Balancing Revenue and Privacy with Signaling Schemes in Online Ad Auctions

Hongtao Liu
Gaoling School of Artificial Intelligence
Renmin University of China
Beijing, China
ht6@ruc.edu.cn

Luxi Chen
Gaoling School of Artificial Intelligence
Renmin University of China
Beijing, China
clx1489@ruc.edu.cn

Yiming Ding
Kuaishou Technology
Beijing, China
dingyiming05@kuaishou.com

Changcheng Li
Kuaishou Technology
Beijing, China
lichangcheng@kuaishou.com

Han Li
Kuaishou Technology
Beijing, China
lihan08@kuaishou.com

Peng Jiang
Kuaishou Technology
Beijing, China
jiangpeng@kuaishou.com

Weiran Shen*
Gaoling School of Artificial Intelligence
Renmin University of China
Beijing, China
shenweiran@ruc.edu.cn

Abstract

In online ad auctions, when an Internet user's certain actions trigger an auction, the auctioneer (the platform) usually sends the information about the user to help the buyers better estimate their valuations. However, by strategically revealing only partial information, we cannot only improve the revenue of the auction, but also help protect the privacy of the user.

In this paper, we propose a privacy measure in the online ad auction setting, and seek to maximize a convex combination of revenue and privacy. We formulate the problem as a convex optimization program and derive structural results and properties of the program. We prove that any combination coefficient achieves a certain fraction of the optimal revenue gain and privacy gain, and that we can trade-off between revenue and privacy by simply tuning the combination coefficient. We also show that the gap between the optimal revenue and the revenue achieved by revealing no information can be bounded by a certain valuation discrepancy between the buyers. We also conduct extensive experiments (on both synthetic and real data) to show the effectiveness of our method.

*Corresponding author.

WSDM '25, March 10–14, 2025, Hannover, Germany.

ACM ISBN 979-8-4007-1329-3/25/03
https://doi.org/10.1145/3701551.3703529

CCS Concepts

• **Security and privacy → Privacy protections**; • **Information systems → Computational advertising**; • **Theory of computation → Algorithmic game theory and mechanism design**.

Keywords

online advertising auctions; privacy protection; f-divergence; signaling scheme.

ACM Reference Format:
Hongtao Liu, Luxi Chen, Yiming Ding, Changcheng Li, Han Li, Peng Jiang, and Weiran Shen. 2025. Balancing Revenue and Privacy with Signaling Schemes in Online Ad Auctions. In *Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining (WSDM '25), March 10–14, 2025, Hannover, Germany*. ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3701551.3703529

1 Introduction

Online advertising has been one of the most important revenue sources of most major Internet companies. The total online ad revenue in the US hit $209.7 billion in 2022[1], and is still rapidly growing. Advertising platforms have developed different types of ads. For example, when an Internet user searches a keyword in a search engine, the resulting page usually includes both sponsored results (ads) and organic results. The space containing sponsored results is divided into several slots, where each slot can be filled with an ad. The seller (the search engine) sells the opportunity for displaying ads through auctions, and interested advertisers compete to win the slots. Another example is ad exchanges, where the platform does not provide slots, but sells slots provided by

[1] https://www.statista.com/statistics/183816/us-online-advertising-revenue-since-2000/

publishers to potential buyers. These publishers are usually third-party websites that have a reserved slot to display ads.

Despite differences in advertising types, these ad platforms share a similar advertising process. An auction is usually triggered by the Internet users' certain actions (e.g., searching a keyword, visiting a webpage). The seller then sends information about the user to the buyers to help them better evaluate this ad display opportunity. In such a process, the platform may share sensitive information (age group, browsing history, cookies, IP address, device ID, etc.) about the user with the advertisers. Even worse, both ad platforms and advertisers have networks that uses technologies such as "cookie syncing" to help one another better identify the Internet user[2]. The use of such technologies has become more and more common these days and may explain why the ads for the same product appear mysteriously on different websites. These privacy issues have led to complaints and lawsuits filed against major ad auction platforms over the years[3].

To tackle this problem, an intuitive idea is to apply privacy preserving concepts and techniques, such as differential privacy [11, 12] and k-anonymity [29]. These techniques basically add noise to the original data so that data receivers can only obtain inaccurate but still useful data. However, they have not been widely integrated into the ad auction process for various reasons. For example, the diversity of Internet users requires very large noise to protect their privacy, which may result in the advertisers getting useless information.

In this paper, we propose to use the so-called "Bayesian persuasion" [22] to protect the privacy of the users. The intuition behind Bayesian persuasion is to send randomized signals to the advertisers based on the actual Internet user. Since the signal is randomized, the advertisers can only get a posterior belief about the user's identity. Hence it is also called a "signaling scheme". The idea of applying signaling schemes in ad auctions is not new [14]. However, to the best of our knowledge, we are the first to consider how to protect privacy with signaling schemes in online ad auctions.

The application of signaling schemes requires a prior belief about the Internet users. Such a prior is quite easy to obtain in online ad auctions as each advertiser participates in many auctions every day. However, both differential privacy and k-anonymity do not consider priors. Compared with differential privacy and k-anonymity, signaling schemes can utilize more information and provide us with fine-grained control over the magnitude of noise added to different locations of the original information space. We aim to optimize the platform's revenue without sacrificing too much privacy. We formulate the problem as a mathematical program. It turns out that the problem is equivalent to optimizing a convex combination of both revenue and privacy.

1.1 Our Contributions

We make the following contributions in this paper:

- We propose to use signaling schemes to protect users' privacy in online ad auctions and formulate the problem as a mathematical program;
- We provide structural results of the optimal signaling scheme and derive theoretic guarantees in terms of revenue gain and privacy gain;
- We conduct extensive experiments to show the effectiveness of our method on both synthetic and real data sets.

1.2 Related Work

Researches on privacy issues in online ad auctions began more than a decade ago. Evans [15] study how private information can benefit advertisers' ad targeting. Goldfarb and Tucker [17], Johnson [20] show that with privacy regulations, ad auctions are less effective in terms of targeting and revenue generation. However, it is shown that narrow ad targeting may reduce competition, and thus harms revenue [3, 4, 7, 25]. This is also reflected in our paper. Baltrusaitis et al. [2] investigate the influence of differential privacy on estimating conversion rates in ad auctions.

Another closely related topic is the so-called "Bayesian persuasion". This concept was first proposed by Kamenica and Gentzkow [22], and later followed by many [6, 8, 18, 30, 32]. More recently, Bergemann et al. [5], Junjie et al. [21] apply Bayesian persuasion to click-through rate (CTR) estimations. The closest to ours is [14]. They consider how to optimize revenue by revealing partial information without taking privacy issues into consideration.

Our paper is also related to entropic regularizations in optimal transport [9, 13, 23] and machine learning [24, 26] problems. For example, Cuturi [9] provides the Sinkhorn algorithm that can solve the entropic optimal transport problem very efficiently. Eckstein and Nutz [13] studies the stability problem of the entropic optimal transport. Recently, entropic regularization has also been applied to reinforcement learning [16, 31]. Most of these studies apply entropic regularizations to obtain an approximate and fast solution, while in our paper, the regularization term is used as a privacy measure.

2 Preliminaries

Consider an ad auction model with one ad slot for sale. Let N be the set of n buyers, i.e., $|N| = n > 1$. In real-time bidding (RTB) actions, when an Internet user searches a keyword in a search engine or visits a certain webpage, the seller (e.g., a search engine or an ad exchange platform) sends information about the user characteristics to the buyers. And the buyers then make use of such information to determine their bids and send them back to the seller.

Let U be the set of possible users with $|U| = m$. The user $u_j \in U$ is clearly a random variable to the buyer before receiving any information from the seller. Suppose that u_j follows a publicly known distribution $q(u_j)$. Without loss of generality, we assume $q(u_j) > 0$ since we can safely ignore u_j if $q(u_j) = 0$. Let v_{ij} and b_{ij} be buyer i's valuation and bid when u_j is sent to the buyer. Denote by c_{ij} the probability that user u_j clicks on buyer i's ad if it is shown to the user, i.e., the click-through rate (CTR). The seller ranks the buyers by $b_{ij}c_{ij}$ and uses the second-price auction to sell the ad slot, i.e., the winner is

$$i^* = \arg\max_i \{b_{ij}c_{ij}\}.$$

[2] https://clearcode.cc/blog/cookie-syncing/
[3] https://www.reuters.com/legal/litigation/google-privacy-lawsuit-over-ad-bidding-process-go-forward-2022-06-14/

The winner only pays when the user clicks on the ad, and the payment is:

$$p_{i^*} = \frac{v_{i_2 j} c_{i_2 j}}{c_{i^* j}},$$

where $i_2 = \arg\max_{i \in N \setminus \{i^*\}} \{b_{ij} c_{ij}\}$. Since the second-price auction is truthful, all the buyers will use their valuations as bids. Instead of sending u_j directly to the buyer, the seller can choose to reveal only partial information about the user. Following the so-called "Bayesian persuasion" model [22], we assume that the seller can design a "signaling scheme" as a means of revealing partial information. Specifically, let S be the set of possible signals that the seller can send to the buyers. Let $\pi(s_k|u_j)$ be the probability of sending signal s_k when the actual auction context is u_j. Therefore, after receiving s_k, the buyer will update their belief about the auction context with the Bayes' rule:

$$q(u_j|s_k) = \frac{\pi(s_k|u_j) q(u_j)}{\sum_{j'=1}^{m} \pi(s_k|u_{j'}) q(u_{j'})}. \tag{1}$$

We consider a public signaling scheme, i.e., the seller uses the same scheme and sends the same information to all buyers. After receiving signal s_k, buyer i will place their expected valuation as their bid for the ad:

$$v_i(s_k) = \sum_{j=1}^{m} v_{ij} q(u_j|s_k). \tag{2}$$

Let $\textsc{Rev}(\pi)$ be the seller's expected revenue when the signaling scheme π is used. Then we have:

$$\begin{aligned} \textsc{Rev}(\pi) &= \mathop{\mathbb{E}}_{u_j \sim q, s \sim \pi(s_k|u_j)} \left[c_{i^* j} \cdot \frac{v_{i_2(s_k)}(s_k) c_{i_2 j}}{c_{i^* j}} \right] \\ &= \mathop{\mathbb{E}}_{u_j \sim q, s \sim \pi(s_k|u_j)} [v_{i_2(s_k)}(s_k) c_{i_2 j}]. \end{aligned}$$

For simplicity, in this paper, we consider the case with $c_{ij} = 1, \forall i, j$. Therefore,

$$\begin{aligned} \textsc{Rev}(\pi) &= \mathop{\mathbb{E}}_{u_j \sim q, s \sim \pi(s_k|u_j)} [v_{i_2(s_k)}(s_k)] \\ &= \sum_j q(u_j) \sum_k \pi(s_k|u_j) v_{i_2(s_k)}(s_k) \\ &= \sum_k v_{i_2(s_k)}(s_k) \sum_j q(u_j) \pi(s_k|u_j) \\ &= \sum_k \left(\sum_j q(u_j|s_k) v_{i_2(s_k), j} \right) \left(\sum_j q(u_j) \pi(s_k|u_j) \right). \end{aligned}$$

Plugging in Equation (1), we get:

$$\textsc{Rev}(\pi) = \sum_k \sum_j \pi(s_k|u_j) q(u_j) v_{i_2(s_k), j}. \tag{3}$$

It is known that the online ad auction platform can increase their revenue by strategically revealing information to the buyers [14]. In fact, using a signaling scheme also helps protect the users' privacy as the scheme only reveals partial information about the user.We utilize the following f-divergence to define our privacy metric.

