Privacy Auctions for Recommender Systems

Pranav Dandekar[1], Nadia Fawaz[2], and Stratis Ioannidis[2]

[1] Stanford University ppd@stanford.edu
[2] Technicolor {nadia.fawaz, stratis.ioannidis}@technicolor.com

Abstract. We study a market for private data in which a data analyst publicly releases a statistic over a database of private information. Individuals that own the data incur a cost for their loss of privacy proportional to the differential privacy guarantee given by the analyst at the time of the release. The analyst incentivizes individuals by compensating them, giving rise to a *privacy auction*. Motivated by recommender systems, the statistic we consider is a linear predictor function with publicly known weights. The statistic can be viewed as a prediction of the unknown data of a new individual, based on the data of individuals in the database. We formalize the trade-off between privacy and accuracy in this setting, and show that a simple class of estimates achieves an order-optimal trade-off. It thus suffices to focus on auction mechanisms that output such estimates. We use this observation to design a truthful, individually rational, proportional-purchase mechanism under a fixed budget constraint. We show that our mechanism is 5-approximate in terms of accuracy compared to the optimal mechanism, and that no truthful mechanism can achieve a $2-\varepsilon$ approximation, for any $\varepsilon > 0$.

1 Introduction

Recommender systems are ubiquitous on the Internet, lying at the heart of some of the most popular Internet services, including Netflix, Yahoo, and Amazon. These systems use algorithms to predict, *e.g.*, a user's rating for a movie, her propensity to click on an advertisement or to purchase a product online. By design, such prediction algorithms rely on access to large training datasets, typically comprising data from thousands (often millions) of individuals. This large-scale collection of user data has raised serious privacy concerns among researchers and consumer advocacy groups. Privacy researchers have shown that access to seemingly non-sensitive data (*e.g.*, movie ratings) can lead to the leakage of potentially sensitive information when combined with de-anonymization techniques [1]. Moreover, a spate of recent lawsuits [2,3,4] as well as behavioral studies [5] have demonstrated the increasing reluctance of the public to allow the unfettered collection and monetization of user data.

As a result, researchers and advocacy groups have argued in favor of legislation protecting individuals, by ensuring they can "opt-out" from data collection if they so desire [6]. However, a widespread restriction on data collection would be detrimental to profits of the above companies. One way to address this tension

between the value of data and the users' need for privacy is through *incentivization*. In short, companies releasing an individual's data ought to appropriately compensate her for the violation of her privacy, thereby incentivizing her consent to the release.

We study the issue of user incentivization through *privacy auctions*, as introduced by Ghosh and Roth [7]. In a privacy auction, a data analyst has access to a database $\mathbf{d} \in \mathbb{R}^n$ of private data d_i, $i = 1, \ldots, n$, each corresponding to a different individual. This data may represent information that is to be protected, such as an individual's propensity to click on an ad or purchase a product, or the number of visits to a particular website. The analyst wishes to publicly release an estimate $\hat{s}(\mathbf{d})$ of a statistic $s(\mathbf{d})$ evaluated over the database. In addition, each individual incurs a privacy cost c_i upon the release of the estimate $\hat{s}(\mathbf{d})$, and must be appropriately compensated by the analyst for this loss of utility. The analyst has a budget, which limits the total compensation paid out. As such, given a budget and a statistic s, the analyst must (a) solicit the costs of individuals c_i and (b) determine the estimate $\hat{s}$ to release as well as the appropriate compensation to each individual.

Ghosh and Roth employ *differential privacy* [8] as a principled approach to quantifying the privacy cost c_i. Informally, ensuring that $\hat{s}(\mathbf{d})$ is ϵ-differentially private with respect to individual i provides a guarantee on the privacy of this individual; a small ϵ corresponds to better privacy since it guarantees that $\hat{s}(\mathbf{d})$ is essentially independent of the individual's data d_i. Privacy auctions incorporate this notion by assuming that each individual i incurs a cost $c_i = c_i(\epsilon)$, that is a function of the privacy guarantee ϵ provided by the analyst.

1.1 Our Contribution

Motivated by recommender systems, we focus in this paper on a scenario where the statistic s takes the form of a *linear predictor*:

$$s(\mathbf{d}) := \langle \mathbf{w}, \mathbf{d} \rangle = \textstyle\sum_{i=1}^{n} w_i d_i, \tag{1}$$

where $\mathbf{w} \in \mathbb{R}^n$, is a publicly known vector of real (possibly negative) weights. Intuitively, the public weights w_i serve as measures of the similarity between each individual i and a new individual, outside the database. The function $s(\mathbf{d})$ can then be interpreted as a prediction of the value d for this new individual.

Linear predictors of the form (1) include many well-studied methods of statistical inference, such as the k-nearest-neighbor method, the Nadaranya-Watson weighted average, ridge regression, as well as support vector machines. We provide a brief review of such methods in Section 5. Functions of the form (1) are thus of particular interest in the context of recommender systems [9,10], as well as other applications involving predictions (*e.g.*, polling/surveys, marketing). In the sequel, we ignore the provenance of the public weights $\mathbf{w}$, keeping in mind that any of these methods apply. Our contributions are as follows:

1. **Privacy-Accuracy Trade-off.** We characterize the accuracy of the estimate $\hat{s}$ in terms of the *distortion* between the linear predictor s and $\hat{s}$ defined

as $\delta(s, \hat{s}) := \max_{\mathbf{d}} \mathbb{E}\left[|s(\mathbf{d}) - \hat{s}(\mathbf{d})|^2\right]$, *i.e.*, the maximum mean square error between $s(\mathbf{d})$ and $\hat{s}(\mathbf{d})$ over all databases $\mathbf{d}$. We define a *privacy index* $\beta(\hat{s})$ that captures the amount of privacy an estimator $\hat{s}$ provides to individuals in the database. We show that any estimator $\hat{s}$ with low distortion must also have a low privacy index (Theorem 1).

