Private Data Manipulation in Optimal Sponsored Search Auction

Xiaotie Deng
Peking University
Beijing, China
xiaotie@pku.edu.cn

Tao Lin
Peking University
Beijing, China
lin_tao@pku.edu.cn

Tao Xiao
Shanghai Jiao Tong University
Shanghai, China
xt_1992@sjtu.edu.cn

ABSTRACT

In this paper, we revisit the sponsored search auction as a repeated auction. We view it as a learning and exploiting task of the seller against the private data distribution of the buyers. We model such a game between the seller and buyers by a Private Data Manipulation (PDM) game: the auction seller first announces an auction for which allocation and payment rules are based on the value distributions submitted by buyers. The seller's expected revenue depends on the design of the protocol and the game played among the buyers in their choice on the submitted (fake) value distributions.

Under the PDM game, we re-evaluate the theory, methodology, and techniques in the sponsored search auctions that have been the most intensively studied in Internet economics.

CCS CONCEPTS

• **Theory of computation → Algorithmic game theory**; *Algorithmic mechanism design*; Computational pricing and auctions.

KEYWORDS

sponsored search auction, private data, Internet economics

ACM Reference Format:

Xiaotie Deng, Tao Lin, and Tao Xiao. 2020. Private Data Manipulation in Optimal Sponsored Search Auction. In *Proceedings of The Web Conference 2020 (WWW '20), April 20–24, 2020, Taipei, Taiwan.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3366423.3380023

1 INTRODUCTION

The sponsored search auction (SSA) refers to the commercialization model of a search engine to sell a click on an advertisement slot by its users. Those advertisements are presented in web-links showing products and their prices of the advertisers along with natural search results. SSA refers to the mechanism chosen by the seller for the advertisers to compete for the chance to pay for the placement of their ads along with the search results.

Google, arguably the most successful search engine, has traditionally employed the generalized second price (GSP) auction for its SSA. The bidders are charged according to the bid (per quality unit) one rank lower than its own in the sorted bid list [22, 41]. Despite many other pricing and allocation mechanisms have been tried [16, 21, 25, 43], the GSP has sustained the test of competition to become one of the key factors for Google's success in the sponsored search market in the early years [8]. It has since become the most commonly used auction mechanism for sponsored search, and can reasonably be named the mother of all sponsored search auctions.

Taking into consideration of the repeated nature of SSA, a data analyst would examine every day's revenue with a report on the income derived from *each* keyword out of millions of clicks. The algorithm takes (bid) data-in and generates (revenue) data-out. In-between, a natural task of the seller is to maximize the expected revenue based on a longer and longer bid sequence from the buyers.

We are interested in such a data engineer's design task for today's market maker. As the auction takes place billions of times a day, the data engineer has a rich recorded history of advertisers' bidding behavior. With this information in relation with the value distribution of each advertiser for the available slots, the market maker would be expected to derive as much revenue as in the optimal auction of the Myerson style.

We model such a process as a two-stage operation. The market maker first announces an auction, such as the Myerson style optimal auction, whose allocation and payment rules may depend on value distributions of buyers. Then buyers submit value distributions. However, the auction that combines these two stages may not work truthfully, because buyers may cheat to report fake value distributions for eventually better expected utilities. This may be a marginal case in general auctions, but it can become a key challenge in the SSA as the game can be repeated billions of times on the major search engines. It is under such a situation we have an environment where the auction participants with their own private value distributions can manipulate their private value data.

In summary, since each advertiser could bid a distribution different from their private value distribution to gain more in an SSA, all their individual efforts result in an equilibrium in such a two-stage game led by the market maker. We call this game a Private Data Manipulation (PDM) game. In this article, we make a first step in understanding the revenue optimization issue against the strategic behavior of advertisers in the second stage of the game.

1.1 Our Results and Techniques

First of all, we consider the repeated auction of multiple items in the sponsored search auctions. We model it by a Private Data Manipulation (PDM) game where the strategy of advertisers (bidders) is to submit distributions to the market maker (auctioneer) who run (prior-dependent) auctions, according to the submitted distributions. We prove that under this PDM game, the originally truthful and revenue-maximizing auction (Myerson's, Mye) is no longer truthful. Every equilibrium in this game has a 1-to-1 mapping to the equilibrium of the untruthful and sub-optimal generalized first-price auction (GFP). The auctioneer obtains the same revenue under the respective equilibrium of Mye and GFP. Then combined with the classical revenue-equivalence theorem, we conclude that Mye, GFP,

WWW '20, April 20–24, 2020, Taipei, Taiwan

ACM ISBN 978-1-4503-7023-3/20/04.
https://doi.org/10.1145/3366423.3380023

GSP, and VCG auctions are equivalent under certain conditions (see Figure 1).

From a technical perspective, the sponsored search auction fits into the single-parameter environment, where the optimal auction, given buyers' distributions, is a Myerson style auction that achieves the maximum achievable expected revenue for the market maker among all truthful auctions. Most importantly, the expected utility of each agent in Mye under PDM turns out to be of the same form as that of GFP, allowing us to establish the equivalence between Mye and GFP in all single-parameter environments, including SSA.

In summary, our results develop another class of revenue equivalence theorem beyond that of single-item auction [34, 40] towards the sponsored search auction, including those for VCG [16, 25, 43] and Forward Looking Nash equilibrium [11] in GFP [21] and GSP [22, 41] to cover a realistic repeated auction setting among Internet scale applications. In this way, our results give a revenue justification for Google's recent switch from GSP to First Price Auctions [38] in the highly repeated auction SSA.

1.2 Related Work

There has been a considerable amount of work on the sponsored search auction, one can refer to [29] for a survey. Among all practically applied auctions for sponsored search, the generalized second-price auction stands out attracting the most attention in research [12, 22, 41]. A significant fraction has focused on social welfare properties of equilibria [14, 24, 30, 31] where the goal is to measure the price of anarchy. Others have studied revenue properties [23, 32, 42] where the goal is to measure the guaranteed revenue out of different equilibrium concepts.

