To Score or Not to Score? Estimates of a Sponsored Search Auctions Model

Yu-Wei Hsieh* USC

Matthew Shum† Caltech

Sha Yang‡ USC

August 10, 2015

Abstract

We estimate a structural model of sponsored search auctions at "Website X", one of the largest online marketplaces in China. We use the estimates to address the distributional and revenue impact of "scoring" the auction, which is a way for shopping websites to reward larger merchants (by penalizing smaller merchants). We find evidence of horizontal differentiation in preferences, as merchants of different qualities differ in their valuations of ad positions. Counterfactual evaluations show that the price trend becomes steeper after moving to a score-weighted generalized second price auction, with much higher prices obtained for the top position but lower prices for the other positions. Overall there is only a very modest change in total revenue from introducing popularity scoring, despite the intent in bid scoring to reward popular merchants with price discounts.

Keywords: Sponsored-search advertising; Auctions; Market design; Two-sided Matching; Bayesian estimation.

JEL classification: D44; D47; C11; C15

*Department of Economics, University of Southern California; `yuwei.hsieh@usc.edu`

†Division of Humanities and Social Sciences, Caltech; `mshum@caltech.edu`

‡Department of Marketing, USC Marshall School of Business; `shayang@marshall.usc.edu`

Acknowledgement: We thank Arie Beresteanu, Baiyu Dong, Hashem Pesaran, Joris Pinkse, Sergio Montero, Roger Moon, Geert Ridder and Guofu Tan for their helpful discussions. Thanks also go to the seminar participants at Colorado-Boulder, CUEB (Beijing), Penn State, Rice, Rochester, USC, UC-Riverside, UCLA, Yale-SOM, 2013 Annual Conference of Taiwan Econometric Society, 2014 California Econometrics Conference and 2014 North American Summer Meeting of the Econometric Society.

1 Introduction

> ***Longzhi521***: "small merchants on WebsiteX are unable to do business anymore, now WebsiteX is just for the big merchants, its really unfair!"
>
> ***ILoveRainyDays***: "[I] recently quit my job to sell on WebsiteX full time, but after the change it is impossible, as a small merchant, to get my ads noticed."
>
> ***HangzhouWuMing*** (quoting an ancient Chinese proverb): "the waters (small merchants) which carry the boat can also capsize it!"

Longzhi521, *ILoveRainyDays*, and *HangzhouWuMing* are three small merchants selling their wares through WebsiteX (not its real name), one of the largest online marketplaces in China (and one of the busiest websites in the world by web traffic). These merchants, among many others, gathered in front of WebsiteX's Beijing headquarters[1] in an impromptu protest against a change in the rules for its sponsored search auctions. Such auctions are widely used by online companies for determining the positions and pricing for advertisements on their websites.[2]

Specifically, these merchants were angry that WebsiteX began to incorporate *scoring* into its sponsored search auctions. With scoring, the ad position obtained by a merchant depends on her bid times a "quality score", which reflects the quality/popularity of a seller or her ads (as measured by factors like past sales or click volumes). By introducing scoring, WebsiteX essentially intended to reward "popular" merchants – those whose previous ads generated many clicks, and sold many products in the past – with price discounts. This idea was originally developed by Google (see Varian, 2007) and subsequent adopted widely by other search engines and online marketplaces including Bing, Amazon, etc. The discontent of the small merchants quoted above underscores a new dimension to the market design problem – that scoring introduces important distributional impacts across different merchants, especially at platforms like WebsiteX, which consists of many small, vulnerable merchants.[3]

Despite its popularity in practice, little is known empirically about introducing scoring into sponsored search auctions. A main contribution of this paper is to measure the allocative and revenue consequences of introducing scoring into sponsored search auctions. Utilizing a large dataset from WebsiteX, we estimate a structural model of sponsored search auctions. Methodologically, we develop a novel econometric model which exploits an equivalence between sponsored search auctions and a classic model of two-sided matching with transferable

[1] "Merchants gather in front of WebsiteX's headquarters to protest new rules", news release from sina.com.cn (date: Sept. 7, 2010)

[2] A report from the Internet Advertising Bureau in 2013 shows that sponsored-search advertising is now generating more than $18 billion per year, nearly half of the annual online advertising spending in US.

[3] Indeed, eBay, a shopping platform used by a large number of small merchants, refrains from using auctions to determine search positions on its webpages, perhaps to avoid such complicated distributional issues.

utility (Shapley-Shubik (1972)).

Results: a snapshot. Our estimation results show that merchants of different qualities differ in their preferences across ad positions. High quality merchants value top positions for informative keywords, and hence bid more aggressively for, and win these positions. Lower quality merchants, on the other hand, value vague keywords more than higher-quality merchants, and are more likely to obtain top positions for these keywords. One explanation is that, for digital camera-related keywords, informative keywords (including specific camera model numbers) are likely to be queried by serious buyers, and high quality merchants are more experienced and have learned that clicks from these buyers are more likely to lead to sales. On the other hand, users who query vague keywords (including brand names and promotional terms) may not be ready to buy, and so experienced high quality merchants are less interested in these auctions. Thus we find evidence of both horizontal and vertical differentiation in these auctions; auction outcomes can be non-assortative, in that higher quality merchants do not always obtain higher positions.

Using these estimates, we perform counterfactual simulations to show that the price trend becomes steeper under scoring, with much higher prices obtained for the top position but lower prices for the other positions. This suggests that while the intention of scoring was to reward popular merchants with price discounts, these discounts were undone (in part) by more aggressive bidding by these merchants. On the one hand, this confirms the fears of the small merchants quoted above that they have to pay higher prices to get top positions. However, since CTRs in our model are not necessarily decreasing in ad position (and indeed are quite random), the effect on the platform's total revenue is ambiguous. Our simulations show that, indeed, there is only a very modest change in total revenue from introducing popularity scoring. Although other long term benefits from scoring – such as encouraging merchants to improve ad quality, thus generating more clicks and enhancing the attractive of the selling platform – are beyond the scope of our analysis, it is remarkable that even the platform's short term profits may not be much affected by rewarding popular merchants with price discounts.

Existing literature. Although sponsored advertising auctions have received great attention in the theoretical literature (beginning with Edelman et al. (2007) and Varian (2007)), empirical research is sparser. Börgers et al. (2013) utilize a revealed-preference approach to test whether bids in Yahoo search auctions satisfied the Nash Equilibrium inequalities for the sponsored search auction model. Yang et al. (2014) studied how competition affects sponsored search advertisers' bidding behavior, and they modeled bids as equilibrium outcomes using the specification in Edelman et al. (2007). Athey and Nekipelov (forthcoming) propose and estimate a structural model tailored to specific features of sponsored search auctions run by US search engines (such as Google or Microsoft). Specifically, they estimate a model characterized by score and entry uncertainty ("SEU"), in which bidders face uncertainty when choosing their bids, due to randomness in a bidder's quality score over time,[4] as well as in the set of competitors bidding in the auction at any time.

Our paper makes several contributions to this literature. First, our paper appears to be the first which aims to measure the distributional and revenue effects of scoring. The existing literature has mainly focused on efficiency properties in sponsored search auctions (e.g. Athey and Nekipelov (forthcoming)). Indeed, the use of bid scoring in sponsored search auctions has similar implications as bid preference policies which have been analyzed in other auction settings, including procurement auctions (cf. McAfee and MacMillan (1989), Krasnokutskaya and Seim (2011)).

Second, our model relaxes some restrictive but convenient specification assumptions made in much of the existing literature. We allow for click through rates which are not necessarily decreasing with position ranks, a well-documented phenomenon which has been called the "position paradox".[5] We also allow bidders to have preferences for positions which are not multiplicative in bidder- and position-specific effects.[6] In addition, our empirical approach easily accommodates the rich auction and bidder heterogeneity observed in our dataset.

Methodologically, we propose a new estimation approach for sponsored search auctions model based on the equivalence of this model with Shapley and Shubik's (1972) assignment game, which has been used to model two-sided matching under transferable utility.[7] Viewing the prices and allocations in the sponsored search auctions as equilibrium outcomes arising from a particular two-sided matching game allows us to derive a likelihood function for these outcomes in a straightforward and computationally tractable manner. This connection between matching and auctions was also exploited in Fox and Bajari's (2011) study of the FCC spectrum auctions, but our estimation approach is quite different in spirit.

In the following section we provide background information about the platform and present the structural model for ad positions and prices from the generalized second price auction. We derive the crucial link between the auction model and two-sided matching models, which

[4] During our sample period, our study website ran un-scored auctions, which eliminates an important source of uncertainty in Athey and Nekipelov's SEU model. Only *after* our sample period did the study website introduce scoring.

[5] See Jerath et al. (2011). The position paradox contradicts a typical assumption made in the literature, that click-through rates are decreasing across positions; e.g., Edelman et al. (2007); Varian (2007); Athey and Nekipelov (forthcoming)).

[6] Jeziorski and Segal (2015), Jeziorski and Moorthy (2014), Goldman and Rao (2014) document empirical evidence which contradict the multiplicative specification.

[7] See also Demange, Gale and Sotomayor (1986). Becker (1973), Choo and Siow (2006), Fox (2013), Galichon and Salanie (2012), and Graham (2011), among others, use this framework to study marriage markets. Other applications of the assignment game include mergers (Akkus et al. (2013)) and hedonic pricing models (Chiappori et al. (2009)).

we exploit in estimation. In section 3, we describe the data set and present some motivating preliminary analysis. Based on this, we specify our empirical model and propose an estimation approach, as discussed in sections 4 and 5. In section 6 we conduct the counterfactual analysis to address the "to score or not to score" question. Section 7 concludes.

2 A Model of Sponsored Search Auctions

2.1 Background: sponsored search auctions at "WebsiteX"

WebsiteX is one of the largest online marketplaces in China and, hence, one of the most prominent websites in the world by traffic. Given the high costs and regulatory and bureaucratic hurdles associated with opening brick-and-mortar businesses in China, many small merchants market and sell their wares mainly using internet shopping platforms like our study website. As such, WebsiteX has no direct American counterpart, but shares features of both eBay and Craigslist. Unlike eBay, goods on WebsiteX are not sold via auction, but rather by merchants posting prices for their products. WebsiteX provides a platform whereby buyers can make secure money transactions to merchants. Sponsored search results typically appear as "tiles" on the right-hand side and bottom margins of each search page. (See Figure 3.)

Since WebsiteX is a marketplace, the content and role of its sponsored ads differ substantially from the ads appearing on search engine result pages (like Google or Yahoo). As Figure 3 shows, WebsiteX's sponsored ads (as well as the non-sponsored "organic" search results) typically contain a picture of the product, price, merchant name and information, shipping details, and product specifications; in contrast, such details are typically absent from Google's sponsored search results, which contain only the URL along with some brief slogans.

Since WebsiteX is a shopping platform, most of its users have a serious intent of purchasing, and are using WebsiteX to find prices and product specifications to suit their needs.[8] Consequently, as we mentioned earlier, the top positions may not always receive more clicks (and indeed, they do not, as we will show below in section 3). For this reason, the common assumption in the existing sponsored search auction models (Edelman et al. (2007) and Varian (2007), Athey and Nekipelov (forthcoming)) that the surplus matrix is supermodular (being the product of a vector of bidder-specific constants and a nonincreasing vector of position-specific click-through rates) seems inadequate for WebsiteX. Accordingly, in our

[8] It has been discussed for a while that many retailers fear becoming Amazon's Showroom. This phenomenon is even more radical in China as running a retail store would incur more tax and fee liability. By contrast, running a online store can sidestep these hidden costs. Consequently, the price gap between online and retail stores in China is even larger than the USA counterpart.

setup, we will allow the surplus to vary arbitrarily among merchants and across ad positions, and also allow click volumes to be non-monotonic in ad position.

2.2 Generalized Second-Price Auction (GSPA) Mechanism

Next we describe the generalized second-price auction framework. There are N available positions and $M \geq N+1$ potential bidders (synonymously, merchants) for a generic keyword auction. If bidder i obtains the j-th position, he obtains valuation (or surplus) V_{ij}, for all bidders i and positions j. In what follows, without loss of generality we will index the positions from top to bottom by $i = 1, \ldots, N$,and similarly we will also label the $N+1$ highest bidders by $i = 1, \ldots, N+1$.

The rules of the generalized second-price auction are as follows. Each bidder submits a single bid and can only obtain at most one ad position. For N positions, the N-highest bidders will be winners, with the i-th ($1 \leq i \leq N$) highest bidder obtaining position i at the per-click price equal to the $i+1$-th bidder's bid. The final payment made to the platform is the per-click price times the click volume. Using the terminology in Varian (2007) and Börgers et al. (2013), we focus on the so-called "symmetric" Nash equilibria in this complete-information bidding game.[9] The equilibrium conditions satisfied by a bid vector $(b_1, \ldots, b_M)$ for $M \geq N+1$ are

$$V_{ii} - \alpha_i b_{i+1} \geq V_{ij} - \alpha_j b_{j+1}, \ \forall i, j \tag{1}$$

where α_j denotes the click volume for position j. This inequality ensures that, at the equilibrium, the bidder who obtains positions i (who obtains valuation V_{ii} and makes a payment equal to the click volume α_i times b_{i+1}, the per-click price submitted by the bidder in position $i+1$), does not wish to deviate to position j, for which the surplus would be equal to the RHS of the inequality.[10] Making the substitution $p_i = b_{i+1}$ (that is, the per-click price for the i-th position equals the bid in the $i+1$-th position), we have

$$V_{ii} - \alpha_i p_i \geq V_{ij} - \alpha_j p_j, \ \forall i, j. \tag{2}$$

[9] These are closely-related to the "locally envy-free" equilibria in Edelman et al. (2007). These equilibria are convenient to analyze, and easy to compute via linear programming; as noted in Börgers et al. (2013), no such characterization is available for the asymmetric Nash equilibria. Nevertheless, it is possible to adapt our estimator to the case of asymmetric Nash equilibria; we leave the details for future research.

[10] In contrast, in asymmetric Nash equilibria, Eq. (1) holds only for $j > i$, but is $V_{ii} - \alpha_i b_{i+1} \geq V_{ij} - \alpha_i b_j$ for $j < i$. This recognizes an asymmetry that in order to switch to a lower position, bidder i only needs to beat the price of that position, but to switch to a higher position, bidder i must beat the *bid* of the winner of that position.

2.3 GSPA as Two-sided Matching

Our estimation approach relies critically on the reinterpretation of the GSPA as a two-sided matching model as in Shapley and Shubik (1972), which Varian (2007) also alluded to. As we will see below, this link between GSPA and matching enables us to derive a likelihood function for the allocations and prices in the GSPA in a straightforward and computationally tractable manner. At the same time, this framework also readily accommodates the substantial auction and bidder heterogeneity which is present in our dataset.

To draw this connection, we consider a "matching" problem where bidders are matched to positions. We denote by

$$u_i \equiv V_{ii} - \alpha_i p_i \tag{3}$$

the equilibrium payoff of bidder i, and

$$t_j \equiv \alpha_j p_j \tag{4}$$

the equilibrium payoff for the platform from the j-th position. Now rewriting the equilibrium inequalities (2) above, we get

$$u_i + t_j \geq V_{ij}$$

with equality (by construction) iff $i = j$. These can be interpreted as "no-blocking" conditions for a matching problem (cf. Roth and Sotomayor (1990; chap. 8)). To see why, consider a bidder i and position j, and assume that $u_i + t_j < V_{ij}$ or, equivalently, $u_i < V_{ij} - t_j$. In this case, since bidder i's payoff u_i is lower than the net surplus that she would obtain from deviating to position j ($=V_{ij} - t_j$), she would not agree to the given allocation, and the equilibrium would break down; since the pair of bidder i and position j would "block" the proposed allocation, the payoffs (u_i, t_j) cannot support this allocation in equilibrium.

