cemmap

centre for microdata methods and practice

Nonparametric estimation of sponsored search auctions and impact of Ad quality on search revenue

Dongwoo Kim
Pallavi Pal

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP16/24

Economic
and Social
Research Council

NONPARAMETRIC ESTIMATION OF SPONSORED SEARCH AUCTIONS AND IMPACT OF AD QUALITY ON SEARCH REVENUE

DONGWOO KIM
Department of Economics, Simon Fraser University

PALLAVI PAL
School of Business, Stevens Institute of Technology

This paper presents an empirical model of sponsored search auctions where advertisers are ranked by bid and ad quality. Our model is developed under the *'incomplete information'* setting with a general quality scoring rule. We establish nonparametric identification of the advertiser's valuation and its distribution given observed bids and introduce novel nonparametric estimators. Using Yahoo! search auction data, we estimate value distributions and study the bidding behavior across product categories. We also conduct counterfactual analysis to evaluate the impact of different quality scoring rules on the auctioneer's revenue. Product-specific scoring rules can enhance auctioneer revenue by at most 24.3% at the expense of advertiser profit (-28.3%) and consumer welfare (-30.2%). The revenue-maximizing scoring rule depends on market competitiveness.

Keywords: Online advertising, digital marketing, sponsored search ads, generalized second price auction, incomplete information, nonparametric estimation, score squashing, user targeting

---

Dongwoo Kim: dongwook@sfu.ca

Pallavi Pal: ppal2@stevens.edu

This paper subsumes "Sponsored Search Auction and the Revenue-Maximizing Number of Ads per Page". We are grateful to Thomas Jeitschko, Kyoo il Kim, Hanzhe Zhang, Roger J. Calantone, Michael Conlin, Jay Pil Choi, Chenhui Guo, Ashish Agarwal, David Croson, Vivek Bhattacharya, Arthur Lewbel, and seminar and conference participants at MSU, MEG 2019, American Marketing Association conference 2019, Boston College, Grand Valley State University, Cemmap/SNU Advances in Econometrics Workshop 2022, IAAE conference 2023, AMES 2023, EEA-ESEM 2023, CESG 2023 for useful discussions and suggestions. We would also like to thank Yahoo! Labs for providing the dataset. Kim gratefully acknowledges financial support from the Social Sciences and Humanities Research Council of Canada under the Insight Development Grant (430-2022-00841).

1. INTRODUCTION

Online advertisements, the primary source of revenue for digital platforms, have become a ubiquitous aspect of daily life. The projected global spending on digital ads in 2024 is $\$740.3$ billion and search advertising is the largest segment ($\$306.7$ billion).[1] Search ads are paid links that appear above organic search results on platforms like Google and Yahoo! and are sold through a weighted generalized second price (GSP^w henceforth) auction. Ad positions are assigned in descending order of the weighted bids. The weight, called "quality score", captures the advertiser's clickability.

This paper investigates the influence of quality scores in sponsored search auctions. We develop a game theoretic model accommodating multiple extensions of previous results in the literature. Our main contribution is to propose a novel nonparametric method for estimating participants' maximum willingness to pay (valuations). Understanding valuations is pivotal for advertising platforms. Firstly, it provides insights into potential profits left on the table and market competitiveness by uncovering bid shading practices. Secondly, it enables platforms to fine-tune market mechanisms within the GSP^w framework by exploring counterfactual experiments. Lastly, optimizing policy parameters, such as the reserve price, hinges on the valuation distribution.

We derive the unique symmetric Bayes-Nash equilibrium (BNE) in the GSP^w auction assuming *incomplete information*. The weights in the GSP^w auction introduce a multidimensional type, which we deal with by showing that the bids only depend on a pseudo-one-dimensional type.[2] Subsequently, we demonstrate that the valuation can be nonparametrically identified based on observed bids and auction parameters. Then we propose novel nonparametric estimators of the value and its distribution and density functions. Our estimators do not involve density estimation thus free from the boundary problem and tuning parameters. Therefore, our estimators enjoy favorable finite sample properties, ensuring precise estimation even with a relatively small sample size.

[1] Statista Market Insights: https://www.statista.com/outlook/dmo/digital-advertising/worldwide#ad-spending.

[2] A similar approach was explored in Che (1993) and Asker and Cantillon (2008). They consider the case where price enters linearly in the scoring rule, which is not true in GSP^w auctions.

The proposed method is applied to large data that cover all search ads displayed on Yahoo! over four months in 2008 across five categories: *cruise travel, car insurance, laptops, cable TV*, and *collectible coins*. We estimate valuations and their empirical distribution in each category. Our findings suggest that advertisers with higher weighted values can shade their bids more. In all categories, bid shading was close to 0 at the median but increased rapidly across higher percentiles. The level of competition and position-specific click rates are closely linked to bidding behaviors.

In counterfactual analysis, we explore alternative quality scoring rules. Firstly, we examine score squashing, a method that modifies the relative importance of quality weights by raising the advertiser's clickability to a power of $\theta \in [0, 1]$.[3] Our results suggest that score squashing improves auction revenue by 8.5%–24.3% at the expense of advertiser profit and consumer welfare. The revenue-maximizing squashing level differs across markets, indicating that the scoring rule should be locally decided. Secondly, we investigate the personalization of ad displays through quality scores. Using more precisely computed scores based on user attributes, known as "user targeting" in the literature, increases click rates but reduces competition among advertisers. This trade-off motivates platforms to manipulate the accuracy of quality scores to maximize revenue. Our results demonstrate that revenue can be increased by 0.4%–3.0% by optimally adjusting the precision of quality scores.

This paper is structured as follows. The remainder of this section discusses the related literature. Section 2 outlines the sponsored search market. Section 3 presents the model, equilibrium analysis, and identification and estimation of valuations. Section 4 describes the data. Section 5 explains the estimation procedure with our dataset. Section 6 shows the empirical results. Section 7 conducts counterfactual experiments. Section 8 concludes. Possible extensions to limited bid data and the reserve price, theoretical proofs, simulation experiments on generated auction data, additional empirical results, robustness analysis, and more details on data are provided in the appendix.

[3] Score squashing is widely recognized and utilized in the industry (Lahaie and Pennock, 2007, Charles et al., 2016, Thompson and Leyton-Brown, 2013).

1.1. *Related Literature*

In the theory literature, Edelman et al. (2007) and Varian (2007) were pioneers in deriving ex-post Nash equilibria in unweighted generalized second price (GSP henceforth) auctions under *complete information*.[4] Börgers et al. (2013) extend this analysis by estimating position-specific valuations. Athey and Ellison (2011) explore the impact of consumer search on auction design. The subsequent research further investigates uncertainty in GSP auctions. Gomes and Sweeney (2014) (GS14 henceforth) derive the Bayes-Nash equilibrium under *incomplete information*. Caragiannis et al. (2015) allow for uncertainty in the opponents' values and quality scores and quantify inefficiency in GSP^{w} auctions. Sun et al. (2014) derive the optimal reserve price in GSP$^{\omega}$ auctions without formally deriving the equilibrium. To our knowledge, we are the first to accommodate both a general scoring rule and a non-zero correlation between value and quality in GSP$^{\omega}$ auctions.

Estimating valuations in the GSP^{w} auctions faces several challenges, including multi-dimensional types, complex equilibrium bid equations, unobserved non-winning bids, and unknown finite sample properties of previously proposed estimators. As a result, empirical works in this area have been limited.[5] Börgers et al. (2013) and Hsieh et al. (2015) estimate the value distribution in the GSP auction using Varian (2007)'s model with Chinese data. Yao and Mela (2011) look at dynamic aspects of the generalized first-price auction.[6] Athey and Nekipelov (2010) (AN10 hereafter) derive a unique ex-post Nash equilibrium under market uncertainty assuming complete information. Mohri and Medina (2015) propose an algorithm to estimate the optimal reserve price. They rely on the equilibrium derived in GS14, which cannot be directly applied in GSP^{w} auctions. We present formal proof of

[4]This assumption seems implausible in this market, featuring many advertisers each of whom receives limited information in each round. For further discussions on incomplete information, see Yan (2019).

[5]Ostrovsky and Schwarz (2011) and Bae and Kagel (2019) studied bidding behaviors in GSP auctions using experiments. Both found that the equilibrium in the static complete information case deviates from the Vickrey–Clark–Groves equilibrium, highlighting the importance of studying the incomplete information setting.

[6]Yao and Mela (2011) conduct a policy simulation to study the impact of a switch from first- to second-price auction. Our paper differs from theirs as our model is empirically applicable to data. Unlike their model, we do not look at the dynamic nature of the market.

the equilibrium in the GSP^w auction. Our approach relies on relatively weaker assumptions than previous studies and can be extended to cases of limited data and reserve prices. Furthermore, our estimator can be more easily computed without a tuning parameter.[7]

Numerous studies (Ghose and Yang, 2009, Jerath et al., 2011, Amaldoss et al., 2015, Blake et al., 2015, Decarolis and Rovigatti, 2021, Svitak et al., 2021) have investigated the impact of auction designs on market outcomes. Our paper extends this literature by empirically examining the effects of score squashing on auction outcomes. Lahaie and Pennock (2007) introduced score squashing, finding that its impact on revenue depends on how it alters the allocation. Subsequent research has explored various modifications of the squashing factor.[8] Furthermore, we use different levels of user targeting to study the platform's motivation to deviate from an efficient match between advertisers and consumers. Prior research has theoretically investigated how the ad match quality leads to a trade-off between click rates and price competition.[9] We empirically analyze how the platform fine-tunes the alignment between users and ads to maximize revenue. Although Hauser et al. (2009) and Joshi et al. (2011) indicate advancements in the quality of online ad matches, recent survey data suggest user dissatisfaction with displayed ads.[10] We contribute to uncovering the platform's motivation to reduce match quality to bolster ad revenue. Existing literature presents mixed empirical findings, which underscore the complexity of revenue dynamics and highlight the need for further research to reconcile these discrepancies.[11]

[7] AN10's valuation estimator relies on a numerical derivative formula with a tuning parameter τ_n. Therefore, the convergence rate of their estimator follows the nonparametric rate $\sqrt{n\tau_n}$, whereas our estimator follows the $\sqrt{n}$ rate. Mohri and Medina (2015)'s estimator also involves nonparametric density estimation.

[8] Relying on Varian (2007)'s full information model, Charles et al. (2016) consider squashing only for the first position, Thompson and Leyton-Brown (2013) look at reserve price methods, and Lahaie and McAfee (2011) investigate noisy quality score.

[9] See Levin and Milgrom (2010), Shin and Shin (2023), Bergemann and Bonatti (2011), and Fu et al. (2012).

[10] https://www.surveymonkey.com/curiosity/74-of-people-are-tired-of-social-media-ads-but-theyre-effective/

[11] Yao and Mela (2011) show that targeting can enhance revenue. However, Rafieian and Yoganarasimhan (2021) find that revenue is maximized without targeting. AN10 demonstrate that less precise quality scores can improve revenue.

2. MARKET OVERVIEW

Search engines serve as key digital platforms enabling interactions between consumers and advertisers. Among various advertising avenues they offer, search ads emerge as an effective tool for advertisers to reach interested consumers with specific needs. For instance, a consumer seeking a dentist in New York might use Yahoo! to search for 'New York dentists' and click one of the links displayed to make an appointment. As illustrated in Figure 1, paid ads (highlighted in red box) are strategically positioned atop organic search results to maximize ad visibility by leveraging established user behaviors, such as the F-shaped scanning pattern.[12]

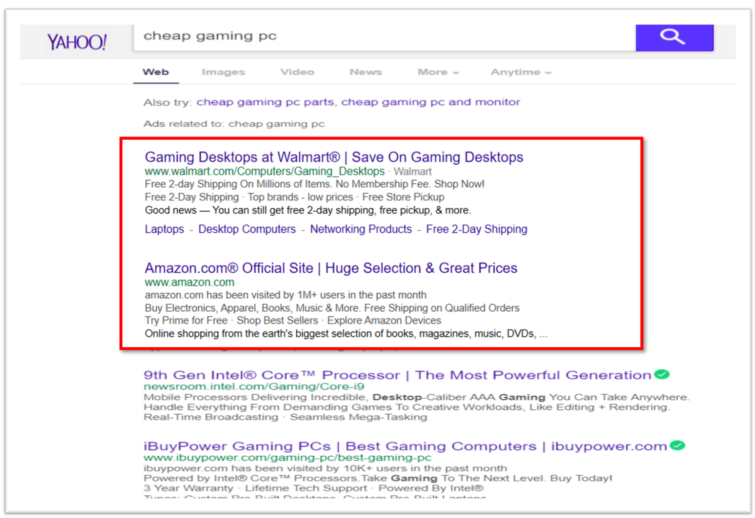

FIGURE 1.—An example of a Yahoo! search result page

A GSP^w auction is a prevalent mechanism for allocating and pricing ad slots on search platforms. Each slot possesses a varying click-through rate (CTR), reflecting the likelihood of a consumer clicking an ad displayed in that position. The GSP^w auction enables advertisers to bid for specific keywords that trigger their ads. Then these ads are assigned to the

[12]Lines at the top and words at the beginning of lines tend to receive more gazes than subsequent lines and words, according to Neilson Norman group research: https://www.nngroup.com/articles/f-shaped-pattern-reading-web-content/.

TABLE I

EXAMPLE OF A GSP^w AUCTION

Advertiser	Bid	Quality Score	Weighted Bid	Position	Price per click
A	2	0.8	1.6	2nd	1.50
B	3	0.6	1.8	1st	2.67
C	1	0.9	0.9	X	X
D	4	0.3	1.2	3rd	3.00

slots based on weighted bids, calculated by multiplying the original bid and quality score. This quality score affects both the ranking and the payment of the ad. Ad positions are allocated in descending order of weighted bids. Each winning advertiser pays a price per click that equals the next highest weighted bid divided by his quality score. The GSP^w auction boasts several advantages: simplicity, speed, and efficiency. It can operate in real time, facilitating ad allocation whenever a user submits a search query. Suppose there are three ad slots and four advertisers who bid for the keyword 'New York dentists'. Table I details their bids, quality scores, weighted bids, assigned positions, and prices per click. The winner of the first position need not have the highest bid or quality score. The weighted bids ultimately dictate the winners and prices.

The search process involves three stages: bidding, auctioning, and feedback as illustrated in Figure 2. During the bidding stage, each advertiser decides his bid per click on his ad and selects the keywords that trigger his ad display. For example, Amazon might bid

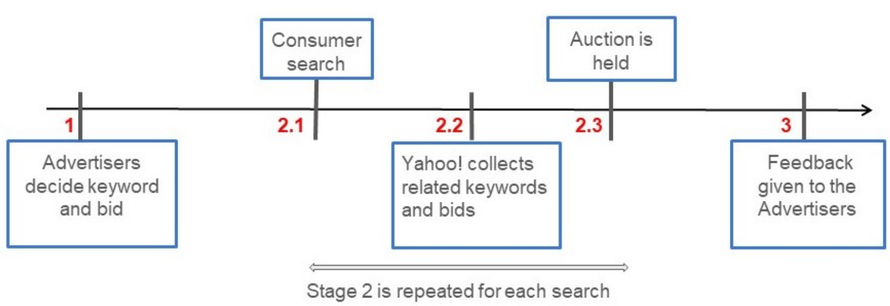

FIGURE 2.—Search process timeline

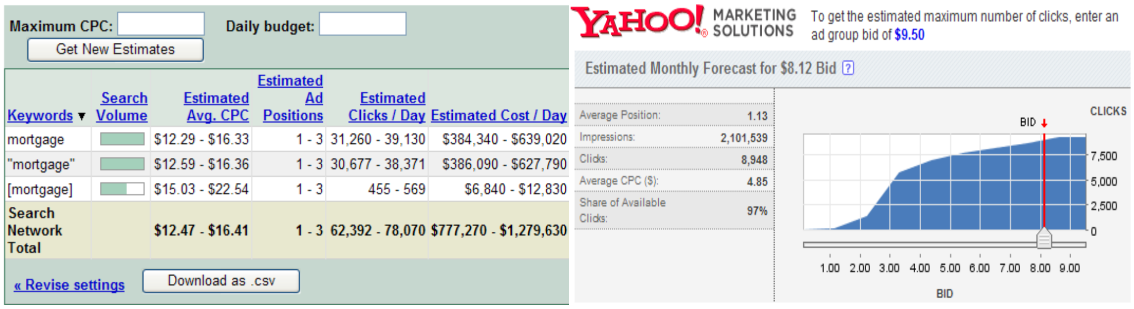

(a) Feedback on keywords (b) Estimated performance of each ad

FIGURE 3.—An example of Yahoo! advertiser account information

$2 for the keyword 'gaming laptop'. In the auctioning stage, when a consumer submits a search query (e.g., "budget gaming laptop"), identifies all relevant keyword-matching ads and ranks them based on weighted bids. Subsequently, the platform assigns each ad to a position on the search results page and determines the corresponding price. Finally, the feedback stage involves the platform providing advertisers with performance data regarding their ads. Figure 3 showcases examples of data accessible to Yahoo advertisers. This data includes display frequency, click frequency, and cost per click (CPC) for each ad position and keyword, mirroring the information utilized in this paper.

3. THEORETICAL MODEL

We present a theoretical model of sponsored search auctions under incomplete information. We first introduce the market environment and outline the model setup and necessary assumptions. Then we derive the unique equilibrium in the auction and establish identification of valuations.

3.1. *Market Environment*

Consumer side.— Each consumer, denoted by i, has a unit demand for a product or service and initiates their search by submitting a query through an online search engine. Once the result page appears with links related to the search query, the consumer decides whether to click on any of the relevant links. The consumer considers the anticipated benefit from clicking on the ad, which is determined by the ad's visible attributes and its position on the results page.[13] Each click incurs a substantial search cost in terms of time spent by the consumer.

Let $U_{i,j}$ denote the expected utility derived by consumer i from clicking on ad j. If the consumer chooses not to click on the ad, she allocates her time to an outside good, resulting in a normalized utility of $\bar{U}_{i,j} \equiv 0$. In the equilibrium, consumers essentially click on all ad links as long as the benefit of a click outweighs the search cost. Let $y^*_{i,j}$ denote the binary indicator capturing consumer i's click decision for ad j as follows:

$$y^*_{i,j} = \begin{cases} 1 & \text{consumer } i \text{ clicks on ad } j \text{ if } U_{i,j} > 0 \\ 0 & \text{if } U_{i,j} \leq 0 \end{cases}$$

The above equation is used to estimate the predicted click probability for ad j in ad-position k. We assume that the ad-specific and position-specific effects are multiplicatively separable. This assumption was adopted by other papers in the literature for identification of the quality of advertiser (for instance, see Varian (2007) and AN10). Thus, for all ad j and position k, the click probability can be written as:

$$\text{Click Probability} = s_j \times c_k \tag{1}$$

where s_j is the effect of advertisement j (used as the quality score in the standard GSP^w auction) and c_k is the effect of ad-position k on the click probability, ceteris paribus. The set of click-through rates for all ad positions on the result page is denoted as $\mathcal{C} = \{c_1, ...c_K\}$. The click probability in (1) is employed in the advertiser's profit maximization problem.

Advertiser side.—Each advertiser, denoted by $j \in \mathcal{J} := \{1, \cdots, N\}$, places a single ad on the search engine, and the subscript j is used interchangeably for the advertiser and the ad. We assume each ad is separately optimized following AN10. The ads sold in the auction are contingent, meaning that advertisers only pay for their ad display if a consumer clicks on the ad. To simplify notation, all ad-related terms such as bids, valuations, and prices are defined on a per-click basis and will be referred to as "value", "bid", and "price" in

[13] Other factors may influence the click decision such as the characteristics of generic links on the page, but their effect on consumer click behavior cannot be determined due to limitations in our data.

the remainder of this article. Advertiser j has an ad value, v_j, drawn independently from a common distribution F_v supported on $[\underline{v}, \bar{v}]$. In addition, the advertiser-specific click probability, s_j, is drawn independently from a common distribution F_s with bounded support on $[\underline{s}, \bar{s}]$. The auctioneer assigns the quality score $q_j \equiv q(s_j)$ to each advertiser according to the scoring rule $q : [\underline{s}, \bar{s}] \rightarrow [\underline{q}, \bar{q}]$, which is known to auction participants.[14] Then advertiser j's type is given by (v_j, s_j). We define the *weighted value* as the product of the ad value and the quality score of the advertiser denoted as $\omega_j \equiv (v_j \times q_j) \sim F_w$. The potential number of advertisers is denoted by N, and the ads are sold through an auction mechanism. Note that the advertisers are symmetric in the sense that they draw their value and clickability from the common distributions and have access to the same information.

