Bridging Central and Local Differential Privacy in Data Acquisition Mechanisms

Alireza Fallah
EECS Department
Massachusetts Institute of Technology
afallah@mit.edu

Ali Makhdoumi
Fuqua School of Business
Duke University
ali.makhdoumi@duke.edu

Azarakhsh Malekian
Rotman School of Management
University of Toronto
azarakhsh.malekian@rotman.utoronto.ca

Asuman Ozdaglar
EECS Department
Massachusetts Institute of Technology
asuman@mit.edu

Abstract

We study the design of optimal Bayesian data acquisition mechanisms for a platform interested in estimating the mean of a distribution by collecting data from privacy-conscious users. In our setting, users have heterogeneous sensitivities for two types of privacy losses corresponding to local and central differential privacy measures. The local privacy loss is due to the leakage of a user's information when she shares her data with the platform, and the central privacy loss is due to the released estimate by the platform to the public. The users share their data in exchange for a payment (e.g., through monetary transfers or services) that compensates for their privacy losses. The platform knows the distribution of privacy sensitivities but not their realizations, and must design a mechanism to solicit their preferences and then deliver both local and central privacy guarantees while minimizing the estimation error plus the expected payment to users. We first establish minimax lower bounds for the estimation error, given a vector of privacy guarantees, and show that a linear estimator is (near) optimal. We then turn to our main goal: designing an optimal data acquisition mechanism. We establish that the design of such mechanisms in a Bayesian setting (where the platform knows the distribution of users' sensitivities and not their realizations) can be cast as a nonconvex optimization problem. Additionally, for the class of linear estimators, we prove that finding the optimal mechanism admits a Polynomial Time Approximation Scheme.

1 Introduction

Users' personal data are currently being utilized for personalized advertising, medical trials, targeted advertising, and recommendation systems, among others. The transaction of individual data is set to grow exponentially in the coming years, with more widespread applications of artificial intelligence (AI) and machine learning techniques. Even though it is widely accepted that users need to own their data (see, e.g., Posner and Weyl [2019], Kushmaro [2021], and WILL.I.AM [2019]), the impact of different market architectures on the design and operation of data markets are not clear: some users prefer to protect their own raw data while others expect companies to protect their data proactively (see, e.g., GDMA [2018]). In this paper, we consider the design of data acquisition mechanisms when users have heterogeneous privacy concerns and ask the following question:

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

What is the optimal data acquisition mechanism when users have heterogeneous privacy concerns regarding access to their raw data and the outcome of the platform's processing?

We use differential privacy to measure the two types of privacy losses. Informally, an estimator is called differentially private if its distribution over outputs is insensitive to the changes in a user's data.

In particular, we consider a platform whose goal is to estimate an underlying parameter of interest by collecting data from a set of users $\mathcal{N} = \{1, \dots, n\}$ who own a noisy version of the underlying parameter. For instance, consider a medical trial in which a hospital wants to collect users' data to estimate the efficacy of a drug. Each user has two types of privacy concerns: (i) local privacy concern that captures how much information their shared data reveal about their raw data, and (ii) central privacy concern that captures how much information the platform's output reveal about their raw data. We adopt local and central Rényi differential privacy to measure these two types of privacy losses. The reason for choosing Rényi differential privacy is twofold. First, our framework can cover a wide range of information measures by varying the Rényi divergence parameter. Second, it can be achieved by a Gaussian mechanism, which simplifies our analysis while capturing the main tradeoffs in the design of two-part data acquisition mechanisms.

Before formulating the platform's data acquisition problem, we derive optimal estimators for a given vector of heterogeneous local privacy loss levels. In particular, we establish a minimax lower bound for the estimation error and prove that first privatizing users' data by adding a properly designed Gaussian noise to them and then using a properly designed weighted average of these privatized data points achieves this lower bound. This result motivates us to consider the design of the optimal data acquisition mechanism for the class of linear estimators.

We then turn to our mechanism design problem. Each user has a heterogeneous preference regarding the importance of the above two privacy concerns. For instance, if a user fully trusts the platform, then the first type of concern lessens, and the main concern would be about the information revealed from the platform's estimate. On the other hand, if a user does not trust the platform at all, the first type of concern would be more than the second one. We model such a setting by assuming each user i has a privacy sensitivity $c_i \in [0, 1]$ that determines the relative weight she puts on the local privacy concern ($1 - c_i$ is the weight she puts on the central privacy concern). The utility of user i is the payment she receives from the platform (in exchange for sharing her data), minus c_i times her local privacy loss, and again, minus $1 - c_i$ times her central privacy loss. The platform does not know the value of c_i and (knowing its distribution) must design a (Bayesian) data acquisition mechanism to elicit the true privacy sensitivities (that guide the optimal choice of local and central privacy losses delivered to each user) and optimize its objective.

In particular, the platform designs a *two-part data acquisition mechanism* that comprises a payment scheme, a local privacy guarantee, and a central privacy guarantee as a function of the reported privacy sensitivity of users. The platform's goal is to minimize the sum of the mean estimation error and the expected total payment to users while satisfying the *incentive compatibility* and *individual rationality* constraints. Incentive compatibility ensures that users have no incentive to misreport their privacy sensitivity, and individual rationality ensures that the payment to users (and the delivered privacy guarantees) are such that users are willing to share their data with the platform.

