ECAI 2016
G.A. Kaminka et al. (Eds.)

doi:10.3233/978-1-61499-672-9-515

Modeling Bounded Rationality for Sponsored Search Auctions

Jiang Rong[1] and **Tao Qin**[2] and **Bo An**[3] and **Tie-Yan Liu**[2]

Abstract. Sponsored search auctions (SSAs) have attracted a lot of research attention in recent years and different equilibrium concepts have been studied in order to understand advertisers' bidding strategies. However, the assumption that advertisers are perfectly rational in these studies is unrealistic in the real world. In this work, we apply the quantal response equilibrium (QRE), which is powerful in modeling bounded rationality, to SSAs. Due to high computational complexity, existing methods for QRE computation have very poor scalability for SSAs. Through exploiting the structures of QRE for SSAs, this paper presents an efficient homotopy-based algorithm to compute the QRE for large-size SSAs, which features the following two novelties: 1) we represent the SSAs as an Action Graph Game (AGG) which can compute the expected utilities in polynomial time; 2) we further significantly reduce redundant calculations by leveraging the underlying relations between advertisers' utilities. We also develop an estimator to infer parameters of SSAs and fit the QRE model into a dataset from a commercial search engine. Our experimental results indicate that the algorithm can significantly improve the scalability of QRE computation for SSAs and the QRE model can describe the real-world bidding behaviors in a very accurate manner.

1 Introduction

Sponsored search has become a major monetization means for commercial search engines (e.g., Google, Yahoo! and Bing) and has shown great success [23, 29, 38]. When a user issues a query to a search engine, in addition to several relevant webpages, a set of selected advertisements will also be displayed on the search result page. To show his/her ad on the search result page, an advertiser (or bidder) is required to submit a bid for the query. Most of the time, there are many more advertisers bidding for the query than the number of available ad slots and the search engines need a mechanism to decide which ads should be shown on the result page, how to allocate the slots to the shown ads, and how to charge an advertiser if his/her ad is clicked by users.

Generalized Second Price (GSP) is the most popular mechanism used in sponsored search auctions (SSAs) and has attracted much research attention recently [16, 35, 38, 40]. Among those studies, equilibrium analysis is a hot topic to understand advertisers' behaviors. Varian [45] studied the concept of symmetric Nash equilibrium for GSP auctions and proved its existence. Edelman et al. [15] defined a subset of Nash equilibria called locally envy-free equilibria which are equivalent to the symmetric Nash equilibria. Borgers et al. [5] further proved the existence of multiple Nash equilibria in GSP auctions. The forward looking Nash equilibrium was studied in [7, 8].

A critical limitation of existing studies on equilibrium analysis is that they assume the full rationality of advertisers. That is, advertisers are very smart; they can find their optimal strategies and take optimal actions. In practice, an advertiser may fail in estimating his/her competitors' bidding strategies and therefore cannot take the "best-response" action [21, 46]. As a result, it is necessary to study the equilibrium for SSAs based on bounded rationality, which is the focus of this work.

In this paper, we introduce the *quantal response equilibrium* (QRE) [17, 20, 33, 34] into SSAs, which can deal with bounded rationality and has demonstrated very good performance in general normal form games (NFGs). Specifically, because of the limited information about the market and opponents in the real world, an advertiser cannot calculate his/her accurate utility, where the error is assumed to follow some distribution (e.g., the extreme value distribution [33, 34]) with a *precision parameter* (i.e., a measurement of a advertiser's rationality). Due to the disturbance of the errors, advertisers can only maximize their inaccurate utilities in each round of the auctions, which makes their outcome policies to form a QRE — a mixed-strategy equilibrium in which strategies with higher utilities are more likely to be chosen than those with lower utilities, but the best one is not chosen with certainty. A higher precision parameter implies that the advertiser is more rational and hence can choose the better strategies with higher probabilities.

We focus on designing an algorithm for the search engine (or the auctioneer) to compute the QRE[4] in SSAs, which can be used to capture advertisers' bidding behaviors. We show that the calculation of a QRE is equivalent to computing the fixed point of Browder's functions [6, 28, 42], the complexity of which is at least PPAD-complete [10, 11, 37]. We further formulate the problem as finding a solution of a continuous non-linear function. Basic Newton-type algorithms are usually locally convergent and work well only when we could provide a good starting point, which, however, is difficult to find in SSAs. To address this problem, we leverage the *homotopy* principle [2, 3, 30], which has been used for equilibrium computation [18, 22, 39, 44]. Advantages of homotopy-based methods include their numerical stability and potential to be globally convergent. We noticed that Gambit [32] used a similar method to compute the QRE for NFGs, which, however, is very time-consuming and cannot be directly applied to large-size SSAs. To tackle this challenge, we leverage some nice properties of the SSAs (as compared to general N-

[1] The Key Lab of Intelligent Information Processing, Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences. Email: rongjiang13@mails.ucas.ac.cn
[2] Microsoft Research. Email: {taoqin, tie-yan.liu}@microsoft.com
[3] School of Computer Science and Engineering, Nanyang Technological University. Email: boan@ntu.edu.sg

[4] There might be multiple QREs and the one computed with our algorithm is called the principal equilibrium, which is mostly studied in the literature [18, 22, 44]

FGs), including the context-specific independence structure and the underlying relations between advertisers' utilities, to refine the computational procedure and significantly speed up the algorithm. The experimental results show that the improved homotopy algorithm can efficiently compute a QRE for SSAs in large sizes. We also investigate how to use the QRE model to infer the parameters, including the values and precisions of bidders and the click-through-rates (CTRs). We develop an estimating algorithm based on the commonly used Maximum Likelihood Estimation (MLE) principle [1, 26, 36]. Our experiments show that the QRE model fits the real data well.

To sum up, this paper makes two major contributions. First, we design an efficient homotopy-based algorithm to compute the QRE for large-size SSAs by utilizing the nice properties of SSAs. Second, we fit the QRE model into the real data and do extensive experiments to show that, comparing with Nash equilibria, the QRE is more practical since it can handle bidders' bounded rationality and model the real world in a very accurate manner.

The rest of this paper is organized as follows. In Section 2, we introduce the model of the GSP mechanism and then define the QRE for SSAs. The homotopy-based algorithm is proposed in Section 3, including the methods for significantly improving the efficiency of the algorithm. Section 4 gives the parameter estimation algorithm. Then we conduct extensive experiments in Section 5 to evaluate our algorithms. Conclusions are given in the last section.

2 Quantal Response Equilibrium

In this section, we first demonstrate our motivation of investigating the QRE for SSAs, following which we give some notations and assumptions and then define the QRE for SSAs.

2.1 Motivation

Given that GSP is not a dominant-strategy mechanism, Nash equilibrium solutions become an important means to understand how bidders behave in SSAs. While there exist quite a few Nash equilibrium concepts proposed and studied for SSAs, symmetric Nash equilibrium [45] and locally envy-free Nash equilibrium [15] are the most famous two: the former captures the notion that there should be no incentive for any pair of bidders to swap their slots, and the latter captures the notion that there should be no incentive for any bidder to exchange bids with the bidder ranked one position above him/her.

While those equilibrium concepts have many nice properties, a common limitation of them is that they assume the perfect rationality of bidders, i.e., bidders have perfect knowledge about their utilities and take optimal actions to maximize their utilities. However, the perfect rationality assumption is too good to be true in real-world SSAs. Therefore, a natural question arises: how can we weaken the perfect-rationality assumption and still obtain some meaningful solution concept for SSAs?

Observing that in real-world SSAs, an advertiser usually has uncertainty about his/her utility and is more likely to choose strategies with higher utilities instead of always choosing the best one. We introduce the quantal response equilibrium to model the bounded rationality of bidders in the following two subsections.