Auction Setup.—A GSP^w is held for each search query to sell ad positions on search results pages. Under the standard assumption of symmetric and independent private values, there are K ad positions to be auctioned and N potential advertisers in each auction. The bidding strategy for advertiser j is defined by a bid function $b_j \equiv b(v_j, s_j)$, where b_j denotes the bid submitted by j based on his type, (v_j, s_j). To account for the effect of click probability on revenue, the auctioneer uses weighted bids instead of original bids. The weighting of bids reflects the impact of the advertiser on the click rate, which is captured by the quality score, q_j. The weighted bid for advertiser j is expressed as $b_w(v_j, s_j) \equiv b_{j,w} = b_j \times q_j$.

Let G and G_w represent the distributions of original and weighted bids, respectively. The k^{th} highest order statistic of the weighted bid is denoted as $b_w^{[k]}$. The ad positions are assigned in descending order of weighted bids, meaning that the k^{th} ad position is awarded to the advertiser with the k^{th} highest weighted bid. The price paid by the advertiser who wins the k^{th} position is equal to the $[k+1]^{th}$ highest weighted bid divided by his quality score.[15] The price for the same ad position may vary across advertisers. The rules of the

[14] The quality score was solely determined by the click rate in the period our empirical application investigates. We maintain this assumption mainly for notational simplicity. Allowing quality scores to depend both on clickability and other factors is straightforward. Given any quality scores the platform defines, all the theoretical results hold because advertiser j's optimization problem depends on the weighted value $v_j \cdot q_j$, not clickability s_j.

[15] For simplicity, it is assumed that there is no reserve price. Yahoo! had a fixed reserve price during the data period but the reserve price was not reported in the data.

auction can be summarized as follows:

$$k^{th} \text{ position allotted to } j \text{ if } \quad b_w^{[k+1]} \leq b_{j,w} \leq b_w^{[k-1]}$$

$$\text{price for } k^{th} \text{ position paid by } j\text{:} \quad p_{k,j} = \frac{b_w^{[k+1]}}{q_j}.$$

In the GSP^w auction, advertisers do not submit their values as bids, unlike in the standard second-price auction. As the bid influences both the winning probability and the price, bidding less than one's value can result in a strategic advantage. Advertisers face a trade-off when determining the bid amount relative to their value. They must balance the benefit of increasing their bid (higher chance of winning a higher position) with the cost of a potentially higher price upon winning. This incentive to bid strategically is absent in the second-price auction, where the bid solely affects allocation and not the price.

The following presents the information setup and assumptions necessary to derive the equilibrium.

ASSUMPTION 1: *(Incomplete information) Each advertiser knows their type, (v, s), and the scoring rule, q, but does not know the opponents' bids, quality scores, and values. They only know the weighted value distribution, F_w. The number of potential advertisers (N), click rates across ad positions ($\mathcal{C}$) and the number of ads per page (K) are common knowledge.*

ASSUMPTION 2: *The advertiser's weighted bid in the GSP^w auction is strictly increasing in his weighted value.*

Assumption 1 describes the standard incomplete information setting. Assumption 2 is required to guarantee the existence of a unique symmetric equilibrium.

Profit maximization problem.—Denote advertiser j's profit from winning k^{th} position as $\pi_{k,j}$. This quantity can be expressed as:

$$\text{profit from } k^{th} \text{ position:} \quad \pi_{k,j} = \underbrace{(c_k \times s_j)}_{\substack{\text{Prob. of click} \\ \text{at position } k}} \times \underbrace{(v_j - p_{j,k})}_{\substack{\text{Per click profit} \\ \text{at position } k}}. \tag{2}$$

Since each advertiser submits a single bid to acquire one of the K available ad positions, the equilibrium bid maximizes the expected total profit, which is defined as the sum of the product of the profit for each ad position and the probability of winning that position.

$$\text{profit from the auction:} \quad \Pi(b_j; v_j, s_j) = \sum_{k=1}^{K} \underbrace{P(b_{j,w} = b_w^{[k]})}_{\text{Prob. of winning position } k} \underbrace{\mathbb{E}(\pi_{k,j} | b_{j,w} = b_w^{[k]})}_{\text{Profit from position } k}$$

Using (2), we obtain

$$\Pi(b_j; v_j, s_j) = \sum_{k=1}^{K} P(b_{j,w} = b_w^{[k]})(c_k \times s_j) \left(v_j - \mathbb{E}[p_{j,k} | b_{j,w} = b_w^{[k]}] \right). \tag{3}$$

We are also interested in "bid shading" $(v_j - b_j)$, which is defined as the difference between the advertiser's bid and their ad value.

3.2. *Identification Analysis*

The primary focus of this paper is to estimate the distribution of advertisers' valuations, denoted as F_v. This distribution conveys advertisers' willingness to pay for an advertisement and serves as a basis for counterfactual analysis. Given Assumptions 1-2, we first demonstrate the existence of a unique symmetric equilibrium in a GSP^w auction. The equilibrium bid is determined by maximizing the profit function in Equation (3):

$$b(v_j, s_j) = \arg\max_{\hat{b}} \sum_{k=1}^{K} P(b_{-j,w}^{[k]} \leq \hat{b} \times q_j \leq b_{-j,w}^{[k-1]})(c_k \times s_j) \left[v_j - \mathbb{E}\left(\frac{b_{-j,w}^{[k]}}{q_j} \middle| b_{-j,w}^{[k]} \leq \hat{b} \times q_j \leq b_{-j,w}^{[k-1]} \right) \right] \tag{4}$$

where $b_{-j,w}^{[k]}$ denotes the ordered ranking for all weighted bids except for advertiser j. In a standard auction, the valuation is identified by inverting the equilibrium bid function. However, in a GSP$^\omega$ auction, the situation is more complex as the weighted bid, $b_w(v_j, s_j)$, depends on both v_j and s_j. As a result, it is not possible to simply invert the bid function to identify the valuation. Therefore, deriving the Bayesian Nash Equilibrium (BNE) in a GSP$^\omega$ auction is a significant extension of the standard BNE derivation in a GSP auction and is considered an important and challenging step forward as stated in GS14.

To address the issue of the non-invertible weighted bid, we prove the equivalence between a GSP^w auction and a *non-weighted* GSP auction in which an advertiser's value is equal to the weighted value. Specifically, we establish that for any advertiser j, the equilibrium weighted bid in the GSP^w auction is equivalent to their equilibrium bid in the GSP auction, where their value is substituted by the weighted value, ω_j. Thus, the bid in the GSP auction performs a similar function as the weighted bid in the GSP^w auction, with the difference being that the bid in the GSP auction can be represented as a univariate function as shown below.

$$b^{GSP}(\omega_j) \to \mathbb{R}_+, \text{ where } \omega_j \equiv v_j \times q_j.$$

In the following theorem, we formally demonstrate that the equilibrium weighted bid in GSP^w only depends on the weighted value ω_j.

THEOREM 1: *Under Assumption 1, if exists, the equilibrium weighted bid in a GSP^w auction, $b_w(v_j, s_j)$, only depends on weighted value ω_j:*

$$b_w(v_j, s_j) = b_w(\omega_j), \quad \forall j \in \mathcal{J}.$$

This theorem establishes that the weighted bid function can be expressed as a function of the advertiser's weighted value, ω_j. Therefore, the decision problem faced by the advertiser is analogous to a non-weighted GSP auction in which the weighted value substitutes the value and the weighted bid $b_w(v_j, s_j)$ can be represented as $b_w(\omega_j)$ at equilibrium. This simplification greatly facilitates the proof of the existence of the unique symmetric equilibrium and identification of the value distribution. We can now derive the unique symmetric equilibrium bidding strategy using Theorem 1 and Assumptions 1–2.

LEMMA 1: *Let Assumptions 1–2 hold. Then, the unique symmetric Bayesian Nash equilibrium of the GSP$^\omega$ auction is given by the following weighted bidding strategy for all $N \geq 2$:*

$$b_w(\omega) = \omega - \Gamma(\omega) - \sum_{n=1}^{\infty} \int_0^{\omega} M_n(\omega, t)\Gamma(t)dt, \quad \forall \omega \sim F_w(.), \tag{5}$$

where

$$\Gamma(\omega) = \frac{\sum_{k=1}^{K} c_k \binom{N-2}{k-1}(k-1)(1-F_w(\omega))^{k-2}\int_0^{\omega} F_w^{N-k}(x)dx}{\sum_{k=1}^{K} c_k \binom{N-2}{k-1}(1-F_w(\omega))^{k-1}F_w^{N-k-1}(\omega)},$$

$$M_1(\omega,t) = \frac{\sum_{k=1}^{K} c_k \binom{N-2}{k-1}(k-1)(1-F_w(\omega))^{k-2}F_w^{N-k-1}(t)f(t)}{\sum_{k=1}^{K} c_k \binom{N-2}{k-1}(1-F_w(\omega))^{k-1}F_w^{N-k-1}(\omega)},$$

$$M_n(\omega,t) = \int_0^{\omega} M_1(\omega,\varepsilon)M_{n-1}(\varepsilon,t)d\varepsilon, \textit{ for } n \geq 2.$$

The equilibrium is efficient if the quality score is equivalent to the advertiser-specific click probability ($q_j = s_j$) and $c_k \geq c_{k+1}$ for all $k \leq K-1$.

Lemma 1 provides a formula for $b_w^{GSP^w}(v_j, s_j)$ that is indeed identical to $b_w^{GSP}(\omega_j)$. The equilibrium weighted bid is calculated based on the weighted value, and distribution and density functions of the weighted values. However, the complexity of the functional form presents challenges in terms of determining the advertiser's valuation through inversion of the equilibrium bid function, due to the presence of an infinite sum and multiple layers of integral. Furthermore, the distribution and density functions of the weighted bid are never known to the researcher. To address these obstacles, an alternative method for deriving the advertiser's valuation is proposed. It does not require inversion of the equilibrium bid, the solution of differential equations, or knowledge of the distribution and density functions of the latent valuation. As demonstrated below, under Assumption 2, the maximization problem (4) can be modified as

$$b_w(\omega_j) = \arg\max_{\hat{b}_w} \sum_{k=1}^{K} P(\omega^{[k+1]} \leq b_w^{-1}(\hat{b}_w) \leq \omega^{[k-1]})c_k\left[\omega_j - \mathbb{E}\left(b_{-j,w}^{[k]} \middle| \omega^{[k+1]} \leq b_w^{-1}(\hat{b}_w) \leq \omega^{[k-1]}\right)\right].$$

where $\omega^{[k]}$ denotes the k^{th} highest weighted value. The following theorem demonstrates the identification of the advertiser's valuation.

THEOREM 2: *Under Assumptions 1–2, the advertiser's value, v, is identified by:*

$$v = b + \Phi(G_w, b, q|\mathcal{C}, K, N), \tag{6}$$

where $\Phi(G_w, b, q|\mathcal{C}, K, N) =$

$$\frac{\sum_{k=1}^{K} c_k \binom{N-1}{k-1}(k-1)(1-G_w(b_w))^{k-2}\int_0^{b_w} G_w(u)^{N-k}du}{q\sum_{k=1}^{K} c_k \binom{N-1}{k-1} G_w(b_w)^{N-k-1}(1-G_w(b_w))^{k-2}\Big[(N-k)(1-G_w(b_w)) - (k-1)G_w(b_w)\Big]},$$

given the quality score (q), the equilibrium bid (b), the distribution function of weighted bids (G_w), the number of available ad positions (K), click rates across ad positions $(\mathcal{C})$, and the number of potential advertisers (N).

A brief overview of the key insights underlying the theorem is given here. The proof employs an indirect approach, utilizing the Revelation Principle, to derive the equilibrium bid. As a result, the following equation for the equilibrium bid is established:

$$\begin{aligned}\sum_{k=1}^{K} c_k \omega \frac{d\xi(\omega)}{d\omega} &= \sum_{k=1}^{K} c_k \binom{N-1}{k-1} b_w(\omega)(N-k)F_w(\omega)^{N-k-1} f_w(\omega)(1-F_w(\omega))^{k-1} \\ &- \sum_{k=1}^{K} c_k \binom{N-1}{k-1}(k-1)(1-F_w(\omega))^{k-2} f_w(\omega)\int_0^{\omega} b_w(x)(N-k)F_w(x)^{N-k-1} f_w(x)dx, \end{aligned} \tag{7}$$

where $\xi(\omega) \equiv P(\omega^{[k+1]} \leq \omega \leq \omega^{[k-1]})$. We apply integration by parts and integration by substitution to the integral part of (7). Then we replace $F_w(\omega)$ with the observed weighted bid distribution $G_w(b_w)$ using Assumption 2. By canceling out the density functions on both sides and rearranging, the desired outcome is obtained.

The term $\Phi(G_w, b, q|\mathcal{C}, K, N)$ represents the bid shading amount of the advertiser. It can also be expressed in terms of the value percentage as follows:

$$\frac{v_j - b_j}{v_j} = \frac{\Phi(G_w, b_j, q_j|\mathcal{C}, K, N)}{v_j}.$$

While the bid shading amount is of interest, the focus is placed on the bid shading percentage above. It is important to note that the actual bids are re-scaled in the data, and therefore, the bid shading percentage provides more accurate information.

REMARK 1: To ensure the existence of the equilibrium in Lemma 1, Assumption 2 is necessary. For structural estimation of valuations, there must exist an equilibrium from which the data are supposed to be generated. Assumption 2 can be verified after estimating valuations by showing that the weighted bids are monotone in the weighted values. Using the empirical distribution of estimated valuations, one can further reveal whether the equilibrium bidding strategy in Lemma 1 is strictly monotone, which implies that the data are rationalizable by the model.[16]

3.3. *The valuation estimator*

In practice, the equilibrium weighted bid distribution G_w is unknown and must be estimated using a consistent estimator, such as the empirical distribution. The following theorem proposes a consistent valuation estimator and proves its convergence properties.

THEOREM 3: *Let Assumptions 1–2 hold. Suppose that there are M number of repeated auctions recorded in the data, each with N number of bidders who draw their valuations i.i.d. from F_v. Let n denote $N \cdot M$. Then given $\left\{\{q_{mj}, b_{mj}\}_{j=1}^{N}\right\}_{m=1}^{M}$, K, $\mathcal{C}$, the valuation is estimated by*

$$\hat{v} = b + \Phi(\hat{G}_w, b, q|\mathcal{C}, K, N), \tag{8}$$

where $\hat{G}_w(b_w) = \frac{1}{n}\sum_{m=1}^{M}\sum_{j=1}^{N} \mathbb{1}[b_{mj,w} \leq b_w]$ and $b_{mj,w} = b_{mj} \cdot q_{mj}$. The valuation estimator $\hat{v}$ uniformly converges to the true valuation v at the $\sqrt{n}$-rate as $n \to \infty$.

Given the auction parameters, the valuations are estimated by substituting the empirical distribution $\hat{G}_w$ into G_w in (6). The integral component $\int_0^{b_{j,w}} \hat{G}_w(u)^{N-k} du$ can be numerically integrated using trapezoidal sums.[18] Unlike the nonparametric estimator in Guerre

[16] Assumption 2 is untestable without having valuations that are obtained only under the assumption. The same assumption has been imposed and verified by data in other empirical auction papers such as Haile et al. (2003) and Guerre et al. (2000).[17] Note that the assumption does not restrict the equilibrium bid function. Proposition 2 in GS14 shows that the formula in Lemma 1 can result in non-monotonic bidding. The symmetric equilibrium does not exist in this case.

[18] The 'trapz' command in MATLAB is used for numerical integration in our calculations.

et al. (2000), our estimator does not require trimming to correct the boundary problem as it does not involve nonparametric density estimation of observed bids. Given the estimated values, distribution and density functions of valuations are estimated by the empirical distribution and the kernel density estimator respectively.[19] The valuation estimator uniformly converges to the true value as fast as the empirical distribution and much faster than the kernel density estimator. Therefore, the empirical distribution and the kernel density of valuations converge to the true distribution and density respectively at standard rates ($\sqrt{n}$-rate for the empirical distribution and $\sqrt{nh}$-rate for the kernel density).

We conduct simulation experiments on known data generating processes to examine the finite sample performances of our estimator in Appendix C. The results show that our estimator can precisely estimate the underlying value distribution with a relatively small sample size. We further investigate the impact of the correlation between value and quality on auction outcomes. Advertisers can shade bids more when the market is less competitive e.g., N is small or value and quality are more positively correlated.

4. YAHOO WEBSCOPE DATA

We apply our method to large search auction data from the Yahoo Webscope Program.[20] The data cover all search queries except those specific to brands from January to April 2008 across five product categories: *Cruise, Car Insurance, Laptop, Cable TV,* and *Collectible Coins*. The data include aggregated information on search keywords, bids, clicks, ad positions, and display frequencies. The keywords consist of a base word with one or more additional words. For instance, 'business laptop' and 'student laptop' fall under the *laptop* category. Identities of advertisers and keywords are masked, though we can track the same advertiser and keyword over time. While the base categories are anonymized, we can make

[19] Given the estimated values $\{\hat{v}_{m1}, \cdots, \hat{v}_{mN}\}_{m=1}^{M}$, the value distribution is estimated by $\frac{1}{n}\sum_{m=1}^{M}\sum_{j=1}^{N}\mathbb{1}[\hat{v}_{mj} \leq v]$ for $v \in [\underline{v}, \bar{v}]$. Similarly, the density is estimated by $\frac{1}{nh}\sum_{m=1}^{M}\sum_{j=1}^{N}K\left(\frac{\hat{v}_{mj}-v}{h}\right)$ for $v \in [\underline{v}, \bar{v}]$ where $K(\cdot)$ is a kernel function and h is a bandwidth that depends on the sample size n.

[20] Advertising & Markets Data, A3. Yahoo! Search Marketing Advertiser Bid-Impression-Click data on competing Keywords (version 1.0). Researchers at accredited universities can request access to the data through Yahoo! Labs upon the Data Sharing Agreement.

credible speculations about the categories based on observed characteristics, as discussed in Appendix E.

The raw data consist of 77,850,272 observations. The data are narrowed down to ads displayed on the first two pages of search results as click rates on subsequent pages are negligible. The bid amounts, scaled by an unknown amount for masking purposes, are recorded in cents. We rescale them by subtracting an amount close to the lowest. We restrict data to February to eliminate extremely high bid values from several advertisers.[21] The dataset is further limited to the top 10% most popular keywords to rule out keywords with distinct value distributions and prevent multiple keywords from entering the same auction.[22] The most frequent searches are likely to share similar ad values. For instance, the most popular keywords in the *laptop* category in 2008 were 'hp laptop', 'dell laptop', 'laptop computer', 'best laptop', and 'laptops' according to Google Trends. Clustering similar keywords is not possible due to keyword masking. Lastly, keywords missing some ranks from 1 to 14 each day are dropped. The restricted data contain 1,610,333 observations.

The aggregate data do not correspond to individual auctions. Each entry represents a combination of category, day, keyword, advertiser, and ad position, for which the data include the bid amount, number of impressions, and number of clicks. This allows us to observe how frequently an advertiser secures a particular ad position for a keyword on a specific day within a category. For instance, the data report that *firm #1*'s ad specifying *business laptop* as a keyword was displayed in the top position 100 times, resulting in 5 clicks on January 1st. For the same keyword-advertiser pair, there is a separate observation for the ad displayed in lower positions. The variations in click rates across different positions

[21] 5 advertisers have recorded bids of \$999.99, which is 80 times greater than the 99th percentile of all bids (\$12.5). These extremely high bids were deemed outliers as these are unusual bidding behaviors and their removal was necessary to obtain a more reasonable bid distribution. These outliers never appeared in February when the data were restricted to the top 10% most popular keywords.