Definition 1 (f-divergence). *Let $P(x)$ and $Q(x)$ be two probability distributions defined over the same finite set X. For any convex function $f : (0, +\infty) \mapsto \mathbb{R}$ with $f(1) = 0$, the f-divergence of P from Q is:*

$$D_f(P\|Q) = \sum_{x \in X} Q(x) f\left(\frac{P(x)}{Q(x)} \right). \tag{4}$$

Intuitively, if a signaling scheme reveals no information about the user's identity, the signal s_k should be independent of the user u_j and the joint distribution of s_k and u_j is the product of the marginal distributions. Therefore, the f-divergence of the joint distribution from the product distribution can be used to measure how much information the signal s_k reveals about u_j:

$$\begin{aligned} &D_f(\pi(u_j, s_k) \| q(u_j) P_\pi(s_k)) \\ &= \sum_{j,k} q(u_j) P_\pi(s_k) f\left(\frac{\pi(u_j, s_k)}{q(u_j) P_\pi(s_k)} \right), \end{aligned} \tag{5}$$

where $\pi(u_j, s_k) = \pi(s_k|u_j) q(u_j)$ is the joint distribution of sending signal s_k and the user being u_j, and $P_\pi(s_k) = \sum_j \pi(u_j, s_k) = \sum_j \pi(s_k|u_j) q(u_j)$ is the marginal distribution of sending signal s_k. We use the negated f-divergence as our privacy metric since a smaller divergence means better privacy protection:

$$\textsc{Privacy}(\pi) = -\sum_{j,k} q(u_j) P_\pi(s_k) f\left(\frac{\pi(u_j, s_k)}{q(u_j) P_\pi(s_k)} \right).$$

Note that by setting $f(x) = x \log x$, Equation (4) becomes the Kullback-Leibler divergence, and Equation (5) becomes the mutual information. For ease of presentation, we also define

$$\begin{aligned} D_f(\pi; u_j, s_k) =& q(u_j) P_\pi(s_k) f\left(\frac{\pi(u_j, s_k)}{q(u_j) P_\pi(s_k)} \right) \\ =& q(u_j) P_\pi(s_k) f\left(\frac{\pi(s_k|u_j)}{P_\pi(s_k)} \right). \end{aligned}$$

Therefore, we have

$$\textsc{Privacy}(\pi) = -\sum_{j,k} D_f(\pi; u_j, s_k).$$

In this paper, we consider maximizing the platform's revenue without sacrificing too much privacy. Specifically, we aim to solve the following mathematical program:

$$\begin{aligned} \text{maximize:} \quad & \textsc{Rev}(\pi) \\ \text{subject to:} \quad & \textsc{Privacy}(\pi) \geq \gamma \end{aligned} \tag{6}$$

In the above program, γ is a parameter that controls how much privacy we allow to sacrifice in exchange for a better revenue.

3 Problem Analysis

In this section, we formulate the problem as a mathematical program, and derive structural results of the optimal solution.

It is known that if $f(x) = x \log x$, the mutual information (Equation (5)) is a convex function of the conditional probability distribution $\pi(s_k|u_j)$ [28]. In fact, any f-divergence shown in Equation (5) is a convex function of $\pi(s_k|u_j)$.

Lemma 1. *$D_f(\pi(u_j, s_k) \| q(u_j) P_\pi(s_k))$ is a convex function of the conditional probability distribution $\pi(s_k|u_j)$, and becomes strictly convex if $f(x)$ is strictly convex.*

The proof is standard and thus deferred to the full version. Lemma 1 immediately implies that Program (6) is a convex optimization problem. This result holds for any signaling scheme π. However, directly solving Program (6) is intractable, as the space of π is too large.

Emek et al. [14] show that when $\alpha = 1$, there exists an optimal signaling scheme that uses at most $n(n-1)$ signals, i.e., $|S| = n(n-1)$. In fact, this result also applies to our setting.

Theorem 1. *It is without loss of generality to consider signaling schemes with $n(n-1)$ signals.*

We defer the lengthy proof of Theorem 1 to the full version. Theorem 1 implies that each signal corresponds to a combination of the top bidder and the second top bidder. Thus, each signal s can actually be indexed by the pair $(i_1(s), i_2(s))$. In the Bayesian persuasion literature, each signal can be interpreted as an "action recommendation", while in this paper, there are multiple receivers, and different signals lead to different expected valuations for each buyer. In both settings, each signal can be viewed as an "outcome indicator". Let $s_{i,i'}$ be the signal that leads to buyer i having the highest expected valuation and i' having the second highest expected valuation. With Theorem 1, we can formulate the problem of maximizing the objective function as the following mathematical program:

$$\begin{aligned}
&\text{maximize:} \\
&\quad J_\alpha(\pi) = \alpha \cdot \textsc{Rev}(\pi) + (1-\alpha) \cdot \textsc{Privacy}(\pi) \\
&\text{subject to:} \\
&\sum_{j=1}^{m} \pi(s_{i,i'}|u_j) q(u_j) v_{i,j} \geq \sum_{j=1}^{m} \pi(s_{i,i'}|u_j) q(u_j) v_{i',j} && \forall i, i' \neq i \\
&\sum_{j=1}^{m} \pi(s_{i,i'}|u_j) q(u_j) v_{i',j} \geq \sum_{j=1}^{m} \pi(s_{i,i'}|u_j) q(u_j) v_{i'',j} && \forall i'' \neq i, i' \\
&\sum_{i,i' \neq i} \pi(s_{i,i'}|u_j) = 1 && \forall j \\
&\pi(s_{i,i'}|u_j) \geq 0 && \forall i, i' \neq i, j
\end{aligned} \tag{7}$$

Note that in the above program, all constraints are linear except the first one. We apply Lagrangian relaxation to the first constraint and obtain a new objective function:

$$\textsc{Rev}(\pi) + \beta(\textsc{Privacy}(\pi) - \gamma).$$

The following result shows that for any γ, there exists an β such that optimizing the relaxed program gives the same optimal solution as the original one.

Lemma 2. *For any γ, there is an β, such that the optimal solution to the relaxed program is also optimal in the original program.*

Due to space limit, the proof of Lemma 2 is also deferred to the full version.

According to Lemma 2, we can simply focus on the relaxed program. Note that the objective function in the relaxed program is $\textsc{Rev}(\pi) + \beta\textsc{Privacy}(\pi) - \beta\gamma$. We can safely ignore the last term as it is constant. Define $\alpha = \frac{1}{1+\beta}$, and the objective function is equivalent to $\alpha\textsc{Rev}(\pi) + (1-\alpha)\textsc{Privacy}(\pi)$, which is a convex combination of $\textsc{Rev}(\pi)$ and $\textsc{Privacy}(\pi)$, and is also convex itself. We can balance between revenue and privacy by changing α in the interval $[0, 1]$.[4] Therefore, from now on, we will only consider the following program

$$\begin{aligned}
&\text{maximize:} \\
&\quad \alpha \cdot \textsc{Rev}(\pi) + (1-\alpha) \cdot \textsc{Privacy}(\pi) \\
&\text{subject to:} \\
&\sum_{j=1}^{m} \pi(s_{i,i'}|u_j) q(u_j) v_{i,j} \geq \sum_{j=1}^{m} \pi(s_{i,i'}|u_j) q(u_j) v_{i',j} && \forall i, i' \neq i \\
&\sum_{j=1}^{m} \pi(s_{i,i'}|u_j) q(u_j) v_{i',j} \geq \sum_{j=1}^{m} \pi(s_{i,i'}|u_j) q(u_j) v_{i'',j} && \forall i'' \neq i, i' \\
&\sum_{i,i' \neq i} \pi(s_{i,i'}|u_j) = 1 && \forall j \\
&\pi(s_{i,i'}|u_j) \geq 0 && \forall i, i' \neq i, j
\end{aligned} \tag{8}$$

Let $J_\alpha(\pi) = \alpha \cdot \textsc{Rev}(\pi) + (1-\alpha) \cdot \textsc{Privacy}(\pi)$. For ease of presentation, we also define

$$J_\alpha(\pi; u_j, s_k) = \alpha\pi(s_k|u_j)q(u_j)v_{i_2(s_k),j} - (1-\alpha)D_f(\pi; u_j, s_k).$$

Therefore, we have $J_\alpha(\pi) = \sum_{j,k} J_\alpha(\pi; u_j, s_k)$.

Theorem 2. *Removing the second constraint does not affect the optimal objective value of Program* (8).

The proof of Theorem 2 is also deferred to the full version. With Theorem 2, we can safely remove the second constraint and obtain the following program:

$$\begin{aligned}
&\text{maximize:} \\
&\quad J_\alpha(\pi) = \alpha \cdot \textsc{Rev}(\pi) + (1-\alpha) \cdot \textsc{Privacy}(\pi) \\
&\text{subject to:} \\
&\sum_{j=1}^{m} \pi(s_{i,i'}|u_j) q(u_j) v_{i,j} \geq \sum_{j=1}^{m} \pi(s_{i,i'}|u_j) q(u_j) v_{i',j} && \forall i, i' \neq i \\
&\sum_{i,i' \neq i} \pi(s_{i,i'}|u_j) = 1 && \forall j \\
&\pi(s_{i,i'}|u_j) \geq 0 && \forall i, i' \neq i, j
\end{aligned} \tag{9}$$

Now we derive structural results of Program (9).

Lemma 3. *For any bidder i, i', i'', and user u_j, if $v_{i'',j} > v_{i,j} > v_{i',j}$, then the optimal scheme π^* satisfies $\pi^*(s_{i'',i}|u_j) \geq \pi^*(s_{i',i}|u_j)$.*

The proof of Lemma 3 is also deferred to the full version. Lemma 3 reveals relative monotonicity relations between valuations and signaling schemes, i.e., for any buyer i and user u_j, if there are both larger-valued buyers and lower-valued buyers, then in the optimal scheme, when user u_j appears, a signal that ranks a larger-valued buyer as the top buyer is more likely to be sent than a signal that ranks a lower-valued buyer as the top buyer.

