

A Structural Model of Sponsored Search Advertising Auctions*

Susan Athey^a and Denis Nekipelov^b.

^a Department of Economics, Harvard University, Cambridge, MA & Microsoft Research

^b Department of Economics, UC Berkeley, Berkeley, CA & Microsoft Research

This Version: January 2012

Abstract

Sponsored links that appear beside Internet search results on the major search engines are sold using real-time auctions. Advertisers place standing bids, and each time a user enters a search query, the search engine holds an auction. Ranks and prices depend on advertiser bids as well as “quality scores” that vary over time, by advertisement and individual user query. Existing models assume that bids are customized for a single user query. In practice queries arrive more quickly than advertisers can change their bids. This paper develops a new model that incorporates these features. In contrast to prior models, which produce multiplicity of equilibria, we provide sufficient conditions for existence and uniqueness of equilibria. In addition, we propose a homotopy-based method for computing equilibria.

We propose a structural econometric model. With sufficient uncertainty in the environment, the valuations are point-identified; otherwise, we consider bounds on valuations. We develop an estimator which we show is consistent and asymptotically normal, and we assess the small sample properties of the estimator using Monte Carlo.

We apply the model to historical data for several search phrases. Our model yields lower implied valuations and bidder profits than those suggested by the leading alternative model. We find that bidders have substantial strategic incentives to reduce their expressed demand in order to reduce the unit prices they pay in the auctions. These incentives are asymmetric, leading to inefficient allocation, which does not arise in models that ignore uncertainty. Although these are among the most competitive search phrases, bidders earn substantial profits: values per click are between 40 and 90 percent higher than prices for low-ranked bidders, but are between 90 and 270 percent higher for high-ranked bidders. Our counterfactual analysis shows that for the search phrases we study, the auction mechanism used in practice is only slightly less efficient than a Vickrey auction, but the revenue effects are ambiguous, with Vickrey leading to substantially less revenue for one search phrase and slightly more for the other. In addition, less efficient quality scoring algorithms can in some cases raise substantially more revenue with relatively small efficiency consequences.

* We acknowledge Anton Schwaighofer, Michael Ostrovsky, Phil Haile, Dmitry Taubinsky, Nikhil Agarwal, Hoan Lee, Daisuke Hirata, Chris Sullivan, Andrew Bacher-Hicks, Daniel Waldinger, and Maya Meidan for helpful comments and assistance with this paper. We also thank the participants at Cowles foundation conference (Yale), INFORMS meetings, NBER winter meeting, the Ad Auctions Workshop at EC 2010, 4th CAPCP Conference, and seminar audiences at Harvard, MIT, UC Berkeley and Microsoft Research for helpful comments.

1 Introduction

Online advertising is a big business. Search advertising is an important way for businesses, both online and offline, to attract qualified leads; revenues from search advertising auctions top \$30 billion per year.

This paper develops and analyzes original theoretical and econometric models of advertiser behavior in the auctions, and applies these models to a real-world dataset. The methods can be used to infer bidder valuations from their observed bids, and to reliably and quickly compute counterfactual equilibrium outcomes for differing economic environments (e.g. different auction format, altered competitive environment). We apply the tools to address economic questions. For example, we quantify the extent to which existing auction rules lead to inefficient allocation as compared to a Vickrey auction, as well as the way in which competition affects the magnitude of the inefficiency.

The model proposed in this paper differs from existing economic models (e.g. [Borgers et al. \(2007\)](#), [Edelman et al. \(2007\)](#), [Ghose and Yang \(2009\)](#), [Mela and Yao \(2009\)](#), [Varian \(2007\)](#)) by incorporating more realistic features of the real-world bidding environment. We show that our more realistic model has several advantages in terms of tractability, ability to rationalize bidding data in an equilibrium framework, and in the specificity of the predictions it generates: it simultaneously avoids the problems of multiplicity of equilibria and lack of point-identification of values that are the focus of much of the existing literature.

Sponsored links that appear beside Internet search results on the major search engines are sold using real-time auctions. Advertisers place standing bids that are stored in a database, where bids are associated with search phrases that form part or all of a user’s search query. Each time a user enters a search query, applicable bids from the database are entered in an auction. The ranking of advertisements and the prices paid depend on advertiser bids as well as “quality scores” that are assigned for each advertisement and user query. These quality scores vary over time, as the statistical algorithms incorporate the most recent data about user clicking behavior on this and related advertisements and search queries.

[Edelman et al. \(2007\)](#) and [Varian \(2007\)](#) assume that bids are customized for a single user query and the associated quality scores; alternatively, one can interpret the models as applying to a situation where quality scores, advertisement texts, and user behavior are static over a long period of time which is known to advertisers. However, in practice quality scores do vary from query to query, queries arrive more quickly than advertisers can change their bids,¹ and advertisers cannot perfectly predict changes

¹Although bids can be changed in real time, the system that runs the real-time auction is updated only periodically based on the state at the time of the update, so that if bids are adjusted in rapid succession, some values of the bids might never be applied.

in quality scores. This paper develops a new model where bids apply to many user queries, while the quality scores and the set of competing advertisements may vary from query to query. In contrast to existing models, which produce multiplicity of equilibria, we provide sufficient conditions for existence and uniqueness of equilibria, and we provide evidence that these conditions are satisfied empirically for the search phrases we study. One requirement is sufficient uncertainty about quality scores relative to the gaps between bids. We show that the necessary conditions for equilibrium bids can be expressed as an ordinary differential equation, and we develop a homotopy-based method for calculating equilibria given bidder valuations and the distribution of uncertainty.

We then propose a structural econometric model. With sufficient uncertainty in the environment, valuations are point-identified; otherwise, we propose a bounds approach. We develop an estimator for bidder valuations, establish consistency and asymptotic normality, and use Monte Carlo simulation to assess the small sample properties of the estimator.

In the last part of the paper, we apply the model to historical data for two search phrases. We start by comparing the estimates implied by our model to those implied by prior approach focusing on envy-free equilibria and ignoring uncertainty, showing that our model yields lower implied valuations and bidder profits. We then use our estimates to examine the magnitude of bidders' incentives to shade their bids and reduce their expressed demands in order to maximize profits, focusing on the degree to which such incentives are asymmetric across bidders with high versus low valuations. We demonstrate that differential bid-shading leads to inefficient allocation.

We analyze the elasticity of residual supply curve for clicks faced by each bidder, which determines the equilibrium gap between each bidder's value and its price per click (the Lerner index, i.e., the profit per click' as a percentage of the price per click). We find that for both search phrases we consider, the Lerner index is generally higher for the top positions than the bottom ones, and for one of our search phrases, the top positions have substantially higher profits. The higher profits correspond to two relatively high-value advertisers who each dominate one of the top two positions with fairly substantial gaps between their bids, while the lower ranks are more competitive in the sense that many advertisers with similar values compete for them.

The incentives for “demand-reduction” are created by the use of a “generalized second-price auction” (GSP), which [Edelman et al. \(2007\)](#) and [Varian \(2007\)](#) show is different from a Vickrey auction. In a model without uncertainty, one of the main results of [Edelman et al. \(2007\)](#) and [Varian \(2007\)](#) is that the GSP auction is outcome-equivalent to a Vickrey auction for a particular equilibrium selection, which we refer to as the “EOS” equilibrium; however, we show that the equivalence breaks down when bidders use differential bid shading and the same bids apply to many user queries with varying quality scores.

A Vickrey auction, run query by query, would lead bidders to bid their values and thus would result in efficient allocation in each auction even when quality scores vary query by query, so long as the quality score is such that it is efficient to rank the advertisements in terms of score-weighted values. Thus, our findings suggest that there is a non-trivial role for auction format to make a difference in this setting, a finding that would not be possible without uncertainty and using the EOS equilibrium, since then, auction format plays no role.

In our model, the revenue difference between the GSP and the Vickrey auctions varies widely between the two search phrases. For our first search phrase, we find that the Vickrey auction raises 22% less revenue than the GSP, even though the efficiency difference is only 0.04%. For this search phrase, the revenue gap arises almost exclusively in the first position, where the model suggests that the GSP incents the second-place bidder to bid closer to his value than the top bidder, extracting more revenue from the top bidder. Efficiency consequences are small because of the large gap between the score-weighted bid (and score-weighted value) of the top bidder and the rest, so that high-value bidder (who generates most of the surplus) is rarely displaced. Reallocations occur at lower positions in the GSP, but they tend to occur when the efficiency consequences of the rerankings are small (bidders have very similar score-weighted values).

For our second search phrase, the Vickrey auction raises 1.2% more revenue, with an efficiency difference of 0.3%. The revenue effects vary non-monotonically with position: GSP raises more revenue in the middle positions but less from the top and bottom positions.

Finally, we explore the incentives of the search advertising platform to use efficient quality scoring algorithms. In our model, if we ran a Vickrey auction, efficiency would be attained by setting each advertisement’s quality score equal to the platform’s best estimate of its “clickability,” which is equal to the clicks the advertisement would attain in the first position, so that the advertisements that that generate the highest value per impression (value per click times expected clicks) are ranked the highest. Even though asymmetric bidding incentives introduce some inefficiency in the GSP, the inefficiency is small when quality scores are equal to clickability. We consider two alternative approaches to quality scores that introduce greater inefficiency. One is an algorithm that uses “squashing” ([Lahaie and Pennock \(2007\)](#)), an approach whereby quality scores are set as a monotone, nonlinear function of clickability that compresses the quality scores. We would expect squashing to increase revenue when the bidders in the highest positions also have high clickability, because squashing lowers their quality scores and thus increases the prices they pay in the GSP. Squashing can be understood in the context of the insight from [Meyerson \(1981\)](#) that biasing auctions in favor of weaker bidders can increase revenue at the expense of efficiency.

A second alternative is “coarsening,” whereby the quality scores are made less precise. For example,

a search advertising platform might face the choice between two clickability estimation algorithms, one which uses more explanatory variables and produces correspondingly more accurate estimates, and one which is less accurate.

We show that “coarsening” the quality scoring algorithm raises revenue for both search phrases, in one case just 2.2% and in the other case, substantially (18%). Squashing lowers revenue slightly (2%) for one search phrase while it increases revenue substantially for the other (9%). This suggests that in practice, search engines might face strong incentives to sacrifice efficiency to raise revenue, so long as the inefficient scoring scheme is carefully designed (since the scores are set by algorithms, in principle the extent of squashing or coarsening can vary query by query).

2 Overview of Sponsored Search Auctions

Auction design for sponsored search auctions has evolved over time; see [Edelman et al. \(2007\)](#) for a brief history. Since the late 1990s, most sponsored search in the U.S. has been sold at real-time auctions. Advertisers enter per-click bids into a database of standing bids. They pay the search engine only when a user clicks on their ad. Each time a user enters a search query, bids from the database of standing bids compete in an auction. Applicable bids are collected, the bids are ranked, per-click prices are determined as a function of the bids (where the function varies with auction design), and advertisements are displayed in rank order, for a fixed number of slots J . Clicks are counted by the ad platform and the advertiser pays the per-click price for each click. For simplicity, we will focus exposition on a single search phrase, where all advertisers place a distinct bid that is applicable to that search phrase alone.²

In this setting, even when there are fewer bidders than positions, bidders are motivated to bid more aggressively in order to get to a higher position and receive more clicks. Empirically, it has been well established that appearing in a higher position on the screen causes the advertisement to receive more clicks. We let $\alpha_{i,j}$ be the ratio of the “click-through rate” (CTR, or the probability that a given query makes a click on the ad) that advertiser i would receive if its ad appears in position j , and the CTR of the ad in position 1. The CTR for the highest slot can be tens to hundreds of times higher than for lower slots. The way in which CTRs diminish with position depends on the search phrase in question.

In 2002, Google introduced the “generalized second price” auction. The main idea of this auction is that advertisements are ranked in order of the per-click bids (say, b_1, \dots, b_N with $b_i > b_{i+1}$) and a bidder

²In general, bidders can place “broad match” bids that apply to any search phrase that includes a specified set of “keywords,” but for very high-value search phrases, such as the ones we study here, most advertisers who appear on the first page use exact match bidding.

pays the minimum per-click price required to keep the bidder in her position (so bidder i has position i and pays b_{i+1}). When there is only a single slot, this auction is equivalent to a second-price auction, but with multiple slots, it differs. Subsequently, Google modified the auction to include weights, called “quality scores,” for each advertisement, where scores are calculated separately for each advertisement on each search phrase. These scores were initially based primarily on the estimated click-through rate the bidder would attain if it were in the first position (the “clickability”). The logic behind this design is straightforward: allocating an advertisement to a given slot yields expected revenue equal to the product of the price charged per click, and the click-through rate. Thus, ranking advertisements by the product of the clickability and the per-click bid is equivalent to ranking them by the expected revenue per impression (that is, the revenue from displaying the ad) they would attain in the top position. Later, Google introduced additional variables into the determination of the weights, including measures of the match between the advertisement and the query. Although the formulas used by each search advertising platform are proprietary information and can change at any time, the initial introduction of quality scores by Microsoft and Yahoo! was described in the industry as a GSP auction using the “click-weighting” version of quality scores, that is, quality scores reflect primarily clickability, and so when needed in our empirical analysis, we will treat the quality scores as estimates of clickability. However, the theoretical analysis does not rely on a particular interpretation of the scores.

In practice, there are also a number of reserve prices that apply for the different advertising platforms. Our empirical application generally has non-binding reserve prices, so we ignore them in the empirical work, but we include the most basic kind of reserve price in the theory.

3 The Model

3.1 A Static Model of a Score-Weighted Generalized Second-Price Auction

We begin with a static model, where each of I advertisers simultaneously place per-click bids b_i on a single search phrase. The bids are then held fixed and applied to all of the users who enter that search phrase over a pre-specified time period (e.g. a day or a week). There is a fixed number of advertising slots J in the search results page.

We model consumer searches as an exogenous process, where each consumer’s clicking behavior is random and $\bar{c}_{i,j}$, the average probability that a consumer clicks on a particular ad in a given position, is the same for all consumers. It will greatly simplify exposition and analysis to maintain the assumption that the parameters $\alpha_{i,j}$ (the ratio of advertisement i ’s CTR in position j to its CTR in position 1) satisfy $\alpha_{i,j} = \alpha_{i',j} \equiv \alpha_j$ for all advertisements i, i' ; we will maintain that assumption throughout the

paper.³ That is, there exists a vector of advertisement effects, γ_i , $i = 1, \dots, I$, and position effects α_j , $j = 1, \dots, J$, with $\alpha_1 = 1$, such that $\bar{c}_{i,j}$ can be written

$$\bar{c}_{i,j} = \alpha_j \gamma_i.$$

The ad platform conducts a click-weighted GSP auction. Each advertisement i is assigned score s_i , and bids are ranked in order of the product $b_i s_i$. In general discussion we will use i to index bidders and j to index positions (slots). We will use the double index notation k_j to denote the bidder occupying slot j . The per-click price p_{k_j} that bidder k_j in position j pays is determined as the minimum price such that the bidder remains in her position

$$p_{k_j} = \min\{b_{k_j} : s_{k_j} b_{k_j} \geq s_{k_{j+1}} b_{k_{j+1}}\} = \frac{s_{k_{j+1}} b_{k_{j+1}}}{s_{k_j}}.$$

Note that advertiser k_j does not directly influence the price that she pays, except when it causes her to change positions, so in effect an advertiser's choice of bid determines which position she attains, where the price per click for each position is exogenous to the bidder and rises with position.

To interpret this auction, observe that if for each i , $s_i = \gamma_i$, then the expected revenue the ad platform receives from placing bidder k_j in position j is $\alpha_j \gamma_{k_{j+1}} b_{k_{j+1}}$ which is what the platform would receive if instead, it had placed bidder in slot $j + 1$ in position j and charged bidder in slot $j + 1$ her per-click bid, $b_{k_{j+1}}$, for each click. So each bidder pays, in expectation, the per-impression revenue that would have been received from the next lowest bidder.

We include the possibility of reserve prices in the auction. The reserve price is assumed to be set in units of score-weighted bids, so that an advertisement is considered only if $s_i b_i > r$. We introduce the reserve price in the model by adding an artificial bidder $I + 1$ such that $b_{I+1} = r$ and $s_{I+1} = 1$.

We assume that advertisers are interested in consumer clicks and each advertiser i has a value v_i associated with a consumer click. The profile of advertiser valuations in a particular market (v_1, \dots, v_I) is fixed, and advertisers know their valuations with certainty. Each click provides the advertiser i with the surplus $v_i - p_i$. The advertisers are assumed to be risk-neutral.

3.2 Equilibrium Behavior with No Uncertainty (NU)

The structure of Nash equilibria in the environment similar to that described in the previous subsection has been considered in [Edelman et al. \(2007\)](#) and [Varian \(2007\)](#). We can write the expected surplus

³Empirically, this assumption can be rejected for many search phrases (see, e.g., [Jeziorski and Segal \(2009\)](#)), but the deviations are often small, and the assumption is more likely to hold when the advertisements are fairly similar, as is the case for the search phrases in our sample.

of advertiser i from occupying the slot j as

$$\bar{c}_{i,j}(v_i - p_j) = \alpha_j \gamma_i \left(v_i - \frac{s_{k_{j+1}} b_{k_{j+1}}}{s_i} \right).$$

The existing literature, including [Edelman et al. \(2007\)](#) and [Varian \(2007\)](#), focus on the case where the bidders know the set of competitors as well as the score-weighted bids of the opponents, and they consider ex post equilibria, where each bidder's score-weighted bid must be a best response to the realizations of $s_{k_{j+1}} b_{k_{j+1}}$ (and recall we have also assumed that the $\bar{c}_{i,j}$ are known). Let us start with this case, which we will refer to as the ‘‘No Uncertainty’’ (NU) case.

The set of bids constituting a full-information Nash equilibrium in the NU model, where each bidder finds it unprofitable to deviate from her assigned slot, are those that satisfy

$$\begin{aligned} \alpha_j \left(v_{k_j} - \frac{s_{k_{j+1}} b_{k_{j+1}}}{s_{k_j}} \right) &\geq \alpha_l \left(v_{k_j} - \frac{s_{k_{l+1}} b_{k_{l+1}}}{s_{k_j}} \right), \quad l > j \\ \alpha_j \left(v_{k_j} - \frac{s_{k_{j+1}} b_{k_{j+1}}}{s_{k_j}} \right) &\geq \alpha_l \left(v_{k_j} - \frac{s_{k_l} b_{k_l}}{s_{k_j}} \right), \quad l < j. \end{aligned}$$

It will sometimes be more convenient to express these inequalities in terms of score-weighted values, as follows:

$$\min_{l < j} \frac{s_{k_l} b_{k_l} \alpha_j - s_{k_{j+1}} b_{k_{j+1}} \alpha_l}{\alpha_l - \alpha_j} \geq s_{k_j} v_{k_j} \geq \max_{l > j} \frac{s_{k_{j+1}} b_{k_{j+1}} \alpha_j - s_{k_{l+1}} b_{k_{l+1}} \alpha_l}{\alpha_j - \alpha_l}.$$

An equilibrium always exists, but it is typically not unique, and equilibria may not be monotone: bidders with higher score-weighted values may not be ranked higher.

[Edelman et al. \(2007\)](#) and [Varian \(2007\)](#) define a refinement of the set of equilibria, which [Edelman et al. \(2007\)](#) call ‘‘envy-free’’: no bidder wants to exchange positions and bids with another bidder. The set of envy-free equilibria is characterized by a tighter set of inequalities:

$$s_{k_j} v_{k_j} \geq \frac{s_{k_{j+1}} b_{k_{j+1}} \alpha_j - s_{k_{j+2}} b_{k_{j+2}} \alpha_{j+1}}{\alpha_j - \alpha_{j+1}} \geq s_{k_{j+1}} v_{k_{j+1}}. \quad (3.1)$$

The term in between the two inequalities is interpreted as the incremental costs divided by the incremental clicks from changing position, or the ‘‘incremental cost per click’’ $ICC_{j,j+1}$:

$$ICC_{j,j+1} = \frac{s_{k_{j+1}} b_{k_{j+1}} \alpha_j - s_{k_{j+2}} b_{k_{j+2}} \alpha_{j+1}}{\alpha_j - \alpha_{j+1}}.$$

Envy-free equilibria are monotone, in that bidders are ranked by their score-weighted valuations, and have the property that local deviations are the most attractive—the equilibria can be characterized by incentive constraints that ensure that a bidder does not want to exchange positions and bids with either the next-highest or the next-lowest bidder.

Edelman et al. (2007) consider a narrower class of envy-free equilibria, the one with the lowest revenue for the auctioneer and the one that coincides with Vickrey payoffs as well as the equilibrium of a related ascending auction game. They require

$$s_{k_j} v_{k_j} \geq ICC_{j,j+1} = s_{k_{j+1}} v_{k_{j+1}}. \quad (3.2)$$

Edelman et al. (2007) show that despite the fact that payoffs coincide with Vickrey payoffs, bidding strategies are not truthful: bidders shade their bids, trading off higher price per click in a higher position against the incremental clicks they obtain from the higher position.

We refer to the equilibrium defined by $ICC_{j,j+1} = s_{k_{j+1}} v_{k_{j+1}}$ as the EOS equilibrium, and the equilibrium defined by $s_{k_j} v_{k_j} = ICC_{j,j+1}$ as the NU-EFLB equilibrium (for “envy-free lower bound”).

3.3 Equilibrium Behavior with Score and Entry Uncertainty (SEU)

In reality, advertiser bids apply to many unique queries by users. Each time a query is entered by a user, the set of applicable bids is identified, scores are computed, and the auction is conducted as described above. In practice, both the set of applicable bids and the scores vary from query to query. This section describes this uncertainty in more detail and analyzes its impact on bidding behavior.

3.3.1 Uncertainty in Scores and Entry in the Real-World Environment

The ad platform produces scores at the advertisement-query level using a statistical algorithm. A key component of quality scores is the click-through rate that the platform predicts the advertisement will attain. In practice, the distribution of consumers associated with a given search query and/or their preferences for given advertisers (or for advertisements relative to algorithmic links) can change over time, and so the statistical algorithms are continually updated with new data. Google has stated publicly that it uses individual search history to customize results to individual users; to the extent that Google continues to use the GSP, ranking ads differently for different users can be accomplished by customizing the quality scores for individual users.

In Appendix B, we illustrate how introducing small amounts of uncertainty affects equilibrium in the NU model. However, because the real-world environment incorporates substantial uncertainty, we focus our exposition in the text on non-trivial uncertainty.