The optimal auction design dates back to the seminal work from [34], which establish the optimal auction under the Bayesian setting of a prior value distribution for each buyer. In practice, the prior information does not come for free. This initiates the study of sample complexity [6, 7, 13, 17, 26, 28, 39], in which the auctioneer has a number of samples from the prior, and the goal is to learn a mechanism with those samples to approximate the optimal revenue. Another line of research focus on repeated auction scenario [3, 9], with the goal of minimizing regret, for bandit settings [4, 5, 20, 35], or settings with strategic bidder [1, 2, 10].

The idea of agent manipulation against learning is not only applicable in multi-agent auction design [18, 19, 36, 40], but can also become a common analytic tool in more general settings. We expect that this approach can find further applications in the future. Besides modeling the real-world practice, our private data manipulation is important for the protection of one's private data against effort trying to extract it for commercial use.

1.3 Organization of the paper

In Section 2, we set up the single parameter auction design under the private data manipulation model. In Section 3, we derive the main technical results showing that the optimal Myerson style auction and the generalized first auction are equivalent. We conclude in Section 4 with an array of relationships (see Figure 1) among GFP, GSP, VCG and Mye under the private data manipulation game.

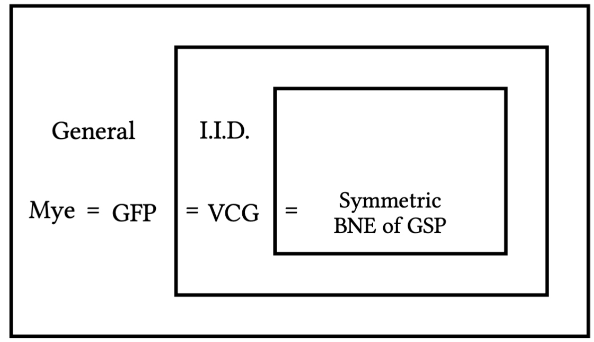

Figure 1: Equivalence between auctions under PDM

2 PRELIMINARY

Consider the setting where the market designer has m advertisement slots, and n bidders who are interested in these slots. Each advertisement slot j has a quality of γ_j. Each bidder i has a value v_i per quality over these slots. Therefore, if bidder i is allocated slot j, her value is $v_i \gamma_j$. Without loss of generality we assume that $\gamma_1 \geq \gamma_2 \geq \cdots \geq \gamma_m$. We use $\vec{v}_{-i} = (v_1, v_2, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n)$ to denote the value vector without v_i.

Agent i has a value of v_i per quality at the point of auction, which is chosen from a distribution F_i by nature. Different from the *traditional Bayesian setting* where the distributions F_is' are known to the auctioneer, here F_i is known to agent i itself but not known to the auctioneer.[1] We use $F = \prod_{i=1}^{n} F_i$ to denote the joint distribution, and $F_{-i} = \prod_{j \neq i} F_j$ to denote the joint distribution without F_i. Also we assume that F_is' are jointly independent distributions.

Prior Information Representation. For a random variable $v \sim G$, we use $G(\cdot)$ to denote the cumulative distribution function, and $g(\cdot)$ to denote the corresponding probability density function. If an agent's value $v \sim G$, we call this v the value representation of the distribution. An alternative representation of prior information is to introduce the concept of *quantile*. Basically, given value distribution G, the quantile $q = 1 - G(v)$ represents the probability that a buyer will have a value at least v. Thus when we use quantile q as the random variable, we map it to the value by the function $v(q) = G^{-1}(1 - q)$. Clearly, $v(q)$ is monotone non-increasing in q. An important property of quantile is that, no matter what the distribution G is, q is always chosen uniformly in $[0, 1]$ interval.

The virtual value representation of a random variable v is defined as follows: take a value v as input, the representation outputs value $v - \frac{1-G(v)}{g(v)}$. Similarly, to represent this mapping from a quantile q to a virtual value representation, we just let the mapping to take value

$$v(q) - \frac{1 - G(v(q))}{g(v(q))} = v(q) - \frac{q}{g(v(q))} = v(q) + qv'(q).$$

An alternative derivation is to define the revenue curve function $R(q) = qv(q)$, then we have:

$$R'(q) = v(q) + qv'(q).$$

[1]We assume that F_i is not known to other bidders, either. As we will discuss later, each bidder will choose a distribution to announce based on what others announce.

So we also denote the virtual value by $\phi(q) = R'(q)$. A distribution is regular if $\phi(q)$ is monotone non-increasing in q, or equivalently, $R(q)$ is concave.

Throughout the paper, we will use the CDF G, the value function $v(\cdot)$, or virtual value function $\phi(\cdot)$ to denote a distribution interchangeably.

2.1 Single-Parameter Environment

For simplicity and generality, we shift from the specific sponsored search setting to the more general *single-parameter environment*. There are n bidders and a set $A \subseteq \mathbb{R}_+^n$ of feasible allocations $\vec{a} = (a_1, \ldots, a_n)$. The allocation to bidder i is a_i, and the allocation to other bidders is denoted by $\vec{a}_{-i}$. Assume that A is *downward-closed*, that is, for each i and any $\vec{a} = (a_i, \vec{a}_{-i}) \in A$, $(0, a_{-i}) \in A$. The sponsored search setting is a special case where a_i stands for the (expected) quality of ad slot given to i.

Each bidder i has a value $v_i \sim F_i$ for receiving one unit of allocation. A mechanism $\mathcal{M}$ associated with the prior F consists of an allocation rule $\vec{X} : \mathbb{R}_+^n \to A$, and a payment rule $\vec{P} : \mathbb{R}_+^n \to \mathbb{R}_+^n$. Bidders submit bids $\vec{b} = (b_i, \vec{b}_{-i})$ to represent their willingness to pay per unit, then bidder i is allocated $X_i(\vec{b})$ units, pays $P_i(\vec{b})$, obtaining a utility of $X_i(\vec{b})v_i - P_i(\vec{b})$. Denote the interim allocation by $x_i(v_i) = \mathbf{E}_{\vec{v}_{-i} \sim F_{-i}}[X_i(v_i, \vec{v}_{-i})]$, and the interim payment by $p_i(v_i) = \mathbf{E}_{\vec{v}_{-i} \sim F_{-i}}[P_i(v_i, \vec{v}_{-i})]$.