2. **Laplace Estimators Suffice.** We show that a special class of *Laplace estimators* [8,11] (*i.e.*, estimators that use noise drawn from a Laplace distribution), which we call Discrete Canonical Laplace Estimator Functions (DCLEFs), exhibits an order-optimal trade-off between privacy and distortion (Theorem 2). This allows us to restrict our focus on privacy auctions that output DCLEFs as estimators of the linear predictor s.
3. **Truthful, 5-approximate Mechanism, and Lower bound.** We design a *truthful*, *individually rational*, and *budget feasible* mechanism that outputs a DCLEF as an estimator of the linear predictor (Theorem 3). Our estimator's accuracy is a 5-approximation with respect to the DCLEF output by an optimal, individually rational, budget feasible mechanism. We also prove a lower bound (Theorem 4): there is no truthful DCLEF mechanism that achieves an approximation ratio $2 - \varepsilon$, for any $\varepsilon > 0$.

In our analysis, we exploit the fact that when $\hat{s}$ is a Laplace estimator minimizing distortion under a budget resembles the knapsack problem. As a result, the problem of designing a privacy auction that outputs a DCLEF $\hat{s}$ is similar in spirit to the knapsack auction mechanism [12]. However, our setting poses an additional challenge because the privacy costs exhibit *externalities*: the cost incurred by an individual is a function of which other individuals are being compensated. Despite the externalities in costs, we achieve the same approximation as the one known for the knapsack auction mechanism [12].

Due to space constraints we omit all proofs from this extended abstract, and refer the interested reader to the full version [13] of the paper.

1.2 Related Work

Privacy of behavioral data. Differentially-private algorithms have been developed for the release of several different kinds of online user behavioral data such as click-through rates and search-query frequencies [14], as well as movie ratings [15]. As pointed out by McSherry and Mironov [15], the reason why the release of such data constitutes a privacy violation is not necessarily that, *e.g.*, individuals perceive it as embarrassing, but that it renders them susceptible to *linkage* and *de-anonymization attacks* [1]. Such linkages could allow, for example, an attacker to piece together an individual's address stored in one database with his credit card number or social security number stored in another database. It is therefore natural to attribute a loss of utility to the disclosure of such data.

Privacy auctions. Quantifying the cost of privacy loss allows one to study privacy in the context of an economic transaction. Ghosh and Roth initiate this study of privacy auctions in the setting where the data is binary and the statistic reported is the sum of bits, *i.e.*, $d_i \in \{0, 1\}$ and $w_i = 1$ for all $i = 1, \ldots, n$ [7].

Unfortunately, the Ghosh-Roth auction mechanism cannot be readily generalized to asymmetric statistics such as (1), which, as discussed in Section 5, have numerous important applications including recommender systems. Our Theorems 1 and 2, which parallel the characterization of order-optimal estimators in [7], imply that to produce an accurate estimate of s, the estimator $\hat{s}$ *must provide different privacy guarantees to different individuals.* This is in contrast to the multi-unit procurement auction of [7]. In fact, as discussed the introduction, a privacy auction outputting a DCLEF $\hat{s}(\mathbf{d})$ has many similarities with a knapsack auction mechanism [12], with the additional challenge of externalities introduced by the Laplacian noise (see also Section 4).

Privacy and truthfulness in mechanism design. A series of interesting results follow an orthogonal direction, namely, on the connection between privacy and truthfulness when individuals have the ability to misreport their data. Starting with the work of McSherry and Talwar [16] followed by Nissim *et al* [17], Xiao [18] and most recently Chen *et al* [19], these papers design mechanisms that are simultaneously truthful and privacy-preserving (using differential privacy or other closely related definitions of privacy). As pointed out by Xiao [18], all these papers consider an *unverified* database, *i.e.*, the mechanism designer cannot verify the data reported by individuals and therefore must incentivize them to report truthfully. Recent work on truthfully eliciting private data through a *survey* [20,21] also fall under the unverified database setting [18]. In contrast, our setting, as well as that of Ghosh and Roth, is that of a *verified* database, in which individuals cannot lie about their data. This setting is particularly relevant to the context of online behavioral data: information on clicks, websites visited and products purchased is collected and stored in real-time and cannot be retracted after the fact.

Correlation between privacy costs and data values. An implicit assumption in privacy auctions as introduced in [7] is that the privacy costs c_i are *not* correlated with the data values d_i. This might not be true if, *e.g.*, the data represents the propensity of an individual to contract a disease. Ghosh and Roth [7] show that when the privacy costs are correlated to the data no individually rational direct revelation mechanism can simultaneously achieve non-trivial accuracy and differential privacy. As discussed in the beginning of this section, the privacy cost of the release of behavioral data is predominantly due to the risk of a linkage attack. It is reasonable in many cases to assume that this risk (and hence the cost of privacy loss) is not correlated to, *e.g.*, the user's movie ratings. Nevertheless, due to its importance in other settings such as medical data, more recent privacy auction models aim at handling such correlation [20,21,22]; we leave generalizing our results to such privacy auction models as future work.

2 Preliminaries

Let $[k] = \{1, \cdots, k\}$, for any integer $k > 0$, and define $\mathrm{I} := [R_{\min}, R_{\max}] \subset \mathbb{R}$ to be a bounded real interval. Consider a database containing the information of $n > 0$ individuals. In particular, the database comprises a vector $\mathbf{d}$, whose entries

$d_i \in \mathrm{I}$, $i \in [n]$, represent the private information of individual i. Each entry d_i is *a priori* known to the database administrator, and therefore individuals do not have the ability to lie about their private data. A data analyst with access to the database would like to publicly release an estimate of the statistic $s(\mathbf{d})$ of the form (1), *i.e.* $s(\mathbf{d}) = \sum_{i\in[n]} w_i d_i$, for some publicly known weight vector $\mathbf{w} = (w_1, \ldots, w_n) \in \mathbb{R}^n$. For any subset $H \subseteq [n]$, we define $w(H) := \sum_{i\in H} |w_i|$, and denote by $W := w([n]) = \sum_{i=1}^n |w_i|$ the ℓ_1 norm of vector $\mathbf{w}$. We denote the length of interval I by $\Delta := R_{\max} - R_{\min}$, and its midpoint by $\bar{R} := (R_{\min} + R_{\max})/2$. Without loss of generality, we assume that $w_i \neq 0$ for all $i \in [n]$; if not, since entries for which $w_i = 0$ do not contribute to the linear predictor, it suffices to consider the entries of $\mathbf{d}$ for which $w_i \neq 0$.

