practice search engines also set a *reserve score* s for each keyword, so that the minimum PPC i can pay is $r_i = s/w_i$. If i bids below r_i, its ad is not shown.

The click-through rate (CTR) of ad i in position j is denoted c_{ij}. We assume that CTRs are *separable* into an advertiser effect a_i and a position effect x_j, meaning that they factor according to $c_{ij} = a_i x_j$. Although separability is only an approximation to actual CTR patterns [1], search engines still estimate ad-specific and position-specific parameters because the ad effect a_i is a key input into the ad's weight w_i. We assume that each agent i has a value per click v_i and that its utility is quasi-linear, meaning that if it obtains slot j at a PPC of p_j then its (expected) utility is:

$$V_j = (v_i - p_j) c_{ij}. \tag{1}$$

In practice the bid space is discretized into increments (e.g., 10 cents), but these are fine enough relative to the range of allowed bids that the bid space can be viewed as continuous. However, note that an agent's utility only depends on the particular position selected, holding the other agents' bids fixed. In a single auction scenario, we can therefore view an agent's bidding decision as a discrete choice problem of selecting which position to appear in. The bid confers no more information about the agent's value beyond the position selected.

2.2 Discrete Choice

From the perspective just developed, we can model an agent's collective rank decisions across the auctions it participates in by using methods of discrete choice analysis from econometrics [11]. The basis of discrete choice analysis is the random utility model. In our context, this model posits that an agent i's utility for slot j decomposes into $U_j = V_j + \epsilon_j$, where ϵ_j is a random error, and V_j is the *representative utility* given by (1), derived from observable features of the chosen alternative—in our case, simply the position effect. Agents act rationally in that they choose the slot j with highest utility U_j. Under the random utility model, an agent's choice of rank can change from auction to auction even if the others' bids are held fixed, as the error terms vary across auctions. In each auction, the random utility induces a distribution over the agent's choice of position.

We use a maximum likelihood approach to fit the representative utility's parameters. In discrete choice modeling the observations take the form $U_{\sigma(t)} \geq U_j$ for $j = 1, \ldots, K$ and $t = 1, \ldots, T$, where t indexes the auctions and $\sigma(t)$ is the slot chosen at auction t. We emphasize that the observations, and therefore the model, do not take into account the bids placed, only the ranks obtained at each auction instance. The actual parametric model we fit is of the form:

$$U_j = \beta_v x_j + \beta_p x_j p_j + \epsilon_j. \tag{2}$$

Because utility can be normalized to any scale, we have dropped the leading a_i term from the equations, and the error variance can also be normalized to some convenient constant C. This follows from the fact that only differences in utility matter when making a choice—we refer to [11] for the technical and conceptual details. Once we fit the model to data, the coefficient $-\beta_p$ corresponds to the marginal utility of money, and hence $-\beta_v/\beta_p$ gives an estimate of v_i. The error variance, which was normalized to C, is proportional to β_p^{-2} on the money scale.

Error Interpretation. A common interpretation of the error term is that it captures unobserved features of the alternatives that impact utility. However, under our value-per-click model, nothing differentiates slots besides their position effects. Instead, we find it more appropriate to interpret the error terms as capturing an agent's *regret*, defined as the amount of foregone utility from choosing one slot over another. If the agent chooses slot j over k, for instance, then $V_k - V_j$ is its regret and $\rho_{kj} = \epsilon_k - \epsilon_j \geq V_k - V_j$ is a bound on this regret (which is binding when errors are minimized). The error distribution induces a distribution over regret. Note that regret can be negative, in which case it indicates the amount by which the chosen slot is preferred over the alternative.

As we will see in Section 3, agents typically hold their bids constant for long periods of time. In this case, variation in rank across auctions comes from exogenous changes such as updates to the advertiser effects, the number of opponents, or the reserve score [10]. The distribution of an agent's regret from keeping its bid fixed is therefore induced by these exogenous changes. Nevertheless, we find it fair to characterize an agent's distribution over ranks and regret as its "behavior", even if it holds its bid fixed, because the distribution captures the extent to which the agent manages its bid to maximize utility. Indeed, the regret (equivalently, error) variance in discrete choice models is sometimes interpreted as a measure of "bounded rationality" [5].