2.2 Notations and Assumptions

We focus on the GSP mechanism. Generally there are N bidders competing for K ad slots ($N \geq K$). We use the symbol $[N]$ to represent the set $\{1, 2, \ldots, N\}$. Let v_i denote the private value of bidder i, which expresses the maximum per-click price he/she is willing to pay, and the vector $v = (v_1, v_2, \ldots, v_N)$ represents the value profile of all bidders. We use b_i to represent the bid submitted by i to participate in the auction. θ_{ik} is the CTR of i's ad when placed at slot k, which is usually assumed to be the product of the ad CTR α_i and the slot CTR β_k [38]. We use $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_N)$ and $\beta = (\beta_1, \beta_2, \ldots, \beta_K)$ to denote the profiles of ad and slot CTRs respectively. Following the common practice [38], without loss of generality, we assume

$$\beta_1 \geq \beta_2 \geq \cdots \geq \beta_K.$$

In the GSP mechanism, bidders are ranked in the descending order of their ranking scores which is defined as

$$s_i = \alpha_i b_i.$$

The top $k \leq K$ bidders whose scores are not less than the reserve price r are allocated at the first k slots. If an ad is clicked, the payment of the corresponding bidder is the minimum amount that maintains his/her current rank position.

2.3 Definition of QRE

Let B_i denote advertiser i's bid space and b_{ij} be the j-th minimal price in B_i. We define the score space of bidder i as

$$S_i = \{s_{ij} | s_{ij} = \alpha_i b_{ij},\ j \in [|B_i|]\},$$

where $|B_i|$ represents the size of B_i, and define the joint score space of all advertisers except i as

$$S_{-i} = \times_{l \in [N] \setminus \{i\}} S_l.$$

Then we have that

$$|S_i| = |B_i|,\ \forall i \in [N].$$

Let $q_{ij}(s_{-i})$ and $p_{ij}(s_{-i})$ denote the slot allocated to and the payment of bidder i, given $s_{-i} \in S_{-i}$ and $s_{ij} \in S_i$. Then the utility of advertiser i is

$$u_{ij}(s_{-i}) = \begin{cases} 0, & s_{ij} < r; \\ (v_i - p_{ij}(s_{-i}))\alpha_i \beta_{q_{ij}(s_{-i})}, & s_{ij} \geq r. \end{cases} \quad (1)$$

Let σ_i be i's mixed strategy over B_i and σ_{ij} denote the probability on b_{ij}. Similarly, we define $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_N)$ and $\sigma_{-i} = (\sigma_1, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_N)$. The expected utility of i with s_{ij}, given σ_{-i}, is

$$\overline{u}_{ij}(\sigma_{-i}) = \sum_{s_{-i} \in S_{-i}} P(s_{-i} | \sigma_{-i}) u_{ij}(s_{-i}), \quad (2)$$

where $P(s_{-i}|\sigma_{-i})$ is the probability that the score profile of other bidders except i is s_{-i} given σ_{-i}. The quantal response π_{ij} of bidder i to others' mixed strategy profile σ_{-i} is defined as[5]

$$\pi_{ij}(\sigma_{-i}) = \frac{1}{\lambda_i} \cdot \frac{1}{|B_i|} + (1 - \frac{1}{\lambda_i}) \frac{e^{\overline{u}_{ij}(\sigma_{-i})\lambda_i}}{\sum_{k \in [|B_i|]} e^{\overline{u}_{ik}(\sigma_{-i})\lambda_i}}$$

[5] In the commonly-used logit form of quantal response [33, 44], multiplying α_i and dividing λ_i for all $i \in [N]$ by the same constant will not change the QRE outcome, which means that they are undistinguishable. To address this problem, we use a slightly different definition which satisfies the principle of QRE, i.e., better strategies are more likely to be chosen.

$$= \frac{1}{\lambda_i|B_i|} + (1-\frac{1}{\lambda_i})\frac{1}{\sum_{k\in[|B_i|]} e^{(\overline{u}_{ik}(\sigma_{-i})-\overline{u}_{ij}(\sigma_{-i}))\lambda_i}}, \quad (3)$$

where $\lambda_i \in [1,+\infty)$ is the *precision parameter* of bidder i. We can easily verify that Eq. (3) is consistent with our expectation that better strategies are more likely to be chosen than worse ones. When $\lambda_i = 1$, we have $\pi_{ij}(\sigma_{-i}) = \frac{1}{|B_i|}$, which means that i uniformly chooses any strategy in B_i; and when $\lambda_i \mapsto +\infty$, the choice probability of the bid strategy with the highest expected utility approaches 1. To sum up, λ_i is a parameter to measure i's "rationality". That is, a larger λ_i suggests that i will choose the best strategy with a higher probability. We use the vector $\lambda = (\lambda_1, \lambda_2, \cdots, \lambda_N)$ to denote the precision profile. A QRE [17, 33] is defined as follows.

Definition 1. *A* quantal response equilibrium *with λ is a mixed strategy profile σ such that for all $i \in [N]$ and $j \in [|B_i|]$, $\sigma_{ij} = \pi_{ij}(\sigma_{-i})$.*

Definition 1 implies that a QRE strategy profile is a fixed point of Browder's functions [6]. Besides, the Nash equilibrium is a special case of the QRE with $\lambda_i \to +\infty$ for all $i \in [N]$.

3 Equilibrium Computation

In this section, we design an efficient homotopy-based algorithm to compute the QRE for SSAs. We first describe the algorithm and then show how to significantly speed it up by exploiting peculiarities of SSAs.

3.1 The Homotopy-based Algorithm

In this subsection, we discuss how to compute a QRE given v, λ, α and β. We define

$$U = \sum_{i\in[N]} |B_i|$$

and obtain a continuous function $F : [0,1]^U \mapsto [0,1]^U$ from Definition 1 as below:

$$F_{ij}(\sigma) = \pi_{ij}(\sigma_{-i}) - \sigma_{ij}, \forall i \in [N], j \in [|B_i|].$$

Now we can see that computing a QRE of SSAs is equivalent to finding a zero point of the nonlinear function $F(\sigma)$. If a good starting point is available, we can directly apply Newton-style iteration methods. However, we have little information about such a good initial point. As pointed out by Allgower and Georg [2, 3], Newton-style iteration methods often fail because poor starting points are very likely to be chosen.

The basic idea of the homotopy is composed of two steps: given a problem we want to solve, first, define a problem with a unique easy-to-compute solution and then build a continuous transformation from the artificial problem into the original problem we want to solve; second, begin with the solution of the easy-to-solve problem and trace solutions of the associated problems of the transformation until finally the solution of the original problem is found.

To design a homotopy-based algorithm, our first step is to propose a degenerate problem which is easy solve. In particular, we can find a degenerated form of F (denoted as G) by letting $\overline{u}_{ij}(\sigma_{-i}), \forall i \in [N]$ and $j \in [|B_i|]$, be zero:

$$G_{ij}(\sigma) = \frac{1}{|B_i|} - \sigma_{ij}, i \in [N], j \in [|B_i|].$$

Obviously, $G(\sigma)$ has a unique zero point: $\sigma_{ij} = 1/|B_i|, \forall i \in [N], j \in [|B_i|]$. Then we define a homotopy function $H : [0,1]^U \times [0,1] \mapsto [0,1]^U$ between $F(\sigma)$ and $G(\sigma)$ as

$$H_{ij}(\sigma,t) = \frac{1}{\lambda_i|B_i|} + (1-\frac{1}{\lambda_i})R(\sigma_{-i},t)^{-1} - \sigma_{ij},$$
$$R(\sigma_{-i},t) = \sum_{k\in[|B_i|]} e^{(\overline{u}_{ik}(\sigma_{-i})-\overline{u}_{ij}(\sigma_{-i}))\lambda_i t},$$

which is a continuous transformation from $H(\sigma,0) = G(\sigma)$ to $H(\sigma,1) = F(\sigma)$ as t grows continuously from 0 to 1.