Moreover, introducing the binary indicators $\mu(i,j) = 1$ if bidder i obtains position j, and zero otherwise,[11] and summing up across all bidders and positions using Eqs. (3,4), we have

$$\sum_{i,j} u_i + t_j = \sum_{i,j} [\mu(i,j)V_{ij} - \alpha_i p_i + \alpha_i p_i] = \sum_{i,j} \mu(i,j)V_{ij}. \tag{5}$$

This "feasibility" condition is the link between the sponsored-search auction and two-sided matching models, as it is implied by the duality theorem of linear programming for that latter model. In the remainder of this section, we flesh out this connection.

[11] Using our indexing convention that bidder i is allocated position i, we have $\mu(i,j) = 1$ for $i = j$, and zero otherwise.

2.4 Optimal allocation in GSPA: matching Positions to Bidders

If bidder i obtains the j-th position, the valuation function is given by

$$V_{ij} = \delta(X_i, Z^j; \beta) + \epsilon_{ij} \tag{4}$$

where X_i is the vector of bidder i's characteristics and Z^j is the j-th position-specific characteristics. The $\delta(\cdot)$ is the deterministic component of the valuation function parametrized by a finite dimensional parameter β, and ϵ_{ij} is the unobservable match-and-auction-specific valuation.[12] We will refer to $\mathbf{V}$ the valuation matrix, where the (i, j) entry of $\mathbf{V}$ is V_{ij}. We further assume the unobserved valuation shocks satisfy the following assumption

Assumption 1. *ϵ_{ij} is a continuous random variable with mean zero and variance $\sigma^2 < \infty$ with unbounded support on $\mathcal{R}$. ϵ_{ij} is mutually independent across index i, j.*

An allocation (or matching) μ, is a binary matrix indicating which bidder acquires which position: its element (ij), $\mu(i, j) = 1$ if bidder i obtains position j, and zero otherwise. In a N-by-N matching game, there are $N!$ allocations. We will refer to Ω as the set of all possible allocations and μ_ω as an generic element of Ω, where $\omega = 1, 2, \ldots, N!$. The total surplus under allocation μ_ω is denoted by $S_{\mu_\omega}(\mathbf{V})$

$$S_{\mu_\omega}(\mathbf{V}) = \sum_{i=1}^{N} \sum_{j=1}^{N} \left[\delta(X_i, Z^j; \beta) + \epsilon_{ij}\right] \cdot \mu_\omega(i, j) \equiv \Delta_{\mu_\omega} + \Xi_{\mu_\omega}, \text{ where}$$

$$\Delta_{\mu_\omega} = \sum_{i=1}^{N} \sum_{j=1}^{N} \delta(X_i, Z^j; \beta) \cdot \mu_\omega(i, j) \text{ and}$$

$$\Xi_{\mu_\omega} = \sum_{i=1}^{N} \sum_{j=1}^{N} \epsilon_{ij} \cdot \mu_\omega(i, j)$$

Based on $\mathbf{V}$, the total surplus $S(\mathbf{V})$ for each allocation $\mu_\omega \in \Omega$ can be calculated. Shapley and Shubik (1972) consider the problem of finding the optimal one-to-one matching that maximizes the total (social) surplus, as well as the stable price systems to support/decentralize the optimal allocation. The social planner's problem, can be formulated as the following linear program (which we denote (P)):

[12] This is a complete information game in which ϵ_{ij} is assumed to be observable to all players within the game but unobservable to the researchers.

$$\begin{aligned} &\max_{\mu(i,j)} \textstyle\sum_{i,j} V_{ij}\mu(i,j) \\ \text{s.t.} \quad &\textstyle\sum_i \mu(i,j) = 1, \forall i \\ &\textstyle\sum_j \mu(i,j) = 1, \forall j \\ &\mu(i,j) \in \{0,1\}, \forall (i,j) \end{aligned} \tag{P}$$

We can easily show that under assumption 1, the optimal allocation (the solution to (P))is unique almost surely.[13]

2.5 Equilibrium prices in the GSPA

The linear program (P) yields the optimal allocation of assigning positions to bidders. Since our goal is to analyze not only allocations, but also prices, we will turn to the dual linear program, which yields the prices supporting the optimal allocation in equilibrium.

Recall that we denote by u_i the equilibrium payoff of bidder i, and $t_j(\equiv \alpha_j p_j)$ the equilibrium payoff of the j-th position. By the duality theorem of linear programming, the dual problem of (P) is given by

$$\begin{aligned} &\min \textstyle\sum_{i=1}^N u_i + \sum_{j=1}^N t_j \\ \text{s.t.} \quad &u_i \geq 0, t_j \geq 0, \forall i,j \\ &u_i + t_j \geq V_{ij}, \forall i,j \end{aligned} \tag{DP}$$

The first set of constraints, $u_i \geq 0, t_j \geq 0$, are individual rationality condition: both bidders and search engine should have non-negative profit. The second set of constraints, corresponds to the incentive compatibility, or no-blocking pair conditions. The set of (u_i, t_j) that solves (DP) is denoted the set of stable matchings (see Roth and Sotomayor, 1990). Shapley and Shubik (1972) further show that $u_i + t_j = V_{ij}$ iff $\mu_{ij} = 1$. By summing this across (i,j), we obtain Eq. (5) above, which provides the link between the GSPA and two-sided matching, as we alluded to before.

[13]Take any two allocations $\mu_l \neq \mu_q$, $(\mu_l, \mu_q) \in \Omega$. The event $\{S_{\mu_l}(\mathbf{V}) = S_{\mu_q}(\mathbf{V})\} = \{\Xi_{\mu_l} - \Xi_{\mu_q} = \Delta_{\mu_q} - \Delta_{\mu_l}\}$ is a set of measure zero under assumption 1. It immediately follows that the ordering of $\{S_{\mu_\omega}(\mathbf{V})\}_{\omega=1,\ldots,N!}$ is strict almost surely.

This uniqueness contrasts with two-sided matching models without transfers, in which the multiplicity of stable allocations becomes a major concern; e.g., Boyd et al. (2006), Logan et al. (2008), Menzel (2011), Hsieh (2011), and Echenique, et al. (2013), among others.

In general, it is well-known that there exist multiple transfers $\mathbf{t} = (t_1, \ldots, t_N)$ that solve (DP), and there exist bidder-optimal $\underline{\mathbf{t}}$ and platform-optimal $\overline{\mathbf{t}}$.[14] This multiplicity in equilibrium prices raises issues for estimation, as we will discuss below. By contrast, given Lemma 1, the corresponding optimal matching μ that solves (P) is unique almost surely.

Within the set of stable matchings $\mathbf{t}$, the generalized second-price auction mechanism selects a subset in which the transfers are monotonically decreasing in ad position:

$$\{\mathbf{t}|\mathbf{t} \text{ solves (DP) \& } p_1 > p_2 > \ldots p_N; t_i = \alpha_i p_i\}. \tag{6}$$

We will refer to this as the set of "stable per-click prices."

Example: (Non-)existence of equilibrium in GSPA. While Shapley and Shubik (1972) proved that the set of stable matchings is always nonempty for arbitrary V_{ij}, this is no longer true once the additional monotonicity condition (6) is imposed. Hence, without extra assumptions on V_{ij} or the click volumes, the generalized second-price mechanism may not necessarily guarantee the existence of a symmetric Nash equilibrium (or equivalently, competitive price system). Consider an example: the valuation matrix is given by

	1	2
bidder1	3	4
bidder2	1	3

and $\alpha_1 = \alpha_2 = 1$. The set of stable prices for the Shapley-Shubik assignment game is shaded in blue in Figure 4. We see that this region lies completely above the 45-degree line, where $p_2 > p_1$ (the price for the second position exceeds that of the top position); the GSPA, on the other hand, requires $p_1 > p_2$. The problem in this example is that both players value the second position more than the first one, making it impossible to sell the first position with a higher market price. In our empirical model, our assumptions on valuations (Assumption 1) allow cases similar to this example to arise with positive probability. For estimation, we wish to restrict attention only to the set of V_{ij} that are consistent with equilibrium in the GSPA, which introduces a complicated truncation problem.[15] ■

3 Data

We obtained a one-month (in 2010) sponsored-link auction data from WebsiteX, which is one of the largest online marketplace in China. The data set includes information on 487

[14]The literature on multi-item auction raises a design problem of how specific auction mechanism may select particular stable matchings. For example, Demange, Gale and Sotomayor (1986) propose an auction mechanism in which the bidder-optimal stable matching is the equilibrium outcome.

[15]This is similar to that arising in Hong and Shum's (2003) econometric study of asymmetric ascending auctions.

keywords of digital camera/camcorder and related accessories. For each keyword, we observe the averaged price paid by each of the participating bidders over that month, as well as the aggregate number of clicks of their ads over that month. The number of ad positions and bidders ranges from 5 to 9.[16] We map these data to the static auction framework of the previous section by sorting the bidders by their average per-click price, and then placing the bidder paying the highest average price in the the top position, the bidder paying the second highest price in the second position, etc.

According to insiders at WebsiteX, merchants there usually review their keyword lists and make purchase decisions infrequently, so we believe that the availability of only aggregated data should not produce too much bias. Indeed, during the sample period, the auction environment at WebsiteX was nowhere as complicated as at Google or Yahoo!, in which bidders often apply some automatic bidding algorithm to dynamically manage their positions and per-click prices. As we remarked before, WebsiteX ran unscored GSPA's during the sample period. The bidding environment of WebsiteX is therefore more closely related to the static model described in Edelman et al. (2007) and Varian (2007).

For each merchant, we also observe "quality ratings" which WebsiteX creates (using their proprietary algorithm) based on buyer feedback and merchants' sales volumes. Specifically, each bidder is rated on an increasing quality scale, which is also broken into quality brackets: high (11-15), medium (6-10) and low (1-5).[17] Third, we construct dummy variables to describe characteristics of each keyword string: 1. *Brand*, whether the keyword includes a brand name such as Nikon or Canon; 2. *Specific*, whether the keyword includes a specific model/series number, such as D300s or 500D; 3. *Promotional*, whether the keyword includes promotional terms such as "cheap" or "sale". Summary statistics for these variables are given in Table 1.

A first look at the data. We begin with some reduced form statistics and tabulations from the data to motivate details of our model specification. From the perspective of search engines it is important to know how advertisers' bidding strategies are related to their quality rankings and keyword characteristics. Do certain types of bidders bid more aggressively for certain type of keywords? What is the sorting pattern between quality and rank? Table 2 contains a contingency table summarizing bidders' quality versus the ad positions they obtained. Overall, the evidence for assortative matching in bidder quality across ad positions is mixed. On the one hand, from the top panel of the table, we see that about one-third of the high-quality bidders (=7.2% of total bidders) get the top ad position in the auctions in which they won a position, and this percentage falls across positions. However, for medium

[16] We do not observe losing bidders. Therefore in an auction with N positions, N bidders are observed.

[17] Subsequently, after our sample period, WebsiteX added an additional classification of "top" merchants, with quality scales running from 15 to 20.

quality bidders, we see that most of them are sorted to position 5 (16.1%) followed by 3 (15.6%). The results are qualitatively stable after doing the analysis separately for different types of keyword queries, as is shown in the remaining panels in Table 2.

Although it seems that the allocation patterns do not vary with keyword characteristics, the per-click prices and click volumes do move dramatically. The boxplot of log click volumes[18] and per-click prices are depicted in Figure (1) and (2) respectively. In Figure (1) we compare the boxplot of log click volumes across the top 5 positions, conditional on different dummies. First, clearly the assumption of decreasing click volume with rank is violated since the middle of the boxplot does not decrease when rank decreases. In fact, only 5 out of 487 keywords have strictly decreasing click volumes. Second, we observe that keyword characteristics shift the click volume distributions. Keywords containing specific model number usually receive more click volumes across all ranks of position (top-left of Figure 1). Keywords containing brand name slightly decrease the click volumes (top-right of Figure 1). It is also interesting to note that keywords containing promotional terms in fact generate smaller click volumes (bottom-left of Figure 1).

Lastly, we turn our attention to the distribution of per-click prices (Figure 2). We find that per-click prices are generally higher (and have more extreme outliers) for keywords containing specific model number (top-left of Figure 2). Adding brand name on average does not change per-click price, but it does create more outliers (top-right of Figure 2). Adding promotional terms does not change per-click price (bottom-left of Figure 2).

4 Estimation

Next we consider the estimation of the sponsored search auction model. Since the equilibrium conditions (in Eqs. (2)) are defined in terms of inequalities involving the observed prices and unobserved valuations, it is not obvious how to form estimating equations based on them. In this section we show how exploit the equivalence between the GSPA and two-sided matching to derive a likelihood function for the prices and allocations in a straightforward manner. Moreover, the characterization of the allocation and prices in the GSPA as optimzers from (dual) linear programming problems makes our estimation procedure quite tractable computationally.

We estimate the parameters using a Bayesian approach. There are a large number of latent variables in this model, relative to the observed variables: in each auction, for the N-dimensional vector of observed bids, there are N^2 corresponding unobservabled: namely, the

[18] As there are some keywords that receive extremely large amount of click volumes, for graphical presentation purpose we depict the boxplots in the log scale.

full set of valuations $\{V_{ij}\}_{i,j=1\ldots N}$ that each merchant has for each position. An important virtue of the Bayesian approach is the use of "data augmentation" (Tanner and Wong (1987)), whereby these latent variables are treated as unknown parameters, and jointly inferred in the estimation procedure.[19] Moreover, by exploiting the two-sided matching analogy, data augmentation coupled with importance sampling allows us to estimate the model without ever needing to solve for the equilibrium allocation and prices in the sponsored search auction, which simplifies the computational burden significantly.

4.1 The likelihood of auction allocations and prices

As the set of bidders and the number of positions vary from auction to auction, defining the random vector $\mathbf{p}$ and random matrix μ requires additional care to avoid a labeling problem. Following the previous section we will sort the element of $\mathbf{p}$, p_i, in decreasing order, and therefore X_i corresponds to the bidder who pay p_i. The econometrician observes T independent keyword auctions indexed by $t = 1, 2, \ldots, T$. In each keyword auction t one observes two sets of dependent variables, the allocation μ_t and the vector of per-click prices $\mathbf{p}_t = (p_{1t}, \ldots, p_{N_t t})$. The exogenous variables are the bidder-specific covariates $\mathbf{X}_t = (X_{1t}, \ldots, X_{N_t t})$, the auction and position-specific covariates $\mathbf{Z}_t = (Z^{1t}, \ldots, Z^{N_t t})$, and click-through rates $\alpha_t = (\alpha_{1t}, \ldots, \alpha_{N_t t})$. Let $\mathbf{V}$ denote the collection of latent valuation matrices for all keyword auctions $(\mathbf{V}_1, \mathbf{V}_2, \ldots, \mathbf{V}_T)$, and similarly we define $\mu = (\mu_1, \mu_2, \ldots, \mu_T)$, $\mathbf{p} = (\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_T)$ and $(\mathbf{X}, \mathbf{Z}, \alpha)$. The posterior is given by

$$f(\theta, \mathbf{V}|\mu, \mathbf{p}, \mathbf{X}, \mathbf{Z}, \alpha) \propto \mathcal{L}(\mu, \mathbf{p}|\theta, \mathbf{V}; \mathbf{X}, \mathbf{Z}, \alpha) p_0(\theta, \mathbf{V}|\mathbf{X}, \mathbf{Z})$$

The specification of the priors is given in the Appendix. Here we focus on the form of the likelihood, given by

$$\mathcal{L}(\mu, \mathbf{p}|\theta, \mathbf{V}; \mathbf{X}, \mathbf{Z}, \alpha) = \Pi_{t=1}^{T} \mathcal{L}(\mu_t, \mathbf{p_t}|\theta, \mathbf{V}_t; \mathbf{X}_t, \mathbf{Z}_t, \alpha_t). \tag{7}$$

Importantly, the likelihood above differs from the likelihood which would be optimized in a frequentist setting (ie. MLE), because the unobserved valuations $\mathbf{V}$ are treated as conditioning variables. This is due to "data augmentation", as discussed above. In the frequentist likelihood, in contrast, $\mathbf{V}$ cannot be conditioned on (since it is unobserved), and

[19] In contrast, in a frequentist framework, these latent valuations must be integrated out of the estimating equations; this is difficult using typical numerical integration methods (quadrature, simulation) due to the large dimensionality of the integration.

would need to be integrated out.[20] The likelihood (7) can be further decomposed into

$$\mathcal{L}(\mu_t, \mathbf{p}_t|\theta, \mathbf{V}_t; \mathbf{X}_t, \mathbf{Z}_t, \alpha_t) = \mathcal{L}_1(\mathbf{p}_t|\mu_t, \theta, \mathbf{V}_t; \mathbf{X}_t, \mathbf{Z}_t, \alpha_t)\mathcal{L}_2(\mu_t|\theta, \mathbf{V}_t; \mathbf{X}_t, \mathbf{Z}_t) \quad (7)$$

the two terms of which are the conditional likelihoods of per-click prices (given the allocation and valuations), and the allocation (given valuations). The explicit forms of $\mathcal{L}_1$ and $\mathcal{L}_2$ are derived from, respectively, the dual and primal LP problems for the assignment game. We consider each component in turn.