The platform's problem is a functional optimization over three functions of the reported privacy sensitivities: payments, local privacy guarantees, and central privacy guarantees. We first find the payment function in terms of the local and central privacy guarantees by using the incentive compatibility and individual rationality constraints. This reduces the space of the platform's decision variables. We then focus on the Gaussian mechanisms and linear estimators (motivated by our minimax optimality result) and show that the platform's problem can be cast as an optimization problem that minimizes a non-convex objective (which depends on the *virtual cost* of users) for any reported vector of privacy sensitivities. This reformulation significantly reduces the space of decision variables that the platform needs to optimize. However, it still involves solving a non-convex optimization problem. We further use the structural properties of this non-convex optimization and use duality theory to develop a polynomial time algorithm to approximate the platform's problem. More precisely, we prove that the design of the optimal two-part data acquisition mechanism admits a Polynomial Time Approximation Scheme (PTAS).

The contribution of our work is threefold. First, we develop a minimax lower bound when users have heterogeneous local privacy losses and establish that a linear estimator (approximately) achieves this bound. Second, we develop a modeling framework for data acquisition mechanisms when users have heterogeneous concerns for both local and central privacy losses. Third, for any estimator and mechanism to deliver privacy guarantees, we characterize the design of the optimal two-part data acquisition mechanism as the solution to a point-wise optimization problem. Additionally, for the class of Gaussian mechanisms to deliver privacy guarantees and linear estimators, we develop an algorithm to approximately find the optimal data acquisition mechanism (despite the fact that the corresponding optimization is non-convex).

Related literature: Our paper relates to the literature on optimal data acquisition from privacy concerned users. There is a large body of work that use differential privacy to measure the privacy loss of users Ghosh and Roth [2011], Nissim et al. [2012], Nissim et al. [2014]. One of the earliest papers in the literature is Ghosh and Roth [2011], which study the design of mechanisms for collecting users' data when users incur some privacy cost from sharing their data. More specifically, Ghosh and Roth [2011] consider binary data (bit) with the platform's goal being to estimate the sum of user's data by using a differentially private and dominant strategy truthful mechanism. They study both the case when the user data and privacy parameter are independent (similar to our paper) and when they are correlated. In the independent case, they propose a mechanism that delivers a single privacy level to all users (as opposed to our setting that delivers heterogeneous privacy levels). For the correlated case, they prove an impossibility result for the existence of a truthful and individually rational mechanism.

Several works build on Ghosh and Roth [2011], extending it to take it or leave it offers Ligett and Roth [2012], strengthening the impossibility results Nissim et al. [2014], and studying the open question posed by Ghosh and Roth [2011] on whether a model with distributional assumption on users' costs and Bayesian mechanism design approach could be used to develop the optimal mechanism for collecting data with privacy guarantees (see, e.g., Liao et al. [2021] and Fallah et al. [2022]). In particular, Roth and Schoenebeck [2012], Chen et al. [2018], and Chen and Zheng [2019] tackle this problem by developing a randomized mechanism in which user's data is randomly included in the final estimator where the inclusion probability depends on the reported privacy costs of the users (as opposed to our setting in which the payments and privacy guarantees depend on the reported privacy sensitivity of all users). These papers do not use differential privacy to model privacy costs and instead use a menu of probability-price pairs to tune the privacy loss and the payment to each user (see also Pai and Roth [2013] for a survey). Similar to the above paper, we consider a setting in which the platform can verify the data of users. For instance, in the context of medical trials, this means that the users decide whether to participate in the medical trial and cannot change the samples they share. A different stream of the literature explores settings in which users can misreport their information Perote and Perote-Pena [2003], Dekel et al. [2010], Meir et al. [2012], Ghosh et al. [2014], Cai et al. [2015], Liu and Chen [2016, 2017].

Our paper differs from these works in three main ways. First, we assume prior information on user privacy sensitivities and focus on characterizing the optimal Bayesian incentive compatible mechanism. Second, we model a setting in which users have both local and central privacy concerns and explore the different privacy guarantees of these two types delivered by an optimal mechanism. Third, we assume that user data are drawn from the same underlying distribution that allows the platform to put differing weights on the data of users depending on their privacy sensitivity, leading to different privacy levels for participating users.

Finally, our paper relates to the literature on differential privacy. Pioneered by the seminal work of Dwork et al. [2006a,b], differential privacy has emerged as a prevalent framework for characterizing the privacy leakage of data oriented algorithms. More specifically, our paper is related to the private mean estimation considered by Duchi et al. [2013], Barber and Duchi [2014], Karwa and Vadhan [2017], Asoodeh et al. [2021], Kamath et al. [2019, 2020], Cummings et al. [2021], and Acharya et al. [2021]. Additionally, our paper relates to the stream of differential privacy literature that studies Rényi differential privacy (RDP) introduced by Bun and Steinke [2016] and Mironov [2017].

2 Problem Formulation

We consider a platform interested in estimating a parameter $\theta \in \mathbb{R}$ by collecting data of n users, indexed by $\mathcal{N} = \{1, \cdots, n\}$. For any $i \in \mathcal{N}$, we denote user i's data by $X_i \in \mathcal{X}$ and we assume X_i

is given by $X_i = \theta + Z_i$ where $Z_1, \cdots, Z_n$ are independent and identically distributed zero-mean random variables with variance VAR. To simplify the exposition, we further assume $|Z_i| \leq 1/2$ for any $i \in \mathcal{N}$.[1] Throughout, we use lower case letters to denote the realization of random variables. The platform's goal is to minimize the estimate's error by collecting data from privacy-concerned users. Therefore, the platform needs to incentivize them to share their data.

2.1 Local and central privacy losses

Before formalizing the utilities/objectives of the platform and the users, we define the notions of privacy losses that we adopt in this paper. In particular, we consider two different types of privacy losses that users suffer from. The first one is the privacy loss of a user when she shares her data (only) with the platform, and the second one is the privacy loss through the released estimate (to the public) by the platform. Depending on how different users trust the platform, they might care differently about these two privacy losses. For instance, if a user fully trusts the platform, then her main privacy concern would be the second one, while a user who does not trust the platform at all would be more concerned with the first one as the public only observes the aggregated estimate, as opposed to the platform which observes each user's (shared) data separately.