We say a mechanism $\mathcal{M}$ is Bayesian-incentive-compatible (BIC), if the following constraints are satisfied:

$$v_i x_i(v_i) - p_i(v_i) \geq v_i x_i(v_i') - p_i(v_i'), \ \forall v_i, v_i', i.$$

We say a mechanism $\mathcal{M}$ is Bayesian-individually-rational (BIR), if the following constraints are satisfied:

$$v_i x_i(v_i) - p_i(v_i) \geq 0, \ \forall v_i, i.$$

In addition, we assume $p_i(0) = 0$.

Lemma 2.1 ([27]). *A mechanism $\mathcal{M}$ is BIC and BIR if the allocation rule and payment rule satisfy:*

- *$X_i(v_i, \vec{v}_{-i})$ is monotone non-decreasing in v_i.*
- *$P_i(\vec{v}) = v_i X_i(\vec{v}) - \int_0^{v_i} X_i(u, \vec{v}_{-i}) \mathrm{d}u$.*

Sometimes we express allocation and payment in the quantile space, and use superscript $(\mathcal{M}, F)$ to indicate that they are determined by mechanism $\mathcal{M}$ with F as prior. Thus we have $x_i^{(\mathcal{M},F)}(q) = x_i(v_i(q))$, $p_i^{(\mathcal{M},F)}(q) = p_i(v_i(q))$, etc. Then the expected payment and the revenue of a BIC mechanism can be conveniently expressed in terms of virtual value $\phi(q)$:

Lemma 2.2 ([27]). *In a BIC and BIR mechanism $\mathcal{M}$, if agents bid their true values, then agent i's expected payment equals to*

$$\mathbf{E}_q[p_i^{(\mathcal{M},F)}(q)] = \mathbf{E}_q[\phi_i(q) x_i^{(\mathcal{M},F)}(q)]. \tag{1}$$

Thus, the revenue is equal to the expected virtual welfare:

$$\sum_{i=1}^{n} \mathbf{E}_{q_i}[p_i^{(\mathcal{M},F)}(q_i)] = \mathbf{E}_{\vec{q}}[\sum_{i=1}^{n} \phi_i(q_i) X_i^{(\mathcal{M},F)}(\vec{q})]. \tag{2}$$

The Myerson's auction. We will show in Section 3.1 that a generalization (hence the abbreviation Mye) of the classical single-item Myerson's optimal auction achieves the maximum revenue under BIC and BIR constraints with respect to the prior F.

The generalized first-price auction. A generalized first-price auction[2] (GFP) is a natural generalization of a first-price auction. The allocation rule is defined by $\vec{X}(\vec{b}) \in \arg\max_{\vec{a} \in A} \sum_{i=1}^{n} b_i a_i$, and the payments are $P_i(\vec{b}) = b_i X_i(\vec{b})$. Contrary to Mye, GFP is prior-independent and not BIC.

2.2 Private Data Manipulation Model

When prior-dependent auctions such as Mye are used in practice, an issue emerges: the *true* prior information F is *not* known to the auctioneer. As an alternative, the auctioneer will observe the possibly manipulated data, reported by bidders. One can interpret the manipulation procedure as follows: given a sample v drawn from F, the manipulation maps it to a bid $\hat{v}$, which follows an underlying distribution $\hat{F}$. With enough data, we assume the auctioneer can reconstruct or approximate well the manipulated distribution $\hat{F}$. Throughout the paper, we will use notions with $\hat{\cdot}$ for the manipulated data, (for example, $\hat{F}$ for the distribution, $\hat{\phi}(\cdot)$ for its virtual value distribution).

For any mechanism $\mathcal{M}$ that takes prior into account, we consider the following two-stage PDM game scenario:

- The auctioneer releases the mechanism $\mathcal{M}$ used in auction before the auction starts.
- Each agent i reports $\hat{F}_i$ by bidding $\hat{v}_i$ such that $\hat{v}_i \sim \hat{F}_i$.
- The auctioneer runs $\mathcal{M}$ with reported manipulated distributions $\hat{F}_i$ as prior information, allocates and charges each agent with respect to the outcome of $\mathcal{M}$.

The two stages are the first and second part of the above scenario. The third part summarizes the outcome in the two-stage PDM game.

In a PDM game, the agents will strategically respond to the mechanism. Their strategy sets are defined as follows:

Definition 2.3 (Strategy Set). *Let $S = \prod_{i=1}^{n} S_i$ be the strategy set of all agents, where S_i is the strategy set of agent i. A strategy $s_i \in S_i$ consists of:*

- *$\hat{v}_i : [0,1] \to \mathbb{R}_+$, the manipulated distribution;*
- *σ_i : a permutation[3] of $[0,1]$ such that agent i bids $\hat{v}_i(q)$ when her true value is $v_i(\sigma_i(q))$. We require that $\sigma_i(\mathrm{Uniform}[0,1]) = \mathrm{Uniform}[0,1]$.*

The introduction of $\sigma_i(\cdot)$ is necessary because $\hat{v}_i(\cdot)$ only defines the overall bid distribution but does not describe how values are mapped to bids, and $\sigma_i(\cdot)$ allows an agent to choose the mapping arbitrarily, while keeping the distribution $\hat{v}_i(\cdot)$ unchanged.

Once the auctioneer has chosen a mechanism $\mathcal{M}$ to release, and every agent has chosen her own manipulation strategy, we can compute each agent i's utility, as follows:

$$U_i(s_1, \ldots, s_n) = \mathbf{E}_q \left[v_i(\sigma_i(q)) \cdot x_i^{(\mathcal{M},\hat{F})}(q) - p_i^{(\mathcal{M},\hat{F})}(q) \right]. \tag{3}$$

We emphasize that $x_i^{(\mathcal{M},\hat{F})}(q)$ is the interim allocation when agent i's bid is $\hat{v}_i(q)$, and others' bids $\hat{\vec{v}}_{-i}$ follows $\hat{F}_{-i}$, under mechanism $\mathcal{M}$ with manipulated distribution $\hat{F} = \prod_{i=1}^{n} \hat{F}_i$ as input.