The component $\mathcal{L}_1$. For expositional simplicity we will drop the auction index t for now. Conditional on μ, the set of stable per-click prices is a convex polyhedron defined by a set of linear inequalities derived from the dual LP problem of the assignment game, which comprise (i) the "no-blocking pair" inequalities (Eq. (1)); (ii) the individual rationality constraints, that $V_{ii} - t_i \geq 0$; and (iii) the monotonicity and non-negativity conditions on per-click prices: $p_1 > p_2 > \ldots > p_N \geq 0$. The details of these linear inequalities are given in the appendix; we let $\mathbf{P}(\mathbf{V}, \alpha)$ denote the polyhedron of equilibrium per-click prices in the sponsored-search auction, given a set of bidder valuations and position-specific click volumes $(\mathbf{V}, \alpha)$. As $\mathbf{P}$ depends on the realization of unobserved ϵ_{ij}, the set of stable per-click prices itself is a random closed convex polyhedron.

Correspondingly, the conditional likelihood for the per-click prices is a distribution supported on a convex polyhedron: $\mathcal{L}_1(\mathbf{p}_t|\mu_t, \theta, \mathbf{V}_t; \mathbf{X}_t, \mathbf{Z}_t, \alpha_t) = \mathcal{L}_1(\mathbf{p}_t|\mu_t, \mathbf{P}(\mathbf{V_t}, \alpha_t), \lambda)$, where λ, a probability measure over $\mathbf{P}$, denotes an equilibrium selection rule whereby $\mathbf{p}_t$ is selected from $\mathbf{P}(\mathbf{V}_t, \alpha_t)$. In this paper, we complete the model by assuming particular parametric equilibrium selection rules, following Bajari et al. (2010).[21] Specifying such a selection rule here is nontrivial, as the support $\mathbf{P}(\mathbf{V}, \alpha)$ depends on the latent variables. We consider two parsimonious alternative specifications.

1. *Uniform selection.* Prices satisfying the no-blocking conditions are drawn with uniform probability, equal to the reciprocal of the volume of the polyhedron of values of prices which satisfy the linear inequalities $\{\mathbf{p} : A(\alpha_t)\mathbf{p} \leq \mathbf{b}(\mathbf{V}_t)\}$:

$$\mathcal{L}_1(\mathbf{p}_t|\mu_t, \mathbf{P}(\mathbf{V_t}, \alpha_t), \lambda) = \mathbb{1}(A(\alpha_t)\mathbf{p}_t \leq \mathbf{b}(\mathbf{V}_t))\frac{1}{vol(\mathbf{P}(\mathbf{V}_t, \alpha_t))}$$

2. *Beta selection.* Each p_{it} is independently drawn from a beta distribution defined on

[20] Specifically, the frequentist likelihood is $\mathcal{L}(\mu, \mathbf{p}|\theta, \mathbf{V}; \mathbf{X}, \mathbf{Z}, \alpha) = \int \cdots \int \mathcal{L}(\mu, \mathbf{p}|\theta, \mathbf{V}; \mathbf{X}, \mathbf{Z}, \alpha) dG(\mathbf{V}|\mathbf{X}, \mathbf{Z})$ which is difficult to evaluate due to the large dimensionality of $\mathbf{V}$.

[21] In Appendix B, we consider an alternative approach, using bounds based upon the structure of the equilibrium price set, which is agnostic as to the equilibrium selection rule. (This is in the spirit of Ciliberto and Tamer (2009); Beresteanu et al. (2011); Galichon and Henry (2011), among others.) Such an approach turns out to be quite computationally challenging, compared to the approach we use in this paper.

$[\underline{p}_{it}, \bar{p}_{it}]$ truncated to $A(\alpha_t)\mathbf{p}_t \leq \mathbf{b}(\mathbf{V}_t)$, where $\Pi_{i=1}^{N_t}[\underline{p}_{it}, \bar{p}_{it}]$ is the smallest bounding box of $\mathbf{P}(\mathbf{V}_t, \alpha_t)$:

$$\mathcal{L}_1(\mathbf{p}_t|\mu_t, \mathbf{P}(\mathbf{V_t}, \alpha_t), \lambda) \propto \mathbb{1}(A(\alpha_t)\mathbf{p}_t \leq \mathbf{b}(\mathbf{V}_t))\Pi_{i=1}^{N_t}\Big\{\frac{(p_{it} - \underline{p}_{it})^{a-1}(\bar{p}_{it} - p_{it})^{b-1}}{B(a,b)(\bar{p}_{it} - \underline{p}_{it})^{a+b-1}}\Big\},$$

where $B(a, b)$ is the Beta function. Under this specification, the shape parameters (a, b) in the beta distribution will be additional parameters to be estimated. Depending on the values of these shape parameters, the equilibrium selection distribution can be skewed towards high or lower prices, which are indicative as to whether equilibrium selection favors bidders (lower prices) or the platform (higher prices).

The component $\mathcal{L}_2$. Conditional on $\mathbf{V}_t$, the likelihood $\mathcal{L}_2$ is binary (0-1) valued:

$$\mathcal{L}_2(\mu_t|\theta, \mathbf{V}_t; \mathbf{X}_t, \mathbf{Z}_t) = \mathcal{L}_2(\mu_t|\mathbf{V}_t) = \begin{cases} 1 & \text{if } \mathbf{V}_t \text{ rationalizes } \mu_t \\ 0 & \text{otherwise} \end{cases}$$

Moreover, at the observed $(\mathbf{p}_t, \mu_t)$, all $\mathbf{V}_t$ which lead to a nonzero value for $\mathcal{L}_1$ automatically rationalize μ_t, and lead to a nonzero value of $\mathcal{L}_2$. As a result, $\mathcal{L}_2$ is redundant and we can simplify

$$\mathcal{L}(\mu_t, \mathbf{p}_t|\theta, \mathbf{V}_t; \mathbf{X}_t, \mathbf{Z}_t, \alpha_t) \propto \mathcal{L}_1(\mathbf{p}_t|\mu_t, \theta, \mathbf{V}_t; \mathbf{X}_t, \mathbf{Z}_t, \alpha_t).$$

4.2 Estimation algorithm

We estimate the structural parameters via a Metropolis-Hastings within Gibbs sampler (Robert and Casella (2005)). The "outer loop" is a Gibbs sampler which loops sequentially over three conditional densities: the conditional density of θ, the conditional density related to equilibrium selection of multiple equilibrium prices, and the conditional density of the latent valuations V_{ij} (the augmented component). Since these are difficult to sample from directly, within each Giibs step we use a Metropolis-Hastings approach to obtain draws from these three conditional densities. (Complete details are given in the Appendix.)[22]

4.3 Specification details

We end this section with some details of our model specification. First, in our model we allow the click volumes (the α's) to vary non-monotonically across positions, in line with the evidence presented earlier in Figure (1). However, in our data, there are instances when the top position may receive zero clicks (cf. Table 1), but still attract bidders. Such an occurence

[22]The procedure closely follows that of Albert and Chib (1994), and Logan et al. (2008).

would not be explicable in our model unless we allow for bidders to be uncertain about the click volume at the time they are submitting their bids, and hence bid based on expectations about the click volume (so that the realized click volume may not coincide with the ex ante expectation).[23] Hence, we allow for this by assuming that bidders' beliefs about α_{jt}, the click volume of the j-th position at auction t, satisfies the following shifted log-normal process (and are independent of their valuations V_{ij}):

$$\log(\alpha_{jt}+1) = \gamma_{j0} + \gamma_{j1}N_t + \gamma_{j2}\text{specific}_t + \gamma_{j3}\text{promotional}_t + \gamma_{j4}\text{brand}_t + \eta_{jt},$$

where N_t is the number of available positions and η_{it} follows $\mathcal{N}(0,\sigma^2)$.[24] The unknown parameters can then be estimated by OLS, and the expected click volume is given by

$$E[\alpha_{jt}] = \exp(\gamma_{j0} + \gamma_{j1}N_t + \gamma_{j2}\text{specific}_t + \gamma_{j3}\text{promotional}_t + \gamma_{j4}\text{branded}_t + \sigma_j^2/2) - 1.$$

We use the estimated $E[\alpha_{jt}]$ instead of α_{jt} in our estimation procedure.[25]

Second, for bidders' valuations V_{ij} we consider two specifications which contain interactions of bidder-specific and position-specific variables:

Model I

$$\begin{aligned} V_{ijt} &= \beta_0 + \beta_1\text{specific}_t \\ &+\beta_2\text{promotional}_t + \beta_3\text{brand}_t \\ &+\beta_4\text{high}_i\tfrac{1}{j} \qquad +\beta_5\text{medium}_i\tfrac{1}{j} \\ &+\beta_6\text{high}_i\tfrac{1}{j}\times\text{specific}_t \qquad +\beta_7\text{medium}_i\tfrac{1}{j}\times\text{specific}_t \\ &+\beta_8\text{high}_i\tfrac{1}{j}\times\text{promotional}_t \qquad +\beta_9\text{medium}_i\tfrac{1}{j}\times\text{promotional}_t \\ &+\beta_{10}\text{high}_i\tfrac{1}{j}\times\text{branded}_t \qquad +\beta_{11}\text{medium}_i\tfrac{1}{j}\times\text{brand}_t + \sigma\epsilon_{ijt} \end{aligned} \tag{11}$$

where ϵ_{ijt} is an i.i.d. standard normal random sequence. In this specification, we impose a hyperbolic decay $(1/j)$ in coefficients across positions. Note that while this specification imposes monotonicity in the coefficients across positions, the trend is allowed to be either decreasing or increasing (depending on the sign of the β's).

[23] Indeed, without this, zero click volume for a position would imply zero transfer for the top position, which would unreasonably imply zero transfer for all positions in the model.

[24] A similar specification is used in Yang et al. (2014), but here we do not impose that α_i decays from high to low positions.

[25] Even after regression smoothing, $E[\alpha_{jt}]$ is still nondecreasing in j in some cases. An important restriction of the assignment game framework is that we cannot allow the (expected) click volume to be bidder-specific. We will return to this point when discussing the counterfactual simulations below.

Model II

$$
\begin{aligned}
V_{ijt} &= \beta_0 + \beta_1 \text{specific}_t \\
&+\beta_2 \text{promotional}_t + \beta_3 \text{brand}_t \\
&+\beta_{1j} \text{high}_i \qquad\qquad +\beta_{2j} \text{medium}_i \\
&+\beta_{3j} \text{high}_i \times \text{specific}_t \qquad +\beta_{4j} \text{medium}_i \times \text{specific}_t \\
&+\beta_{5j} \text{high}_i \times \text{promotional}_t \qquad +\beta_{6j} \text{medium}_i \times \text{promotional}_t \\
&+\beta_{7j} \text{high}_i \times \text{brand}_t \qquad +\beta_{8j} \text{medium}_i \times \text{brand}_t + \sigma\epsilon_{ijt} \\
&\text{if } j \leq 5 \text{ and } = \sigma\epsilon_{ijt} \text{ if } j > 5
\end{aligned} \tag{12}
$$

This is a more flexible specification which allows each of the β coefficients to be position specific, and no longer imposes a monotonic trend as in Model I. (For tractability, we assume these coefficients to be zero for positions lower than 5, as click volumes and per-click prices for such lower positions are generally so small that they can be ignored.)[26]

Comparison: model specification in existing papers. Our specification details here contrast with those in Varian (2007), Edelman et al. (2007), and Athey and Nekipelov (forthcoming), who consider a multiplicative specification of bidder's valuations:

$$V_{ij} = v_i \alpha_j. \tag{5}$$

In tandem with the assumption that click volumes decrease in ad position ($\alpha_1 > \alpha_2 \cdots > \alpha_N$), this multiplicative specification (5) implies that every bidder has exactly the same preference ordering over positions: everyone prefers a higher position. Subsequently, it is easy to show that the optimal allocation is perfectly assortative matching: the k-th position is assigned to the bidder having the k-th highest valuation. In contrast, our specifications of valuations and click volume described earlier allow us to accommodate the richer patterns in allocations and click volumes which were evidenced in the data. ■

5 Estimation Results

Table (3) contains the results of the Model I specification. Since the results are similar using both the uniform and beta equilibrium selection rules, we will focus on the uniform rule (results in third column of Table 3) in the discussion. We find evidence for positive assortative matching between bidders' quality and ad positions. The posterior mean of coefficients of high and medium quality dummies are all positive (688.96 and 479.75, respectively), and distinct from zero (larger than 3 times the standard deviation). Furthermore, the magnitude

[26] See the appendix for a discussion of parameter identification for these specifications of preferences.

of the coefficient of high quality is larger than that of medium, implying that the top quality bidders have the highest valuations, all else equal.

However, there is substantial heterogeneity. We see that this positive assortative matching pattern is further amplified in auctions containing product-specific keywords (the interactions with "specific" are 539.45 and 516.41 for high and medium bidders, respectively). Since the positive assortative matching is extremely strong in product-specific keyword auctions, it is less likely that low quality bidder would win top positions here. In contrast, the interactions with "brand" are strongly negative (-341.05 for high-quality bidders, -395.04 for medium quality), and offset the positive coefficients described earlier. This finding suggests that high quality online (camera) merchants have a relatively low assessment of keyword strings containing brand names, so that lower quality merchants stand a higher chance of winning such auctions. (The interactions are likewise negative for the promotional keywords, but not as large in magnitude.)

These results imply horizontal differentiation across different types of keywords, as high quality merchants have relatively higher valuations for keywords including specific model names, and relatively lower valuations for other types of keywords. This difference may be explained by heterogeneity in the consumers who use the different keyword queries. For cameras, keyword queries with model-specific keywords result in the most narrow range of search results. Major camera manufacturers usually use unique model numbers to distinguish their products from others. For example, Nikon's DSLR (digital single-lens reflexive) camera models typically start with a "D" followed by numbers; e.g., D3, D90, D300s, etc. Similarly, Canon models typically begin with numbers followed by "D"; e.g., 550D and 5D.[27] Hence, shoppers querying with model-specific keywords are probably well-informed consumers who have a clear idea which specific products they are interested in, and are searching with a strong intention of purchasing. Our results indicate, then, that high quality online merchants, who are typically also more experienced, gravitate towards more narrowly defined keyword queries which are likely to be made by serious buyers.

In contrast, shoppers who use brand names to search may be more interested in browsing and collecting information about different camera models; searching for, say, "Nikon" will return a wide variety of models and accessories across many price points. These consumers may have a more muted intention of purchasing, and our results imply that high quality online merchants – again, those who are more experienced – have correspondingly lower assessments of these keywords. Finally, the design of WebsiteX ads partially neutralizes the effect of additional promotional terms. Most of the sponsored links directly contain price information (Figure 3), making it extremely easy for consumers to compare different prices

[27]Fujifilm uses the combination of "X" and numbers, Sony uses "A" and numbers, Pentax use "K" and numbers, and Olympus usually starts with "E".

on the internet. Hence, their purchase probabilities are unlikely to be swayed by purely marketing terms such as "big sale", rendering these promotional keywords of little value to the more-experienced high-quality merchants.