Lemma 4. *Let π^* be the optimal scheme. For any bidder i, i', i'', and user u_j, if $\pi^*(s_{i'',i}|u_j) > \pi^*(s_{i',i}|u_j)$, then we have either*

$$\sum_{j'} \pi^*(s_{i',i}|u_{j'}) q(u_{j'}) v_{i',j'} = \sum_{j'} \pi^*(s_{i',i}|u_{j'}) q(u_{j'}) v_{i,j'}$$

$$v_{i',j} < v_{i,j},$$

[4] By definition, α cannot be 0. However, we can still set α to include the case where the platform does not have privacy constraints.

or

$$\sum_{j'} \pi^*(s_{i'',i}|u_{j'})q(u_{j'})v_{i'',j'} = \sum_{j'} \pi^*(s_{i'',i}|u_{j'})q(u_{j'})v_{i,j'}$$
$$v_{i'',j} > v_{i,j}.$$

The proof of Lemma 4 is deferred to the full version. Lemma 4 shows that, among all the signals where bidder i ranks second, if two signals have different conditional probabilities, then bidder i's bid is the same as the top bidder under at least one of the signals.

4 Theoretic Guarantees

In this section, we analyze the performance of Program (9). Specifically, we seek to understand how each performance metric is compared to their best possible value, i.e., the approximation of each performance metric. However, our privacy metric is a negated f-divergence, which is know to be non-positive. Therefore, we analyze how much the solution to Program (9) improves each performance metric instead.

Let $P(\alpha)$ be the above program parameterized by α, and π_α the optimal solution to $P(\alpha)$. Different α's leads to a different solution. Since $P(\alpha)$ is a convex program for any α, optimizing $P(\alpha)$ gives a Pareto optimal solution. (i.e., there exists no signaling scheme that leads to both a larger revenue and a larger privacy).

Specifically, π_1 and π_0 are the signaling schemes that maximizes only $\textsc{Rev}(\pi)$ and $\textsc{Privacy}(\pi)$, respectively[5].Note that π_0 reveals no information at all (e.g., always sending the same signal). The problem becomes uninteresting if $\textsc{Rev}(\pi_0) = \textsc{Rev}(\pi_1)$ or $\textsc{Privacy}(\pi_0) = \textsc{Privacy}(\pi_1)$, since in these cases, simply using π_0 or π_1 simultaneously maximizes revenue and privacy. Therefore, from now on, we assume $\textsc{Rev}(\pi_1) > \textsc{Rev}(\pi_0)$ and $\textsc{Privacy}(\pi_0) > \textsc{Privacy}(\pi_1)$. Intuitively, by increasing α, the objective function of program $P(\alpha)$ focuses more on revenue but less on privacy, and thus leads to higher revenue and lower privacy.

Lemma 5. *$\textsc{Rev}(\pi_\alpha)$ is an increasing function of α, and $\textsc{Privacy}(\pi_\alpha)$ is a decreasing function of α.*

The proof is straightforward and thus deferred to the full version. For ease of presentation, we write $\textsc{Rev}_M = \textsc{Rev}(\pi_1)$ and $\textsc{Privacy}_M = \textsc{Privacy}(\pi_0)$ to be the largest possible revenue and privacy value achieved by any signaling scheme. We also define $\textsc{Rev}_{base} = \textsc{Rev}(\pi_0)$ and $\textsc{Privacy}_{base} = \textsc{Privacy}(\pi_1)$ to be the revenue and privacy achieved when optimizing only the other performance metric.

The following result shows that we can achieve a certain fraction of the best revenue and privacy gain by simply tuning α

Theorem 3. *For any α, there exists $\lambda \in [0,1]$ such that π_α achieves λ fraction of the optimal revenue gain and $1-\lambda$ fraction of the optimal privacy gain, i.e.,*

$$\frac{\textsc{Rev}(\pi_\alpha) - \textsc{Rev}_{base}}{\textsc{Rev}_M - \textsc{Rev}_{base}} \geq \lambda, \tag{10}$$

$$\frac{\textsc{Privacy}(\pi_\alpha) - \textsc{Privacy}_{base}}{\textsc{Privacy}_M - \textsc{Privacy}_{base}} \geq 1-\lambda. \tag{11}$$

[5]There may be multiple schemes that maximize $\textsc{Rev}(\pi)$ or $\textsc{Privacy}(\pi)$, We choose π_1 (π_0) to be the one with the largest privacy (revenue) among all schemes that maximize $\textsc{Rev}(\pi)$ ($\textsc{Privacy}(\pi)$). However, all our theoretic results apply to any choice of π_1 and π_0.

Furthermore, for any $\lambda \in [0,1]$, there exists $\alpha \in [0,1]$, such that π_α satisfies the above inequalities.

The proof of Theorem 3 is deferred to the full version. Theorem 3 shows that by tuning the parameter α, we can easily balance between optimizing different objectives.

Let $V = (v_{ij})_{i,j}$ be the valuation profile of all buyers. Define $d_{i,i'}(V) = \max_j |v_{ij} - v_{i'j}|$ to be the maximum valuation discrepancy between buyer i and i' among all users. Define the "diameter" of V to be $d(V) = \max_{i,i'} d_{i,i'}(V)$, i.e., the maximum $d_{i,i'}(V)$ among all buyer pairs. In other words, if we define $v_i = (v_{i1}, v_{i2}, \ldots, v_{im})$ to be the valuation profile of buyer i, then $d(V) = \max_{i,i'} \|v_i - v_{i'}\|_\infty$ is the largest distance among all vectors v_i, where the distance is induced by ∞-norm. Then we have the following result:

Theorem 4. *π_0 satisfies $\textsc{Rev}(\pi_0) \geq \textsc{Rev}_M - d(V)$.*

The proof is straightforward and thus deferred to the full version. Theorem 4 shows that if the buyers' valuations align well ($d(V)$ is small) then there is enough competition and revealing no information almost simultaneously maximizes privacy and revenue. In fact, it is often the case in real-world applications since similar advertisers usually target similar users and thus place similar bids.

Another simple yet interesting observation is that compared with the standard industry practice where the platform reveals all information to the buyers to help them better evaluate the ads, revealing only partial information can actually improve both revenue and privacy simultaneously.

5 Experiments

In this section, we conduct experiments on both synthetic and real data sets and report the results of our methodWe choose to experiment with the following f-divergence:

- Kullback-Leibler divergence: $f(x) = x \log x$;
- Jensen-Shannon divergence: $f(x) = x \log \frac{2x}{x+1} + \log \frac{2}{x+1}$;
- Total variation divergence: $f(x) = \frac{1}{2}|x-1|$.

Note that when the Kullback-Leibler divergence is used, $D_f(\cdot)$ actually becomes the mutual information. All the above f-divergences are widely used in the literature as measures of information.

5.1 Independent Valuations

We first conduct experiments in the setting where the bidders' values are independent. This setting is standard in most existing papers on auction design.

5.1.1 Synthetic Data Set. In the synthetic data set, the number of bidders is set to 5 ($n = 5$), and the number of users is set to 10 ($m = 10$) and 20 ($m = 20$), respectively. The results with $m = 20$ show similar patterns to those with $m = 10$. Due to space limit, these results are deferred to the full version. We generate 50 different problem instances for both $m = 10$ and $m = 20$. And for each problem instance, v_{ij} is drawn independently from the uniform distribution $U[0,1]$, for all i and j. The distribution $q(u_j)$ is also randomly generated. We choose 500 values uniformly from the interval $[0,1]$ and use these values as α in Program (9). All reported results are averaged over these 50 problem instances. All the curves in the figures are Pareto frontiers of the corresponding settings. To

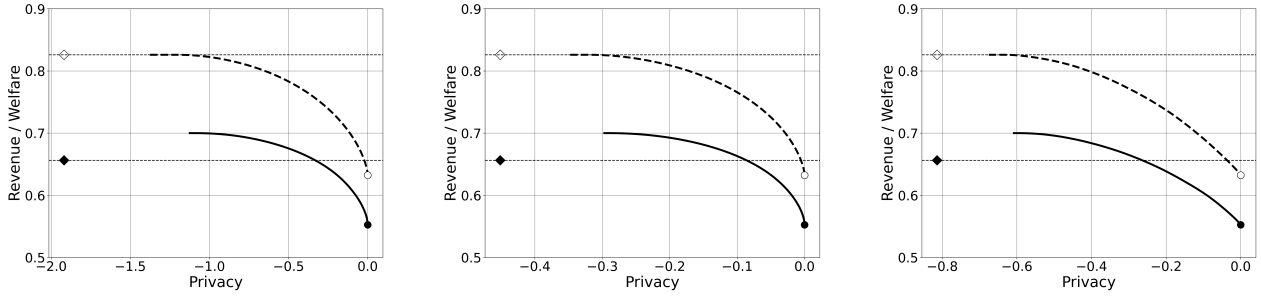

(a) KL divergence (b) JS divergence (c) TV divergence

Figure 1: Performance with $m = 10, n = 5$ on the synthetic data set. The solid and dashed lines represent the curves for revenue and welfare. The points indicated by diamond markers and round markers correspond to the results by revealing full information and no information (π_0), respectively.

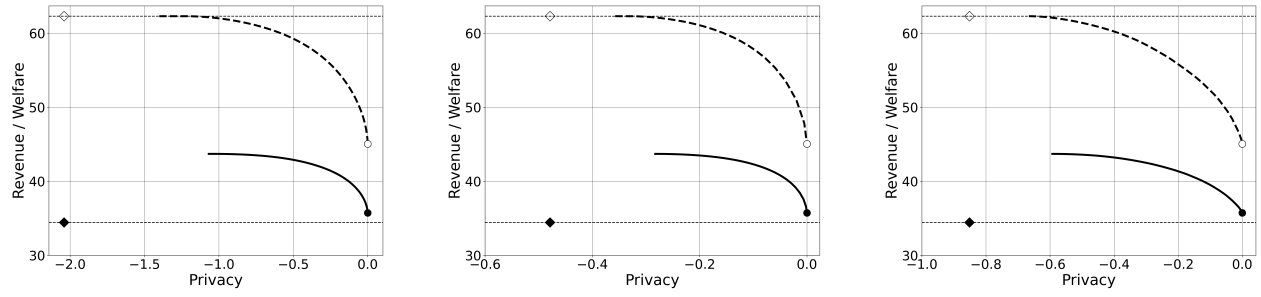

(a) KL divergence (b) JS divergence (c) TV divergence

Figure 2: Performance with $m = 10, n = 5$ on the iPinYou data set. The solid and dashed lines represent the curves for revenue and welfare. The points indicated by diamond markers and round markers correspond to the results by revealing full information and no information (π_0), respectively.