We further define the solution set of $H(\sigma,t) = 0$ as

$$H^{-1}(0) = \{(\sigma,t)|H(\sigma,t) = 0\}.$$

It follows from Browder's fixed point theorem [6] that, for a given $t \in [0,1]$, there is a $\sigma(t)$ such that $H(\sigma(t),t) = 0$. From the definition of H we know that $\sigma(0)$ and $\sigma(1)$ correspond to the zero point of $G(\sigma)$ and $F(\sigma)$ respectively. The remaining problem is to trace out a path consisting of $(\sigma(t),t) \in H^{-1}(0)$, which starts at $(\sigma(0),0)$ and ends at $(\sigma(1),1)$. Considering the possibility of the existence of turning points [27], increasing t monotonically when tracing the path may lead to points far away from the path. A common practice to avoid the disturbance of turning points is to view the σ and t as functions of an implicit parameter a simultaneously and to compute a parametric path

$$c(a) = (\sigma(t(a)), t(a)),$$

which satisfies

$$H(c(a)) = 0. \quad (4)$$

The method we use to trace the path is called predictor-corrector (PC) [2, 3], the basic idea of which is to numerically trace the path $c(a)$ by generating a sequence of points $c_i = (\sigma,t)_i, i = 1,2,\ldots$ along the path satisfying $\|H(c_i)\| \le \varepsilon$ for some $\varepsilon > 0$. In particular, given that we have found a point c_i on the path $c(a)$, an Euler predictor step is used to predict the next point c_{i+1} on $c(a)$:

$$c_{i+1} = c_i + \Delta \cdot \frac{c'(a)\big|_{c(a)=c_i}}{\| c'(a)\big|_{c(a)=c_i} \|}, \quad (5)$$

where $c'(a)$ is the derivative of $c(a)$ with respect to a and $\Delta > 0$ is the step length. Then a corrector phase is necessary to refine the accuracy of c_{i+1}. We make use of the Gauss-Newton method as presented below:

$$\widehat{c}_{i+1} = c_{i+1} - H'(c_{i+1})^+ H(c_{i+1}), \quad (6)$$

where $H'(c_{i+1})^+$ is the Moore-Penrose inverse[6] of the Jacobian matrix $H'(c_{i+1})$ of $H(\cdot)$ at point c_{i+1}, and $\widehat{c}_{i+1}$ is the refined point of c_{i+1}. If $\|H(\widehat{c}_{i+1})\| > \varepsilon$, we will substitute $\widehat{c}_{i+1}$ into the right side of Eq. (6) to further refine it. The corrector procedure may be performed several times until we find the satisfactory point which will be used in the predictor phase to infer the next point. The PC method, starting with $(\sigma(0),0)$, is applied step by step until $(\sigma(1),1)$ is reached.

Now we discuss how to compute the derivative $c'(a)$ in Eq. (5) and the Jacobian matrix $H'(c_{i+1})$ in Eq. (6). We first consider the calculation of $c'(a)$. By differentiating Eq. (4) we get the following equation:

$$H'(c(a))c'(a) = 0. \quad (7)$$

[6] The Moore-Penrose inverse is defined by $A^+ = A^T(AA^T)^{-1}$.

The solution of Eq. (7) is

$$c'_d(a) = \mu \cdot (-1)^d \cdot det(H'_{-d}(c(a))) \quad (8)$$

where $c'_d(a)$, $d=1,\cdots,U+1$, denotes the d-th[7] component of $c'(a)$ and $H'_{-d}(c(a))$ is $H'(c(a))$ with the d-th column removed; $det(\cdot)$ is the determinant operation; $\mu=\pm 1$ is the sign of $c'(a)$ to be chosen. We know from Eq. (8) that once $H'(c(a))$ is given, $c'(a)$ can be obtained directly. Eq. (6) also involves computing $H'(\cdot)$. So next we will concentrate on how to compute $H'(\cdot)$. We use $\widetilde{u}_{ij}$ to represent $\overline{u}_{ij}(\sigma_{-i})\lambda_i$ for simplicity in the remaining part of this paper. Since the function H is a mapping from $[0,1]^U \times [0,1]$ to $[0,1]^U$, its Jacobian matrix contains $U\cdot(U+1)$ partial derivatives that can be divided into four cases:

Case 1. $i \in [N]$ and $j \in [|B_i|]$:

$$\frac{\partial H_{ij}}{\partial \sigma_{ij}} = -1;$$

Case 2. $i \in [N]$, j and $k \in [|B_i|]$, $j \neq k$:

$$\frac{\partial H_{ij}}{\partial \sigma_{ik}} = 0;$$

Case 3. i and $l \in [N]$, $j \in [|B_i|]$, $m \in [|B_l|]$, $i \neq l$:

$$\frac{\partial H_{ij}}{\partial \sigma_{lm}} = -(1-\frac{1}{\lambda_i})R(\sigma_{-i},t)^{-2}\frac{\partial R(\sigma_{-i},t)}{\partial \sigma_{lm}},$$
$$\frac{\partial R(\sigma_{-i},t)}{\partial \sigma_{lm}} = te^{-\widetilde{u}_{ij}t}\sum_{k\in[|B_i|]} e^{\widetilde{u}_{ik}t}(\frac{\partial \widetilde{u}_{ik}}{\partial \sigma_{lm}} - \frac{\partial \widetilde{u}_{ij}}{\partial \sigma_{lm}});$$

Case 4. $i \in [N]$, $j \in [|B_i|]$:

$$\frac{\partial H_{ij}}{\partial t} = -(1-\frac{1}{\lambda_i})R(\sigma_{-i},t)^{-2}\frac{\partial R(\sigma_{-i},t)}{\partial t},$$
$$\frac{\partial R(\sigma_{-i},t)}{\partial t} = e^{-\widetilde{u}_{ij}t}\sum_{k\in[|B_i|]} e^{\widetilde{u}_{ik}t}(\widetilde{u}_{ik}-\widetilde{u}_{ij}).$$

We choose μ to ensure that the derivative of t with respect to a, which corresponds to the $(U+1)$-th component in $c'(a)$, is positive at $(\sigma(0),0)$, i.e.,

$$\mu \cdot (-1)^{U+1} det(H'_{-(U+1)}(\sigma(0),0)) > 0.$$

Substituting $t=0$ into this inequality and combining with cases 1-3, we get

$$\mu \cdot (-1)^{U+1} \cdot (-1)^U = (-1)^{2U+1}\mu > 0,$$

which indicates that $\mu = -1$ at $(\sigma(0),0)$.

By now we have discussed how to compute Eqs. (5) and (6). Then we can use the PC method to find the point $(\sigma(1),1) \in H^{-1}(0)$ step by step, the convergence property of which is analyzed in [2, 3].

The complete process of our proposed method is summarized in Algorithm 1. In line 1 we assign t with 0 and the starting point c_1 with $(\sigma(0),0)$. Line 2 initializes Δ and ϵ. Lines 3-8 use the PC method to generate a set of points $c_i, i=2,3,\ldots$ along the path until eventually the point $(\sigma(1),1)$ is found. Line 4 is the Euler predictor step which computes the next point c_{i+1} given c_i according to Eq. (5). The Gauss-Newton corrector step is performed repeatedly in lines 5-6 to improve the accuracy of the point predicted in line 4. The value of t is updated in line 7 and the step length is adjusted in line 8 based on the Asymptotic Expansion method proposed in [2, 3]. The result is returned in line 10.

[7] σ_{ij}'s are assumed to be assigned to $c(a)$ in lexicographic order of their subscripts. The last component of the vector corresponds to t.

Algorithm 1: Computing a QRE

1 $t \leftarrow 0$, $c \leftarrow (\sigma(0),0)$;
2 initialize Δ and ε;
3 **while** $t \neq 1$ **do**
4 $\quad c \leftarrow c + \Delta \cdot \frac{c'(a)|_{c(a)=c}}{\|(c'(a)|_{c(a)=c})\|}$;
5 $\quad$ **while** $\|H(c)\| > \varepsilon$ **do**
6 $\quad\quad c \leftarrow c - H'(c)^+ H(c)$;
7 $\quad t \leftarrow$ the last component of c;
8 $\quad$ Adjust the step length Δ;
9 $(\sigma, t) \leftarrow c$;
10 **return** σ;

3.2 Efficient Computation for SSAs

Algorithm 1 indicates that we need to compute the Jacobian matrix $H'(\cdot)$ at each predictor and corrector step when tracing the path. Thus the efficiency of calculating $H'(\cdot)$ will significantly affect the scalability of the algorithm. In this subsection, we discuss how to reduce the complexity of Algorithm 1 through efficient calculation of $H'(\cdot)$ by leveraging the properties of SSAs. First. we represent the SSAs as an Action Graph Game (AGG) [24, 25, 43] which can compute the components of $H'(\cdot)$ in polynomial time, while general NFGs cost exponential time to calculate them. Second, we further significantly reduce redundant calculations based on the analyses of the relations between advertisers' utilities.