[22] Yahoo! prioritizes matching user queries with ads containing the most relevant keywords. When insufficient competition exists for a specific keyword, the platform may broaden the candidate pool to include ads with less precise keyword matches. To mitigate this concern, we focus on the most popular keywords that have a significantly higher average number of advertisers than less prevalent keywords.

TABLE II

DESCRIPTIVE STATISTICS

Variable	Cruise	Car Insurance	Laptop	Cable TV	Coin
Number of advertisers[1]	45	403	124	142	55
Number of Keywords	67	33	110	68	49
Keylength (mean)[2]	2.28	3.02	2.33	2.16	2.07
Keylength (max)	(4)	(5)	(4)	(4)	(4)
CTR (mean)	0.87%	0.62%	0.99%	0.75%	0.94%
Bid (mean)	1.15	6.94	0.83	0.90	0.46
Keyword popularity (mean)[3]	1355	9664	5514	1820	1857

Note: 1) Computed as the maximum number of bidders across day-keyword combinations. 2) The average number of words. 3) Measured by the daily average of the number of times an ad was matched using the keyword in the first position.

for the same advertiser identify the advertiser-specific click probability and the position-specific click rates separately. Most advertisers maintain a consistent bid throughout the day, thus we observe the set of all bids placed by the advertisers per day.[23]

The click-through rate (CTR) is determined by the ratio of clicks to the number of displays for each ad.[24] The average CTR across all ads is 0.8%. The keyword associated with each ad provides insights into the type of search query. The ads matched with longer keywords typically correspond to more specific search queries and are therefore considered more valuable for advertisers (Ramaboa and Fish, 2018). We define the number of words in the keyword as keylength. There are 327 distinct keywords in the data, with a maximum of 5 words and an average of 2.4 words per keyword. We measure the popularity of keywords as the daily average of the number of times an ad was matched using that keyword in the first position. The descriptive statistics are reported in Table II.

[23] A small portion of advertisers (5.56% of the restricted dataset) adjusted their bids during the day. To account for this behavior, the bids for these advertisers were averaged for each day.

[24] 27 entries have click rates greater than 1, which may be due to recording error or consumer mistake. The data do not provide information to examine this behavior so we impute the click rate 1 if it exceeded 1.

REMARK 2: Our data do not reveal which ads were shown together on a search page. For estimating valuations, however, we only need to know the observed bid distribution and the number of potential advertisers (N) in each auction given the position-specific click-through rates and quality scores. We need not know the exact identities of participating advertisers. Most advertisers placed their fixed daily bids for each day so the bid distribution is identified from aggregated data. We also assume that advertisers set their equilibrium bids based on the potential number of competitors (N). Following Guerre et al. (2000), we identify N as the maximum number of advertisers observed across all markets (day-keyword combinations).

5. ESTIMATION PROCEDURE WITH YAHOO DATA

We define the market, denoted by m, as each combination of category, day, and keyword. Variables are indexed by the market, reflecting that observations are collected from multiple markets. The estimation procedure involves two steps: 1) Estimating advertiser- and position-specific effects on click probability, and 2) Calculating the advertiser's value based on the observed bids.

Step 1: Estimation of advertiser- and position-specific effects on click probability

The position-specific effect is used to quantify the click rate for each ad position. The advertiser's effect on click probability is used to compute the quality score so that $s_j = q_j$.[25] To analyze consumer click decisions, we use the following linear probability model (LPM) with a log transformation and incorporate the aggregate CTR for each category-day-keyword-ad position-advertiser combination:

$$\log(ctr_{k,j,m}) = \alpha_k + \beta_j + \gamma_m + \epsilon_{k,j,m} \tag{9}$$

where $ctr_{k,j,m}$ is the CTR, α_k is the position fixed effect, β_j is the advertiser fixed effect, γ_m is the market fixed effects estimated by including keyword dummies and the weekend

[25]When our data were collected, the quality score captured the advertiser's impact on click probability. The current practice of calculating the quality score factors in consumer and ad-display characteristics.

dummy, and $\epsilon_{i,j,m}$ is the idiosyncratic shock. To prevent log transformation from producing $-\infty$, we replace zero values with the minimum of observed CTRs. This specification respects the non-negative nature of the CTR and is fully saturated as only dummy variables are included. Thus, the specification is fully nonparametric and does not require distributional assumptions on $\epsilon_{i,j,m}$. Furthermore, the specification implies multiplicativity of click probability and is aligned with our theoretical model because taking exponential on both sides of (9) yields

$$ctr_{k,j,m} = exp(\alpha_k + \beta_j + \gamma_m + \epsilon_{k,j,m}) = \underbrace{exp(\alpha_k)}_{=c_k} \cdot \underbrace{exp(\beta_j)}_{=s_j} \cdot exp(\gamma_m) \cdot exp(\epsilon_{k,j,m}). \quad (10)$$

The parameters are estimated by weighted least squares, where the weight is determined by the number of impressions because the click probabilities are more precisely estimated for more frequently displayed ads.

REMARK 3: The separability condition (1), commonly imposed in the literature, may not universally hold in empirical contexts. Jeziorski and Moorthy (2018) show that lesser-known advertisers might witness a larger increase in click rates than well-known counterparts when their ads appear in top positions. We conduct a robustness analysis to check whether the multiplicative separability assumption imposed on the click rate holds in our data. Specifically, we generalize the regression model (9) by allowing the fixed effects specific to each advertiser-position pair. Our analysis, provided in Appendix G, demonstrates that heterogeneity among advertisers in position-specific click-through rates is limited, and multiplicative separability offers a very reasonable approximation for click probability within our data.

Given the estimated parameters, two key parameters for the second step estimation are computed. First, we compute the position-specific click rate ($\hat{c}_k$) using the estimated position-specific fixed effect, $exp(\hat{\alpha}_k)$. We normalize the click rates by dividing them by the click rate of the first position in each category:

$$\hat{c}_k = exp(\hat{\alpha}_k - \hat{\alpha}_1)$$

The estimated fixed effects for positions on the first page ($k = 1, \cdots, 7$) are highly significant (p-value < 0.0000) but not for the second page ($k = 8, \cdots, 14$). As a result, we limit our analysis to only the first page of ad positions ($K = 7$), assuming positions beyond the first page have negligible effects on advertiser payoffs. There are substantial gaps between the click rates of the first and second positions in all categories. A decrease in click rate was also noted as the ad position became lower, though the decline was not as steep as the drop between the first and second positions. In *car insurance* and *coin* categories, the decrease was not strictly monotonic. This pattern was consistent with the simple means of $ctr_{k,j,m}$ across ad positions. Second, we determine the quality score ($\hat{s}_j$) using the estimated advertiser fixed effect, $exp(\hat{\gamma}_j)$. This quantity is then normalized by dividing with the highest value of $exp(\hat{\gamma}_j)$. The estimated position-specific click rates and quantiles of quality scores for each product category are presented in Table III.

Step 2: Estimating the advertiser's value and its distribution

In this step, we obtain the weighted bids by multiplying the observed bids with the estimated quality scores ($\hat{b}_{j,w} = b_j \hat{s}_j$). The distribution of weighted bids is then estimated by the empirical distribution ($\hat{G}_w$). The valuations (v_j) are computed using Equation (6) presented in Theorem 2. In this calculation, the unknown parameters G_w, $\mathcal{C}$, and q_j are

TABLE III

POSITION-SPECIFIC CLICK RATES AND QUANTILES OF QUALITY SCORES ACROSS PRODUCT CATEGORIES

Click Rates						Quality Scores					
rank	Cruise	Car Ins.	Laptop	Cable	Coins	qtile	Cruise	Car Ins.	Laptop	Cable	Coins
2^{th}	0.490	0.843	0.627	0.522	0.646	25%	0.002	0.002	0.001	0.005	0.006
3^{th}	0.327	0.480	0.519	0.428	0.537	50%	0.006	0.006	0.004	0.011	0.020
4^{th}	0.142	0.257	0.310	0.269	0.489	75%	0.010	0.030	0.013	0.020	0.045
5^{th}	0.097	0.180	0.421	0.224	0.351	90%	0.022	0.065	0.028	0.025	0.084
6^{th}	0.040	0.310	0.219	0.185	0.246	99%	0.108	0.175	0.174	0.072	0.431
7^{th}	0.037	0.218	0.181	0.149	0.138						

replaced by their respective estimates $\hat{G}_w$, $\hat{\mathcal{C}}$, and $\hat{s}_j$. Finally, the distribution of values is estimated by the empirical distribution of the estimated values.

6. EMPIRICAL RESULTS

We first estimate the advertiser's maximum willingness to pay (valuation) for an ad. The relationship between the weighted value and weighted bid is then verified by drawing the weighted bid as a function of the estimated weighted value in Figure 4. We also compute the equilibrium bid functions given the empirical distribution of estimated values using Lemma 1. For all product categories, the weighted bid is strictly monotone in the weighted value (Assumption 2 holds), ensuring the existence of the equilibrium in each category and implying the data are *rationalized* by our model. The results show that the equilibrium weighted bid function deviates significantly from the 45-degree line as the weighted value increases. The bid shading amount (the distance between the 45-degree line and the weighted bid function), also increases monotonically with the weighted value, implying that advertisers with higher weighted values can shade their bids more. We further compute the correlation between estimated values and quality scores by category in Figure 5. All the categories exhibit a weak positive correlation. The correlation is the strongest in *Car insurance* and near 0 in *cruise*.

Table IV summarizes the quantiles of bid shading percentages across categories. The median bid shading percentage is close to 0 in all categories, with the highest being less than 0.8%. The extent to which advertisers can shade their bids varies across product categories, with advertisers in the *laptop*, *cruise*, and *coins* categories shading their bids more than their counterparts in the *car insurance* and *cable TV* categories. We find that market characteristics such as the number of advertisers, position-specific click-through rates, and the distribution of quality scores are closely related to different bid shading behaviors. For example, while *laptop* and *cable TV* have similar numbers of advertisers (124 and 142 respectively), the steeper drop in position-specific click rates and less skewered quality score distribution in *cable TV* leads to less bid shading compared to *laptop*. Similarly, the low number of advertisers (55), smaller decrease in click rates across positions, and more skewed score distribution in *coins* result in the highest bid shading percentages at high per-

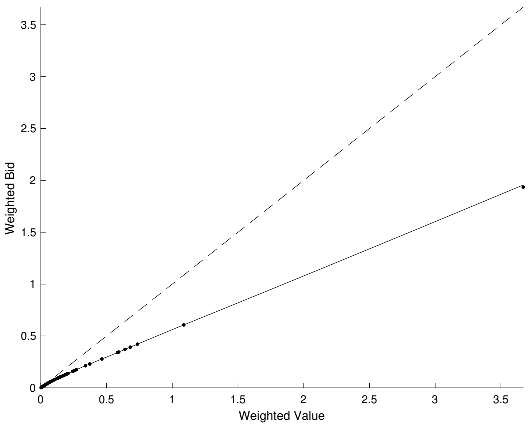

(a) Cruise

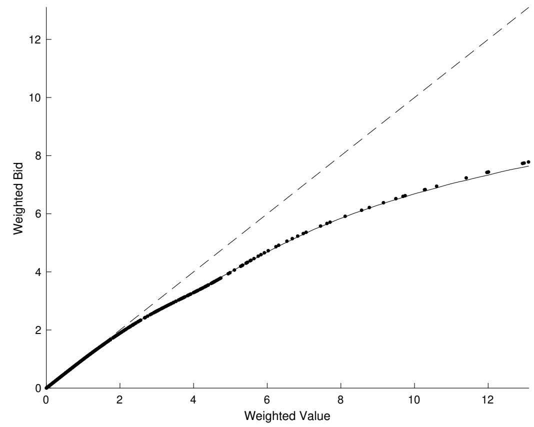

(b) Car Insurance

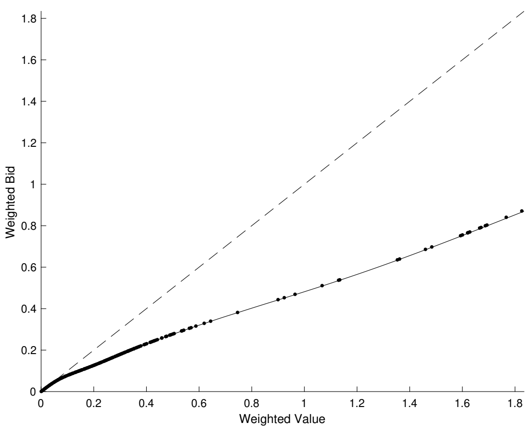

(c) Laptop

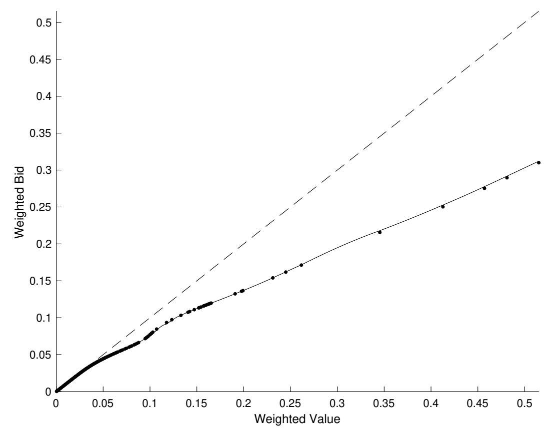

(d) Cable TV

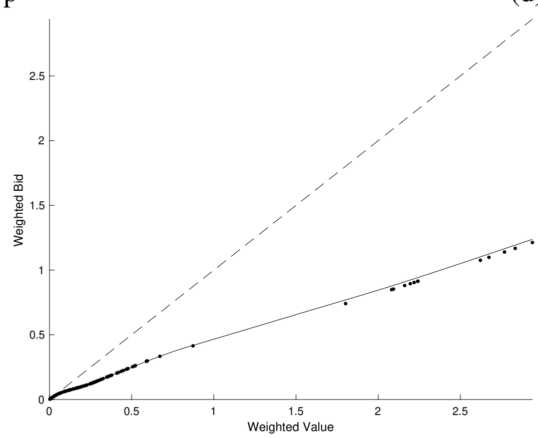

(e) Coins

FIGURE 4.—Relationship between estimated weighted values and observed weighted bids. The solid lines are the equilibrium bid functions given the estimated value distributions computed by the iterative approximation procedure in Appendix D. The dashed line is a 45-degree line which represents truthful bidding.

centiles. *car insurance* shows the lowest bid shading level despite the most gradual drop in click rates across positions because of the largest number of competing advertisers (403).

TABLE IV

QUANTILES OF BID SHADING PERCENTAGE ACROSS PRODUCT CATEGORY

	Cruise	Car Ins.	Laptop	Cable TV	Coins
25%	0.164%	0.003%	0.025%	0.010%	0.084%
50%	0.769%	0.011%	0.153%	0.068%	0.605%
75%	6.367%	0.069%	0.745%	0.270%	2.878%
90%	13.339%	0.379%	4.825%	1.792%	20.161%
99%	29.366%	17.381%	40.256%	22.927%	51.050%

We also examine the distribution of valuations for each category. We compare the empirical distributions of estimated values with the observed bid distributions. Figure 6 demonstrates that *car insurance* has the highest per-click values, followed by *cruise*, while *coins* has the lowest. *Cable TV* has a similar value distribution to *cruise* below the median, but with fewer advertisers with high per-click values. The value distribution in *laptop* falls between *coins* and *cable TV*, but this category has more advertisers with high values. Our results further suggest that ad values in general follow a log-normal distribution. Appendix F includes the approximated value distributions using log-normal distributions, which fit particularly well for *cruise*, *cable TV*, and *coins*. This supports the extensive use of log-normal specifications in the theoretical auction literature for Monte Carlo simulations.

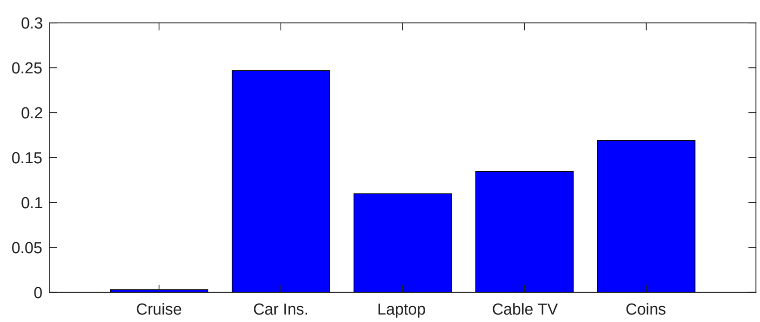

FIGURE 5.—Correlation between value and quality score by category.

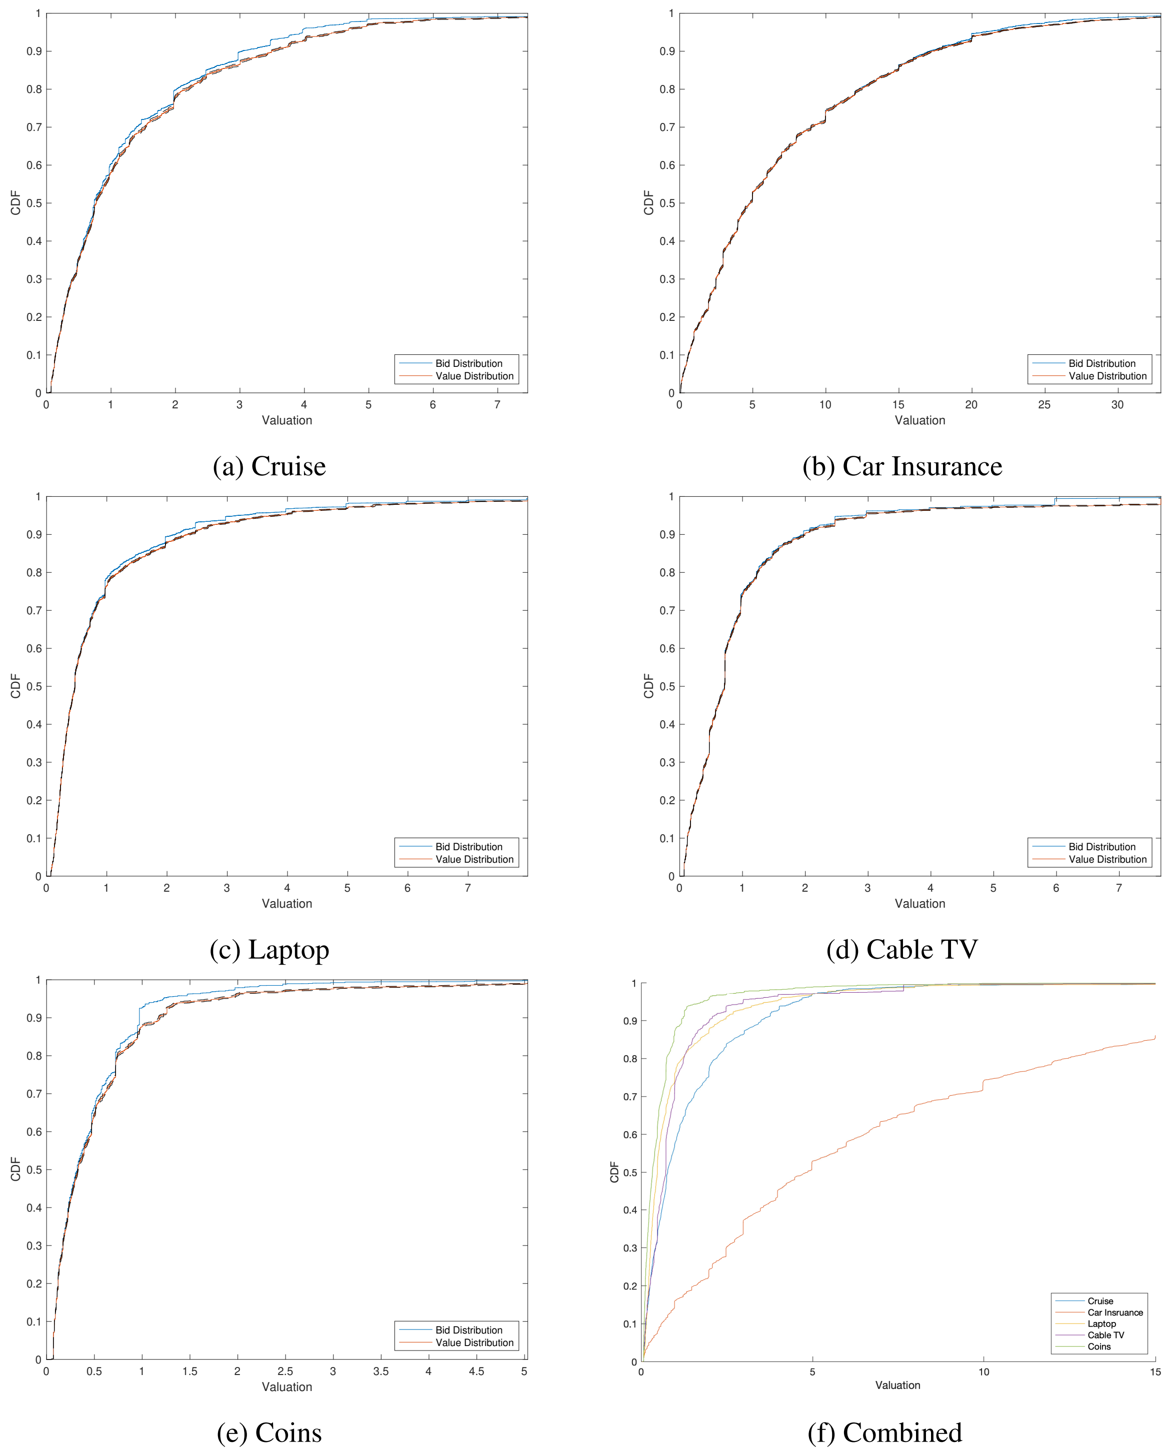

(a) Cruise

(b) Car Insurance

(c) Laptop

(d) Cable TV

(e) Coins

(f) Combined

FIGURE 6.—The cumulative distribution of advertiser's ad value. The black dashed lines are the 90% confidence bands of the estimated valuation distribution obtained by bootstrap with 200 replications.