We assume that the score of a particular bidder i for a user query is a random variable, denoted s_i , which is equal to

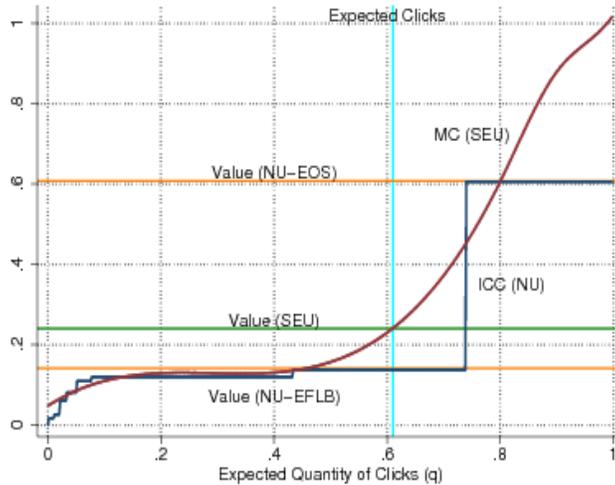
$$s_i = \bar{s}_i \varepsilon_i,$$

where ε_i is a shock to the score induced by random variation in the algorithm’s estimates.⁴

Now consider uncertainty in bidder entry. There are many sources of variation in the set of advertisements that are considered for the auction for a particular query. First, some bidders specify budgets (limits on total spending at the level of the account, campaign, or keyword), which the ad platforms respect in part by spreading out the advertiser’s participation in auctions over time, withholding participation in a fraction of auctions. Bidders may also “pause” and “reactivate” their campaigns. Second, bidders experiment with multiple advertisements and with different ad text. These advertisements will have distinct click-through rates, and so will appear to other bidders as distinct competitors. For new advertisements, it takes some time for the system to learn the click-through rates; and the ad platform’s statistical algorithm may “experiment” with new ads in order to learn. Third, some bidders may target their advertisements at certain demographic categories, and they may enter different bids for those categories (platforms make certain demographic categories available for customized bidding, such as gender, time of day, day of week, and user location).

For these and other reasons, it is typical for the configuration of ads to vary on the same search phrase; this variation is substantial for all three major search ad platforms in the U.S., as can be readily verified by repeating the same query from different computers or over time.

Figure 1: Marginal and Incremental Cost and Implied Valuations for Alternative Models



The role of the score and entry uncertainty can be illustrated by Figure 1. The x-axis gives the (expected) click-through rate a bidder receives (the “click share”), relative to the click-through rate it

⁴In the subsequent discussion we often refer to \bar{s}_i as the “mean score” of the bidder. We will further make an identifying assumption $E[\log \varepsilon_i] = 0$ which may not imply that the mean score is equal to \bar{s}_i . We will indicate specific points where this distinction is important.

would attain in the top position (that is, the average of α_j over the positions a bidder with a given bid experiences). The step function in the figure shows the relationship between the incremental cost per click and expected number of clicks for a single user query, with a commonly observed configuration of advertisements and associated bids, and assuming that each advertisement is assigned a score equal to its average score from the week. As the bidder in question's score-weighted bid increases and crosses the score-weighted bid of each opponent, the bidder moves to a higher position, receiving a higher average CTR. Given a value of $\alpha \in [\alpha_{j+1}, \alpha_j]$, the associated incremental cost per click is $ICC_{j,j+1}$.

The smooth curve shows how uncertainty affects the incremental cost per click. The curve is constructed by varying the bid of a given advertisement. For each value of the bid, we calculate the expectation of the share of possible clicks the advertisement receives, where the expectation is taken over possible realizations of quality scores, using the distribution of these scores we estimate below. Corresponding to each expected click share, we calculate the marginal cost of increasing the click share and plot that on the y-axis (details of the computation are provided below). The marginal cost curve increases smoothly rather than in discrete steps because the same advertisement with the same bid would appear in different positions for different user queries, and changing the bid slightly affects outcomes on a small but non-zero share of user queries.

This smoothness reflects the degree of uncertainty faced by the advertisers. For the search phrases we consider, the most commonly observed advertisements have a standard deviation of their position number close to 1 position.

3.3.2 Formalizing the Score and Entry Uncertainty (SEU) Model

Start with the NU model, and consider the following modifications. Bids are fixed for a large set of user queries on the same search phrase, but the game is still a simultaneous-move game: bidders simultaneously select their bids, and then they are held fixed for a pre-specified period of time. Let \tilde{C}^i be a random subset of advertisers excluding advertiser i , with typical realization C^i , and consider shocks to scores as defined in the last subsection.

We use the solution concept of ex post Nash equilibrium. In the environment with uncertainty, we need to specify bidder beliefs. Since our environment has private values (bidders would not learn anything about their own values from observing the others' information) and we model the game as static, an ex post Nash equilibrium merely requires that each bidder correctly anticipates the mapping from his own bids to expected quantities and prices, taking as given opponent bids. Note that the major search engines provide this feedback to bidders through advertiser tools (that is, bidders can enter a hypothetical bid on a keyword and receive estimates of clicks and cost).

Despite these weak information requirements, for simplicity of exposition, we endow the bidders with information about the primitive distributions of uncertainty in the environment. That is, we assume that advertisers correctly anticipate the share of user queries where each configuration of opposing bidders C^i will appear; the mean of each opponent's score-weighted bid, $b_i \bar{s}_i$; and the distribution of shocks to scores, $F_\varepsilon(\cdot)$.

Define Φ_{ik}^j to be an indicator for the event that bidder i is in slot j and bidder k is in slot $j + 1$. Let $\mathcal{C}_{j,k}^i$ be the set of all subsets of C^i with cardinality j that contains k , with typical element $C_{j,k}^i$, representing the set of bidders above bidder i as well as k . Also recall that we introduce the reserve price as an additional bidder $I + 1$ with bid $b_{I+1} = r$, $\bar{s}_{I+1} = 1$ and $\varepsilon_{I+1} = 1$. Let b, \bar{s}, ε be vectors of bids, mean scores, and shocks to scores, respectively. Then:

$$\Phi_{ik}^j(b, \bar{s}, \varepsilon; C^i) = \sum_{C_{j,k}^i \in \mathcal{C}_{j,k}^i} \prod_{m \in C_{j,k}^i \setminus \{k\}} \mathbf{1}\{b_m \bar{s}_m \varepsilon_m > b_i \bar{s}_i \varepsilon_i\} \prod_{m \in C^i \setminus C_{j,k}^i} \mathbf{1}\{b_m \bar{s}_m \varepsilon_m < b_k \bar{s}_k \varepsilon_k\} \mathbf{1}\{b_i \bar{s}_i \varepsilon_i > b_k \bar{s}_k \varepsilon_k\}.$$

We can then write the expected number of clicks a bidder will receive as a function of her bid b_i as follows:

$$Q_i(b_i; b_{-i}, \bar{s}) = \mathbb{E}_{\tilde{C}^i, \varepsilon} \left[\sum_{j=1, \dots, J} \sum_{k \in \tilde{C}^i} \Pr(\Phi_{ik}^j(b, \bar{s}, \varepsilon; \tilde{C}^i) = 1) \cdot \alpha_j \cdot \gamma_i \right].$$

The expected total expenditure of the advertiser for the clicks received with bid b_i can be written

$$TE_i(b_i; b_{-i}, \bar{s}) = \mathbb{E}_{\tilde{C}^i, \varepsilon} \left[\sum_{j=1, \dots, J} \sum_{k \in \tilde{C}^i} \Pr(\Phi_{ik}^j(b, \bar{s}, \varepsilon; \tilde{C}^i) = 1) \cdot \alpha_j \cdot \gamma_i \cdot \frac{\bar{s}_k \varepsilon_k b_k}{\bar{s}_i \varepsilon_i} \right].$$

Then, the bidder's problem is to choose b_i to maximize

$$EU_i(b_i; b_{-i}, \bar{s}) \equiv v_i \cdot Q_i(b_i; b_{-i}, \bar{s}) - TE_i(b_i; b_{-i}, \bar{s}). \quad (3.3)$$

We let $EU(b, \bar{s})$ and $TE(b, \bar{s})$ be vector functions where the i th elements are $EU_i(b_i; b_{-i}, \bar{s})$ and $TE_i(b_i; b_{-i}, \bar{s})$, respectively, for $i = 1, \dots, I$.

We assume that the distributions of the scores have bounded supports. In general, this can lead to a scenario where expected clicks, expenditures and thus profits are constant in bids over certain ranges, since there can be a range of bids that maintain the same average position.

3.4 Existence, Uniqueness, and Computation of Equilibria in the SEU Model

In this section, we derive a particularly convenient representation of the conditions that characterize equilibria in the SEU model, and then we show that standard results from the theory of ordinary differ-

ential equations can be used to provide necessary and sufficient conditions for existence and uniqueness of equilibrium. We start by making the following assumption, which we maintain throughout the paper.

ASSUMPTION 1. *The vector of shocks to the scores $\varepsilon = (\varepsilon_1, \dots, \varepsilon_{I+1})$ has the following properties: the components are independent; the distribution of ε_{I+1} is degenerate at 1; the remaining I components are identically distributed with distribution $F_\varepsilon(\cdot)$, which does not have mass points and has an absolutely continuous density $f_\varepsilon(\cdot)$ that has a finite second moment and that is twice continuously differentiable and strictly positive on its support.*

Many of the results in the paper carry over if this assumption is relaxed, but they simplify the analysis substantially.

To begin, we present a simple but powerful identity, proved below in 3.4:

$$\frac{d}{d\tau} EU_i(\tau b, \bar{s})|_{\tau=1} = -TE_i(b, \bar{s}), \quad \text{for } i < I + 1, \quad (3.4)$$

that is, a proportional increase in all bids decreases bidder i 's utility at the rate $TE_i(b, \bar{s})$, the amount bidder i is spending. The intuition is that ranks and prices depend on the ratios of bids, so a proportional change in all bids simply increases costs proportionally.

The system of first-order conditions that are necessary for equilibrium when all bidders win a strictly positive click share is given by

$$v_i \frac{\partial}{\partial b_i} Q_i(b, \bar{s}) = \frac{\partial}{\partial b_i} TE_i(b, \bar{s}) \quad \text{for all } i. \quad (3.5)$$

Our next result works by combining (3.4) with the first-order conditions, to conclude that a proportional increase in *opponent bids only* decreases utility at the rate $TE_i(b, \bar{s})$; this follows because when bidder i is optimizing, a small change in her own bid has negligible impact.

LEMMA 1 *Assume that $\frac{\partial}{\partial b_{-(I+1)}} EU(b, \bar{s})$ and TE are continuous in b . Then a vector of strictly positive bids b satisfies the first order necessary conditions for equilibrium (3.5) if and only if*

$$\frac{d}{d\tau} EU_i(b_i, \tau b_{-i}, \bar{s})|_{\tau=1} = -TE_i(b, \bar{s}) \quad \text{for all } i < I + 1. \quad (3.6)$$

Proof: Denote the probability of bidders i and k from configuration $C^i \cup \{i\}$ being in positions j and $j + 1$ by $G_{ik}^j(b, \bar{s}, C^i)$. Then $G_{ik}^j(b, \bar{s}, C^i) = \int \Phi_{ik}^j(b, \bar{s}, \varepsilon; C^i) dF_\varepsilon(\varepsilon)$, recalling that Φ_{ik}^j is an indicator

for the event that bidder i is in slot j and bidder k is in slot $j + 1$. The total quantity of clicks for bidder i can be computed as

$$Q_i(b_i, b_{-i}; \bar{s}) = \sum_{C^i} \sum_{k \in C^i} \sum_{j=1}^J \alpha_j \gamma_i G_{ik}^j(b, \bar{s}, C^i).$$

The total expenditure can be computed as

$$TE_i(b_i, b_{-i}; \bar{s}) = b_i \sum_{C^i} \sum_{k \in C^i} \sum_{j=1}^J \alpha_j \gamma_i \int \frac{\bar{s}_k b_k \varepsilon_k}{\bar{s}_i b_i \varepsilon_i} \Phi_{ik}^j(b, \bar{s}, \varepsilon; C^i) dF_\varepsilon(\varepsilon).$$

Note that $TE_i(b_i, b_{-i}; \bar{s})/b_i$ is homogeneous of degree zero in b .

The function $G_{ik}^j(b, \bar{s}, C^i)$ is homogeneous of degree zero in b as well. As a result, $\sum_{k'=1}^K b_{k'} \frac{\partial}{\partial b_{k'}} G_{ik}^j(b, \bar{s}, C^i) = 0$. Then, the following identity holds

$$\frac{\partial}{\partial b'} EU(b, \bar{s}) b = -TE(b, \bar{s}), \quad (3.7)$$

which can in turn be rewritten as, for each $i < I + 1$,

$$\frac{d}{d\tau} EU_i(b_i, \tau b_{-i}, \bar{s})|_{\tau=1} + b_i \frac{\partial}{\partial b_i} EU_i(b_i, b_{-i}, \bar{s}) = -TE_i(b, \bar{s}).$$

Thus, (3.6) is equivalent to $\frac{\partial}{\partial b_i} EU_i(b_i, b_{-i}, \bar{s}) = 0$ whenever $b_i > 0$.

Q.E.D.

We now build on 3.4 to analyze existence and uniqueness of equilibrium. To do so, we introduce some additional notation.

Let $EU(b, \bar{s})$ be the vector of bidder expected utilities, and let $D(b, \bar{s})$ the matrix of partial derivatives

$$D(b, \bar{s}) = \frac{\partial}{\partial b'_{-(I+1)}} EU(b, \bar{s}).$$

Let $D_0(b, \bar{s})$ be the matrix obtained by replacing the diagonal elements of $D(b, \bar{s})$ with zeros. Then, (3.6) can be rewritten in matrix notation as

$$D_0(b, \bar{s}) b_{-(I+1)} = -TE(b, \bar{s}) - r \frac{\partial EU(b, \bar{s})}{\partial b_{I+1}}.$$

We can then implicitly define a mapping $\beta : [0, 1] \rightarrow \mathbb{R}^{I+1}$ where, under some regularity conditions imposed on the payoff function, the solution $\beta(\tau)$ will exist in some neighborhood of $\tau = 1$:

$$\tau \frac{d}{d\tau} EU_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s}) = -TE_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s}), \quad b_{I+1} = r \quad (3.8)$$

for all bidders $i < I + 1$. The next theorem establishes the conditions under which a solution $\beta(\tau)$ exists locally around $\tau = 1$ and globally for $\tau \in [0, 1]$. To state the theorem, let $V = [0, v_1] \times \cdots \times [0, v_I]$ be the support of potential bids when bidders bid less than their values, as will be optimal in this game.

THEOREM 1. *Consider a GSP auction in the SEU environment with a reserve price $r > 0$. Assume that D_0 and TE are continuous in b and D_0 is non-singular and Lipschitz at the solution of (3.5). Suppose that for each $i=1, \dots, I$, $Q_i(b, \bar{s}) > 0$, and that each EU_i is quasi-concave in b_i on V . Then:*

(i) *An equilibrium exists if and only if for some $\delta > 0$ the system of equations (3.8) has a solution on $\tau \in [1 - \delta, 1]$.*

(ii) *The conditions from part (i) are satisfied for all $\delta \in [0, 1]$, and so an equilibrium exists, if $D_0(b, \bar{s})$ is Lipschitz and non-singular for all $b \in V$ except a finite number of points.*

(iii) *There is a unique equilibrium if and only if for some $\delta > 0$ the system of equations (3.8) has a unique solution on $\tau \in [1 - \delta, 1]$.*

(iv) *The conditions from part (iii) are satisfied for all $\delta \in [0, 1]$, so that there is a unique equilibrium, if each element of $\frac{\partial}{\partial b} EU(b, \bar{s})$ is Lipschitz in b and non-singular for $b \in V$.*

The full proof of this theorem is provided in the Appendix. Quasi-concavity is assumed to ensure that solutions to the first-order condition are always global maxima; it is not otherwise necessary.

Theorem 1 makes use of a high-level assumption that the matrix D_0 is non-singular. In the following lemma we provide more primitive conditions outlining empirically relevant cases where this assumption is satisfied.

LEMMA 2 (i) *For a given $b \in V$ and a given \bar{s} , suppose that the bidders are arranged according to their mean score weighted values $\bar{s}_i b_i \geq \bar{s}_{i+1} b_{i+1}$ for $i = 1, \dots, I - 1$. Then $D_0(b, \bar{s})$ is non-singular if for each bidder her utility is strictly locally monotone in the bid of either the bidder above or below her in the ranking or both.*

(ii) *Given V , there exist values $\underline{\varepsilon}$ and $\bar{\varepsilon}$ such that (i) is satisfied for all $b \in V$ if the support of ε contains $[\underline{\varepsilon}, \bar{\varepsilon}]$.*

Part (i) of 3.4 is satisfied, e.g., if $\frac{\partial EU_i}{\partial b_{i-1}} \neq 0$ for $i = 2, \dots, I$ and $\frac{\partial EU_1}{\partial b_2} \neq 0$. To see this, note that the diagonal elements of the matrix $D_0(b, \bar{s})$ are zero. Therefore, we can compute the determinant

$\det(D_0(b, \bar{s})) = -\frac{\partial EU_1}{\partial b_2} \prod_{i>2} \frac{\partial EU_i}{\partial b_{i-1}} \neq 0$. For part (ii) we note that we can find $\underline{\varepsilon}$ and $\bar{\varepsilon}$ such that for each bidder i we can find bidder i' such that $b_i \bar{s}_i \underline{\varepsilon} > b_{i'} \bar{s}_{i'} \bar{\varepsilon}$ and $b_i \bar{s}_i \bar{\varepsilon} < b_{i'} \bar{s}_{i'} \underline{\varepsilon}$. Then the probability that bidder i' is ranked below bidder i is positive and depends on the bid of bidder i' . Thus, the derivative of bidder i 's utility with respect to the bid of bidder i' is not equal to zero.

Equation (3.8) plays a central role in determining the equilibrium bid profile. Now we show that it can be used as a practical device to compute the equilibrium bids. Suppose that functions TE_i and EU_i are known for all bidders. Then, initializing $\beta(0) = 0$, we treat the system of equations (3.8) as a system of ordinary differential equations for $\beta(\tau)$. We can use standard methods for numerical integration of ODE if a closed-form solution is not available. Then the vector $\beta(1)$ will correspond to the vector of equilibrium bids.

This suggests a computational approach, which can be described as follows. Suppose that one needs to solve a system of non-linear equations

$$\mathbf{H}(\mathbf{b}) = \mathbf{0},$$

where $\mathbf{H} : \mathbb{R}^N \mapsto \mathbb{R}^N$ and $\mathbf{b} \in \mathbb{R}^N$. This system may be hard to solve directly because of significant non-linearities. However, suppose that there exists a function $\mathbf{F}(\mathbf{b}, \tau)$ such that $\mathbf{F} : \mathbb{R}^N \times [0, 1] \mapsto \mathbb{R}^N$ with the following properties. If $\tau = 0$, then the system

$$\mathbf{F}(\mathbf{b}, 0) = 0$$

has an easy-to-find solution, and if $\tau = 1$ then

$$\mathbf{F}(\mathbf{b}, 1) = \mathbf{H}(\mathbf{b}) = 0.$$

Denote the solution of the system $\mathbf{F}(\mathbf{b}, 0) = 0$ by \mathbf{b}_0 . If \mathbf{F} is smooth and has a non-singular Jacobi matrix, then the solution of the system

$$\mathbf{F}(\mathbf{b}, \tau) = 0$$

will be a smooth function of τ . As a result, we can take the derivative of this equation with respect to τ to obtain

$$\frac{\partial \mathbf{F}}{\partial \mathbf{b}'} \dot{\mathbf{b}} + \frac{\partial \mathbf{F}}{\partial \tau} = 0,$$

where $\dot{\mathbf{b}} = \left(\frac{db_1}{d\tau}, \dots, \frac{db_N}{d\tau} \right)'$. This expression can be finally re-written in the form

$$\dot{\mathbf{b}} = - \left(\frac{\partial \mathbf{F}}{\partial \mathbf{b}'} \right)^{-1} \frac{\partial \mathbf{F}}{\partial \tau}. \quad (3.9)$$

Equation (3.9) can be used to solve for $\beta(\tau)$. $\beta(0) = \mathbf{b}_0$ is assumed to be known, and $\beta(1)$ corresponds to the solution of the system of equations of interest. Systems of ordinary differential equations are usually easier to solve than non-linear equations.

The computational approach we propose is to define \mathbf{F} using (3.8). If the payoff function is twice continuously differentiable and the equilibrium existence conditions are satisfied, then \mathbf{F} has the desired properties. Details of the application of this method to our problem are in Appendix G.

3.5 Bidder Incentives in the SEU Model

It is easier to understand the bidder's incentives in terms of general economic principles if we introduce a change of variables. When bidding, the advertiser implicitly selects an expected quantity of clicks, and a total cost for those clicks. Fix b_{-i}, \bar{s} and suppress them in the notation, and define

$$Q_i^{-1}(q_i) = \inf\{b_i : Q_i(b_i) \geq q_i\},$$

and define

$$\begin{aligned} TC_i(q_i) &= TE_i(Q_i^{-1}(q_i)). \\ AC_i(q_i) &= TE_i(Q_i^{-1}(q_i))/q_i. \end{aligned}$$

Then, the bidder's objective can be rewritten as

$$\max_{q_i} q_i(v_i - AC_i(q_i)).$$

This is isomorphic to the objective function faced by an oligopsonist in an imperfectly competitive market. As usual, the solution will be to set marginal cost equal to marginal value, when the average cost curve is differentiable in the relevant range (assume it is for the moment).

$$v_i = q_i AC_i'(q_i) + AC_i(q_i) \equiv MC_i(q_i). \quad (3.10)$$

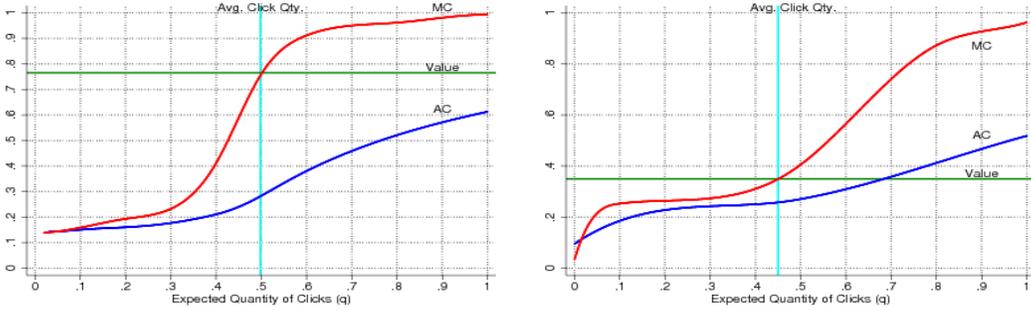
The bidder trades off selecting a higher expected CTR (q_i) and receiving the average per-click profit $v_i - AC_i(q_i)$ on more units, against the increase in the average cost per click that will be felt on all existing clicks. The optimality conditions can be rewritten in the standard way:

$$\frac{v_i - AC_i(q_i)}{AC_i(q_i)} = \frac{d \ln AC_i(q_i)}{d \ln(q_i)}.$$

The bidder's profit as a percentage of cost depends on the elasticity of the average cost per click curve.