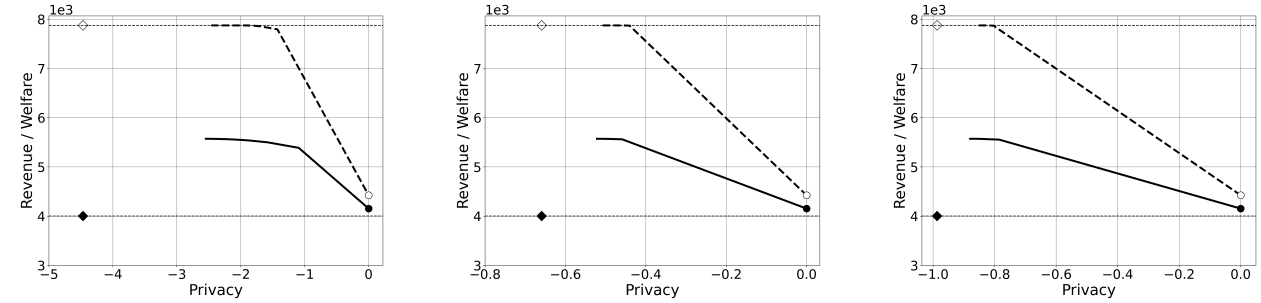

(a) KL divergence (b) JS divergence (c) TV divergence

Figure 3: Performance with $m = 100, n = 10$ Kuaishou data set. The solid and dashed lines represent the curves for revenue and welfare. The points indicated by diamond markers and round markers correspond to the results by revealing full information and no information (π_0), respectively.

solve Program (9), we use the cvxpy package [1, 10] and set the solver to be SCS (Splitting Conic Solver [27]).

The results of the first experiment are shown in Figure 1, where we also calculate the welfare in our experiments. The expected

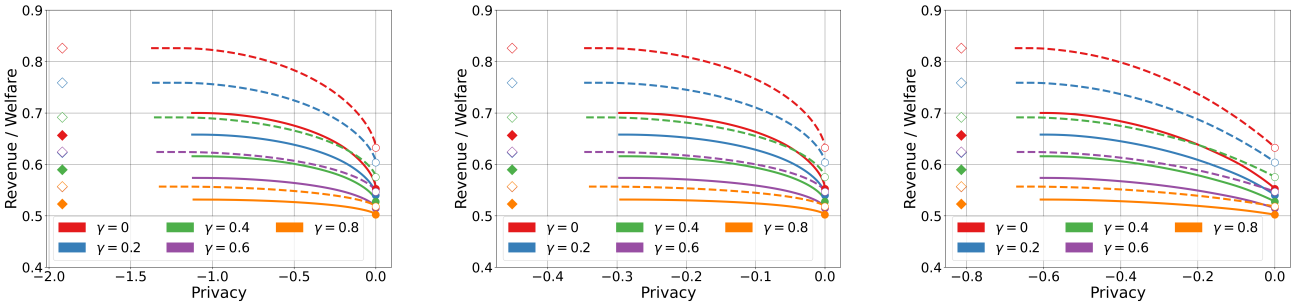

(a) KL divergence (b) JS divergence (c) TV divergence

Figure 4: Performance with $m = 10$ and five different weight coefficients on the synthetic data set.

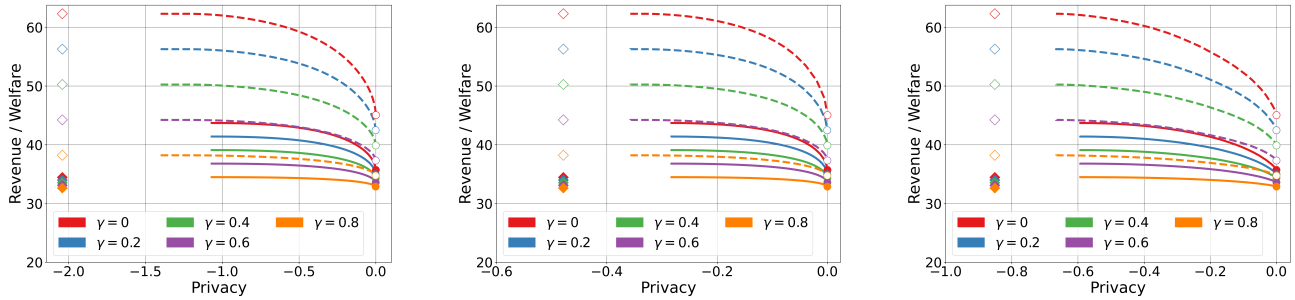

(a) KL divergence (b) JS divergence (c) TV divergence

Figure 5: Performance with $m = 10$ and five different weight coefficients on the iPinYou data set.

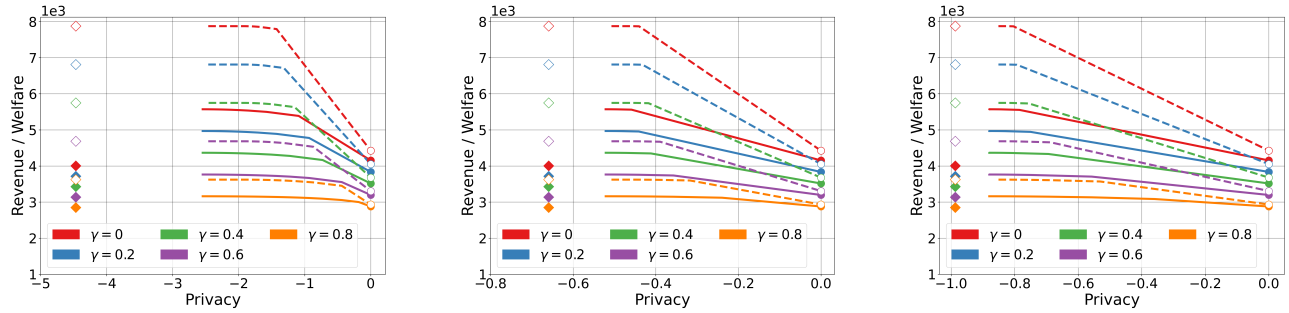

(a) KL divergence (b) JS divergence (c) TV divergence

Figure 6: Performance with $m = 100, n = 10$ and five different weight coefficients on Kuaishou data set.

welfare based scheme π is:

$$\textsc{Wel}(\pi) = \sum_k \sum_j \pi(s_k|u_j)q(u_j)v_{i_1(s_k),j}.$$

The welfare curve can be derived by solving programs obtained by replacing the objective in Program (9) with the above function.

All the 3 figures indicate that compared with the case of revealing full information (sending the identity of u_j directly to the bidders) , using a scheme π_1 can indeed improve both revenue and privacy by a large margin. In addition, to achieve the same revenue as revealing full information, we can in fact use a scheme that performs much better in terms of privacy protection to achieve the same revenue.

Both the welfare curve and the revenue curve are almost flat when privacy is low, which corresponds to cases with large α's. This means we can significantly improve privacy by not sacrificing too much revenue and welfare. Compared with the revenue-maximizing

scheme (π_1) and welfare-maximizing scheme (revealing full information), revealing no information (using π_0) can achieve 79% and 77% of the optimal revenue and optimal welfare, respectively.

In all the 3 plots in Figure 1, the leftmost point of the dashed curve can be obtained by maximizing privacy among all schemes that maximize welfare. Therefore, we can conclude that the revenue-maximizing scheme better protects the users' privacy compared to any welfare-maximizing scheme.

5.1.2 Real Data Set. We also conduct experiments using the iPinYou data set [19]. This data set is released by a demand-side platform that places bids on ad auction platforms on behalf of the advertisers. The data set includes 3 auction seasons (24 days) of data with 78 million records. Each bid record contains information on bidder id, bid, paying pricing, etc. There are 6 different ad platforms including Google, Baidu, Alibaba and Tencent. We use data of the second season in our experiment as the data of this season includes user characteristics such as hobbies and buying preferences. We extracted 533,184 users and 5 bidders from the data set. However, the bid records are set to be very high only to collect data, and thus cannot be used. But the paying prices are the true bids of other bidders, and we use these prices as bids in our experiments.

To simplify our experiment process, we use the K-means algorithm to cluster the users into 10 ($m = 10$) and 20 ($m = 20$) meta-users. We also defer the results for $m = 20$ to the full version. For each meta-user and each bidder, we fit the bids to a lognormal distribution. The probability $q(u_j)$ is also calculated based on the number of records in the data set.

As shown in Figure 2, the results for this experiment show similar patterns as those for the first experiment. In this experiment, the revenue and welfare of revealing no information achieve 82% of the optimal revenue and 72% of the optimal welfare. Interestingly, the revenue of revealing full information is even smaller than that of revealing no information under all privacy measures, which indicates that using any α leads to a scheme that simultaneously improves revenue and privacy.

The second data set is from Kuaishou, a major short-form video and live-stream company in China. We extract data from a week's worth of ad auctions, in which we select 10 advertisers (n=10) to serve as bidders. We then categorize the platform's users into 100 different groups (m=100) based on characteristics such as age, gender, and interests. We consider the average bid made by each advertiser for the same type of user within a week as v_{ij}.

As illustrated in Figure 3, the results of this experiment align closely with those of our first and second experiments. The revenue and welfare of revealing no information reached 75% of the optimal revenue and 56% of the optimal welfare. Interestingly, we also discovered that the revenue of revealing full information is less than the revenue of revealing no information. This result is consistent with the conclusions drawn from the iPinYou data set.

5.2 Correlated Valuations

In this section, we report the results of experiments conducted in the setting with correlated valuations. This setting is also called the "inter-dependent valuation" setting in the literature. The inter-dependence among the bidders' valuations captures the cases where similar bidders have similar preferences over the users. For example, bidders from the same industry may target a similar set of users and place similar bids when a user arrives.

We re-conduct all the experiments in Section 5.1 with the same parameters except the joint valuation distribution. We posit that a bidder's bid can be attributed to two components: the bidder's intrinsic valuation of the user and an aggregate of interrelated valuations among bidders. This can be expressed as the bidder's new bid, denoted as $v_{i,j} = (1-\gamma)v^0_{i,j} + \gamma\bar{v}_j$, where $v^0_{i,j}$ is an independent random variable that follows the same distribution as in Section 5.1, and $\bar{v}_j = \frac{1}{n}\sum_i^n v^0_{i,j}$ is the averaged valuation of all bidders, which is the same among all bidders, and γ serves as the weight coefficient that control the degree of correlation. Clearly, $\gamma = 0$ leads to the same setting in Section 5.1. In our experiments, we explored several different weight coefficients: $\gamma = 0.2, \gamma = 0.4, \gamma = 0.6, \gamma = 0.8$. Similar to Section 5.1, we evaluated the relation between welfare and privacy, as well as revenue and privacy, under each weight coefficient. We report results for $m = 10$ of the first two experiments in Figure 4 and 5, and the results for the Kuaishou data set in Figure 6. The results for $m = 20$ are deferred to the full version.