Error Distribution. To complete the model specification we need to detail the error distribution. In this work we assume that errors are independently and identically distributed according to an extreme value distribution with mean 0, which implies that regrets are distributed according to a logistic distribution with mean 0. This is known as the *logit* model, and with this specification there is a closed-form formula for the choice probabilities of different slots given their representative utilities [11]. Once we have fit $v_i = -\beta_v/\beta_p$ and therefore obtain $V_1, \ldots, V_K$ for agent i in a given auction, and the *inverse* estimated error variance is $\lambda = \beta_p^{-2}/C$, then the choice probabilities are given by the familiar logit formula:

$$\Pr(i,j) = \frac{e^{\lambda V_j}}{\sum_{k=1}^{K} e^{\lambda V_k}}.$$

Observe that as λ increases (error variance decreases), the choice probabilities put increasing mass on the slot with highest representative utility, and the agent is utility maximizing in the limit. As λ decreases (error variance increases), the choice probabilities become increasingly uniform over slots, and the agent's choices of ranks across auctions are less consistent with a fixed value per click.

3 Data Description

Our empirical analysis is based on Yahoo's sponsored search logs for a one month period in the first half of 2010. We randomly sampled 20 keywords from each of the top 5 keyword deciles by volume. These 100 keywords together yield a data set of nearly three million auctions that involve 15699 unique advertisers. We used the first three weeks of data for training and the last week for testing. As our study only examines advertisers present in both the training and testing data sets, the number of included advertisers drops to 2603.

To estimate our logit model as specified in (2) we used the `mlogit` package in R [13]. We fit a model to each advertiser separately. The construction of the logit model entails computing the PPC p_j of each position j that ad i sees as an option, for every auction. Because ad i only occupied a certain position j, yielding one p_j value, we computed p_k for every $k \neq j$, by applying the second-price rule described in Section 2.1 and using available data on its opponents' weights and bids, as well as its reserve price.[3]

We needed to further filter the data in order to avoid regression problems such as collinearity. Estimating the PPC of the last position requires information about the ad immediately below the last ranked ad, which is unavailable in our data—we only have records on ads that were shown. Therefore we discarded the final slot in each auction as an alternative. We also filtered out advertisers that were almost always charged their reserve price, and consequently appeared mainly in the bottom positions, because such ads saw the same PPC for multiple slots (i.e., the reserve price) which created singularity issues for the regression. Although the amount of discarded data due to this latter issue accounts for more than 20% of the remaining data, these advertisers' rank decisions, which overwhelmingly focus on getting the bottom slots, would provide little information

[3] Yahoo maintains two different reserve scores: one for the mainline (ads shown at the top) and the sidebar (ads shown on the right). The mainline reserve price was not available for this analysis. However, we found that using the second-price rule together with the sidebar reserve price alone was enough to reproduce observed PPC's to within 0.02% accuracy on average.

about bidders' behavior in general. Moreover, as the inclusion of ads displayed infrequently would add significant noise to our analysis, while providing few insights about bidding behavior, we further removed more than 1500 ads that were shown less than once a day on average, leaving us with a dataset of 197 ads.