3.2.1 Representing SSAs as AGG

The elements in $H'(\cdot)$ are classified into four cases as shown in Section 3.1, the last two cases of which involve computing $\widetilde{u}_{ij}$ and $\frac{\partial \widetilde{u}_{ij}}{\partial \sigma_{lm}}$. We can rewrite Eq. (2) as

$$\overline{u}_{ij}(\sigma_{-i}) = \sum_{m\in[|S_l|]} \sigma_{lm} \sum_{s_{-il}\in S_{-il}} P(s_{-il}|\sigma_{-il}) u_{ij}(s_{-il}, s_{lm})$$
$$= \sum_{m\in[|S_l|]} \sigma_{lm} \frac{\partial \overline{u}_{ij}(\sigma_{-i})}{\partial \sigma_{lm}}, \forall l \neq i, \quad (9)$$

where $S_{-il} = \times_{i'\in[N]\setminus\{i,l\}} S_{i'}$ and s_{-il} is an element of S_{-il}. So the main effort on calculating $H'(\cdot)$ is to compute a set of partial derivatives, i.e.,

$$D = \{\frac{\partial \overline{u}_{ij}(\sigma_{-i})}{\partial \sigma_{lm}} | i,l \in [N]; l \neq i; j \in [|B_i|]; m \in [|B_l|]\}.$$

We need to traverse the $\prod_{i'\in[N]\setminus\{i,l\}} |B_{i'}|$ realizations in S_{-il} to compute $\frac{\partial \overline{u}_{ij}(\sigma_{-i})}{\partial \sigma_{lm}}$ if we view SSAs as general NFGs. Clearly, this traversal method (TM for short) is exponential, i.e., $O(M^N)$, where

$$M = \max\{|B_{i'}| \big| i' \in [N]\}.$$

Fortunately, expected utilities in SSAs with GSP mechanism have many special properties that could be utilized to reduce the computational complexity. Here we take bidder i with s_{ij} and bidder l with s_{lm} as an example ($s_{ij} \geq r$) and define

$$I_G = \{n | s_n > s_{ij}, n \neq i\}.$$

Similarly, we further define

$$I_E = \{n | s_n = s_{ij}, n \neq i\}$$

and

$$I_L = \{n | s_n < s_{ij}, n \neq i\}.$$

We assume the tie is broken randomly. It thus follows that:

1. When $|I_G| \geq K$, i's utility is zero.
2. When $|I_G| < K$ and $|I_G| + |I_E| \geq K$, bidder i has a probability $\frac{1}{|I_E|+1}$ to be allocated at a slot ranging from $|I_G| + 1$ to K and his/her payment is $p_i = b_{ij}$. According to our assumption on the tie, i' utility in such case is

$$\frac{1}{|I_E|+1}(v_i - b_{ij}) \sum_{k=|I_G|+1}^{K} \alpha_i \beta_k.$$

3. When $|I_G| + |I_E| < K$, i's location will be any one from slot $|I_G|+1$ to slot $|I_G|+|I_E|+1$ with an identical probability $\frac{1}{|I_E|+1}$. His/her payment is $p_i = b_{ij}$ if ranked at the slot between $|I_G| + 1$ and $|I_G| + |I_E|$. On the other hand, his/her payment is p_{max}/α_i if allocated at slot $|I_G| + |I_E| + 1$, where

$$p_{max} = \max\{r, s_n | n \in I_L\}.$$

Bidder i's utility in this case is

$$\frac{1}{|I_E|+1}((v_i - b_{ij}) \sum_{k=|I_G|+1}^{|I_G|+|I_E|} \alpha_i \beta_k + (v_i - p_{max})\alpha_i \beta_{|I_G|+|I_E|+1}).$$

The above properties indicate that given bidder i's score $s_i = s_{ij}$, his/her utility only depends on $|I_G|$, $|I_E|$ and p_{max}, but not on who are in I_G and I_E or exactly what their bids are, and not on whose ranking score is p_{max}. That is SSAs have considerable context-specific independence structure and can be represented compactly by an AGG. An action graph (AG) is a trie[8], each leaf of which corresponds to a tuple $(|I_G|, |I_E|, p_{max})$. Specifically, when computing $\frac{\partial \overline{u}_{ij}(\sigma_{-i})}{\partial \sigma_{lm}}$, because 1) we just consider the cases where $|I_G| < K$, 2) $|I_E| \leq N$, and 3) p_{max} has at most NM different values, the AG has $O(KN^2M)$ leaves, which can be built in time $O(KN^2MN) = O(KN^3M)$ by a dynamic program [24]. Compared with that of TM ($O(M^N)$), it is a significant improvement.

3.2.2 Reducing Redundancy

The computation of the set D with AGG involves two steps: 1) building the AGs and 2) calculating the partial derivatives based on these AGs. Intuitively, we need to apply the two procedures to each of the $N(N-1)M^2$ elements of D. Actually, we can significantly reduce the calculations by utilizing the properties of SSAs.

We first focus on the process of AG construction (i.e., step 1). The AG for $\frac{\partial \overline{u}_{ij}(\sigma_{-i})}{\partial \sigma_{lm}}$ is built by a dynamic program which traverses the union set

$$\{s_{nk} | n \in [N] \setminus \{i, l\}, k \in [|B_n|]\} \cup \{s_{lm}\},$$

the first part of which is the same for all $m \in [|B_l|]$. Thus we can just build one trie with $\{s_{nk} | n \in [N] \setminus \{i, l\}, k \in [|B_n|]\}$, from which the AGs for $\frac{\partial \overline{u}_{ij}(\sigma_{-i})}{\partial \sigma_{lm}}$, $\forall m \in [|B_l|]$, can be directly derived by further taking s_{lm} into consideration. This observation implies that we can compute D by building at most $N(N-1)M$ AGs. Next we show how to further refine step 1 based on the following propositions.

[8] Trie is an ordered tree data structure. More details can be found at `https://en.wikipedia.org/wiki/Trie` and [4].

Proposition 1. *For all $l \neq i$, $\frac{\partial \overline{u}_{ij}(\sigma_{-i})}{\partial \sigma_{lm}} = 0$ if $s_{ij} < r$.*

This proposition is straightforward since Eq. (1) indicates that $u_{ij}(s_{-il}, s_{lm}) = 0$ if $s_{ij} < r$.

Proposition 2. *For all $l \neq i$,*

$$\frac{\partial \overline{u}_{ij}(\sigma_{-i})}{\partial \sigma_{lm}} = \frac{\partial \overline{u}_{ij+1}(\sigma_{-i})}{\partial \sigma_{lm}} = \ldots = \frac{\partial \overline{u}_{ij+k}(\sigma_{-i})}{\partial \sigma_{lm}},$$

if there is no $s_{l'm'}$, $l' \neq i$ and $m' \in [|B_{l'}|]$, satisfying $s_{ij} \leq s_{l'm'} \leq s_{ij+k}$.