2.1 Differential Privacy and Distortion

Similar to [7], we use the following generalized definition of differential privacy:

Definition 1 (Differential Privacy). *A (randomized) function $f : \mathrm{I}^n \to \mathbb{R}^m$ is $(\epsilon_1, \ldots, \epsilon_n)$-differentially private if for each individual $i \in [n]$ and for any pair of data vectors $\mathbf{d}, \mathbf{d}^{(i)} \in \mathrm{I}^n$ differing in only their i-th entry, ϵ_i is the smallest value such that $\mathbb{P}[f(\mathbf{d}) \in S] \leq e^{\epsilon_i}\mathbb{P}[f(\mathbf{d}^{(i)}) \in S]$ for all $S \subset \mathbb{R}^m$.*

This definition differs slightly from the usual definition of ϵ-differential privacy [11], as the latter is stated in terms of the *worst case* privacy across all individuals. More specifically, according to the notation in [11], an $(\epsilon_1, \ldots, \epsilon_n)$-differentially private function is ϵ-differentially private, where $\epsilon = \max_i \epsilon_i$.

Given a deterministic function f, a well-known method to provide ϵ-differential privacy is to add random noise drawn from a Laplace distribution to this function [11]. This readily extends to $(\epsilon_1, \ldots, \epsilon_n)$-differential privacy.

Lemma 1 ([11]) *Consider a deterministic function $f : \mathrm{I}^n \to \mathbb{R}$. Define $\hat{f}(\mathbf{d}) := f(\mathbf{d}) + Lap(\sigma)$, where $Lap(\sigma)$ is a random variable sampled from the Laplace distribution with parameter σ. Then, $\hat{f}$ is $(\epsilon_1, \ldots, \epsilon_n)$-differentially private, where $\epsilon_i = S_i(f)/\sigma$, and $S_i(f) := \max_{\mathbf{d},\mathbf{d}^{(i)}\in\mathrm{I}^n} |f(\mathbf{d}) - f(\mathbf{d}^{(i)})|$, is the sensitivity of f to the i-th entry d_i, $i \in [n]$.*

Intuitively, the higher the variance σ of the Laplace noise added to f, the smaller ϵ_i, and hence, the better the privacy guarantee of $\hat{f}$. Moreover, for a fixed σ, entries i with higher sensitivity $S_i(f)$ receive a worse privacy guarantee (higher ϵ_i).

There is a natural tradeoff between the amount of noise added and the accuracy of the perturbed function $\hat{f}$. To capture this, we introduce the notion of *distortion* between two (possibly randomized) functions:

Definition 2 (Distortion). *Given two functions $f : \mathrm{I}^n \to \mathbb{R}$ and $\hat{f} : \mathrm{I}^n \to \mathbb{R}$, the* distortion, *$\delta(f, \hat{f})$, between f and $\hat{f}$ is given by*

$$\delta(f, \hat{f}) := \max_{\mathbf{d}\in\mathrm{I}^n} \mathbb{E}\left[|f(\mathbf{d}) - \hat{f}(\mathbf{d})|^2\right].$$

In our setup, the data analyst wishes to disclose an *estimator function* $\hat{s} : \mathrm{I}^n \to \mathbb{R}$ of the linear predictor s. Intuitively, a good estimator $\hat{s}$ should have a small distortion $\delta(s, \hat{s})$, while also providing good differential privacy guarantees.

2.2 Privacy Auction Mechanisms

Each individual $i \in [n]$ has an associated cost function $c_i : \mathbb{R}_+ \to \mathbb{R}_+$, which determines the cost $c_i(\epsilon_i)$ incurred by i when an $(\epsilon_1, \ldots, \epsilon_n)$-differentially private estimate $\hat{s}$ is released by the analyst. As in [7], we consider linear cost functions, *i.e.*, $c_i(\epsilon) = v_i\epsilon$, for all $i \in [n]$. We refer to v_i as the *unit-cost* of individual i. The unit-costs v_i are not *a priori* known to the data analyst. Without loss of generality, we assume throughout the paper that $v_1 \leq \ldots \leq v_n$.

Given a weight vector $\mathbf{w} = (w_1, \ldots, w_n) \in \mathbb{R}^n$, let M_s be a mechanism compensating individuals in $[n]$ for their loss of privacy from the release of an estimate $\hat{s}$ of the linear predictor $s(\mathbf{d})$. Formally, M_s takes as input a vector of reported unit-costs $\mathbf{v} = (v_1, \ldots, v_n) \in \mathbb{R}^n_+$ and a budget B, and outputs

1. a payment $p_i \in \mathbb{R}_+$ for every $i \in [n]$, and
2. an estimator function $\hat{s} : I^n \to \mathbb{R}_+$.