To see how this works in practice, consider the following figure, which illustrates the average cost curve $AC_i(q_i)$ and the marginal cost curve $MC_i(q_i)$ for a given search phrase. We select a particular bidder, call it i . Given the actual bid of the advertiser, b_i , we calculate $q_i = Q_i(b_i; b_{-i}, \bar{s})$. We then calculate $MC_i(q_i)$. If the bidder selects q_i optimally, then $v_i = MC_i(q_i)$, as illustrated in the figure. Thus, under the assumption of equilibrium bidding, we infer that the bidder's valuation must have been $MC_i(q_i)$. In Figure 2 we illustrate the structure of the marginal cost and average cost functions for bidders for

Figure 2: Average cost, marginal cost, and value for frequent bidders



(a) Search phrase #1

(b) Search phrase #2

high-value search phrases that most frequently appear in the top position as compared to other bidders. The horizontal line corresponds to the value per click of the considered bidders and the vertical line reflects the quantity of clicks received by the bidder provided her actual bid.

This approach to inferring a bidder's valuation from her bid and the average cost curve she faces is the main approach we use in our empirical work. The case where the average cost curve is not differentiable is considered below.

4 Identification of Valuations under Alternative Models

In this section, we consider identification and inference in the following environment. We assume that position-specific click-through rates α_j are known; identification of these is discussed in the appendix.

For the SEU model, we consider observing a large number of queries for a given set of potential bidders, and consider the question of whether the valuations of the bidders can be identified. For each query, we assume that we observe bids, the set of entrants, and the scores. For the NU model, it is more subtle to define the problem, given the disconnect between the model and the real-world bidding environment. The model treats each query as separate, and so in principle, we could imagine that a bidder's valuation changes query to query. In that case, we consider identification of the valuation for each query.

4.1 The No Uncertainty Model

Recall the condition for envy-free Nash equilibrium in the NU model, that the score weighted value for bidder j is bounded by incremental cost per clicks $ICC_{j-1,j}$ and $ICC_{j,j+1}$. This implies that observed scores, bids and α_j 's are consistent with envy-free Nash equilibrium bidding for some valuations, if and only if

$$ICC_{j,j+1} = \frac{s_{k_{j+1}} b_{k_{j+1}} \alpha_j - s_{k_{j+2}} b_{k_{j+2}} \alpha_{j+1}}{\alpha_j - \alpha_{j+1}} \text{ is nonincreasing in } j. \quad (4.11)$$

This is a testable restriction of the envy-free Nash equilibrium.

Following [Varian \(2007\)](#), we can illustrate the requirements of the envy-free equilibrium with a figure. Recall [Figure 1](#). The envy-free equilibrium refinement requires that a bidder j selects the position (that is, a feasible click share α_j) that yields the highest value of $s_j v_j \alpha_j - s_j TC_j(\alpha_j)$. This is equivalent to requiring that the score-weighted value is bounded by $ICC_{j,j-1}$ and $ICC_{j,j+1}$.

The requirement that $ICC_{j,j+1}$ is nonincreasing in j corresponds to the total expenditure curve being convex. If (4.11) holds, then we can solve for valuations that satisfy (3.2): we can find score-weighted valuations for each bidder that lie between the steps of the ICC curve. In general, if the inequalities in (3.2) are strict, there will be a set of valuations for the bidder in each position.

Thus, (3.2) determines bounds on the bidder's valuation, as follows:

$$s_{k_j} v_{k_j} \in [ICC_{j,j+1}, ICC_{j-1,j}].$$

For the highest excluded bidder, $v_{k_{J+1}} = b_{k_{J+1}}$, and for the highest position,

$$s_{k_1} v_{k_1} \in \left[\frac{s_{k_2} b_{k_2} \alpha_1 - s_{k_3} b_{k_3} \alpha_2}{\alpha_1 - \alpha_2}, \infty \right).$$

The EOS equilibrium selection requires $s_{k_j} v_{k_j} = ICC_{j-1,j}$.

4.2 The Score and Entry Uncertainty Model

For the case where $Q_i(\cdot)$ and $TE_i(\cdot)$ are strictly increasing and differentiable in the bid, we can recover the valuation of each bid using the necessary condition for optimality

$$v_i = MC_i(Q_i(b_i)), \quad (4.12)$$

given that all of the distributions required to calculate $MC_i(q_i)$ are assumed to be observable. Note that the local optimality condition is necessary but not sufficient for b_i to be a best response bid for a bidder with value v_i ; a testable restriction of the model is that the bid is globally optimal for the valuation that satisfies local optimality. One requirement for global optimality is that the objective function is

locally concave at the chosen bid: $MC'_i(Q_i(b_i)) \geq 0$. A sufficient (but not necessary) condition for global optimality is that MC_i is increasing everywhere, since this implies that the bidder's objective function (given opponent bids) is globally concave, and we can conclude that indeed, b_i is an optimal bid for a bidder with value v_i . If MC_i is nonmonotone, then global optimality of the bid should be verified directly.

Now consider the case where $TE_i(\cdot)$ is not differentiable everywhere. This occurs when score uncertainty has limited support, and when there is not too much uncertainty about entry. This analysis parallels the ‘‘kinked demand curve’’ theory from intermediate microeconomics. Note that $Q_i(\cdot)$ is nondecreasing and continuous from the left, so it must be differentiable almost everywhere. If $Q_i(\cdot)$ is constant on $[b'_i, b''_i)$ and then increasing at b''_i , then $Q_i^{-1}(q_i) = b'_i$ for $q_i \in [Q_i(b'_i), Q_i(b''_i))$, while $Q_i^{-1}(Q_i(b''_i)) = b''_i$. This implies in turn that $TC_i(\cdot)$ is non-differentiable at $Q_i(b''_i)$, and that $MC_i(\cdot)$ jumps up at that point. Thus, if we observe any b_i on $[b'_i, b''_i)$, the assumption that this bid is a best response implies only that

$$v_i \in [MC_i(Q_i(b'_i)), MC_i(Q_i(b''_i))]. \quad (4.13)$$

Summarizing:

THEOREM 2. *Consider the SEU model, where bids are fixed over a large number of queries. Suppose that we observe the bids of I bidders (b_1, \dots, b_I) , the joint distribution of their scores s , and the set of entrants in each query. Then:*

(i) *Bounds on the valuation for bidder i are given by (4.13), where $b'_i = Q_i^{-1}(Q_i(b_i); b_{-i}, \bar{s})$, and $b''_i = \sup\{b'''_i : Q_i(b'''_i; b_{-i}, \bar{s}) = Q_i(b_i; b_{-i}, \bar{s})\}$.*

(ii) *A necessary and sufficient condition for the observed bids to be consistent with ex post equilibrium is that for some (v_1, \dots, v_I) within the bounds from (i), the observed bids (b_1, \dots, b_I) are globally optimal for a bidder solving (3.3). A sufficient condition is that $MC_i(\cdot)$ is nondecreasing for each i .*

The proof follows directly from the discussion above and the fact that the functions Q_i and MC_i are uniquely defined from the observed bids and the distribution of scores and entrants.

Equilibria in the SEU environment are not necessarily envy-free, and further, they are not necessarily monotone in the sense that bidders with higher score-weighted valuations place higher score-weighted bids. However, if there are many bidders and substantial uncertainty, each bidder will face a similar marginal cost curve, and monotonicity will follow.

4.3 Comparing Inferences From Alternative Models

A natural question concerns how the inferences from the NU and SEU models compare, given the same auction environment. In this subsection, we show that if the NU model gives bounds on valuations that are consistent across queries (that is, the intersection of the bounds are non-empty), then those bounds will be contained in the bounds from the SEU model. However, in practice, we find that consistency typically fails—the bounds implied by the NU model for one query do not intersect with the bounds from another.

THEOREM 3. *Consider a dataset with repeated observations of search queries, where bids are constant throughout the sample. Consider two alternative models for inference, the NU model and the SEU model. Assume that the NU model produces bounds on valuations that are consistent for a given bidder across the different observations of search queries in the dataset where the advertiser’s bid is entered, and consider the intersection of these bounds for each advertiser. This intersection is contained in the bounds on valuations obtained using the SEU model.*

Proof: Fix a vector of bids b and the distributions of scores and entrants. Let $u_i(v_i, b'_i; b_{-i}, \bar{s}_i, \varepsilon, C)$ be the bidder’s utility for a particular user query when the bidder’s valuation is v_i and bids b'_i , and for this proof include explicitly each bidder’s valuation as an argument of EU_i . Let v^{NU} be a vector of valuations that is consistent with b being a Nash equilibrium bidding profile in the NU model for all possible realizations of scores and participants. Suppose that v^{NU} is not in the bounds for valuations in the SEU model. Then,

$$\begin{aligned} EU_i(v_i^{NU}, b_i; b_{-i}, \bar{s}) &= \mathbb{E}_{\varepsilon, C}[\max_{b'_i} u_i(v_i^{NU}, b'_i; b_{-i}, \bar{s}_i, \varepsilon, C)] \\ &\geq \max_{b'_i} \mathbb{E}_{\varepsilon, C}[u_i(v_i^{NU}, b'_i; b_{-i}, \bar{s}_i, \varepsilon, C)] \\ &> EU_i(v_i^{NU}, b_i; b_{-i}, \bar{s}). \end{aligned}$$

This is a contradiction. Thus, we conclude that v_i^{NU} is in the bounds for valuations in the SEU model.

5 Estimation of Bidder Valuations

In this section we demonstrate how the expected payoff of a bidder in a position auction can be recovered from the data. The structure of the data for position auctions makes the estimation procedure different from the standard empirical analyses of auctions. In the setup of online position auctions the same set of bids will be used in a set of auctions. In our historical data sample most bidders keep their bids unchanged during the considered time period.

In our data sample a portion of advertisers have multiple advertisements. Bidders submit a separate bid for each ad. Our analysis will be facilitated by the fact that the search engine has a policy of not showing two ads by the same advertiser on the same page. We will use a simplifying assumption that bidders maximize an expected profit from bidding for each ad separately. We associate the bidder with a unique combination of the advertisement and the bid and the goal of our structural estimation will be to recover the latent values per click of the advertisers bidding to place the ads.

Previously we assumed that de-measured scores have the same distribution across advertisers. We use an additional subscript t to indicate individual user queries with bidder configurations. We assume that configurations C_t of the bidders who were considered for user query t are observed. We assume that the number of advertisers I is fixed and denote $N_i = \sum_t \mathbf{1}\{i \in C_t\}$ the number of queries for which advertiser i was considered. Our further inference is based on the assumption that $N_i \rightarrow \infty$ for all bidders $i = 1, \dots, I$. We denote the total number of user queries in the dataset by T .

We impose the following assumption regarding the joint distribution of shocks to the scores and configurations.

ASSUMPTION 2. *The shocks to the scores are independent from the configurations: $\varepsilon_{it} \perp C_t$ for $i = 1, \dots, I$. Configurations of advertisers are i.i.d. across queries and the shocks to the scores are i.i.d. across queries and advertisers with expectation $E[\log \varepsilon_{it}] = 0$.*

Assumption 2 is used in the identification and estimation of the uncertainty in the score distribution.⁵

To analyze the uncertainty of the scores we use their empirical distribution. In our model for bidder i the score in query t is determined as $s_{it} = \bar{s}_i \varepsilon_{it}$. We note that from Assumption 2 it follows that $E[\log s_{it}] = \log \bar{s}_i$. Using this observation, we estimate the mean score from the observed realizations of scores for bidder i for impressions t as $\hat{\bar{s}}_i = \exp\left(\frac{1}{N_i} \sum_t \mathbf{1}\{i \in C_t\} \log s_{i,t}\right)$. Then by Assumption 2 and the Slutsky theorem it follows that $\frac{1}{T} N_i \xrightarrow{p} P(i \in C_t)$. Similarly, we find that $\frac{1}{T} \sum_t \mathbf{1}\{i \in C_t\} \log s_{i,t} \xrightarrow{p} P(i \in C_t) \log \bar{s}_i$. Then the consistency of the mean score estimator follows from the continuous mapping theorem.

Then we form the sample of estimated shocks to the scores using $\hat{\varepsilon}_{it} = \frac{s_{it}}{\hat{\bar{s}}_i}$. As an estimator for the

⁵The assumption that shocks to scores are i.i.d. across advertisers is actually stronger than what is necessary. There could be a component of the shocks to scores that is common to all advertisements on a specific user query, and the distribution of shocks to scores would still be identified following the literature on measurement error. See Appendix F for details. The assumption that $\varepsilon_{it} \perp C_t$ is explored further from the perspective of the empirical application in Appendix H.

distribution of the shocks to the scores we use the empirical distribution

$$\widehat{F}_\varepsilon(\varepsilon) = \frac{1}{I} \sum_{i=1}^I \frac{1}{N_i} \sum_t \mathbf{1}\{i \in C_t\} \mathbf{1}\{\widehat{\varepsilon}_{it} \leq \varepsilon\}.$$

Using Assumption 2 and stochastic equicontinuity of the empirical distribution function, the estimator can be expressed by

$$\begin{aligned} & \frac{1}{T} \sum_t \mathbf{1}\{i \in C_t\} \mathbf{1}\left\{\frac{\bar{s}_i \varepsilon_{it}}{\widehat{s}_i} \leq \varepsilon\right\} = \frac{1}{T} \sum_t \mathbf{1}\{i \in C_t\} \mathbf{1}\{\varepsilon_{it} - \varepsilon \leq 0\} \\ & + f_\varepsilon(0) \frac{\bar{s}_i - \widehat{s}_i}{\bar{s}_i^2} \frac{1}{T} \sum_t \mathbf{1}\{i \in C_t\} \varepsilon_{it} + o_p(1) = E[\mathbf{1}\{\varepsilon_{it} \leq \varepsilon\}] P(i \in C_t) + o_p(1). \end{aligned}$$

Combining this with our previous result we find that $\widehat{F}_\varepsilon(\varepsilon)$ is a consistent estimator for the distribution of the shocks to the scores $F_\varepsilon(\cdot)$.

In the case where the expected payoff function has a unique maximum for each value of the bidder (as we find empirically) the bidder's first-order condition can be used to derive the bidder's valuation, as follows:

$$v_i = \frac{\frac{\partial TE_i(b_i, b_{-i}, \bar{s})}{\partial b_i}}{\frac{\partial Q_i(b_i, b_{-i}, \bar{s})}{\partial b_i}}.$$

Each of the functions needed to recover the value can be estimated from the data. We use the empirical distribution of the scores to approximate the uncertainty in the scores and use the observed bidder configurations to approximate the uncertainty in bidder configurations. To compute the approximation we take independent samples from the empirical sample of observed configurations and estimated shocks to the scores $\{C_t^i, \widehat{\varepsilon}_{kt}\}_{t,k=1,\dots,I}$ excluding the bidder of interest i from the sample (recall that we denoted by C^i the configuration excluding bidder i). Following the literature on bootstrap we index the draws from this empirical sample by t^* and denote the simulated sample size T^* . A single draw t^* will include the configuration $C_{t^*}^i$ and the shocks to the scores for all bidders $\widehat{\varepsilon}_{1t^*}, \dots, \widehat{\varepsilon}_{It^*}$. For consistent inference we require that $\frac{N_i}{T^*} \rightarrow 0$ for all $i = 1, \dots, I$. Then for each such draw we compute the rank of the bidder of interest i as

$$\text{rank}_i(C_{t^*}^i) = \text{rank}\{b_i \widehat{s}_i \widehat{\varepsilon}_{it^*}; b_k \widehat{s}_k \widehat{\varepsilon}_{kt^*}, \forall k \in C_{t^*}^i\}.$$

We also compute the price paid by bidder i as

$$\text{Price}_i(C_{t^*}^i) = \frac{b_k \widehat{s}_k \widehat{\varepsilon}_{kt^*}}{\widehat{s}_i \widehat{\varepsilon}_{it^*}}, \text{ for the } k \text{ such that } \text{rank}_k(C_{t^*}^i) = \text{rank}_i(C_{t^*}^i) + 1.$$

Then we estimate the total expenditure function as

$$\widehat{TE}_i(b_i, b_{-i}, \bar{s}) = \frac{1}{T^*} \sum_{t^*=1}^{T^*} \widehat{\alpha}_{\text{rank}_i(C_{t^*}^i)} \text{Price}_i(C_{t^*}^i),$$

and the expected quantity of clicks as

$$\widehat{Q}_i(b_i, b_{-i}, \bar{s}) = \frac{1}{T^*} \sum_{t^*=1}^{T^*} \widehat{\alpha}_{\text{rank}_i(C_{t^*}^i)}.$$

At the next step we estimate the derivatives. To do that we use a higher-order numerical derivative formula. For a step-size τ_N , depending on the sample size, we compute the implied value as

$$\widehat{v}_i = \frac{-\widehat{T}E_i(b_i - 2\tau_N, b_{-i}, \bar{s}) + 8\widehat{T}E_i(b_i - \tau_N, b_{-i}, \bar{s}) - 8\widehat{T}E_i(b_i + \tau_N, b_{-i}, \bar{s}) + \widehat{T}E_i(b_i + 2\tau_N, b_{-i}, \bar{s})}{-\widehat{Q}_i(b_i - 2\tau_N, b_{-i}, \bar{s}) + 8\widehat{Q}_i(b_i - \tau_N, b_{-i}, \bar{s}) - 8\widehat{Q}_i(b_i + \tau_N, b_{-i}, \bar{s}) + \widehat{Q}_i(b_i + 2\tau_N, b_{-i}, \bar{s})}.$$

The choice of τ_N such that $\sqrt{N_i} \tau_N \rightarrow \infty$, $\sqrt{N_i} \tau_N^3 \rightarrow 0$ and $\tau_N \rightarrow 0$ for all $i = 1, \dots, I$ assures that the empirical numerical derivative above converges to the slope of the population marginal cost function. We use this formula to recover the implied valuations. The following result is based on the derivation in [Hong et al. \(2010\)](#) and its proof is given in the Appendix.

THEOREM 4. *Under the sufficient conditions of Theorem 1 and Assumption 2, the derivative of the total expenditure function with respect to the bid vector satisfies the Lindeberg condition and has a finite variance in the limit, while if the derivative of the total quantity of clicks with respect to the bid vector is non-vanishing in the limit our estimator of valuations is asymptotically normal:*

$$\sqrt{N_i} \tau_N (\widehat{v}_i - v_i) \xrightarrow{d} N \left(0, \frac{324 \Omega}{(Q'_i(b_i, b_{-i}, \bar{s}))^2} \right),$$

where

$$\Omega = \text{Var} \left(\frac{u_i(v_i, b_i + \tau_N; b_{-i}, \bar{s}_i, \varepsilon_{it}, C_t) - u_i(v_i, b_i - \tau_N; b_{-i}, \bar{s}_i, \varepsilon_{it}, C_t)}{\sqrt{\tau_N}} \right)$$

This shows that with the increasing number of impressions, the estimates of advertiser's valuations will be asymptotically normal and their asymptotic variance will be determined by the variance of the profit per click for the advertiser of interest.

Our analysis extends to the case where the objective function of the bidder can have a set of optimal points. An empirical approach to this case is discussed in Appendix D.

6 Data

For estimation we use a sample of data of auctions for two high-value search phrases (these are among the highest revenue search phrases on the advertising platform). The data is historical, for a three-month period sometime between 2006 and 2009, and it has been preserved for research purposes. The specific

time period and the specific search phrases are kept confidential to avoid revealing any proprietary information, and all bids are normalized to a single scale in order to avoid revealing information about the specific revenue of the search phrases. We analyze each search phrase entirely separately, and we compare the results.

We begin with describing the main dataset. There are more than 7,500 searches per week between the two search phrases. We focus on impressions from the first page of advertising results. In the page showing the results of the consumer's search query up to 8 ads are displayed: some in the space above the algorithmic search results and some to the side. In our empirical analysis we control for the position of the advertisement. For consistency of the bidding data with our static analytic framework, we use the data only from one week at a time. However, we compare results across weeks for various specification tests and to validate our general approach.

The following variables are observed for each user query (individual auction): the advertiser account associated with each advertisement; the specific advertisement (characterized by ad text, a bid and a landing page where a user is redirected after clicking on the ad); the positions in which the advertisements were displayed on the screen; the per-click bids and system-assigned scores for the advertisements on the individual query; the per-click prices charged for each advertisement; and the clicks received by each of the advertisements.

A complication that we did not emphasize in the theoretical section is that each advertiser can have multiple active advertisements (possibly with distinct bids) on a given search phrase, while the advertising platform only allows one advertisement per bidder to appear. The different advertisements receive different scores by the system, and thus even if advertiser bids are the same across advertisements, the rotation among different advertisements will create fluctuations in outcomes for opposing bidders. Thus, the variation in advertisements is an important source of uncertainty. They also create complications for thinking about bidder optimization. Why does a bidder have multiple active advertisements, and do the motivations conflict with our assumptions about optimal bidding? In practice, bidders tend to test out variations on advertisements to see whether different ad texts perform better and/or are scored better by the advertising platform.

We chose to handle the multiple advertisements by first treating them separately, and assuming that the advertiser takes the existence of multiple advertisements (and the system's selection among them) as exogenous. Since two advertisements by the same advertiser cannot appear in the same auction, it is possible to treat the advertiser's objective function as additively separable in its bids. We estimate separately the valuations for the different advertisements. We find that valuations and profits are very close for different advertisements by the same bidder. In the empirical application, nearly three quarters of advertisers have only one advertisement, while 78% to 92% have less than three. About 90% of ad

impressions come from advertisers with less than three advertisements, and more than three quarters of advertisers have most of their ad impressions with the same advertisement (Herfindahl indices of more than .85, measuring the “market shares” of the advertisements in terms of user queries).

Another complication that arises with our data is that in our research data set, we observe only the ads that were actually displayed. We also can infer the product of the bid and the quality score for the last ad that was **not** displayed, because it is used to set the price for the last displayed ad. The missing data potentially creates problems for our estimation, because shocks to the quality scores of excluded ads can potentially put them onto the screen, but without knowledge of these advertisements, we exclude that possibility. Another problem is that it can bias our estimation of the distribution of quality scores, because very low draws of quality scores might push an advertisement off of the page, removing it from our sample. We evaluate these biases in a Monte-Carlo study replicating the modelled environment and find that they are economically and statistically insignificant.

7 Estimates from Alternative Models

We use several alternative models for estimation: NU-EOS, NU-EFLB, and SEU.

7.1 Modeling Details for NU Models

In the NU Models, we treat each auction as separate, envisioning that bids and valuations might change from auction to auction. We then empirically characterize whether bounds on valuations are consistent across auctions, and how implied valuations change over time. In particular, for each auction we recover the bounds on valuations using the constructed $ICC_{j,j+1}$ curves for positions. We notice that in a large fraction of cases the ICC curve fails to be monotone auction-by-auction. [Varian \(2007\)](#) suggested computing an approximate weighted solution. We consider a weighted ICC as

$$ICC_{j,j+1}^d = \frac{s_{k_{j+1}} b_{k_{j+1}} \alpha_j d_j - s_{k_{j+2}} b_{k_{j+2}} \alpha_{j+1} d_{j+1}}{\alpha_j d_j - \alpha_{j+1} d_{j+1}},$$

where weights minimize $\sum_{j=1}^J (1 - d_j)^2$ such that a weighted ICC is monotone. In the empirical study we perform this procedure for all considered search phrases. We recover the values of the advertisers from the re-weighted ICC curve as

$$s_{k_j} v_{k_j} \in \left[ICC_{j,j+1}^{d^*}, ICC_{j-1,j}^{d^*} \right],$$

where weights d^* solve the minimization problem above. We abuse notation and omit the index of user query t that should subscript the weights and the score. The selected weights are tailored to each specific auction and vary auction by auction.