[2] Usually, "generalized-first-price auction" refers specifically to the auction in the sponsored search setting. Here we use the term in a more general sense.

[3] Similarly we can define random permutation where a true value can be mapped to a distribution among bids, although it can be seen from Lemma 3.1 that the randomization in unnecessary.

In such a scenario of manipulation against private data elicitation, we assume that agents are strategic and they will end up playing a Nash equilibrium, which is defined as follows:

Definition 2.4 (Nash Equilibrium). *We say a manipulation $\vec{s} = (s_1, s_2, \ldots, s_n)$ is a Nash equilibrium of mechanism $\mathcal{M}$ under the Private Data Manipulation model if:*

$$U_i(s_i, s_{-i}) \geq U_i(t_i, s_{-i})$$

for any strategy $t_i \in S_i$.

We show in the following lemma that for *prior independent auctions*, in which the allocation and payment of $\mathcal{M}$ do not depend on input distribution $\hat{F}$, the behaviors of bidders under the PDM setting and the traditional Bayesian setting are strategically equivalent. In the traditional Bayesian setting, the strategy of each bidder is a value-to-bid mapping $\pi_i : \mathbb{R}_+ \to \mathbb{R}_+$. Let $\pi_i^s : \mathbb{R}_+ \to \mathbb{R}_+$ denote the value-to-bid mapping induced by the strategy s_i in the PDM game (i bids $\pi_i^s(v_i(q)) = \hat{v}_i(\sigma_i^{-1}(q))$ when her true value is $v_i(q)$).

Lemma 2.5. *For any prior independent auctions $\mathcal{M}$, $(s_1, ..., s_n)$ is a Nash equilibrium of $\mathcal{M}$ under PDM if and only if $(\pi_1^s, ..., \pi_n^s)$ is a Bayes-Nash equilibrium of $\mathcal{M}$ under the traditional Bayesian setting.*

This lemma holds since manipulating distributions does not affect the allocation and payment rules of $\mathcal{M}$.

3 EQUIVALENCE BETWEEN AUCTIONS IN PRIVATE DATA MANIPULATION MODEL

In this section, we show that Myerson's optimal auction and the generalized first-price auction are equivalent when agents are able to manipulate their value distributions. We consider Myerson's auction in PDM because it is natural for the auctioneer to choose it, given the reported value distributions, since it extracts the most revenue based on all available information. However, it turns out that the originally optimal auction degenerates into a generalized first-price auction in which no private data is used at all in a PDM game. We start by reviewing how the optimal auction works.

3.1 The Myerson's Auction

The Myerson's optimal auction Mye in a single-parameter environment is a generalization of Myerson's single-item auction in [34]. Denote the given prior by F.

If F_i's are regular, by Lemma 2.2, the revenue can be maximized by choosing an allocation that maximizes the virtual welfare for each bid vector $\vec{v}(\vec{q})$. Formally,

$$\vec{X}^{(\text{Mye},F)}(\vec{q}) \in \arg\max_{\vec{a} \in A} \sum_{i=1}^{n} a_i \phi_i(q_i).$$

Payments are calculated according to Lemma 2.1. Such a mechanism is BIC by Lemma 2.1, because when v_i increases, its virtual value $\phi(q_i(v_i))$ does not decrease by definition, then the allocation X_i is monotone non-decreasing.

In addition, in order to maximize the virtual welfare, Mye will never allocate positive units to a bidder with non-positive virtual value; otherwise, setting her allocation to be zero weakly increases the virtual welfare, without violating the feasibility constraint since A is downward-closed. That is, we set a reserve price r_i at which the virtual value of bidder i is zero, and only allocate to agents whose bids are above their respective reserve prices. We choose the maximum reserve price r_i when there are multiple solutions to $\phi_i^v(r_i) = 0$. As a result,

$$\phi_i(q_i) \leq 0 \implies X_i^{(\text{Mye},F)}(q_i, \vec{q}_{-i}) = 0, \ \forall \vec{q}_{-i}. \quad (4)$$

For irregular distribution F_i, Myerson's auction will first "iron" the distribution [34], transforming it into a regular distribution $\overline{F}_i$, and then run Mye with respect to $\overline{F}_i$.

3.2 Equivalence between Mye and GFP

Now we are ready to analyze how agents will behave under the Nash equilibrium of Mye under PDM. Recall that the true distribution F is no longer available, and Mye has to rely on agents' strategic distribution $\hat{F}$ to allocate items and charge prices.

First we argue that without loss of generality, the set of best responses can be limited as follows:

Lemma 3.1. *For any s_{-i}, choosing the strategy $s_i = (\hat{v}_i(\cdot), \sigma_i)$ that satisfy the following properties maximizes agent i's utility:*

1. *Identity permutation: $\sigma_i(q) = q$.*
2. *Monotonicity (of virtual value): $\forall q_1 < q_2$, $\hat{\phi}_i(q_1) \geq \hat{\phi}_i(q_2)$.*
3. *Non-negative virtual value: $\hat{\phi}_i(q) \geq 0$.*

Proof of Property 1 and 2. For 1, since the payment term in (3) does not depend on σ_i, to maximize utility we only need to maximize the first term, which is $\int_0^1 \left[v_i(\sigma_i(q)) \cdot x_i^{(\text{Mye},\hat{F})}(q) \right] \mathrm{d}q$ in integral form.

According to the allocation rule of Mye, $x_i^{(\text{Mye},\hat{F})}(q)$ is monotone non-increasing in q, so by the rearrange inequality, we have:

$$\int_0^1 \left[v_i(\sigma_i(q)) \cdot x_i^{(\text{Mye},\hat{F})}(q) \right] \mathrm{d}q \leq \int_0^1 \left[v_i(q) \cdot x_i^{(\text{Mye},\hat{F})}(q) \right] \mathrm{d}q.$$

Thus $\sigma_i(q) = q$ maximizes expected value.