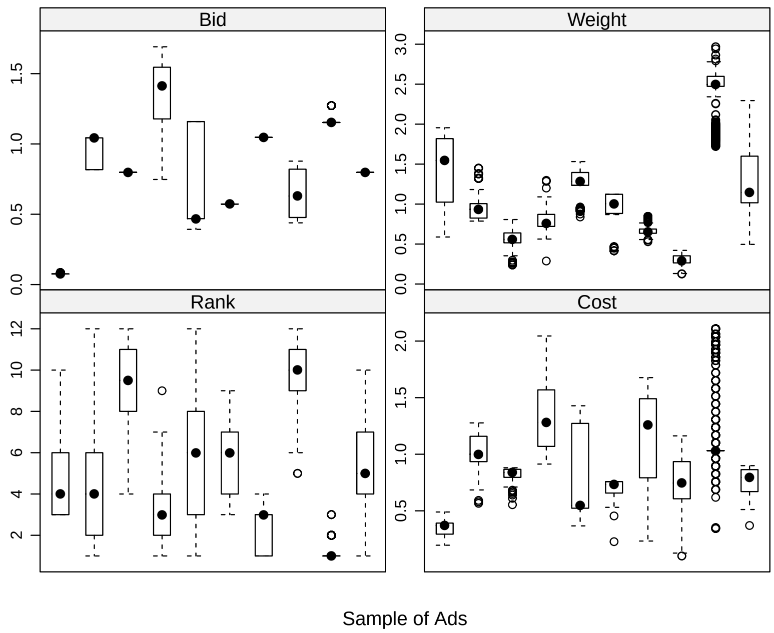

Fig. 1. Variation in bid, rank, weight (closely related to ad effect), and cost (i.e., PPC) for a representative set of 10 ads sampled uniformly at random from our dataset. The center dot gives the median; the box gives the lower and upper quartiles; the whiskers give the minimum and maximum; and any remaining dots indicate outliers. The ads' bids, weights, and costs have been normalized by the mean bid, weight, and cost for the ads' respective keywords to enable variation comparisons across panels.

Preliminary Analysis. Figure 1 presents a simple summary of bidding behavior for a representative sample of 10 ads from our dataset of 197. In this figure an ad's bids were normalized by the average bid (over opponents) of its associated keyword; we also normalized weight and cost (i.e., PPC) in the same way. A normalized bid of 1 means that it matches the average bid on the keyword.

We observe that six out of the ten ads barely vary their bid at all, and only the fourth and fifth substantially vary their bid. Pin and Key [10] also found that bids changed very little in Microsoft's sponsored search market, so this is a general feature of sponsored search and not just Yahoo's market. On the other hand, note that there is substantial variation in rank for all ads. In particular, the first ad takes on positions between 3 and 10, and the sixth ad between 3 and 9, even though their bids stay constant. For the first ad, some of the rank variation can be attributed to its changing weight (i.e., the estimate of its ad effect changes across auctions). There is less weight variation for the sixth ad, but other factors can change the rank such as variation in the number of opponents.

The variation in rank and cost here makes it possible for our discrete choice approach to identify values for the advertisers. Revealed preference approaches based on bid changes, as in [2], would fare poorly because of the dearth of bid update observations for most ads. On the other hand, for those ads whose bids do vary a lot, such as the third and fourth, approaches based on static bids need to estimate a different value for each bid, while our logit model also handles this case and estimates a single value for each advertiser.

4 Regression Results

We first report on the regression coefficients of the fitted models for each ad to confirm that they take on sensible values. Among the 197 ads we examine, 178 (or 90%) have nonzero β_v and β_p coefficients significant at the 5% level. All of these 178 ads have positive β_v coefficients and negative β_p coefficients, implying that utility is increasing in CTR and decreasing in PPC, as expected.