Similar to [Varian \(2007\)](#) we find a large number of violations from monotonicity, and we correct them using weighting. We also find that the bounds on valuations fluctuate substantially in the NU models. The fluctuations occur from query to query, oscillating back and forth between bounds for commonly observed sets of entrants and scores, so that it is difficult to imagine rationalizing the fluctuations on the basis of changing valuations (and bids do not change at that frequency, and often don't change at all). The median standard deviation of the recovered value for a single bid across queries is approximately 10% for the NU-EFLB model and approximately 7% to 8% for the NU-EOS model. Moreover, the number of auctions that violate the value monotonicity auction-by-auction exceeds 25% with most violations occurring in the middle and the lowest positions.

The weights aimed at making the ICC curves monotone vary from auction to auction, depending on how far a particular configuration is from the configurations with the monotone ICC.

Throughout, in our analysis of the NU models, we use the variant where ICC curves are adjusted to be monotone using this approach.

Now consider the estimation of values. Across all of the advertisements that have auctions for which the monotonicity of the implied score-weighted values is not violated, we could not find any examples where the bounds have a non-empty intersection. Even restricting the dataset to a very limited period of time (8 hours) allows us to find only 6% of advertisements where the bounds intersect. In this case the number of observations per advertisement ranges from 1 to a few hundred.

To address the issue of non-overlapping bounds, we take the following approach. The set of recovered values corresponds to the bounds constructed from the incremental cost curve, with the lower bound corresponding to NU-EFLB and the upper bound corresponding to NU-EOS. For each bidder we can collect a set of values corresponding to different user queries. We use the median over different values of the lower and upper bounds for each bidder as estimates of valuations from the full-information models NU-EFLB and NU-EOS, respectively (the median is used rather than the mean, to reduce the sensitivity of our results to assumptions we make about the upper bound of valuations for the NU-EOS model for the bidder in the top position, where the upper bound is not identified from the data). We should note that this approach may result in negative implied profit per click for some queries, since we infer values using the median values across user queries.

7.2 Modeling Details for the SEU Model

The exact procedure for estimation of the first-order condition has been described in the previous section. We observe empirically that the SEU model yields very tight bounds or point estimates for almost all advertisers. As a result, we will focus on the lower bound of SEU valuations, and refer to them as if

they are unique. We already illustrated in Figure 2 the estimated average cost curves, marginal cost curves, and implied valuations for an individual bidder for a given search phrase. The figures illustrate how valuations are inferred from bids: the vertical line shows the expected CTR the bidder attains with the bid she places in the data, and the place where that line intersects the marginal cost curve defines the implied valuation for this bidder on this search phrase.

7.3 Empirical Results

We find empirically that the expected profits are strictly quasi-concave for each of the observed advertisements on both search phrases, which implies that the implied valuations and bids comprise an ex post Nash equilibrium in the SEU model. We formally test this by considering a grid of bids (with 100 grid points). At each point we test the hypothesis that the marginal cost for a sample of score realizations is strictly above the value per click for the bids smaller than the actual bids and strictly below the value per click for larger bids. This test did not reject the null at 5% level for all grid points and phrases (we constructed the grid such that the maximum bid guarantees that the bidder always attains the highest position, corresponding to the maximum achievable clickthrough rate).

Using our three alternative models, we recover valuations for all advertisements featured in the auctions in the selected week of data for a selected search phrase. We present our results normalizing recovered valuations and profits per click using the maximum bid for search phrase # 1 as a numeraire. We use the same normalizing factor for the values and profits per click recovered for both considered search phrases. In Table 1 we display basic statistics for valuations for the two analyzed search phrases.⁶ We notice that the search phrase #2 is the higher value phrase out of the two phrases that we analyze. However, the range of valuations remains comparable across both search phrases. Assuming that the valuations corresponding to different advertisements are stable within the period of analysis, we can compute the standard errors for the recovered values using the asymptotic formula. The standard errors of the values are fairly tight, with the average standard errors ranging from about 1-2% for the high-ranked bidders, to 3-4% for the bottom bidders.

Not surprisingly, the values computed from the NU-EFLB model in columns 1-2 of Table 1 tend to under-estimate the values recovered in the SEU framework while the values computed from the NU-EOS framework over-estimate the SEU values. We also notice that a substantial disadvantage of the NU-EOS framework is that it does not provide a meaningful upper for the bidders that are in the top position. To compute the numbers in the first two columns of Table 1 we imputed an upper bound on valuations equal to the largest bid in the sample for each search phrase. The results turn out to be sensitive to the definition of the upper bound on valuations in the NU-EOS framework and in columns 3-4 of Table

⁶We report the standard errors for the estimated values in Tables 12 and 13 in the Appendix.

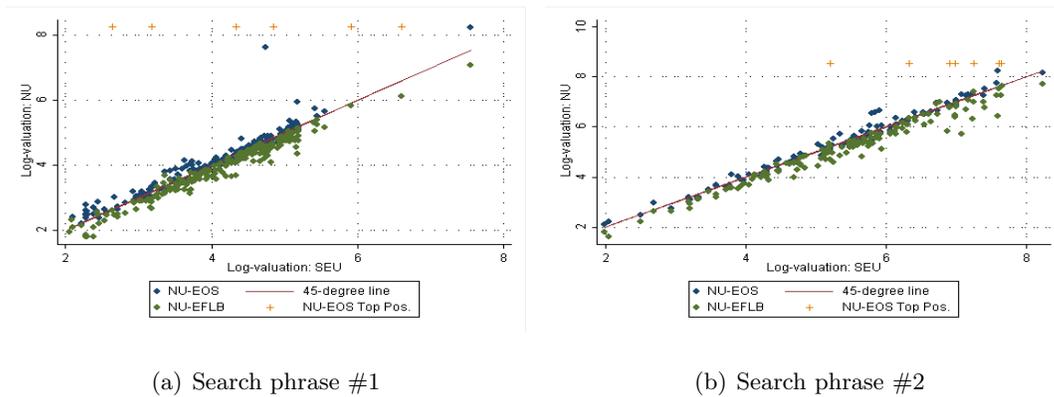
Table 1: Mean of valuations for different models across search phrases

Model; Search phrase	All bidders		Top bidders in >50% queries excluded	
	#1	#2	#1	#2
NU-EFLB	0.047	0.217	0.045	0.189
NU-EOS	0.136	0.451	0.070	0.276
SEU	0.059	0.292	0.057	0.261

We report the means of recovered valuations across search phrases and bidders for both of the models, where for the NU models we use the modified (monotone) ICC curves. The values are normalized by the highest observed bid for search phrase #1.

1 we report the mean log-values omitting the bidders who appear in the top position in more than 50% of cases. As expected, the average values for this group of bidders are smaller than overall mean values, while the average estimated values in NU-EFLB and NU-EOS frameworks are closer to the estimated values in the SEU framework.

Figure 3: Log of Valuations in NU-models Versus SEU Model (with indicated imputed NU-EOS values)



Co-location of estimated values can be represented graphically. Figure 3 displays the implied valuations for alternative models (or their bounds) for all advertisements against the implied valuations from the SEU model, in logarithmic scale, for both search phrases. NU-EFLB underestimates the values for most of the advertisements, while the NU-EOS overestimates the values. To understand why, recall first from Theorem 3 that if the NU models have an interval of valuations that is in the bounds on valuations across all queries, then those valuations will also be within the bounds for the SEU model. Since there is no such interval in our data for any advertiser, thus it is an empirical question as to how the NU model bounds will relate to the SEU valuations. We separately indicate the values where the median value was implied from the upper bound on the values for the top bidders in NU-EOS framework (this

occurs for bidders in the top position for more than half of the queries where they appear). Figure 3 demonstrates that these imputed values can induce large biases in the estimated average values.

Combining the recovered values with the data, we can compute the implied ex-post profits per click across the bidders by averaging the per-impression profit per click across different impressions. In the NU framework, the value obtained under the NU-EFLB assumption under-estimates the valuation and the value obtained under the NU-EOS assumption over-estimates the valuation; the same comparisons hold for profit per click.

We begin with the analysis of the profit per click per advertisement. Table 2 illustrates per-query profit per click relative to the average cost per click for each of the different models (the Lerner Index for the advertiser), weighting each advertisement equally. We compute each advertiser’s Lerner Index by weighting both the profit per click and the cost per click by the score and the position-specific clickthrough rate.

We can note from columns 1-4 of Table 2 that the surplus predicted by SEU and NU framework can be substantially different. The largest deviations from the surplus in SEU framework is demonstrated by the surplus evaluated from NU-EOS framework. We have recognized that this result is very sensitive to the value estimates for the top bidders (whose NU-EOS values are inferred from the upper bound on the distribution of the values). To isolate this effect, in columns 5-8 of Table 2 we report the profit PC excluding the bidders that appear in the top position in more than 50% of queries. It clear that the last four columns of Table 2 provide a more reasonable base for comparison between the frameworks, showing that NU-EOS framework tends to over-estimated the Lerner Index relative to SEU framework, while NU-EFLB tends to underestimate the Lerner Index. However, NU-EOS Lerner Indices are more spread out across advertisers. The magnitude of the Lerner Index is comparable between the considered search phrases.

8 Counterfactual experiments

8.1 Alternative Models and the Role of Uncertainty

We begin by looking at how the alternative models do in terms of predicting behavior out of sample. We proceed by taking implied valuations from each model using one week of data (taking the valuations corresponding to the mean values for the NU models), and then predicting revenue on an auction-by-auction basis in the next week of data using the same model to generate counterfactual predictions. For the advertisements which do not appear in the first week of data, we hold the counterfactual bids equal to the observed bids. Figure 4 illustrates the results, where the x-axis is the expected revenue given the

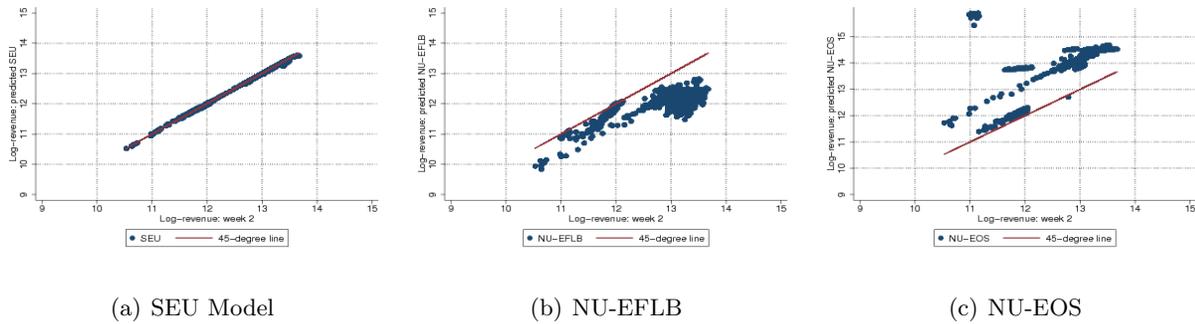
Table 2: Lerner Index $\frac{\text{Avg.}(value-CPC)}{\text{Avg.}(CPC)}$ (advertisement is unit of analysis)

Model	All bidders				Top bidders in >50% queries excluded			
	Mean	25%	50%	75%	Mean	25%	50%	75%
Search phrase #1, advertisers equally weighted								
NU-EFLB	0.143	0.020	0.090	0.219	0.134	0.019	0.086	0.217
NU-EOS	4.365	0.197	0.379	0.666	0.632	0.195	0.371	0.647
SEU	0.420	0.259	0.352	0.503	0.400	0.259	0.347	0.485
Search phrase #1, advertisers weighted by expected clicks								
NU-EFLB	0.681	0.141	0.157	1.536	0.153	0.141	0.141	0.141
NU-EOS	10.921	0.686	4.456	20.753	2.652	0.314	4.456	4.456
SEU	1.810	0.468	1.566	3.111	1.018	0.289	1.566	1.566
Search phrase #2, advertisers equally weighted								
NU-EFLB	0.289	0.076	0.176	0.449	0.266	0.071	0.161	0.384
NU-EOS	1.715	0.259	0.514	1.035	0.810	0.244	0.488	0.871
SEU	0.735	0.319	0.506	0.879	0.709	0.312	0.483	0.829
Search phrase #2, advertisers weighted by expected clicks								
NU-EFLB	0.644	0.559	0.718	0.817	0.563	0.232	0.711	0.730
NU-EOS	14.310	1.108	1.603	13.876	1.245	0.831	1.436	1.603
SEU	1.173	0.679	1.066	1.731	1.227	0.776	1.184	1.731

actual bids and prices in the auction, and the y-axis shows the predictions (or bounds on predictions) for the SEU model. Note that the SEU model provides a very good fit for the data. One reason for that is that the sample of advertisers and their bids do not change substantially from week to week, creating a similar competitive environment for the bidders in each week. Despite the fact that the realizations of score and entry uncertainty do change query to query and thus week to week, the SEU model predicts bids based on the distribution of scores and entry, so that our model predicts very similar bids for the same advertisers in the second week of data.

Details of the computational algorithm for the SEU model are given in Appendix G.

Figure 4: Log of Predicted Revenues v. Log of Actual Revenues for Search Phrase 1



Note that there are two ways to consider the NU-EOS and NU-EFLB models. If we wish to consider the predictions we can make without making any selection from the set of Nash equilibria, then when we attempt to first infer valuations and then make predictions, the bounds will be quite wide: we would want to combine the lower-bound values (NU-EFLB) with lower bound counterfactual bids (NU-EOS), and vice versa. However, the models can also be taken independently as particular selections from the set of Nash equilibria. Here, we treat them as particular selections.

To predict the revenues in the NU model we used the median of the advertiser’s valuation inferred from the corresponding equilibrium and then computed the bids from the same (upper or lower bound) equilibrium. The revenue predicted by the NU-EFLB model tends to under-state the actual revenue. On the other hand, the revenue predicted by the NU-EOS model tends to over-state the revenue. In most cases, however, when the NU bounds straddle the actual revenue, the revenue in the SEU case is within the NU bounds.

Table 3 describes the distribution of the standard deviations of the predicted revenues from the actual revenue in week 2 normalized by the mean actual revenues.

Table 3: Mean squared deviation of the predicted revenues per query from the true revenues (normalized by true mean revenue per query)

Model for equilibrium (Model for values)	Mean	25%	50%	75%
Search phrase #1				
SEU (SEU)	0.00011	0.00000	0.00001	0.00002
NU-EFLB (SEU)	0.34054	0.26660	0.38197	0.44344
NU-EOS (SEU)	0.01593	0.00229	0.00446	0.00875
NU-EFLB (NU-EFLB)	0.15941	0.10394	0.18185	0.21074
NU-EOS (NU-EOS)	0.24965	0.13932	0.22225	0.28404
NU-EFLB (NU-EOS)	5.89371	5.29678	5.51929	5.75895
NU-EOS (NU-EFLB)	0.45432	0.37673	0.51380	0.57772
Search phrase #2				
SEU (SEU)	0.00095	0.00003	0.00022	0.00108
NU-EFLB (SEU)	0.02870	0.00214	0.01135	0.03830
NU-EOS (SEU)	0.15713	0.09096	0.14865	0.21837
NU-EFLB (NU-EFLB)	0.02073	0.00262	0.00972	0.02126
NU-EOS (NU-EOS)	0.16492	0.00151	0.01211	0.05811
NU-EFLB (NU-EOS)	37.47682	0.14496	2.63420	23.88841
NU-EOS (NU-EFLB)	0.12617	0.04462	0.09855	0.17046

8.2 Competition, Elasticities and Profits Per Click

In this section, we examine the properties and implications of the estimated elasticities for the average cost curve for clicks. First, we observe that there is substantial variation across bidders and between search phrases in the elasticity of the average cost curve. Table 4 provides summary statistics on the inverse of the elasticity faced by all the advertisements, grouping bidders together by the average ranking the advertisements received. The table also shows the gaps between value and bid, and between bid and payment, each normalized by the bid, for bidders in each category (recall that the Lerner Index $\frac{\text{Value}-\text{CPC}}{\text{CPC}}$ will be equal to the inverse of the elasticity). The gap between bid and payment is large and implies that the bids substantially exceed the payment. The bid tends to be between .55 and .90 of the value for all positions, typically around .70 of the value on search phrase #2. The results for the first position for search phrase #1 are skewed by an unusually high value estimate for a bidder who dominates the first position.

8.3 Auction Design: Comparing Vickrey Auctions and the Generalized Second Price Auction

In a model without uncertainty, EOS and Varian have shown that the EOS equilibrium implements the same allocation and the same prices as a Vickrey auction. Thus, the choice of auction design does not matter. However, once real-world uncertainty is incorporated, this equivalence breaks down. If the auctioneer held a separate Vickrey auction for each user query, it would be optimal for each advertiser to bid its value, even if the same bid was to be applied across many different user queries. If we take the quality scores calculated for each impression as the best estimate of the efficient scores (that is, efficiency requires ads to be ranked according to the product of value and quality score), then the ads will always be ranked efficiently, query by query, in the Vickrey auction, even as quality scores change over time.

In contrast, in the GSP auction used in practice, if bids apply to many queries (as in the SEU model) and scores and entry vary across queries, then different bidders will have different gaps between their bids and values. This implies that the ads will not be ranked efficiently in many cases. Therefore, the GSP auction is strictly less efficient than the Vickrey auction, so long as there is sufficient uncertainty in the environment.

Table 5 shows the results of a counterfactual comparison of the two mechanisms. We used the values estimated in the SEU model, and computed counterfactual equilibria in each auction format: Vickrey and GSP auction.⁷ To simplify, we ignored reserve prices, which were rarely binding in any case. Note

⁷We report the standard errors for welfare and revenue in Table 16 in the Appendix. As the test for the difference of

Table 4: Characteristics of competition

Avg. ranking	Mean	Mean	Mean	Inverse Elasticity= $\frac{\text{Value-CPC}}{\text{CPC}}$			
	$\frac{\text{Value-Bid}}{\text{Value}}$	$\frac{\text{Value-Bid}}{\text{CPC}}$	$\frac{\text{Bid-CPC}}{\text{CPC}}$	Mean	25%	50%	75%
Search phrase #1							
[1, 1.5)	0.445	1.665	1.030	2.695	2.769	2.769	2.769
[1.5, 2.5)	0.160	0.336	0.779	1.115	1.197	1.197	1.197
[2.5, 4)	0.099	0.144	0.288	0.432	0.372	0.372	0.517
[4, 5.5)	0.097	0.143	0.254	0.397	0.303	0.385	0.421
[5.5, 8)	0.131	0.189	0.239	0.428	0.377	0.421	0.449
Search phrase #2							
[1, 1.5)	0.267	0.522	0.396	0.918	0.710	0.871	0.871
[1.5, 2.5)	0.354	0.758	0.365	1.123	1.048	1.048	1.210
[2.5, 4)	0.393	0.967	0.410	1.377	1.040	1.230	1.833
[4, 5.5)	0.308	0.599	0.302	0.901	0.654	1.029	1.063
[5.5, 8)	0.224	0.374	0.269	0.643	0.544	0.605	0.669

We report mean elasticities of the MC curve from the SEU model corresponding to bidders whose average position is in the displayed bracket. The statistics are first computed at the advertisement level, and then weighted by the frequency of user queries in which they appear in the relevant position buckets.

that the Vickrey auction gives the same results as if the NU-EOS model is used, since in a world where bidders change their bids to play the NU-EOS equilibrium in each query, the allocation and prices are the same as Vickrey prices. The SEU model equilibrium gives the outcome of the generalized second price auction under uncertainty.

We see that the Vickrey auction always gives higher efficiency, which is necessarily the case, and the efficiency differences are very small (.04% and .3%). It is perhaps surprising that the efficiency differences are so small despite the asymmetric markups of the bidders. The result follows because most of the deviations from efficiency occur when the score-weighted values are similar (the realization of the shock to the score is high for the bidder with the lower average score-weighted value). Further, in each search phrase the gaps between the score-weighted bids are large enough in the top two positions that re-rankings are relatively rare (an inefficient bidder appears in the first position only 3% of the time for search phrase 1, and 16% where the efficiency gap between the top two bidders is smaller).

The revenue comparison between the mechanisms is theoretically ambiguous, so it is an empirical question as to which one performs better. We see that for search phrase 1, the revenue difference is much larger in magnitude than the efficiency difference, and in the opposite direction: the Vickrey auction raises 22% less revenue. Vickrey performs worse in the first position, while revenue is very close in the middle positions and 25% higher in the lowest positions in the Vickrey auction. The GSP raises more revenue in the first position because the gap between bid and value is substantially smaller for the second-highest bidder than the top one on this search phrase, making the second-highest bidder a more aggressive competitor and driving the price for the top bidder well above the Vickrey price. Since most of the clicks come from the top position, this leads to a large revenue gap.

For search phrase 2, the Vickrey auction raises 1.2% more revenue than GSP. Vickrey raises more revenue in the top and bottom positions, and less in the middle. For this search phrase, middle ranked bidders have larger gaps between bids and values than the top bidder in the GSP, depressing the revenue extracted from the top bidder.

Summarizing, the structural model allows us to compare auction designs in a scenario where theory is ambiguous about the revenue comparison. Our estimates show that the efficiency gains from a Vickrey auction are small for the search phrases we study, and that the revenue comparison will likely vary from search phrase to search phrase. Thus, further research is required to assess the best choice for the platform as a whole from a revenue perspective, while from an efficiency perspective, Vickrey auctions appear to offer the potential for only modest improvements.

revenues and welfare between the NU-EOS vs SEU model shows, the difference in welfare is insignificant at the 1% level between the two models, while the difference in revenues is significant as seen in Table 5.

Table 5: Predicted revenues and welfare for the SEU generalized second price auction model versus the Vickrey auction, using SEU values and actual bidder configurations: with decomposition by bidders in different positions

Model (values)	Positions			
	All	1	2-5	6-8
Search phrase #1				
Revenue SEU (SEU)	2.490	1.583	0.886	0.021
Revenue Vickrey=NU-EOS(SEU)	1.947***	1.032***	0.890**	0.026
Welfare SEU (SEU)	8.228	6.358	1.838	0.032
Welfare Vickrey=NU-EOS (SEU)	8.231	6.365	1.834	0.032
Search phrase #2				
Revenue SEU (SEU)	3.114	1.846	1.203	0.066
Revenue Vickrey=NU-EOS(SEU)	3.153***	1.910***	1.152***	0.091
Welfare SEU (SEU)	6.513	3.433	2.957	0.122
Welfare Vickrey=NU-EOS (SEU)	6.531	3.459	2.950	0.122

This table represents the expected total per query revenue of the auction platform and the expected per query social welfare. For all counterfactuals, we used SEU estimated values. To compute the expected clicks we used our estimated position clickthrough rates and the scores of the advertisers. Values and cost per click are normalized by the maximum bid for the search phrase # 1. To group the SEU welfare and SEU revenue by positions we used the positions based on SEU equilibrium bids. To group the EOS welfare and EOS revenue by positions we used the positions based on EOS equilibrium bids. The table also represents the results of the hypothesis test for the difference in expected total per query revenue of the auction platform and the expected per query social welfare in NU-EOS model versus SEU model, using asymptotic standard errors . * - corresponds to significance at 10% level, ** - significance at 5% level, and *** - significance at 1% level.