We use the differential privacy framework to quantify these privacy losses. Since differential privacy was introduced by Dwork et al., several variants of it have been also proposed. In particular, a popular one in the machine learning literature is Rényi differential privacy (RDP), introduced by Mironov [2017], which we also adopt in this paper. Let us first recall the definition of Rényi divergence.

Definition 1. *Let P and Q be two distributions over $\mathbb{R}$ with densities p and q. For any $\alpha \in (1, \infty]$, the Rényi α-divergence between P and Q is denoted by $D_\alpha(P||Q)$ and is given by*

$$D_\alpha(P||Q) := \frac{1}{\alpha - 1} \log \int \left(\frac{p(x)}{q(x)}\right)^\alpha q(x)dx.$$

For two random variables X and Y, $D_\alpha(X||Y)$ denotes the α-divergence between their distributions.

We next define two notions of differential privacy, known as central and local, to capture the two aforementioned types of privacy losses. Local differential privacy corresponds to the privacy loss of a user when she shares her data with the platform through a randomized mapping, known as a *channel*.

Definition 2. *Let $\varepsilon \geq 0$ and $\alpha \in (1, \infty]$. A randomized channel $\mathcal{C} : \mathcal{X} \to \mathbb{R}$ is locally (ε, α)-Rényi (differentially) private if for any $x, x' \in \mathcal{X}$,*

$$D_\alpha(\mathcal{C}(x)||\mathcal{C}(x')) \leq \varepsilon.$$

Central differential privacy corresponds to the other privacy loss mentioned above. It bounds the change in the distribution of the platform's output, i.e., the released estimate, by changing one user's data. We next provide the formal definition.

Definition 3. *Let $\boldsymbol{\varepsilon} = (\varepsilon_i)_{i=1}^n \in \mathbb{R}_+^n$ and $\alpha \in (1, \infty]$. A randomized algorithm $\mathcal{A} : \mathcal{X}^n \to \mathbb{R}$ is $(\boldsymbol{\varepsilon}, \alpha)$-Rényi (differentially) private if for any two datasets $x_{1:n}, x'_{1:n} \in \mathcal{X}^n$ that only differ in i-th coordinate (data of user i),*

$$D_\alpha(\mathcal{A}(x_{1:n})||\mathcal{A}(x'_{1:n})) \leq \varepsilon_i.$$

The customary approach to guarantee RDP is Gaussian mechanism in which a properly tuned zero-mean Gaussian noise is added to fulfill the required condition. The following lemma, adapted from Mironov [2017], allows us to characterize the Gaussian noise's variance for a privacy loss level.

Lemma 1. *For a function $f : \mathcal{X}^n \to \mathbb{R}$, we define its sensitivity with respect to i-th coordinate as*

$$L_i(f) := \sup \left\{|f(x_{1:n}) - f(x'_{1:n})| \ : \text{for all } x_{1:n} \text{ and } x'_{1:n} \text{ differing only at } i\text{-th coordinate}\right\}.$$

For any $\alpha \in (1, \infty]$, $\mathcal{A}(x_{1:n}) = f(x_{1:n}) + W$ with $W \sim \mathcal{N}(0, \sigma^2)$ is $\left((\frac{\alpha L_i(f)^2}{2\sigma^2})_{i=1}^n, \alpha\right)$-RDP.

For a given vector of local privacy losses $(\varepsilon_1^{(l)}, \ldots, \varepsilon_n^{(l)})$, a natural way to privately estimate the mean is by using a linear estimator with Gaussian mechanism which is given by

$$\hat{\theta}(x_1, \ldots, x_n) := \sum_{i=1}^n w_i \hat{x}_i \text{ where } \sum_{i=1}^n w_i = 1 \text{ and } \hat{x}_i = x_i + \mathcal{N}(0, \frac{\alpha}{2\varepsilon_i^{(l)}}) \text{ for all } i \in \mathcal{N}. \quad (1)$$

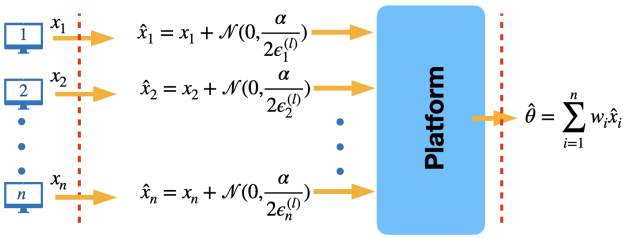

Figure 1: The interaction between the users and the platform in the two-part private data acquisition.

Figure 1 depicts this estimator. In particular, using Lemma 1, the local privacy delivered to user $i \in \mathcal{N}$ is $\varepsilon_i^{(l)}$ and the central privacy delivered to user $i \in \mathcal{N}$ is $\varepsilon_i^{(c)} = w_i^2 / \sum_{j=1}^n w_j^2 / \varepsilon_j^{(l)}$. In the next section, we focus on linear estimators in designing the optimal private data acquisition mechanism. We motivate such specification by showing that, for a given vector of local privacy losses $(\varepsilon_i^{(l)})_{i=1}^n$, a linear estimator is *optimal* with respect to mean square error. To formalize this statement, we first need to define the *minimax* estimation error as the notion of optimality. Let $\mathcal{P}$ be a class of distributions over $\mathcal{X}$. For any $P \in \mathcal{P}$, we denote its mean by $\theta(P)$. A $(\varepsilon_i^{(l)})_{i=1}^n$-locally RDP estimator can be cast as $\hat{\theta}((\mathcal{C}_i(x_i))_{i=1}^n)$, where $\mathcal{C}_i(.)$ is the randomized channel corresponding to user i. Let $\mathcal{Q}((\varepsilon_i^{(l)})_{i=1}^n)$ be the class of such $(\varepsilon_i^{(l)})_{i=1}^n$-locally RDP estimators. The minimax estimation error is