Assume that the estimator $\hat{s}$ satisfies $(\epsilon_1, \ldots, \epsilon_n)$-differential privacy. A mechanism is *budget feasible* if $\sum_{i\in[n]} p_i \leq B$, *i.e.*, the payments made by the mechanism are within the budget B. Moreover, a mechanism is *individually rational* if for all $i \in [n]$, $p_i \geq c_i(\epsilon_i) = v_i\epsilon_i$, *i.e.*, payments made by the mechanism exceed the cost incurred by individuals. Finally, a mechanism is *truthful* if for all $i \in [n]$, $p_i(v_i, v_{-i}) - v_i\epsilon_i(v_i, v_{-i}) \geq p_i(v'_i, v_{-i}) - v_i\epsilon_i(v'_i, v_{-i})$, *i.e.*, no individual can improve her utility by misreporting her private unit-cost.

2.3 Outline of our approach

We denote by $\delta_{M_s} := \delta(s, \hat{s})$ the distortion between s and the function output by the mechanism M_s. Ideally, a mechanism should output an estimator that has small distortion. However, the smaller the distortion, the higher the privacy violation and, hence, the more money the mechanism needs to spend. As such, the objective of this paper is to design a mechanism with minimal distortion, subject to the constraints of truthfulness, individual rationality, and budget feasibility.

To address this question, in Section 3, we first establish a privacy-distortion tradeoff for differentially-private estimators of the linear predictor. We then introduce a family of estimators, Discrete Canonical Laplace Estimator Functions (DCLEFs), and show that they achieve a near-optimal privacy-distortion tradeoff. This result allows us to limit our attention to DCLEF privacy auction mechanisms, *i.e.*, mechanisms that output a DCLEF $\hat{s}$. In Section 4, we present a mechanism that is truthful, individually rational, and budget feasible, while also being near-optimal in terms of distortion.

3 Privacy-Distortion Tradeoff and Laplace Estimators

Recall that a good estimator should exhibit low distortion and simultaneously give good privacy guarantees. In this section, we establish the privacy-distortion tradeoff for differentially-private estimators of the linear predictor. Moreover, we introduce a family of estimators that exhibits a near-optimal tradeoff between privacy and distortion. This will motivate our focus on privacy auction mechanisms that output estimators from this class in Section 4.

3.1 Privacy-Distortion Tradeoff

There exists a natural tension between privacy and distortion, as highlighted by the following two examples.

Example 1. Consider the estimator $\hat{s} := \bar{R}\sum_{i=1}^{n} w_i$, where recall that $\bar{R} = (R_{\min} + R_{\max})/2$. This estimator guarantees perfect privacy (*i.e.*, $\epsilon_i = 0$), for all individuals. However, $\delta(s, \hat{s}) = (W\Delta)^2/4$.

Example 2. Consider the estimator function $\hat{s} := \sum_{i=1}^{n} w_i d_i$. In this case, $\delta(s, \hat{s}) = 0$. However, $\epsilon_i = \infty$ for all $i \in [n]$.

In order to formalize this tension between privacy and distortion, we define the *privacy index* of an estimator as follows.

Definition 3 *Let $\hat{s} : \mathrm{I}^n \to \mathbb{R}$ be any $(\epsilon_1, \ldots, \epsilon_n)$-differentially private estimator function for the linear predictor. We define the* privacy index, *$\beta(\hat{s})$, of $\hat{s}$ as*

$$\beta(\hat{s}) := \max\left\{ w(H) : H \subseteq [n] \text{ and } \sum_{i \in H} \epsilon_i < 1/2 \right\}. \tag{2}$$

$\beta(\hat{s})$ captures the weight of the individuals that have been guaranteed good privacy by $\hat{s}$. Next we characterize the impossibility of having an estimator with a low distortion but a high privacy index. Note that for Example 1, $\beta(\hat{s}) = W$, *i.e.*, the largest value possible, while for Example 2, $\beta(\hat{s}) = 0$. We stress that the selection of $1/2$ as an upper bound in (2) is arbitrary; Theorems 1 and 2 still hold if another value is used, though the constants involved will differ.

Our first main result establishes a trade-off between the privacy index and the distortion of an estimator.

Theorem 1 (Trade-off between Privacy-index and Distortion) *Let $0 < \alpha < 1$. Let $\hat{s} : \mathrm{I}^n \to \mathbb{R}$ be an arbitrary estimator function for the linear predictor. If $\delta(s, \hat{s}) \leq (\alpha W\Delta)^2/48$ then $\beta(\hat{s}) \leq 2\alpha W$.*

In other words, if an estimator has low distortion, the weight of individuals with a good privacy guarantee (*i.e.*, a small ϵ_i) can be at most an α fraction of $2W$.

3.2 Laplace Estimator Functions

Consider the following family of estimators for the linear predictor $\hat{s} : \mathrm{I}^n \to \mathbb{R}$:

$$\hat{s}(\mathbf{d}; \mathbf{a}, \mathbf{x}, \sigma) := \sum_{i=1}^{n} w_i d_i x_i + \sum_{i=1}^{n} w_i a_i (1 - x_i) + \mathrm{Lap}(\sigma) \tag{3}$$

where $x_i \in [0,1]$, and each $a_i \in \mathbb{R}$ is a constant independent of the data vector $\mathbf{d}$. This function family is parameterized by $\mathbf{x}, \mathbf{a}$ and σ. The estimator $\hat{s}$ results from distorting s in two ways: (a) a randomized distortion by the addition of the Laplace noise, and (b) a deterministic distortion through a linear interpolation between each entry d_i and some constant a_i. Intuitively, the interpolation parameter x_i determines the extent to which the estimate $\hat{s}$ depends on entry d_i. Using Lemma 1 and the definition of distortion, it is easy to characterize the privacy and distortion properties of such estimators.