Lastly, we document the average estimated values and quality scores across average ad positions in Table V. The average per-click values in general drop steeply across positions, except in *car insurance* where the level of competition is the fiercest. The average values for the top 3 positions in *car insurance* are very close to each other and the average value in the lowest rank range is close to half of that in the highest rank range. The average quality scores are not always monotone across ad positions.

TABLE V

QUALITY SCORE AND VALUE ACROSS AVERAGE RANKS

	Quality Score					Value				
Average	Cruise	Car Ins.	Laptop	Cable	Coin	Cruise	Car Ins.	Laptop	Cable	Coin
$[1,2)$	0.0095	0.0276	0.0351	0.0404	0.0664	4.75	13.57	3.78	3.61	1.31
$[2,3)$	0.0087	0.0281	0.0299	0.0152	0.0357	2.26	12.97	2.07	2.41	1.05
$[3,4)$	0.0082	0.0308	0.0200	0.0155	0.0360	2.39	12.65	1.83	1.68	1.08
$[4,5)$	0.0118	0.0286	0.0164	0.0168	0.0379	2.05	10.61	1.43	1.43	0.91
$[5,6)$	0.0159	0.0249	0.0155	0.0172	0.0370	1.62	8.52	1.12	1.25	0.72
$[6,7)$	0.0148	0.0234	0.0130	0.0150	0.0417	1.25	6.92	0.84	1.12	0.51

Note: The table summarizes the average quality scores and values across average ranks.

REMARK 4: **(Keyword heterogeneity)** We use the full keyword sample to examine keyword heterogeneity. The results in Appendix G show substantial keyword heterogeneity in the full sample. The perceived value of a click may vary depending on the specific keyword. For example, in *laptop*, an advertiser may expect higher profit from a click on an ad associated with "high-performance laptop" than "budget laptop". Since keywords in our data are masked, we cannot cluster similar keywords. Instead, we assume the most commonly searched keywords share the same valuation distribution. One could apply our framework after clustering keywords using unsupervised machine learning techniques if keyword identities are available.

7. COUNTERFACTUAL ANALYSIS

"When someone has a really high ad click probability, they're very hard to beat, so it's not a really competitive auction. So that they don't just win [every auction], we do squashing. This makes the auction more competitive. It's like handicapping. We handicap the people with the high click probability." – Preston McAfee[26]

In this section, we conduct counterfactual experiments to investigate the impacts of alternative quality scoring schemes on auction outcomes such as the search engine's revenue, the advertiser's profit, and consumer welfare. We first consider score squashing, which involves attaching monotonic nonlinear transformations of advertiser-specific click probability to bids. Secondly, we examine user targeting which personalizes ad displays through computing quality scores based on user characteristics. These alternative scoring rules aim to increase revenue at the cost of efficiency.

7.1. *Score squashing*

We begin by investigating score squashing under which the squashed quality score is expressed as follows:

$$\text{New score} \rightarrow q_j^{SQ} = s_j^{\theta} \quad \forall j.$$

$\theta \in [0, 1]$ is called the squashing factor.[27] Lahaie and Pennock (2007) show that the impact of squashing depends on the rankings of advertisers based on quality scores, weighted bids, and squashed weighted bids. Revenue increases if all three rankings are the same in the top $K + 1$ positions because squashing increases prices for high-quality advertisers without altering allocations. Conversely, revenue decreases if both weighted bid and squashed weighted bid rankings differ from the quality score ranking. The revenue consequence is ambiguous if the weighted and squashed weighted bid rankings differ. These insights provide an understanding of how squashing changes revenue. However, drawing

---

[26]Chief economist of Yahoo! in 2010. See https://www.theregister.com/2010/09/16/yahoo_does_squashing/ for details on score squashing practice at Yahoo!.

[27]s^{θ} inflates to 1 for all s and the relative importance of the quality score shrinks as θ approaches 0.

practical implications is challenging. In practice, it is more useful to examine how the outcome changes with auction characteristics at the search-topic level, such as the shape of value and quality score distributions, position-specific click rates, and the competition level in the market.

We aim to address the knowledge gap surrounding the effects of score squashing through a comprehensive counterfactual analysis. While AN10 has explored the impact of score squashing with a fixed squashing factor of $\theta = \frac{1}{2}$, our study extends this work by considering a wider range of squashing factors and examining the role of market-specific factors in determining the effects of squashing on auction outcomes. The equilibrium bid in the case of score squashing is determined by the distribution of the squashed weighted value, $\omega_j^{SQ} = v_j \times q_j^{SQ}$. We calculate the equilibrium bid and compare auction outcomes under different squashing schemes. We examine the impact of θ on auction outcomes. The squashing factor θ varies in $\{0, 0.1, 0.2, 0.3, \cdots, 0.9, 1\}$. We simulate 2000 auctions for each category-squashing factor combination and collect the resulting auction outcomes. To evaluate the performance of the auction, we calculate the average values of revenue, profit, and consumer welfare (measured by ad quality) across all simulated auctions.

The observed bids in the data cannot be used in our analysis, as the equilibrium bidding strategy varies with the squashing factor. We calculate the equilibrium bid function given the population distribution and density functions of the weighted value in each scenario using the iterative approximation procedure in Appendix D. Given that the value and quality score distributions in the data are well approximated by log-normal distributions, we assume $v_j \sim L(\mu_v, \sigma_v^2)$ and $s_j \sim L(\mu_s, \sigma_s^2)$.[28] We allow the value and quality score to be correlated. As a result, the weighted value with the squashing factor θ is also log-normally distributed, $\omega_j^{SQ} \sim L(\mu_v + \theta\mu_s, \sigma_v^2 + \theta^2\sigma_s^2 + 2\theta\sigma_{vs})$, where σ_{vs} denotes the covariance between logarithms of v_j and s_j. To estimate the means and variances of the log-normal distributions, we use maximum likelihood estimation (MLE). The estimated distributions approximate the empirical weighted value distributions well.

[28] Suppose $Z \sim N(0, 1)$. Then, $X \equiv exp(\mu + \sigma Z)$ is log-normally distributed. μ and σ represent the mean and standard deviation of $\log(X)$, respectively. We denote the log-normal distribution as $X \sim L(\mu, \sigma^2)$.

The algorithm determining the auction outcomes is outlined in the steps below. Given each category and the squashing factor θ, the following steps are executed:

1. Use MLE to estimate $(\mu_v, \mu_s, \sigma_v^2, \sigma_s^2)$ of the log-normal distributions of the valuation and quality score.
2. Solve for the equilibrium bid function using the weighted value distribution and position-specific click rates $\{\hat{c}_1, \cdots, \hat{c}_K\}$.
3. Draw N values and advertiser-specific click rates from log-normal distributions.
4. Compute the weighted bids and determine the auction outcome. Calculate auction revenue, advertiser profit, and ad quality for top 7 positions.
5. Repeat Steps 3 to 4 for 2000 replications and compute the average revenue, profit, per-click price, and ad quality.

Figure 7 presents the counterfactual results. The revenue-maximizing level of squashing (θ^*) differs across categories, implying that a localized decision for squashing is more effective in maximizing revenue. All categories show strong improvement in revenue through squashing. Per-click price is maximized without quality weighting ($\theta = 0$). However, consumer welfare (measured by average ad quality) reduces monotonically as the squashing parameter decreases. Squashing reduces average ad quality but improves the revenue of the search platform, which implies that the per-click price increase induced by squashing has a greater impact on revenue than the decrease in quality. On the advertiser side, squashing reduces their profit due to higher prices and reduced ad quality of winning ads. Overall, score squashing improves revenue at the cost of advertiser profit and consumer welfare. The maximum sum of revenue and profit is achieved with no squashing, which is naturally the case because the equilibrium with no squashing guarantees efficient allocation. This suggests that the search engine has to consider a higher value of θ than θ^* when caring about the long-term relationship with the advertisers and consumers.

Further, we investigate heterogeneity in the impact of squashing by examining the change in θ^* when switching one of the market parameters. We take *cruise* as the base category and analyze the impact of replacing its position-specific click-through rates, the number of advertisers, value distribution, score distribution, or correlation between value and score.

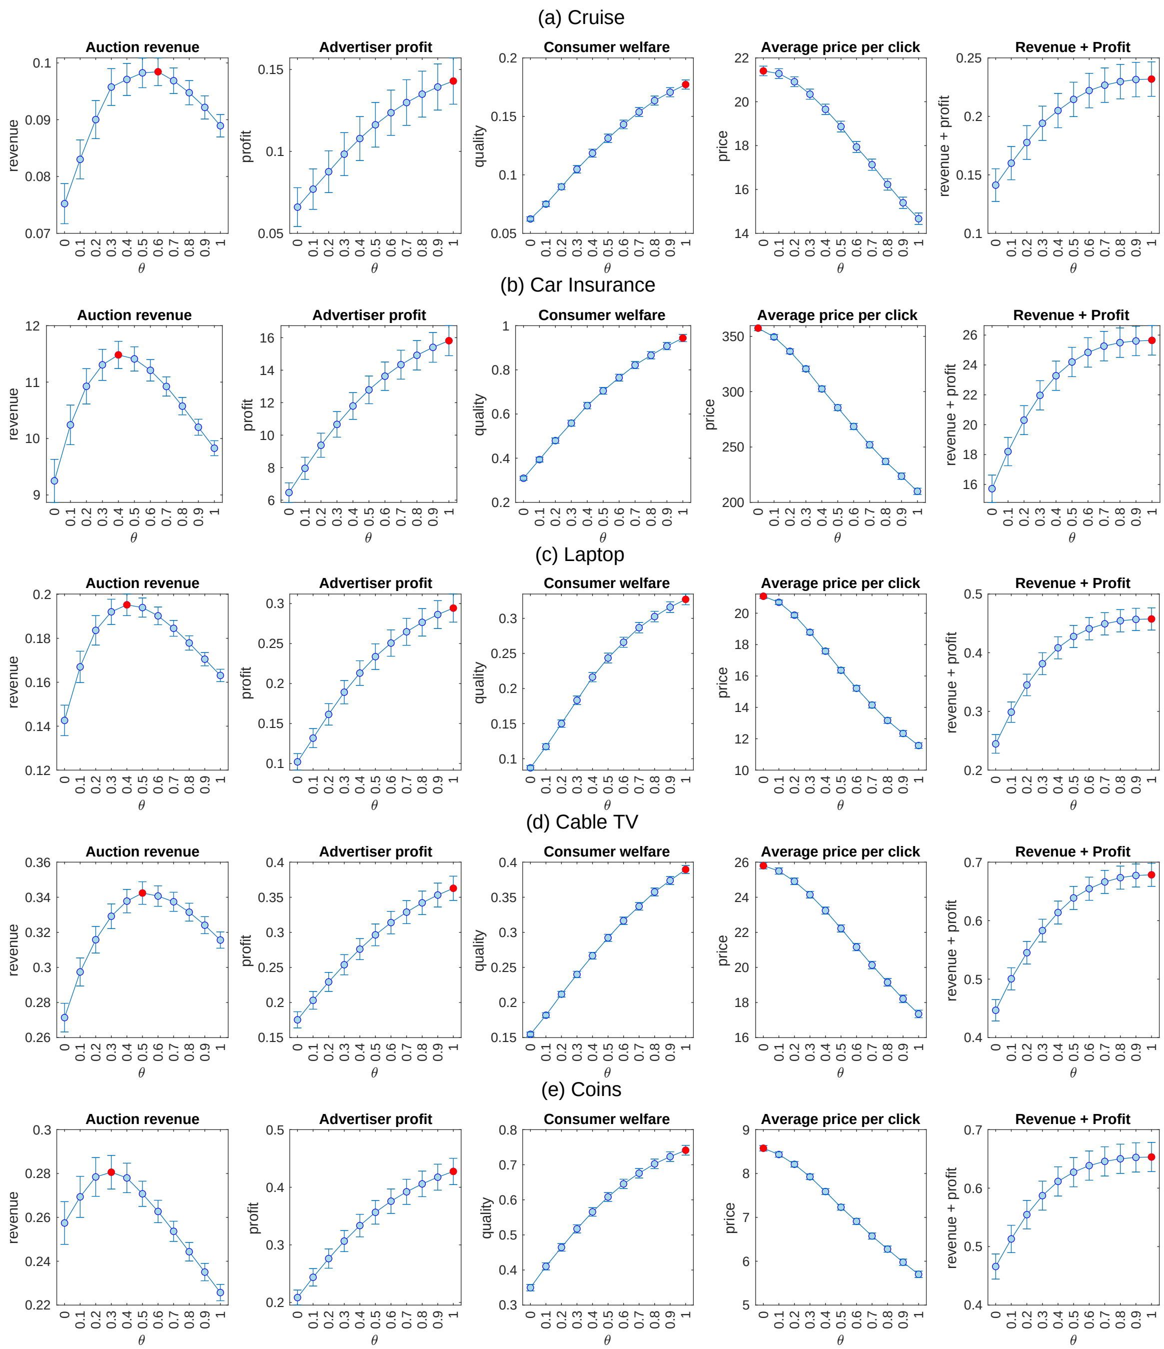

FIGURE 7.—Counterfactual simulations with score squashing. The dots are averages across 2000 simulated auctions. The red dots are maximums across values of θ. The ranges are 95% confidence intervals.

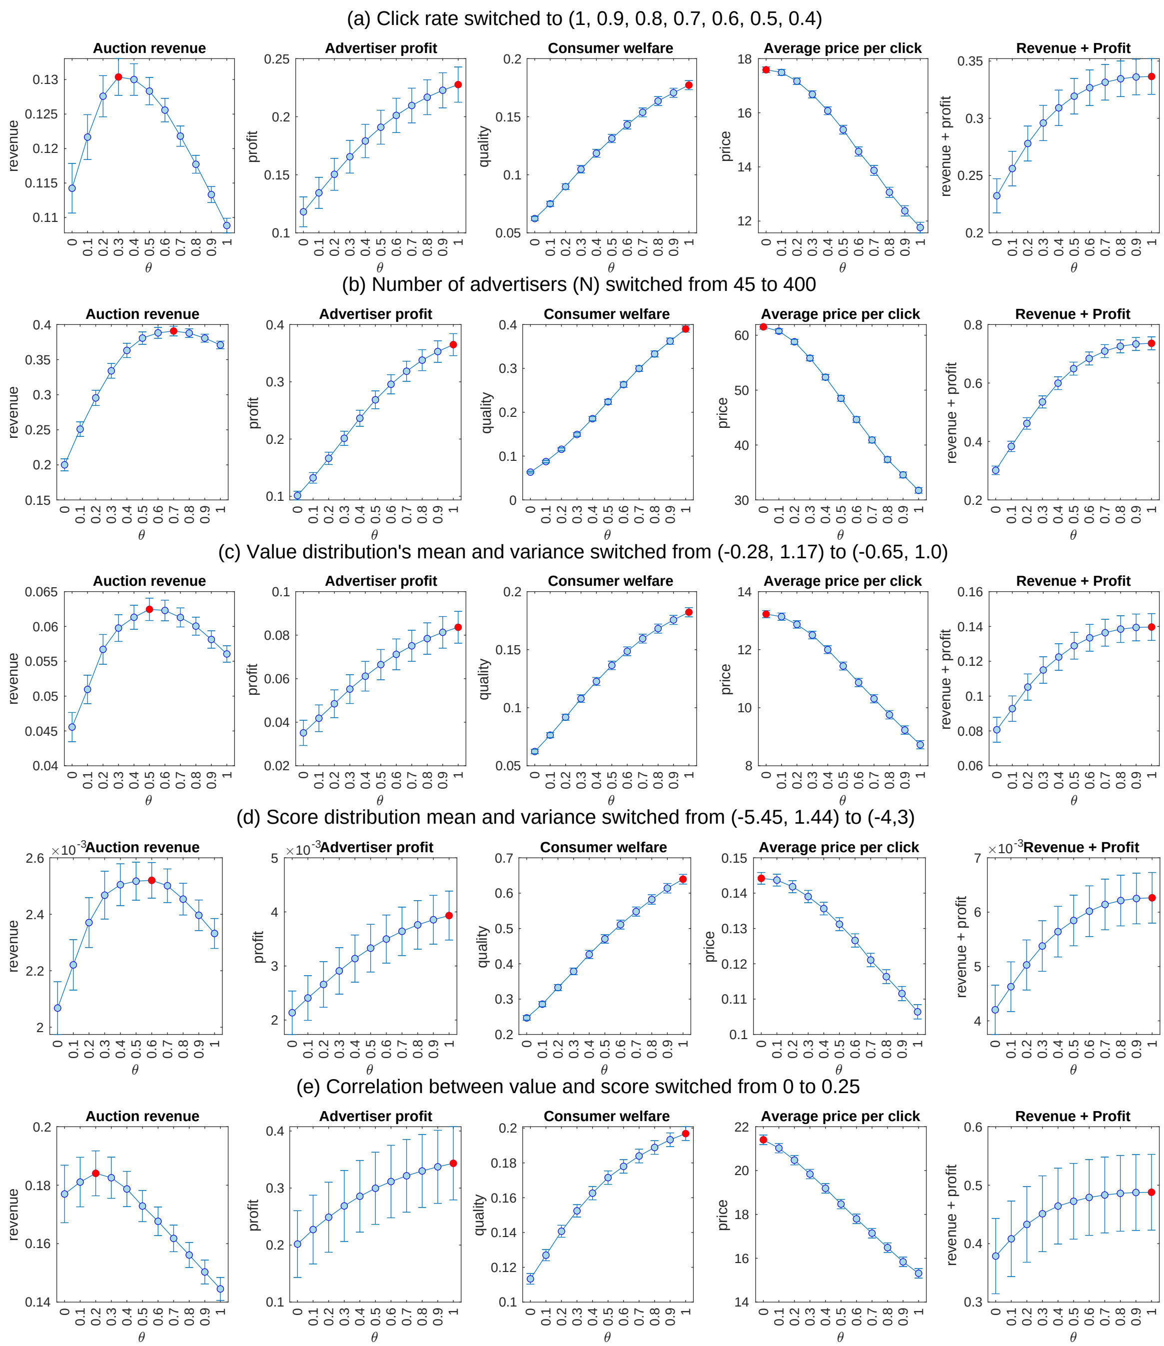

FIGURE 8.—Counterfactual simulations with switching auction characteristics of *Cruise*. The dots are averages across 2000 simulated auctions. The red dots are maximums across values of θ. The ranges are 95% confidence intervals.