8.4 The Role of Quality Score for Welfare, Profits and Revenue of the Search Engine

Next we present the results of counterfactual experiments showing the impact of the quality scores used for ranking on the social welfare, profits of the bidders and revenue of the search engine. As motivated in the introduction, theory suggests a possible tradeoff between the efficiency of quality scores and the revenue extracted.

We consider two alternatives for the quality scores, both motivated by real-world examples. Researchers from Yahoo! [Lahaie and Pennock \(2007\)](#) proposed using an advertiser’s clickability raised to the power θ , with $\theta < 1$, and this approach came to be known in the industry as “squashing.” The second alternative we refer to as “coarsening,” or using a less accurate estimator for “clickability.”

In each case, we use a similar methodology. We conduct the experiments on synthetic data, for both the baseline case and the counterfactual scenarios of squashing and coarsening. Using synthetic data allows us to eliminate some of the least frequent advertisers in order to reduce the computational burden and reduce the variance of our predictions. It also makes it less likely that variance in bids introduced by the computation creates an asymmetric effect on one regime or the other. Using synthetic data also has the advantage that we can incorporate the identity of the highest losing bidder (this was unknown in the actual data) and predict how those bids change when quality scores change.

More precisely, we create a sample of simulated bidder configurations, where each configuration contains 9 advertisers. Each of the 9 bidders is randomly drawn without replacement from the sample of actual advertisers with the probability equal to the proportion of actual queries where this advertiser was observed. To construct the scores of advertisers we compute the mean log-scores of advertisers in the actual configurations and then generate the shocks from the empirical distribution of shocks to the scores via random sampling with replacement.

In our analysis we assume that the scores in the data correspond to the true advertiser-specific “clickability.” We use these clickability estimates in combination with our estimated position discounts to predict the clicks an advertiser receives in each position. When we consider counterfactual quality score models in the auction, we continue to use the original “clickability” estimates and position discounts to predict clicks that an advertiser receives once it is assigned a position.

8.4.1 Score Squashing

We begin by considering “squashing.” “Squashed” scores are generated as $\tilde{s}_i = s_i^\theta$ with $\theta = \frac{1}{2}$. [Table 6](#) compares the outcomes (advertiser profit, platform revenue, and welfare) with and without squashing. The table also highlights the role of bid changes to the results: it shows the impact of each score regime

when bids are taken as the equilibrium bids from the other score regime.⁸

The welfare effects of squashing are relatively small for these search phrases: a loss of .5% for search phrase 1 and 4.5% for search phrase 2. The small loss for the first search phrase is due to the large gap between score-weighted values between the top bidder and lower-ranked bidders. For the first search phrase, revenue goes down by 2%, due to the fact that the highest-ranked bidder’s original quality score is lower than that of the bidder below. Squashing thus lowers the price per click for the top bidder (which is equal to $\frac{b_2 \cdot s_2}{s_1}$, since squashing reduces $\frac{s_2}{s_1}$). For the second search phrase, revenue increases by 9%; for this search phrase, the most common bidder in position 1 has a higher score than the most frequent position 2 bidder, and squashing increases $\frac{s_2}{s_1}$ and thus price per click.

Interestingly, the advertiser bidding response to the change in quality score regimes affects the sign of the revenue effect for search phrase 2. Without the advertiser response, squashing hurts revenue, but with it, revenue increases substantially. For search phrase 1, revenue falls even farther after bids adjust.⁹

8.4.2 Score Coarsening

In practice, quality scores are produced by statistical algorithms. A choice for the platform is how accurate the algorithm should attempt to be. Another choice is how tailored the scores should be to individual users. To highlight the tradeoff, suppose that the search query is for clothing and that there are two clothing advertisers, one that sells men’s clothes and one that sells women’s clothes, and that each is equally appealing to its own target gender and has the same value per click, but generates no clicks from the other gender. Suppose that there is only one advertising slot. Suppose that the search advertising platform observes the gender of the user. One algorithm (a “coarsened” algorithm) might ignore that information, and give each advertiser the same score. This yields only half of the efficient level of clicks, but the price per click is equal to the value per click. A second algorithm uses the gender information. But then, the score-weighted bid of second-place store is always zero. The algorithm is more efficient (double the clicks), but there is no competition.

⁸In practice, changes in score regimes are often tested in experiments where a small fraction of users experiences an alternative score regime. Advertiser do not notice individual, small experiments and so their bids do not adjust (and indeed, they often take some time to adjust). Thus, the table can be used to predict what would happen in an experiment where the score regime is changed but advertiser bids are fixed.

⁹Computed welfare, profit and revenue changes in Table 6 are significant for the second search phrase on 1% level. For the search phrase #1 the welfare change is insignificant overall and for the top position on 1% level, while the revenue change is significant. We report the standard errors in Table 14 in the Appendix.

Table 6: Counterfactual effect of score squashing on social welfare, auction revenue, and advertiser profits decomposed by positions

Outcome	Scenario	Positions			
		All	1	2–5	6–8
Search phrase #1					
<i>Welfare</i>	Before Squashing	6.956	5.371	1.556	0.030
	After Squashing	6.921	5.350	1.541**	0.030***
	Original Scores, Bids After Squashing	6.958	5.377	1.552	0.029
	After Squashing, Original Bids	6.907	5.360	1.517***	0.030***
<i>Revenue</i>	Before Squashing	2.091	1.210	0.861	0.020
	After Squashing	2.049***	1.060***	0.966***	0.023***
	Original Scores, Bids After Squashing	2.100	1.178***	0.901***	0.021***
	After Squashing, Original Bids	2.052***	1.120***	0.910***	0.022***
<i>Advertiser Profit</i>	Before Squashing	4.865	4.161	0.694	0.010
	After Squashing	4.872	4.290	0.575***	0.007***
	Original Scores, Bids After Squashing	4.858	4.198	0.651***	0.009***
	After Squashing, Original Bids	4.855	4.240	0.608***	0.008***
Search phrase #2					
<i>Welfare</i>	Before Squashing	6.781	3.747	2.905	0.129
	After Squashing	6.476***	3.354***	2.989***	0.132***
	Original Scores, Bids After Squashing	6.784	3.758	2.90	0.130**
	After Squashing, Original Bids	6.410***	3.325***	2.951***	0.134***
<i>Revenue</i>	Before Squashing	3.223	1.799	1.354	0.069
	After Squashing	3.513***	1.773***	1.654***	0.086***
	Original Scores, Bids After Squashing	3.509***	1.919***	1.508***	0.082***
	After Squashing, Original Bids	3.211**	1.674***	1.466***	0.071***
<i>Advertiser Profit</i>	Before Squashing	3.559	1.948	1.550	0.060
	After Squashing	2.962***	1.581***	1.335***	0.046***
	Original Scores, Bids After Squashing	3.275***	1.840***	1.389***	0.047***
	After Squashing, Original Bids	3.200***	1.651***	1.486***	0.063***

This table represents the expected total per query revenue of the auction platform and the expected per query social welfare. To compute the expected revenue we used the estimated position clickthrough rates and the scores of the advertisers. To compute the expected welfare we used the values estimated from the SEU model and the estimated position clickthrough rates and scores of the advertisers. Values and cost per click are normalized by the maximum bid for the search phrase # 1. Position ranges are determined by the SEU bids in each scenario. The table also represents the results of the hypothesis test for the difference in expected profit of advertisers, total per query revenue of the auction platform and the expected per query social welfare relative to those quantities before the squashing, using asymptotic standard errors. * - corresponds to significance at 10% level, ** - significance at 5% level, and *** - significance at 1% level.

We choose to model “coarsening” rather than increasing accuracy, since we have in our data the best estimate of the true clickability of each advertiser, and can use that to predict clicks; we then add noise to the observed scores to simulate coarsening. More precisely, we assume that the shock to the mean quality score has two components: $\log \varepsilon_i = \log \varepsilon_{1,i} + \log \varepsilon_{2,i}$. We assume that the component $\log \varepsilon_{2,i}$ has variance equal to half of the variance of $\log \varepsilon_i$. For example, if a regression model were used to generate scores as a function of observed user characteristics, each component of the shock to scores might be a linear function (with the coefficient determined by a regression) of a user characteristic (e.g. $\log \varepsilon_{1,i} = x_{1,i}\beta_1$ and $\log \varepsilon_{2,i} = x_{2,i}\beta_2$). For coarsening, we remove one of the characteristics from the model, which corresponds to setting the score shock equal to $\varepsilon_{1,i}$ (assuming that the underlying user characteristics are independent).

Table 7 gives the results. As a result of coarsening, welfare goes down by 1-2%. Revenue is affected more substantially: it goes up 2% for the first search phrase and 18% for the second. Bid adjustments are a crucial part of the revenue gain. For the first search phrase, the revenue actually goes down after coarsening with bids fixed, but goes up once bids adjust. For the second search phrase, the revenue gains are close to zero when bids are fixed, but substantial once bids adjust. This means that the benefits of coarsening for revenue would be missed in short-term experiments that didn’t incorporate advertiser behavior.¹⁰ Structural models like this one can be used to predict the impact when longer-term experiments on advertisers are too time-consuming or expensive.

9 Conclusion

In this paper we develop and estimate a new model of advertiser behavior under uncertainty in the sponsored search advertising auctions. Unlike the existing models which assume that bids are customized for a single user query, we model the fact that queries arrive more quickly than advertisers can change their bids, and advertisers cannot perfectly predict quality scores. We present theoretical characterizations of existence and uniqueness, and propose a computational algorithm for computing equilibria given primitives of the model. We develop an estimator for bidder valuations, establish its properties, and apply it to historical data from Microsoft’s search advertising auctions. Our model yields lower implied valuations and bidder profits than approaches that ignore uncertainty.

The empirical application provides insight into the economics of search advertising auctions. We find that bidders have substantial profits per click, even on some of the industry’s most competitive search phrases: bidder values are typically 40% to 90% higher than their cost per click, even on very

¹⁰Computed welfare and revenue changes in Table 7 are significant for the second search phrase on 1% level. For the search phrase #1 the welfare and advertiser profit change are insignificant overall and for the top position on 1% level, while the revenue change is significant. We report the standard errors in Table 15 in the Appendix.

Table 7: Counterfactual effect of score coarsening on social welfare, auction revenue, and advertiser profits decomposed by positions

Outcome	Scenario	Positions			
		All	1	2-5	6-8
Search phrase #1					
<i>Welfare</i>	Before Coarsening	6.956	5.371	1.555	0.030
	After Coarsening	6.903	5.363	1.510***	0.031***
	Original Scores, Bids After Coarsening	6.954	5.362	1.563	0.030
	After Coarsening, Original Bids	6.903	5.362	1.510***	0.031***
<i>Revenue</i>	Before Coarsening	2.091	1.209	0.862	0.020
	After Coarsening	2.138***	1.251***	0.865	0.022***
	Original Scores, Bids After Coarsening	2.187***	1.278***	0.888***	0.020**
	After Coarsening, Original Bids	2.036***	1.174***	0.841***	0.022***
<i>Advertiser Profit</i>	Before Coarsening	4.865	4.162	0.693	0.010
	After Coarsening	4.765	4.112	0.644***	0.009***
	Original Scores, Bids After Coarsening	4.768	4.084	0.674***	0.009***
	After Coarsening, Original Bids	4.867	4.188	0.670***	0.009***
Search phrase #2					
<i>Welfare</i>	Before Coarsening	6.780	3.746	2.905	0.129
	After Coarsening	6.665***	3.594***	2.939***	0.132***
	Original Scores, Bids After Coarsening	6.732*	3.680***	2.922*	0.130**
	After Coarsening, Original Bids	6.712 **	3.656***	2.925**	0.131***
<i>Revenue</i>	Before Coarsening	3.223	1.799	1.355	0.069
	After Coarsening	3.800***	2.147***	1.576***	0.077***
	Original Scores, Bids After Coarsening	3.785***	2.150***	1.561***	0.075***
	After Coarsening, Original Bids	3.232	1.792*	1.369***	0.071***
<i>Profit</i>	Before Coarsening	3.557	1.946	1.550	0.060
	After Coarsening	2.865***	1.447***	1.363***	0.055***
	Original Scores, Bids After Coarsening	2.947***	1.531***	1.362***	0.055***
	After Coarsening, Original Bids	3.480***	1.863***	1.556	0.060

This table represents the expected total per query revenue of the auction platform and the expected per query social welfare. To compute the expected revenue we used the estimated position clickthrough rates and the scores of the advertisers. To compute the expected welfare we used the values estimated from the SEU model and the estimated position clickthrough rates and scores of the advertisers. Values and cost per click are normalized by the maximum bid for the search phrase # 1. Position ranges are determined by the SEU bids in each scenario. The table also represents the results of the hypothesis test for the difference in expected profit of advertisers, total per query revenue of the auction platform and the expected per query social welfare relative to those quantities before the score coarsening, using asymptotic standard errors. * - corresponds to significance at 10% level, ** - significance at 5% level, and *** - significance at 1% level.

competitive search phrases. Further, profits vary with the equilibrium position on the page, and on the search phrases we study, they appear to be higher in the top positions for a few high-value advertisers. Profits are determined by the rate at which clicks decline with the position on the page, and by the dispersion of competing bidder values around a given bidder’s value.

We find that bidders have substantial strategic incentives to reduce their expressed demand in order to reduce the unit prices they pay in the auctions, and these incentives are asymmetric across bidders, leading to inefficient allocation. We quantify the inefficiency as being very small (less than 0.3%). We show that a Vickrey auction would eliminate the inefficiency, but the impact of switching to a Vickrey auction for revenue is ambiguous and can be much larger. Since equilibrium Vickrey bids (equal to bidder values) are estimated to be much higher than GSP bids, unless a mechanism was devised to transition GSP bids automatically, there could be substantial transition costs from changing the mechanism as well.

We also find that search advertising platforms might have substantial incentives to modify quality scoring algorithms to extract more revenue, which can be accomplished at modest efficiency costs. Finally, advertiser bid changes in response to changes in quality scoring algorithms can change the direction of the revenue impact. This means that the advertising platform may need to rely on longer-term experiments on advertisers or structural models to evaluate the tradeoffs, since the industry-standard short-term experiments on a small fraction of users will not necessarily guide policy in the desired direction.

References

- Agarwal, N. and Athey, S. and Yang, D. (2009): “Skewed Bidding in Pay-per-Action Auctions for Online Advertising”, *American Economic Review: Paper and Proceedings*, 441–447.
- Athey, S. and Ellison, G. (2011): “Position auctions with consumer search,” *Quarterly Journal of Economics*.
- Athey, S., and P. Haile (2002): “Identification of standard auction models”, *Econometrica*, 2107—2140.
- Borgers, T. and Cox, I. and Pesendorfer, M. and Petricek, V. (2007): “Equilibrium bids in sponsored search auctions: Theory and evidence”, Working paper.
- Boltyanskii, V., R. Gamkrelidze, and L. Pontryagin (1960): “Theory of optimal processes. I. The maximum principle”, *Izv. Akad. Nauk SSSR Ser. Mat.*, 24(1), 3—42.
- Chernozhukov, V., H. Hong, and E. Tamer (2007): “Parameter set inference in a class of econometric models”, *Econometrica*, 75, 1243—1284.

- Edelman, B., M. Ostrovsky, and M. Schwarz (2007): “Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords,” *American Economic Review*, 97(1), 242—259.
- Edelman, B., and M. Ostrovsky (2007): “Strategic bidder behavior in sponsored search auctions”, *Decision support systems*, 43(1), 192—198.
- Ghose, A., and S. Yang. (2009): “An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets”, *Management Science*, 55 (10), 1605—1622.
- Haile, P., and E. Tamer (2003): “Inference with an incomplete model of English auctions”, *Journal of Political Economy*, 111(1), 1—51.
- Hendricks, K. and Paarsch, H.J. (1995): “A survey of recent empirical work concerning auctions”, *The Canadian Journal of Economics*, 28(2), 403—426.
- Hashimoto, T. (2009): “Equilibrium Selection and Inefficiency in Internet Advertising Auctions”, SSRN Working paper.
- Hong, H., Mahajan, A., and D. Nekipelov (2010): “Extremum Estimation and Numerical Derivatives”, Stanford and UC Berkeley Working paper.
- Imbens, G., and C. Manski (2004): “Confidence intervals for partially identified parameters”, *Econometrica*, 1845—1857.
- Jeziorski, P. and I. Segal, “What Makes them Click: Empirical Analysis of Consumer Demand for Internet Search Advertising,” working paper, Stanford University, February 2009.
- Kosorok, M. (2008): “Introduction to empirical processes and semiparametric inference”, Springer.
- Lahaie, S., and D. Pennock (2007): “Revenue Analysis of a Family of Ranking Rules for Keyword Auctions”, *ACM Conference on Electronic Commerce*, 50-56.
- McFadden, D., and W. Newey (1994): “Large Sample Estimation and Hypothesis Testing”, *Handbook of econometrics*, 4, 2111—2245.
- Mela, C.F., and Yao, S. (2009): “A dynamic model of sponsored search advertising”, SSRN Working paper.
- Myerson, R. (1981): “Optimal Auction Design”, *Mathematics of Operations Research*, 6(1), 58-73.
- Paarsch, H.J. (1997): “Deriving an estimate of the optimal reserve price: an application to British Columbian timber sales”, *Journal of Econometrics*, 78(2), 333—357.
- Pontryagin, L. (1966): “On the theory of differential games”, *Russian Mathematical Surveys*, 21(4), 193—246.

- Pollard, D. (1990) “Empirical processes: theory and applications”, *Institute of Mathematical Statistic*.
- Varian, H. (2007): “Position auctions”, *International Journal of Industrial Organization*, 25(6), 1163—1178.

Appendix

A Proof of Theorem 1

Throughout the proof, we abuse notation by writing $\frac{\partial}{\partial b_j} EU_i(\beta_i(\tau), \tau\beta_{-i}(\tau), \bar{s})$ for $\frac{\partial}{\partial b_j} EU_i(b_i, \tau b_{-i}, \bar{s}) \Big|_{b=\beta(\tau)}$.

Before proceeding with the proof we prove the following lemma.

LEMMA 3. $TE_i(\cdot, \bar{s})$, $Q_i(\cdot, \bar{s})$, and $D_0(\cdot, \bar{s})$ are differentiable.

Proof:

Considered functions are represented by the integrals over the product of a finite number of indicator functions $\mathbf{1}\{\cdot > 0\}$ and the differentiable density function $f_\epsilon(\cdot)$. Therefore, they are also differentiable.

Q.E.D.

Now we start with proving parts (i) and (iii). First, we prove the sufficiency of these conditions. Thus, we start with the assumption:

$$\text{For some } \delta > 0 \text{ with } \tau \in [1 - \delta, 1], \text{ there exists a (unique) solution } \beta(\tau) \text{ to equation (3.8)} \quad (\text{A.14})$$

Lemma 1 (in Section 3.4) establishes that (3.8) holds at $\tau = 1$ if and only if the bidders' first-order conditions (3.5) hold for $b = \beta(1)$.

Since Theorem 1 maintains the assumption of quasi-concavity of the objective function, a solution to (3.5) will correspond to the maximum of the payoff function. This means that there will exist an equilibrium in the considered auction b^* corresponding to $\beta(1)$ as required for sufficiency.

Second, we prove necessity. Under our maintained assumption of quasi-concavity, we need to show that (A.14) holds, whenever there exists (a unique) b^* solving the system of the first-order conditions

$$\frac{\partial EU_i(b_i^*, b_{-i}^*, \bar{s})}{\partial b_i} = 0.$$

We prove this result in the following lemma.

LEMMA 4. *There exists $\delta > 0$ such that the system of nonlinear equations (3.8) has a (unique) solution in the interval $\tau \in [1 - \delta, 1]$*

Proof:

We show that (3.8) has solutions on $\tau \in [1 - \delta, 1]$ with the initial condition $\beta(1) = b^*$ for some $\delta > 0$ in three steps.

Step 1: We write an alternative expression for the objective function by constructing a modified total expenditure function $\widetilde{TE}_i(\cdot)$. Equation (3.7) can be expressed in the form

$$\frac{\partial}{\partial b'_{-(I+1)}} EU(b, \bar{s}) b_{-(I+1)} + \frac{\partial}{\partial b_{I+1}} EU(b, \bar{s}) b_{I+1} = -TE(b, \bar{s}). \quad (\text{A.15})$$

As a result, we can re-write (3.7) as

$$\frac{d}{d\tau} EU_i(b_i, \tau b_{-i, I+1}, b_{I+1} \bar{s})|_{\tau=1} = -TE_i(b, \bar{s}) - b_{I+1} \frac{\partial}{\partial b_{I+1}} EU_i(b, \bar{s}). \quad (\text{A.16})$$

Define

$$\widetilde{TE}_i(b_{-(I+1)}, b_{I+1}, \bar{s}) = TE_i(b, \bar{s}) + b_{I+1} \frac{\partial}{\partial b_{I+1}} EU_i(b, \bar{s}, r).$$

Then we can re-write (A.15) as

$$\frac{d}{d\tau} EU_i(b_i, \tau b_{-i, I+1}, b_{I+1} \bar{s})|_{\tau=1} = -\widetilde{TE}_i(b_{-(I+1)}, b_{I+1} \bar{s}).$$

Step 2: Using the definition of \widetilde{TE}_i , we can re-cast equation (3.8) defining $\beta(\tau)$ as

$$\tau \frac{d}{d\tau} EU_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s}) = -\widetilde{TE}_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s}), \quad \text{for all } i = 1, \dots, I, \quad (\text{A.17})$$

by moving the term corresponding to $\frac{\partial EU_i}{\partial b_{I+1}}$ to the right-hand side and taking into account the constraint $b_{I+1} = r$.

Using the chain rule we find that

$$\begin{aligned} \tau \frac{d}{d\tau} EU_i(\beta_i(\tau), \tau \beta_{-i, (I+1)}(\tau), r, \bar{s}) &= \sum_j \left(\left(1 + \frac{1-\tau}{\tau} \mathbf{1}_{j=i} \right) \frac{\partial}{\partial b_j} EU_i(\beta_i(\tau), \tau \beta_{-i, (I+1)}(\tau), r, \bar{s}) \dot{\beta}_j \right. \\ &\quad \left. + \frac{\partial}{\partial b_j} EU_i(\beta_i(\tau), \tau \beta_{-i, (I+1)}(\tau), r, \bar{s}) \beta_j(\tau) \right), \end{aligned} \quad (\text{A.18})$$

where $\dot{\beta}_j$ stands for $\frac{d\beta_j(\tau)}{d\tau}$.