Lemma 2 *Given w_i, $i \in [n]$, let $s(\mathbf{d})$ be the linear predictor given by (1), and $\hat{s}$ an estimator of s given by (3). Then,*

1. *$\hat{s}$ is $(\epsilon_1, \ldots, \epsilon_n)$-differentially private, where $\epsilon_i = \frac{\Delta |w_i|\, x_i}{\sigma}$, $i \in [n]$.*
2. *The distortion satisfies $\delta(s, \hat{s}) \geq \left(\frac{\Delta}{2} \sum_{i=1}^{n} |w_i|(1 - x_i)\right)^2 + 2\sigma^2$, with equality attained when $a_i = \bar{R}$, for all $i \in [n]$.*

Note that the constants a_i do not affect the differential privacy properties of $\hat{s}$. Moreover, among all estimators with given $\mathbf{x}$, the distortion $\delta(s, \hat{s})$ is minimized when $a_i = \bar{R}$ for all $i \in [n]$. In other words, to minimize the distortion without affecting privacy, it is always preferable to interpolate between d_i and $\bar{R}$. This motivates us to define the family of Laplace estimator functions as follows.

Definition 4 *Given w_i, $i \in [n]$, the* Laplace estimator function family *(LEF) for the linear predictor s is the set of functions $\hat{s} : \mathrm{I}^n \to \mathbb{R}$, parameterized by $\mathbf{x}$ and σ, such that*

$$\hat{s}(\mathbf{d}; \mathbf{x}, \sigma) = \sum_{i=1}^{n} w_i d_i x_i + \bar{R} \sum_{i=1}^{n} w_i (1 - x_i) + Lap(\sigma) \tag{4}$$

We call a LEF *discrete* if $x_i \in \{0, 1\}$. Furthermore, we call a LEF *canonical* if the Laplace noise added to the estimator has a parameter of the form

$$\sigma = \sigma(\mathbf{x}) := \Delta \sum_{i=1}^{n} |w_i|(1 - x_i) \tag{5}$$

Recall that x_i controls the dependence of $\hat{s}$ on the entry d_i; thus, intuitively, the standard deviation of the noise added in a canonical Laplace estimator is proportional to the "residual weight" of data entries. Note that, by Lemma 2, the distortion of a canonical Laplace estimator $\hat{s}$ has the following simple form:

$$\delta(s, \hat{s}) = \frac{9}{4}\Delta^2 \Big(\sum_{i=1}^{n} |w_i|(1 - x_i) \Big)^2 = \frac{9}{4}\Delta^2 \big(W - \sum_{i=1}^{n} |w_i| x_i \big)^2. \tag{6}$$

Our next result establishes that there exists a discrete canonical Laplace estimator function (DCLEF) with a small distortion and a high privacy index.

Theorem 2 (DCLEFs suffice) *Let* $0 < \alpha < 1$. *Let*

$$\hat{s}^* := \operatorname*{argmax}_{\hat{s}:\delta(s,\hat{s}) \leq (\alpha W \Delta)^2/48} \beta(\hat{s})$$

be an estimator with the highest privacy index among all $\hat{s}$ *for which* $\delta(s,\hat{s}) \leq (\alpha W\Delta)^2/48$. *There exists a DCLEF* $\hat{s}^\circ : \mathrm{I}^n \to \mathbb{R}$ *such that* $\delta(s,\hat{s}^\circ) \leq (9/4)(\alpha W\Delta)^2$, *and* $\beta(\hat{s}^\circ) \geq \frac{1}{2}\beta(\hat{s}^*)$.

In other words, there exists a DCLEF that is within a constant factor, in terms of both its distortion and its privacy index, from an optimal estimator $\hat{s}^*$. Theorem 2 has the following immediate corollary:

Corollary 1 *Consider an arbitrary estimator* $\hat{s}$ *with distortion* $\delta(s,\hat{s}) < (W\Delta)^2/48$. *Then, there exists a DCLEF* $\hat{s}^\circ$ *such that* $\delta(s,\hat{s}^\circ) \leq 108\delta(s,\hat{s})$ *and* $\beta(\hat{s}^\circ) \geq \frac{1}{2}\beta(\hat{s})$.

Proof. Apply Theorem (2) with $\alpha = \sqrt{48\delta(s,\hat{s})}/(W\Delta)$. In particular, for this α and $\hat{s}$ as in the theorem statement, we have that $\hat{s}^* := \operatorname{argmax}_{\hat{s}':\delta(s,\hat{s}') \leq \delta(s,\hat{s})} \beta(\hat{s}')$, hence $\beta(\hat{s}^*) \geq \beta(\hat{s})$. Therefore, there exists a DCLEF $\hat{s}^\circ$ such that $\delta(s,\hat{s}^\circ) \leq (9/4)(\alpha W\Delta)^2 \leq 108\delta(s,\hat{s})$, and $\beta(\hat{s}^\circ) \geq \frac{1}{2}\beta(\hat{s}^*) \geq \frac{1}{2}\beta(\hat{s})$.

Theorems 1 and 2 imply that, when searching for estimators with low distortion and high privacy index, it suffices (up to constant factors) to focus on DCLEFs. Similar results were derived in [7] for estimators of unweighted sums of bits.

4 Privacy Auction Mechanism

Motivated by Theorems 1 and 2, we design a truthful, individually rational, budget-feasible DCLEF mechanism (*i.e.*, a mechanism that outputs a DCLEF) and show that it is 5-approximate in terms of accuracy compared with the optimal, individually rational, budget-feasible DCLEF mechanism. Note that a DCLEF is fully determined by the vector $\mathbf{x} \in \{0,1\}^n$. Therefore, we will simply refer to the output of the DCLEF mechanisms described below as $(\mathbf{x},\mathbf{p})$, as the latter characterize the released estimator and the compensations to individuals.