All experimental results show that when bidders' bids are correlated, our previous theoretic results still hold. At the same time, as γ increases, both revenue and welfare decrease. Another observation revealed from Figure 4, 5 and 6 is that a higher γ leads to smaller revenue gain and welfare gain. This means the more correlated valuations the bidders have, the less revenue gain and welfare gain can be extracted from revealing partial information with a signaling scheme. This aligns with our intuition and confirms our theoretic result Theorem 4. Consider the extreme case where the valuations are fully correlated, e.g., all bidders have the same valuation for the same user ($\gamma = 1$). In this case, no matter how much information the platform reveals to the bidders, all the bidders have the same expected valuation, which leads to the same revenue and welfare (no revenue and welfare gain).

6 Conclusion

In this paper, we study how to apply signaling scheme to help protect the privacy of the users in online ad auctions. We use an f-divergence to measure how much information is leaked to the advertiser. We show that this problem can be formulated as a convex optimization problem and thus can be efficiently solved. Our theoretic results show that we can easily balance between revenue and privacy by simply tuning a single parameter. We also prove that the revenue of revealing no information can be bounded by how well the advertisers' valuations are aligned. This means if the market is competitive enough, using a signaling scheme can improve both revenue and privacy, which is also confirmed by our experiment results. Moreover, our experiment results also show that compared with the revenue maximizing or welfare maximizing scheme, we can set an appropriate parameter α to significantly improve privacy without sacrificing too much revenue and welfare.

Acknowledgments

We thank all the anonymous reviewers for their helpful comments. This work was supported in part by the National Natural Science Foundation of China (No. 72192805), and a research fund from Kuaishou.

References

[1] Akshay Agrawal, Robin Verschueren, Steven Diamond, and Stephen Boyd. 2018. A rewriting system for convex optimization problems. *Journal of Control and Decision* 5, 1 (2018), 42–60.

[2] Laurynas Baltrusaitis, TJ Klein, and Kim van Asten MSc. 2022. The Effect of Implementing Differential Privacy on the Second-Price Auction for Display Ads. (2022).

[3] Dirk Bergemann and Alessandro Bonatti. 2023. Data, competition, and digital platforms. *arXiv preprint arXiv:2304.07653* (2023).

[4] Dirk Bergemann, Alessandro Bonatti, and Nicholas Wu. 2023. How Do Digital Advertising Auctions Impact Product Prices? (2023).

[5] Dirk Bergemann, Paul Duetting, Renato Paes Leme, and Song Zuo. 2022. Calibrated click-through auctions. In *Proceedings of the ACM Web Conference 2022*. 47–57.

[6] Dirk Bergemann and Stephen Morris. 2016. Information design, Bayesian persuasion, and Bayes correlated equilibrium. *American Economic Review* 106, 5 (2016), 586–591.

[7] Alessandro Bonatti. 2022. *The Platform Dimension of Digital Privacy.* Technical Report. National Bureau of Economic Research.

[8] Matteo Castiglioni, Andrea Celli, Alberto Marchesi, and Nicola Gatti. 2020. Online bayesian persuasion. *Advances in Neural Information Processing Systems* 33, 16188–16198.

[9] Marco Cuturi. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems* 26 (2013).

[10] Steven Diamond and Stephen Boyd. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* 17, 83 (2016), 1–5.

[11] Cynthia Dwork. 2006. Differential privacy. In *International colloquium on automata, languages, and programming*. Springer, 1–12.

[12] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*. Springer, 265–284.

[13] Stephan Eckstein and Marcel Nutz. 2022. Quantitative Stability of Regularized Optimal Transport and Convergence of Sinkhorn's Algorithm. *SIAM Journal on Mathematical Analysis* 54, 6 (2022), 5922–5948.

[14] Yuval Emek, Michal Feldman, Iftah Gamzu, Renato Paes Leme, and Moshe Tennenholtz. 2014. Signaling schemes for revenue maximization. *ACM Transactions on Economics and Computation (TEAC)* 2, 2 (2014), 1–19.

[15] David S. Evans. 2009. The Online Advertising Industry: Economics, Evolution, and Privacy. *Journal of Economic Perspectives* 23, 3 (September 2009), 37–60. https://doi.org/10.1257/jep.23.3.37

[16] Benjamin Eysenbach and Sergey Levine. 2019. If MaxEnt RL is the answer, what is the question? *arXiv preprint arXiv:1910.01913* (2019).

[17] Avi Goldfarb and Catherine E Tucker. 2011. Privacy regulation and online advertising. *Management science* 57, 1 (2011), 57–71.

[18] Xu Haifeng. 2020. On the tractability of public persuasion with no externalities. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2708–2727.

[19] Liao Hairen, Peng Lingxiao, Liu Zhenchuan, and Shen Xuehua. 2014. iPinYou global rtb bidding algorithm competition dataset. In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*. 1–6.

[20] Garrett Johnson. 2013. The impact of privacy policy on the auction market for online display advertising. (2013).

[21] Chen Junjie, Li Minming, Xu Haifeng, and Zuo Song. 2023. Bayesian Calibrated Click-Through Auction. *ArXiv. /abs/2306.06554* (2023).

[22] Emir Kamenica and Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review* 101, 6 (2011), 2590–2615.

[23] Charlotte Laclau, Ievgen Redko, Basarab Matei, Younes Bennani, and Vincent Brault. 2017. Co-clustering through optimal transport. In *International conference on machine learning*. PMLR, 1955–1964.

[24] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436–444.

[25] Jonathan Levin and Paul Milgrom. 2010. Online advertising: Heterogeneity and conflation in market design. *American Economic Review* 100, 2 (2010), 603–607.

[26] Tom M Mitchell. 1997. Machine learning.

[27] Brendan O'Donoghue, Eric Chu, Neal Parikh, and Stephen Boyd. 2016. Conic Optimization via Operator Splitting and Homogeneous Self-Dual Embedding. *Journal of Optimization Theory and Applications* 169, 3 (June 2016), 1042–1068. http://stanford.edu/~boyd/papers/scs.html

[28] Anup Rao. 2010. Information Theory in Computer Science.

[29] Pierangela Samarati and Latanya Sweeney. 1998. Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression. (1998).

[30] Dughmi Shaddin and Xu Haifeng. 2016. Algorithmic bayesian persuasion. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. 412–425.

[31] Kefan Su and Zongqing Lu. 2022. Divergence-regularized multi-agent actor-critic. In *International Conference on Machine Learning*. PMLR, 20580–20603.

[32] Babichenko Yakov and Barman Siddharth. 2017. Algorithmic aspects of private Bayesian persuasion. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

Appendix

A Omitted Proofs

A.1 Proof of Lemma 1

Proof. For any conditional distributions $\pi(s_k|u_j)$ and $\pi'(s_k|u_j)$, define $\bar{\pi}(s_k|u_j) = \frac{1}{2}[\pi(s_k|u_j) + \pi'(s_k|u_j)]$.

Let $P_\pi(s_k) = \sum_j \pi(s_k|u_j)q(u_j)$. It suffices to show that for any j and k, $D_f(\pi; u_j, s_k)$ is a convex function of $\pi(s_k|u_j)$, since $D_f(\pi(u_j, s_k)\|q(u_j)P(s_k)) = \sum_{i,j} D_f(\pi; u_j, s_k)$ is a linear function of $D_f(\pi; u_j, s_k)$.

$P_\pi(s_k)$ is a linear function of $\pi(s_k|u_j)$, Thus we have $P_{\bar{\pi}}(s_k) = \frac{1}{2}P_\pi(s_k) + \frac{1}{2}P_{\pi'}(s_k)$. Therefore,

$$
\begin{aligned}
&\frac{1}{2}\left[D_f(\pi; u_j, s_k) + D_f(\pi'; u_j, s_k)\right] \\
=&\frac{1}{2}P_\pi(s_k) f\left(\frac{\pi(s_k|u_j)}{P_\pi(s_k)}\right) + \frac{1}{2}P_{\pi'}(s_k) f\left(\frac{\pi'(s_k|u_j)}{P_{\pi'}(s_k)}\right) \\
=&P_{\bar{\pi}}(s_k)\left[\frac{\frac{1}{2}P_\pi(s_k)}{P_{\bar{\pi}}(s_k)} f\left(\frac{\pi(s_k|u_j)}{P_\pi(s_k)}\right)\right. \\
&\left. + \frac{\frac{1}{2}P_{\pi'}(s_k)}{P_{\bar{\pi}}(s_k)} f\left(\frac{\pi'(s_k|u_j)}{P_{\pi'}(s_k)}\right)\right] \\
\geq&P_{\bar{\pi}}(s_k) f\left(\frac{\frac{1}{2}P_\pi(s_k)}{P_{\bar{\pi}}(s_k)}\frac{\pi(s_k|u_j)}{P_\pi(s_k)} + \frac{\frac{1}{2}P_{\pi'}(s_k)}{P_{\bar{\pi}}(s_k)}\frac{\pi'(s_k|u_j)}{P_{\pi'}(s_k)}\right) \\
=&P_{\bar{\pi}}(s_k) f\left(\frac{\frac{1}{2}\pi(s_k|u_j)}{P_{\bar{\pi}}(s_k)} + \frac{\frac{1}{2}\pi'(s_k|u_j)}{P_{\bar{\pi}}(s_k)}\right) \\
=&P_{\bar{\pi}}(s_k) f\left(\frac{\bar{\pi}(s_k|u_j)}{P_{\bar{\pi}}(s_k)}\right) \\
=&D_f(\bar{\pi}; u_j, s_k),
\end{aligned}
$$

where the inequality is due to the Jensen's inequality.

And if $f(x)$ is strictly convex, the above inequality also becomes strict, making $D_f(\pi(u_j, s_k)\|q(u_j)P_\pi(s_k))$ a strictly convex function. □

A.2 Proof of Theorem 1

Proof. We prove the theorem by showing that for any signaling scheme with strictly more than $n(n-1)$ signals, we can construct a new signaling scheme with $n(n-1)$ signals that achieve a higher objective value.

For any signal s, it induces an expected valuation $v_i(s)$ for buyer i. Define

$$
\begin{aligned}
i_1(s) &= \arg\max_{i \in N} v_i(s), \\
i_2(s) &= \arg\max_{i \in N \setminus \{i_1(s)\}} v_i(s)
\end{aligned}
$$

to be the buyers with the highest and second highest expected valuations. For any signaling scheme π with strictly more than $n(n-1)$ signals, there must exists at least two signals s_{k_1} and s_{k_2} with $i_1(s_{k_1}) = i_1(s_{k_2})$ and $i_2(s_{k_1}) = i_2(s_{k_2})$, since there are only $n(n-1)$ different combinations of the top two buyers. Now we construct a new signaling scheme S' and π' as follows:

$$
\begin{aligned}
S' &= S \setminus \{s_{k_1}, s_{k_2}\} \cup \{s_{k'}\}, \\
\pi'(s|u_j) &= \begin{cases} \pi(s|u_j) & \text{if } s \in S \setminus \{s_{k_1}, s_{k_2}\} \\ \pi(s_{k_1}|u_j) + \pi(s_{k_2}|u_j) & \text{if } s = s_{k'} \end{cases}.
\end{aligned}
$$

Simply put, in the new signaling scheme, we merge signals s_{k_1} and s_{k_2} to a single signal $s_{k'}$, and let the seller send signal $s_{k'}$ whenever they send s_{k_1} or s_{k_2} in the original scheme.