For 2, if $\hat{\phi}_i$ is regular, then the monotonicity of $\hat{\phi}_i$ comes from the regularity of $\hat{F}_i$. If $\hat{\phi}_i$ is irregular, recall that the ironing procedure will first "iron" $\hat{F}_i$ to regular distribution $\overline{F}_i$, and then run the Myerson's auction. Note that here the virtual value function for agent i is no longer defined on $\hat{F}_i$, but defined on $\overline{F}_i$ which is a regular distribution. Thus the monotonicity property still holds. □

Before proving property 3, we characterize the strategy set and simplify some notations. By 1, we can omit σ_i and assume that s_i is uniquely determined by $\hat{v}_i(\cdot)$, that is, bidders will always bid $\hat{v}_i(q)$ for value $v_i(q)$. Then in the opposite direction to property 2, we will show that each monotone virtual value function $\hat{\phi}_i(\cdot)$ determines a unique valid strategy, as follows:

Lemma 3.2. *For any non-increasing function $\hat{\phi} : [0, 1] \to \mathbb{R}$, there exists a unique distribution $\hat{v}(\cdot)$ whose virtual value function is $\hat{\phi}(\cdot)$.*

Proof. Since $\hat{\phi}$ is Riemann integrable, we can reconstruct its revenue curve by integrating: $q\hat{v}(q) = \int_{x=0}^{q} \hat{\phi}(x)\mathrm{d}x$. Dividing q gives the desired $\hat{v}(\cdot)$.[4] □

[4]Although $[q\hat{v}(q)]'$ is undefined at the discontinuous point of $\hat{\phi}(q)$, we can define it to be $\hat{\phi}(q)$ because discontinuous points have measure zero and have no effect on the expected utility.

Consequently, no two different strategies s_i, s_i' have manipulation distributions with the same virtual value, given property 1. Now we can use $\hat{\phi}_i$ solely to represent a strategy. Lemma 3.2 also implies that "strategy $\hat{\phi}_i$" always exists as long as $\hat{\phi}_i$ is monotone. Using Lemma 2.2 to compute the expected payment, the expected utility of agent i becomes:

$$U_i(\hat{\phi}_i, \hat{\phi}_{-i}) = \mathbf{E}_q \left[x_i^{(\text{Mye},\hat{F})}(q) \left(v_i(q) - \hat{\phi}_i(q) \right) \right]. \quad (5)$$

Proof of Property 3. Assume that $\hat{\phi}_i(\cdot)$ is monotone non-increasing. Suppose for some strategy, $\hat{\phi}_i(\cdot)$ takes a negative value in $[w_i, 1]$ (or $(w_i, 1]$), then we define another virtual value function $\widetilde{\phi}_i(\cdot)$, which takes the following value:

$$\widetilde{\phi}_i(q) = \begin{cases} \hat{\phi}_i(q) & \text{if } q < w_i, \\ \hat{\phi}_i(q) \cdot \mathbb{I}[\hat{\phi}_i(q) > 0] & \text{if } q = w_i, \\ 0 & \text{if } q > w_i. \end{cases}$$

The new strategy $\widetilde{\phi}_i(\cdot)$ is valid because $\widetilde{\phi}_i(\cdot)$ is also monotone.

If $\hat{\phi}_i(q_i) \le 0$, then both $X_i^{(\text{Mye},\hat{F})}(\vec{q})$ and $X_i^{(\text{Mye},\widetilde{F}_i \times \hat{F}_{-i})}(\vec{q})$ are 0 according to the allocation rule of Mye (by (4)). The two allocations are also the same for all agents when $\hat{\phi}_i(q) = \widetilde{\phi}_i(q) > 0$. This implies that no agent gets an altered allocation, thus they obtain unchanged expected utility. This concludes the proof. □

Thus we can assume without loss of generality that $\hat{\phi}_i(\cdot)$ always takes a non-negative value, which enables us to relate it with the bidding strategy in a generalized first-price auction, as discussed next.

Recall the allocation rule $X_i^{\text{GFP}}(\vec{b})$ in Section 2.1 for the generalized first-price auction, which maximizes $\sum_{i=1}^n a_i b_i$. Let

$$x_i^{\text{GFP}}(b_i(q_i)) = \mathbf{E}_{\vec{q}_{-i}} X_i^{\text{GFP}} \left(\vec{b}_i(\vec{q}_i), \vec{b}_{-i}(q_{-i}) \right)$$

denote the corresponding interim allocation.

Then we write down the utility of agent i in GFP. Since GFP is not BIC, we use a quantile-to-bid mapping $b_i : [0,1] \to \mathbb{R}_+$ to represent bidder i's strategy. With value $v_i(q_i)$, agent i bids $b_i(q_i)$, obtaining an allocation of $X_i\left(b_i(q_i), \vec{b}_{-i}(\vec{q}_{-i})\right)$ units, paying $b_i(q_i)$ for each unit. Thus the expected utility of agent i is:

$$\begin{aligned} U_i^{\text{GFP}}(b_i, b_{-i}) &= \mathbf{E}_{q_i, \vec{q}_{-i}} \left[X_i^{\text{GFP}} \left(b_i(q_i), \vec{b}_{-i}(\vec{q}_{-i}) \right) (v_i(q_i) - b_i(q_i)) \right] \\ &= \mathbf{E}_{q_i} \left[x_i^{\text{GFP}}(b_i(q_i))(v_i(q_i) - b_i(q_i)) \right]. \quad (6) \end{aligned}$$

A key observation is: the expected utility (5) of agent i in a PDM with Mye can be written exactly in the same way:

$$U_i^{\text{Mye}}(\hat{\phi}_i, \hat{\phi}_{-i}) = \mathbf{E}_{q_i} \left[x_i^{\text{GFP}}(\hat{\phi}_i(q_i)) \left(v_i(q_i) - \hat{\phi}_i(q_i) \right) \right]. \quad (7)$$

This is because $X_i^{(\text{Mye},\hat{F})}(\vec{q}) = X_i^{\text{GFP}}(\hat{\phi}(\vec{q}))$ by definition[5].