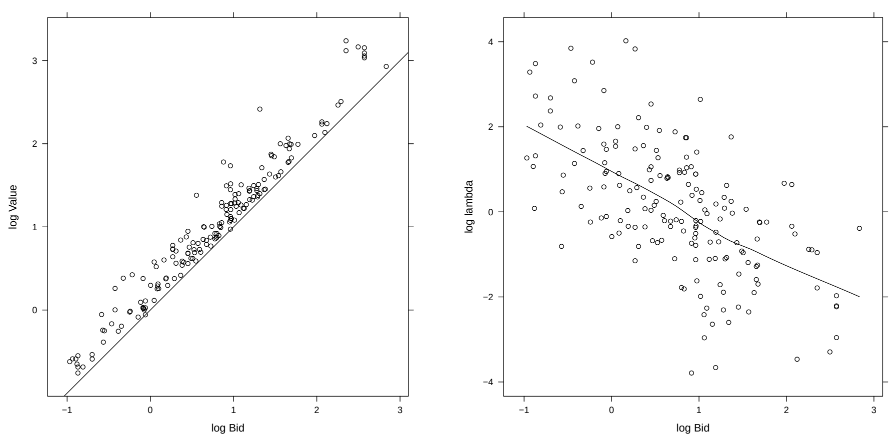

Fig. 2. Regressions results (value and error variance) against average bid (over the testing period) over 178 ads. Axes have been arbitrarily normalized for confidentiality reasons. In the left panel we see that bids lie uniformly below values. The Loess curve in the right panel confirms that inverse error variance λ decreases with bid in general.

Figure 2 provides a more detailed look at the estimated values per click v_i and inverse error variances λ_i, derived from the regression coefficients as explained in Section 2.2. According to the theory on sponsored search [4, 12], bidding above one's value is a dominated strategy, so we would expect estimated values to exceed bids. This is corroborated in the figure, where values uniformly lie above the agents' average bids over the training period. We find this result striking because our model imposes no such constraints on agents' values, and indeed with little training data (one week rather than three) we do observe a few estimated values falling below bids. In fact, recall that bids are not an input to our model: we only rely on the observed position effects, PPC's, and an agent's rank at each auction. According to these estimated values agents shade their bids 20% below their value, on average (in terms of median and mean). The agents' return on investment (ROI), defined as profit per click over PPC, had a median of 48% and a mean of 95%, indicating a skewed distribution.

In Figure 2 we also see how the inverse error variances λ_i correlate with average bids. Recall that a high λ_i suggests that i is behaving more 'rationally', in the sense that its choice of slots across auctions is almost consistent with a fixed value per click, whereas a low λ_i indicates a more 'irrational' agent because its choices imply a high regret variance no matter what value per click is ascribed. According to the figure bid is negatively correlated with 'rationality', or stated more formally, correlated with high regret variance. We see the following possible reason. Low bidders tend to compete either on low-competition keywords or for low-ranked slots, and in those cases the slots are similar in terms of both CTR and PPC. Therefore even if the agent's position varies among these bottom slots (as it holds it bid fixed), its regret stays low and varies little. The situation is the opposite for high bidders that appear on high-competition keywords, where the top slots are highly differentiated in terms of CTR and PPC.

5 Model Evaluation

In this section we first describe the baseline models against which we compare our logit model, and then proceed to evaluate their performance in predicting future ranks and realized clicks.

5.1 Baseline Models

We first compare our logit model, denoted as $\mathsf{M}_{\mathsf{logit}}$, against two simple baseline models that provide predictions about bidders' positions and number of clicks using empirical distributions constructed directly from training data. The first simple baseline model, the *constant rank* model ($\mathsf{M}_{\mathsf{rank}}$), specifically focuses on rank predictions. In particular, $\mathsf{M}_{\mathsf{rank}}$ assumes that each advertiser seeks to have its ad i displayed at a targeted position j^*, and treats the most frequent observed position for ad i in the training data as its targeted position j^* by assigning a probability value of 90% to j^*. In order to account for variation in agents' positions, the model allows positions other than j^* to appear with equal probabilities that sum up to the remaining $(100\% - 90\%) = 10\%$.

The second simple baseline model, called *historical click* ($\mathsf{M}_{\mathsf{click}}$), is tailored for click predictions, assuming that agents expect to receive a constant click through rate for each auction. Given the training data, $\mathsf{M}_{\mathsf{click}}$ computes $\bar{c}_i$, the average number of clicks per auction that ad i received during the training data's timespan, and uses $\bar{c}_i$ to estimate the number of clicks i will receive in the future by multiplying $\bar{c}_i$ with the number of auctions in which i has a slot.