Table (4) and (5) summarize the estimation results of model II. Again, we focus on the results for the uniform equilibrium selection rule (in Table 4). The main differences between the Model I and Model II specifications is that the latter allows the parameters in bidder's valuations to be completely flexible vis-a-vis position rank. But even after allowing this flexibility, we find qualitatively similar patterns in valuations compared with the more restrictive model I results. As before, we find that high-quality bidders have relatively higher valuations for top positions in keyword queries involving product-specific keywords, but they have relatively lower valuations for top positions in queries involving promotional of brand-specific keywords. Thus our finding of horizontal differentiation in preferences appears robust to different specifications of bidders' valuations.

Quantitatively, the more flexible Model II specification does yield some additional findings. In some cases, it is the second position that generates the largest value for high quality merchants, not the first position; for specific keywords, as an example, high quality merchants value position 2 most highly: the coefficient for position 2 is 203.12=230.99-27.87, while for position 1 it is only 112.22=16.39+95.83. Similarly, medium-quality merchants also value the second position most highly in specific keyword queries. This phenomena may be related to the empirical fact we found in the click volume data, that often it is the second position that generates the largest click volume, even after the regression smoothing. From the merchants' perspective, the second position is almost as good as the first slot, because the click volume is comparable with the top but the per-click price is lower.

However, when the valuation matrix is not positively assortative, which is the case in our results, then it is unclear whether the GSP mechanism is socially optimal. As pointed out by Athey and Ellison (2011) and Chen and He (2011), the sponsored-link auction also plays the role of information intermediary. If the links are sorted according to merchants' quality, then it allows the consumer to efficiently search for merchants who fit their quality needs. If the valuation matrix is positively assortative between quality and ranking, then high quality merchants will bid aggressively and hence GSP is an efficient way to convey information to online shoppers.

On the other hand, as we noted before, WebsiteX merchants typically post their prices on the sponsored ads, thus providing consumers with the most important piece of purchase-related information without requiring further clicking behavior. In this setting, consumer search may be less relevant and, hence, the click volume in WebsiteX is less regular than described in the literature motivated from Yaoo!, Google and Microsoft; e.g., Varian (2006), Athey and Nekipelov (forthcoming).[28]

Finally, the estimated coefficients of the Beta distribution parameters a and b are positive, with $a >> b$, for both the Model I and model II estimates (bottom of tables 3 and 5). This implies that the equilibrium selection density function is left-skewed, so that higher prices are much more likely to be chosen. This suggests that the observed prices are better explained as optimal for WebsiteX, rather than for the merchants.

6 Counterfactual: The Effects of Bid Scoring

Using our estimation results, we now turn to the main policy question of this paper, which is assessing the effects of the bid-scoring policy which WebsiteX enacted only shortly after our sample period. This bid scoring was implemented using a score-weighted version of the GSPA (which we will call WGSPA hereinafter); that is, the positions are allocated by ranking the product of each bid times a "popularity score" for this bidder. Specifically, letting $\kappa_i, i = 1, \ldots, N$ denote the bidder-specific popularity scores, the positions are assigned according to the weighted bids: $\kappa_1 b_1 > \kappa_2 b_2 > \cdots > \kappa_N b_N$. Furthermore, the bidder winning position i pays a per-click price p_{ii} such that bidder i's score $\kappa_i p_{ii}$ is exactly equal to the score of the bidder in the $i+1$-th position:

$$\kappa_i p_{ii} = \kappa_{i+1} b_{i+1} \quad \text{or} \quad p_{ii} = \frac{\kappa_{i+1}}{\kappa_i} b_{i+1}. \tag{13}$$

The total payment of a bidder in the i-th position is $\alpha_i p_{ii}$. This mechanism essentially rewards the high quality advertisers with price discounts (if $\kappa_i > \kappa_{i+1}$, then $p_{ii} < b_{i+1}$, while $p_{ii} = b_{i+1}$ under the unscored GSPA rule); at the same time, this also incentivizes online merchants to improve their quality. Intuitively, offering price discounts may reduce the platform's revenue.[29] We show, however, it is possible that the per-click price can be even higher under WGSPA.

Importantly, in the WGSPA, the implicit "price" that the bidder in position i must pay to obtain position j differs from bidder to bidder (depending on their score κ_i), and hence the resulting game cannot be formulated as an assignment game à la Shapley-Shubik (in which agents are essentially price-takers). Suppose a generic bidder is indexed by i, and a generic position and the bid paid for that position is indexed by j. An allocation μ is a one-to-one function that maps each bidder's index to the corresponding position index. $\mu(i) = j$ means

[28] While it is difficult to collect price data from WebsiteX, we collected a limited sample of screen captures and found no noticeably trends in product prices across ad positions.

[29] Indeed, Myerson's (1981) classic work suggests that, for standard auctions, a platform may be able to raise expected revenues by discriminating *against* high-valuation bidders.

that bidder i is assigned to the j-th position, and the inverse mapping, $\mu^{-1}(j) = i$, would identify who is assigned to the j-th position.

Since the WGSPA cannot be formulated as an assignment game, the equilibrium allocation may not be unique. An allocation and a sequence of bids $(\mu; b_1, \ldots, b_N, b_{N+1})$ constitutes a symmetric Nash equilibrium if the following inequalities are satisfied[30]:

1. $\kappa_{\mu^{-1}(1)} b_1 > \kappa_{\mu^{-1}(2)} b_2 > \cdots > \kappa_{\mu^{-1}(N+1)} b_{N+1}$ (Allocation Rule)
2. $\alpha_i p_{\mu^{-1}(i)i} \leq V_{\mu^{-1}(i)i}$ for all i (Individual Rationality)
3. Incentive compatibility:

$$V_{\mu^{-1}(i)i} - \alpha_i p_{\mu^{-1}(i)i} \geq V_{\mu^{-1}(i)j} - \alpha_j p_{\mu^{-1}(i)j}, \quad \text{for all } (i, j)$$

where the counterfactual deviation prices are defined by: $p_{\mu^{-1}(i)j} = \dfrac{\kappa_{\mu^{-1}(j+1)}}{\kappa_{\mu^{-1}(i)}} b_{j+1}$

Sources at WebsiteX tell us that one key element in determining their score index κ_i is the historical performance and click volume of the advertiser i, which is highly correlated with WebsiteX's own quality rating system (see section 4). Therefore, we perform the counterfactual analysis under two alternative scoring systems (1) a *coarser* one with $\kappa_i \in \{1, 2, 3\} \equiv$ {low, medium, high quality}; and (2) a *finer* one with $\kappa_i \in \{3, 4, \ldots, 15\}$.[31]

We simulate the per-click prices $\tilde{\mathbf{p}}_t$ under the new mechanism. We also compute the corresponding componentwise upper bound $\bar{\mathbf{p}}_t$, and the componentwise lower bound $\underline{\mathbf{p}}_t$. Due to the computational burden of computing SNE of WGSPA, we only consider auctions with no more than 7 positions. Complete implementation details are given in the appendix. The summary statistics of the simulated (cross-sectional) per-click price distribution under different model specifications are summarized in Tables 6 and 7.

Results. For simplicity, we focus on the results for the coarser scoring rule, in Table 6, in our discussion. First, bid scoring appears to "steepen" the price gradient across positions, with the top position increasing in price but lower positions decreasing in price. Specifically, in the Model I-Uniform results, the price per click for the top position increases, on average, by 12 RMB (around 2 USD) relative to the baseline unscored scenario, while price for the second position decreases by 3 RMB (around \$0.50). Results are similar for Model II: there, the price for the top position increases by 6 RMB (around \$1) and the price for the second position also increases (by 2 RMB = \$0.30), but the prices for all lower positions decrease relative to the baseline scenario.

The price increases for the top positions are striking, especially as the bid-scoring rules were intended to reward popular merchants (those with high scores) with price discounts. To

[30] See Varian (2007) for more details.

[31] As we do not have the data for the losers, we assume κ_{N+1} to be the lowest score.

assess the extent of price discounting under the bid-scoring system, in Table 6 we also provide summary statistics for the "bid-discount ratios" $\frac{\kappa_{j+1}}{\kappa_j}$ for position j, which are a measure of the price discount (cf. Eq. (13)). When this ratio is less than one, then the bidder winning position j was given a price discount, while if it exceeds one, then the bidder paid a price premium. We see that, across all four specifications, and across the top four positions, this ratio was less than one in over 80% of the simulations. This implies that the prices changed in response to bid scoring *despite* the winning bidders being given price discounts. Apparently the price increases for the top positions were triggered by more aggressive bidding in response to the introduction of bid scoring.

In Table 6, we also provide summary statistics for the bid-discount ratios $\frac{\kappa_{j+1}}{\kappa_j}$ for the baseline unscored scenario. There, we see that these statistics are not much different from those generated in the bid scoring scenarios. This implies that bid-scoring does not increase the degree of assortative matching in the allocation positions – it does not appear that larger, more popular merchants were systematically more likely to end up in top positions after the move to bid scoring. Thus while the fears of the small merchants quoted in this paper's introduction – that they have to pay much higher prices to get their ads in top positions – seem justified to some extent by our simulation results, at the same time popular merchants are not more likely (relative to the baseline scenario) to get these top positions.

Finally, the right-hand side of Table 6 shows that total platform revenue from the auctions remains unchanged from bid scoring. At the least, the finding that revenue does not decrease upon the introduction of bid scoring may justify WebsiteX's move toward WGSPA: it can promote long term benefits beyond the scope of our analysis (such as improvements in ad quality, and buyer responsiveness to ads) without sacrificing short term revenue.

7 Conclusion

We conclude with a brief summary of some main points from our study. Empirically, we uncover some horizontal differentiation between different types of keywords, as high quality merchants have relatively higher valuations for informative keywords, and relatively lower valuations for vague keywords. This suggests that high quality bidders, who are typically also more experienced merchants, gravitate towards more narrowly defined keyword queries which are likely to be made by serious buyers, and leave the top positions in broader keyword queries for less experienced (lower quality) merchants. Counterfactual evaluations show that the price trend becomes steeper under the score weighted generalized second price auction, with much higher prices obtained for the top position but lower prices for the other positions. Apparently, while scoring the auction grants price discounts to popular merchants, our findings suggest that scoring also heightens the bid competition, thus leading to higher prices for top positions.

Overall, we do not find large effects on platform revenue and sorting patterns from shifting to a scoring rule.

Our methodological approach is also novel as it is motivated by the equivalence between symmetric Nash equilibria in GSPA, and stable outcomes in the classic assignment game of Shapley and Shubik (1972). To accommodate some stylized empirical facts in sponsored search auctions, our specification generalizes previous work by allowing bidders to have preferences for positions which are not multiplicative in bidder- and position-specific effects, and click volumes to be nondecreasing with position ranks. For estimation, we utilize a Bayesian procedure and develop a Metropolis-Hastings within Gibbs sampler.

Inevitably, some strong assumptions underlie our analytical framework in this paper. It is possible, albeit at the cost of computational expense, to extend our Bayesian estimation procedure to a more general setup, in order to accommodate asymmetric Nash equilibria (as studied in Börgers et al (2013)), or bidder-specific click volumes (as in Jeziorski and Moorthy (2014)). We leave these for future inquiry.

References

Akkus, O., J. A. Cookson and A. Hortacsu (2013). The Determinants of Bank Mergers: A Revealed Preference Analysis, working paper.

Albert, J. H. and Chib, S. (1993). Bayesian Analysis of Binary and Polychotomous Response Data, *Journal of the American Statistical Association*, **88**, 669–679.

Athey, S. and G. Ellison (2011). Position Auctions with Consumer Search, *Quarterly Journal of Economics*, **126**, 1213–1270.

Athey, S. and D. Nekipelov (forthcoming). A Structural Model of Sponsored Search Advertising Auctions, *Journal of Political Economy*, to appear.

Bajari, P., H. Hong, and S. Ryan (2010). Identification and Estimation of Discrete Games of Complete Information, *Econometrica*, **78**, 1529–1568.

Becker, G. S. (1973). A theory of marriage: part I, *Journal of Political Economy*, **81**, 813–846.

Beresteanu, A. and I. Molchanov and F. Molinari (2011). Sharp Identification Regions in Models with Convex Moment Predictions, *Econometrica*, **79**, 1785-1821.

Börgers, T. and I. Cox and M. Pesendorfer and V. Petricek (2013). Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence, *American Economic Journal: Microeconomics*, **5**, 61-71.

Boyd, D., H. Lankford, S. Loeb and J. Wyckoff (2006). Analyzing the Determinants of the Matching of Public School Teachers to Jobs: Estimating Compensating Differentials in Imperfect Labor Markets, NBER working paper.

Choo, E. and A. Siow (2006). Who Marries Whom and Why, *Journal of Political Economy*, **114**, 175–201.

Chen, Y., and C. He (2011). Paid Placement: Advertising and Search on the Internet, *Economic Journal*, **121**, 309–328.

Chernozhukov, V., H. Hong, and E. Tamer (2007). Estimation and Inference on Identified Parameter Sets, *Econometrica*, **75**, 1243–1284.

Chiappori, P.-A., R. J. McCann and L. P. Nesheim (2009). Hedonic Price Wquilibria, Stable matching, and Optimal Transport: Equivalence, Topology, and Uniqueness, *Economic Theory*, **42**, 317–354.

Ciliberto, F. and E. Tamer (2009). Market Structure and Multiple Equilibria in Airline Markets. *Econometrica*, **77**, 1791–1828.

Dantzig, G. B. (1968). *Linear Programming and Extensions*, Princeton University Press, Princeton, NJ.

Demange, D., D. Gale and M. Sotomayor (1986). Multi-Item Auctions, *Journal of Political Economy*, **94**, 863–872.

Echenique, F., S. Lee, M. Shum, and M. Yenmez (2013). The Revealed Preference Theory of Stable and Extremal Stable Matchings, *Econometrica*, **81**, 153–171.

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**, 242–259.

Fox, J. T. (2013). Estimating Matching Games with Transfers, working paper.

Fox, J. T. and P. Bajari (2013). Measuring the Efficiency of an FCC Spectrum Auction, *American Economic Journal: Microeconomics*, **5**, 100–146.

Galichon, A. and M. Henry (2011). Set Identification in Models with Multiple Equilibria, *Review of Economic Studies*, **78**, 1264-1298.

Galichon, A. and B. Salanié (2012). Cupids Invisible Hand: Social Surplus and Identification in Matching Models. *Working Paper*.

Goldman, M. and J. Rao (2014). Experiments as Instruments: Heterogeneous Position Effects in Sponsored Search Auctions, working paper, UCSD.

Graham, B. S. (2011). Econometric Methods for the Analysis of Assignment Problems in the Presence of Complementarity and Social Spillovers, *Handbook of Social Economics*, **1B**, 965–1052, Amsterdam, North-Holland.

Haile, P. A. and E. Tamer (2003). Inference with an Incomplete Model of English Auctions, *Journal of Political Economy*, **111**, 1–51.

Hong, H. and M. Shum. (2003). Econometric Models of Asymmetric Ascending Auctions, *Journal of Econometrics*, **112**, 327-358.

Hsieh, Y.-W. (2011). Understanding Mate Preferences from Two-Sided Matching Markets: Identification, Estimation and Policy Analysis, manuscript, New York University.

Jerath, K., L. Ma, Y.-H. Park and K. Srinivasan (2011). A Position Paradox in Sponsored Search Auctions, *Marketing Science*, **30**, 612–627.

Jeziorski, P. and S. Moorthy (2014). Advertiser Prominence Effects in Search Advertising, working paper, UC Berkeley, Haas School of Business.

Jeziorski, P. and I. Segal (2015). What makes them Click: Empirical Analysis of Consumer Demand for Search Advertising, *American Economic Journal: Microeconomics*, **7**, 24-53.

Krasnokutskaya, E. and K. Seim (2011). Bid Preference Programs and Participation in Procurement, *American Economic Review*, **101**, 2653-2686.

Logan, J. A., P. D. Hoff, and M. A. Newton (2008). Two-Sided Estimation of Mate Preferences for Similarities in Age, Education, and Religion, *Journal of the American Statistical Association*, **103**, 559–569.