Proof. I_E is empty under the assumption. Given b_{-il}, bidders except i are either in I_G or in I_L and will not change their positions when i's score changes from s_{ij} to s_{ij+n}, $n \in [k]$. Then we have that, for all $s_{-il} \in S_{-il}$, $u_{ij}(s_{-il}, s_{lm}) = u_{ij+n}(s_{-il}, s_{lm}), \forall n \in [k]$, because they have the same I_G, I_E and I_L. Taking expectation over S_{-il} completes the proof according to Eq. (9). □

Proposition 3. *For all $l \neq i$, the AGs for $\frac{\partial \overline{u}_{ij}(\sigma_{-i})}{\partial \sigma_{lm}}$ and $\frac{\partial \overline{u}_{lj}(\sigma_{-i})}{\partial \sigma_{im}}$ are the same if $S_i = S_l$.*

Proof. Since the AG is indeed a trie, we only need to show that they have the same leaves. Given $s_{-il} \in S_{-il}$, because $s_{ij} = s_{lj}$ and $s_{im} = s_{lm}$, the profiles $s_{-i} = (s_{-il}, s_{lm})$ for s_{ij} and $s_{-l} = (s_{-il}, s_{im})$ for s_{lj} can be mapped to the same tuple $(|I_G|, |I_E|, p_{max})$. By traversing S_{-il} and considering that there is a one-to-one correspondence between the tuples and leaves, we prove the proposition. □

Proposition 2 shows that if a subset $\{s_{ij}, s_{ij+1}, \ldots, s_{ij+k}\}$ of S_i satisfies the constraint, then $\frac{\partial \overline{u}_{ij+1}(\sigma_{-i})}{\partial \sigma_{lm}}, \ldots, \frac{\partial \overline{u}_{ij+k}(\sigma_{-i})}{\partial \sigma_{lm}}$ can be obtained directly through $\frac{\partial \overline{u}_{ij}(\sigma_{-i})}{\partial \sigma_{lm}}$ without building the AGs. Proposition 2 cannot be applied to the case where bidders have the same score space. By contrast, Proposition 3 is particularly useful for this case, which can reduce the number of AGs to $\frac{N(N-1)M}{2}$ in step 1. The next proposition is used to speed up step 2.

Proposition 4. *For all $l \neq i$,*

$$\frac{\partial \overline{u}_{ij}(\sigma_{-i})}{\partial \sigma_{lm_1}} = \frac{\partial \overline{u}_{ij}(\sigma_{-i})}{\partial \sigma_{lm_2}},$$

if 1) $s_{lm_1}, s_{lm_2} \leq r$ or 2) $s_{lm_1}, s_{lm_2} > s_{ij}$.

Proof. Similar to the proof of Proposition 2, we just need to prove that for all $s_{-il} \in S_{-il}$, $u_{ij}(s_{-il}, s_{lm_1}) = u_{ij}(s_{-il}, s_{lm_2})$, which is true due to the fact that (s_{-il}, s_{lm_1}) and (s_{-il}, s_{lm_2}) correspond to the same tuple $(|I_G|, |I_E|, p_{max})$ for s_{ij} on each of the two conditions. □

We use an example where bidders have the same score space to analyze the effects of Propositions 3 and 4 in step 2. We assume $r = 0$ for ease of analysis. Given i and l, we only need to compute $\frac{\partial \overline{u}_{ij}(\sigma_{-i})}{\partial \sigma_{lm}}$ for $j < M$ and $m \leq j + 1$, and $\frac{\partial \overline{u}_{iM}(\sigma_{-i})}{\partial \sigma_{lm}}$ for $m \leq M$, because those for $j < M$ and $m > j + 1$ are equal to $\frac{\partial \overline{u}_{ij}(\sigma_{-i})}{\partial \sigma_{lj+1}}$ according to Proposition 4. Then the amount of components calculated with AGs is

$$(2 + 3 + \ldots + M + M) = \frac{M^2 + 3M - 2}{2}.$$

On the other hand, given j and m, Proposition 3 implies that $\frac{\partial \overline{u}_{ij}(\sigma_{-i})}{\partial \sigma_{lm}}$ and $\frac{\partial \overline{u}_{lj}(\sigma_{-i})}{\partial \sigma_{im}}$ as a pair can be computed simultaneously

with one AG. There are $\frac{N(N-1)}{2}$ such pairs. As a result, the number of calculations reduced in total is

$$N(N-1)M^2 - \frac{N(N-1)}{2}\frac{M^2+3M-2}{2} > \frac{3}{4}N(N-1)M(M-1).$$

Since D has $N(N-1)M^2$ elements, the efficiency for step 2 is improved by about 75%.

4 Parameter Estimation

In this subsection, we propose an algorithm based on the principle of MLE and QRE to estimate v, α and β of SSAs from real data. It is pointed out in [19] that the auctioneer itself can not accurately estimate the CTRs and it is not rare to get a 50% error. Our model provides an alternative way for the search engine to infer these parameters when they are unknown.

Given a QRE strategy σ, the logarithmic likelihood of the unknown parameters v, λ, α and β is

$$L(v,\lambda,\alpha,\beta|\sigma) = \log(\prod_{i\in[N]}\prod_{j\in[|B_i|]}\pi_{ij}(\sigma_{-i})^{\sigma_{ij}}) \tag{10}$$

Then the parameters can be estimated by maximizing the likelihood as shown in the following optimization problem:

$$\max_{v,\lambda,\alpha,\beta} L(v,\lambda,\alpha,\beta|\sigma) \tag{11}$$

$$s.t.\begin{cases} v_i \geq 0,\ \forall i\in[N];\\ \lambda_i \geq 1,\ \forall i\in[N];\\ 1>\beta_s\geq\beta_{s+1}>0,\ \forall s\in[K-1];\\ 0<\alpha_i<1,\ \forall i\in[N]. \end{cases} \tag{12}$$

However, $q_{ij}(s_{-i})$ is not a continuous function of $\alpha_i, \forall i\in[N]$, nor are the utilities of bidders. As a result, the likelihood defined in Eq. (10) is not continuous with respect to α.

To address the discontinuity of the likelihood function, we split the unknown parameters into two groups and sequentially optimize them: we treat v, λ, β as a group and α as the other group; in each iteration, we first optimize v, λ, β and then α.

The function $L(v,\lambda,\alpha,\beta|\sigma)$ in Eq. (11) is continuous with respect to the parameters in the first group. We can learn a better set of v, λ, β by solving the following sub optimization problem:

$$\max_{v,\lambda,\beta} L(v,\lambda,\alpha,\beta|\sigma) \tag{13}$$

$$s.t.\begin{cases} v_i \geq 0,\ \forall i\in[N];\\ \lambda_i \geq 1,\ \forall i\in[N];\\ 1>\beta_s\geq\beta_{s+1}>0,\ \forall s\in[K-1]. \end{cases} \tag{14}$$

Since the above optimization problem is non-convex, it is difficult to find the global maximum. We turn to find a set of local maxima with different starting points and then choose the best one to improve the possibility of reaching the global maximum of the sub problem.

As aforementioned, the likelihood function is not continuous with respect to α. Here we do not optimize bidders' ad CTRs simultaneously. Instead, we deal with them one by one. Let us take the ad CTR α_i of bidder i as an example and keep $\alpha_j, \forall j\neq i$ fixed. Given that all the other parameters are fixed, it is easy to know that the likelihood has the following discontinuous points:

$$\{\frac{\alpha_j b_j}{b_i} | 0<\frac{\alpha_j b_j}{b_i}<1, j\in[N]\setminus\{i\}, b_i\in B_i, b_j\in B_j\}.$$

Then we can partition the feasible domain of α_i into several intervals where the likelihood function L is continuous with respect to α_i, and then by solving the optimization problem defined in Eq. (15) in each interval we can find a better α_i given all the other parameters:

$$\max_{\alpha_i}\ L(v,\lambda,\alpha,\beta|\sigma) \quad s.t.\ \alpha_i \text{ in the continuous interval.} \tag{15}$$

Similarly, the above optimization problem is not convex. To avoid being tracked into a bad local maximum, we can also find a set of local maxima with different starting points and choose the best one.