We further evaluate the estimates produced by $\mathsf{M}_{\mathsf{logit}}$ against those obtained from the *stochastic* model ($\mathsf{M}_{\mathsf{stoch}}$) of Pin and Key [10]. They model an agent called Agent 0 with known value v_0 and weight w_0 submitting bid b_0 against n opponents, who submit random i.i.d bids. Note that in this context, each agent is associated with an ad and a bid value, which diverges from the viewpoint previously used in the other models, $\mathsf{M}_{\mathsf{rank}}$, $\mathsf{M}_{\mathsf{click}}$, and $\mathsf{M}_{\mathsf{logit}}$, that only view each ad as a different agent, who may place multiple bids across time.

$\mathsf{M}_{\mathsf{stoch}}$ assumes that the agents' weighted bids $w_i b_i / w_0$, from the perspective of Agent 0, are drawn from a known probability distribution, whose cumulative distribution function (c.d.f) is denoted as F. As the number of opponents may vary from one auction to another, $\mathsf{M}_{\mathsf{stoch}}$ incorporates a discrete probability distribution on the number of opponents, q_n, where $\sum_{n=0}^{N-1} q_n = 1$.

When Agent 0 bids b_0 greater than the reserved price, the probability that it gets the j-th position given that it faces n opponents is:

$$\Pr(j; n) = \binom{n}{j} q(n) F(b_0)^{n-j} (1 - F(b_0))^j. \tag{3}$$

Let $\psi(b_0)$ be the CTR of the slot Agent 0 receives when it bids b_0. The expected number of clicks that Agent 0 receives per auction is computed as:

$$\mathrm{E}[\psi(b_0)] = \sum_{n=0}^{N-1} \sum_{j=0}^{n} \Pr(j; n) c_{0j}. \tag{4}$$

Given the training data set, we can construct the distribution of number of opponents q_n and the distribution of weighted bids F for each agent. These distributions allow us to estimate their expected ranks and expected number of clicks per auctions via (3) and (4) respectively.

Recall that an ad may appear in $\mathsf{M}_{\mathsf{stoch}}$ as different agents, each of which corresponds to a different bid value submitted for the same ad. In order to compare the predictions of $\mathsf{M}_{\mathsf{stoch}}$ with those of the other models, we apply $\mathsf{M}_{\mathsf{stoch}}$ to each pair of ad and bid value, and subsequently average over all same-ad pairs to compute the estimates for each ad.

5.2 Estimation Results

We evaluate the models $\mathsf{M}_{\mathsf{rank}}$, $\mathsf{M}_{\mathsf{click}}$, $\mathsf{M}_{\mathsf{stoch}}$, and $\mathsf{M}_{\mathsf{logit}}$ based on their predictions of ads' ranks and clicks they receive. Note that the $\mathsf{M}_{\mathsf{rank}}$ baseline only applies to rank prediction, and the $\mathsf{M}_{\mathsf{click}}$ baseline only applies to click prediction.

Rank Distribution. We measure the predictive power of $\mathsf{M}_{\mathsf{rank}}$, $\mathsf{M}_{\mathsf{stoch}}$, and $\mathsf{M}_{\mathsf{logit}}$ with respect to ads' ranks by the likelihood of the testing data induced by each model. In particular, given a model $\mathsf{M} \in \{\mathsf{M}_{\mathsf{rank}}, \mathsf{M}_{\mathsf{stoch}}, \mathsf{M}_{\mathsf{logit}}\}$ learned from the

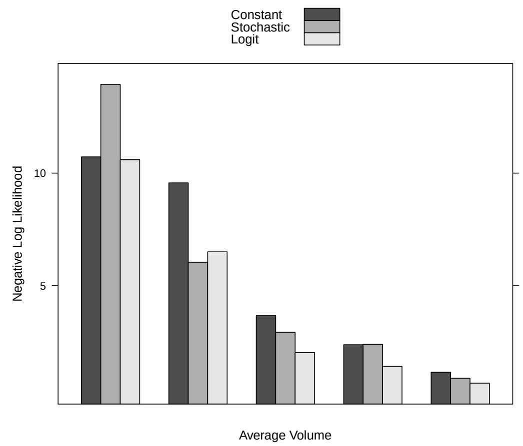

Fig. 3. Rank prediction results, with ads divided into 5 bins according to average volume. Volume increases exponentially towards the right.