McAfee, P. and J. McMillan (1989). Government Procurement and International Trade, *Journal of International Economics*, **26**, 291-308.

Menzel, K. (2011). MCMC Methods for Incomplete Structural Models of Social Interactions, *Working Paper.*

Myerson, R. (1981). Optimal Auction Design. *Mathematics of Operations Research*, **6**, 58-73.

Roth, A. E. and M. A. O. Sotomayor (1990). *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Econometric Society Monographs No. 18, Cambridge University Press.

Shapley, L. S. and M. Shubik (1972). The Assignment Game, I: The Core *International Journal of Game Theory*, **1**, 111–130.

Tanner, M. and W. Wong (1987). The calculation of posterior distributions by data augmentation, *Journal of the American Statistical Association*, **82** 528-540.

Uetake, K. and Y. Watanabe (2012). Entry by merger: Estimates from a Two-sided Matching Model with Externalities, working paper, Yale School of Management.

Varian, H. R. (2007). Position Auction, *International Journal of Industrial Organization*, **25**, 1163–1178.

Yang, S., S. Lu and X. Lu (2014). Modeling Competition and Its Impact in Paid-Search Advertising, *Marketing Science*, **33**, 134-153.

Figure 1: Boxplot of Log Click Volume

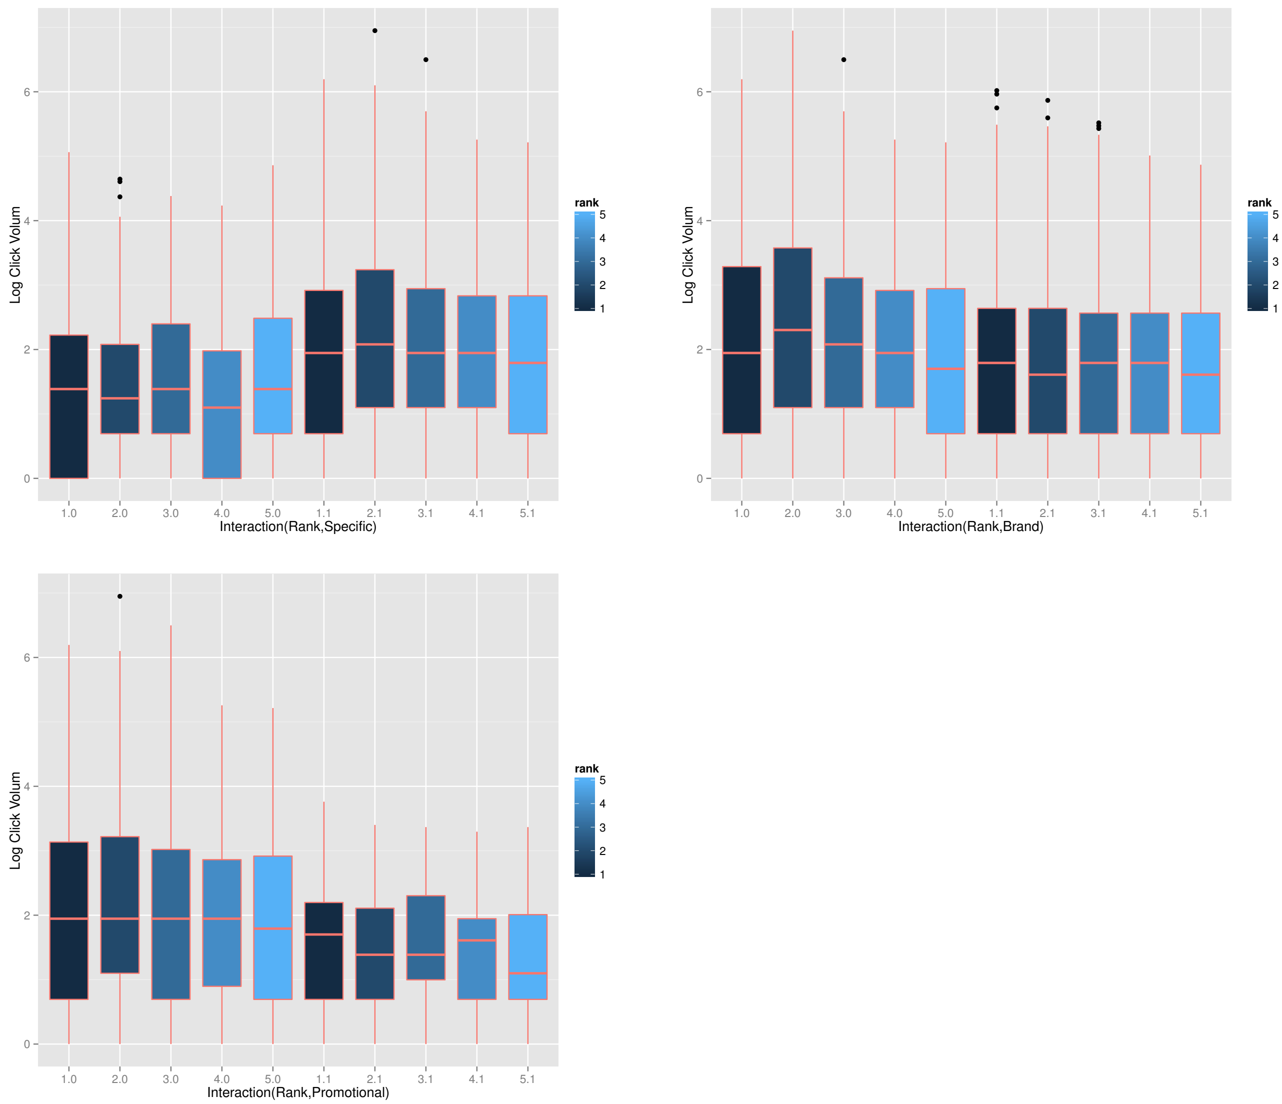

*The pair of numbers on the horizontal axis represents the values of two variables. For example, Interaction(Rank,Specific)= 3.1 on the x-axis of the top-left graph means "3rd position, Specific dummy = 1"

Figure 2: Boxplot of Per-Click Price

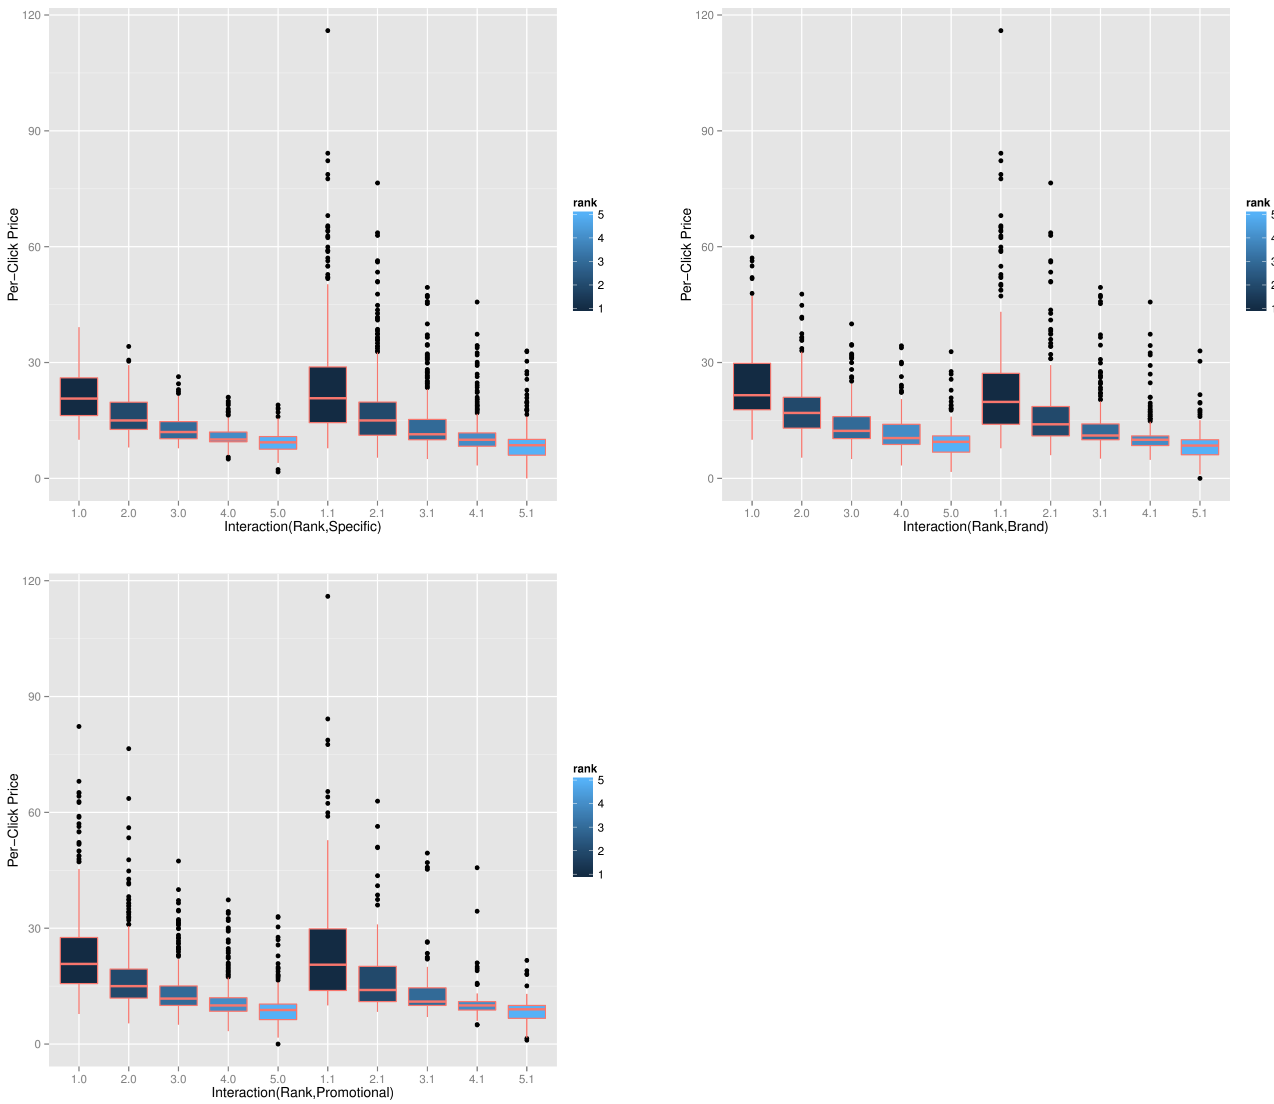

*The pair of numbers on the horizontal axis represents the values of two variables. Per-click prices are given in Chinese Renminbi (RMB, or "yuan"), with exchange rate roughly 6RMB=1USD. For example, Interaction(Rank,Specific)= 3.1 on the x-axis of the top-left graph means "3rd position, Specific dummy = 1"