Algorithm 2: Parameter estimation

```
1  L* ← −∞;
2  Randomly generate an ad CTR profile α;
3  while True do
4      Fix α and update v, λ, β by solving the problem shown in
       Eqs. (13) and (14);
5      for i ← 1, 2, ··· , N do
6          Fix α_j, ∀j ≠ i, v, λ, β and update α_i by solving the sub
           problem as shown in Eq. (15) in each continuous
           interval;
7      L̂ ← L(v, λ, α, β|σ);
8      if L̂ > L* then
9          L* ← L̂;
10     else return the learned parameters v, λ, α, β;
```

The complete procedure is presented in Algorithm 2. In line 1 we initialize the likelihood of the original optimization problem with negative infinity. Line 2 sets an initial α. Lines 3-9 iteratively optimize the two groups of parameters. Line 4 fixes α and updates (v, λ, β). Lines 5-6 fix (v, λ, β) and update α. Lines 7-10 control the optimization process: if we make progress in this iteration, we continue the optimization; otherwise, we return the latest parameters. Again, to avoid a bad local maximum, we run the algorithm for multiple times with different initial α's in Line 2 and choose the best parameters as the final output in our experiments.

5 Experimental Evaluation

We conduct extensive experiments to evaluate the algorithms for QRE computation and parameter estimation.

5.1 Effectiveness of the Homotopy Algorithm

We first evaluate the runtime of the three different approaches for computing D: TM, AGG, and AGG combined with our speed-up methods (AGGSU for short). The experiments are divided into two groups based on whether bidders have the same score space. In each group, we test the three methods with different game sizes where $N = 5, 6, \ldots, 20$ and $M = 10, 15, 20$. In all the experiments, α, β and B_i, $\forall i \in [N]$, are sampled from uniform distributions with supports $(0.1, 1)^N$, $(0.1, 1)^K$ and $(0, M)^M$ respectively. We let $K = \lfloor N/2 \rfloor$ and $v_i = \max_{b_i\in B_i} b_i$, $\forall i \in [N]$. The runtime for each method in each setting is averaged over 100 experiments. The

results are depicted in Figure 1 with logarithmic y-axis, where data greater than 10^4 seconds are not displayed and the symbol "-S" ("-D") denotes the same (different) score space group. Since the runtimes for TM in the two groups are almost the same, we just plot the average of them.

We see from Figure 1 that the runtime of TM increases exponentially as N grows. The 8×10 (i.e., $N = 8$, $M = 10$) game cannot be solved by TM within 1 hour. As a comparison, both AGG and AGGSU are much more efficient than TM. We observe that AGGSU-D (AGGSU-S) is about ten times faster than AGG-D (AGG-S), which confirms the efficiency of our speed-up methods proposed in Section 3.2.2. We further notice that AGGSU-S (AGG-S) always runs slower than AGGSU-D (AGG-D). That is because I_E is almost an empty set for AGGSU-D (AGG-D) and thus its AGs have $O(KNM)$ leaves, while the AGs for AGGSU-S (AGG-S) contain $O(KN^2M)$ leaves. Overall, our AGGSU method performs the best in all the settings.

Next we evaluate the performance of Algorithm 1. The parameters of the experiments are generated as above and λ is uniformly sampled from the support $(0, 10)^N$. We do not assume the identical score space and the set D is computed with AGGSU-D. The experiments are based on three different settings (game sizes). We use dynamic (Dy) and various fixed step lengths to test Algorithm 1. The results are depicted in Table 1, which shows the runtime of the algorithm in seconds (time), the number of Euler steps (#E), and the averaged amount of Gauss-Newton steps (#G) in one Euler step.

Table 1. Performance of Algorithm 1

	$N=10$, $M=5$			$N=10$, $M=10$			$N=15$, $M=10$		
Δ	time	#E	#G	time	#E	#G	time	#E	#G
Dy	0.55	5.10	2.40	5.66	8.00	2.44	19.5	5.42	2.38
0.1	1.92	26.0	0.73	18.6	47.0	0.76	82.1	29.1	0.68
0.3	0.85	10.9	1.03	8.58	18.1	1.05	37.8	11.2	1.03
0.5	0.82	7.45	1.71	7.67	12.3	1.83	33.0	8.03	1.37
0.7	0.81	6.76	2.40	7.68	10.7	2.32	26.8	5.45	2.44
0.9	0.82	5.44	2.60	7.92	8.12	3.56	30.1	4.96	2.75
1.1	0.85	4.34	3.25	10.8	10.2	3.91	37.7	5.02	3.03

We learn from Table 1 that larger step lengths usually lead to fewer Euler procedures, but more Gauss-Newton processes are needed to correct the zero point predicted in the Euler phase. The Gauss-Newton corrector often fails to converge when $\Delta \geq 1.3$, which implies that Newton-style methods cannot be directly used to compute the QRE. When Δ is small, the predicted zero points are often accurate enough and the corrector step is not needed, hence the averaged number of Gauss-Newton steps may be less than 1. Another interesting finding is that the numbers of steps for the predictor and corrector phases do not increase with the game size. We can verify that, given the numbers of bidders (N) and strategies (M), the runtime of Algorithm 1 is positively correlated to the total amount of steps and the dynamic step length strategy outperforms those strategies with fixed step lengths. The results indicate that our algorithm can efficiently compute the QRE for large-size SSAs.

5.2 Evaluate the Estimation Algorithm

We first used Algorithm 2 to infer the parameters of the QREs computed in Section 5.1 and found that the estimated parameters are always equal to the generated ones, which verifies the effectiveness of Algorithm 2 in parameter estimation. Next we conduct experiments based on Yahoo's public data on advertising and market [14, 41, 47], which contain the information about advertisers' bids and ranks over 4 months. More than 89% of the queries[9] in the dataset have less than 5 bidders. We find that bidders' information in many queries are very incomplete, i.e., there are only several records about a bidder over the 4 months. As in some related work like [13], we pick out 70 queries with almost complete information in the log (which do not include the queries containing just one bidder), and further remove the bidders who give a very high or very low bid and never make a change since these bidders will create singularity issues for the estimation and provide little information about bidders' behaviors.

We fit the QRE model into the processed dataset, in which the distribution of the number of bidders is shown in Table 2. Specifically, for each query, we first compute bidders' real mixed strategy profile σ with the log file, then we use Algorithm 2 to infer parameters with the QRE model, next by substituting the estimated parameters into Eq. (3), we compute bidders' quantal responses π_i for all $i \in [N]$. We first evaluate whether π_i is equal to σ_i, $\forall i \in [N]$, for each query, i.e., whether advertisers' bidding strategies (σ) form a QRE. To do this, we calculate the error $\frac{1}{\sum_{i\in[N]}|B_i|}\sum_{i\in[N],j\in|[B_i]|}|\pi_{ij} - \sigma_{ij}|$ for each query, based on which we compute the maximum, minimum and average of the errors over each of the four scenarios. The results are depicted in Table 2.

Table 2. Fitting Accuracy Evaluation

scenario	No. of bidders	distribution	maximum	minimum	average
1	3	77.14%	0.0732	0.0002	0.0373
2	4	15.71%	0.0813	0.0002	0.0388
3	5	4.29%	0.0922	0.0011	0.0432
4	≥ 6	2.86%	0.1078	0.0106	0.0592

We see from Table 2 that for some queries, the minimal errors are at the magnitude of 10^{-4}, and on average the errors are less than 0.06, which indicates that the QRE can model advertisers' behaviors in the real world well. In most cases, the worst-case errors are around 0.08. Thus it is reasonable to assume that bidders were playing the QRE. Next, following the practice in [45], we use two specific queries to show details about the parameters estimated with QRE model and then make a comparison with the mixed strategy Nash equilibrium (MSNE) model[10].

The parameters estimated with the QRE are depicted in Table 3, by substituting which into Eq. (3) we compute bidders' quantal responses (π_i) and expected utilities ($\overline{u}_i$), as shown in Tables 4 and 5. We see that better strategies are chosen with higher probabilities, which is consistent with the principle of QRE. We find that the quantal responses (π_i) are very close to the real mixed strategies (σ_i) in both tables, which implies that the QRE model can accurately describe bidders' bidding behaviors in the real world.