training data for an ad i, we compute the log likelihood of the ad's positions in the testing data set D of m auctions in which i won a slot, as follows:

$$L_i(D \mid \mathsf{M}) = \sum_{t=1}^{m} \log \Pr_i(\sigma(t) \mid M), \quad (5)$$

where $\Pr_i(\sigma(t) \mid M)$ is the probability that i gets the slot $\sigma(t)$ in auction t, as specified by model M. In order to investigate these models' robustness to data availability, we also varied the training data set's timespan. Due to space constraints we only report in detail on the results from models trained on the whole first three weeks of the data.

We divide the ads into 5 bins according to the volume of auctions in which they are present. For each bin, we compute the average negative log likelihood of ads' ranks per auction and per ad. Figure 3 shows that the logit model $\mathsf{M}_{\mathsf{logit}}$ consistently outperforms the simple baseline $\mathsf{M}_{\mathsf{rank}}$ model in every bin, and also provides better rank predictions than the stochastic model $\mathsf{M}_{\mathsf{stoch}}$ in most bins. As expected, prediction performance improves as the average volume (and hence amount of training data) increases.

Realized Clicks. We next compare the models $\mathsf{M}_{\mathsf{click}}$, $\mathsf{M}_{\mathsf{stoch}}$, and $\mathsf{M}_{\mathsf{logit}}$, by evaluating the estimated number of clicks each ad would receive against the number of realized clicks in the testing data. Given a model M's estimated number of clicks received by ad i per auction, $\hat{c}_i^{\mathsf{M}}$, and the number of realized clicks per auction over the testing data for ad i, c_i^*, we can calculate the relative error[4] for model M as follows:

$$\mathrm{err}_i(D \mid \mathsf{M}) = \frac{\mid \hat{c}_i^{\mathsf{M}} - c_i^* \mid}{c_i^*}. \quad (6)$$

[4] We can only compute this relative error for the 51 ads in our dataset that received at least one click.

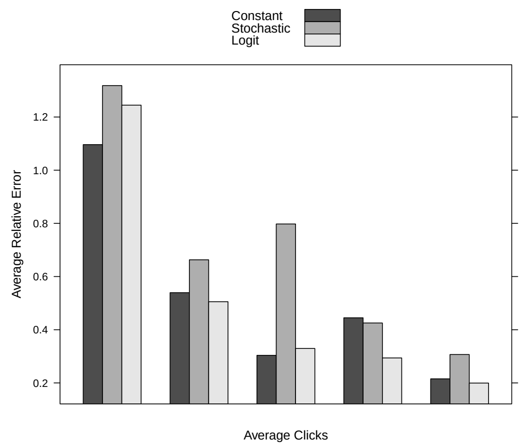

Fig. 4. CTR prediction results, with ads divided into 5 bins according to average number of clicks. Clicks increase exponentially towards the right.

We again split the ads into 5 bins, this time based on their average number of clicks, and then compute each bin's relative error as the average of $\text{err}_i(D \mid \mathsf{M})$ over all ads i in the bin for each model $\mathsf{M} \in \{\mathsf{M}_{\mathsf{click}}, \mathsf{M}_{\mathsf{stoch}}, \mathsf{M}_{\mathsf{logit}}\}$. Figure 4 demonstrates that the simple baseline $\mathsf{M}_{\mathsf{click}}$ model performs particularly well for ads that attract fewer hits, beating both $\mathsf{M}_{\mathsf{stoch}}$ and $\mathsf{M}_{\mathsf{logit}}$ in the least-clicked ad bin. The logit model predicts clicks noticeably better than $\mathsf{M}_{\mathsf{click}}$ for ads that receive more clicks, and moreover, consistently outperforms $\mathsf{M}_{\mathsf{stoch}}$ in all bins. The predictive power of each model improves as the average number of realized clicks increases, as observed in a different study [10]. That study examined only ads that received at least as much actual clicks as the ads in our two most clicked bins, namely the two right-most bins in Figure 4. They also incorporated results from a baseline model similar to $\mathsf{M}_{\mathsf{click}}$, but trained this baseline model on less data than the baseline $\mathsf{M}_{\mathsf{click}}$ we employed.