Figure 3: Sample Search Results in WebsiteX

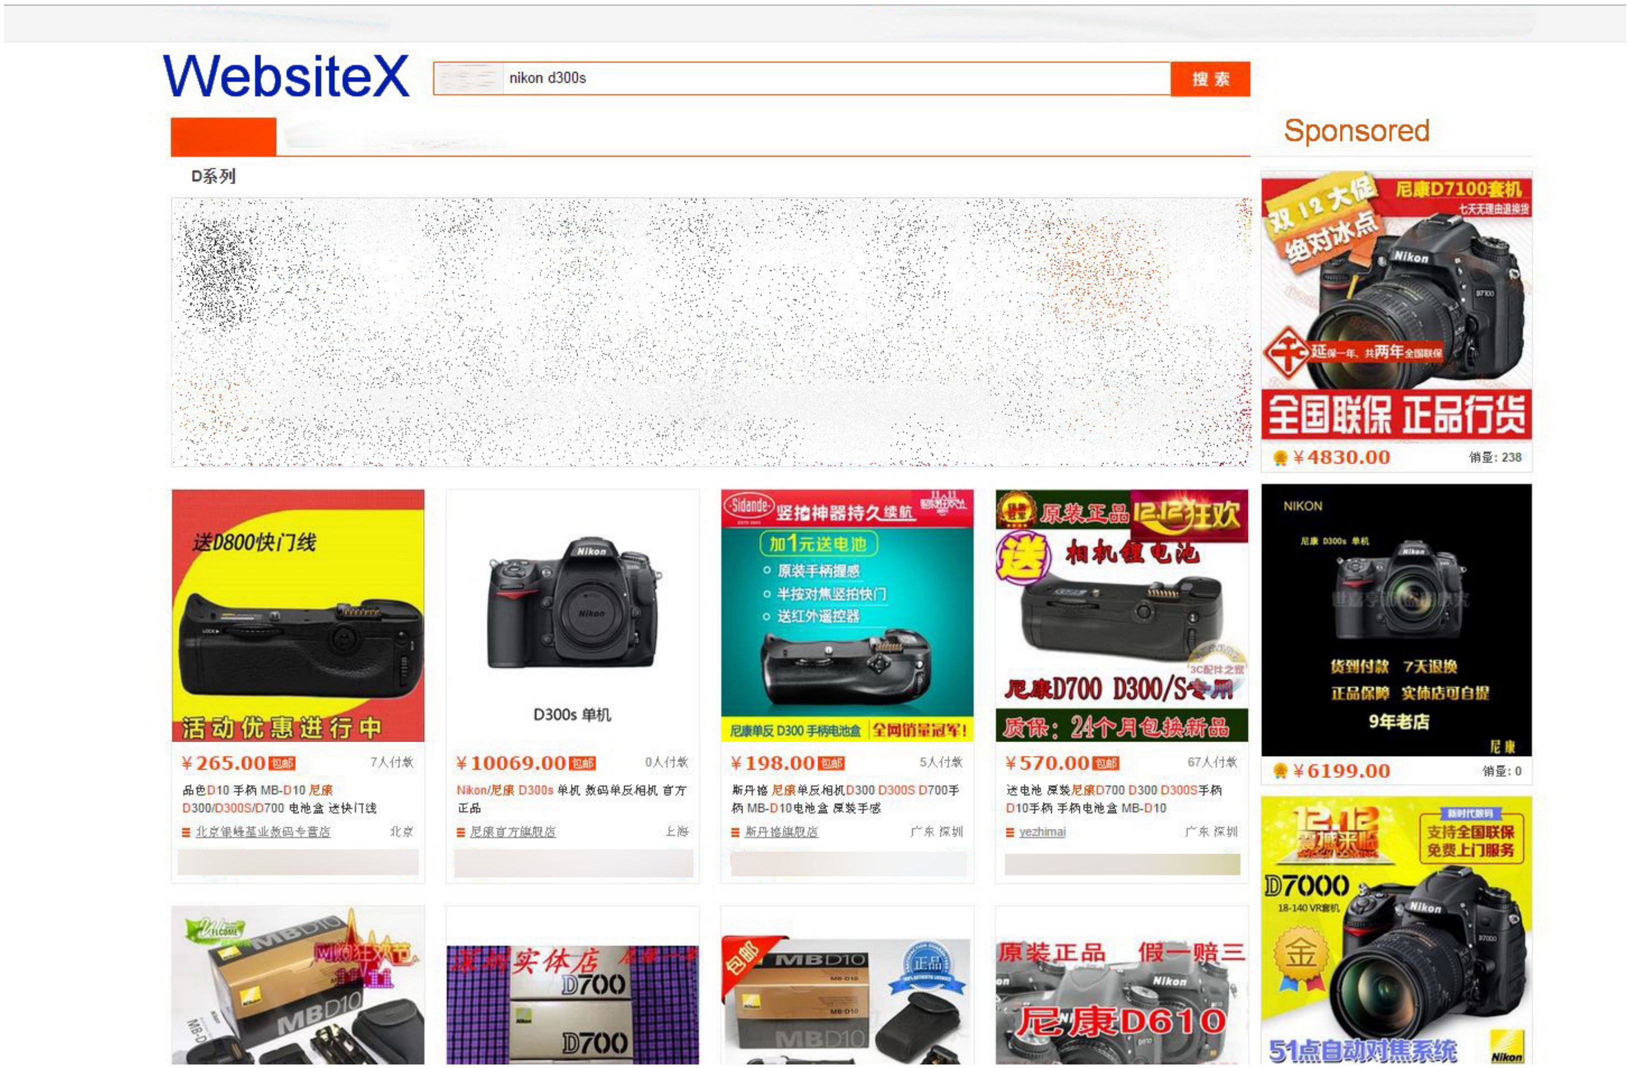

Figure 4: Non-Existence of Equilibrium Price System of GSPA

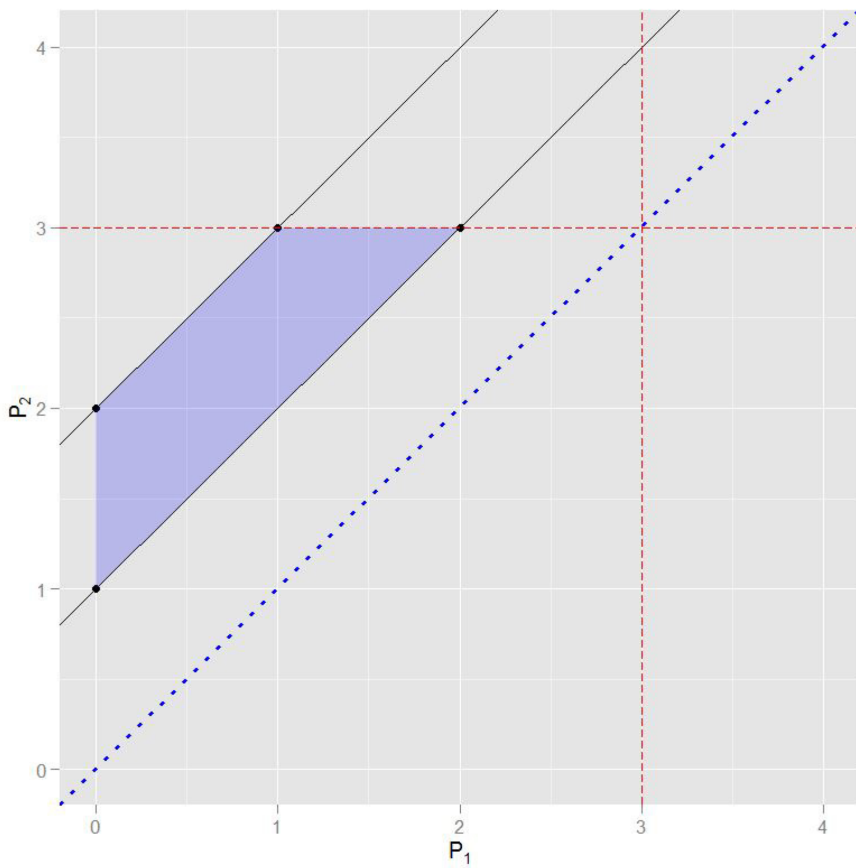

solid black line: no-blocking pair condition; dashed red line: individual rationality condition; blue dotted line: GSPA condition $p_1 > p_2$; shaded area: set of stable matching without imposing $p_1 > p_2$.

Table 1: Variable List and Summary Statistics

Variable	Definition	mean	min	max
Auction Characteristics				
Per-Click Price (RMB)		13.25	0	115.95
		(9.25)		
Click Volume		16.37	0	1040
		(42.98)		
Keyword Characteristics				
Brand	= 1 if keyword includes the brand name	0.63	0	1
		(0.48)		
Specific	= 1 if keyword includes a specific model number	0.77	0	1
		(0.42)		
Promotional	= 1 if keyword includes promotional terms	0.18	0	1
		(0.39)		
Bidder Characteristics				
High Quality	= 1 for high quality merchant	0.2	0	1
		(0.03)		
Medium Quality	= 1 for medium quality merchant	0.74	0	1
		(0.03)		

*Standard error in parentheses.

Table 2: Contingency Table of Bidders' Quality versus Position Ranks

All Keywords					
	Position 1	2	3	4	5
High	7.2*	4.9	3.7	3.8	2.6
Medium	12.6	14.6	15.6	14.9	16.1
Low	0.2	0.5	0.7	1.3	1.3
Keywords with Specific= 1					
	Position 1	2	3	4	5
High	7.1	4.7	3.7	4.1	2.9
Medium	12.7	14.7	15.8	15.0	16.2
Low	0.2	0.6	0.5	0.9	0.9
Keywords with Brand= 1					
	Position 1	2	3	4	5
High	7.5	4.6	3.7	3.7	2.2
Medium	12.3	14.7	15.5	15.0	16.6
Low	0.2	0.7	0.7	1.4	1.2
Keywords with Promotional= 1					
	Position 1	2	3	4	5
High	6.6	3.4	3.2	2.7	2.0
Medium	13.2	16.1	16.6	15.9	17.0
Low	0.2	0.5	0.2	1.4	0.9

*All numbers in the table are given as percentage of all bidders.

Table 3: MCMC Estimation: Model I

Dummy Regressor	Posterior	Model I Uniform Eq. Sel.	Model I Beta Eq. Sel.
constant	mean	-151.57	-214.89
	s.d.	(5.92)	(8.00)
specific	mean	-63.42	-21.96
	s.d.	(6.82)	(7.25)
promotional	mean	-14.58	-35.65
	s.d.	(10.10)	(12.69)
brand	mean	39.65	40.89
	s.d.	(6.27)	(9.98)
high	mean	688.96	704.04
	s.d.	(19.41)	(26.86)
medium	mean	479.75	563.92
	s.d.	(12.75)	(16.26)
high×specific	mean	539.45	382.34
	s.d.	(21.65)	(21.84)
medium×specific	mean	516.41	422.27
	s.d.	(16.60)	(15.21)
high×promotional	mean	-42.53	-139.99
	s.d.	(23.43)	(38.72)
medium×promotional	mean	-161.82	-86.42
	s.d.	(20.68)	(26.22)
high×brand	mean	-341.05	-307.86
	s.d.	(18.08)	(25.68)
medium×brand	mean	-395.04	-405.03
	s.d.	(14.67)	(18.03)
σ^2	mode	40484	49429
	mean	40385	49513
	s.d.	(793)	(366)
left parameter of beta distr. a	mean	N/A	3.27
	s.d.	N/A	(0.59)
right parameter of beta distr. b	mean	N/A	0.83
	s.d.	N/A	(0.09)

Table 4: MCMC Estimation: Model II-Uniform Equilibrium Selection

		Rank of Positions				
Dummy Interaction Terms	Posterior	1	2	3	4	5
high	mean	95.83	-27.87	-360.55	-431.62	-463.62
	s.d.	(15.34)	(26.93)	(16.00)	(22.21)	(24.45)
medium	mean	-175.62	-140.20	-281.94	-424.62	-359.59
	s.d.	(26.16)	(23.98)	(19.24)	(31.72)	(25.66)
high×specific	mean	16.39	230.99	-125.74	-157.83	-294.49
	s.d.	(34.01)	(28.67)	(38.07)	(44.10)	(52.92)
medium×specific	mean	-15.96	124.32	-190.90	-236.87	-332.36
	s.d.	(24.92)	(24.99)	(28.15)	(13.95)	(23.43)
high×promotional	mean	-356.05	-336.30	-153.56	-241.04	-78.22
	s.d.	(22.89)	(26.42)	(25.57)	(33.76)	(21.71)
medium×promotional	mean	-288.32	-324.80	-264.19	-305.36	-313.11
	s.d.	(21.04)	(14.11)	(18.97)	(31.80)	(21.93)
high×brand	mean	-109.84	-454.61	-47.33	-5.62	-16.25
	s.d.	(18.72)	(27.74)	(30.65)	(27.37)	(37.19)
medium×brand	mean	-114.13	-385.46	-50.34	96.47	67.63
	s.d.	(35.35)	(22.66)	(19.39)	(26.89)	(23.22)
Non-Interaction Terms		Constant	Specific	Promotional	Brand	
	mean	366.29	306.39	208.95	-70.1	
	s.d.	(20.82)	(21.94)	(14.89)	(19.79)	
σ^2	mode	36579				
	mean	36849				
	s.d.	(330)				

Table 5: MCMC Estimation: Model II-Beta Equilibrium Selection

		Rank of Positions				
Dummy Interaction Terms	Posterior	1	2	3	4	5
high	mean	19.82	-102.09	-400.17	-483.75	-522.48
	s.d.	(19.67)	(28.49)	(17.07)	(25.14)	(24.15)
medium	mean	-229.56	-186.11	-338.50	-479.67	-416.46
	s.d.	(15.04)	(15.60)	(20.58)	(14.98)	(20.22)
high×specific	mean	-217.75	-15.15	-361.31	-401.81	-490.52
	s.d.	(24.10)	(54.45)	(52.33)	(31.95)	(25.76)
medium×specific	mean	-247.19	-125.35	-426.39	-471.28	-565.03
	s.d.	(44.74)	(35.69)	(50.40)	(40.62)	(29.00)
high×promotional	mean	-334.95	-363.30	-179.40	-174.20	-124.02
	s.d.	(34.52)	(40.65)	(30.01)	(26.91)	(57.01)
medium×promotional	mean	-323.63	-351.31	-304.38	-320.22	-345.52
	s.d.	(31.75)	(31.43)	(27.12)	(27.11)	(19.50)
high×brand	mean	-189.93	-495.68	-134.79	-71.69	-101.03
	s.d.	(20.33)	(20.66)	(16.70)	(24.47)	(23.43)
medium×brand	mean	-202.43	-459.60	-117.65	19.30	-3.05
	s.d.	(17.07)	(17.93)	(26.27)	(24.02)	(22.43)
Non-Interaction Terms		Constant	Specific	Promotional	Brand	
	mean	397.12	549.03	240.05	10.14	
	s.d.	(13)	(34.03)	(26.28)	(17.96)	
σ^2	mode	32767				
	mean	32911				
	s.d.	(668)				
left parameter of beta distr. a	mean	2.62				
	s.d.	(0.43)				
right parameter of beta distr. b	mean	0.73				
	s.d.	(0.06)				

Table 6: Counterfactual per-click prices for Symmetric Nash Equilibrium of WGSPA: Finer Score

(Units of prices are Chinese RMB.)

		Rank of Positions					Total Revenue
Price-Raw Data	Statistics	1	2	3	4	5	
	mean	22.32	16.13	12.68	10.27	8.12	1133
	median	19.6	13.67	11	9.74	7.73	910
	variance	139.62	62.8	32.15	18.95	12.16	616476
$Pr(\frac{\kappa_{i+1}}{\kappa_i} \leq 1)$		0.70	0.61	0.63	0.65	0.80	
Price-Model I-Uniform	mean	35.17	13.02	8.12	5.10	2.74	1120
	median	34.55	11.78	7.45	4.34	2.01	1107
	variance	118.22	46.53	24.59	15.02	6.71	174682
	$[E(\underline{p}), E(\bar{p})]$	$[27.51, 41.02]$	$[7.18, 16.81]$	$[3.14, 11.56]$	$[1.15, 8.72]$	$[0.28, 6.68]$	
$Pr(\frac{\kappa_{i+1}}{\kappa_i} \leq 1)$		0.67	0.63	0.61	0.60	0.83	
Price-Model I-Beta	mean	36.75	14.72	9.84	7.13	4.88	1267
	median	36.01	13.30	8.98	6.28	3.92	1251
	variance	139.84	62.22	35.26	25.25	16.64	196707
	$[E(\underline{p}), E(\bar{p})]$	$[25.69, 39.42]$	$[6.55, 16.37]$	$[2.73, 11.21]$	$[1.02, 8.49]$	$[0.22, 6.43]$	
$Pr(\frac{\kappa_{i+1}}{\kappa_i} \leq 1)$		0.68	0.62	0.6	0.59	0.84	
Price-Model II-Uniform	mean	28.32	17.89	10.37	6.00	2.92	1137
	median	27.32	17.65	9.80	5.08	2.12	1088
	variance	103.90	50.03	33.68	21.97	8.75	219406
	$[E(\underline{p}), E(\bar{p})]$	$[19.38, 34.60]$	$[11.60, 22.25]$	$[4.18, 14.64]$	$[1.10, 10.49]$	$[0.10, 7.86]$	
$Pr(\frac{\kappa_{i+1}}{\kappa_i} \leq 1)$		0.59	0.65	0.62	0.59	0.9	
Price-Model II-Beta	mean	30.29	19.42	12.14	8.03	5.19	1290
	median	29.11	19.06	11.56	7.01	3.97	1250
	variance	113.44	50.93	39.81	34.88	23.73	243226
	$[E(\underline{p}), E(\bar{p})]$	$[18.56, 33.09]$	$[11.15, 21.28]$	$[3.9, 13.79]$	$[0.95, 9.61]$	$[0.07, 7.10]$	
$Pr(\frac{\kappa_{i+1}}{\kappa_i} \leq 1)$		0.59	0.66	0.61	0.59	0.91	

Table 7: Counterfactual per-click prices for Symmetric Nash Equilibrium of WGSPA: Coarser Score

(Units of prices are Chinese RMB.)

		Rank of Positions					Total Revenue
Price-Raw Data	Statistics	1	2	3	4	5	
	mean	22.32	16.13	12.68	10.27	8.12	1133
	median	19.6	13.67	11	9.74	7.73	910
	variance	139.62	62.8	32.15	18.95	12.16	616476
$Pr(\frac{\kappa_{i+1}}{\kappa_i} \leq 1)$		0.86	0.84	0.86	0.84	0.91	
Price-Model I-Uniform	mean	34.98	13.02	8.13	5.10	2.73	1119
	median	34.28	11.84	7.45	4.38	1.99	1101
	variance	118.54	46.76	24.09	14.37	6.55	170717
	$[E(\underline{p}), E(\bar{p})]$	$[27.23, 40.88]$	$[7.15, 16.82]$	$[3.08, 11.6]$	$[1.13, 8.77]$	$[0.27, 6.28]$	
$Pr(\frac{\kappa_{i+1}}{\kappa_i} \leq 1)$		0.87	0.86	0.85	0.85	0.93	
Price-Model I-Beta	mean	37.09	14.88	10.03	7.16	4.99	1277
	median	36.07	13.39	9.16	6.34	4.09	1262
	variance	147.48	64.55	36.24	25.55	17.01	197976
	$[E(\underline{p}), E(\bar{p})]$	$[25.74, 39.78]$	$[6.54, 16.60]$	$[2.76, 11.41]$	$[0.98, 8.53]$	$[0.22, 6.59]$	
$Pr(\frac{\kappa_{i+1}}{\kappa_i} \leq 1)$		0.86	0.84	0.86	0.86	0.93	
Price-Model II-Uniform	mean	28.36	17.98	10.45	6.05	2.92	1142
	median	27.39	17.66	9.93	5.09	2.13	1091
	variance	105.26	49.80	32.91	22.00	8.26	220235
	$[E(\underline{p}), E(\bar{p})]$	$[19.34, 34.63]$	$[11.65, 22.39]$	$[4.19, 14.77]$	$[1.09, 10.61]$	$[0.11, 7.88]$	
$Pr(\frac{\kappa_{i+1}}{\kappa_i} \leq 1)$		0.80	0.86	0.85	0.83	0.95	
Price-Model II-Beta	mean	30.22	19.52	12.12	8.13	5.21	1286
	median	29.12	19.15	11.61	7.15	4.07	1243
	variance	110.84	52.13	37.75	33.41	21.10	240496
	$[E(\underline{p}), E(\bar{p})]$	$[18.57, 33.05]$	$[11.23, 21.39]$	$[3.85, 13.79]$	$[0.99, 9.77]$	$[0.08, 7.16]$	
$Pr(\frac{\kappa_{i+1}}{\kappa_i} \leq 1)$		0.80	0.85	0.84	0.83	0.95	