Table 3. Estimated parameters with QRE

query	i	v_i	λ_i	α_i	β	L^*
1	1	66.5	85	.04	.07	-3.2804
	2	43.0	1078	.03	.06	
	3	61.4	1602	.03	.05	
2	1	6.00	6251	.13	.15	-2.6098
	2	1.65	222	.65	.10	
	3	5.67	4700	.18	.05	

[9] For simplicity, we only consider exact match between a query and keywords. For broad match, please refer to [9, 12, 31].

[10] We do not compare the QRE with the symmetric Nash equilibrium [45] because the latter is a pure-strategy equilibrium which cannot explain advertisers' mixed-strategy behaviors.

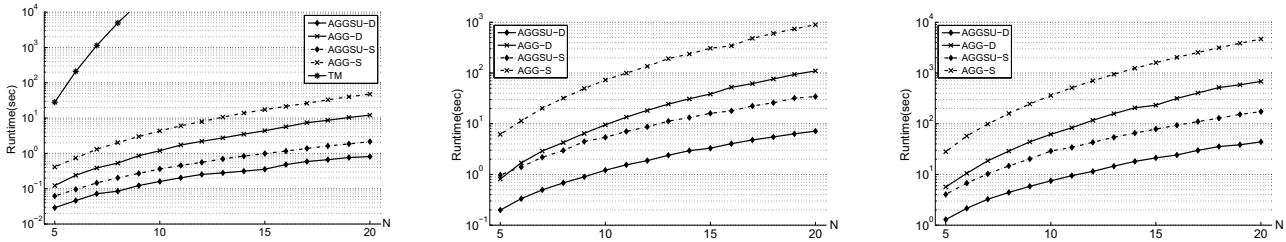

(a) M=10 (b) M=15 (c) M=20

Figure 1. Runtime of TM, AGG and AGGSU

Table 4. Strategies comparison for query 1

σ_1	π_1	$\overline{u}_1$	σ_2	π_2	$\overline{u}_2$	σ_3	π_3	$\overline{u}_3$
.1261	.1239	.1202	.2975	.2979	.0479	.2609	.2607	.0906
.3025	.3563	.1328	.5950	.5948	.0485	.3565	.3697	.0908
.4202	.3714	.1333	.1074	.1073	.0469	.3826	.3697	.0908
.1513	.1484	.1224						

Table 5. Strategies comparison for query 2

σ_1	π_1	$\overline{u}_1$	σ_2	π_2	$\overline{u}_2$	σ_3	π_3	$\overline{u}_3$
.2927	.2682	.0339	.6897	.6881	.0633	.0522	.0524	.0581
.2439	.2682	.0339	.1466	.1040	.0548	.1304	.1304	.0583
.4634	.4635	.0340	.1207	.1040	.0548	.8174	.8172	.0587
			.0431	.1040	.0548			

We learn from Table 3 that bidders' precision parameters differ from each other significantly. Note that Eq. (3) implies that besides λ_i, the magnitude of the difference between bidder i's expected utilities has a strong impact on his/her quantal response. To see this impact, we take bidder 1 in query 2 for an example. We know from Table 3 that $\lambda_1 = 6251$ for query 2, which seems to indicate that bidder 1 should be very rational because λ_1 is large. However, the quantal response $\pi_1 = (.2682, .2682, .4635)$ for query 2 in Table 5 indicates that bidder 1 is not very rational since he/she does not choose the optimal strategy with very high probability. That is because the difference between the components of $\overline{u}_1$ is at the magnitude of 10^{-4}, which reduces the effect of $\sum_{k\in[|B_1|]} e^{(\overline{u}_{1k}(\sigma_{-1})-\overline{u}_{1j}(\sigma_{-1}))\lambda_1}$ in Eq. (3) even though λ_1 is at a magnitude of thousands.

Next we fit the MSNE into the dataset, where each player' expected utilities by choosing different pure strategies are the same. Hence we solve the following optimization problem with Algorithm 2:

$$\widetilde{L}(v,\alpha,\beta|\sigma) = \log(\prod_{i\in[N]}\prod_{j\in[|B_i|]}(\frac{\overline{u}_{ij}(\sigma_{-i})}{\sum_{k\in[|B_i|]}\overline{u}_{ik}(\sigma_{-i})})^{\frac{1}{|B_i|}}) \quad (16)$$

$$s.t.\begin{cases} v_i \geq 0,\ \forall i\in[N]; \\ 1>\beta_s\geq\beta_{s+1}>0,\ \forall s\in[K-1]; \\ 0<\alpha_i<1,\ \forall i\in[N]. \end{cases} \quad (17)$$

which get its maximal value when

$$\frac{\overline{u}_{ij}(\sigma_{-i})}{\sum_{k\in[|B_i|]}\overline{u}_{ik}(\sigma_{-i})} = \frac{1}{|B_i|},\ \forall i\in[N],\ j\in[|B_i|],$$

or equivalently, when

$$\overline{u}_{ij}(\sigma_{-i}) = \overline{u}_{ik}(\sigma_{-i}),\ \forall i\in[N],\ j\in[|B_i|],\ k\in[|B_i|].$$

The estimated results are in Table 6. It shows that the maximal likelihoods of the queries with MSNE ($\widetilde{L}^*$) are less than those with QRE (L^*), which implies that QRE is more accurate than MSNE for modeling advertisers' pricing policies. We learn from the log that $B_1 = \{10, 15, 20, 30\}$, $B_2 = \{10, 20, 25\}$ and $B_3 = \{15, 25, 35\}$ for query 1. The values estimated by QRE are all larger than bids, which is consistent with the experience that bidders usually do not overbid [21]. As a comparison, the values predicted by MSNE are not very reasonable. Besides, the estimated CTRs (θ_{ik}) of MSNE for query 1 are overly large, e.g., $\alpha_1\beta_1 = 0.42$ and $\alpha_2\beta_1 = 0.76$, while the CTRs in the real world are generally lower than 10%. Furthermore, the log shows that the maximal bid in query 2 is not greater than 10, while the estimated values are at the magnitude of 10^6, which does not make sense. The slot CTRs (β) estimated by MSNE for query 2 also seem strange since in the real world CTRs usually decrease from the top position to the bottom one. As a comparison, those estimated with QRE match the real world well. Overall, QRE can fit the real data much better than MSNE.

Table 6. Estimated parameters with MSNE

query	i	v_i	α_i	β	$\widetilde{L}^*$
1	1	18.1	.51	.83	-3.5835
	2	12.8	.91	.02	
	3	112.2	.16	.02	
2	1	8.4×10^6	.02	.07	-4.7594
	2	4.8×10^6	.05	.07	
	3	3.08×10^6	.09	.07	

6 Conclusion

In this paper, we introduced the solution concept of QRE into SSAs to model the bounded rationality of advertisers' bidding behaviors. Along this line, we made two key technical contributions. First, we designed an efficient homotopy-based algorithm to compute the QRE for SSAs. By further utilizing the special structure of advertisers' expected utilities, we significantly improved the efficiency of our algorithm which can be applied to large-size SSAs. Second, we developed an estimation algorithm and fitted the QRE model into real data to infer values, precisions and CTRs of the SSAs. In addition, we conducted extensive experiments to evaluate the performance of our algorithms, which show that the proposed homotopy algorithm for computing QRE is very efficient and the QRE model can fit the real data much better than previous models.

Acknowledgement

This paper is supported by 2015 Microsoft Research Asia Collaborative Research Program.

REFERENCES

[1] Hirotogu Akaike, 'Information theory and an extension of the maximum likelihood principle', in *Selected Papers of Hirotugu Akaike*, 199–213, (1998).

[2] Eugene L Allgower and Kurt Georg, *Introduction to numerical continuation methods*, volume 45, SIAM, 2003.

[3] Eugene L Allgower and Kurt Georg, *Numerical continuation methods: an introduction*, volume 13, Springer Science & Business Media, 2012.

[4] Paul E Black, *Dictionary of algorithms and data structures*, National Institute of Standards and Technology, 2004.