We first show that for the combined signal $s_{k'}$, the top two buyers are the same as those for the original signals s_{k_1} and s_{k_2}. In the original scheme, for any signal $s \in S$, we have

$$
\begin{aligned}
v_{i_1(s)}(s) &\geq v_{i_2(s)}(s), \\
v_{i_2(s)}(s) &\geq v_{i'}(s), \forall i \neq i_1(s), i_2(s).
\end{aligned}
$$

Combining with Equation (1) and (2) yields:

$$
\sum_{j=1}^{m} \pi(s|u_j)q(u_j)v_{i_1(s),j} \geq \sum_{j=1}^{m} \pi(s|u_j)q(u_j)v_{i_2(s),j}. \tag{12}
$$

Also, for any $i \neq i_1(s), i_2(s)$:

$$
\sum_{j=1}^{m} \pi(s|u_j)q(u_j)v_{i_2(s),j} \geq \sum_{j=1}^{m} \pi(s|u_j)q(u_j)v_{i,j}. \tag{13}
$$

In the new scheme, for any $s \in S \setminus \{s_{k_1}, s_{k_2}\}$, we have:

$$
q(u_j|s) = \frac{\pi'(s|u_j)q(u_j)}{\sum_{j'} \pi'(s|u_{j'})q(u_{j'})} = \frac{\pi(s|u_j)q(u_j)}{\sum_{j'} \pi(s|u_{j'})q(u_{j'})},
$$

which is the same as the posterior belief in the original scheme. Thus the order of the buyers are also the same as in the original scheme. If $s = s_{k'}$, we have

$$
\begin{aligned}
q(u_j|s_{k'}) =& \frac{\pi'(s_{k'}|u_j)q(u_j)}{\sum_{j'} \pi'(s_{k'}|u_{j'})q(u_{j'})} \\
=& \frac{[\pi(s_{k_1}|u_j) + \pi(s_{k_2}|u_j)]q(u_j)}{\sum_{j'} \pi'(s_{k'}|u_{j'})q(u_{j'})}.
\end{aligned}
$$

Recall that $i_1(s_{k_1}) = i_1(s_{k_2})$ and $i_2(s_{k_1}) = i_2(s_{k_2})$. Writing Equation (12) for both s_{k_1} and s_{k_2}, adding them together gives:

$$
\begin{aligned}
&\sum_j \frac{[\pi(s_{k_1}|u_j) + \pi(s_{k_2}|u_j)]q(u_j)}{\sum_{j'} \pi'(s_{k'}|u_{j'})q(u_{j'})} v_{i_1(s_{k_1}),j} \\
\geq &\sum_j \frac{[\pi(s_{k_1}|u_j) + \pi(s_{k_2}|u_j)]q(u_j)}{\sum_{j'} \pi'(s_{k'}|u_{j'})q(u_{j'})} v_{i_2(s_{k_1}),j}.
\end{aligned}
$$

Similarly, for any $i \neq i_1(s_{k_1}), i_2(s_{k_1})$, we have:

$$
\begin{aligned}
&\sum_j \frac{[\pi(s_{k_1}|u_j) + \pi(s_{k_2}|u_j)]q(u_j)}{\sum_{j'} \pi'(s^*|u_{j'})q(u_{j'})} v_{i_2(s_{k_1}),j} \\
\geq &\sum_j \frac{[\pi(s_{k_1}|u_j) + \pi(s_{k_2}|u_j)]q(u_j)}{\sum_{j'} \pi'(s_{k'}|u_{j'})q(u_{j'})} v_{i,j}.
\end{aligned}
$$

The above equations mean that when signal $s_{k'}$ is sent, buyer $i_1(s_{k_1})$ (or $i_1(s_{k_2})$) and $i_2(s_{k_1})$ (or $i_2(s_{k_2})$) are still the top two buyers, i.e., $i_1(s_{k_1}) = i_1(s_{k_2}) = i_1(s_{k'})$ and $i_2(s_{k_1}) = i_2(s_{k_2}) = i_2(s_{k'})$.

Next, we show that the new scheme S' and π' achieves a higher objective value. The intuition is that the new scheme gives the seller the same expected revenue, while performs better in protecting the users' privacy since the new scheme reveals less information.

For any $s_k \in S \setminus \{s_{k_1}, s_{k_2}\}$, since the top two buyers are the same in both schemes and $\pi'(s_k|u_j) = \pi(s_k|u_j)$, we have $J_\alpha^{jk}(\pi') = J_\alpha^{jk}(\pi)$.

For $s_{k'}$, combining signal s_{k_1} and s_{k_2} does not change the first term since for any u_j,

$$\begin{aligned}
&q(u_j)\pi'(s_{k'}|u_j)v_{i_2(s_{k'}),j}\\
=&q(u_j)[\pi(s_{k_1}|u_j)+\pi(s_{k_2}|u_j)]v_{i_2(s_{k'}),j}\\
=&q(u_j)\pi(s_{k_1}|u_j)v_{i_2(s_{k_1}),j}+q(u_j)\pi(s_{k_2}|u_j)v_{i_2(s_{k_2}),j}. \qquad (14)
\end{aligned}$$

As for the second term, we have:

$$\begin{aligned}
&D_f(\pi;u_j,s_{k_1})+D_f(\pi;u_j,s_{k_2})\\
=&q(u_j)P_\pi(s_{k_1})f\left(\frac{\pi(s_{k_1}|u_j)}{P_\pi(s_{k_1})}\right)+P_\pi(s_{k_2})f\left(\frac{\pi(s_{k_2}|u_j)}{P_\pi(s_{k_2})}\right)\\
=&q(u_j)P_{\pi'}(s_{k'})\left[\frac{P_\pi(s_{k_1})}{P_{\pi'}(s_{k'})}f\left(\frac{\pi(s_{k_1}|u_j)}{P_\pi(s_{k_1})}\right)\right.\\
&\qquad\left.+\frac{P_\pi(s_{k_2})}{P_{\pi'}(s_{k'})}f\left(\frac{\pi(s_{k_2}|u_j)}{P_\pi(s_{k_2})}\right)\right]\\
\geq&q(u_j)P_{\pi'}(s_{k'})f\left(\frac{P_\pi(s_{k_1})}{P_{\pi'}(s_{k'})}\frac{\pi(s_{k_1}|u_j)}{P_\pi(s_{k_1})}\right.\\
&\qquad\left.+\frac{P_\pi(s_{k_2})}{P_{\pi'}(s_{k'})}\frac{\pi(s_{k_2}|u_j)}{P_\pi(s_{k_2})}\right)\\
=&q(u_j)P_{\pi'}(s_{k'})f\left(\frac{\pi(s_{k_1}|u_j)}{P_{\pi'}(s_{k'})}+\frac{\pi(s_{k_2}|u_j)}{P_{\pi'}(s_{k'})}\right)\\
=&q(u_j)P_{\pi'}(s_{k'})f\left(\frac{\pi(s_{k'}|u_j)}{P_{\pi'}(s_{k'})}\right)\\
=&D_f(\pi';u_j,s_{k'}),
\end{aligned}$$

where the inequality is due to the Jensen's inequality. Combining the above equations gives $J_\alpha(\pi;u_j,s_{k_1})+J_\alpha(\pi;u_j,s_{k_2})\leq J_\alpha(\pi';u_j,s_{k'})$. Therefore,

$$\begin{aligned}
J_\alpha(\pi)=&\sum_j\left[J_\alpha(\pi;u_j,s_{k_1})+J_\alpha(\pi;u_j,s_{k_2})\right]\\
&+\sum_j\left[\sum_{k\neq k_1,k_2}J_\alpha(\pi;u_j,s_k)\right]\\
\leq&\sum_j\left[J_\alpha(\pi';u_j,s_{k'})+\sum_{k\neq k'}J_\alpha(\pi';u_j,s_k)\right]\\
=&\sum_{j,k}J_\alpha(\pi';u_j,s_k)\\
=&J_\alpha(\pi').
\end{aligned}$$

□

A.3 Proof of Lemma 2

Proof. Denote by $\mathcal{D}$ the dual program of Program 6. Let π^* be the optimal solution to the original program. Let η be the dual variable associated to the first constraint in $\mathcal{D}$, and η^* be the value in the optimal solution to $\mathcal{D}$. We set $\beta=\eta^*$. Since the original program is convex, we have the complementary slackness condition: $\eta^*(\textsc{Privacy}(\pi^*)-\gamma)=0$. This means that in the relaxed program, setting $\pi=\pi^*$ achieves the same objective value as in the original one, as the second term becomes 0. However, it is known that the optimal objective of the relaxed program provides an upper bound to the original one, which implies that π^* is already optimal in the relaxed program. Therefore, solving the relaxed program also gives the same optimal solution as in the original one. □

A.4 Proof of Theorem 2

Proof. Let π^* be the optimal solution to Program (8). We prove Theorem 2 by contradiction. Assume, on the contrary, that removing the second constraint leads to a different optimal objective value. The new objective value can only be larger since removing a constraint results in a larger feasible region. This means the solution π to the new program is infeasible in the original program 8. And the constraint that it violates can only be the second constraint. Or equivalently, there exists a signal $s_{i,i'}$ and a bidder i'', such that:

$$\sum_{j=1}^m\pi(s_{i,i'}|u_j)q(u_j)v_{i',j}<\sum_{j=1}^m\pi(s_{i,i'}|u_j)q(u_j)v_{i'',j}.$$

This means bidders i and i' are no long the top two bidders. Let $\hat{i}$ and $\hat{i}'$ be the top two bidders in this case. We construct a new signaling scheme as follows:

$$S'=S\setminus\{s_{i,i'}\}$$

$$\pi'(s|u_j)=\begin{cases}\pi(s_{\hat{i},\hat{i}'}|u_j)+\pi(s_{i,i'}|u_j) & \text{if } s=s_{\hat{i},\hat{i}'}\\ \pi(s_{i,i'}|u_j) & \text{otherwise}\end{cases}.$$

Similar to the proof of Theorem 1, we can show that the new signaling scheme achieves a weakly higher objective value than π, and thus also higher than that achieved by π^*. It is easy to check that the new signaling scheme satisfies the second constraint of Program (8). This means π' is feasible in the original Program (8), but leads to a higher objective value, contradicting the optimality of π^*. □