Note that in order to make rank and click predictions $\mathsf{M}_{\mathsf{stoch}}$ has to examine the actual bids placed by agents in the testing data. In contrast, $\mathsf{M}_{\mathsf{logit}}$ examines an agent's opponents' bids in the testing data to predict realized rank (or more precisely, rank distribution) and clicks; it does not directly draw on the agent's behavior (i.e., bids) to predict. Despite this seeming disadvantage, $\mathsf{M}_{\mathsf{logit}}$ performs very well against $\mathsf{M}_{\mathsf{stoch}}$.

6 Conclusion

We have introduced a novel discrete-choice approach to modeling the bidding behavior of both active and inactive bidders in sponsored search. Our logit model of advertisers' rank decisions produces bidder value estimates that are consistent with basic theory on how values relate to bids, even though these constraints are not incorporated into the regressions. Our empirical evaluation showed that the logit model predicts realized ranks and clicks well, against both simple baselines and a more sophisticated baseline that even draws on agents' actual bidding behavior to make predictions, in contrast to our approach.

The parametric form of utility given in (2) is one of several potential options that we hope to investigate in future work. For instance, we could add position-specific intercepts to the utility specification in order to see whether advertisers value higher slots more than lower slots, all else (i.e., click-through rate) held equal, which would indicate utility for the "branding effect" of slots. We could use a nested logit model [11] that not only relaxes the i.i.d. assumption of the error term, but can also incorporate the variation in the number of bidders competing in an auction. Finally, as our data filtering process left us with a much smaller sample than the original set, we would like to scale up our empirical analysis and include more high-click rather than high-volume ads.

References

[1] Susan Athey and Denis Nekipelov. A structural model of sponsored search advertising auctions. Technical report, Microsoft Research, May 2010.

[2] Tilman Borgers, Ingemar Cox, Martin Pesendorfer, and Vaclav Petricek. Equilibrium bids in sponsored search auctions: Theory and evidence. Working paper, 2008.

[3] Benjamin Edelman and Michael Ostrovsky. Strategic bidder behavior in sponsored search auctions. *Decision Support Systems*, 43:192–198, 2070.

[4] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the Generalized Second Price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1), March 2007.

[5] Jacob K. Goeree, Charles A. Holt, and Thomas R. Palfrey. Regular quantal response equilibrium. *Experimental Economics*, 8:347–367, 2005.

[6] Dustin Hillard, Stefan Schroedl, Eren Manavoglu, Hema Raghavan, and Chris Leggetter. Improving ad relevance in sponsored search. In *Proceedings of the third ACM international conference on Web search and data mining*, pages 361–370, New York, NY, 2010. ACM.

[7] Sébastien Lahaie and David M. Pennock. Revenue analysis of a family of ranking rules for keyword auctions. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, pages 50–56, 2007.

[8] Sébastien Lahaie, David M. Pennock, Amin Saberi, and Rakesh V. Vohra. Sponsored search auctions. In Noam Nisan, Tim Roughgarden, Éva Taros, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, pages 699–716. Cambridge University Press, 2007.

[9] Michael Ostrovsky and Michael Schwarz. Reserve prices in Internet advertising auctions: A field experiment. Working paper, 2009.

[10] Furcy Pin and Peter Key. Stochastic variability in sponsored search auctions: Observations and models. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 61–70, 2011.

[11] Kenneth E. Train. *Discrete Choice Methods with Simulation*. Cambridge University Press, 2009.

[12] Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 25:1163–1178, 27.

[13] Achim Zeileis and Yves Croissant. Extended model formulas in R: Multiple parts and multiple responses. *Journal of Statistical Software*, 34(1):1–13, 2010.