The results are presented in Figure 8. We find that more squashing (smaller θ) is revenue-maximizing when the market is less competitive. We first change the click-through rates from $(1, 0.49, 0.33, 0.14, 0.1, 0.04, 0.04)$ to $(1, 0.9, 0.8, 0.7, 0.6, 0.5, 0.4)$ so that they drop much more gradually across ad positions. This configuration decreases θ^* (from 0.6 to 0.3) as it makes the first position much less attractive, implying stronger incentives for bid shading. The same pattern is observed when we switch the number of advertisers from 45 to 400. θ^* jumps up to 0.7 as the market becomes more competitive.

Next, we examine the impact of the value and quality score distributions. By switching the mean and variance of the value distribution from $(-0.28, 1.17)$ to $(-0.65, 1.0)$ (similar to *laptop*), we find a lower θ^* (0.5). This switch results in less right-skewed value distribution, meaning less competitive auctions. Switching the quality score distribution's mean and variance from $(-5.45, 1.44)$ to $(-4, 1.3)$ (similar to *coins*) does not change θ^*. Lastly, we switch the correlation level from 0 to 0.25 (close to that of *car insurance*), which results in a much lower θ^* (0.2) as a higher correlation between value and quality score makes the auction less competitive.

TABLE VI

CHANGES IN AUCTION OUTCOMES WITH REVENUE-MAXIMIZING SQUASHING (θ^*)

	Cruise	Car Ins.	Laptop	Cable	Coins
θ^*	0.6	0.4	0.4	0.5	0.3
Revenue	10.7%	16.8%	19.7%	8.5%	24.3%
Profit	-13.5%	-25.4%	-27.6%	-18.4%	-28.3%
Welfare	-19.2%	-32.4%	-33.9%	-25.0%	-30.2%

Note: The table summarizes the impact of squashing with θ^* on auction outcomes compared to the case without squashing.

Our experiments shed light on the impact of various factors on the revenue-maximizing level of θ in advertising markets. More competitive market environments in general lead to lower θ^* (the stronger penalty for high-quality advertisers). Score squashing can lead to an increase in auction revenue, ranging from 8.5% (for *cable TV*) to 24% (for *coins*), albeit at the cost of advertiser profit and consumer welfare. Counterfactual changes in auction outcomes given θ^* are summarized in Table VI. Our framework provides a comprehensive approach to determining the most appropriate level of squashing for a given auction. Search engines have the flexibility to adopt an objective function that reflects their priorities by choosing a weighted average of the auction outcomes.

7.2. *User targeting (personalization)*

The search platform can utilize user information in calculating quality scores for user targeting, which matches ads to different types of users to improve click-through rates. In practice, quality scores are computed by statistical algorithms. Search platforms must decide on the algorithm's precision and the level of personalization to individual users to maximize revenue. More precise and personalized scores drive higher click rates but diminish advertiser competition. To illustrate this trade-off, consider a scenario involving gender-specific advertisers, with half appealing to males and the other half appealing to females. For instance, in the clothing category, we might have male-oriented brands and female-oriented brands. With user targeting, the platform considers user gender while deciding the quality score. The level of personalization/accuracy increases with higher weight put on the best ad match for each gender type. This approach yields more clicks but attracts fewer advertisers to the auction, compared to the case where quality scores are computed without factoring in gender information. Thus, to understand this accuracy-competition trade-off, we introduce flexibility in how the search engine weights user type when calculating quality scores, allowing it to differ from the actual click rates.

AN10 investigate the impact of employing a less precise quality scoring algorithm (referred to as "score coarsening" in their paper) on auction outcomes by introducing noise to observed quality scores. They found that this approach increases revenue by 2%–18% depending on the search phrase. Our analysis differs from AN10 in the sense that we consider varying degrees of the user targeting level in calculating quality scores. To analyze the platform's incentive to reduce the precision of quality scores, we conduct the following counterfactual analysis. Consider a case where the search engine has access to the gender information of users, with an equal gender ratio of 1:1. Advertisers offer gender-specific products with 50% of them specializing in male-oriented brands and the remaining half in female-oriented brands. A correct match, when an advertiser is paired with their intended target user, yields a click rate of $(2-p)\cdot s_j$, while an incorrect match results in a click rate of $p\cdot s_j$, where p lies within the range of [0, 1] and s_j is the average click rate of j across genders. When $p=0$, advertisers receive clicks solely from the correct gender, while $p=1$ indicates no gender difference in click rates. Here we allow the search engine to decide the weight assigned to the correct gender match. Gender-specific quality scores are calculated by multiplying the average click rate s_j with $(2-\theta)$ for the correct match and with θ for the incorrect match. The targeting factor θ lies in $[0,1]$, where $\theta=0$ leads to the highest level of accuracy/personalization and $\theta=1$ completely ignores user attributes. Now the quality score under user targeting is as follows:

$$q_j(\theta)=\begin{cases} s_j\times(2-\theta) & \text{for a correct match,}\\ s_j\times(\theta) & \text{for an incorrect match.}\end{cases}$$

Note that, if $\theta=p$, the quality score uses the same weight as the actual click rate, which we define as the efficient level. We denote the profit generated from correct gender matches as Π^c and from incorrect gender matches as Π^{ic}, defined as

$$\Pi^c(b_j;v_j,s_j,q_j)=\sum_{k=1}^{K}P(b_{j,w}=b_w^{[k]})(c_k\times(2-p)\times s_j)\left(v_j-\mathbb{E}[p_{j,k}|b_{j,w}=b_w^{[k]}]\right),$$

$$\Pi^{ic}(b_j;v_j,s_j,q_j)=\sum_{k=1}^{K}P(b_{j,w}=b_w^{[k]})(c_k\times p\times s_j)\left(v_j-\mathbb{E}[p_{j,k}|b_{j,w}=b_w^{[k]}]\right).$$

Thus the profit and consequently the bid depend on gender. Note that we look at the case where the search engine can modify the assigned quality score of the advertiser based on the advertiser-user match type. For varying values of θ, we use the following algorithm to determine counterfactual outcomes:

1. Use MLE to estimate $(\mu_v,\mu_s,\sigma_v^2,\sigma_s^2)$ of the log-normal distributions of the valuation and quality score.
2. Solve for the equilibrium bid function using the weighted value distribution and position-specific click rates $\{\hat{c}_1,\cdots,\hat{c}_K\}$.

3. Draw N number of advertisers with the value (v_j and average click rate (s_j) from log-normal distributions. Each advertiser is assigned a gender indicator g_j that takes a value of 1 if the user is right-gender. The gender indicator follows a Bernoulli distribution with $P(g_j = 1) = 0.5$
4. Compute the weighted bids and determine the auction outcome. Calculate the auction revenue, advertiser profit, and ad quality for the top 7 ad positions.
5. Repeat Steps 3 to 4 for 2000 replications and compute the average revenue, profit, per-click price, and ad quality.

The actual gender penalty on click rate, i.e., p, is set to be 0.25. The quality score distribution in this case becomes a mixture of score distributions for male-oriented and female-oriented advertisers.[29] Thus the weighted value distribution in Step 2 now varies with the value of θ.

The findings, illustrated in Figure 9, offer valuable insights into the impact of user targeting. Consistent with the theoretical expectations, the efficient allocation, where the sum of advertiser profit and platform revenue is maximized, is achieved at $\theta = p = 0.25$. However, for revenue maximization, the search engine goes for neither the highest (i.e., $\theta = 0$) nor the efficient level. The revenue-maximizing level (θ^*) is higher than the efficient level p, with the level varying across product categories from 0.35 (for *car insurance*) to 0.7 (for *coins*). This implies that targeting with a lower level of accuracy/personalization is revenue-maximizing, with gains ranging from 0.4% (*car insurance*) to 3.0% (*coins*). Advertiser profit is also maximized with an inefficient allocation. The profit-maximizing θ is below p, suggesting that increasing personalization enhances advertiser profit. Per-click prices rise with θ, while average ad quality declines, reflecting the trade-off between quality accuracy and price. The average ad quality is maximized when we apply the highest penalty level for mismatch as expected.

[29] Given $s_j \sim L(\mu_s, \sigma_s^2)$, the advertisers draw quality scores from $F_s^1 \equiv L(\mu_s + log\theta, \sigma_s^2)$ when they are the wrong gender. Likewise, if they are the right gender, quality scores are drawn from $F_s^2 \equiv L(\mu_s + log(2-\theta), \sigma_s^2)$. Then the quality score distribution unconditional on gender is the mixture of F_s^1 and F_s^2 with equal weight.

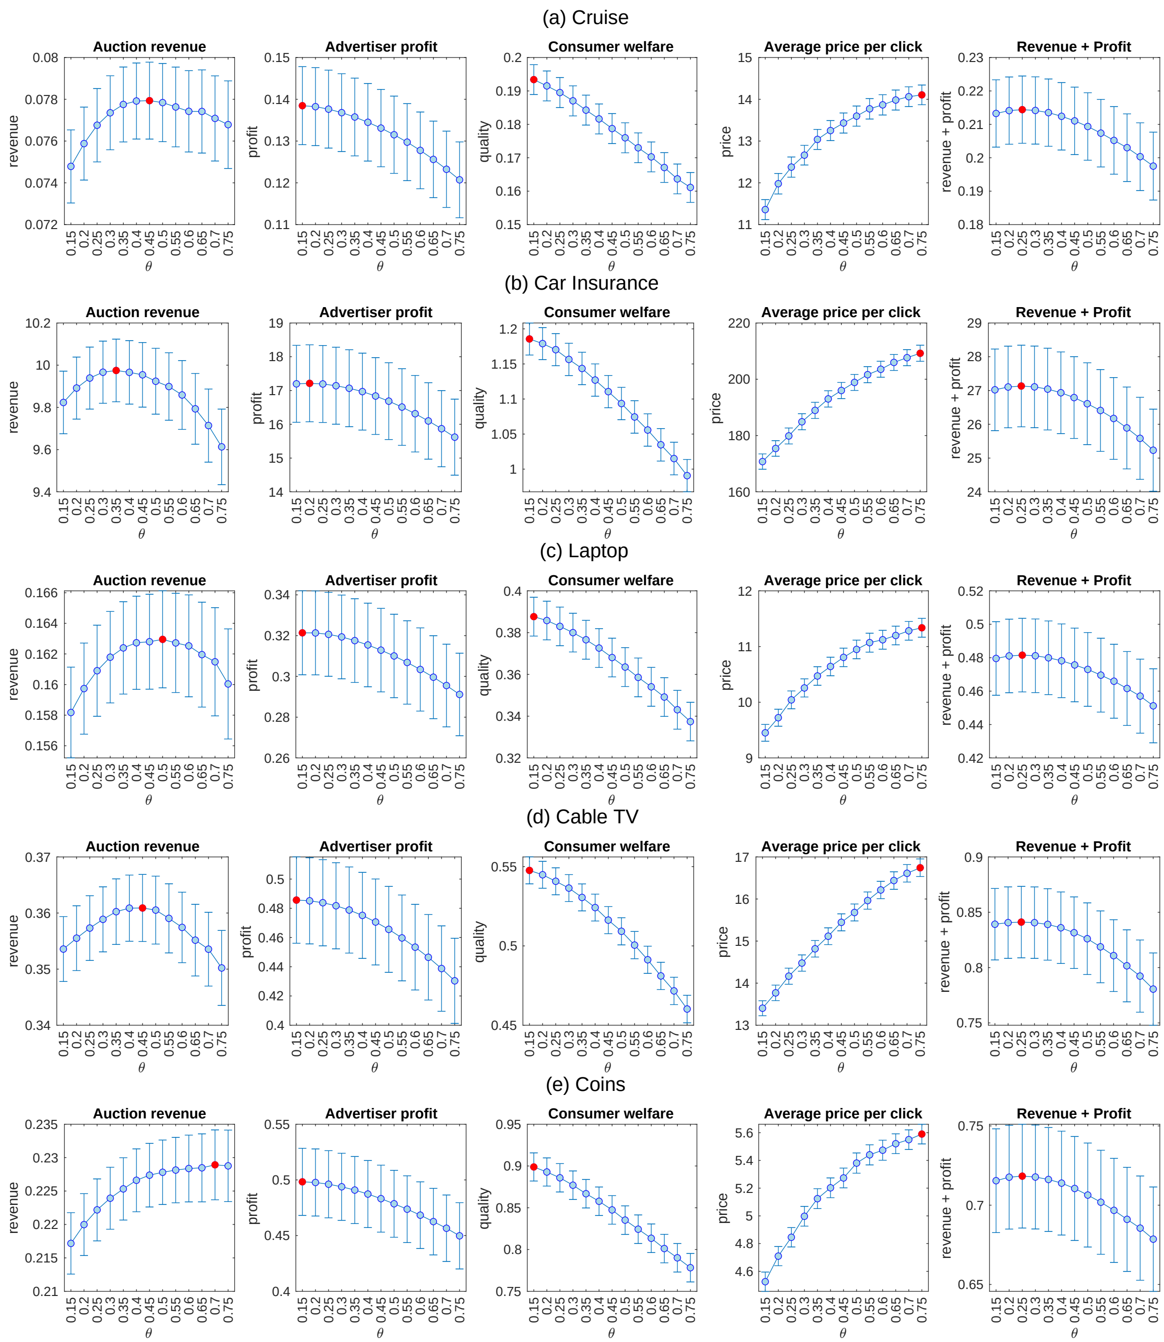

FIGURE 9.—Counterfactual simulations for gender-specific keywords with user targeting. The dots are averages across 2000 simulated auctions. The red dots are maximums across values of θ. The ranges are 95% confidence intervals.

TABLE VII

CHANGES IN AUCTION OUTCOMES WITH THE REVENUE-MAXIMIZING TARGETING LEVEL (θ^*)

	Cruise	Car Ins.	Laptop	Cable	Coins
θ^*	0.45	0.35	0.5	0.45	0.7
Revenue	1.5%	0.4%	1.3%	1.0%	3.0%
Profit	-3.3%	-0.7%	-3.3%	-2.7%	-8.0%
Welfare	-5.7%	-2.3%	-5.1%	-4.5%	-10.8%

Note: The table summarizes the impact of user targeting with θ^* on auction outcomes compared to the case with $\theta = p$.

Our results suggest that platforms may lean towards inefficient, and less accurate ad-user matching. While this approach can optimize revenue, it comes at the expense of advertiser profit and consumer welfare as shown in Table VII. In a competitive landscape, platforms may not necessarily prioritize inefficient matching in the long term, as it risks displeasing their customer base. Beyond its implications for revenue and welfare, our analysis sheds light on the delicate balance between privacy concerns and efficient user-ad matching, offering valuable insights for policymakers and regulators.[30] Despite search engines advocating for user targeting to enhance ad quality,[31] our findings suggest that platforms have incentives to prioritize inefficient matches with lower ad quality when targeting is permitted. Hence, the purported benefits of targeting in ad quality may not be as substantial.

Recently, Larsen and Proserpio (2023) demonstrated that introducing a large language model (LLM) improved revenue and efficiency in search auctions by better interpreting search queries and identifying more relevant advertisers (query-based targeting). Our findings align with theirs, although we focus solely on user-based targeting, which leads to a thinner market. Investigating revenue-optimal query-based targeting would be a fruitful avenue for future research.

[30]The use of user information has raised privacy concerns. See https://www.ftc.gov/news-events/news/press-releases/2012/03/ftc-issues-final-commission-report-protecting-consumer-privacy.

[31]See https://safety.google/privacy/ads-and-data/.

8. CONCLUSION

The framework proposed in this paper provides a valuable tool for estimating the value distribution in GSP^{ω} auctions under more realistic assumptions than those used in prior literature. The estimated value distribution can be used for counterfactual analyses, such as examining the impact of changes in auction design on advertisers, determining revenue-maximizing levels of score squashing and user targeting, and calculating the optimal reserve price. These counterfactual experiments offer a deeper understanding of how advertisers respond to changes in market mechanisms. Unlike the consumer side, where search engines can conduct randomized controlled trials to examine the effects of potential changes, it is challenging to implement such experiments on the advertiser side. This difficulty arises because advertisers' responses to changes in market factors, such as pricing mechanisms, are typically slower, and frequent changes in the market environment may prompt advertisers to leave the platform.

Our model assumes symmetric bidders. In practice, however, advertisers also face budget constraints, and budget smoothing can lead them to shade bids further. Bidders can also exhibit additional dimensions of heterogeneity, such as advertiser prominence. Future research could extend the model to include bidder asymmetry and additional sources of uncertainty, such as entry decisions and position-specific valuations. In practice, the number of competitors changes throughout the day. Incorporating advertisers' entry and exit decisions would be a promising extension, especially when data recording these dynamics are available. Additionally, relaxing the separability assumption on click probability could allow for differential position effects across advertisers. This approach would result in advertiser-specific valuations differing across ad positions, transforming the auction dynamics into selling multiple unordered heterogeneous objects within a single auction. This would necessitate a significantly more intricate theoretical framework. Lastly, future research could extend the model to a dynamic setting, thereby accommodating a multi-period game. It would be interesting to investigate whether the Bayes-Nash Equilibrium of our static model serves as a good approximation to the dynamic model in this market, as speculated by Gomes and Sweeney (2014).