[5] Tilman Börgers, Ingemar Cox, Martin Pesendorfer, and Vaclav Petricek, 'Equilibrium bids in sponsored search auctions: Theory and evidence', *American economic Journal: microeconomics*, **5**(4), 163–187, (2013).

[6] Felix E Browder, 'On continuity of fixed points under deformations of continuous mappings', *Summa Brasiliensis Mathematicae*, **4**, 183–191, (1960).

[7] Tian-Ming Bu, Xiaotie Deng, and Qi Qi, 'Forward looking nash equilibrium for keyword auction', *Information Processing Letters*, **105**(2), 41–46, (2008).

[8] Tian-Ming Bu, Li Liang, and Qi Qi, 'On robustness of forward-looking in sponsored search auction', *Algorithmica*, **58**(4), 970–989, (2010).

[9] Wei Chen, Di He, Tie-Yan Liu, Tao Qin, Yixin Tao, and Liwei Wang, 'Generalized second price auction with probabilistic broad match', in *Proceedings of the 15th ACM Conference on Economics and Computation*, pp. 39–56, (2014).

[10] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou, 'The complexity of computing a nash equilibrium', *SIAM Journal on Computing*, **39**(1), 195–259, (2009).

[11] Constantinos Daskalakis, Aranyak Mehta, and Christos Papadimitriou, 'Progress in approximate nash equilibria', in *Proceedings of the 8th ACM conference on Electronic commerce*, pp. 355–358, (2007).

[12] Peerapong Dhangwatnotai, *Auction Design with Robust Guarantees*, Ph.D. dissertation, Stanford University, 2012.

[13] Quang Duong and Sébastien Lahaie, 'Discrete choice models of bidder behavior in sponsored search', in *Internet and Network Economics*, 134–145, (2011).

[14] Benjamin Edelman and Michael Ostrovsky, 'Strategic bidder behavior in sponsored search auctions', *Decision support systems*, **43**(1), 192–198, (2007).

[15] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz, 'Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords', Technical report, National Bureau of Economic Research, (2005).

[16] Nicola Gatti, Alessandro Lazaric, and Francesco Trovò, 'A truthful learning mechanism for contextual multi-slot sponsored search auctions with externalities', in *Proceedings of the 13th ACM Conference on Electronic Commerce*, pp. 605–622, (2012).

[17] Jacob K Goeree, Charles A Holt, and Thomas R Palfrey, 'Regular quantal response equilibrium', *Experimental Economics*, **8**(4), 347–367, (2005).

[18] Paul W Goldberg, Christos H Papadimitriou, and Rahul Savani, 'The complexity of the homotopy method, equilibrium selection, and lemke-howson solutions', *ACM Transactions on Economics and Computation*, **1**(2), 9, (2013).

[19] Thore Graepel, Joaquin Q Candela, Thomas Borchert, and Ralf Herbrich, 'Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft's bing search engine', in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 13–20, (2010).

[20] Philip A Haile, Ali Hortaçsu, and Grigory Kosenok, 'On the empirical content of quantal response equilibrium', *The American Economic Review*, **98**(1), 180–200, (2008).

[21] Di He, Wei Chen, Liwei Wang, and Tie-Yan Liu, 'A game-heoretic machine learning approach for revenue maximization in sponsored search', in *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pp. 206–212, (2013).

[22] P Jean-Jacques Herings and Ronald Peeters, 'Homotopy methods to compute equilibria in game theory', *Economic Theory*, **42**(1), 119–156, (2010).

[23] Bernard J Jansen and Tracy Mullen, 'Sponsored search: an overview of the concept, history, and technology', *International Journal of Electronic Business*, **6**(2), 114–131, (2008).

[24] Albert Xin Jiang and Kevin Leyton-Brown, 'A polynomial-time algorithm for action-graph games', in *Proceedings of the 21st AAAI Conference on Artificial Intelligence*, pp. 679–684, (2006).

[25] Albert Xin Jiang, Kevin Leyton-Brown, and Navin AR Bhat, 'Action-graph games', *Games and Economic Behavior*, **71**(1), 141–173, (2011).

[26] Soren Johansen and Katarina Juselius, 'Maximum likelihood estimation and inference on cointegration with applications to the demand for money', *Oxford Bulletin of Economics and statistics*, **52**(2), 169–210, (1990).

[27] Kenneth L Judd, *Numerical methods in economics*, MIT press, 1998.

[28] R Bruce Kellogg, Tien-Yien Li, and James Yorke, 'A constructive proof of the brouwer fixed-point theorem and computational results', *SIAM Journal on Numerical Analysis*, **13**(4), 473–483, (1976).

[29] Sébastien Lahaie, David M Pennock, Amin Saberi, and Rakesh V Vohra, 'Sponsored search auctions', *Algorithmic game theory*, 699–716, (2007).

[30] Shijun Liao, 'On the homotopy analysis method for nonlinear problems', *Applied Mathematics and Computation*, **147**(2), 499–513, (2004).

[31] Mohammad Mahdian and Grant Wang, 'Clustering-based bidding languages for sponsored search', in *Algorithms-ESA 2009*, 167–178, Springer, (2009).

[32] Richard D McKelvey, Andrew M McLennan, and Theodore L Turocy, 'Gambit: Software tools for game theory, version 14.1.0. http://www.gambit-project.org.', (2014).

[33] Richard D McKelvey and Thomas R Palfrey, 'Quantal response equilibria for normal form games', *Games and economic behavior*, **10**(1), 6–38, (1995).

[34] Richard D McKelvey and Thomas R Palfrey, 'Quantal response equilibria for extensive form games', *Experimental economics*, **1**(1), 9–41, (1998).

[35] Paul Milgrom, 'Simplified mechanisms with an application to sponsored-search auctions', *Games and Economic Behavior*, **70**(1), 62–70, (2010).

[36] Garib N Murshudov, Alexei A Vagin, and Eleanor J Dodson, 'Refinement of macromolecular structures by the maximum-likelihood method', *Acta Crystallographica Section D: Biological Crystallography*, **53**(3), 240–255, (1997).

[37] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani, *Algorithmic game theory*, volume 1, Cambridge University Press Cambridge, 2007.

[38] Tao Qin, Wei Chen, and Tie-Yan Liu, 'Sponsored search auctions: Recent advances and future directions', *ACM Transactions on Intelligent Systems and Technology (TIST)*, **5**(4), 60, (2015).

[39] Jiang Rong, Tao Qin, and Bo An, 'Computing quantal response equilibrium for sponsored search auctions', in *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pp. 1803–1804, (2015).

[40] Jiang Rong, Tao Qin, Bo An, and Tie-Yan Liu, 'Optimal sample size for adword auctions', in *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pp. 1459–1460, (2016).

[41] Savannah Wei Shi and Xiaojing Dong, 'The effects of bid pulsing on keyword performance in search engines', *International Journal of Electronic Commerce*, **19**(2), 3–38, (2015).

[42] David Roger Smart, *Fixed point theorems*, volume 66, CUP Archive, 1980.

[43] David Robert Martin Thompson and Kevin Leyton-Brown, 'Computational analysis of perfect-information position auctions', in *Proceedings of the 10th ACM conference on Electronic commerce*, pp. 51–60, (2009).

[44] Theodore L Turocy, 'A dynamic homotopy interpretation of the logistic quantal response equilibrium correspondence', *Games and Economic Behavior*, **51**(2), 243–263, (2005).

[45] Hal R Varian, 'Position auctions', *International Journal of Industrial Organization*, **25**(6), 1163–1178, (2007).

[46] Haifeng Xu, Bin Gao, Diyi Yang, and Tie-Yan Liu, 'Predicting advertiser bidding behaviors in sponsored search by rationality modeling', in *Proceedings of the 22nd international conference on World Wide Web*, pp. 1433–1444, (2013).

[47] Yahoo, 'Webscope for advertising and market data. A3: ydata-ysm-keyphrase-bid-imp-click-v1_0. `http://webscope.sandbox.yahoo.com/`', (2013).