Using expression (A.18) we can transform equation (A.17) to

$$\begin{aligned} \sum_{j \neq I+1} \left(1 + \frac{1-\tau}{\tau} \mathbf{1}_{j=i}\right) \frac{\partial}{\partial b_j} EU_i(\beta_i(\tau), \tau \beta_{-i, (I+1)}(\tau), r, \bar{s}) \dot{\beta}_j &= \frac{\widetilde{TE}_i(\beta_i(\tau), \tau \beta_{-i, (I+1)}(\tau), b_{I+1}, \bar{s})}{\tau} \\ &- \sum_{j \neq i, I+1} \frac{\partial}{\partial b_j} EU_i(\beta_i(\tau), \tau \beta_{-i, (I+1)}(\tau), r, \bar{s}) \beta_j(\tau). \end{aligned} \quad (\text{A.19})$$

We can treat this expression as a system of differential equations defining the function $\beta(\cdot)$ with the initial condition $\beta(1) = b^*$.

Step 3: We finalize the proof by establishing the existence of the solution of (A.19) by using the following result.

LEMMA 5. *Let $T = [0, 1]$ and suppose that functions $Z : \mathbb{R}^{n \times n} \times T \mapsto \mathbb{R}^{n \times n}$ and $Y : \mathbb{R}^{n \times n} \times T \mapsto \mathbb{R}^{n \times n}$ are fixed. Suppose that $Z(M, t)$ is non-singular $\forall M \in \mathbb{R}^{n \times n}$ and $t \in [0, 1]$ except for a finite number of points and both $Y(M, t)$ and $Z(M, t)$ are continuous in t and Lipschitz-continuous in M . Then the system of equations*

$$Z(M, t) \dot{M} = Y(M, t) M \quad (\text{A.20})$$

with the boundary condition $M(1) = I_{n \times n}$ (identity matrix) has a unique non-singular solution.

The proof of this lemma can be found in [Boltyanskii et al. \(1960\)](#) and [Pontryagin \(1966\)](#).

By Lemma 3 for $\tau \in [1 - \delta, 1]$ functions $\frac{\widetilde{TE}_i(b_i, \tau b_{-i}, \bar{s})}{\tau}$ and $D_0(b, \bar{s})$ are differentiable at b^* and, therefore, Lipschitz-continuous at that point. Given that matrix D_0 is non-singular and Lipschitz-continuous at b^* , there exists some neighborhood $U_{2\delta}(b^*)$ of b^* of diameter 2δ where this matrix is non-singular. Then the matrix with ij -element

$$\left(1 + \frac{1-\tau}{\tau} \mathbf{1}_{j=i}\right) \frac{\partial}{\partial b_j} EU_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s})$$

is non-singular for $\beta(\tau) \in U_\delta(b^*)$ and $1 - \delta \leq \tau \leq 1$. Thus, making the change of variables $t = (\tau + \delta - 1)/\delta$ and defining $M(t) = \beta(1 - \delta + \delta t)/b^*$ if we let $Z_{ij}(M, t) = \left(1 + \frac{1-\tau}{\tau} \mathbf{1}_{j=i}\right) \frac{\partial}{\partial b_j} EU_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s})$ and

$$Y(M, t) = \frac{\widetilde{TE}_i(\beta_i(\tau), \tau \beta_{-i, (I+1)}(\tau), b_{I+1}, \bar{s})}{\tau} - \sum_{j \neq i, I+1} \frac{\partial}{\partial b_j} EU_i(\beta_i(\tau), \tau \beta_{-i, (I+1)}(\tau), r, \bar{s}) \beta_j(\tau)$$

we note that (A.19) can be re-written as (A.20) and all conditions of Lemma 5 are satisfied. Therefore, the existence of an equilibrium b^* leads to the existence of the function $\beta(\cdot)$ that for some $\tau \in [1 - \delta, 1]$ is the solution of (3.8) with the initial condition b^* . We note that if there were multiple equilibria, each would lead to a system of differential equations with different starting point b^* . If the equilibrium is unique, there will be only one such system of differential equations.

This concludes the proof of Lemma 4 and of parts (i) and (iii).

Now we proceed with proving (ii) and (iv) and establish the result for the global existence of the solution to (3.8) under stronger conditions for the payoff functions. Parts (ii) and (iv) assume that $D_0(b, \bar{s})$ is Lipschitz for all $b \in V$ and non-singular. We need to show that for each τ there exist bid vectors $\beta(\tau)$ which solve the system (3.8), which will transform to the system of equilibrium first-order conditions for $\tau = 1$. Theorem 1 requires that the reserve price is strictly positive, which implies that at $b_i = 0$, $TE_i(b, \bar{s}) = 0$. However $\frac{\partial}{\partial b_i} EU_i(b_i, \tau b_{-i}, \bar{s}) \neq 0$ at $\tau = 0$ for $b_i \downarrow 0$. Since $\beta_i(\tau)$ has to solve (3.8) for $\tau \downarrow 0$, then $\beta_i(0) \neq 0$. Then defining $\lim_{\tau \downarrow 0} \beta_i(\tau) = \bar{b}_i$, we can find \bar{b}_i as a solution to

$$\widetilde{TE}_i(\beta_i(0), 0, r, \bar{s}) = 0.$$

Following our arguments for parts (i) and (iii), the solution exists (is unique) iff the system of differential equations (A.19) has a (unique) solution on $(0, 1]$ with initial condition $\beta(0) = \bar{b}$. We note that the only ‘‘problematic’’ point is $\tau = 0$ because the matrix of coefficients for $\dot{\beta}(\tau)$ in system (A.19) can potentially become singular in the vicinity of $\tau = 0$.

First of all, provided that $D_0(b, \bar{s})$ is Lipschitz, then so is $\frac{\partial \widetilde{TE}_i(b_i, b_{-i, (I+1)}, b_{I+1}, \bar{s})}{\partial b'_i}$. Then note that by L’Hospital’s rule,

$$\lim_{\tau \rightarrow 0} \frac{\widetilde{TE}_i(\beta_i(\tau), \tau \beta_{-i, (I+1)}(\tau), b_{I+1}, \bar{s})}{\tau} = \frac{\partial \widetilde{TE}_i(\bar{b}_i, 0, b_{I+1}, \bar{s})}{\partial b_j} \dot{\beta}_i + \sum_{i \neq j} \frac{\partial \widetilde{TE}_i(\bar{b}_i, 0, b_{I+1}, \bar{s})}{\partial b_j} \bar{b}_j. \quad (\text{A.21})$$

Provided that limit (A.21) exists, the right-hand side of (A.19) is continuous from the right at $\tau = 0$.

As a result, we established the following. First, note that in (3.8) we define the vector of bids as a function of parameters τ . This means that we can represent the given system of differential equations (A.19) with initial condition $\beta(0) = \bar{b}$. The set of bids satisfying the first-order condition will correspond to the solution of this equation $\beta(\tau)$ when $\tau = 1$.

Second, we just verified that system of equations (A.19) satisfies condition of Lemma 5 for $t = 1 - \tau$ and

$M(t) = \beta(1 - t)/\bar{b}$. As a result, we can apply Lemma 5 which establishes the sufficient condition for the uniqueness of the solution and proves the results (ii) and (iv) in Theorem 1.

B The Impact of Vanishing Uncertainty on Bidding and Identification

To gain some further intuition for how a model with uncertainty differs from the NU model, consider some limiting cases that are close to the NU model, where a small amount of uncertainty is added that serves as a refinement to the set NU equilibria. (In the empirical application, uncertainty is not small, so this exercise is intended to build intuition only.) First, consider what we call the random entry refinement. Suppose that there is no score uncertainty, but that with probability ϕ , a new advertiser enters with a random bid, and the distribution of the advertiser's score-weighted bid has full support over the relevant region. This is a realistic model of a new entrant or a new advertiser: the initial scores assigned by the system will not stay constant, and an advertiser may appear with a number of different score-weighted bids, each with low probability.

Now consider taking the limit as ϕ approaches zero. Then, taking into account that the entry of the random bidder affects marginal incentives only when it ties with the bidders score-weighted bid, it will be optimal for each advertiser to submit a bid that is an ex post equilibrium in the NU model, and in addition, where the bidder is indifferent between her current position when paying exactly her bid, or taking the next-lower position and paying the bid of the next-lowest bidder. Formally, the equilibrium conditions are the original equilibrium conditions (3.1), plus

$$s_{k_j} v_{k_j} \geq \frac{s_{k_{j+1}} b_{k_{j+1}} \alpha_{j+1} - s_{k_{j+2}} b_{k_{j+2}} \alpha_{j+2}}{\alpha_{j+1} - \alpha_{j+2}} = s_{k_{j+1}} v_{k_{j+1}},$$

except for the lowest-ranked bidder who bids her valuation. This contrasts with the [Edelman et al. \(2007\)](#) refinement, that satisfies

$$s_{k_j} v_{k_j} \geq \frac{s_{k_{j+1}} b_{k_{j+1}} \alpha_j - s_{k_{j+2}} b_{k_{j+2}} \alpha_{j+1}}{\alpha_j - \alpha_{j+1}} = s_{k_{j+1}} v_{k_{j+1}}.$$

The random entry strategies are envy-free if and only if $\alpha_j/\alpha_{j-1} \leq \alpha_{j+1}/\alpha_j$ for all $1 < j < J$ and the equilibrium is monotone. However, in general the random entry equilibrium may not exist in pure strategies. Intuitively, the auction has a ‘‘first-price’’ flavor: with some probability, each bidder pays her bid. Then, two bidders with similar score-weighted valuations will also place similar score-weighted bids; but when an opponent's bid is too close, a bidder's best response may be to drop down a position and take a lower price. This in turn might induce the opponent to change her bid, leading to cycling.

It is somewhat more subtle to consider the effects of small amounts of score uncertainty. We provide some intuition for a special case. Assume that $v_1 s_1 > v_2 s_2 > v_3 s_3$, and suppose there are two slots. Assume that \tilde{s}_2 is the stochastic score for bidder 2, and that the scores of the other bidders are fixed at their means. Let f_{1/\tilde{s}_2} be the PDF of $1/\tilde{s}_2$. The local indifference condition defining the optimal bid b_2 (given the bids b_1, b_3) is

$$\alpha_2(v_2 - b_2)f_{1/\tilde{s}_2}\left(\frac{b_3 s_3}{b_2}\right) + \left[\alpha_1(v_2 - b_2) - \alpha_2\left(v_2 - \frac{b_3 s_3 b_2}{b_1 s_1}\right)\right] f_{1/\tilde{s}_2}\left(\frac{b_1 s_1}{b_2}\right) = 0 \quad (\text{B.22})$$

Suppose for a moment that $f_{1/\tilde{s}_2}\left(\frac{b_3 s_3}{b_2}\right) = 0$, so bidder 2 is not at risk for dropping a position. If $\gamma_2^* = \frac{b_1 \gamma_1}{b_2}$ is the critical value of the quality score that makes bidder 2 tie for the top position, the indifference condition reduces to

$$\alpha_1(v_2 - b_2) = \alpha_2\left(v_2 - \frac{b_3 \gamma_3}{\gamma_2^*}\right),$$

which is the EOS condition in the contingency where bidder 2 is tied with bidder 1. In contrast, if $f_{1/\tilde{s}_2}\left(\frac{b_1 s_1}{b_2}\right) = 0$ (no chance of moving up a position), the bidder is always better off by increasing her bid until $b_2 = v_2$, for standard reasons: the bid only matters if it causes the bidder to go from losing to winning. So, a small amount of quality score uncertainty puts upward pressure on bids if a bid is far from moving up to the next position, and we should generally expect to see the lowest position bidder place bids in a region where the bidder has some chance of moving up.¹¹

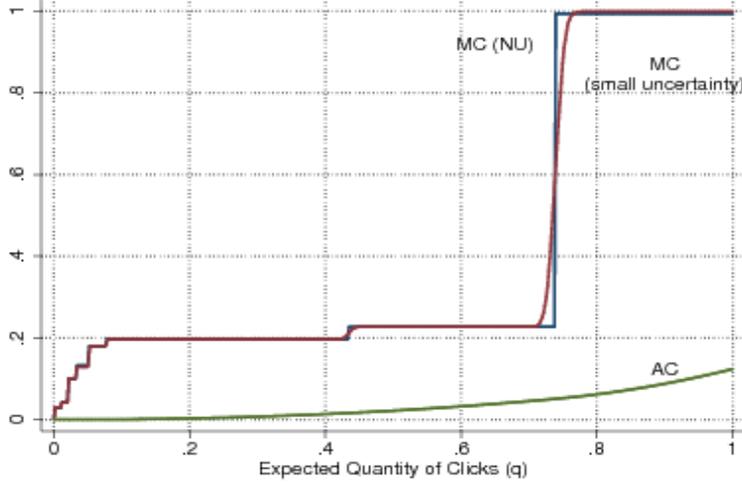
We can also consider a refinement where the bidders face uncertainty, but the probability of a change in score or configuration is very small. Figure 5 below shows an effect of the small noise on the marginal and total cost. We use the actual bid and score data from a top configuration in a particular market. In this picture we assume that the score has a distribution with a mass point in the mean score. The sample for computation is generated by picking the score equal to the mean with probability $1 - \varepsilon$ and equal to a random draw from the empirical distribution of scores with probability ε .

C Proof of Theorem 4

To analyze the properties of the estimate for valuation we use the fact that the empirical profit function converges in probability to the population expected payoff function uniformly in valuation and the bid. Moreover, by our assumption regarding the distribution of the score, the score has a continuous density with a finite support. This implies that the numerical derivative will converge to the true derivative for the population analog of the

¹¹A similar result has been independently obtained in Hashimoto (2009).

Figure 5: Marginal cost and total cost curves for bidder in a frequent configuration



considered functions. In particular, using Taylor's expansion and assuming that considered functions are twice differentiable with a Lipschitz-continuous residual of the second-order Taylor's expansion we can write:

$$\begin{aligned} & \frac{-TE_i(b_i - 2\tau_N, b_{-i}, \bar{s}) + 8TE_i(b_i - \tau_N, b_{-i}, \bar{s}) - 8TE_i(b_i + \tau_N, b_{-i}, \bar{s}) + TE_i(b_i + 2\tau_N, b_{-i}, \bar{s})}{-Q_i(b_i - 2\tau_N, b_{-i}, \bar{s}) + 8Q_i(b_i - \tau_N, b_{-i}, \bar{s}) - 8Q_i(b_i + \tau_N, b_{-i}, \bar{s}) + Q_i(b_i + 2\tau_N, b_{-i}, \bar{s})} \\ &= \frac{TE'_i(b_i, b_{-i}, \bar{s}) + L_1 \tau_N^3}{Q'_i(b_i, b_{-i}, \bar{s}) + L_2 \tau_N^3} = \frac{TE'_i(b_i, b_{-i}, \bar{s})}{Q'_i(b_i, b_{-i}, \bar{s})} + L_1 \tau_N^3 + L_2 \tau_N^3 + o(\tau_N^3), \end{aligned}$$

where L_1 and L_2 are Lipschitz constants. Next we consider the difference:

$$\begin{aligned} \hat{v}_i - v_i &= \frac{-\widehat{TE}_i(b_i - 2\tau_N, b_{-i}, \bar{s}) + 8\widehat{TE}_i(b_i - \tau_N, b_{-i}, \bar{s}) - 8\widehat{TE}_i(b_i + \tau_N, b_{-i}, \bar{s}) + \widehat{TE}_i(b_i + 2\tau_N, b_{-i}, \bar{s})}{-\widehat{Q}_i(b_i - 2\tau_N, b_{-i}, \bar{s}) + 8\widehat{Q}_i(b_i - \tau_N, b_{-i}, \bar{s}) - 8\widehat{Q}_i(b_i + \tau_N, b_{-i}, \bar{s}) + \widehat{Q}_i(b_i + 2\tau_N, b_{-i}, \bar{s})} - \frac{TE'_i(b_i, b_{-i}, \bar{s})}{Q'_i(b_i, b_{-i}, \bar{s})} \\ &= D_1 + D_2 + D_3 + o_p\left(\frac{1}{\sqrt{T} \tau_T}\right). \end{aligned}$$

Here we use the following decomposition:

$$\begin{aligned} D_1 &= \frac{18}{Q'_i(b_i, b_{-i}, \bar{s})} \left[\widehat{TE}_i(b_i, b_{-i}, \bar{s}) - TE_i(b_i, b_{-i}, \bar{s}) \right], \\ D_2 &= -\frac{18 TE'_i(b_i, b_{-i}, \bar{s})}{(Q'_i(b_i, b_{-i}, \bar{s}))^2} \left[\widehat{Q}_i(b_i, b_{-i}, \bar{s}) - Q_i(b_i, b_{-i}, \bar{s}) \right], \end{aligned}$$

and

$$D_3 = \frac{-TE_i(b_i - 2\tau_N, b_{-i}, \bar{s}) + 8TE_i(b_i - \tau_N, b_{-i}, \bar{s}) - 8TE_i(b_i + \tau_N, b_{-i}, \bar{s}) + TE_i(b_i + 2\tau_N, b_{-i}, \bar{s})}{-Q_i(b_i - 2\tau_N, b_{-i}, \bar{s}) + 8Q_i(b_i - \tau_N, b_{-i}, \bar{s}) - 8Q_i(b_i + \tau_N, b_{-i}, \bar{s}) + Q_i(b_i + 2\tau_N, b_{-i}, \bar{s})} - \frac{TE'_i(b_i, b_{-i}, \bar{s})}{Q'_i(b_i, b_{-i}, \bar{s})}.$$

We omitted all the terms of the smaller order than $o_p((T\tau_T)^{-1/2})$ using the assumption regarding the rate of the numerical differentiation. We can use the result in [Pollard \(1990\)](#) to argue that

$$\sup_{|\delta| < \epsilon} \frac{1}{\sqrt{T^*}} \sum_{t^*} [u_i(v_i, b_i + \delta; b_{-i}, \bar{s}_i, \hat{\varepsilon}_{it^*}, C_{t^*}) - u_i(v_i, b_i + \tau_T; b_{-i}, \bar{s}_i, \hat{\varepsilon}_{it^*}, C_{t^*})] = o_p(\sqrt{\delta}).$$

Finally, using the structure of total expenditure and expected quantity of clicks, we can write:

$$\sqrt{T\tau_T}(\hat{v}_i - v_i) = -18 \frac{1}{Q'_i(b_i, b_{-i}, \bar{s})} \frac{1}{\sqrt{T}} \sum_t \frac{u_i(v_i, b_i + \tau_T; b_{-i}, \bar{s}_i, \hat{\varepsilon}_{it}, C_t) - u_i(v_i, b_i - \tau_T; b_{-i}, \bar{s}_i, \hat{\varepsilon}_{it}, C_t)}{\sqrt{\tau_T}},$$

Then if $\Omega = \text{Var} \left(\frac{u_i(v_i, b_i + \tau_T; b_{-i}, \bar{s}_i, \hat{\varepsilon}_{it}, C_t) - u_i(v_i, b_i - \tau_T; b_{-i}, \bar{s}_i, \hat{\varepsilon}_{it}, C_t)}{\sqrt{\tau_T}} \right)$, it follows that the and i.i.d. Assumption 2, bootstrap is valid by [Kosorok \(2008\)](#) and

$$\sqrt{T\tau_T}(\hat{v}_i - v_i) \xrightarrow{d} N \left(0, \frac{324 \Omega}{(Q'_i(b_i, b_{-i}, \bar{s}))^2} \right)$$

D Estimation of valuations in case of set-valued best response correspondences

Even though we can consistently estimate the payoff of the bidder for each valuation and the score, there is no guarantee that for each bid there will be a single valuation which makes this bid consistent with the first-order condition. General results for set inference in the auction settings have been developed for instance in [Haile and Tamer \(2003\)](#), while general results for identification in the auction settings are given in [Athey and Haile \(2002\)](#). This result will display most likely in the situation where score-weighted bids have limited overlap, i.e. for a fixed set of bids we can find positions such that some bidders will never have their ads displayed in these positions. In this case local bid changes may not affect the payoff as they will not affect the relative ranking of the bidders. If $b_k \bar{s}_k \bar{\varepsilon} < b_i \bar{s}_i \underline{\varepsilon}$, then the score-weighted bid of bidder k will always be below the bid of bidder i . Similarly, if $b_k \bar{s}_k \underline{\varepsilon} > b_i \bar{s}_i \bar{\varepsilon}$ then the bid of bidder k will always be ranked higher than the bid of bidder i . In the extreme case where for each pair of bidders j and k we have

$$(b_k \bar{s}_k \underline{\varepsilon} - b_j \bar{s}_j \bar{\varepsilon})(b_j \bar{s}_j \underline{\varepsilon} - b_k \bar{s}_k \bar{\varepsilon}) > 0$$

(i.e. the ranked bids never overlap), then the model substantially simplifies. Assume that the bids are ordered by their ranks using the mean scores: $b_j \bar{s}_j > b_{j-1} \bar{s}_{j-1}$. Also assume that $\pi = 0$ so that all bidders are always present in the auction. A selected bidder will be placed in position k and pay $b_k \bar{s}_k E[s^{-1}]$ if $b_k \frac{\bar{s}_k \bar{\varepsilon}}{\bar{s}_i \underline{\varepsilon}} < b < b_{k-1} \frac{\bar{s}_{k-1} \underline{\varepsilon}}{\bar{s}_i \bar{\varepsilon}}$. If the bid is $b_k \frac{\bar{s}_k \underline{\varepsilon}}{\bar{s}_i \underline{\varepsilon}} < b < b_k \frac{\bar{s}_k \bar{\varepsilon}}{\bar{s}_i \underline{\varepsilon}}$ or $b_k \frac{\bar{s}_k \underline{\varepsilon}}{\bar{s}_i \bar{\varepsilon}} < b < b_k \frac{\bar{s}_k \bar{\varepsilon}}{\bar{s}_i \bar{\varepsilon}}$, then the probability of being placed in position k is

$$\int F_\varepsilon \left(\frac{bs}{b_k \bar{s}_k} \right) f_\varepsilon \left(\frac{s}{\bar{s}_i} \right) ds,$$

and the expected payment is

$$\int \int \mathbf{1}\{b_k s' < bs\} \frac{b_k s'}{s} f_\varepsilon \left(\frac{s'}{\bar{s}_k} \right) f_\varepsilon \left(\frac{s}{\bar{s}_i} \right) ds ds'.$$

Similarly if $b_{k-1} \frac{\bar{s}_{k-1}\underline{\varepsilon}}{\bar{s}_i\underline{\varepsilon}} < b < b_{k-1} \frac{\bar{s}_{k-1}\bar{\varepsilon}}{\bar{s}_i\bar{\varepsilon}}$ or $b_{k-1} \frac{\bar{s}_{k-1}\underline{\varepsilon}}{\bar{s}_i\bar{\varepsilon}} < b < b_{k-1} \frac{\bar{s}_{k-1}\bar{\varepsilon}}{\bar{s}_i\underline{\varepsilon}}$, then the probability of being placed in position k is

$$\int \left(1 - F_\varepsilon\left(\frac{bs}{b_{k-1}\bar{s}_{k-1}}\right)\right) f_\varepsilon\left(\frac{s}{\bar{s}_i}\right) ds,$$

and the expected payment is

$$\int \int \mathbf{1}\{b_{k-1}s' > bs\} \frac{b_k \bar{s}_k}{s} f_\varepsilon\left(\frac{s'}{\bar{s}_{k-1}}\right) f_\varepsilon\left(\frac{s}{\bar{s}_i}\right) ds ds'.$$

Then the objective function of the bidder i will be not strictly monotone. It will have “flat spots” where there is no bid overlap and it will be smooth where score-weighted bids overlap. We can explicitly compute the marginal utility from bidding b as

$$\frac{\partial}{\partial b} \mathbb{E}_{\varepsilon, C} [u_i(v_i, b_i = b, b_{-i}; \varepsilon_{it}, C_t^i)] = \begin{cases} 0, & \text{if } b_k \frac{\bar{s}_k\bar{\varepsilon}}{\bar{s}_i\underline{\varepsilon}} < b < b_{k-1} \frac{\bar{s}_{k-1}\underline{\varepsilon}}{\bar{s}_i\bar{\varepsilon}}, \\ \bar{\alpha}_k \int \left(\frac{v_i s}{b_k} - b\right) f_\varepsilon\left(\frac{s}{\bar{s}_i}\right) f_\varepsilon\left(\frac{bs}{b_k \bar{s}_k}\right) \\ - \bar{\alpha}_{k+1} \int \left(\frac{v_i s}{b_k} - \frac{b_{k+1} \bar{s}_{k+1}}{s}\right) f_\varepsilon\left(\frac{s}{\bar{s}_i}\right) f_\varepsilon\left(\frac{bs}{b_k \bar{s}_k}\right) ds, & \\ & \text{if } b_k \frac{\bar{s}_k\underline{\varepsilon}}{\bar{s}_i} < b < b_k \frac{\bar{s}_k\bar{\varepsilon}}{\bar{s}_i}. \end{cases}$$

In the limited overlap case the numerical algorithm for computation of the best responses will contain 3 steps.