REFERENCES

AMALDOSS, WILFRED, PREYAS S DESAI, AND WOOCHOEL SHIN (2015): "Keyword search advertising and first-page bid estimates: A strategic analysis," *Management Science*, 61 (3), 507–519. [5]

ASKER, JOHN AND ESTELLE CANTILLON (2008): "Properties of scoring auctions," *The RAND Journal of Economics*, 39 (1), 69–85. [2]

ATHEY, SUSAN AND GLENN ELLISON (2011): "Position auctions with consumer search," *The Quarterly Journal of Economics*, 126 (3), 1213–1270. [4]

ATHEY, SUSAN AND DENIS NEKIPELOV (2010): "A structural model of sponsored search advertising auctions," in *Sixth ad auctions workshop*, vol. 15. [4, 1]

BAE, JINSOO AND JOHN H KAGEL (2019): "An experimental study of the generalized second price auction," *International Journal of Industrial Organization*, 63, 44–68. [4]

BERGEMANN, DIRK AND ALESSANDRO BONATTI (2011): "Targeting in advertising markets: implications for offline versus online media," *The RAND Journal of Economics*, 42 (3), 417–443. [5]

BLAKE, THOMAS, CHRIS NOSKO, AND STEVEN TADELIS (2015): "Consumer heterogeneity and paid search effectiveness: A large-scale field experiment," *Econometrica*, 83 (1), 155–174. [5]

BÖRGERS, TILMAN, INGEMAR COX, MARTIN PESENDORFER, AND VACLAV PETRICEK (2013): "Equilibrium bids in sponsored search auctions: Theory and evidence," *American economic Journal: microeconomics*, 5 (4), 163–87. [4]

CARAGIANNIS, IOANNIS, CHRISTOS KAKLAMANIS, PANAGIOTIS KANELLOPOULOS, MARIA KYROPOULOU, BRENDAN LUCIER, RENATO PAES LEME, AND ÉVA TARDOS (2015): "Bounding the inefficiency of outcomes in generalized second price auctions," *Journal of Economic Theory*, 156, 343–388. [4]

CHARLES, DENIS, NIKHIL R DEVANUR, AND BALASUBRAMANIAN SIVAN (2016): "Multi-score position auctions," in *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*, 417–425. [3, 5]

CHE, YEON-KOO (1993): "Design competition through multidimensional auctions," *The RAND Journal of Economics*, 668–680. [2]

DEBNATH, LOKENATH AND PIOTR MIKUSINSKI (2005): *Introduction to Hilbert spaces with applications*, Academic press. [6]

DECAROLIS, FRANCESCO AND GABRIELE ROVIGATTI (2021): "From mad men to maths men: Concentration and buyer power in online advertising," *American Economic Review*, 111 (10), 3299–3327. [5]

DVORETZKY, ARYEH, JACK KIEFER, AND JACOB WOLFOWITZ (1956): "Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator," *The Annals of Mathematical Statistics*, 642–669. [8]

EDELMAN, BENJAMIN, MICHAEL OSTROVSKY, AND MICHAEL SCHWARZ (2007): "Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords," *American economic review*, 97 (1), 242–259. [4]

FU, HU, PATRICK JORDAN, MOHAMMAD MAHDIAN, URI NADAV, INBAL TALGAM-COHEN, AND SERGEI VASSILVITSKII (2012): "Ad auctions with data," in *International Symposium on Algorithmic Game Theory*, Springer, 168–179. [5]

GHOSE, ANINDYA AND SHA YANG (2009): "An empirical analysis of search engine advertising: Sponsored search in electronic markets," *Management science*, 55 (10), 1605–1622. [5]

GOMES, RENATO AND KANE SWEENEY (2014): "Bayes–nash equilibria of the generalized second-price auction," *Games and Economic Behavior*, 86, 421–437. [4, 39, 6]

GUERRE, EMMANUEL, ISABELLE PERRIGNE, AND QUANG VUONG (2000): "Optimal nonparametric estimation of first-price auctions," *Econometrica*, 68 (3), 525–574. [16, 20]

HAILE, PHILIP A, HAN HONG, AND MATTHEW SHUM (2003): "Nonparametric tests for common values at first-price sealed-bid auctions," Tech. rep., National Bureau of Economic Research. [16]

HAILE, PHILIP A AND ELIE TAMER (2003): "Inference with an incomplete model of English auctions," *Journal of Political Economy*, 111 (1), 1–51. [2]

HAUSER, JOHN R, GLEN L URBAN, GUILHERME LIBERALI, AND MICHAEL BRAUN (2009): "Website morphing," *Marketing Science*, 28 (2), 202–223. [5]

HSIEH, YU-WEI, MATTHEW SHUM, AND SHA YANG (2015): "To score or not to score? Estimates of a sponsored search auction model," *Estimates of a Sponsored Search Auction Model (February 1, 2015). USC-INET Research Paper*, (15-09). [4]

JERATH, KINSHUK, LIYE MA, YOUNG-HOON PARK, AND KANNAN SRINIVASAN (2011): "A "position paradox" in sponsored search auctions," *Marketing Science*, 30 (4), 612–627. [5]

JEZIORSKI, PRZEMYSŁAW AND SRIDHAR MOORTHY (2018): "Advertiser prominence effects in search advertising," *Management Science*, 64 (3), 1365–1383. [21]

JOSHI, AMRUTA, ABRAHAM BAGHERJEIRAN, AND ADWAIT RATNAPARKHI (2011): "User demographic and behavioral targeting for content match advertising," in *Proceedings of the Fifth International Workshop on Data Mining and Audience Intelligence for Advertising (ADKDD 2011)*, Citeseer, 53–60. [5]

KRISHNA, VIJAY (2009): *Auction theory*, Academic press. [11]

LAHAIE, SÉBASTIEN AND R PRESTON MCAFEE (2011): "Efficient ranking in sponsored search," in *Internet and Network Economics: 7th International Workshop, WINE 2011, Singapore, December 11-14, 2011. Proceedings 7*, Springer, 254–265. [5]

LAHAIE, SÉBASTIEN AND DAVID M PENNOCK (2007): "Revenue analysis of a family of ranking rules for keyword auctions," in *Proceedings of the 8th ACM Conference on Electronic Commerce*, 50–56. [3, 5, 28]

LARSEN, POET AND DAVIDE PROSERPIO (2023): "The Impact of Large Language Models on Search Advertising: Evidence from Google's BERT," *USC Marshall School of Business Research Paper Sponsored by iORB*. [38]

LEVIN, JONATHAN AND PAUL MILGROM (2010): "Online advertising: Heterogeneity and conflation in market design," *American Economic Review*, 100 (2), 603–607. [5]

MASSART, PASCAL (1990): "The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality," *The annals of Probability*, 1269–1283. [8]

MILGROM, PAUL AND ILYA SEGAL (2002): "Envelope theorems for arbitrary choice sets," *Econometrica*, 70 (2), 583–601. [4, 6]

MILGROM, PAUL ROBERT (2004): *Putting auction theory to work*, Cambridge University Press. [6]

MOHRI, MEHRYAR AND ANDRÉS MUNOZ MEDINA (2015): "Non-parametric revenue optimization for generalized second price auctions," *arXiv preprint arXiv:1506.02719*. [4, 5]

OSTROVSKY, MICHAEL AND MICHAEL SCHWARZ (2011): "Reserve prices in internet advertising auctions: A field experiment," in *Proceedings of the 12th ACM conference on Electronic commerce*, 59–60. [4]

RAFIEIAN, OMID AND HEMA YOGANARASIMHAN (2021): "Targeting and privacy in mobile advertising," *Marketing Science*, 40 (2), 193–218. [5]

RAMABOA, KUTLWANO KKM AND PETER FISH (2018): "Keyword length and matching options as indicators of search intent in sponsored search," *Information Processing & Management*, 54 (2), 175–183. [19]

SHIN, JIWOONG AND WOOCHOEL SHIN (2023): "A theory of irrelevant advertising: An agency-induced targeting inefficiency," *Management Science*, 69 (8), 4481–4497. [5]

SUN, YANG, YUNHONG ZHOU, AND XIAOTIE DENG (2014): "Optimal reserve prices in weighted GSP auctions," *Electronic Commerce Research and Applications*, 13 (3), 178–187. [4]

SVITAK, JAN, JAN TICHEM, AND STEFAN HAASBEEK (2021): "Price effects of search advertising restrictions," *International Journal of Industrial Organization*, 77, 102736. [5]

THOMPSON, DAVID RM AND KEVIN LEYTON-BROWN (2013): "Revenue optimization in the generalized second-price auction," in *Proceedings of the fourteenth ACM conference on Electronic commerce*, 837–852. [3, 5]

VARIAN, HAL R (2007): "Position auctions," *international Journal of industrial Organization*, 25 (6), 1163–1178. [4, 5, 9]

YAN, HAOMIN (2019): "Position Auctions with Interdependent Values," Tech. rep. [4]

YAO, SONG AND CARL F MELA (2011): "A dynamic model of sponsored search advertising," *Marketing Science*, 30 (3), 447–468. [4, 5]

APPENDIX A: EXTENSIONS

The proposed method possesses a key advantage in its versatility in incorporating additional features present in online search auctions. This study specifically examines two noteworthy cases. In the first case, we only observe winning bids, while the latter case involves the presence of a fixed reserve price that is observable.

A.1. *Case 1: Limited bid data*

In sponsored search auctions, a full set of bids is sometimes not available to the researcher. For instance, Athey and Nekipelov (2010) (AN10 hereafter) discusses a potential missing data problem, as they only observe a subset of weighted bids in each auction.[32] This can potentially lead to bias in the estimation of the empirical weighted bid distribution. In this section, we provide one possible way to resolve this bias. The advantage of search bid data is that we observe the weighted bids as well as their order statistic. We can infer the order of an observed weighted bid by the winning position. Under this setting, we need very limited bid information for our proposed method, as long as we know the total number of advertisers. Consider the case where we only observe the weighted bid of the advertiser who won the first position. Recall, in terms of the order statistic notation, we observe $b_w^k \sim G_{w,k:N}(\cdot)$, where $G_{w,k:N}(\cdot)$ is the distribution of the kth highest order statistic.

To further exploit the order statistic, we use the following equivalence for the order statistic

$$G_{w,k:N}(b_w) = \int_0^{G_w(b)} \frac{N!}{(N-k)!(k-1)!} t^{N-k}(1-t)^{k-1} dt$$
$$= \Gamma(G_w(b_w); N-k+1, k). \tag{11}$$

[32] AN10 observes only the winning bids and the highest (in terms of weighted bid) non-winning bid. In our data, we observe a large set of non-winning bids (beyond the first search result page) so this problem is negligible.

We can use the above equation to see the relationship between the order statistic distribution and the primitive distribution given an incomplete beta function, represented here as $\Gamma(z;a,b)$ with $a = N - k + 1$ and $b = k$.

Next, we can use Equation (11) and a well-known property of order statistic i.i.d variables, which states that given a known number (N) of i.i.d. draws of weighted bids from $G_w(\cdot)$, the order statistic distribution can be used to identify the parent distribution $G_w(\cdot)$ using the following equation:

$$G_w(b_w) = \Gamma^{-1}\Big(G_{w,k:n}(b_w); N - k + 1, k\Big). \tag{12}$$

where Γ^{-1} is the inverse of the incomplete beta function defined in Equation (11).[33] Once we have estimated the weighted bid distribution, all the other steps follow the same as our main results.

A.2. *Case 2: Reserve Price*

We explore the case in which the data have a known fixed reserve price, denoted by r. Assume that the search engine sets a fixed reserve price observed by everyone. In such a case, we can still derive the value given observable quantities. However, the solution for the value is now obtained by modifying the equation in Theorem 2. The following theorem proves the identification of the valuation in the presence of a reserve price.

THEOREM 4: *Under Assumptions 1–2, the advertiser value, v, is identified by:*

$$v = b + \Phi(G_w, b, q|\mathcal{C}, K, N, r) \tag{13}$$

where

$\Phi(G_w, b, q|\mathcal{C}, K, N, r) =$

[33] For further details of using this property in an auction context, refer to Lemma 2 of Haile and Tamer (2003).

$$\frac{\sum_{k=1}^{K} c_k \binom{N-1}{k-1}(k-1)(1-G_w(b_w))^{k-2}\left(\int_r^{b_w} G_w(u)^{N-k}du - rG_w(r)^{N-k}\right)}{q_j \sum_{k=1}^{K} c_k \binom{N-1}{k-1} G_w(b_w)^{N-k-1}(1-G_w(b_w))^{k-2}\left[(N-k)(1-G_w(b_w)) - (k-1)G_w(b_w)\right]}$$

given the quality score (q)*, the equilibrium bid* (b)*, the distribution function of equilibrium weighted bids* (G_w)*, the number of available ad positions* (K)*, click rates across ad positions* $(\mathcal{C})$*, the number of advertisers* (N)*, and the fixed reserve price* (r)*.*

We have now a lower bound on the price which is the reserve price. More details are provided in the proof in appendix B. Note that bids below the reserve price are not observed. In such a case, one can use the strategy proposed in the previous subsection (A.1) to derive the weighted bid distribution.

APPENDIX B: THEORETICAL PROOFS

Proof of Theorem 1: We begin by noting that in (4), the advertiser's profit maximization problem can be redefined as the advertiser choosing a weighted bid that maximizes the payoff, instead of the bid. Thus, we can express the equilibrium outcome in terms of weighted bids by rewriting (4) as:

$$b_w(v_j, s_j) = \underset{\hat{b}_w}{\operatorname{argmax}}\ \Pi(\hat{b}_w | v_j, s_j). \tag{14}$$

Consider any symmetric equilibrium $\{b^*_{1,w}, \cdots, b^*_{n,w}\}$ of a GSP^w auction. We show that the equilibrium strategy only depends on the weighted value ω_j for any arbitrary advertiser j. The equilibrium strategy of advertiser j in the GSP^w is given by

$$\begin{aligned} b^*_{j,w} &= \underset{\hat{b}_w}{\operatorname{argmax}} \sum_{k=1}^{K} s_j c_k \left[v_j - \frac{\mathbb{E}\left(b^{*[k]}_{-j,w} \middle| b^{*[k]}_{-j,w} \leq \hat{b}_w \leq b^{*,[k-1]}_{-j,w} \right)}{q_j} \right] \times P(b^{*[k]}_{-j,w} \leq \hat{b}_w \leq b^{*,[k-1]}_{-j,w}), \\ &= \underset{\hat{b}_w}{\operatorname{argmax}} \sum_{k=1}^{K} \frac{s_j}{q_j} c_k \left[\omega_j - \mathbb{E}\left(b^{*[k]}_{-j,w} \middle| b^{*[k]}_{-j,w} \leq \hat{b}_w \leq b^{*,[k-1]}_{-j,w} \right) \right] \times P(b^{*[k]}_{-j,w} \leq \hat{b}_w \leq b^{*,[k-1]}_{-j,w}), \end{aligned}$$

where $\omega_j = v_j \times q_j$. Here the advertiser's auction outcome, i.e., the price paid and the allocation probability is not affected by the term $\frac{s_j}{q_j}$. Thus, the above maximization problem is the same as the following:

$$b^*_{j,w} = \underset{\hat{b}_w}{\operatorname{argmax}} \sum_{k=1}^{K} c_k \left[\omega_j - \mathbb{E}\left(b^{*[k]}_{-j,w} \middle| b^{*[k]}_{-j,w} \leq \hat{b}_w \leq b^{*,[k-1]}_{-j,w} \right) \right] \cdot P(b^{*[k]}_{-j,w} \leq \hat{b}_w \leq b^{*,[k-1]}_{-j,w}). \quad (15)$$

This implies that the equilibrium strategy only depends on the weighted value ω_j for advertiser j and therefore

$$b_w(v_j, s_j) = b_w(\omega_j).$$

Q.E.D.

Proof of Lemma 1: Suppose an equilibrium exists in this auction. Under Assumption 2, in the equilibrium allocation, an advertiser with weighted value ω wins the k-th ad position with probability:

$$\zeta_k(\omega) \equiv P(w^{[k+1]} \leq w \leq w^{[k-1]}) = \binom{N-1}{k-1} (1 - F_w(\omega))^{k-1} F_w^{N-k}(\omega). \quad (16)$$

Using the Revelation Principle, an advertiser with weighted value ω has payoff that satisfies

$$\omega = \underset{\hat{\omega}}{\operatorname{argmax}} \sum_{k=1}^{K} \zeta_k(\hat{\omega}) c_k \left[\omega - \mathbb{E}\left(b_w(\omega^{[k+1]}) \middle| \omega^{[k+1]} \leq \hat{\omega} \leq \omega^{[k-1]} \right) \right] \quad (17)$$

Applying the envelop theorem (see Milgrom and Segal (2002)) in the payoff function in Equation (17), we have:

$$\frac{d}{d\omega} \Pi(\omega) = \sum_{k=1}^{K} c_k \zeta_k(\omega) \quad (18)$$

and also using the Fundamental Theorem of Calculus, we get

$$\Pi(\omega) = \Pi(\underline{\omega}) + \sum_{k=1}^{K} c_k \int_0^{\omega} \zeta_k(x) dx. \quad (19)$$

As a bidder with type $\underline{\omega}$ never has a non-zero payoff – $\Pi(\underline{\omega}) = 0$, we have

$$\Pi(\omega) = \sum_{k=1}^{K} c_k \int_0^{\omega} \zeta_k(x) dx. \quad (20)$$

Furthermore, using Equations (17) and (20), we obtain

$$\sum_{k=1}^{K} c_k \int_0^\omega \zeta_k(x)dx = \sum_{k=1}^{K} c_k \zeta_k(\omega) \left[\omega - \mathbb{E}\left(b_w(\omega^{[k+1]}) \middle| \omega^{[k+1]} \leq \omega \leq \omega^{[k-1]} \right) \right]$$

which can be rearranged as

$$\sum_{k=1}^{K} c_k \left[\zeta_k(\omega)\omega - \int_0^\omega \zeta_k(x)dx \right] = \sum_{k=1}^{K} c_k \zeta_k(\omega) \mathbb{E}\left(b_w(\omega^{[k+1]}) \middle| \omega^{[k+1]} \leq \omega \leq \omega^{[k-1]} \right).$$

Using integration by parts on the left-hand side we derive

$$\sum_{k=1}^{K} c_k \int_0^\omega x \frac{d\zeta_k(x)}{dx} dx = \sum_{k=1}^{K} c_k \zeta_k(\omega) \mathbb{E}\left(b_w(\omega^{[k+1]}) \middle| \omega^{[k+1]} \leq \omega \leq \omega^{[k-1]} \right).$$

Opening up the expectation on the right-hand side,

$$\sum_{k=1}^{K} c_k \int_0^\omega x \frac{d\zeta_k(x)}{dx} dx = \sum_{k=1}^{K} c_k \zeta_k(\omega) \frac{\int_0^\omega b_w(x)(N-k)F_w(x)^{N-k-1} f_w(x)dx}{F_w(\omega)^{N-k}}$$

and substituting $\zeta_k(\omega)$ using Equation (16) in the right hand side yields,

$$\sum_{k=1}^{K} c_k \int_0^\omega x \frac{d\zeta_k(x)}{dx} dx =$$
$$\sum_{k=1}^{K} c_k \binom{N-1}{k-1} \int_0^\omega b_w(x)(N-k)F_w(x)^{N-k-1} f_w(x)(1-F_w(\omega))^{k-1} dx.$$

Differentiating both sides we get:

$$\sum_{k=1}^{K} c_k \omega \frac{d\zeta_k(\omega)}{d\omega} = \sum_{k=1}^{K} c_k \binom{N-1}{k-1} b_w(\omega)(N-k)F_w(\omega)^{N-k-1} f_w(\omega)(1-F_w(\omega))^{k-1}$$
$$- \sum_{k=1}^{K} c_k \binom{N-1}{k-1} (k-1)(1-F_w(\omega))^{k-2} f_w(\omega) \int_0^\omega b_w(x)(N-k)F_w(x)^{N-k-1} f_w(x)dx. \tag{21}$$

Gomes and Sweeney (2014) show that the above equation can be expressed as the following Volterra equation of the second kind:

$$\omega - b_w(\omega) = \Gamma(\omega) + \int_0^{\omega} M_1(\omega, x)(x - b_w(x))dx. \tag{22}$$

Debnath and Mikusinski (2005) proved that (22) has a unique solution given by (5) given that $\Gamma(\omega)$ and $M_1(\omega, t)$ are square-integrable.

Now we show the uniqueness and existence of the equilibrium weighted bid. Similar to the equivalence shown in (19), by using Integral-form Envelope Theorem (Milgrom and Segal, 2002), we can show that if the weighted bid is as defined in (5), the profit function satisfies the following:

$$\Pi(b_w(\omega), \omega) = \Pi(b_w(\underline{\omega}), \underline{\omega}) + \sum_{k=1}^{K} c_k \int_0^{\omega} \zeta_k(x)dx. \tag{23}$$

Then using the constraint simplification theorem (Milgrom, 2004), we can show that $b_w(\omega)$ is a selection from $B_w^*(\omega) = arg\max_{\hat{b}_w} \Pi(\hat{b}_w(\omega), \omega)$ if and only if the envelop formula (23) holds and the weighted bid defined in (5) is strictly monotonic. This implies that under Assumption 2, for each bidder j, $b_w(\omega)$ is the best response to all possible bids of competitors and thus constitutes a unique equilibrium weighted bid which exists if and only if Assumption 2 is satisfied.