We are ready to show our main theorem. We say a strategy s_i of PDM is *normal* if it satisfies the three properties in Lemma 3.1 (identity permutation, monotone and non-negative virtual value).

[5]Assume that ties in GFP are broken in the same way as ties are broken in the virtual welfare maximization of Mye.

Theorem 3.3 (**Main Theorem**). *Myerson's auction under PDM is equivalent to the generalized first-price auction. Specifically, there is a bijection between all normal strategy vectors $\vec{s} = (s_i, s_{-i})$ of PDM and all non-increasing strategy vectors $\vec{b} = (b_i, b_{-i})$ of GFP, such that $\vec{s}$ and $\vec{b}$ produce the same expected utility and expected payment for each bidder, and the same revenue for the auctioneer.*

Proof. We can equate s_i and $\hat{\phi}_i$ because: each normal strategy s_i has a non-increasing and non-negative virtual value function $\hat{\phi}_i$; and by Lemma 3.2, each non-increasing and non-negative $\hat{\phi}_i$ determines a unique normal strategy s_i. Thus by setting $\hat{\phi}_i(\cdot) = b_i(\cdot)$ for each i, the bijection between normal $\vec{s}$ and monotone $\vec{b}$ is established.

Then by the above observation (6) and (7), we have:

$$U_i^{\text{GFP}}(b_i, b_{-i}) = U_i^{\text{Mye}}(\hat{\phi}_i, \hat{\phi}_{-i}).$$

Furthermore, for any $\vec{v}(\vec{q})$, we have the same allocation in the two auctions: $X_i^{(\text{Mye},\hat{F})}(\vec{q}) = X_i^{\text{GFP}}(\vec{b}(\vec{q}))$. By subtracting the value term from the expected utility, we conclude that the expected payment and revenue are the same. □

Corollary 3.4. *A normal strategy vector $\vec{s}$ is a Nash equilibrium of Mye under PDM if and only if the corresponding monotone strategy vector $\vec{b}$ is a Bayes-Nash equilibrium of GFP.*

4 EQUIVALENCE IN SPONSORED SEARCH AUCTIONS

In this section, we focus our attention on the sponsored search auction. We first introduce some necessary notions.

Let b_i be the bid of willingness to pay per quality by agent i. Let $\mathbf{X}_{ij}(\vec{b})$ denote the probability that item j is allocated to agent i, and let $P_i(\vec{b})$ denote the payment for agent i given bidding vector $\vec{b}$. We use $\mathcal{M} = (\mathbf{X}, \vec{P})$ to denote a mechanism $\mathcal{M}$ that consists of an allocation matrix and payment vector (both as functions of bid vector), where $\mathbf{X} = (X_{ij})_{n \times m}$ and $\vec{P} = (P_1, P_2, \ldots, P_n)$. The allocation vector to agent i is $\vec{X}_{i\cdot}$. We assume here the value across items is additive. Let $X_i = \sum_j X_{ij}\gamma_j$ be the total quality allocated to bidder i. Let $v_i(\mathbf{X}) = v(\vec{X}_{i\cdot}) = \sum_j X_{ij} \cdot v_i \gamma_j = v_i \cdot \sum_j X_{ij}\gamma_j = v_i X_i$. Basically, $v_i(\mathbf{X})$ is the total value of i in the auction outcome specified by $\mathbf{X}, \vec{P}$.

In a sponsored search auction, we say the auction is feasible, if each bidder is allocated with at most one ad slot, and each slot is allocated to at most one bidder. This can be interpreted by the following constraints:

$$\text{(unit demand)} \sum_j X_{ij}(\vec{b}) \le 1 \qquad \forall \vec{b}, i$$

$$\text{(avoid over allocation)} \sum_i X_{ij}(\vec{b}) \le 1 \qquad \forall \vec{b}, i$$

$$\text{(Indivisible ad slots)} \; X_{ij}(\vec{b}) \in \{0,1\} \qquad \forall \vec{b}, i, j$$

We will call them *feasibility constraints* for the rest of the paper.

It is not hard to verify that the sponsored search auction scenario defined above is a single-parameter setting: $X_i(\vec{b}) = \sum_j X_{ij}\gamma_j$ is the equivalent allocation rule based on quality; the feasibility constraints form a feasible domain that is downward-closed. Thus the optimal auction in sponsored search scenario can be derived naturally: bidders with the i-th highest positive ironed virtual value gets the i-th ad slot (one can easily verify its monotonicity), and payments follow from the Myerson style payment rule in Lemma 2.1.

4.1 Prior-Independent Auctions

Here we introduce some prior-independent sponsored search auctions that are widely used in practice. Sort the bids of agents non-increasingly, $b_1 \geq b_2 \geq \cdots \geq b_n$. Let $\gamma_1 \geq \gamma_2 \geq \cdots \geq \gamma_m$ denote the qualities of ad slots.

Generalized first-price auction. In a generalized first-price auction, each bidder $i \in [m] = \{1, \ldots, m\}$ gets slot i, pays $b_i \cdot \gamma_i$.

Generalized second-price auction. In a generalized second-price auction, each bidder $i \in [m]$ gets slot i, pays $b_{i+1} \cdot \gamma_i$.

VCG auction. In a VCG auction, each bidder $i \in [m]$ gets slot i, pays $\sum_{j=i+1}^{m+1} b_j \cdot (\gamma_{j-1} - \gamma_j)$. (Here we denote $\gamma_{m+1} = 0$.)

4.2 PDM Game in Sponsored Search Auction

Now we focus on PDM game played in sponsored search auction, and study the revenue equivalence phenomenon concerning various sponsored search auctions: GSP, VCG, GFP, Mye.