$$\mathcal{L}(\mathcal{P}, \mathcal{Q}, (\varepsilon_i^{(l)})_{i=1}^n) := \inf_{\hat{\theta}, \{\mathcal{C}_i\}_{i=1}^n \in \mathcal{Q}((\varepsilon_i^{(l)})_{i=1}^n)} \sup_{P \in \mathcal{P}} \mathbb{E}_{(X_i \sim P)_{i=1}^n, \hat{\theta}}[|\hat{\theta}((\mathcal{C}_i(X_i))_{i=1}^n) - \theta(P)|^2], \quad (2)$$

where the expectation is taken over both randomness of data and estimator (including private channels). In other words, the optimal estimator is the one that has the lowest worst case error among all estimators that satisfy the privacy requirements. We next state the optimality result.

Theorem 1. *Assume $\alpha \geq 2$ and $\varepsilon_i^{(l)} \leq 1$ for all i. Let $\mathcal{P}_1$ be the family of distributions over $[-\frac{1}{2}, \frac{1}{2}]$ and $\mathcal{C}_1, \cdots, \mathcal{C}_n$ be independent channels. Then, there exists a universal constant c such that*

$$\mathcal{L}(\mathcal{P}, \mathcal{Q}, (\varepsilon_i^{(l)})_{i=1}^n) \geq c \min(\frac{1}{\sum_{i=1}^n \varepsilon_i^{(l)}}, 1).$$

Furthermore, there exists a linear estimator with Gaussian mechanism such that

$$\mathbb{E}_{(X_i \sim P)_{i=1}^n, \hat{\theta}}[|\hat{\theta}((\mathcal{C}_i(X_i))_{i=1}^n) - \theta(P)|^2] \leq \mathcal{O}(1) \frac{\alpha}{\sum_{i=1}^n \varepsilon_i^{(l)}}.$$

We prove the lower bound by using the Le Cam's method Yu [1997] that reduces the problem of finding lower bounds to a hypothesis testing problem and requires bounding the total variation distance between the estimates when the input data points change. We then find it more convenient to bound the total variation in terms of Hellinger distance and use a series of inequalities to bound it.

2.2 Data acquisition mechanism with two-part privacy guarantees

We next describe the utility functions of the users and the platform and then formulate the platform's optimal data acquisition mechanism. As we described earlier, each user suffers from two privacy losses when sharing her data. The first one is a central privacy loss because of the leakage of her information through the platform's output. The second one is a local privacy loss because of the leakage of her information through the raw data that she shares with the platform. Each user has a heterogeneous *privacy sensitivity* for these two types of privacy losses. To model such heterogeneity, for each $i \in \mathcal{N}$, we let $c_i \in [0, 1]$ be her *relative local privacy sensitivity*, representing the relative weight that user i assigns to the (per unit cost of) local privacy loss. We also let $1 - c_i$ be her *relative central privacy sensitivity*, representing the relative weight that user i assigns to the (per unit cost of) central privacy loss. Therefore, $c_i \approx 1$ implies that user i suffers a higher loss of privacy by sharing her raw data with the platform (local privacy loss) compared to her loss from the platform's output

[1]This assumption does not have any fundamental impact on the results and is made to simplify the notations.

(central privacy loss). Differently, $c_i \approx 0$ implies that the user suffers a smaller loss of privacy by sharing her raw data with the platform compared to her loss from the platform's output. In what follows, we use the term *privacy sensitivity* instead of relative local privacy sensitivity.

For each $i \in \mathcal{N}$, the privacy sensitivity c_i is independently drawn from a publicly known distribution whose support is $[0, 1]$ with cumulative distribution and probability density functions $F_i(\cdot)$ and $f_i(\cdot)$. We also let $\mathbf{c} = (c_1, \ldots, c_n)$ be the vector of privacy sensitivities. The privacy sensitivity of each user is her private information, i.e., the platform does not know it. This is because individuals have different views regarding how trustworthy the platform is in protecting their raw data.

The platform's objective is to design a mechanism to collect users' data by paying them to compensate for their privacy losses without knowing the privacy sensitivity of users. To introduce the platform's objective formally, we adopt the formalism of Bayesian mechanism design pioneered by Myerson [1981]. More specifically, the platform designs and announces a payment function, a local privacy loss function, and a central privacy loss function that are mappings from the reported privacy sensitivities of users. The users then report their privacy sensitivities (which may or may not be truthful). Based on the payment function, the platform compensates the users (the compensation could be monetary or some free or discounted service provided to the user). Based on the local and central privacy functions, the platform designs randomized channels and randomized estimation algorithms that deliver the guaranteed local and central privacy losses while minimizing the sum of the mean squared error and the total expected payments. Given this interaction, we next formally introduce a data acquisition mechanism with two-part data privacy guarantees.

Definition 4 (two-part private data acquisition mechanism)**.** We call the tuple $(\hat{\theta}, \boldsymbol{\varepsilon}^{(l)}, \boldsymbol{\varepsilon}^{(c)}, \mathbf{t})$ a *two-part private data acquisition mechanism* where

1. For all $i \in \mathcal{N}$, $\varepsilon_i^{(l)} : \mathbb{R}_+^n \to \mathbb{R}_+$ is a function that maps the vector of privacy sensitivities $\mathbf{c}$ to a local privacy loss for user i, $\varepsilon_i^{(l)}(\mathbf{c})$, with $\boldsymbol{\varepsilon}^{(l)}(\cdot) = (\varepsilon_i^{(l)}(\cdot))_{i=1}^n$.