A Complete details of estimation procedure

A.1 Specifying the Priors

We assume the prior distribution of β follows normal distribution $\mathcal{N}(\beta_0, B_0)$ and σ^2 follows inverted Gamma distribution $\mathcal{IG}(\alpha_0/2, \delta_0/2)$. β and σ^2 are independent and their joint distribution will be denoted by $\pi(\theta)$. We further assume that $V_{ij} = \delta(X_i, Z^j; \beta) + \epsilon_{ij}$ with ϵ_{ij} follows i.i.d. $\mathcal{N}(0, \sigma^2)$[32] across the index (i, j) and different auction t. Although there always exists a stable matching for an arbitrary draw of V_{ij}, it is not the case once the GSP restriction $p_1 > p_2 > \cdots > p_N$ being imposed. The intersection of $p_1 > p_2 > \cdots > p_N$ and the set of stable matchings of Shapley and Shubik may be empty. When estimating the model we shall restrict our attention to the set of $\mathbf{V}_t$ that can guarantee the existence of an equilibrium[33]:

$$\{\mathbf{V}_t \in \text{GSP}\} \equiv \{\exists(t_1, \ldots, t_N)|(t_1, \ldots, t_N) \text{ solves (DP) \& } p_1 > p_2 > \ldots p_N; t_i = \alpha_i p_i\}.$$

The probability of this region is $Pr\{\mathbf{V}_t \in \text{GSP}\} = c(\theta; \mathbf{X}_t, \mathbf{Z}_t)$ and can be approximated by simulation.[34] This specification, together with the restriction of the existence of equilibrium, implies that the joint distribution of $\mathbf{V}_t$ is a multivariate truncated normal distribution

$$f(\mathbf{V}_t|\theta; \mathbf{X}_t, \mathbf{Z}_t) = \mathbb{1}(\mathbf{V}_t \in \text{GSP}) \frac{\Pi_{i=1}^N \Pi_{j=1}^N \frac{1}{\sigma} \phi\left(\frac{V_{ijt} - \delta(X_{it}, Z^{jt}; \beta)}{\sigma}\right)}{c(\theta; \mathbf{X}_t, \mathbf{Z}_t)},$$

where $\phi(\cdot)$ is the pdf of standard normal distribution. The specification of $p_0(\theta, \mathbf{V}|\mathbf{X}, \mathbf{Z})$ is now completed after specifying $\pi(\theta)$ and $f(\mathbf{V}|\theta; \mathbf{X}, \mathbf{Z})$ since $p_0(\theta, \mathbf{V}|\mathbf{X}, \mathbf{Z}) = f(\mathbf{V}|\theta; \mathbf{X}, \mathbf{Z})\pi(\theta)$.[35] Finally, for the Beta equilibrium selection model, we assume uniform priors over the shape paraeters (a, b): $\pi(a, b) \propto \frac{1}{c}^2, 0 < c < \infty$.

A.2 Algorithm

We propose a Metropolis-Hastings within Gibbs sampler to draw $(\theta, \lambda, \mathbf{V})$ from the posterior. Gibb's sampling loops over the three conditional distributions $f_1(\mathbf{V}|\theta, \lambda, \mu, \mathbf{p}, \mathbf{X}, \mathbf{Z}, \alpha) = \Pi_{t=1}^T f_1(\mathbf{V}_t|\theta, \lambda, \mu_t, \mathbf{p}_t, \mathbf{X}_t, \mathbf{Z}_t, \alpha_t)$, $f_2(\theta|\mathbf{V}, \lambda, \mu, \mathbf{p}, \mathbf{X}, \mathbf{Z}, \alpha)$, and $f_3(\lambda|\theta, \mathbf{V}, \mu, \mathbf{p}, \mathbf{X}, \mathbf{Z}, \alpha)$. As it is difficult to draw directly from the above conditional densities, one can instead using

[32]While we use normality assumption in the empirical study, the proposed algorithm here can be easily applied to other distributions.

[33]Here the indices $1, \ldots, N$ refer only to the positions, and not to specific bidders.

[34]See A.3 for the implementation detail.

[35]By the independence assumption $f(\mathbf{V}|\theta; \mathbf{X}, \mathbf{Z}) = \Pi_{t=1}^T f(\mathbf{V}_t|\theta; \mathbf{X}_t, \mathbf{Z}_t)$

Metropolis-Hastings sampler. For f_1, we simulate $\mathbf{V}_t$ for each t from normal distribution, truncated to $A(\alpha_t)\mathbf{p}_t \leq \mathbf{b}(\mathbf{V}_t)$. This step amounts to impose equilibrium restriction, but without explicitly adjusting for the effect of multiple equilibria. We then use an independent M-H step to correct for it later, essentially weighted by the equilibrium selection probability. For f_2, it is a truncated normal likelihood function and hence can be simulated using standard M-H procedure too. Below is the implementation detail:

1. Conditional on $(\theta^{(\tau)}, \lambda^{(\tau)}, \mathbf{V}^{(\tau)})$, update $\mathbf{V}^{(\tau+1)}$ via independence M-H chain

1.1 Simulate

$$\begin{aligned}\tilde{\mathbf{V}}_t &\sim q_t(\tilde{\mathbf{V}}_t|\mathbf{V}_1^{(\tau+1)}, \ldots, \mathbf{V}_{t-1}^{(\tau+1)}, \mathbf{V}_t^{(\tau)}, \mathbf{V}_{t+1}^{(\tau)}, \ldots, \mathbf{V}_T^{(\tau)}, \theta^{(\tau)}, \lambda^{(\tau)}) \\ &= q_t(\tilde{\mathbf{V}}_t|\theta^{(\tau)}) \propto \mathbb{1}(A(\alpha_t)\mathbf{p}_t \leq \mathbf{b}(\tilde{\mathbf{V}}_t)) f(\tilde{\mathbf{V}}_t|\theta^{(\tau)}; \mathbf{X}_t, \mathbf{Z}_t).\end{aligned}$$

Following Logan, et al. (2008) we suggest the following steps to simulate $\tilde{V}_{ijt}$[36]: 1. simulate the diagonal elements $\tilde{V}_{iit}$ from normal distribution with mean $\delta(X_{it}, Z^{it}; \beta^{(\tau)})$ and variance $\sigma^2_{(\tau)}$, left-truncated at $\alpha_{it}p_{it}$ (individual rationality). 2. Given the simulated $\tilde{V}_{iit}$, one then proceed to simulate the off-diagonal elements $\tilde{V}_{ijt}, i \neq j$. The NBP condition implies that $\tilde{V}_{ijt}$ follows normal distribution with mean $\delta(X_{it}, Z^{jt}; \beta^{(\tau)})$ and variance $\sigma^2_{(\tau)}$, right-truncated at $\tilde{V}_{iit} - (\alpha_{it}p_{it} - \alpha_{jt}p_{jt})$. 3. Notice that the monotonicity condition of per-click price does not directly affect the simulation of V_{ij}. Because the observed data already satisfies the monotonicity condition, the simulated V_{ij} will automatically lead to an equilibrium polytope of per-click prices that intersects with the set $p_1 > p_2 > \ldots, p_N$.

1.2 Take

$$\mathbf{V}_t^{(\tau+1)} = \begin{cases} \mathbf{V}_t^{(\tau)} & \text{with probability } 1-\rho, \\ \tilde{\mathbf{V}}_t & \text{with probability } \rho, \end{cases}$$

[36]The standard GHK simulator does not apply in this case.

where

$$\rho = 1 \wedge \Big[\frac{f_1(\tilde{\mathbf{V}}_t|\theta^{(\tau)}, \lambda^{(\tau)}, \mu_t, \mathbf{p}_t, \mathbf{X}_t, \mathbf{Z}_t, \alpha_t)}{f_1(\mathbf{V}_t^{(\tau)}|\theta^{(\tau)}, \lambda^{(\tau)}, \mu_t, \mathbf{p}_t, \mathbf{X}_t, \mathbf{Z}_t, \alpha_t)}\Big]\Big[\frac{q_t(\mathbf{V}_t^{(\tau)}|\theta^{(\tau)})}{q_t(\tilde{\mathbf{V}}_t|\theta^{(\tau)})}\Big]$$

$$= 1 \wedge \Big[\frac{\mathcal{L}_1(\mathbf{p}_t|\tilde{\mathbf{V}}_t, \theta^{(\tau)}, \lambda^{(\tau)}, \mu_t, \mathbf{X}_t, \mathbf{Z}_t, \alpha_t) f(\tilde{\mathbf{V}}_t|\theta^{(\tau)}, \mathbf{X}_t, \mathbf{Z}_t)}{\mathcal{L}_1(\mathbf{p}_t|\mathbf{V}_t^{(\tau)}, \theta^{(\tau)}, \lambda^{(\tau)}, \mu_t, \mathbf{X}_t, \mathbf{Z}_t, \alpha_t) f(\mathbf{V}_t^{(\tau)}|\theta^{(\tau)}, \mathbf{X}_t, \mathbf{Z}_t)}\Big]\Big[\frac{q_t(\mathbf{V}_t^{(\tau)}|\theta^{(\tau)})}{q_t(\tilde{\mathbf{V}}_t|\theta^{(\tau)})}\Big]$$

$$= 1 \wedge \Big[\frac{\mathcal{L}_1(\mathbf{p}_t|\mathbf{P}(\tilde{\mathbf{V}}_t, \alpha_t), \lambda^{(\tau)}, \mu_t) f(\tilde{\mathbf{V}}_t|\theta^{(\tau)}, \mathbf{X}_t, \mathbf{Z}_t)}{\mathcal{L}_1(\mathbf{p}_t|\mathbf{P}(\mathbf{V}_t^{(\tau)}, \alpha_t), \lambda^{(\tau)}, \mu_t) f(\mathbf{V}_t^{(\tau)}|\theta^{(\tau)}, \mathbf{X}_t, \mathbf{Z}_t)}\Big]\Big[\frac{\mathbb{1}(A(\alpha_t)\mathbf{p}_t \leq \mathbf{b}(\mathbf{V}_t^{(\tau)})) f(\mathbf{V}_t^{(\tau)}|\theta^{(\tau)}; \mathbf{X}_t, \mathbf{Z}_t)}{\mathbb{1}(A(\alpha_t)\mathbf{p}_t \leq \mathbf{b}(\tilde{\mathbf{V}}_t)) f(\tilde{\mathbf{V}}_t|\theta^{(\tau)}; \mathbf{X}_t, \mathbf{Z}_t)}\Big]$$

$$= 1 \wedge \frac{\mathcal{L}_1(\mathbf{p}_t|\mathbf{P}(\tilde{\mathbf{V}}_t, \alpha_t), \lambda^{(\tau)}, \mu_t)}{\mathcal{L}_1(\mathbf{p}_t|\mathbf{P}(\mathbf{V}_t^{(\tau)}, \alpha_t), \lambda^{(\tau)}, \mu_t)}.$$

In particular, if the equilibrium selection rule is uniform then $\rho = 1 \wedge \frac{vol(\mathbf{P}(\mathbf{V}_t^{(\tau)}, \alpha_t))}{vol(\mathbf{P}(\tilde{\mathbf{V}}_t, \alpha_t))}$. In this case, $\tilde{\mathbf{V}}_t$ will be accepted with probability 1 if the resulting polyhedron has smaller volume relative to $\mathbf{V}_t^{(\tau)}$. To sum up, one would need to independently simulate T valuation matrices for T keyword auctions, and then run T independent M-H steps to decide whether to accept the new draws or not.

2. Conditional on $(\theta^{(\tau)}, \lambda^{(\tau)}, \mathbf{V}^{(\tau+1)})$, update $\theta^{(\tau+1)}$ via random walk M-H chain

2.1 Simulate

$$\tilde{\theta} \sim q(\tilde{\theta}|\theta^{(\tau)}) = \begin{cases} A^{-1} & \text{for } \tilde{\theta} \in \theta^{(\tau)} \pm a \\ 0 & \text{otherwise,} \end{cases}$$

where a is a vector of the same dimension as θ and A is the volume of the box spanned by $\theta_i^{(\tau)} \pm a_i$. As $q(\cdot|\cdot)$ is a random-walk proposal density, it is also symmetric: $q(\tilde{\theta}|\theta^{(\tau)}) = q(\theta^{(\tau)}|\tilde{\theta})$.

2.2 Take

$$\theta^{(\tau+1)} = \begin{cases} \theta^{(\tau)} & \text{with probability } 1-\rho, \\ \tilde{\theta} & \text{with probability } \rho, \end{cases}$$

where

$$\rho = 1 \wedge \frac{f_2(\tilde{\theta}|\mathbf{V}^{(\tau+1)}, \lambda^{(\tau)}, \mu, \mathbf{p}, \mathbf{X}, \mathbf{Z}, \alpha)}{f_2(\theta^{(\tau)}|\mathbf{V}^{(\tau+1)}, \lambda^{(\tau)}, \mu, \mathbf{p}, \mathbf{X}, \mathbf{Z}, \alpha)}$$

$$= 1 \wedge \frac{\Pi_{t=1}^T f_1(\mathbf{V}_t^{(\tau+1)}|\tilde{\theta}, \lambda^{(\tau)}, \mu_t, \mathbf{p}_t, \mathbf{X}_t, \mathbf{Z}_t, \alpha_t)\pi(\tilde{\theta})}{\Pi_{t=1}^T f_1(\mathbf{V}_t^{(\tau+1)}|\theta^{(\tau)}, \lambda^{(\tau)}, \mu_t, \mathbf{p}_t, \mathbf{X}_t, \mathbf{Z}_t, \alpha_t)\pi(\theta^{(\tau)})}$$

$$= 1 \wedge \frac{\Pi_{t=1}^T \mathcal{L}_1(\mathbf{p}_t|\mathbf{V}_t^{(\tau+1)}, \tilde{\theta}, \lambda^{(\tau)}, \mu_t, \mathbf{X}_t, \mathbf{Z}_t, \alpha_t) f(\mathbf{V}_t^{(\tau+1)}|\tilde{\theta}, \mathbf{X}_t, \mathbf{Z}_t)}{\Pi_{t=1}^T \mathcal{L}_1(\mathbf{p}_t|\mathbf{V}_t^{(\tau+1)}, \theta^{(\tau)}, \lambda^{(\tau)}, \mu_t, \mathbf{X}_t, \mathbf{Z}_t, \alpha_t) f(\mathbf{V}_t^{(\tau+1)}|\theta^{(\tau)}, \mathbf{X}_t, \mathbf{Z}_t)} \frac{\pi(\tilde{\theta})}{\pi(\theta^{(\tau)})}$$

Because $\mathcal{L}_1$ only depends on $(\mathbf{V}_t, \alpha_t, \lambda)$ and by construction $\mathbf{V}_t^{(\tau+1)} \in \text{GSP}$, the above equation can be further simplified:

$$= 1 \wedge \frac{\Pi_{t=1}^T \left[\Pi_{i=1}^N \Pi_{j=1}^N \frac{1}{\tilde{\sigma}} \phi\left(\frac{V_{ijt}^{(\tau+1)} - \delta(X_{it}, Z^{jt}; \tilde{\beta})}{\tilde{\sigma}}\right)\right]}{\Pi_{t=1}^T \left[\Pi_{i=1}^N \Pi_{j=1}^N \frac{1}{\sigma^{(\tau)}} \phi\left(\frac{V_{ijt}^{(\tau+1)} - \delta(X_{it}, Z^{jt}; \beta^{(\tau)})}{\sigma^{(\tau)}}\right)\right]} \frac{\Pi_{t=1}^T c(\theta^{(\tau)}; \mathbf{X}_t, \mathbf{Z}_t)}{\Pi_{t=1}^T c(\tilde{\theta}; \mathbf{X}_t, \mathbf{Z}_t)} \frac{\pi(\tilde{\theta})}{\pi(\theta^{(\tau)})}$$

This step is nothing but treating $\mathbf{V}^{(\tau+1)}$ as the data, and then evaluate the likelihood ratio of the truncated normal density.

3. Conditional on $(\theta^{(\tau+1)}, \lambda^{(\tau)}, \mathbf{V}^{(\tau+1)})$, update $\lambda^{(\tau+1)}$ via random walk M-H chain

3.1 Simulate

$$\tilde{\lambda} \sim q(\tilde{\lambda}|\lambda^{(\tau)}) = \begin{cases} A^{-1} & \text{for } \tilde{\lambda} \in \lambda^{(\tau)} \pm a \\ 0 & \text{otherwise,} \end{cases}$$

where a is a vector of the same dimension as λ and A is the volume of the box spanned by $\lambda_i^{(\tau)} \pm a_i$.