We say a Bayes-Nash equilibrium (BNE) is efficient if for any value vector $\vec{v}$ with corresponding bid vector $\vec{b} = \vec{b}(\vec{v})$, the resulting allocation $\mathbf{X}(\vec{b})$ is feasible and maximizes the social welfare $\sum_i v_i X_i$. When buyers have i.i.d. value distribution, we say a BNE is symmetric if buyers use a same bidding function. We first introduce some useful lemmas.

Lemma 4.1 ([15]). *In a generalized first-price auction, when there are n agents with i.i.d. continuous distribution on values, there is only one Bayes-Nash equilibrium that is symmetric and efficient.*

Lemma 4.2 ([24]). *In a generalized second-price auction, when there are n agents with i.i.d. continuous distribution on values, if symmetric Bayes-Nash equilibrium exists, then it is efficient.*

Lemma 4.3 ([37]). *For any two mechanisms that gives efficient Bayes-Nash equilibrium, if for some value $\vec{v}^0$, the expected payment for each agent is the same in the two mechanisms, then the two mechanisms have the same revenue.*

We are now ready to show our equivalence results.

First, since the sponsored search auction in this paper is single-parameter, by Theorem 3.3 we immediately conclude that:

Theorem 4.4. *In a PDM game, Mye is revenue-equivalent to GFP under sponsored search auction scenario.*

We then show that when buyers have i.i.d. values, GFP, and VCG are equivalent while GSP are equivalent to them in some cases.

Theorem 4.5. *When buyers have i.i.d. values, Mye, GFP, and VCG are revenue-equivalent in their respective equilibria under PDM.*

Proof. We just need to show the equivalence between GFP and VCG. By Lemma 2.5, we know that the equilibrium of GFP under PDM is exactly the same under traditional Bayesian setting. By Lemma 4.1 we know that GFP only has symmetric and efficient equilibrium. By Lemma 4.3 we know that GFP is revenue-equivalent to VCG. This concludes the proof. □

However, GFP and VCG may not be equivalent if buyers have independent but non-identical distributions.

Example 4.6 (Example 2 in [33]). Suppose there is one item and two buyers. Buyer 1 has a uniformly random value v_1 on $[0, \frac{1}{1+z}]$, buyer 2 has a uniformly random value v_2 on $U[0, \frac{1}{1-z}]$, for $z \geq 0$. In equilibrium, the bidders' inverse bidding functions are:

$$v_1 = b_1^{-1}(b) = \frac{2b}{1 + z(2b)^2}, \qquad v_2 = b_2^{-1}(b) = \frac{2b}{1 - z(2b)^2}.$$

The CDF $G_{FP}(b)$ of the larger bid is:

$$G_{FP}(b) = \Pr[b_1(v_1) \leq b] \cdot \Pr[b_2(v_2) \leq b] = \frac{(1 - z^2)(2b)^2}{1 - z^2(2b)^4},$$

which is decreasing in z. Thus the revenue of the first-price auction is increasing in z. However, for the second-price auction, the CDF $G_{SP}(b)$ of the second value is:

$$G_{SP}(b) = 1 - \Pr[v_1 > b] \cdot \Pr[v_2 > b] = 2b - (1 - z^2)b^2,$$

which increases with z. Since the two auctions have the same revenue when $z = 0$ (i.i.d. case), the revenue of the first-price auction is strictly better than that of the second-price auction if $z > 0$.

Theorem 4.7. *When the i.i.d. value distribution of buyers and the qualities of slots satisfy certain conditions such that GSP has a symmetric BNE, Mye, GFP, VCG, and GSP are revenue-equivalent in their respective symmetric equilibria under PDM.*

Proof. We just need to show the equivalence between GSP and VCG. By Lemma 2.5, we know that the equilibrium of GSP under PDM is the same as the traditional Bayes-Nash equilibrium. By Lemma 4.2, the symmetric equilibrium of GSP must be efficient. By Lemma 4.3 we know that GSP is revenue-equivalent to VCG. □

Finally, we demonstrate by the following two examples that GSP may not have symmetric BNE, and may admit asymmetric BNE, even if buyers have i.i.d. values. Our equivalence results concerning GSP do not hold in these two situations.

Example 4.8 (Example 1 in [24]). Consider three buyers with i.i.d. values from Uniform$[0, 1]$ and two slots with qualities $(1, \gamma_2)$. A symmetric equilibrium exists if and only if $\gamma_2 \leq 0.75$.

Example 4.9 (Section 3.1 in [32]). Consider three buyers with i.i.d. values from Uniform$[1, 2]$ and three slots with qualities $(1, \frac{1}{2}, \frac{1}{2})$. There is an asymmetric equilibrium: $b_1(v_1) = v_1$, $b_2(v_2) = b_3(v_3) = 0$. Clearly buyer 1 is in equilibrium. To see that buyer $i = 2, 3$ are in equilibrium, suppose i has valuation $v_i > 0$. If i bids $b \leq 1$, then her utility is $\frac{1}{2}v_i - 0$, and if she bids $b > 1$, her utility would be:

$$\begin{aligned} &v_i \Pr[v_1 \leq b] + \frac{1}{2}v_i \Pr[v_1 > b] - \int_1^b v_1 \mathrm{d}v_1 \\ =\ &\frac{1}{2}bv_i - \frac{b^2 - 1}{2} \leq \frac{1}{2}v_i. \end{aligned}$$

ACKNOWLEDGMENTS

This work is supported by Science and Technology Innovation 2030 - "New Generation Artificial Intelligence" Major Project (No. 2018AAA0100901) and by National Natural Science Foundation of China (No. 61761146005, 61632017).

REFERENCES

[1] Kareem Amin, Afshin Rostamizadeh, and Umar Syed. 2013. Learning prices for repeated auctions with strategic buyers. In *Advances in Neural Information Processing Systems*. 1169–1177.

[2] Kareem Amin, Afshin Rostamizadeh, and Umar Syed. 2014. Repeated contextual auctions with strategic buyers. In *Advances in Neural Information Processing Systems*. 622–630.

[3] Susan Athey and Ilya Segal. 2013. An efficient dynamic mechanism. *Econometrica* 81, 6 (2013), 2463–2485.