4.1 An Optimal DCLEF Mechanism

Consider the problem of designing a DCLEF mechanism M that is individually rational and budget feasible (but not necessarily truthful), and minimizes δ_M. Given a DCLEF $\hat{s}$, define $H(\hat{s}) := \{i : x_i = 1\}$ to be the set of individuals that receive non-zero differential privacy guarantees. Eq. (6) implies that $\delta(s,\hat{s}) = \frac{9}{4}\Delta^2(W - w(H(\hat{s})))^2$. Thus, minimizing $\delta(s,\hat{s})$ is equivalent to maxi-

mizing $w(H(\hat{s}))$. Let $(\mathbf{x}_{opt}, \mathbf{p}_{opt})$ be an optimal solution to the following problem:

$$\begin{aligned} \text{maximize} \quad & S(\mathbf{x}; \mathbf{w}) = \sum_{i=1}^{n} |w_i| x_i \\ \text{subject to:} \quad & p_i \geq v_i \epsilon_i(\mathbf{x}), \quad \forall i \in [n], \quad \text{(individual rationality)} \\ & \sum_{i=1}^{n} p_i \leq B \qquad\qquad \text{(budget feasibility)} \\ & x_i \in \{0, 1\}, \ \forall i \in [n] \quad \text{(discrete estimator function)} \end{aligned} \tag{7}$$

where, by Lemma 2 and (5),

$$\epsilon_i(\mathbf{x}) = \frac{\Delta |w_i| x_i}{\sigma(\mathbf{x})} = \frac{|w_i| x_i}{\sum_i |w_i|(1 - x_i)} \text{ (canonical property).} \tag{8}$$

A mechanism M_{opt} that outputs $(\mathbf{x}_{opt}, \mathbf{p}_{opt})$ will be an optimal, individually rational, budget feasible (but not necessarily truthful) DCLEF mechanism. Let $OPT := S(\mathbf{x}_{opt}; \mathbf{w})$ be the optimal objective value of (7). We use OPT as the benchmark to which we compare the (truthful) mechanism we design below. Without loss of generality, we make the following assumption:

Assumption 5 *For all* $i \in [n]$, $|w_i| v_i / (W - |w_i|) \leq B$.

Observe that if an individual i violates this assumption, then $c_i(\epsilon_i(\mathbf{x})) > B$ for any $\mathbf{x}$ output by a DCLEF mechanism that sets $x_i = 1$. In other words, no DCLEF mechanism (including M_{opt}) can compensate this individual within the analyst's budget and, hence, will set $x_i = 0$. Therefore, it suffices to focus on the subset of individuals for whom the assumption holds.

4.2 A Truthful DCLEF Mechanism

To highlight the challenge behind designing a truthful DCLEF mechanism, observe that if the privacy guarantees were given by $\epsilon_i(\mathbf{x}) = x_i$ rather than (8), the optimization problem (7) would be identical to the budget-constrained mechanism design problem for knapsack studied by Singer [12]. In the reverse-auction setting of [12], an auctioneer purchases items valued at fixed costs v_i by the individuals that sell them. Each item i is worth $|w_i|$ to the auctioneer, while the auctioneer's budget is B. The goal of the auctioneer is to maximize the total worth of the purchased set of items, *i.e.*, $S(\mathbf{x}; \mathbf{w})$. Singer presents a truthful mechanism that is 6-approximate with respect to OPT. However, in our setting, the privacy guarantees $\epsilon_i(\mathbf{x})$ given by (8) introduce *externalities* into the auction. In contrast to [12], the ϵ_i's couple the cost incurred by an individual i to the weight of other individuals that are compensated by the auction, making the mechanism design problem harder. This difficulty is overcome by our mechanism, which we call FairInnerProduct, described in Algorithm 1.

The mechanism takes as input the budget B, the weight vector $\mathbf{w}$, and the vector of unit-costs $\mathbf{v}$, and outputs a set $O \subset [n]$, that receive $x_i = 1$ in the

Algorithm 1 FairInnerProduct($\mathbf{v}, \mathbf{w}, B$)

Let k be the largest integer such that $\frac{B}{w([k])} \geq \frac{v_k}{W - w([k])}$.
Let $i^* := \text{argmax}_{i \in [n]} |w_i|$.
Let $\hat{p}$ be as defined in (9).
if $|w_{i^*}| > \sum_{i \in [k] \setminus \{i^*\}} |w_i|$ **then**
  Set $O = \{i^*\}$.
  Set $p_{i^*} = \hat{p}$ and $p_i = 0$ for all $i \neq i^*$.
else
  Set $O = [k]$.
  Pay each $i \in O,\ p_i = |w_i| \min\{\frac{B}{w([k])}, \frac{v_{k+1}}{W - w([k])}\}$, and for $i \notin O,\ p_i = 0$.
end if
Set $x_i = 1$ if $i \in O$ and $x_i = 0$ otherwise.

DCLEF, as well as a set of payments for each individual in O. Our construction uses a greedy approach similar to the Knapsack mechanism in [12]. In particular, it identifies users that are the "cheapest" to purchase. To ensure truthfulness, it compensates them within budget based on the unit-cost of the last individual that was not included in the set of compensated users. As in greedy solutions to knapsack, this construction does not necessarily yield a constant approximation w.r.t. OPT; for that, the mechanism needs to sometimes compensate only the user with the highest absolute weight $|w_i|$. In such cases, the payment of the user of the highest weight is selected so that she has no incentive to lie about here true unit cost.

Recall that $v_1 \leq \ldots \leq v_n$. The mechanism defines $i^* := \text{argmax}_{i \in [n]} |w_i|$ as the individual with the largest $|w_i|$, and k as the largest integer such that $\frac{B}{w([k])} \geq \frac{v_k}{W - w([k])}$. Subsequently, the mechanism either sets $x_i = 1$ for the first k individuals, or, if $|w_{i^*}| > \sum_{i \in [k] \setminus \{i^*\}} |w_i|$, sets $x_{i^*} = 1$. In the former case, individuals $i \in [k]$ are compensated *in proportion to their absolute weights* $|w_i|$. If, on the other hand, only $x_{i^*} = 1$, the individual i^* receives a payment $\hat{p}$ defined as follows: Let

$$S_{-i^*} := \Big\{ t \in [n] \setminus \{i^*\} : \frac{B}{\sum_{i \in [t] \setminus \{i^*\}} |w_i|} \geq \frac{v_t}{W - \sum_{i \in [t] \setminus \{i^*\}} |w_i|} \text{ and} \sum_{i \in [t] \setminus \{i^*\}} |w_i| \geq |w_{i^*}| \Big\}.$$

If $S_{-i^*} \neq \emptyset$, then let $r := \min\{i : i \in S_{-i^*}\}$. Define

$$\hat{p} := \begin{cases} B, & \text{if } S_{-i^*} = \emptyset \\ \frac{|w_{i^*}| v_r}{W - |w_{i^*}|}, & \text{otherwise} \end{cases} \tag{9}$$

The next theorem states that FairInnerProduct has the properties we desire.