2. For all $i \in \mathcal{N}$, $\varepsilon_i^{(c)} : \mathbb{R}_+^n \to \mathbb{R}_+$ is a function that maps the vector of privacy sensitivities $\mathbf{c}$ to a central privacy loss for user i, $\varepsilon_i^{(c)}(\mathbf{c})$, with $\boldsymbol{\varepsilon}^{(c)}(.) = (\varepsilon_i^{(c)}(\cdot))_{i=1}^n$.

3. $\hat{\theta} : \mathcal{X}^n \times \mathbb{R}_+^n \times \mathbb{R}_+^n \to \mathbb{R}$ is a $(\boldsymbol{\varepsilon}^{(c)}(\mathbf{c}), \alpha)$-Rényi differentially private estimator that maps acquired locally $(\varepsilon_i^{(l)}(\mathbf{c}), \alpha)$-Rényi differentially private data of user i for $i \in \mathcal{N}$ to an estimate $\hat{\theta}(\mathbf{x}, \boldsymbol{\varepsilon}^{(l)}(\mathbf{c}), \boldsymbol{\varepsilon}^{(c)}(\mathbf{c}))$.

4. For all $i \in \mathcal{N}$, $t_i : \mathbb{R}_+^n \to \mathbb{R}_+$ is a function that maps the vector of privacy sensitivities $\mathbf{c}$ to a payment for user i, $t_i(\mathbf{c})$, with $\boldsymbol{t}(.) = (t_i(\cdot))_{i=1}^n$.

Notice that we have not specified the estimator and the mechanisms that delivers (local and central) Rényi differential privacy. In the rest of this subsection, we introduce the utilities and the platform's problem for a general estimators and mechanisms to deliver differential privacy. Later, we focus on linear estimator and Gaussian mechanisms and explicitly solve the platform's problem.

Each user that participates in a two-part private data acquisition mechanism suffers from both the local and central privacy losses and need to be compensated by the platform. In particular, the utility of user i from participation when her privacy sensitivity is c_i and she reports c_i' is given by

$$u_i(\boldsymbol{\varepsilon}^{(l)}(c_i', \mathbf{c}_{-i}), \boldsymbol{\varepsilon}^{(c)}(c_i', \mathbf{c}_{-i}), \mathbf{t}, \hat{\theta}) = \mathbb{E}_{\mathbf{c}_{-i}}[t_i(\mathbf{c}_{-i}, c_i')) - c_i \varepsilon_i^{(l)}(\mathbf{c}_{-i}, c_i') - (1 - c_i)\varepsilon_i^{(c)}(\mathbf{c}_{-i}, c_i')],$$

where the term $t_i(\mathbf{c}_{-i}, c_i'))$ is the payment from the platform, the term $c_i \varepsilon_i^{(l)}(\mathbf{c}_{-i}, c_i')$ is the relative local privacy sensitivity of the user multiplied by her local privacy loss, and the term $(1-c_i)\varepsilon_i^{(c)}(\mathbf{c}_{-i}, c_i')$ is her relative central privacy sensitivity multiplied by her central privacy loss. A user $i \in \mathcal{N}$ that does not participate in the mechanism neither compromises her privacy nor gets a compensation. Therefore, the utility of a user who does not participate in the mechanism becomes 0.

The goal of the platform is to minimize the sum of the mean squared error and the overall payment to users. We let $\gamma \in \mathbb{R}_+$ represents the relative weight of the mean estimation error and the payments in

the platform's objective.[2] Therefore, the platform's objective is

$$\mathbb{E}_{\mathbf{c}}[\gamma \text{MSE}(\boldsymbol{\varepsilon}^{(l)}(\mathbf{c}), \boldsymbol{\varepsilon}^{(c)}(\mathbf{c}), \hat{\theta}) + \sum_{i=1}^{n} t_i(\mathbf{c})],$$

where the first term is the mean square error of estimator $\hat{\theta}$ given reported vector of privacy sensitivity and resulting local and central privacy losses $\boldsymbol{\varepsilon}^{(l)}$ and $\boldsymbol{\varepsilon}^{(c)}$, i.e.,

$$\text{MSE}(\boldsymbol{\varepsilon}^{(l)}(\mathbf{c}), \boldsymbol{\varepsilon}^{(c)}(\mathbf{c}), \hat{\theta}) = \mathbb{E}_{\mathbf{x}}[|\hat{\theta}(\hat{\mathbf{x}}, \boldsymbol{\varepsilon}^{(l)}, \boldsymbol{\varepsilon}^{(c)}) - \theta|^2].$$

Also, each summand of the second term is the compensation that the platform gives to a user to incentivize her to participate and report her privacy sensitivity truthfully.