Furthermore, the efficient allocation maximizes the overall auction surplus $\sum_{k=1}^{K} c_k s^{[k]} v^{[k]}$ where $s^{[k]}$ and $v^{[k]}$ are the advertiser-specific click rate and the valuation of the winner of the k^{th} position. By Theorem 1 and Assumption 2, the allocation in the equilibrium will be monotonic in the weighted value, which leads to efficiency if the weighted value is equal to the product of the value and the advertiser-specific click rate and $c_k \geq c_{k+1}$ for all $k \leq K - 1$. *Q.E.D.*

Proof of Theorem 2: We begin by using (21) in the proof of Lemma 1. The integral part in Equation (21) can be re-written using integration by parts:

$$\int_0^{\omega} b_w(x)(N-k)F_w(x)^{N-k-1} f_w(x)dx = b_w(\omega)F_w(\omega)^{N-k} - \int_0^{\omega} b_w'(x)F_w(x)^{N-k}dx \tag{24}$$

Under Assumption 2, the integral in the last term of Equation (24) can be replaced by

$$\int_0^{\omega} b'_w(x) G_w(b_w(x))^{N-k} dx$$

where G_w is the distribution function of weighted bids. Integration by substitution yields:

$$\int_0^{\omega} b'_w(x) G_w(b_w(x))^{N-k} dx = \int_0^{b_w(\omega)} G_w(u)^{N-k} du$$

and therefore plugging the above expression into Equation (21) using Equation (24),

$$\begin{aligned}
\sum_{k=1}^{K} c_k \omega \frac{d(\zeta_k(\omega))}{d\omega} =& \sum_{k=1}^{K} c_k \binom{N-1}{k-1} b_w(\omega)(N-k) F_w(\omega)^{N-k-1} f_w(\omega)(1-F_w(\omega))^{k-1} \\
&- \sum_{k=1}^{K} c_k \binom{N-1}{k-1} (k-1)(1-F_w(\omega))^{k-2} f_w(\omega) b_w(\omega) F_w(\omega)^{N-k} \\
&+ \sum_{k=1}^{K} c_k \binom{N-1}{k-1} (k-1)(1-F_w(\omega))^{k-2} f_w(\omega) \int_0^{b_w(\omega)} G_w(u)^{N-k} du.
\end{aligned}$$

Now the last step is to open up $\frac{d(\zeta_k(\omega))}{d\omega}$ and to rearrange.

$$\omega = b_w(\omega) + \frac{\sum_{k=1}^{K} c_k \binom{N-1}{k-1} (k-1)(1-F_w(\omega))^{k-2} \int_0^{b_w(\omega)} G_w(u)^{N-k} du}{\sum_{k=1}^{K} c_k \binom{N-1}{k-1} F_w(\omega)^{N-k-1} (1-F_w(\omega))^{k-2} \Big[(N-k)(1-F_w(\omega)) - (k-1) F_w(\omega) \Big]}$$

Divide both sides by quality score q:

$$v = b + \frac{\sum_{k=1}^{K} c_k \binom{N-1}{k-1} (k-1)(1-F_w(\omega))^{k-2} \int_0^{b_w(\omega)} G_w(u)^{N-k} du}{q \sum_{k=1}^{K} c_k \binom{N-1}{k-1} F_w(\omega)^{N-k-1} (1-F_w(\omega))^{k-2} \Big[(N-k)(1-F_w(\omega)) - (k-1) F_w(\omega) \Big]}$$

The only unobservable object in the solution is $F_w(x)$. By replacing it with $G_w(b_w(x))$, the desired expression in the theorem is obtained. *Q.E.D.*

Proof of Theorem 3: By the Glivenko-Cantelli Theorem, the empirical distribution $\hat{G}_w(b_w)$ uniformly converges to $G_w(b_w)$. Furthermore, the Dvoretzky–Kiefer–Wolfowitz inequality (Dvoretzky et al., 1956, Massart, 1990) provides that for every $\epsilon > 0$,

$$P\left(\sup_{b_w} |\hat{G}_w(b_w) - G_w(b_w)| > \epsilon\right) \leq 2e^{-2n\epsilon^2}.$$

By setting $\epsilon = k/\sqrt{n}$ for a constant $k > 0$, the inequality becomes

$$P\left(\sup_{b_w} |\hat{G}_w(b_w) - G_w(b_w)|\sqrt{n} > k\right) \leq 2e^{-2k^2}.$$

For any small $\varepsilon > 0$, there exists a constant $k = \sqrt{-\frac{\log(\varepsilon/2)}{2}}$ such that

$$P\left(\sup_{b_w} |\hat{G}_w(b_w) - G_w(b_w)|\sqrt{n} > k\right) \leq \varepsilon,$$

which implies that $\sup_{b_w} |\hat{G}_w(b_w) - G_w(b_w)|$ converges in probability to 0 and is $O_p(1/\sqrt{n})$. Therefore, $\hat{G}_w(b_w)$ uniformly converges to $G_w(b_w)$ at the $\sqrt{n}$-rate.

Now we show that $\Phi(\hat{G}_w, b, q|\mathcal{C}, K, N)$ in the valuation estimator uniformly converges to $\Phi(G_w, b, q|\mathcal{C}, K, N)$ at the $\sqrt{n}$-rate as $n \to \infty$. Recall that

$$\Phi(\hat{G}_w, b, q|\mathcal{C}, K, N) = \frac{\sum_{k=1}^{K} c_k \binom{N-1}{k-1}(k-1)(1-\hat{G}_w(b_w))^{k-2}\int_0^{b_w} \hat{G}_w(u)^{N-k}du}{q\sum_{k=1}^{K} c_k \binom{N-1}{k-1}\hat{G}_w(b_w)^{N-k-1}(1-\hat{G}_w(b_w))^{k-2}\left[(N-k)(1-\hat{G}_w(b_w)) - (k-1)\hat{G}_w(b_w)\right]}.$$

Notice that addition, subtraction, multiplication, and division are rate-preserving. Thus we only need to show that a power function of $\hat{G}_w(b_w)$ and its integral uniformly converge at the $\sqrt{n}$-rate. First, consider $\hat{G}_w(b_w)^l$ for any integer l. By the mean value theorem, there exists some $c(b_w)$ between $\hat{G}_w(b_w)$ and $G_w(b_w)$ such that

$$|\hat{G}_w(b_w)^l - G_w(b_w)^l| = |lc(b_w)^{l-1}(\hat{G}_w(b_w) - G_w(b_w))| \leq |lc(b_w)^{l-1}||(\hat{G}_w(b_w) - G_w(b_w))|.$$

As $|lc(b_w)^{l-1}|$ is bounded and $\sup_{b_w} |(\hat{G}_w(b_w) - G_w(b_w))| = O_p(1/\sqrt{n})$, $\sup_{b_w} |\hat{G}_w(b_w)^l - G_w(b_w)^l|$ is also $O_p(1/\sqrt{n})$, meaning that $\hat{G}_w(b_w)^l$ uniformly converges to $\hat{G}_w(b_w)^l$ at the

$\sqrt{n}$-rate. Now consider $\int_0^{b_w} \hat{G}_w(u)^l du$ for any integer l. We can show that

$$\left|\int_0^{b_w} \left(\hat{G}_w(u)^l - G_w(u)^l\right) du\right| \leq \int_0^{b_w} \left|\hat{G}_w(u)^l - G_w(u)^l\right| du \leq b_w \sup_u \left|\hat{G}_w(u)^l - G_w(u)^l\right|.$$

Given that b_w is bounded and that the uniform convergence of $\hat{G}_w(u)^l$,

$$\sup_{b_w} \left|\int_0^{b_w} \left(\hat{G}_w(u)^l - G_w(u)^l\right) du\right| = O_p(1/\sqrt{n}).$$

Therefore, $\Phi(\hat{G}_w, b, q|\mathcal{C}, K, N)$ uniformly converges to $\Phi(G_w, b, q|\mathcal{C}, K, N)$ at the $\sqrt{n}$-rate as $n \to \infty$ and so does the valuation estimator $\hat{v}$. *Q.E.D.*

Proof of Theorem 4: Using the same argument as used in the proof of Theorem 2, given the fixed reserve price r, the advertiser with the weighted value ω has a payoff that satisfies

$$\omega = \underset{\hat{\omega}}{\operatorname{argmax}} \sum_{k=1}^{K} \zeta_k(\hat{\omega}) c_k \left[\omega_j - \mathbb{E}\left(Max\{b_w(\omega^{[k+1]}), r\}\middle|\omega^{[k+1]} \leq \hat{\omega} \leq \omega^{[k-1]}\right)\right] \tag{25}$$

Additionally, similar to Equation (20) derived in the proof of Theorem 2, here again we derive the advertiser's payoff function to be equal to

$$\Pi(\omega) = \sum_{k=1}^{K} c_k \int_r^{\omega} \zeta_k(x) dx. \tag{26}$$

Using Equations (25) and (26), we obtain

$$\sum_{k=1}^{K} c_k \int_r^{\omega} \zeta_k(x) dx = \sum_{k=1}^{K} c_k \zeta_k(\omega) \left[\omega - \mathbb{E}\left(Max\{b_w(\omega^{[k+1]}), r\}\middle|\omega^{[k+1]} \leq \hat{\omega} \leq \omega^{[k-1]}\right)\right]$$

After differentiating the above equation and algebraic manipulations we yield

$$\sum_{k=1}^{K} c_k \omega \frac{d\zeta_k(\omega)}{d\omega} = \sum_{k=1}^{K} c_k \binom{N-1}{k-1} b_w(\omega)(N-k) F_w(\omega)^{N-k-1} f_w(\omega)(1-F_w(\omega))^{k-1}$$

$$-\sum_{k=1}^{K} c_k \binom{N-1}{k-1} (k-1)(1-F_w(\omega))^{k-2} f_w(\omega) \int_r^{\omega} b_w(x)(N-k) F_w(x)^{N-k-1} f_w(x) dx$$

$$-\sum_{k=1}^{K} c_k \binom{N-1}{k-1}(k-1)(1-F_w(\omega))^{k-2} f_w(\omega) \int_0^r r(N-k) F_w(x)^{N-k-1} f_w(x) dx \tag{27}$$

Then the same steps used in the proof of Theorem 2 obtain the desired result. *Q.E.D.*

APPENDIX C: SIMULATIONS

We examine the finite sample performances of our estimator and the impact of the correlation between value and quality on auction outcomes via the following simulation designs. We consider auctions with 5 ad positions. For $K > 2$ and arbitrary N, no analytic representation for the equilibrium weighted bid exists. Due to the complexity of numerically implementing the equilibrium bid function in Lemma 1, an iterative approximation procedure is employed instead to solve the equilibrium strategy. The specifics of this procedure can be found in Appendix D. The solution to the equilibrium bid equation depends on the parameters $\{c_1, \cdots, c_5\}$, N, and F_w. For sample sizes $n = 1200$, we simulate the auctions n/N times with $N = 10, 25, 50$ and 100, resulting in a consistent number of observations for each value of N. We draw the value v_j and the advertiser-specific click rate s_j from the following joint log-normal distribution:

$$\begin{bmatrix} v_j \\ s_j \end{bmatrix} \sim L\left(\begin{bmatrix} \mu_v \\ \mu_s \end{bmatrix}, \begin{bmatrix} \sigma_v^2 & \sigma_{vs} \\ \sigma_{vs} & \sigma_s^2 \end{bmatrix}\right),$$

so that v_j and s_j can be correlated. We consider the standard GSP^w auction where the scoring rule is simply $q_j = s_j$. As a result, the weighted value is also log-normally distributed ($\omega_j \sim L(\mu_v + \mu_s, \sigma_v^2 + \sigma_s^2 + 2\sigma_{vs})$). We use parameter values $\mu_v = -0.5, \mu_s = -3.5, \sigma_v^2 = 0.2, \sigma_{vs} = 0, \sigma_s^2 = 0.1$. We consider cases where v and s are uncorrelated ($\sigma_{vs} = 0$), negatively correlated ($\sigma_{vs} = -0.1$), or positively correlated ($\sigma_{vs} = 0.1$) here. We set $c_k = (1/2)^{k-1}$ so that the position-specific click rate quickly drops across ad positions. We estimate valuations using the proposed method. This process is repeated 1000 times for each combination of n and N.

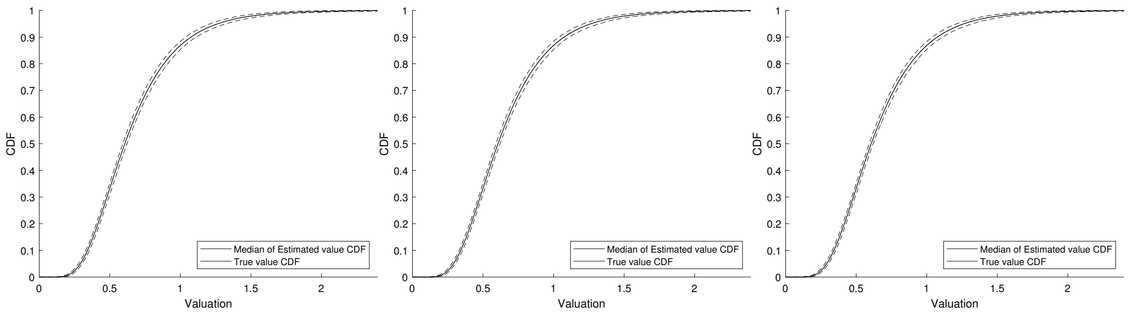

(a) $\sigma_{vs} = -0.1$ (b) $\sigma_{vs} = 0$ (c) $\sigma_{vs} = 0.1$

FIGURE C.1.—Finite sample performances when $n = 1200$. The dashed lines are the 5th and 95th percentiles of the estimated value distributions across 1000 replications.

Figure C.1 displays the median, 5th, and 95th percentiles of the estimated value distributions in comparison to the true distribution.[34] The median of the estimated distribution coincides with the true distribution. It also demonstrates the finite sample performances of the proposed estimator. The 5th and 95th percentiles are very close to the true distribution even with a relatively small sample size. The availability of extensive observations across repeated auctions in online auction data renders the proposed estimator capable of accurately estimating the unobserved values from bids.

Figure C.2 displays the equilibrium weighted bid as a function of the weighted value. We find that as N increases, bid shading approaches zero. This aligns with what is observed in other auction designs, such as first-price auctions.[35] Table C.I presents a summary of the selected quantiles of bid shading percentage. We find that the decrease in bid shading percentage with an increase in N is much slower for higher quantiles. Furthermore, for each quantile, the bid shading percentage becomes higher as the value and the quality score are more positively correlated. This is because the advertisers with higher values tend to also have higher quality scores in the positive correlation case so they are much more likely to win ad positions when they shade bids more compared to the negative correlation case. Note that lower bid shading in the negative correlation case does not necessarily mean that

[34]The underlying value distribution remains constant regardless of the value of N, resulting in identical estimation results for different values of N.

[35]For further details on effects of N on bids in the first price auction, refer to Krishna (2009), Ch 2.3.

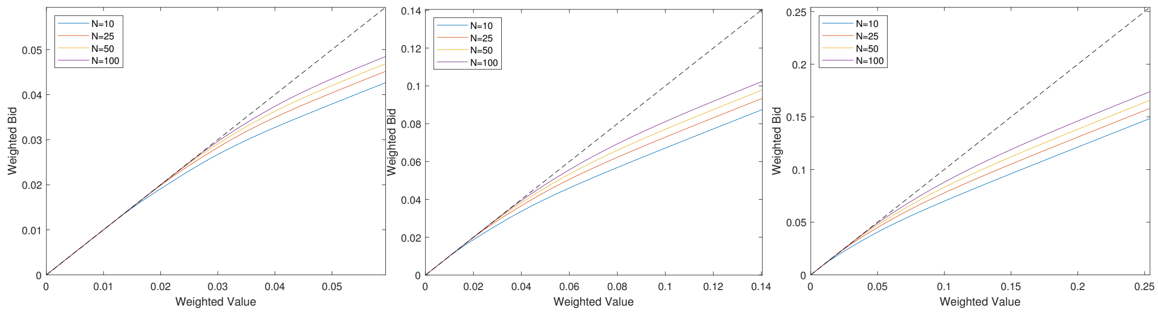

(a) $\sigma_{vs} = -0.1$ (b) $\sigma_{vs} = 0$ (c) $\sigma_{vs} = 0.1$

FIGURE C.2.—Equilibrium weighted bids given weighted values. The dashed line in each plot is a 45-degree line

the search engine will have higher revenue. To further investigate this, our next figure plots various auction outcomes across different degrees of correlation.

TABLE C.I

QUANTILES OF BID SHADING PERCENTAGE (%)

	Negative correlation				Zero correlation				Positive correlation			
	# of advertisers (N)				# of advertisers (N)				# of advertisers (N)			
qtile	10	25	50	100	10	25	50	100	10	25	50	100
25%	1.07	0.08	0.02	0.00	1.79	0.14	0.03	0.01	2.26	0.18	0.03	0.00
50%	3.29	0.38	0.08	0.02	5.42	0.64	0.13	0.03	6.76	0.82	0.17	0.04
75%	6.24	1.66	0.43	0.10	10.05	2.80	0.73	0.17	12.37	3.55	0.94	0.22
90%	9.22	4.33	1.92	0.60	14.46	7.12	3.24	1.02	17.48	8.88	4.10	1.30
99%	17.01	11.39	8.09	5.43	25.01	17.63	12.91	8.86	29.13	21.16	15.83	11.04

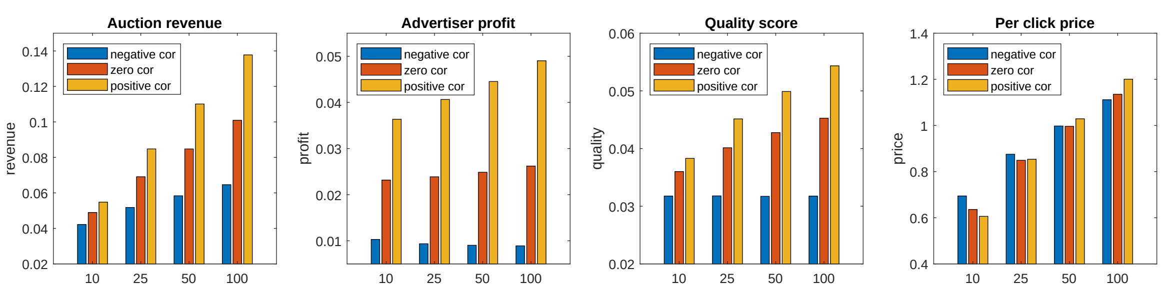

FIGURE C.3.—Auctioneer revenue, advertiser profit, quality score, and per click price

Figure C.3 reports the means of auction revenue, advertiser profit, quality of winning advertisers, and per click price across 1000 Monte Carlo samples for different values of N (the number of advertisers) and DGPs. The search engine's revenue is influenced by both the per-click price and the click rate (captured by the quality score). For lower N, the negative correlation case yields a higher per-click price, possibly due to lower bid shading, while also resulting in the lower average quality of the winning ads. We find that the impact of the quality score dominates, leading to auction revenue being positively correlated with the degree of correlation between v and s. Additionally, as competition increases with larger N, the negative correlation case no longer has a higher per-click price due to the gap in bid shading diminishing, and the average value of winning advertisers being lower compared to the other cases. Lastly, we find that advertiser profit is positively correlated with the degree of correlation. In the case of negative correlation, advertiser profit slightly decreases as N becomes larger, while the opposite occurs in the zero correlation case. We see a more rapid increase in the positive correlation case. Winning advertisers have higher weighted values when N is larger but at the same time, they can shade their bids less. The impact of bid shading is more dominating as the degree of correlation becomes lower.