Theorem 3 *FairInnerProduct is truthful, individually rational and budget feasible. It is 5-approximate with respect to OPT. Further, it is 2-approximate when all weights are equal.*

We note that the truthfulness of the knapsack mechanism in [12] is established via Myerson's characterization of truthful single-parameter auctions (*i.e.*, by

showing that the allocation is monotone and the payments are threshold). In contrast, because of the coupling of costs induced by the Laplace noise in DCLEFs, we are unable to use Myerson's characterization and, instead, give a direct argument about truthfulness.

We prove a 5-approximation by using the optimal solution of the fractional relaxation of (7). This technique can also be used to show that the knapsack mechanism in [12] is 5-approximate instead of 6-approximate. FairInnerProduct generalizes the Ghosh-Roth mechanism; in the special case when all weights are equal FairInnerProduct reduces to the Ghosh-Roth mechanism, which, by Theorem 3, is 2-approximate with respect to OPT. In fact, our next theorem states that the approximation ratio of a truthful mechanism is at least 2.

Theorem 4 (Hardness of Approximation) *For all $\varepsilon > 0$, there is no truthful, individually rational, budget feasible DCLEF mechanism that is also $2-\varepsilon$-approximate with respect to OPT.*

Our benchmark OPT is stricter than that used in [7]. In particular, Ghosh and Roth show that their mechanism is optimal among all truthful, individually rational, budget-feasible, and *envy-free* mechanisms. In fact, the example we use to show hardness of approximation is a uniform weight example, implying that the lower-bound also holds for uniform weight case. Indeed, the mechanism in [7] is 2-approximate with respect to OPT, although it is optimal among individually rational, budget feasible mechanisms that are also truthful *and* envy free.

5 Discussion on Linear Predictors

As discussed in the introduction, a statistic $s(\mathbf{d})$ of the form (1) can be viewed as a *linear predictor* and is thus of particular interest in the context of recommender systems. We elaborate on this interpretation in this section. Assume that each individual $i \in [n] = \{1, \ldots, n\}$ is endowed with a public vector $\mathbf{y}_i \in \mathbb{R}^m$, which includes m publicly known features about this individual. These could be, for example, demographic information such as age, gender or zip code, that the individual discloses in a public online profile. Note that, though features $\mathbf{y}_i$ are public, the data d_i is perceived as private.

Let $\mathbf{Y} = [\mathbf{y}_i]_{i\in[n]} \in \mathbb{R}^{n\times m}$ be a matrix comprising public feature vectors. Consider a new individual, not belonging to the database, whose public feature profile is $\mathbf{y} \in \mathbb{R}^m$. Having access to $\mathbf{Y}$, $\mathbf{d}$, and $\mathbf{y}$, the data analyst wishes to release a prediction for the unknown value d for this new individual. Below, we give several examples where this prediction takes the form $s(\mathbf{d}) = \langle \mathbf{w}, \mathbf{d}\rangle$, for some $\mathbf{w} = \mathbf{w}(\mathbf{y}, \mathbf{Y})$. All examples are textbook inference examples; we refer the interested reader to, for example, [23] for details.

k-Nearest Neighbors. In k-Nearest Neighbors prediction, the feature space $\mathbb{R}^m$ is endowed with a distance metric (*e.g.*, the ℓ_2 norm), and the predicted value is given by an average among the k nearest neighbors of the feature vector $\mathbf{y}$ of the new individual. *I.e.*, $s(\mathbf{d}) = \frac{1}{k}\sum_{i\in\mathcal{N}_k(\mathbf{y})} d_i$ where $N_k(\mathbf{y}) \subset [n]$ comprises the k individuals whose feature vectors y_i are closest to $\mathbf{y}$.

Nadaranya-Watson Weighted Average. The Nadaranya-Watson weighted average leverages all data in the database, weighing more highly data closer to $\mathbf{y}$. The general form of the prediction is $5s(\mathbf{d}) = \sum_{i=1}^{n} K(\mathbf{y}, \mathbf{y}_i) d_i / \sum_{i=1}^{n} K(\mathbf{y}, \mathbf{y}_i)$ where the *kernel* $K : \mathbb{R}^m \times \mathbb{R}^m \to \mathbb{R}_+$ is a function decreasing in the distance between its argument (*e.g.*, $K(\mathbf{y}, \mathbf{y}') = e^{-\|\mathbf{y}-\mathbf{y}'\|^2}$).

Ridge Regression. In ridge regression, the analyst first fits a linear model to the data, *i.e.*, solves the optimization problem

$$\min_{\mathbf{b}\in\mathbb{R}^m} \sum_{i=1}^{n} \left(d_i - \langle \mathbf{y}_i, \mathbf{b}\rangle\right)^2 + \lambda \|\mathbf{b}\|_2^2, \tag{10}$$

where $\lambda \geq 0$ is a regularization parameter, enforcing that the vector $\mathbf{b}$ takes small values. The prediction is then given by the inner product $\langle \mathbf{y}, \mathbf{b}\rangle$. The solution to (10) is given by $\mathbf{b} = (\mathbf{Y}^T\mathbf{Y} + \lambda\mathbf{I})^{-1}\mathbf{Y}^T\mathbf{d}$; as such, the predicted value for a new user with feature vector $\mathbf{y}$ is given by $s(\mathbf{d}) = \langle \mathbf{y}, \mathbf{b}\rangle = \mathbf{y}^T(\mathbf{Y}^T\mathbf{Y} + \lambda\mathbf{I})^{-1}\mathbf{Y}^T\mathbf{d}$.