- Step 1 Compute $\frac{\partial}{\partial b} E [u_i(v_i, b_i = b, b_{-i}, \varepsilon_{it}, C_t^i)]$ at each of $4(N-1)$ points $b_k \frac{\bar{s}_k(\times/\div)\underline{\varepsilon}}{\bar{s}_i(\times/\div)\underline{\varepsilon}}$
- Step 2 If for some k there are 2 points out of 4 $b_k \frac{\bar{s}_k(\times/\div)\underline{\varepsilon}}{\bar{s}_i(\times/\div)\underline{\varepsilon}}$ where the marginal utility has different signs, solve the non-linear equation

$$\bar{\alpha}_k \int \left(\frac{v_i s}{b_k} - b\right) f_\varepsilon(s - \bar{s}_i) f_\varepsilon\left(\frac{bs}{b_k \bar{s}_k}\right) - \bar{\alpha}_{k+1} \int \left(\frac{v_i s}{b_k} - \frac{b_{k+1} \bar{s}_{k+1}}{s}\right) f_\varepsilon\left(\frac{s}{\bar{s}_i}\right) f_\varepsilon\left(\frac{bs}{b_k \bar{s}_k}\right) ds = 0.$$

Obtain solution b^* .

- Step 3 Compare $\bar{\alpha}_k (v_i - \bar{s}_k b_k E[s_{it}^{-1}])$ for all k and $E[u_i(v_i, b_i = b^*, b_{-i}, \varepsilon_{it}, C_t^i)]$ where the latter were computed. If the maximum value is $\bar{\alpha}_k (v_i - \bar{s}_k b_k E[s_{it}^{-1}])$, then the best response is set valued with $b \in [b_k \frac{\bar{s}_k\bar{\varepsilon}}{\bar{s}_i\underline{\varepsilon}}, b_{k-1} \frac{\bar{s}_{k-1}\underline{\varepsilon}}{\bar{s}_i\bar{\varepsilon}}]$. Otherwise, the best response is unique and equal to b^* .

To recover valuations in case of limited overlap of the score-ranked bids, we fix the set of observed bids. We also fix the grid which contains the support of valuations. Then for each bidder and each value on the grid we solve for the set of best responses. Given the produced set of best responses we pick the set of valuations for which the

set of best responses contains the actually observed best response. Technically this implies that we recover the set:

$$S_i = \left\{ (b, v) \mid b \in \text{BR}_i(v, b_{-i}), v \in \mathcal{V} \right\}.$$

The estimated valuation is the cut of this set such that

$$(\widehat{v}_i, \bar{b}_i) \in S_i,$$

where \bar{b}_i is observed in the data.

The structure of our empirical procedure allows us to formulate the following result.

THEOREM 5. *Under Assumption 2 the estimation procedure following the outlined steps 1-3 is numerically equivalent to the statistics inversion procedure in Chernozhukov et al. (2007). As a result, the estimates of identified set of valuations will be described by Theorem 2.1 in Chernozhukov et al. (2007).*

To provide the argument, we consider the following scheme.

1. Consider the sample of all observed bidder configurations over queries $t \{C_t\}_{t=1}^T$ where T is the total number of queries. Uniformly over these sets draw a set C_{t^*} . Select a particular bidder i Construct a set $C_{t^*}^i = C_{t^*} \setminus \{i\}$. In total we construct T^* subsamples of collections of sets of configurations.
2. For a fixed position j make K^* random subsamples $\{C_{t^*,k,j-1}^i\}_{k=1}^{K^*}$ of $j-1$ bidders out of set C_{t^*} . The number of subsamples K^* needs to grow such that $K^*/\sqrt{T} \rightarrow \infty$. For configuration C_{t^*} compute the payoff of bidder i from being placed in position j

$$\begin{aligned} u_{t^*,k}^{i,j}(b_i, v_i) &= \bar{\alpha}_j \sum_{k=1}^{K^*} \int \int \left(v_i F_s \left(\frac{s'}{\bar{s}_k} \right) \right. \\ &\times \prod_{m \in C_{t^*,k,j-1}^i} \left(\frac{1 - F_s \left(\frac{sb}{\bar{s}_m b_m} \right)}{F_s \left(\frac{sb}{\bar{s}_m b_m} \right)} \right) \prod_{n \in C_{t^*}^i} F_s \left(\frac{sb}{\bar{s}_n b_n} \right) \\ &\quad - \sum_{k \in C_{t^*}^i \setminus C_{t^*,k,j-1}^i} \frac{b_k s'}{s} \mathbf{1}_{\{b_k s' < b s\}} \frac{F_s \left(\frac{s}{\bar{s}_i} \right)}{F_s \left(\frac{s' b_k}{\bar{s}_i b} \right)} \\ &\times \prod_{m \in C_{t^*,k,j-1}^i} \left(\frac{1 - F_s \left(\frac{sb}{\bar{s}_m b_m} \right)}{F_s \left(\frac{s' b_k}{\bar{s}_m b_m} \right)} \right) \prod_{n \in C_{t^*}^i} F_s \left(\frac{s' b_k}{\bar{s}_n b_n} \right) \Big) d \log F_s \left(\frac{s}{\bar{s}_i} \right) d \log F_s \left(\frac{s'}{\bar{s}_k} \right). \end{aligned}$$

If we use T^* draws of configurations of bidders in the first stage, and K^* draws in the second stage, we need to compute the approximated payoff by rescaling as

$$\widehat{EU}_i(b_i = b; b_{-i}, \bar{s}) = \sum_{j=1}^J \frac{1}{T^*} \sum_{t=1}^{T^*} \frac{\binom{\#C_{t^*}^i}{j}}{K^*} \sum_{k=1}^{K^*} u_{t^*,k}^{i,j}(b_i, v_i).$$

This procedure allows us to evaluate the payoff function of a single bidder using $T^* \times K^*$ total draws. Note that we can “recycle” the draws of sets of configurations to compute the payoff functions for different bidders. We then can compute the numerical derivative

$$\frac{\partial}{\partial b} \widehat{EU}_i(b_i = b; b_{-i}, \bar{s}) = \sum_{j=1}^J \frac{1}{T^*} \sum_{t=1}^{T^*} \frac{\binom{\#C_{t^*}^i}{j}}{K^*} \sum_{k=1}^{K^*} \frac{u_{t^*,k}^{i,j}(b + \tau, v_i) - u_{t^*,k}^{i,j}(b - \tau, v_i)}{2\tau}.$$

Given the assumption that bidders set their bids optimally, we can write the condition

$$\frac{\partial}{\partial b} \widehat{EU}_i(b_i = \bar{b}_i, b_{-i}) = \sum_{j=1}^J \frac{1}{T^*} \sum_{t^*} \frac{\binom{\#C_{t^*}^i}{j}}{K^*} \sum_{k=1}^{K^*} \frac{u_{t^*,k}^{i,j}(\bar{b}_i + \tau, v_i) - u_{t^*,k}^{i,j}(\bar{b}_i - \tau, v_i)}{2\tau} = o_p(1),$$

at the observed bid. Then we can recover the set of values that correspond to the observable bid. To do so we form the grid over v and minimize

$$\left(\sum_{j=1}^J \frac{1}{T^*} \sum_{t^*} \frac{\binom{\#C_{t^*}^i}{j}}{K^*} \sum_{k=1}^{K^*} \frac{u_{t^*,k}^{i,j}(\bar{b}_i + \tau, v_i) - u_{t^*,k}^{i,j}(\bar{b}_i - \tau, v_i)}{2\tau} \right)^2,$$

with respect to v . The set of minimizers will deliver the identified set of valuations $\widehat{\mathcal{F}}_{v,T^*,J}$. This procedure allows estimation similar to that offered in [Chernozhukov et al. \(2007\)](#). The confidence sets can be recovered using the tools developed in [Imbens and Manski \(2004\)](#).

E Algorithm and description of Monte-Carlo Simulations

In the Monte-Carlo simulations we analyze the stability of our estimation procedure with respect to the sampling noise in the data as well as the width of the support of valuations. The first set of Monte-Carlo simulations was designed to analyze the robustness of the suggested computational procedure to the sampling noise in the observed configurations of advertisers. The setup of the Monte-Carlo simulation was the following. We considered the case where there are 5 advertisers competing for 2 slots. The click-through rates of these slots were fixed at levels 1 and 0.5. The valuations have support on $[0, 1]$ and the scores for all advertisers are uniformly distributed on $[0, .1]$. We consider the cases where the reserve price was equal to 0.1, 0.2 and 0.3. We use the same probability of a binding budget constraint for all bidders. This probability was selected at the levels 0, 0.01, and 0.05. We used

Table 8: Results of Monte-Carlo Analysis (no binding budget constraints)

	Profits					Valuations				
Player#	1	2	3	4	5	1	2	3	4	5
Sample size =500										
	.654	.622	.788	.501	.714	.220	.124	.150	.221	.250
Sample size =1000										
	.311	.355	.330	.341	.318	.110	.098	.101	.118	.106
Sample size =2000										
	.122	.110	.114	.164	.142	.055	.068	.060	.071	.062

2000 Monte-Carlo replications. Each iteration was organized in the following way. First, we sample valuations for each bidder from $U[0,1]$. Second, for the set of valuations we computed the equilibrium of the model. In case of the uniform distribution of the scores, the problem of computing the equilibrium is equivalent to solving a system of polynomial equations (of order 4 for 5 players) with linear constraints. Then for each bidder we generated uniform random variables and removed the bidders for whom the uniform draw was below the probability of a binding budget constraint. Then we fixed the bids and generated each set of Monte-Carlo draws using the algorithm

- Using uniform draws, remove bidders with binding budget constraint
- Record equilibrium bids for remaining bidders
- Generate scores for the bidders from the uniform distribution
- Allocate bidders to slots and compute the prices

We used three setups where each Monte-Carlo sample had 500, 1000 and 2000 individual draws. For each sample we computed the payoff function, and computed the valuations of the participating bidders by inverting the first-order condition. In the table below we report our results. We report standard deviations of the difference between exact and estimated profits for players from 1 to 5 and the standard deviations for recovered valuations for players from 1 to 5. The following table reports the estimates for the case where the probability of players dropping out due to budget constraints is zero.

Table 9: Results of Monte-Carlo Analysis (probability of reaching the budget constraint 1%)

Player#	Profits					Valuations				
	1	2	3	4	5	1	2	3	4	5
Sample size =500										
	1.034	.1.507	1.142	.980	1.450	.320	.215	.345	.318	.343
Sample size =1000										
	.890	1.079	1.120	.760	1.235	.250	.201	.305	.285	.299
Sample size =2000										
	.530	.511	.595	.544	.645	.176	.129	.201	.148	.187

Table 10: Results of Monte-Carlo Analysis (probability of reaching the budget constraint 5%)

Player#	Profits					Valuations				
	1	2	3	4	5	1	2	3	4	5
Sample size =500										
	2.003	3.790	3.202	2.254	2.990	.269	.235	.130	.021	.189
Sample size =1000										
	1.840	1.089	2.044	2.011	2.940	.336	.218	.238	.299	.201
Sample size =2000										
	1.188	1.112	2.230	1.970	1.450	.096	.128	.130	.199	.160

This table shows a significant decline in the standard errors of estimation when the Monte-Carlo sample size increases. This supports the formal argument of consistency of our estimation procedure.

F Recovering distributions of scores and clickthrough rates from the data

Now we will provide a more formal argument for identification of the CTR. First, we consider identification of the distribution of noise in the click-through rates, and subsequently, the distribution of estimated click-through rates. The distribution of the estimated advertiser-specific rate is denoted $F_{\gamma,i}(\cdot|z)$ and the distribution of

the estimated slot-specific click-through rate is denoted $F_{\alpha,j}(\cdot | z)$. The distribution of bidder valuations is also a common knowledge among bidders. The following proposition establishes the fact that we can recover distributions of the bidder-specific and the slot-specific CTR from observable frequencies of clicks $G_{ij}(\cdot)$ for bidder i in slot j .

THEOREM 6. *Assume that the distribution of the estimated slot-specific CTR is degenerate at α in slot 1 (where α is a known constant), and the distribution of the noise in the advertiser-specific CTR $F_\gamma(\cdot)$ is the same across advertisers. Moreover, assume that the noise in the estimated slot-specific CTR ε_j^α is independent from the noise in the estimated advertiser-specific CTR ε_i^γ for all advertisers and all slots. Then both the distribution of advertiser-specific CTR and the distribution of slot-specific CTR $F_{\alpha,j}(\cdot)$ for all slots j are identified.*

Proof:

Given that $G_{c,i,j}(x) = E[\mathbf{1}\{C_{ij} < x\}]$, then for slot 1

$$G_{c,i,1}(x) = E[\mathbf{1}\{\alpha\Gamma_i < x\}] = F_\gamma\left(\frac{x}{\alpha}\right),$$

meaning that the distribution of Γ_i is identified. Denote the distribution of $\log C_{ij}$ by $G_{c,i,j}^l(\cdot)$ and the distribution of $\log A_j$ and $\log \Gamma_i$ by $F_{\alpha,i}^l$ and F_γ^l correspondingly. Then the density of the logarithm of the CTR is expressed through the density of slot-specific CTR and advertiser-specific CTR by the convolution formula

$$g_{c,i,j}^l(x) = \int_{\log \underline{\gamma}}^{\log \bar{\gamma}} f_\gamma^l(\gamma) f_{\alpha,j}^l(x - \gamma) d\gamma.$$

Then the characteristic function for the distribution of A_j can be expressed using deconvolution

$$\chi_{\alpha,j}^l(t) = \frac{\chi_{c,i,j}^l(t)}{\chi_\gamma^l(t)}.$$

The characteristic function is computed as

$$\chi_\gamma^l(t) = \int_{-\infty}^{+\infty} e^{itx} f_\gamma^l(x) dx,$$

where $i = \sqrt{-1}$. Then we can recover the distribution of slot-specific CTR for slot j using the inverse Fourier transformation

$$F_{\alpha,j}(x) = \int_{-\infty}^{\log x} dz \int_{-\infty}^{+\infty} e^{-itz} \chi_{\alpha,j}^l(t) dt.$$

As a result, for each slot $j = 1, \dots, J$ starting from the second one we can find the distribution of its slot-specific conversion rate.

Q.E.D.

G Computing equilibria via numerical continuation

For $\tau \in [0, 1]$ the system (3.8) can be re-written as

$$\sum_{j \neq i} \frac{\partial EU_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s})}{\partial b_j} \tau b_j(\tau) = -TE_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s}), \quad i = 1, \dots, N. \quad (\text{G.23})$$

If the payoff function is twice continuously differentiable and the equilibrium existence conditions are satisfied, then $\beta(\tau)$ is a smooth function of τ . As a result, we can further differentiate both sides of this expression with respect to τ . For the left-hand side we can obtain

$$\begin{aligned} & \sum_{j, k \neq i} \frac{\partial^2 EU_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s})}{\partial b_j \partial b_k} \left[\tau^2 b_j \dot{b}_k + \tau b_j b_k \right] + \frac{\partial^2 EU_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s})}{\partial b_j \partial b_i} \tau b_j \dot{b}_i \\ & + \sum_{j \neq i} \frac{\partial EU_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s})}{\partial b_j} \left[\tau \dot{b}_j + b_j \right], \end{aligned} \quad (\text{G.24})$$

where $\dot{b} = \frac{db}{d\tau}$. Then using the notation δ_{kj} for the Kronecker symbol, we can re-write the expression of interest as

$$\sum_k a_k^i \dot{b}_k = c^i, \quad (\text{G.25})$$

and

$$\begin{aligned} a_k^i &= \left[\tau^2 (1 - \delta_{ik}) + \tau \delta_{ik} \right] \sum_{j \neq i} \frac{\partial^2 EU_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s})}{\partial b_j \partial b_k} b_j b_k + \tau (1 - \delta_{ik}) v_i \frac{\partial Q_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s})}{\partial b_k} b_k \\ & + \delta_{ik} \frac{\partial TE_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s})}{\partial b_i} b_i \end{aligned}$$

and

$$c^i = - \sum_k \tau (1 - \delta_{ik}) \sum_{j \neq i} \frac{\partial^2 EU_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s})}{\partial b_j \partial b_k} b_j b_k + (1 - \delta_{ik}) v_i \frac{\partial Q_i(\beta_i(\tau), \tau \beta_{-i}(\tau), \bar{s})}{\partial b_k} b_k$$

We make an inverse transformation and express the system of equations of interest in the form

$$A(\mathbf{b}, \tau) \dot{\mathbf{b}} = c(\mathbf{b}, \tau),$$

where the elements of matrix $A(\mathbf{b}, \tau)$ can be computed as $A_{ik}(\mathbf{b}, \tau) = a_k^i$. We know that the original system of non-linear equations has the solution $\beta(0) = 0$ corresponding to the point $\tau = 0$. We solve the problem by constructing a grid over $\tau \in [0, 1]$ and choosing the tolerance level Δ accordingly to the step of the grid. The set of grid point is $\{\tau_N\}_{t=1}^T$ where $\Delta = \max_{t=2, \dots, T} \|\tau_N - \tau_{t-1}\|$. The solution at each grid point τ_N will be a vector of bids b_t . Then we can use the modified Euler integration scheme to compute the solution on the extended interval. We can note that the system of differential equation has a singularity of order one at the origin. We use a simple regularization scheme which allows us to avoid the singularity at a cost of an additional approximation error of order Δ^α , where α is the power such that $\lim_{\delta \rightarrow +0} \delta^{-\alpha} \frac{\partial^2 EU_i(b_i, b_{-i})}{\partial b_i \partial b_j} \Big|_{\|b\|=\delta} < \infty$ for all i . Note that this condition is

satisfied if the Hessian matrix of the payoff function is non-degenerate at the origin. We initialize the system at $b_0 = \Delta/4$ and make a preliminary inverse Euler step by solving

$$\mathbf{b}_{1/2} = b_0 + A(\mathbf{b}_{1/2}, \Delta/2)^{-1} c(\mathbf{b}_{1/2}, \Delta/2) \Delta/2 \quad (\text{G.26})$$

with respect to $\mathbf{b}_{1/2}$. Such an inverse step enhances the stability of the algorithm and it will be the most time-consuming part. Then the algorithm proceeds from step t to step $t + 1$ in the steps of $1/2$. Suppose that \mathbf{b}_t is the solution at step t . Then we make a preliminary Euler step

$$\mathbf{b}_{t+1/2} = \mathbf{b}_t + \frac{\Delta}{2} A(\mathbf{b}_t, \tau_N)^{-1} c(\mathbf{b}_t, \tau_N). \quad (\text{G.27})$$

Then using this preliminary solution we make the final step

$$\mathbf{b}_{t+1} = \mathbf{b}_t + \Delta A\left(\mathbf{b}_{t+1/2}, \tau_N + \frac{1}{2}\Delta\right)^{-1} c\left(\mathbf{b}_{t+1/2}, \tau_N + \frac{1}{2}\Delta\right).$$

Note that the values that are updated only influence the evaluated derivative, while the final step size is still equal to Δ . We can use standard numerical derivative approximation to compute the elements of $A(\mathbf{b}, \tau)$ and $c(b, \tau)$. For the first derivative we use the third-order formula such that

$$\frac{\partial EU_i(\mathbf{b}, \tau, \bar{s})}{\partial b_j} = \frac{EU_i(b_j - 2\delta, b_{-j}, \tau, \bar{s}) - 8EU_i(b_j - \delta, b_{-j}, \tau, \bar{s}) + 8EU_i(b_j + \delta, b_{-j}, \tau, \bar{s}) - EU_i(b_j + 2\delta, b_{-j}, \tau, \bar{s})}{12\delta} + o(\delta^5),$$

where δ is the step size in the domain of bids¹². For the second cross-derivatives we can use the “diamond” formula

$$\begin{aligned} \frac{\partial^2 EU_i(\mathbf{b}, \tau)}{\partial b_j \partial b_k} = & \frac{1}{12\delta^2} \left[EU_i(b_j - 2\delta, b_{-j}, \tau, \bar{s}) - EU_i(b_k - 2\delta, b_{-k}, \tau, \bar{s}) \right. \\ & - 8EU_i(b_j - \delta, b_{-j}, \tau, \bar{s}) + 8EU_i(b_k - \delta, b_{-k}, \tau, \bar{s}) \\ & + 8EU_i(b_j + \delta, b_{-j}, \tau, \bar{s}) - 8EU_i(b_k + \delta, b_{-k}, \tau, \bar{s}) \\ & \left. - EU_i(b_j + 2\delta, b_{-j}, \tau, \bar{s}) + EU_i(b_k + 2\delta, b_{-k}, \tau, \bar{s}) \right] + o(\delta^4), \end{aligned}$$

Then the order of approximation error on the right-hand side is $o(\delta^4)$. For stability of the computational algorithm it is necessary that $\delta^4 = o(\Delta)$. This can be achieved even if one chooses $\delta = \Delta$ (up to scale of the grid). This condition becomes essential if in the sample the function EU_i is not smooth. In that case the minimal step size δ is determined by the granularity of the support of the payoff function. The step size for τ should be chosen appropriately and cannot be too small to avoid the accumulation of numerical error.