APPENDIX D: ITERATIVE PROCEDURE FOR THE EQUILIBRIUM BID

This section explains the iterative procedure used in our counterfactual analysis to approximate the equilibrium bid function. From Equation (21) in the proof of Theorem 2, one can derive the following equation:

$$b_w(\omega) = \omega - \frac{\sum_{k=1}^{K} c_k \binom{N-1}{k-1}(k-1)(1-F_w(\omega))^{k-2}\delta(\omega|b_w, F_w, f_w, N, k)}{\sum_{k=1}^{K} c_k \binom{N-1}{k-1} F_w(\omega)^{N-k-1}(1-F_w(\omega))^{k-2}\Big[(N-k)(1-F_w(\omega)) - (k-1)F_w(\omega)\Big]}$$

where $\delta(\omega|b_w, F_w, f_w, N, k) = b_w(\omega)F_w(\omega)^{N-k} - \int_0^\omega b_w(x)(N-k)F_w(x)^{N-k-1}f_w(x)dx$. We do not know about b_w so we start with our initial guess $b_w^{(1)}(\omega) = \omega$. We plug the initial guess in the RHS of the above equation and check whether the resulting function is sufficiently close to the guess. If not, we update our guess and repeat the step. More formally, for each n stage, using the guess $b_w^{(n)}$, we compute the RHS of the equation to obtain $b_w^{(n*)}$ as follows.

$$b_w^{(n*)}(\omega) = \omega - \frac{\sum_{k=1}^{K} c_k \binom{N-1}{k-1} (k-1)(1-F_w(\omega))^{k-2} \delta(\omega | b_w^{(n)}, F_w, f_w, N, k)}{\sum_{k=1}^{K} c_k \binom{N-1}{k-1} F_w(\omega)^{N-k-1}(1-F_w(\omega))^{k-2} \Big[(N-k)(1-F_w(\omega)) - (k-1)F_w(\omega) \Big]}$$

At the end of each stage, we check whether $\sup_\omega |b_w^{(n*)}(\omega) - b_w^{(n)}(\omega)| < \varepsilon$ for some small $\varepsilon > 0$. If the condition is met, we stop and use $b_w^{(n*)}$ as the approximated equilibrium bid function. Otherwise, we update our guess $b_w^{(n+1)} = ab_w^{(n)} + (1-a)b_w^{(n*)}$ for some $a \in (0,1)$ and iterate the steps until convergence. As $n \to \infty$, $b_w^{(n)}$ converges to β_w from above. We set $\varepsilon = 10^{-4}$ and $a = 0.9$. In every case we consider, the convergence is quickly achieved with n less than 100.

APPENDIX E: IDENTIFICATION OF PRODUCT CATEGORIES

The raw dataset we use in the paper has multiple product categories, namely *cruise, car insurance, laptop, cable*, and *coins*. Additionally, the keywords are declassified and thus the product categories are also declassified (stated by numbers 0-4).[36] To overcome this problem, we analyze the differences in the categories and match each of them to the closest possible category among *cruise, car insurance, laptop, cable*, and *coins* according to the observed features. Table E.I gives a summary of how variables differ across categories. This table also shows the corresponding means of all features in different categories.

First, look at the features of category 1. This category is the easiest to identify because there are no keywords with one word. The base keyword consists of two words which must be '*car insurance*'. We can see that all the keywords in this category share the same two-word base keyword. This category is characterized by the very high average bid as well as the relatively small number of competitors compared to the other categories. This is consistent with the *car insurance* category. They are known to be the industry with one

[36] The categories are identified through the base keywords. The data have four single-word base keywords which identify the four categories, and one category is identified by a two-word base keyword.

of the highest prices per click. This is due to the high profit margins in the auto insurance industry which is a highly concentrated market.

The next category that stands out is category 2, which is characterized by a high number of advertisers and a high number of search queries per day. Due to its high volume of consumer searches, this is likely a consumer good. Therefore, it is closest to the '*laptop*' category as that is the only consumer good category in the data. Another category that is easy to identify is category 0. This category has a high number of advertisers and a high average click-through rate. A key feature of this category is more detailed searches that have longer keyword lengths. This is again a popular category with detailed search, and thus it is best matched with the '*cruise*' category.

Category 4 has the lowest average bid and search volume, as well as the fewest number of advertisers. It is most likely to be the least popular category in the data and thus likely '*coins*'. Lastly, '*cable TV*' is also a less popular and less expensive category but it is relatively more popular than '*coins*'. Therefore, we match '*cable TV*' with category 3. The table below summarizes the findings. Although these claims are just from our speculation, we use this classification for the main analysis in the paper. Even if there is some error in identifying the categories, we can still use the features of the category and interpret how and why the results might differ for categories with different features.

TABLE E.I

FEATURES OF DIFFERENT CATEGORIES

Category	**Description**	**CTR** (%)	**bid** (cent)	**adv**	**search** (million)	**length**
Cruise (Cat 0)	high competition & detailed search	1.28	0.51	6223	1320	3.18
Car insurance (Cat 1)	highest bids & high concentration	0.44	3.59	3815	2509	3.75
Laptop(Cat 2)	popular & high competition	1.33	0.45	4764	2913	3.03
Cable (Cat 3)	less popular & high bids	0.64	0.77	4703	1874	3.02
Coins (Cat 4)	low value across variables	1.36	0.36	3330	784	2.83

Note: This table shows the summary of key features of product categories in the raw dataset. 'CTR' is the average click-through rates, 'bid' is the average bid, 'adv' is the number of advertisers, 'search' is the daily average search volume, and 'length' is the average word counts across keywords within each category.

APPENDIX F: LOG-NORMAL APPROXIMATION OF THE VALUE DISTRIBUTION

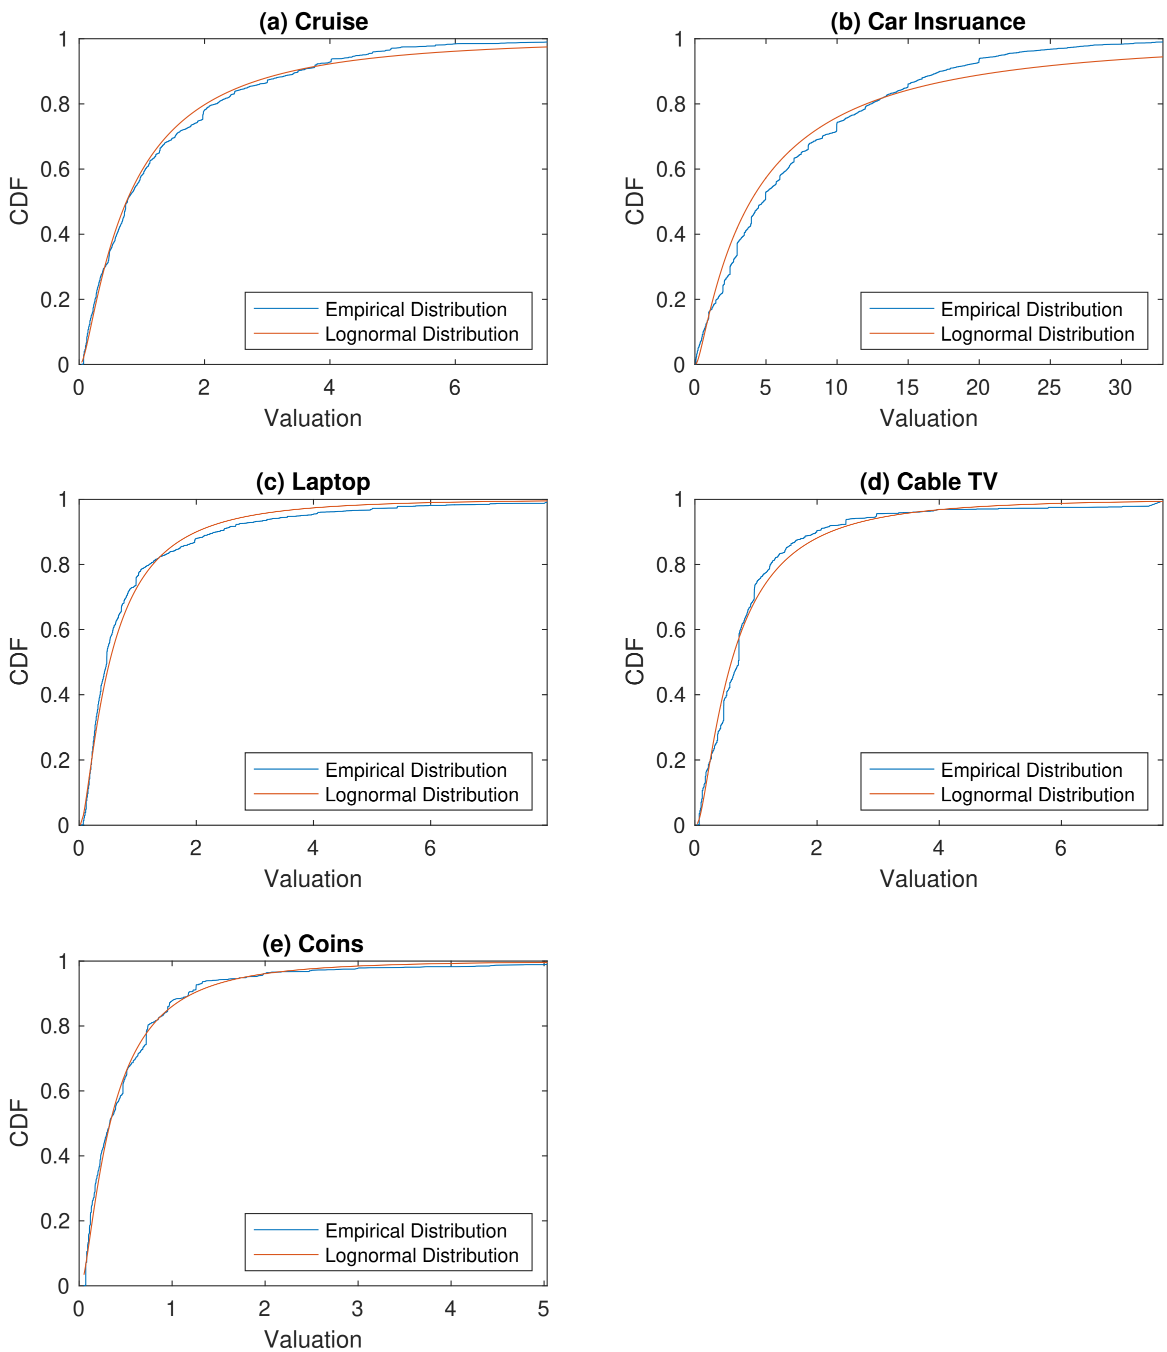

FIGURE F.1.—The fitted log-normal distribution of advertiser's ad value. The blue lines are the empirical distributions and the red lines are fitted log-normal distributions.

APPENDIX G: ROBUSTNESS ANALYSIS

G.1. *Click rate regression with interactive fixed effects*

To justify the log-linear regression specification (9) in Section 5, we compare it with the following specifications:

$$ctr_{k,j,m} = \alpha_k + \beta_j + \gamma_m + \epsilon_{k,j,m}, \tag{28}$$

$$log(ctr_{k,j,m}) = \delta_{kj} + \gamma_m + \epsilon_{k,j,m}. \tag{29}$$

The latter model incorporates interactive (advertiser-position specific) fixed effects, thereby allowing for heterogeneity among advertisers in click rate changes across positions. Results in Table G.I reveal that the log-linear model (9) offers significant improvements in most categories in terms of both adjusted R^2 and F-statistics compared to the linear model. Although the most flexible specification (29) marginally improves adjusted R^2, it is accompanied by substantially lower F-statistics than (9). These findings suggest that heterogeneity among advertisers in position-specific click-through rates is limited and the log-linear model is a good approximation for click probability within our dataset.

TABLE G.I

ADJUSTED R^2 AND F-STATISTIC FOR DIFFERENT MODELS

	Adjusted R^2 (F-statistics in the parentheses)				
	Cruise	Car Ins.	Laptop	Cable TV	Coins
Linear model	0.457	0.177	0.325	0.254	0.489
	(225.0)	(47.2)	(173.4)	(131.2)	(249.5)
Log-linear Model (additive FE)	0.533	0.526	0.403	0.482	0.459
	(304.2)	(238.6)	(243.0)	(356.4)	(221.7)
Log-linear Model (interactive FE)	0.563	0.585	0.483	0.508	0.496
	(37.4)	(32.5)	(35.3)	(40.4)	(26.1)

G.2. *Full sample analysis and keyword heterogeneity*

We restrict the data to the top 10% most popular keywords only to make sure that the underlying valuation distributions across markets are similar. Here we consider the full sets of keywords in our estimation. This empirical strategy results in much larger sample sizes so that the underlying valuation distributions are more precisely estimated. Note that the included keywords are now much more heterogeneous and hence the common value distribution assumption is more likely to be violated. Nonetheless, this exercise can shed some light on keyword heterogeneity. We expect that the popular keywords correspond to higher valuations than less popular keywords.

TABLE G.II

POSITION SPECIFIC CLICK RATES WITH THE FULL SAMPLE

Ad-position	Cruise	Car Insurance	Laptop	Cable TV	Coins
2^{th}	0.457	0.837	0.599	0.522	0.618
3^{th}	0.293	0.479	0.482	0.417	0.484
4^{th}	0.129	0.256	0.279	0.253	0.384
5^{th}	0.091	0.180	0.339	0.208	0.287
6^{th}	0.042	0.289	0.171	0.163	0.196
7^{th}	0.038	0.205	0.143	0.133	0.114

Note: The click rate of the first position is 1 due to re-scaling and hence omitted in the table.

Tables G.II-G.III show the position-specific click rates, quantiles of quality scores, and quantiles of bid shading percentages estimated on the full sample. Figure G.1 compares the empirical distributions of the estimated valuations for the popular keywords and the full sample. The estimated bid shading percentages on the full sample are overall smaller than those for the popular keywords. This is partly due to the lower position-specific click rates. The full sample contains many keywords that are rarely searched by consumers. The results show that ads related to unpopular keywords are less likely to be clicked when they are displayed. Another reason for less bid shading is the lower weighted value. Quality scores

TABLE G.III

QUANTILES OF QUALITY SCORE AND BID SHADING PERCENTAGE ACROSS PRODUCT CATEGORIES

	Quality score					Bid shading percentage (%)				
qtile	Cruise	Car Ins.	Laptop	Cable TV	Coins	Cruise	Car Ins.	Laptop	Cable TV	Coins
25%	0.0000	0.0002	0.0002	0.0111	0.0001	0.1	0.0	0.0	0.0	0.1
50%	0.0000	0.0004	0.0004	0.0169	0.0002	0.6	0.0	0.1	0.1	0.5
75%	0.0001	0.0010	0.0009	0.0242	0.0003	5.7	0.1	0.7	0.2	3.4
90%	0.0001	0.0021	0.0018	0.0299	0.0005	16.5	0.3	5.2	1.6	20.6
99%	0.0007	0.0054	0.0092	0.0772	0.0019	27.0	13.0	38.4	24.3	48.0

Note: The table summarizes the quantiles of the estimated quality scores and bid shading in terms of the percentage of the corresponding estimated values.

and valuations are overall lower on the full sample. As advertisers with higher weighted values can shade their bids more, the overall bid shading is smaller on the full sample.

The results on the full sample clearly show heterogeneity across keywords. The valuation distributions of the popular keywords are strictly below their counterparts from the full samples. The gaps are larger for '*cruise*' and '*car insurance*' categories. This implies that keyword heterogeneity is more outstanding in these categories. We speculate that the two categories have more diverse underlying products in terms of price. For instance, the price gap between 'luxury cruise' and 'budget cruise' would be much larger than the gap between 'high-end laptop' and 'budget laptop.' These results support our baseline empirical strategy that focuses on popular keywords.

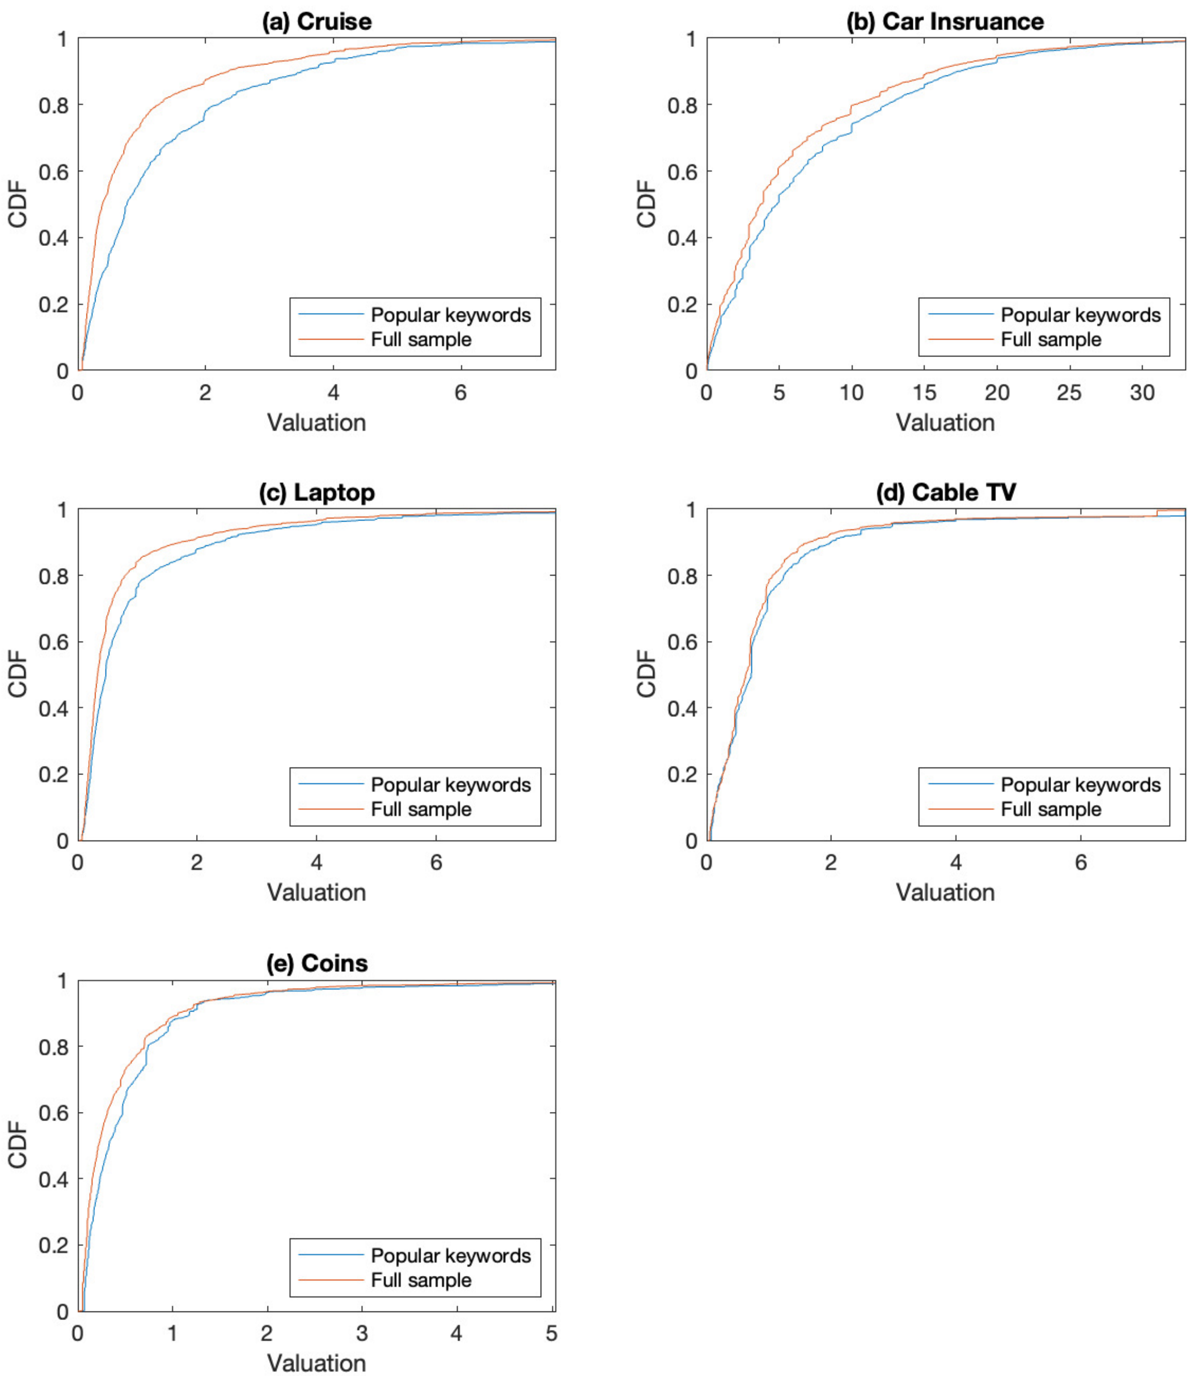

FIGURE G.1.—The valuation distributions for the top 10% popular keywords and the full sample.