In all three examples, the prediction $s(\mathbf{d})$ is indeed of the form (1). Note that the weights are non-negative in the first two examples, but may assume negative values in the last one.

6 Conclusion and Future Work

We considered the setting of an auction, where a data analyst wishes to buy, from a set of n individuals, the right to use their private data $d_i \in \mathbb{R}$, $i \in [n]$, in order to *cheaply* obtain an *accurate* estimate of a statistic. Motivated by recommender systems and, more generally, prediction problems, the statistic we consider is a linear predictor with publicly known weights. The statistic can be viewed as a prediction of the unknown data of a new individual based on the database entries. We formalized the trade-off between privacy and accuracy in this setting; we showed that obtaining an accurate estimate necessitates giving poor differential privacy guarantees to individuals whose cumulative weight is large. We showed that DCLEF estimators achieve an order-optimal trade-off between privacy and accuracy, and, consequently, it suffices to focus on DCLEF mechanisms. We use this observation to design a truthful, individually rational, budget feasible mechanism under the constraint that the analyst has a fixed budget. Our mechanism can be viewed as a proportional-purchase mechanism, *i.e.*, the privacy ϵ_i guaranteed by the mechanism to individual i is proportional to her weight $|w_i|$. We show that our mechanism is 5-approximate in terms of accuracy compared to an optimal (possibly non-truthful) mechanism, and that no truthful mechanism can achieve a $2 - \varepsilon$ approximation, for any $\varepsilon > 0$.

Our work is the first studying privacy auctions for asymmetric statistics, and can be extended in a number of directions. An interesting direction to investigate is characterizing the most general class of statistics for which truthful privacy auctions that achieve order-optimal accuracy can be designed. An orthogonal direction is to study the release of asymmetric statistics in other settings such as (a) using a different notion of privacy, (b) allowing costs to be correlated with the data values, and (c) survey-type settings where individuals first decide whether to participate and then reveal their private data.

References

1. Narayanan, A., Shmatikov, V.: Robust De-anonymization of Large Sparse Datasets. In: IEEE Symposium on Security and Privacy. (2008) 111–125
2. Netflix Privacy Litigation: www.videoprivacyclass.com.
3. Mello, J.P.: Facebook hit with lawsuit alleging privacy wrongs. PCWorld (May 18 2012)
4. Ribeiro, J.: Google faces class-action lawsuits over new privacy policy. PCWorld (Mar 22 2012)
5. Joseph, J., King, J., Hoofnagle, C.J., Bleakley, A., Hennessy, M.: Americans reject tailored advertising and three activities that enable it (2009) http://ssrn.com/abstract=1478214.
6. Mayer, J., Narayanan, A., Stamm, S.: Do not track: A universal third-party web tracking opt out. IETF Internet-Draft (March, 7th 2011)
7. Ghosh, A., Roth, A.: Selling privacy at auction. In: Proc. ACM EC. (2011) 199–208
8. Dwork, C., Mcsherry, F., Nissim, K., Smith, A.: Calibrating noise to sensitivity in private data analysis. In: Proc. Theory of Cryptography Conference. (2006)
9. Sarwar, B., Karypis, G., Konstan, J., Riedl, J.: Item-based collaborative filtering recommendation algorithms. In: Proceedings of the 10th international conference on World Wide Web. WWW '01, New York, NY, USA, ACM (2001) 285–295
10. Linden, G., Smith, B., York, J.: Amazon.com recommendations: item-to-item collaborative filtering. Internet Computing, IEEE **7**(1) (jan/feb 2003) 76 – 80
11. Dwork, C.: Differential privacy. In: Proc. ICALP. (2006) 1–12
12. Singer, Y.: Budget feasible mechanisms. In: Proc. FOCS. (2010)
13. Dandekar, P., Fawaz, N., Ioannidis, S.: Privacy auctions for recommender systems. CoRR **abs/1111.2885** (2012)
14. Korolova, A., Kenthapadi, K., Mishra, N., Ntoulas, A.: Releasing search queries and clicks privately. In: WWW. (2009)
15. McSherry, F., Mironov, I.: Differentially private recommender systems: building privacy into the net. In: Proc. ACM KDD. (2009) 627–636
16. McSherry, F., Talwar, K.: Mechanism design via differential privacy. In: Proc. FOCS. (2007)
17. Nissim, K., Smorodinsky, R., Tennenholtz, M.: Approximately optimal mechanism design via differential privacy. (2010)
18. Xiao, D.: Is privacy compatible with truthfulness? Technical Report 2011/005, Cryptology ePrint Archive (2011)
19. Chen, Y., Chong, S., Kash, I.A., Moran, T., Vadhan, S.P.: Truthful mechanisms for agents that value privacy. CoRR **abs/1111.5472** (2011)
20. Ligett, K., Roth, A.: Take it or Leave it: Running a Survey when Privacy Comes at a Cost. In: Proceedings of the 8th Workshop on Internet and Network Economics. WINE '12 (2012) To appear
21. Roth, A., Schoenebeck, G.: Conducting truthful surveys, cheaply. In: Proceedings of the 13th ACM Conference on Electronic Commerce. EC '12, New York, NY, USA, ACM (2012) 826–843
22. Fleischer, L., Lyu, Y.H.: Approximately optimal auctions for selling privacy when costs are correlated with data. CoRR **abs/1204.4031** (2012)
23. Hastie, T., Tibshirani, R., Friedman, J.: The Elements of Statistical Learning. 2nd edn. Springer (2009)