A.5 Proof of Lemma 3

Proof. Suppose on the contrary that in the optimal scheme, we have $\pi^*(s_{i',i}|u_j)>\pi^*(s_{i'',i}|u_j)$. Let $c=\frac{1}{2}\left[\pi^*(s_{i',i}|u_j)+\pi^*(s_{i'',i}|u_j)\right]$ and define

$$\bar{\pi}(s|u)=\begin{cases}c & \text{if } s\in\{s_{i',i},s_{i'',i}\} \text{ and } u=u_j\\ \pi^*(s|u) & \text{otherwise}\end{cases}.$$

We claim that $\bar{\pi}$ is a feasible scheme. We only show that $\bar{\pi}$ satisfies the first constraint, as the other two constraints clearly hold. For the first constraint, $\bar{\pi}$ satisfies the inequality for any $s\notin\{s_{i',i},s_{i'',i}\}$

since in this case $\bar{\pi}(s|u) = \pi^*(s|u)$. For signal $s_{i',i}$, we have:

$$\begin{aligned}
&\sum_j \bar{\pi}(s_{i',i}|u_j)q(u_j)v_{i',j} \\
=&\bar{\pi}(s_{i',i}|u_j)q(u_j)v_{i',j} + \sum_{j'\neq j} \bar{\pi}(s_{i',i}|u_{j'})q(u_{j'})v_{i',j'} \\
=&cq(u_j)v_{i',j} + \sum_{j'\neq j} \bar{\pi}(s_{i',i}|u_{j'})q(u_{j'})v_{i',j'} \\
=&[c - \pi^*(s_{i',i}|u_j)]q(u_j)v_{i',j} + \sum_{j'} \pi^*(s_{i',i}|u_{j'})q(u_{j'})v_{i',j'} \\
\geq&[c - \pi^*(s_{i',i}|u_j)]q(u_j)v_{i,j} + \sum_{j'} \pi^*(s_{i',i}|u_{j'})q(u_{j'})v_{i,j'} \\
=&cq(u_j)v_{i,j} + \sum_{j'\neq j} \bar{\pi}(s_{i',i}|u_{j'})q(u_{j'})v_{i,j'} \\
=&\sum_j \bar{\pi}(s_{i',i}|u_j)q(u_j)v_{i,j},
\end{aligned}$$

where the inequality holds since π^* is feasible and $c < \pi^*(s_{i',i}|u_j)$ and $v_{i,j} > v_{i',j}$. Similarly, we can also show that signal $s_{i'',i}$ satisfies:

$$\sum_j \bar{\pi}(s_{i'',i}|u_j)q(u_j)v_{i'',j} \geq \sum_j \bar{\pi}(s_{i'',i}|u_j)q(u_j)v_{i,j}.$$

Following arguments similar to the proof of Lemma 1, we know that $J_\alpha(\bar{\pi}; u_j, s) \geq J_\alpha(\pi^*; u_j, s), \forall s$. Therefore, $\bar{\pi}$ achieves a higher objective value than π^*, contradicting to the optimality of π^*. □

A.6 Proof of Lemma 4

Proof. Let $\epsilon > 0$ be any sufficiently small positive number. We define π' and $\bar{\pi}$ as follows:

$$\pi'(s|u) = \begin{cases} \pi^*(s_{i',i}|u_j) + \epsilon & \text{if } s = s_{i',i} \text{ and } u = u_j \\ \pi^*(s_{i'',i}|u_j) - \epsilon & \text{if } s = s_{i'',i} \text{ and } u = u_j \\ \pi^*(s|u) & \text{otherwise} \end{cases},$$

$$\bar{\pi}(s|u) = \begin{cases} c & \text{if } s \in \{s_{i',i}, s_{i'',i}\} \text{ and } u = u_j \\ \pi^*(s|u) & \text{otherwise} \end{cases},$$

where $c = \frac{1}{2}\left[\pi^*(s_{i',i}|u_j) + \pi^*(s_{i'',i}|u_j)\right]$. Clearly, we have $\textsc{Rev}(\pi') = \textsc{Rev}(\pi^*)$. As for $\textsc{Privacy}(\pi')$, we define

$$\beta = \frac{c - \pi^*(s_{i',i}|u_j) - \epsilon}{c - \pi^*(s_{i',i}|u_j)}.$$

One can easily check that $\beta\pi^* + (1-\beta)\bar{\pi} = \pi'$.

Since $D_f(\pi; u, s)$ is a convex function in π, we have that for any signal s and user u,

$$\begin{aligned}
&\beta D_f(\pi^*; u, s) + (1-\beta) D_f(\bar{\pi}; u, s) \\
\geq& D_f(\beta\pi^* + (1-\beta)\bar{\pi}; u, s) \\
=& D_f(\pi'; u, s).
\end{aligned}$$

According to the proof of Lemma 3, we know that $D_f(\bar{\pi}; u, s) \leq D_f(\pi^*; u, s)$. Thus we have:

$$\begin{aligned}
D_f(\pi^*; u, s) =& \beta D_f(\pi^*; u, s) + (1-\beta) D_f(\pi^*; u, s) \\
\geq& \beta D_f(\pi^*; u, s) + (1-\beta) D_f(\bar{\pi}; u, s) \\
=& D_f(\pi'; u, s).
\end{aligned}$$

This immediately implies that π' leads to a higher objective value than π^*. However, π' is not optimal. The only reason is that π' is not feasible. Compared with π^*, π' only changes the conditional probabilities of sending signals $s_{i',i}$ and $s_{i'',i}$. The second and third constraints are clearly satisfied by π'. Thus the only affected constraints are the following two:

$$\sum_{j'} \pi^*(s_{i',i}|u_{j'})q(u_{j'})v_{i',j'} \geq \sum_{j'} \pi^*(s_{i',i}|u_{j'})q(u_{j'})v_{i,j'} \quad (15)$$

$$\sum_{j'} \pi^*(s_{i'',i}|u_{j'})q(u_{j'})v_{i'',j'} \geq \sum_{j'} \pi^*(s_{i'',i}|u_{j'})q(u_{j'})v_{i,j'} \quad (16)$$

If changing from π^* to π' violates Inequality (15), Then we have

$$\sum_{j'} \pi'(s_{i',i}|u_{j'})q(u_{j'})v_{i',j'} < \sum_{j'} \pi'(s_{i',i}|u_{j'})q(u_{j'})v_{i,j'},$$

or equivalently,

$$\begin{aligned}
&\pi'(s_{i',i}|u_j)q(u_j)v_{i',j} + \sum_{j'\neq j} \pi'(s_{i',i}|u_{j'})q(u_{j'})v_{i',j'} \\
<&\pi'(s_{i',i}|u_j)q(u_j)v_{i,j} + \sum_{j'\neq j} \pi'(s_{i',i}|u_{j'})q(u_{j'})v_{i,j'}.
\end{aligned}$$

Plugging in the definition of π' yields:

$$\begin{aligned}
&\epsilon q(u_j)v_{i',j} + \sum_{j'} \pi^*(s_{i',i}|u_{j'})q(u_{j'})v_{i',j'} \\
<&\epsilon q(u_j)v_{i,j} + \sum_{j'} \pi^*(s_{i',i}|u_{j'})q(u_{j'})v_{i,j'}.
\end{aligned}$$

With Inequality (15), the above can happen for any small positive number ϵ only if Inequality (15) is actually an equation and $v_{i',j} < v_{i,j}$.

Similarly, if changing from π^* to π' violates Inequality (16), we must have that Inequality (16) holds as equality and that $v_{i'',j} > v_{i,j}$. □

A.7 Proof of Lemma 5

Proof. For any $\alpha, \alpha' \in [0, 1]$, we have

$$\begin{aligned}
&\alpha\textsc{Rev}(\pi_\alpha) + (1-\alpha)\textsc{Privacy}(\pi_\alpha) \\
\geq&\alpha\textsc{Rev}(\pi_{\alpha'}) + (1-\alpha)\textsc{Privacy}(\pi_{\alpha'})
\end{aligned}$$

and

$$\begin{aligned}
&\alpha'\textsc{Rev}(\pi_{\alpha'}) + (1-\alpha')\textsc{Privacy}(\pi_{\alpha'}) \\
\geq&\alpha'\textsc{Rev}(\pi_\alpha) + (1-\alpha')\textsc{Privacy}(\pi_\alpha).
\end{aligned}$$

This is because π_α and $\pi_{\alpha'}$ are the solutions to $P(\alpha)$ and $P(\alpha')$, respectively. With slight re-arrangement, the above inequalities can be written as:

$$\begin{aligned}
&\alpha\left[\textsc{Rev}(\pi_\alpha) - \textsc{Rev}(\pi_{\alpha'})\right] \\
&+ (1-\alpha)\left[\textsc{Privacy}(\pi_\alpha) - \textsc{Privacy}(\pi_{\alpha'})\right] \geq 0
\end{aligned} \quad (17)$$

$$\begin{aligned}
&\alpha'\left[\textsc{Rev}(\pi_{\alpha'}) - \textsc{Rev}(\pi_\alpha)\right] \\
&+ (1-\alpha')\left[\textsc{Privacy}(\pi_{\alpha'}) - \textsc{Privacy}(\pi_\alpha)\right] \geq 0
\end{aligned} \quad (18)$$

Multiplying Equation (17) by $1-\alpha'$ and Equation (18) by $1-\alpha$, and then adding them together gives:

$$(\alpha - \alpha')\left[\textsc{Rev}(\pi_\alpha) - \textsc{Rev}(\pi_{\alpha'})\right] \geq 0.$$

It follows that $\textsc{Rev}(\pi_\alpha)$ is monotone increasing in α. With similar arguments, we can also show that $\textsc{Privacy}(\pi_\alpha)$ is monotone decreasing in α. □

A.8 Proof of Theorem 3

Proof. We first prove the first part of the statement. For any α, let π_α be the optimal solution to $P(\alpha)$. Define

$$\lambda = \frac{\textsc{Rev}(\pi_\alpha) - \textsc{Rev}_{base}}{\textsc{Rev}_M - \textsc{Rev}_{base}}. \tag{19}$$

The choice of λ clearly satisfies Equation (10). Now we show it also satisfies Equation (11). Define $\hat{\pi} = \lambda\pi_1 + (1-\lambda)\pi_0$. According to the choice of λ (Equation (19)) and the linearity of $\textsc{Rev}(\pi)$, we have

$$\begin{aligned}\textsc{Rev}(\pi_\alpha) =& \lambda\textsc{Rev}(\pi_1) - \lambda\textsc{Rev}(\pi_0) + \textsc{Rev}(\pi_0)\\ =& \lambda\textsc{Rev}(\pi_1) + (1-\lambda)\textsc{Rev}(\pi_0)\\ =& \textsc{Rev}(\hat{\pi}).\end{aligned}$$

Therefore,

$$\begin{aligned}\textsc{Privacy}(\pi_\alpha) =& \frac{J_\alpha(\pi_\alpha) - \alpha\textsc{Rev}(\pi_\alpha)}{1-\alpha}\\ \geq& \frac{J_\alpha(\hat{\pi}) - \alpha\textsc{Rev}(\pi_\alpha)}{1-\alpha}\\ =& \frac{\alpha\textsc{Rev}(\hat{\pi}) + (1-\alpha)\textsc{Privacy}(\hat{\pi}) - \alpha\textsc{Rev}(\pi_\alpha)}{1-\alpha}\\ =& \textsc{Privacy}(\hat{\pi})\\ \geq& \lambda\textsc{Privacy}(\pi_1) + (1-\lambda)\textsc{Privacy}(\pi_0)\\ =& \lambda\textsc{Privacy}_{base} + (1-\lambda)\textsc{Privacy}_M,\end{aligned}$$

where the first inequality is due to the optimality of $J_\alpha(\pi_\alpha)$, and the second inequality is because of the concavity of the function $\textsc{Privacy}(\pi)$. One can easily check that the above inequality is equivalent to Equation (11).