[4] Moshe Babaioff, Robert D Kleinberg, and Aleksandrs Slivkins. 2015. Truthful mechanisms with implicit payment computation. *Journal of the ACM (JACM)* 62, 2 (2015), 1–37.

[5] Moshe Babaioff, Yogeshwer Sharma, and Aleksandrs Slivkins. 2014. Characterizing truthful multi-armed bandit mechanisms. *SIAM J. Comput.* 43, 1 (2014), 194–230.

[6] Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik. 2018. A general theory of sample complexity for multi-item profit maximization. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. ACM, 173–174.

[7] Maria-Florina F Balcan, Tuomas Sandholm, and Ellen Vitercik. 2016. Sample complexity of automated mechanism design. In *Advances in Neural Information Processing Systems*. 2083–2091.

[8] Andrew Beattie. 2018. *The Story Behind Google's Success*. https://www.investopedia.com/articles/personal-finance/042415/story-behind-googles-success.asp

[9] Dirk Bergemann and Juuso Välimäki. 2010. The dynamic pivot mechanism. *Econometrica* 78, 2 (2010), 771–789.

[10] Mark Braverman, Jieming Mao, Jon Schneider, and Matt Weinberg. 2018. Selling to a no-regret buyer. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. 523–538.

[11] Tian-Ming Bu, Xiaotie Deng, and Qi Qi. 2008. Forward looking Nash equilibrium for keyword auction. *Inform. Process. Lett.* 105, 2 (2008), 41–46.

[12] Tian-Ming Bu, Xiaotie Deng, and Qi Qi. 2012. Multi-bidding strategy in sponsored search auctions. *Journal of combinatorial optimization* 23, 3 (2012), 356–372.

[13] Yang Cai and Constantinos Daskalakis. 2017. Learning multi-item auctions with (or without) samples. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 516–527.

[14] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. 2011. On the efficiency of equilibria in generalized second price auctions. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, 81–90.

[15] Shuchi Chawla and Jason D Hartline. 2013. Auctions with unique equilibria. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 181–196.

[16] Edward H Clarke. 1971. Multipart pricing of public goods. *Public choice* 11, 1 (1971), 17–33.

[17] Richard Cole and Tim Roughgarden. 2014. The sample complexity of revenue maximization. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*. ACM, 243–252.

[18] Xiaotie Deng, Tao Xiao, and Keyu Zhu. 2017. Learn to play maximum revenue auction. *IEEE Transactions on Cloud Computing. Early access, June 05* (2017).

[19] Xiaotie Deng and Keyu Zhu. 2019. On Bayesian Epistemology of Myerson Auction. *IEEE Transactions on Cloud Computing. Early access, Jan 31* (2019).

[20] Nikhil R Devanur and Sham M Kakade. 2009. The price of truthfulness for pay-per-click auctions. In *Proceedings of the 10th ACM conference on Electronic commerce*. 99–106.

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

[22] Benjamin Edelman, 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 (2007), 242–259.

[23] Benjamin Edelman and Michael Schwarz. 2010. Optimal auction design and equilibrium selection in sponsored search auctions. *American Economic Review* 100, 2 (2010), 597–602.

[24] Renato Gomes and Kane Sweeney. 2014. Bayes–nash equilibria of the generalized second-price auction. *Games and economic behavior* 86 (2014), 421–437.

[25] Theodore Groves. 1973. Incentives in Teams. *Econometrica* 41, 4 (1973), 617–31.

[26] Chenghao Guo, Zhiyi Huang, and Xinzhi Zhang. 2019. Settling the sample complexity of single-parameter revenue maximization. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 662–673.

[27] Jason D Hartline. 2013. Mechanism design and approximation, chapter 3. *Book draft. October* 122 (2013).

[28] Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. 2018. Making the most of your samples. *SIAM J. Comput.* 47, 3 (2018), 651–674.

[29] S Lagaie, DMPA Saberi, and RV Vohra. 2007. Algorithmic Game Theory, chapter Sponsored Search Auctions.

[30] Renato Paes Leme and Eva Tardos. 2010. Pure and Bayes-Nash price of anarchy for generalized second price auction. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. IEEE, 735–744.

[31] Brendan Lucier and Renato Paes Leme. 2011. GSP auctions with correlated types. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, 71–80.

[32] Brendan Lucier, Renato Paes Leme, and Eva Tardos. 2012. On revenue in the generalized second price auction. In *Proceedings of the 21st international conference on World Wide Web*. ACM, 361–370.

[33] Eric Maskin and John Riley. 2000. Asymmetric auctions. *The Review of Economic Studies* 67, 3 (2000), 413–438.

[34] Roger B Myerson. 1981. Optimal auction design. *Mathematics of operations research* 6, 1 (1981), 58–73.

[35] Hamid Nazerzadeh, Amin Saberi, and Rakesh Vohra. 2008. Dynamic cost-per-action mechanisms and applications to online advertising. In *Proceedings of the 17th international conference on World Wide Web*. 179–188.

[36] Thomas Nedelec, Noureddine El Karoui, and Vianney Perchet. 2019. Learning to Bid in Revenue Maximizing Auction. In *Companion Proceedings of The 2019 World Wide Web Conference*. ACM, 934–935.

[37] Noam Nisan. 2007. Algorithmic Game Theory, chapter 9.

[38] Ronan Shields. 2019. *After Months of Preparation, Google Set to Roll Out First-Price Auctions*. https://www.adweek.com/programmatic/after-months-of-preparation-google-rolls-out-first-price-auctions/

[39] Vasilis Syrgkanis. 2017. A sample complexity measure with applications to learning optimal auctions. In *Advances in Neural Information Processing Systems*. 5352–5359.

[40] Pingzhong Tang and Yulong Zeng. 2018. The price of prior dependence in auctions. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. ACM, 485–502.

[41] Hal R Varian. 2007. Position auctions. *international Journal of industrial Organization* 25, 6 (2007), 1163–1178.

[42] Hal R Varian. 2009. Online ad auctions. *American Economic Review* 99, 2 (2009), 430–34.

[43] William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance* 16, 1 (1961), 8–37.