In Appendix we prove that, similar to the classical mechanism design setting, *revelation principle* holds. This means that there is no loss of generality in focusing on the class of direct incentive compatible mechanisms, meaning the platform's optimization problem can be written as

$$\min_{\boldsymbol{\varepsilon}^{(l)}(\cdot), \boldsymbol{\varepsilon}^{(c)}(\cdot), \mathbf{t}(\cdot)} \quad \mathbb{E}_{\mathbf{c}}[\gamma \text{MSE}(\boldsymbol{\varepsilon}^{(l)}(\mathbf{c}), \boldsymbol{\varepsilon}^{(c)}(\mathbf{c}), \hat{\theta}) + \sum_{i=1}^{n} t_i(\mathbf{c})] \tag{3}$$

$$u_i(\boldsymbol{\varepsilon}^{(l)}(\mathbf{c}), \boldsymbol{\varepsilon}^{(c)}(\mathbf{c}), \mathbf{t}, \hat{\theta}) \geq u_i(\boldsymbol{\varepsilon}^{(l)}(c_i', \mathbf{c}_{-i}), \boldsymbol{\varepsilon}^{(c)}(c_i', \mathbf{c}_{-i}), \mathbf{t}, \hat{\theta}) \tag{4}$$

$$u_i(\boldsymbol{\varepsilon}^{(l)}(\mathbf{c}), \boldsymbol{\varepsilon}^{(c)}(\mathbf{c}), \mathbf{t}, \hat{\theta}) \geq 0 \quad \text{for all } i \in \mathcal{N}, c_i, \tag{5}$$

where the constraints in (4) represent the *incentive compatibility*. These constraints guarantee that that each user i has no incentive to misrepresent her privacy sensitivity when others report truthfully (reporting truthfully is an equilibrium of the game among the users). Also, the constraints in (5) represent *individual rationality*, which ensures that each user receives a non-negative utility from participating in the platform's mechanism and sharing her data.

3 From the mechanism design problem to an optimization problem

For a given estimator $\hat{\theta}$, the platform's decision comprises the local and central privacy loss functions $\boldsymbol{\varepsilon}^{(l)}(\cdot)$ and $\boldsymbol{\varepsilon}^{(c)}(\cdot)$ together with the payment functions $\mathbf{t}(\cdot)$. We next show that this problem can be equivalently formulated as an optimization problem over the vector of local privacy losses and central privacy losses (as opposed to functions). In the rest of the paper, we impose the following assumption which is well-known in the mechanism design literature and simplifies the analysis.[3]

Assumption 1. *For any user $i \in \mathcal{N}$, the* virtual cost *defined as $\psi_i(c) = c + \frac{F_i(c)}{f_i(c)}$ is increasing in c, where $f_i(\cdot)$ and $F_i(\cdot)$ are probability density and cumulative distribution functions of c_i, respectively.*

The above assumption holds for a wide class of distributions such as the ones with log-concave density functions (e.g., uniform).

Theorem 2. *Suppose Assumption 1 holds. For a given estimator $\hat{\theta} : \hat{\mathcal{X}}^n \times \mathbb{R}_+^n \times \mathbb{R}_+^n \to \mathbb{R}$, in the optimal two-part data acquisition mechanism, for a given vector of reported privacy sensitivities $\mathbf{c}$, the local and central privacy losses are the solution of*

$$\min_{\{\boldsymbol{\varepsilon}^{(l)}\}_{i=1}^n, \{\boldsymbol{\varepsilon}^{(c)}\}_{i=1}^n} \quad \gamma \text{MSE}(\boldsymbol{\varepsilon}^{(l)}, \boldsymbol{\varepsilon}^{(c)}, \hat{\theta}) + \sum_{i=1}^{n} \varepsilon_i^{(l)} \psi_i(c_i) + \sum_{i=1}^{n} \varepsilon_i^{(c)} (1 - \psi_i(c_i)). \tag{6}$$

Proof Sketch of Theorem 2: We introduce the following *interim functions*

$$t_i(c_i) = \mathbb{E}_{\mathbf{c}_{-i}}[t(c_i, \mathbf{c}_{-i})], \quad \varepsilon_i^{(l)}(c_i) = \mathbb{E}_{\mathbf{c}_{-i}}[\varepsilon_i^{(l)}(c_i, \mathbf{c}_{-i})], \quad \text{and } \varepsilon_i^{(c)}(c_i) = \mathbb{E}_{\mathbf{c}_{-i}}[\varepsilon_i^{(c)}(c_i, \mathbf{c}_{-i})].$$

We first establish a *payment identity* that determines the optimal payment in terms of the optimal local and central privacy losses. In particular, by evaluating the first order condition corresponding to the incentive compatibility constraint (4), we establish that this constraint holds if and only if

$$t_i(c_i) = t_i(0) + \varepsilon_i^{(c)}(c_i) - \varepsilon_i^{(c)}(0) + c_i(\varepsilon_i^{(l)}(c_i) - \varepsilon_i^{(c)}(c_i)) - \int_0^{c_i} (\varepsilon_i^{(l)}(z) - \varepsilon_i^{(c)}(z)) dz,$$

[2] Notice that changing the parameter γ enables us to study a wide range of platform's objectives with differing relative weights between the estimation error and the total payments.

[3] Without this assumption, extending the results requires ironing technique of Myerson [1981].

and $\varepsilon_i^{(l)}(z) - \varepsilon_i^{(c)}(z)$ is weakly decreasing in z. We then plug in this payment identity back to the platform's objective, use the individual rationality constraint, and rewrite the platform's expected utility in terms of the privacy loss functions and the virtual cost of users. This is still a functional optimization problem in terms of $\varepsilon^{(l)}(\cdot)$ and $\varepsilon^{(c)}(\cdot)$. However, we establish that, under Assumption 1, we can solve this functional optimization point-wise (i.e., for any given $\mathbf{c}$). ■

Theorem 2 highlights the tradeoff in the platform's problem: by decreasing the local privacy loss, the second term of the objective decreases (this term corresponds to the payment to users) while the first term (i.e., the mean squared error) increases. The role of the central privacy loss is more nuanced, and there are two cases. If the coefficient $1 - \psi_i(c_i)$ is non-negative, by decreasing the central privacy loss, the third term of the objective decreases while the first term increases. If the coefficient $1 - \psi_i(c_i)$ is negative, increasing the central privacy loss decreases both the third term and the first term. However, we cannot increase the central privacy loss level without limits because the central privacy loss level is always below the local privacy loss level. Therefore, the platform's optimal mechanism should find the "right" balance between these terms.