Now we prove the second part of the statement. Let D denote the feasible region described by Program (9). Consider the following program:

$$\begin{aligned}\text{maximize:} \quad & \textsc{Privacy}(\pi)\\ \text{subject to:} \quad & \textsc{Rev}(\pi) \geq \textsc{Rev}(\hat{\pi})\\ & \pi \in D\end{aligned} \tag{20}$$

The program is clearly feasible ($\hat{\pi}$ is a feasible solution) and convex ($\textsc{Rev}(\pi)$ is linear and $\textsc{Privacy}(\pi)$ is concave). Therefore, the solution to the above program satisfies the two conditions described in the theorem. It suffices to show that the solution can also be obtained by solving Program 9 with a certain α.

We apply Lagrangian relaxation to the first constraint of the above program and obtain

$$\begin{aligned}\text{maximize:} \quad & \textsc{Privacy}(\pi) + \beta(\textsc{Rev}(\pi) - \textsc{Rev}(\hat{\pi}))\\ \text{subject to:} \quad & \pi \in D\end{aligned} \tag{21}$$

Similar to the proof of Lemma 2, we know that if β is set to be the optimal dual variable, the optimal solution to the relaxed program is also the optimal solution to Program 20. □

A.9 Proof of Theorem 4.

Proof. We compare both sides with $\sum_j q(u_j)\max_i v_{ij} - d(V)$.

We first show $\textsc{Rev}(\pi_0) \geq \sum_j q(u_j)\max_i v_{ij} - d(V)$. π_0 maximizes privacy and reveals no information. In this case, each buyer's posterior belief over user u_j is just the same as their prior $q(u_j)$. All the buyers just bid their expected valuation and the winner is always the same bidder, which we denote by $\hat{i}$. Then we have

$$\begin{aligned}&\sum_j q(u_j)\max_i v_{ij} - \textsc{Rev}(\pi_0)\\ =&\sum_j q(u_j)\max_i v_{ij} - \sum_k\sum_j \pi_0(s_k|u_j)q(u_j)v_{\hat{i},j}\\ =&\sum_j q(u_j)\left[\max_i v_{ij} - v_{\hat{i},j}\sum_k \pi_0(s_k|u_j)\right]\\ =&\sum_j q(u_j)\left[\max_i v_{ij} - v_{\hat{i},j}\right]\\ \leq&\sum_j q(u_j)\left[\max_i v_{ij} - \min_i v_{ij}\right]\\ =&\sum_j q(u_j)\max_{i,i'}|v_{ij} - v_{i'j}|\\ \leq&\sum_j q(u_j)\left[\max_j\{\max_{i,i'}|v_{ij} - v_{i'j}|\}\right]\\ =&\max_{i,i'}\{\max_j|v_{ij} - v_{i'j}|\}\sum_j q(u_j)\\ =&d(V).\end{aligned}$$

Now we prove $\textsc{Rev}_M \leq \sum_j q(u_j)\max_i v_{ij}$.

$$\begin{aligned}\textsc{Rev}_M =& \textsc{Rev}(\pi_1)\\ =&\sum_k\sum_j \pi_1(s_k|u_j)q(u_j)v_{i_2(s_k),j}\\ \leq&\sum_k\sum_j \pi_1(s_k|u_j)q(u_j)\max_i v_{ij}\\ =&\sum_j q(u_j)\max_i v_{ij}\sum_k \pi_1(s_k|u_j)\\ =&\sum_j q(u_j)\max_i v_{ij}.\end{aligned}$$

□

B Additional Experiment Results

In this section, we provide the experiment results for $m = 20$ on both synthetic and real data sets. The results for the independent valuation setting are shown in Figure 7 and 8. The results for the correlated valuation setting are shown in Figure4, Figure5,Figure 6, Figure9 and 10. All the figures within $m = 20$ show similar patterns to the cases with $m = 10$.

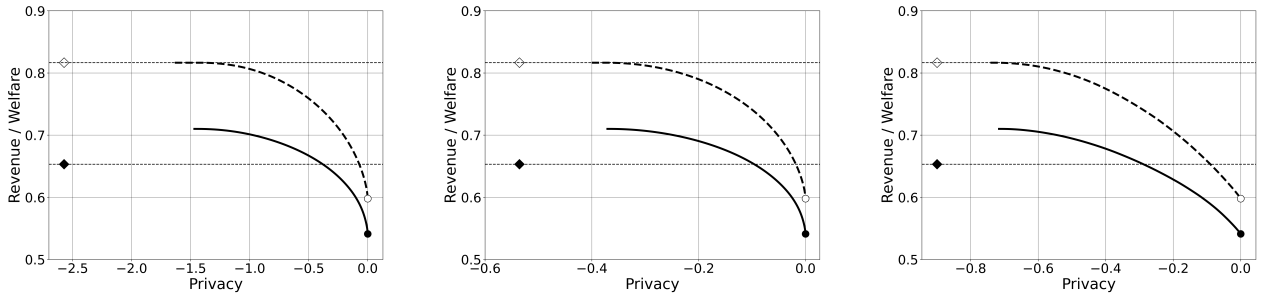

(a) KL divergence (b) JS divergence (c) TV divergence

Figure 7: Performance with $m = 20$ on the synthetic data set. The points indicated by diamond markers and round markers correspond to the results by revealing full information and no information (π_0), respectively.

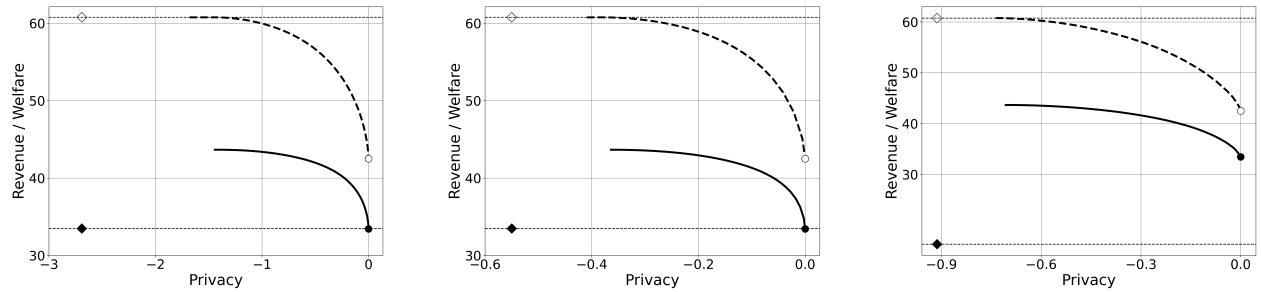

(a) KL divergence (b) JS divergence (c) TV divergence

Figure 8: Performance with $m = 20$ on the iPinYou data set. The points indicated by diamond markers and round markers correspond to the results by revealing full information and no information (π_0), respectively.

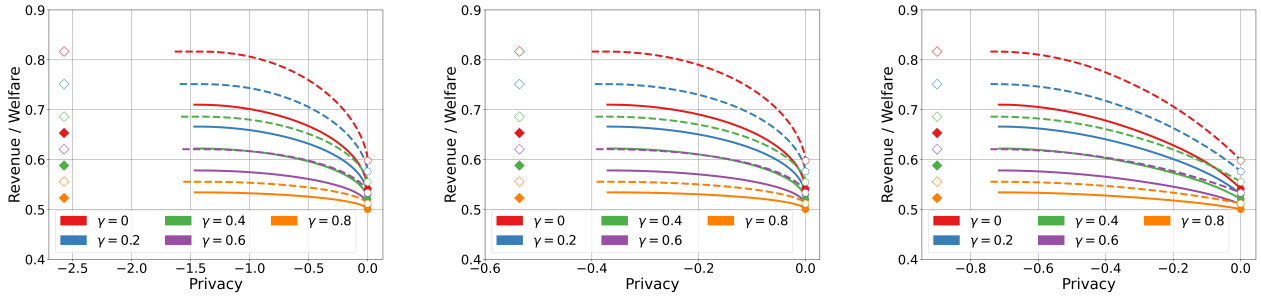

(a) KL divergence (b) JS divergence (c) TV divergence

Figure 9: Performance with $m = 20$ and five different weight coefficients on the synthetic data set.

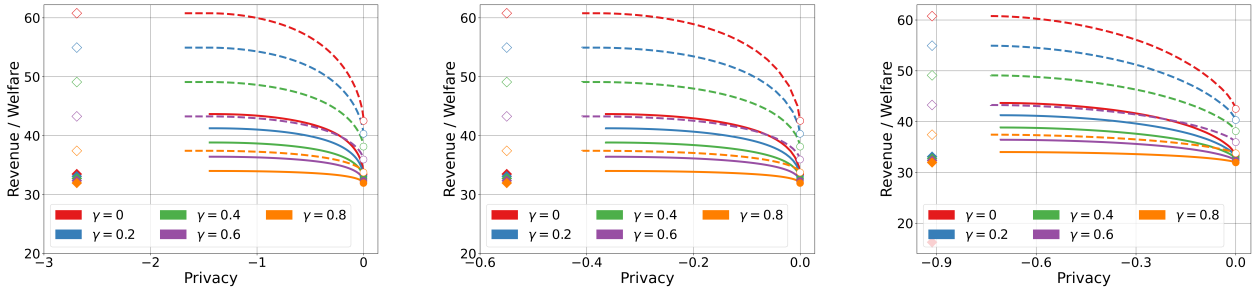

(a) KL divergence (b) JS divergence (c) TV divergence

Figure 10: Performance with $m = 20$ and five different weight coefficients on the iPinYou data set.