Initialization of the system simplifies when the auction has a reserve price. When the reserve price is equal to r , then both the expected utility and the total expenditure become functions of r . Homogeneity of the utility

¹²We need to emphasize that for the appropriate quality of approximation, when using the fourth-order formula for the numerical derivative, one needs to assure that $\Delta \gg \delta^5$. In other words, the step size for numerical integration should be larger than the step size for numerical derivative.

function will also be preserved when we consider the vector of bids accompanied by r . As a result, the system of equilibrium equations will take the form

$$\frac{\partial}{\partial \mathbf{b}'} EU(\mathbf{b}, \bar{s}, r) \mathbf{b} + \frac{\partial}{\partial r} EU(\mathbf{b}, \bar{s}, r) r = -TE(\mathbf{b}, \bar{s}, r). \quad (\text{G.28})$$

As a result, we can re-write our main result as

$$\frac{d}{d\tau} EU_i(b_i, \tau \mathbf{b}_{-i}, \bar{s})|_{\tau=1} = -TE_i(\mathbf{b}, \bar{s}) - r \frac{\partial}{\partial r} EU_i(\mathbf{b}, \bar{s}, r). \quad (\text{G.29})$$

Our results for τ in the neighborhood of $\tau = 1$ will apply with total expenditure function corrected by the influence of the reserve price. In the case where the vector of the payoff functions has a non-singular Jacobi matrix globally in the support of bids, we can also extend the results for $\tau \in [0, 1]$ to the case with the reserve price. In this case, the initial condition for $\tau = 0$ will solve

$$-TE_i(b_i(0), 0, \bar{s}) - r \frac{\partial}{\partial r} EU_i(b_i(0), 0, \bar{s}, r) = 0.$$

Note that for all bidders $i = 1, \dots, N$ this is a non-linear equation with a scalar argument $b_i(0)$, which can be solved numerically. This will allow us to construct a starting value for the system of differential equations. Note that in this case equilibrium computations simplify because there is no need in the “inverse” Euler step which we used to stabilize the system of differential equations at the origin. The algorithm will start from the standard preliminary Euler step $\frac{1}{2}\Delta$.

H The sources of estimation bias and robustness check

In this Appendix, we discuss the modeling choices we made in light of the data limitations, and we present the empirical results that establish the robustness of our estimation approach to these modeling choices.

There are three main elements used to estimate the marginal cost for a particular advertiser: (i) the distribution of quality scores (mean values and the distribution of shocks); (ii) the set of user queries where the advertisement of the advertiser of interest was considered; (iii) the set of competing advertisements that was considered for a user query. The feature of our historical research dataset is that we do not observe bids and quality scores for advertisements that did not appear on the page. As a result, we do not know the full set of advertisements that was considered for a particular user query. There could be several reasons why an advertisement did not appear in a particular user query. First, the random draw of the quality score was too low and the score-weighted bid of the advertiser was either outbid by other bidders or did not exceed the reserve price. Second, the advertiser has

set budget limit for the ad campaign and the budget has been exceeded. Third, the advertiser has set exclusion targeting and a particular user query does not satisfy targeting restrictions.

In our empirical analysis we assume that the observed sets of ads coincides with the sets of ads considered for user queries. This creates several potential problems for our analysis, which can be discussed in the context of the three components of the marginal cost estimation. First, a selection problem may arise, in that we only observe quality scores that were high enough so that the product of the advertiser’s per-click bid and their quality score ranked in the top set of advertisements. This could potentially impact our estimates of the mean quality scores as well as the shape of the quality score distribution. Second, we may over-estimate the uncertainty in rival configuration by exposing the ad to the queries for which it was not eligible due to exclusion targeting.

Now consider how we handle these problems. Our approach is loosely motivated by a model (although this is not a completely accurate description of the setting) where advertisers submit multiple advertisements and have budget constraints that determine the fraction of user queries the advertisements might appear on, and the system randomly selects which advertisement is chosen as well as which user queries to assign the advertisement to. We first discuss the choices and provide a comparison between the outcomes of the following empirical analyses. (i) We use our baseline methodology and estimate the distribution of quality scores, ignoring the selection problem and treat the data as if it came from the population of quality scores rather than a selected sample. In addition, we focus only on the first page of advertisements viewed by the user, which account for a very large share of the clicks and revenue for each advertisement. (ii) We assume that each advertisement’s bid was considered for all user queries in the sample (that is, even though in practice the advertisement did not appear on many user queries, we assume that a priori the advertisement could have appeared on any of them and the advertiser did not anticipate in advance which subset would be selected). (iii) We assume that the empirical distribution of competing advertisements is the distribution that advertiser anticipates.

We discussed the first approach in our main empirical section. We will now compare the results obtained using our baseline approach with the results obtained under the second and the third sets of assumptions.

Estimation of the score distribution: To study the effect of the sample selection on the estimate of the distribution of the shocks to the scores, we adapt our estimation methodology to the second assumption, that each advertisement’s bid was considered for all user queries in the sample. Therefore, the ads that did not appear in some user queries received low draws of quality scores. Our goal is to assess the robustness of our estimate of the empirical distribution of shocks to the scores to this assumption.

To estimate the marginal cost of advertisers under this assumption we created an additional dataset that contains user sessions where user queries contain pages with the search results beyond the first one. Then the ads that were considered for the first page of the search results but were not placed because of low draws of the scores can be considered for placement in the lower pages of the search results. By creating a database of the ads within the same user session we construct an approximation to the set of ads considered for a certain query. Then we use the sample of such long sessions to construct the empirical distribution of shocks to the scores.

We were not able to reject the null in the Kolmogorov-Smirnov test while testing the difference between the estimated cdf of the original sample of shocks to the scores and in the new dataset. This similarity between the empirical distributions of shocks to the scores translates into the similarity of the estimated values per click for advertisers demonstrated in Table 11. The deviation of the estimated value with the adjustment for ad eligibility is the largest for the bidders in the bottom positions, which is a feature inherited from the approach taking into account eligibility of ads. The impact of the bias in estimation of the distribution of shocks to the scores is small. We show the histogram for the estimated distribution of logarithm of shocks excluding top and bottom 1% quantiles. As one can see, even though the distribution has long right and left “tails”, most of the distribution mass is concentrated about zero with a much larger kurtosis than the normal distribution. This means that even though the scores may take very small values, the probability of such extreme draws is small and is not sufficient to create large biases in the estimates of values.

Selection of user queries where the advertisement is considered: To study the effect of eligibility of ads for queries, we adapted our empirical methodology to the third assumption that the empirical distribution of observed competing advertisements is the distribution that advertiser anticipates.

We use the additional dataset on long user queries to estimate the scores. Then when we estimate the expected cost per click for the advertisers, we only use rival ad configurations where the ad of the advertiser of interest was observed. Note that the disadvantage of this approach (one reason why we did not adopt it for our baseline methodology) is that we ignore the fact that ads did not appear in certain user queries because their quality scores were low. Therefore, we may underestimate the impact of bids on participation since a higher bid may lead to a higher probability of participation.

We estimate the marginal cost for each advertiser by using only user queries where the advertisement of this advertiser was displayed. Then using the finite-point approximation to the derivative, we estimate the marginal cost for each bidder and recover valuations. The results of the analysis across three analyzed search phrases are

demonstrated in Table 11. We show the mean log-values for all bidders and also separate the results for the top bidders (those whose average position is above 2) and the bottom bidders. One can see that the impact of the imposed change in the procedure on the overall mean is below 1%. A bidder-by-bidder analysis shows that for all bidders the confidence intervals for the valuations obtained using our main method and the method adjusted to the ad eligibility overlap. The deviation of the estimated value with the adjustment for ad eligibility is the largest for the bidders in the bottom positions. The main explanation for this result is that many ads that are at the bottom positions are appearing infrequently. This means that the sample sizes that can be used for the method taking into account the ad eligibility are small, leading to larger error in the estimated values.

Table 11: Log-values recovered from alternative estimation procedures

Baseline estimator			
Search phrase	Mean	Avg. position < 2	Avg. position > 2
#1	-3.253	-2.673	-3.479
#2	-1.988	-1.329	-2.655
Adjustment for ad eligibility			
Search phrase	Mean	Avg. position < 2	Avg. position > 2
#1	-3.330	-2.849	-3.689
#2	-2.024	-1.367	-2.731
Adjustment for selection bias			
Search phrase	Mean	Avg. position < 2	Avg. position > 2
#1	-3.253	-2.657	-3.465
#2	-1.992	-1.321	-2.671

I Standard errors of estimated values and predicted welfare and revenue

To estimate the standard errors we use the result in Appendix C. We note that we estimated the valuation as

$$\hat{v}_i = \frac{\frac{\partial}{\partial b} \widehat{TE}_i(b_i)}{\frac{\partial}{\partial b} \widehat{Q}_i(b_i)}$$

Using the Delta method, we can attribute the lead component of the variance to the numerator in this formula. Recalling that

$$\frac{\partial \widehat{TE}_i(b_i)}{\partial b} = \frac{\widehat{TE}_i(b_i + 2\tau_n) - 8\widehat{TE}_i(b_i + \tau_n) + 8\widehat{TE}_i(b_i - \tau_n) - \widehat{TE}_i(b_i - 2\tau_n)}{12\tau_n}.$$

The central element of the asymptotic variance formula is

$$\Omega = \text{Var} \left(\frac{u(v_i, b_i + \tau_n, \varepsilon_t, C_t) - u(v_i, b_i - \tau_n, \varepsilon_t, C_t)}{\sqrt{\tau_n}} \right).$$

Thus, to estimate the variance we need to find an estimator for Ω . To do that we, first find the conditional variance given the configuration, which we estimate by averaging the payoff differences over the empirical distribution of score shocks (by making the draws from that empirical distribution):

$$\hat{\Omega}(C_t) = \frac{1}{N_s} \sum_{s=1}^{N_s} \left(\frac{u(v_i, b_i + \tau_n, \varepsilon_s, C_t) - u(v_i, b_i - \tau_n, \varepsilon_s, C_t)}{\sqrt{\tau_n}} - \sqrt{\tau_n} \frac{\partial \widehat{EU}_i(b_i)}{\partial b} \right)^2,$$

where N_s is the number of simulated draws from the empirical distribution of score shocks. In our computations we used the simulated sample of size 20,000, which allows us to neglect the error of simulation as compared to the statistical noise in the data. In the last step, we estimated $\hat{\Omega}$ by taking the average over the sample of observed bidder configurations. The final expression for the variance formula corresponds to that in Appendix C.

We demonstrate the standard errors of the estimated bidder values in Table 12 tabulated by the average bidder positions. One can see that the standard errors tend to increase towards the top positions for search phrase # 1 which has the bidders who occupy top positions in a large fraction of queries, thus not creating a large variation for value estimation. However, the standard errors remain small, amount to 2-5% of the value for most bidders.

To estimate the standard errors of estimates of welfare and revenue changes in the counterfactual experiments we assumed that the standard errors of estimated are uncorrelated across advertisers. In that case we can evaluate the asymptotic standard error of the computed welfare, profits and revenue by combining the obtained estimate for the standard error of bidder valuations and using the delta-method to compute the standard error of computed equilibrium bid vector:

$$\frac{\partial b(v)}{\partial v_j} = (H_1 - H_2)^{-1} h_j,$$

where $H_1 = \left(v_i \frac{\partial^2 Q_i}{\partial b_i \partial b_k} \right)_{i,k}^I$ and $H_2 = \left(\frac{\partial^2 TE_i}{\partial b_i \partial b_k} \right)_{i,k}^I$ are square matrices and $h_i = \left(0, \dots, 0, \frac{\partial Q_i}{\partial b_i}, 0, \dots, 0 \right)'$. This allows us to compute the asymptotic formulas for the standard errors of bids as

$$\sigma_{b_i} = \sqrt{\sum_{k=1}^I \left(\frac{\partial b_i}{\partial v_k} \right)^2 \sigma_{v_k}},$$

Table 12: Standard Errors of Estimated Values, by Average Position Range

Average Position Range	Mean Value	Mean Std. Err.	Std. Err. 25%	Std. Err. 50%	Std. Err. 75%
Search phrase #1					
[1, 1.5)	0.14075	0.00411	0.00037	0.00109	0.00132
[1.5, 2.5)	0.30559	0.02661	0.00116	0.00265	0.00544
[2.5, 4)	0.04080	0.00150	0.00044	0.00078	0.00304
[4, 5.5)	0.04942	0.00093	0.00045	0.00061	0.00091
[5.5, 8]	0.05157	0.00152	0.00057	0.00105	0.00208
Search phrase #2					
[1, 1.5)	0.74798	0.01904	0.00437	0.01887	0.03254
[1.5, 2.5)	0.93997	0.02392	0.01306	0.01905	0.03341
[2.5, 4)	0.62005	0.01495	0.00313	0.00930	0.01638
[4, 5.5)	0.22394	0.00594	0.00178	0.00463	0.00698
[5.5, 8]	0.11117	0.00342	0.00110	0.00217	0.00389

The table represents the mean estimated valuations of advertisers along with the mean and quantiles of asymptotic standard errors of estimated values.

The numbers in the table are computed at the advertiser level.

Table 13: Standard Errors of Estimated Values as a % of Value, by Average Position Range

Average Position Range	Mean Value	Mean Std. Err.	Std. Err. 25%	Std. Err. 50%	Std. Err. 75%
(% of value)					
Search phrase #1					
[1, 1.5)	0.141	2.460	2.139	2.265	3.119
[1.5, 2.5)	0.306	3.787	1.599	2.000	4.110
[2.5, 4)	0.041	4.123	1.920	4.626	5.608
[4, 5.5)	0.049	2.578	0.742	1.803	3.162
[5.5, 8]	0.052	3.275	1.981	3.321	4.151
Search phrase #2					
[1, 1.5)	0.748	2.295	1.162	2.454	3.208
[1.5, 2.5)	0.940	2.429	1.740	2.111	3.037
[2.5, 4)	0.620	2.004	1.314	1.731	2.375
[4, 5.5)	0.224	3.375	1.189	1.938	4.262
[5.5, 8]	0.111	3.843	1.384	2.979	5.839

The table represents the mean estimated valuations of advertisers along with the mean and quantiles of asymptotic standard errors of estimated values relative to the values (in percentage terms). The numbers in the table are computed at the advertiser level.

given that the valuation estimates were assumed independent. Similarly, we can compute covariances between the bids and the values, which delivers the expressions for the variances of the estimated profits, welfare and revenue.

We performed the counterfactual experiments using large samples of simulated scores and configurations. Provided that the simulation error has stochastic order $O(N_s^{-1/2})$, where N_s is the total number of simulations, provided our simulation sample size of 20,000 for the score shocks and 500,000 bidder configurations, simulation error is small relative to the statistical estimation error. To estimate the standard errors of revenue and welfare, we use the delta-method. For instance, to evaluate the standard error of estimated welfare we consider the first-order Taylor expansion:

$$W(\hat{v}_1, \dots, \hat{v}_I) = \sum_{i=1}^I \frac{\partial W}{\partial v_i} (\hat{v}_i - v_i) + O(\max_i |\hat{v}_i - v_i|^2)$$

Then we construct the standard error for the welfare as

$$\sigma_W = \sqrt{\sum_{i=1}^I \left(\frac{\partial W}{\partial v_i} \right)^2 \sigma_{v_i}^2}.$$

When computing the standard error of the advertiser's profit, we note that the profit is maximized at a given bid. This means that the first term of the Taylor expansion of the profit with respect to the own bid will vanish. As a result, the standard error corresponding to the own bidder value will only have a direct effect on the advertiser profit (without the indirect effect through the bid) and only opponent bids will impact the standard error of advertiser's profit. As a result, we observe comparable standard errors for the evaluated welfare and the aggregate profits of the bidders.

The standard errors for predicted welfare, revenue and profits are given in Tables 14 and 15. As one can see, standard errors are small relative to the predicted changes in revenue, profits and welfare. This allows us to claim that the predicted effects are significant. In Tables 6 and 7 we show the results of the test for a significant change in the welfare, revenue and advertiser profits during the counterfactuals. While for the first search phrase, we were not able to reject the null hypothesis that the welfare changes during the counterfactual, we find that the change of revenue and profits of advertisers are significant. On the other hand, the change in the welfare, revenue and profits of advertisers are significant for the second search phrase.

We also demonstrate the standard errors of evaluated revenue and welfare for the experiment comparing the Vickrey and the SEU setup. Table 16 gives the standard errors of evaluated revenue and welfare by the SEU position buckets. The table clearly shows that standard errors do not exceed 10% of revenue and welfare figures in all position intervals. In Table 5 we provide the results of the test for the difference in revenue and welfare

between the NU-EOS and SEU model. As the table shows, the predicted differences in welfare are insignificant between the two models. On the other hand, the difference in revenues is significant.

Table 14: Standard errors for social welfare, auction revenue, and advertiser profits in case of squashing in Table 6

Outcome	Scenario	Positions			
		All	1	2–5	6–8
Search phrase #1					
<i>Welfare</i>	Before Squashing	0.06385	0.06262	0.00456	0.00006
	After Squashing	0.06434	0.06402	0.00487	0.00007
	Original Scores, Bids After Squashing	0.06421	0.06373	0.00470	0.00006
	After Squashing, Original Bids	0.06432	0.06394	0.00480	0.00007
<i>Revenue</i>	Before Squashing	0.00388	0.00047	0.00340	0.00008
	After Squashing	0.00294	0.00033	0.00261	0.00012
	Original Scores, Bids After Squashing	0.00376	0.00047	0.00327	0.00009
	After Squashing, Original Bids	0.00324	0.00034	0.00292	0.00011
<i>Advertiser Profit</i>	Before Squashing	0.06360	0.06238	0.00457	0.00007
	After Squashing	0.06407	0.06378	0.00460	0.00010
	Original Scores, Bids After Squashing	0.06396	0.06349	0.00467	0.00007
	After Squashing, Original Bids	0.06405	0.06370	0.00453	0.00009
Search phrase #2					
<i>Welfare</i>	Before Squashing	0.01865	0.01278	0.00718	0.00027
	After Squashing	0.01784	0.01278	0.00783	0.00031
	Original Scores, Bids After Squashing	0.01907	0.01336	0.00680	0.00026
	After Squashing, Original Bids	0.01743	0.01261	0.00772	0.00031
<i>Revenue</i>	Before Squashing	0.00421	0.00279	0.00144	0.00005
	After Squashing	0.00310	0.00108	0.00206	0.00005
	Original Scores, Bids After Squashing	0.00437	0.00318	0.00131	0.00005
	After Squashing, Original Bids	0.00332	0.00132	0.00204	0.00005
<i>Advertiser Profit</i>	Before Squashing	0.01689	0.01154	0.00674	0.00023
	After Squashing	0.01703	0.01267	0.00683	0.00027
	Original Scores, Bids After Squashing	0.01709	0.01159	0.00651	0.00023
	After Squashing, Original Bids	0.01648	0.01244	0.00671	0.00027

This table represents the asymptotic standard errors for expected total per query revenue of the auction platform and the expected per query social welfare. To compute the standard errors for the welfare and revenue we use the delta-method to compute the overall standard error of the predicted welfare, revenues and profit using the asymptotic standard errors for the estimated valuations of advertisers. Values and cost per click are normalized by the maximum bid for the search phrase # 1. Position ranges are determined by the SEU bids in each scenario.

Table 15: Standard errors for social welfare, auction revenue, and advertiser profits in case of score coarsening in Table 7

Outcome	Scenario	Positions			
		All	1	2-5	6-8
Search phrase #1					
<i>Welfare</i>	Before Coarsening	0.06386	0.06263	0.00456	0.00006
	After Coarsening	0.06435	0.06403	0.00489	0.00006
	Original Scores, Bids After Coarsening	0.06422	0.06374	0.00469	0.00006
	After Coarsening, Original Bids	0.06430	0.06381	0.00477	0.00006
<i>Revenue</i>	Before Coarsening	0.00390	0.00047	0.00341	0.00008
	After Coarsening	0.00345	0.00033	0.00314	0.00010
	Original Scores, Bids After Coarsening	0.00377	0.00047	0.00329	0.00009
	After Coarsening, Original Bids	0.00365	0.00034	0.00333	0.00009
<i>Advertiser Profit</i>	Before Coarsening	0.06361	0.06240	0.00457	0.00007
	After Coarsening	0.06410	0.06379	0.00482	0.00008
	Original Scores, Bids After Coarsening	0.06397	0.06350	0.00467	0.00007
	After Coarsening, Original Bids	0.06405	0.06357	0.00474	0.00008
Search phrase #2					
<i>Welfare</i>	Before Coarsening	0.01864	0.01276	0.00719	0.00027
	After Coarsening	0.01899	0.01315	0.00696	0.00027
	Original Scores, Bids After Coarsening	0.01907	0.01337	0.00681	0.00026
	After Coarsening, Original Bids	0.01858	0.01263	0.00735	0.00028
<i>Revenue</i>	Before Coarsening	0.00421	0.00278	0.00145	0.00005
	After Coarsening	0.00435	0.00313	0.00135	0.00005
	Original Scores, Bids After Coarsening	0.00437	0.00317	0.00132	0.00005
	After Coarsening, Original Bids	0.00422	0.00278	0.00147	0.00006
<i>Advertiser Profit</i>	Before Coarsening	0.01688	0.01153	0.00675	0.00023
	After Coarsening	0.01701	0.01142	0.00666	0.00024
	Original Scores, Bids After Coarsening	0.01709	0.01160	0.00652	0.00023
	After Coarsening, Original Bids	0.01681	0.01139	0.00690	0.00024

This table represents the asymptotic standard errors for expected total per query revenue of the auction platform and the expected per query social welfare. To compute the standard errors for the welfare and revenue we use the delta-method to compute the overall standard error of the predicted welfare, revenues and profit using the asymptotic standard errors for the estimated valuations of advertisers. Values and cost per click are normalized by the maximum bid for the search phrase # 1. Position ranges are determined by the SEU bids in each scenario.

Table 16: Standard errors for predicted revenues and welfare for the SEU generalized second price auction model versus the NU-EOS (equivalent to query-by-query Vickrey auctions) model in Table 5

Model (values)	Positions			
	All	1	2-5	6-8
Search phrase #1				
Revenue SEU (SEU)	0.00134	0.00030	0.00120	0.00020
Revenue Vickrey=NU-EOS(SEU)	0.00129	0.00031	0.00115	0.00020
Welfare SEU (SEU)	0.08203	0.08128	0.00747	0.00009
Welfare Vickrey=NU-EOS (SEU)	0.08217	0.08181	0.00758	0.00009
Search phrase #2				
Revenue SEU (SEU)	0.00417	0.00322	0.00107	0.00007
Revenue Vickrey=NU-EOS(SEU)	0.00428	0.00389	0.00139	0.00007
Welfare SEU (SEU)	0.01915	0.01329	0.00851	0.00035
Welfare Vickrey=NU-EOS (SEU)	0.02010	0.01531	0.00825	0.00035

This table represents the asymptotic standard errors for expected total per query revenue of the auction platform and the expected per query social welfare. To compute the standard errors for the welfare and revenue we use the delta-method to compute the overall standard error of the predicted welfare, revenues and profit using the asymptotic standard errors for the estimated valuations of advertisers. Values and cost per click are normalized by the maximum bid for the search phrase # 1. To group the SEU welfare and SEU revenue by positions we used the positions based on SEU equilibrium bids. To group the EOS welfare and EOS revenue by positions we used the positions based on EOS equilibrium bids.