4 Optimal mechanism with two-part privacy guarantees for linear estimators

For the rest of the paper, we focus on linear estimators with Gaussian mechanism described in Section 2.1. The following is a direct corollary of Theorem 2.

Corollary 1. *Suppose Assumption 1 holds. For any reported vector of privacy sensitivities* $\mathbf{c}$*, the optimal local privacy loss levels are* $\varepsilon_i^{(l)}(\mathbf{c}) = y_i^*$ *and the optimal central privacy loss levels are*

$$\varepsilon_i^{(c)} = w_i^{*2} / \sum_{j=1}^{n} \frac{w_j^{*2}}{y_j^*} \text{ where } (w_1^*, \ldots, w_n^*) \text{ and } (y_1^*, \ldots, y_n^*) \text{ are the optimal solution of}$$

$$\min_{\mathbf{w},\mathbf{y}} \ \text{VAR}\gamma \sum_{i=1}^{n} w_i^2 + \frac{\gamma\alpha}{2} \sum_{i=1}^{n} \frac{w_i^2}{y_i} + \sum_{i=1}^{n} (1 - \psi_i(c_i)) \frac{w_i^2}{\sum_{j=1}^{n} \frac{w_j^2}{y_j}} + \sum_{i=1}^{n} \psi_i(c_i) y_i \tag{7}$$

$$\text{s.t. } w_i, y_i \geq 0, \text{ for all } i \in \mathcal{N} \text{ and } \sum_{i=1}^{n} w_i = 1.$$

Let us highlight the difference between our characterization and that of classic mechanism design (e.g., Myerson [1981]). In classic mechanism design, the designer's problem becomes linear optimization. However, in our setting, the designer's problem is a non-linear and non-convex optimization. This makes the problem of finding the optimal two-part data acquisition mechanism challenging. Before addressing this computational challenge, let us revisit the form of the Gaussian mechanism that we have adopted: the platform adds Gaussian noise locally and then outputs a convex combination of the privatized users' data without adding any noise centrally. More specifically, one may guess that the platform may benefit by having a central noise added to the final output in addition to the local noises. In the following subsection, we establish that there is another Gaussian mechanism for any Gaussian mechanism that only adds local noises and achieves a weakly lower cost.

4.1 Optimality of having only local noises in the Gaussian mechanism

The platform has the opportunity of adding Gaussian noise to both the raw data of each user and the final estimator and ex-ante one may guess that it is optimal to use both of these instruments. However, as we establish next, interestingly, in the optimal two-part data acquisition mechanism, it is always optimal to only add noises locally.

For a given vector of local privacy losses $(\varepsilon_1^{(l)}, \ldots, \varepsilon_n^{(l)})$ and central privacy losses $(\varepsilon_1^{(c)}, \ldots, \varepsilon_n^{(c)})$, a Gaussian mechanism with both local and central noises is of the form

$$\hat{\theta}(x_1, \ldots, x_n) := \sum_{i=1}^{n} w_i \hat{x}_i + \mathcal{N}(0, \frac{\alpha}{2\varepsilon}) \text{ where } \sum_{i=1}^{n} w_i = 1 \text{ and } \hat{x}_i = x_i + \mathcal{N}(0, \frac{\alpha}{2\varepsilon_i^{(l)}}) \forall i \in \mathcal{N}.$$

Proposition 1. *In the optimal two-part data acquisition mechanism that adopts a Gaussian mechanism with both local and central noises, we have* $\varepsilon = \infty$*.*

Algorithm 1: Computing the optimal two-part private data acquisition mechanism

Input: The vector of privacy sensitivities $(c_1, \ldots, c_n)$

for $S \in \text{Grid}\left([\underline{S}, \bar{S}], \delta\right)$ **do**

Let

$$\nu_i = \frac{1}{\gamma \text{VAR} + (1 - \psi_i(c_i))/S}, \; \zeta_i = \frac{\nu_i}{\sum_j \nu_j}, \; \xi_i = \zeta_i \left(\sum_{j=1}^{n} \nu_j (\sqrt{\psi_j(c_j)} - \sqrt{\psi_i(c_i)}) \right).$$

Let $p = \left(\frac{\sum_{i=1}^n \zeta_i \sqrt{\psi_i(c_i)}}{S - \sum_{i=1}^n \sqrt{\psi_i(c_i)} \xi_i} \right)^2$.

Let

$$w_i(S) = \frac{\nu_i + \nu_i \sum_j \nu_j \sqrt{\psi_j(c_j) p}}{\sum_j \nu_j} - \nu_i \sqrt{\psi_i(c_i) p}, \quad y_i(S) = w_i(S) \sqrt{\frac{p}{\psi_i(c_i)}},$$

and $\text{OBJ}(S)$ be the objective of Problem (7) evaluated for this solution.

end

Output: $\{y_i(S^*), w_i(S^*)\}_{i=1}^n$, where $(S^*) = \arg\min_{(S)} \text{OBJ}(S)$.

Proposition 1 has an important implication in terms of the design of data market architecture when users have both central and local privacy costs: it is optimal to add noise locally! Adding a noise centrally to the final estimator has an advantage because the weights in the final estimator give the platform a lever to deliver heterogeneous central privacy guarantees to users. Despite this advantage, we establish that adding noise centrally is never optimal. This is because the platform prefers to add the noise locally to contribute to both central and local privacy guarantees delivered to users.