3.2 Take

$$\lambda^{(\tau+1)} = \begin{cases} \lambda^{(\tau)} & \text{with probability } 1-\rho, \\ \tilde{\lambda} & \text{with probability } \rho, \end{cases}$$

where

$$\rho = 1 \wedge \frac{f_3(\tilde{\lambda}|\theta^{\tau+1}, \mathbf{V}^{\tau+1}, \mu, \mathbf{p}, \mathbf{X}, \mathbf{Z}, \alpha)}{f_3(\lambda^{\tau}|\theta^{\tau+1}, \mathbf{V}^{\tau+1}, \mu, \mathbf{p}, \mathbf{X}, \mathbf{Z}, \alpha)}$$

$$= 1 \wedge \frac{\Pi_{t=1}^T \mathcal{L}_1(\mathbf{p}_t|\mathbf{P}(\mathbf{V}_t^{\tau+1}, \alpha_t), \tilde{\lambda}, \mu_t)}{\Pi_{t=1}^T \mathcal{L}_1(\mathbf{p}_t|\mathbf{P}(\mathbf{V}_t^{\tau+1}, \alpha_t), \lambda^{\tau}, \mu_t)}.$$

Under uniform equilibrium selection, there is no need to perform step 3.

A.3 Other Implementation Details

A.3.1 Characterizing the set of equilibrium prices

As discussed in section 3.1 above, the set of equilibrium per-click prices is a convex polyhedron defined by various linear inequalities. First, we consider the no-blocking pair conditions in Eq. (1). one can substitute $u_i + t_i = V_{ii} \forall i$ into the rest of NBP conditions conditions: The resulting system of inequalities are

$$V_{ii} - V_{ij} \geq t_i - t_j \geq V_{ji} - V_{jj}. \tag{8}$$

Equivalently, equation (8) can be re-written using the matrix notation

$$\{\mathbf{p} | A_1 \mathbf{p} \leq b_1\}.$$

Take a 3-by-3 case as an example,

$$A_1 = \begin{bmatrix} \alpha_1 & -\alpha_2 & 0 \\ -\alpha_1 & \alpha_2 & 0 \\ \alpha_1 & 0 & -\alpha_3 \\ -\alpha_1 & 0 & \alpha_3 \\ 0 & \alpha_2 & -\alpha_3 \\ 0 & -\alpha_2 & \alpha_3 \end{bmatrix}, \quad \text{and} \quad b_1 = \begin{bmatrix} V_{11} - V_{12} \\ V_{22} - V_{21} \\ V_{11} - V_{13} \\ V_{33} - V_{31} \\ V_{22} - V_{23} \\ V_{33} - V_{32} \end{bmatrix}. \tag{9}$$

The second set of inequalities are the individual rationality conditions; i.e., bidders' payoff should be positive

$$\{\mathbf{p} | A_2 \mathbf{p} \leq b_2\},$$

where A_2 is the diagonal matrix with $(\alpha_1, \alpha_2, \ldots, \alpha_N)$ on the main diagonal and $b_2 = (V_{11}, V_{22}, \ldots, V_{NN})'$. Finally, the last set of inequalities states that $p_1 > p_2 > \cdots > p_N \leq 0$, corresponding to the non-negative price systems that are consistent with the generalized second price auction.

$$\{\mathbf{p} | A_4 \mathbf{p} \leq b_4\},$$

For a 3-by-3 case,

$$A_3 = \begin{bmatrix} -1 & 1 & 0 \\ 0 & -1 & 1 \end{bmatrix}, \quad \text{and} \quad b_3 = \begin{bmatrix} 0 \\ 0 \end{bmatrix}. \tag{10}$$

along with the nonegativity constraints

$$\{\mathbf{p}|A_4\mathbf{p} \leq b_4\},$$

where $A_4 = -\mathbf{I}_{N\times N}$ and $b_4 = -\underline{0}\mathbf{1}_{N\times 1}$.

vThe set of equilibrium prices is given by $\mathbf{P}(\mathbf{V},\alpha) = \{\mathbf{p}|A(\alpha)\mathbf{p} \leq \mathbf{b}(\mathbf{V})\}$, where $A = (A_1^{'}|A_2^{'}|A_3^{'}|A_4^{'})'$ and $\mathbf{b} = (b_1^{'}|b_2^{'}|b_3^{'}|b_4^{'})'$.

A.3.2 Parameter Setup

For the first specification of bidder valuations (Model I), the prior for β is assume to be joint normal distribution. The mean vector β_0 equals to the zero vector, and the covariance matrix B_0 equals to $100000I$, where I is an 12-by-12 identify matrix. The prior for σ^2 is assumed to be inverted Gamma distribution with shape parameter $\alpha_0 = 2$ and scale parameter $\delta_0 = 5$. We choose the uniform distribution on the half-line as the prior for the shape parameters (γ_1, γ_2) of the beta distribution when estimating the equilibrium selection density. This set of prior values only impose minimum prior information on the parameters. For example, when the shape parameter is 2, the variance of inverted Gamma distribution does not exist. Moreover, given the "large sample" feature in the Bayesian updating step for (β, σ),[37]the prior specification only plays a negligible role in determining the posterior. We make 30,000 draws and the first 10,000 draws are treated as the burn-in. We then keep every 50th draw of the remaining 20,000 draws to estimate the posterior mean and standard deviations. The radius for the random walk proposal for β, σ^2, and (γ_1, γ_2) are respectively 100, 1000 and 0.2.

For Model II, the prior for β is assume to be joint normal distribution. The mean vector β_0 equals to the zero vector, and the covariance matrix B_0 equals to $100000I$, where I is an 44-by-44 identify matrix. We make 50,000 draws and the first 30,000 draws are treated as the burn-in. We then keep every 50th draw of the remaining 20,000 draws to estimate the posterior mean and standard deviations. The radius for the random walk proposal for β, σ^2, and (γ_1, γ_2) are respectively 5, 100 and 0.2.

A.3.3 Calculating the Volume of Equilibrium Polytope of Prices

In order to evaluate the likelihood, one has to compute the volume of $\mathbf{P}(\mathbf{V},\alpha)$ if uniform selection is imposed. We do it via simulation: first, we draw 1000 independent (multivariate) uniform random numbers from the smallest bounding box of $\mathbf{P}(\mathbf{V},\alpha)$. Second, the volume

[37]In a 5-player-5-position game, 25 V_{ij} will be drawn in MCMC. If there are 100 keyword auctions, (β, σ^2) will be estimated by 2500 simulated V_{ij}.

of $\mathbf{P}(\mathbf{V}, \alpha)$ is approximately the volume of the bounding box times the proportion of the previous draws that belong to $\mathbf{P}(\mathbf{V}, \alpha)$. If beta selection is imposed, one instead draw beta random numbers from the bounding box. We also try 5000 draws and the accuracy seems to be similar.

A.3.4 Calculating the Truncation Probability $c(\theta; \mathbf{X}_t, \mathbf{Z}_t)$

We use a simulation-extrapolation strategy to compute the truncation probability, as we find that $c(\theta; \mathbf{X}_t, \mathbf{Z}_t)$ is a relatively smooth function of θ. Given θ, and for each keyword auction t we can simulate N valuation matrices $\mathbf{V}_i$ to approximate $c(\theta; \mathbf{X}_t, \mathbf{Z}_t)$ by $\frac{1}{N}\sum_{i=1}^{N} \mathbb{1}(\mathbf{V}_i \in \text{GSP})$. This step would require solving (DP). We first make 10,000 MCMC draws by ignoring this truncation probability. The first 1000 draws are discarded, and we keep every 10th draw of the remaining chain. We simulate $c(\theta_j; \mathbf{X}_t, \mathbf{Z}_t)$ under these simulated θ_j, $j = 1, 2, \ldots, 900$. Finally, we regress $c(\theta_j; \mathbf{X}_t, \mathbf{Z}_t)$ on θ_j using beta regression[38], with the logistic function being the link function. The estimated regression coefficients are then used to calculate $c(\theta; \mathbf{X}_t, \mathbf{Z}_t)$ in MCMC. One important fact is that if $\mathbf{V}_i \notin \text{GSP}$, then $a + b\mathbf{V}_i \notin \text{GSP}$, where (a, b) are some scalar constants. It is the relative size of each component within $\mathbf{V}$ that leads to nonexistence of equilibrium, not because of its scale. As a result, the scale parameter and the non-interacted location parameters (see Appendix B) does not affect $c(\theta; \mathbf{X}_t, \mathbf{Z}_t)$. When running the regression, we discard the non-interacted location parameters, and normalized the interacted location parameters by the scale parameter σ.

B Identification

The specification of the latent valuation V_{ij} implicitly assumes that the distribution of V_{ij} belongs to the location-scale family. However, the meaning of *location* and *scale* should be carefully interpreted in matching models, as location parameter for V_{ij} may actually possess scale effect on the dependent variable $\mathbf{p}$.

There are three types of parameters in the specification of V_{ij}: interacted location parameter $(\beta_4, \ldots, \beta_{11})$, non-interacted location parameter $(\beta_0, \ldots, \beta_3)$ and scale parameter σ. First, the interacted location parameters characterize the preference over position ranks. Or equivalently, they characterize the complementarity between quality and ranks, and hence will determine the (cross-sectional) distribution of allocation. They are also related to the price distribution through the channel of the no-blocking-pair conditions, because they essentially determine the shape of the equilibrium polytope of price. The point identification

[38] Beta regression is a flexible regression model to handel the cases when the dependent variable is proportion.

results for the interacted location parameters (up to scale normalization) from the allocation data μ have been established in Choo-Siow model and Fox (2010) under different model assumptions.

The price data essentially provides identification power for parameters with scale effect on the valuation matrix. There are two types of such parameters: First, the non-interacted location parameters cannot be identified through the no-blocking-pair conditions, as they are differenced out. Instead, the individual rationality condition can be used to learn information about them. Such parameters do not affect the preference over ranks, and consequently they have no effect on the allocation. However, as they would shift the scale of the valuation matrix, they will also shift the size of the equilibrium polytope of prices. Similarly, the scale parameter σ have no effect on the allocation, but it will also affect the size of the equilibrium polytope of price. Both non-interacted location and scale parameters have scale effect on the price distribution. The difference is, the non-interacted location parameters only shift the size of the equilibrium polytope in *certain* directions (through individual rationality), while the scale parameter shift the size of the equilibrium polytope in *all* directions. By looking at the price distribution, and the shape/size of the support of prices one can then learn information about these parameters. The intuition is simple: if on average the price is \$5, the scale of $\mathbf{V}$ cannot be around \$1, as it would violate the individual rationality. On the other hand, it cannot be around \$100, as \$5 would be too cheap under competitive bidding.

Remark: a partial identification approach. It is possible to derive a bound for the structural parameters in terms of bounds of CDF (similarly to Haile and Tamer, 2003).

Theorem 1. *Suppose the econometrician observes T independent auctions indexed by $t = 1, 2, \ldots, T$. Let $F(\mathbf{p}|\mathbf{X}, \mathbf{Z}, \alpha)$ represents the conditional distribution of observed per-click prices. We denote by $\mathbf{B}(\mathbf{X}, \mathbf{Z}, \alpha, \epsilon_{ij}; \theta)$, the smallest (random) bounding box that covers $\mathbf{P}(\mathbf{V}, \alpha) = \mathbf{P}(\mathbf{X}, \mathbf{Z}, \alpha, \epsilon_{ij}; \theta)$. $\mathbf{B}(\mathbf{X}, \mathbf{Z}, \alpha, \epsilon_{ij}; \theta) = \Pi_{i=1}^{N}[p_{il}, p_{iu}]$, where (p_{il}, p_{iu}) are respectively the smallest and largest elements of the i-th coordinate of $\mathbf{P}$. Define two random vectors $\bar{\mathbf{p}} = \max \mathbf{B}(\mathbf{X}, \mathbf{Z}, \alpha, \epsilon_{ij}; \theta) = (p_{1u}, p_{2u}, \ldots, p_{Nu})$ and $\underline{\mathbf{p}} = \min \mathbf{B}(\mathbf{X}, \mathbf{Z}, \alpha, \epsilon_{ij}; \theta) = (p_{1l}, p_{2l}, \ldots, p_{Nl})$. The identified set of θ is given by*

$$\Theta_0 = \left\{\theta | F(\underline{\mathbf{p}}|\mathbf{X}, \mathbf{Z}, \alpha; \theta) \geq F(\mathbf{p}|\mathbf{X}, \mathbf{Z}, \alpha) \geq F(\bar{\mathbf{p}}|\mathbf{X}, \mathbf{Z}, \alpha; \theta)\right\},$$

Proof. Conditional on $(\mathbf{X}, \mathbf{Z})$, the (joint) distribution function of $\mathbf{p}$, $F(\mathbf{p}|\mathbf{X}, \mathbf{Z})$, is identified by the sampling process. Given $(\mathbf{X}, \mathbf{Z}, \alpha, \epsilon_{ij}; \theta)$, $\mathbf{p}$ is a measurable selection from the set of stable prices $\mathbf{P}$. By construction, $\underline{\mathbf{p}} \leq \mathbf{p} \leq \bar{\mathbf{p}}$, for almost-all ϵ_{ij}, where the inequality are defined componentwise. This implies that the distributions of $\underline{p}$, p, and $\bar{p}$ are ordered in the sense of first-order stochastic dominance, which immediately implies the bound on CDF. ■

Although the bound approach is appealing as it does not require an equilibrium selection assumption, it is extremely computationally demanding to compute such bound since for each draw of latent valuation matrix, one has to solve (DP).[39] By contrast, as we discussed earlier, the Bayesian approach does not require solving for game for each draw of the latent valuations, and hence is computatinoally more appealing. See also Uetake and Watanabe (2012) for a bounds approach for estimating a two-sided matching model applied to bank mergers.

C Complete details of counterfactual simulation

To solve for the set of SNE in the WGSPA, we consider each possible allocation in turn; for each candidate allocation, we use the inequalities characterizing WGSPA to determine the set of equilibrium bids (which may be empty if this candidate allocation is not an equilibrium), and then repeat this routine for all possible allocations. Consequently, the problem of solving SNE of WGSPA is combinatorial. Parallel to the standard GSPA, the above inequalities can be written as matrix form. The set of equilibrium bid under allocation μ will be referred as $B_\mu \equiv \{\mathbf{b} | D_\mu(\alpha, \kappa)\mathbf{b} \leq c_\mu(\mathbf{V})\}$.

We then simulate the structural parameters θ from the posterior distribution (by directly using the MCMC output), and for each θ we draw the utility shocks ϵ_{ijt} to obtain the valuation matrix $\mathbf{V}_t$. Given $\mathbf{V}_t$ one can solve the game to obtain equilibrium per-click prices. As the number of SNE may be huge, we do not attempt to solve all SNE. Instead, we employ a simple routine to perform the counterfactual analysis. As long as an equilibrium allocation is found, we then stop searching for another equilibrium allocation[40] and simulate b_i from the equilibrium polyhedron of bids B_μ[41] according to uniform distribution or the estimated beta distribution. Finally, we apply the scored pricing rule in Eq. (13) to obtain the per-click price $\tilde{\mathbf{p}}_t$.

[39] Using the approach of Chernozhukov et al. (2007), the total number of times which the game must be solvedgame solving is at least equal to the numberθ times the number of keyword auctions (= 487) times the number of $\mathbf{V}$ draws. Moreover, depending on $c(\theta; \mathbf{X}_t, \mathbf{Z}_t)$, one has to discard many $\mathbf{V}$ draws that do not satisfy the GSPA restriction.

[40] To avoid the concentration on a particular type of allocation, we perform random search.

[41] There is no need to simulate the maximum bid b_1 for the following two reasons: First, the per-click prices are irrelevant of b_1. Second, if $(b_2, \ldots, b_{N+1})$ satisfy all the inequalities, one can always choose b_1 large enough to meet the allocation rule.