4.2 Computing the optimal privacy loss function

The implementation of the optimal two-part private data acquisition mechanism requires solving Problem (7), which is a non-convex program. However, we use the structure of the problem to develop a polynomial time algorithm to solve it approximately. To do so, we first replace $\sum_{i=1}^n w_i^2/y_i$ by an auxiliary variable S. Next, we consider the corresponding lagrangian problem. Using Karush-Kuhn-Tucker (KKT) conditions, we establish a number of relations between problems' parameters, S, and p, the lagrangian coefficient corresponding to $S = \sum_{i=1}^n w_i^2/y_i$. Furthermore, we develop upper and lower bounds for S. Finally, we do a grid search to find the approximate optimal solution.

Theorem 3. *For any vector of reported privacy sensitivities and $\epsilon > 0$, Algorithm 1 finds local privacy loss levels and the differentially private linear estimator of the two-part data acquisition mechanism whose cost (i.e., platform's objective) is at most $1 + \delta$ of the optimal cost in time* $\text{poly}(n, \frac{1}{\delta})$.

Notice that the approximation factor in Theorem 3 depends on the underlying parameters. Therefore, we have a Polynomial Time Approximation Scheme (PTAS) for finding the optimal two-part data acquisition mechanism in the class of linear estimators.

In the Appendix, we provide a case study with two users to illustrate the performance of the optimal two-part data acquisition mechanism in terms of the guaranteed privacy levels and payments as functions of the reported privacy sensitivities.

5 Conclusion

In this paper we develop a unified framework to study the design of data acquisition mechanisms when users have both local and central privacy concerns and are heterogeneous in how they value these two privacy concerns. We use Rényi differential privacy to measure the privacy loss of users and first establish a minimax lower bound that motivates us to focus on linear estimators. We then establish a point-wise optimization problem whose solution fully characterizes the optimal data acquisition mechanism that constitute a payment scheme to compensate users for their privacy losses, a local privacy guarantee, and a central privacy guarantee all as a function of users' preferences for

local and central privacy concerns. We then focus on linear estimators, motivated by our optimality results, and establish that, even though the corresponding optimization problem is non-convex, the platform's problem admits a Polynomial Time Approximation Scheme. Finally, we focused on data acquisition to estimate mean population. However, our framework is more general and allows for considering other (potentially vector) estimates. In particular, our Theorem 2 converts the data acquisition mechanism design problem into a (potentially) non-convex optimization problem.

6 Acknowledgement

Alireza Fallah acknowledges support from the Apple Scholars in Ai/ML PhD fellowship.

References

J. Acharya, Z. Sun, and H. Zhang. Differentially private assouad, fano, and le cam. In *Algorithmic Learning Theory*, pages 48–78. PMLR, 2021.

S. Asoodeh, M. Aliakbarpour, and F. P. Calmon. Local differential privacy is equivalent to contraction of an f-divergence. In *2021 IEEE International Symposium on Information Theory (ISIT)*, pages 545–550, 2021.

R. F. Barber and J. C. Duchi. Privacy and statistical risk: Formalisms and minimax bounds. *arXiv preprint arXiv:1412.4451*, 2014.

M. Bun and T. Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pages 635–658. Springer, 2016.

Y. Cai, C. Daskalakis, and C. Papadimitriou. Optimum statistical estimation with strategic data sources. In *Conference on Learning Theory*, pages 280–296. PMLR, 2015.

Y. Chen and S. Zheng. Prior-free data acquisition for accurate statistical estimation. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 659–677, 2019.

Y. Chen, N. Immorlica, B. Lucier, V. Syrgkanis, and J. Ziani. Optimal data acquisition for statistical estimation. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 27–44, 2018.

R. Cummings, V. Feldman, A. McMillan, and K. Talwar. Mean estimation with user-level privacy under data heterogeneity. In *NeurIPS 2021 Workshop Privacy in Machine Learning*, 2021.

O. Dekel, F. Fischer, and A. D. Procaccia. Incentive compatible regression learning. *Journal of Computer and System Sciences*, 76(8):759–777, 2010.

J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Local privacy and statistical minimax rates. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 429–438. IEEE, 2013.

C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor. Our data, ourselves: Privacy via distributed noise generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 486–503. Springer, 2006a.

C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006b.

A. Fallah, A. Makhdoumi, A. Malekian, and A. Ozdaglar. Optimal and differentially private data acquisition: Central and local mechanisms. *arXiv preprint arXiv:2201.03968*, 2022.

GDMA. Global data privacy: What the consumer really thinks. 2018. URL `https://dma.org.uk/uploads/misc/5b0522b113a23-global-data-privacy-report---final-2_5b0522b11396e.pdf`.

A. Ghosh and A. Roth. Selling privacy at auction. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 199–208, 2011.

A. Ghosh, K. Ligett, A. Roth, and G. Schoenebeck. Buying private data without verification. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 931–948, 2014.

G. Kamath, J. Li, V. Singhal, and J. Ullman. Privately learning high-dimensional distributions. In *Conference on Learning Theory*, pages 1853–1902. PMLR, 2019.

G. Kamath, V. Singhal, and J. Ullman. Private mean estimation of heavy-tailed distributions. In *Conference on Learning Theory*, pages 2204–2235. PMLR, 2020.

V. Karwa and S. Vadhan. Finite sample differentially private confidence intervals. *arXiv preprint arXiv:1711.03908*, 2017.

P. Kushmaro. Why data privacy is a human right (and what businesses should do about it). 2021. URL `https://www.forbes.com/sites/forbescommunicationscouncil/2021/06/07/why-data-privacy-is-a-human-right-and-what-businesses-should-do-about-it/?sh=c2a9684ec3ca`.

G. Liao, Y. Su, J. Ziani, A. Wierman, and J. Huang. The privacy paradox and optimal bias-variance trade-offs in data acquisition. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 689–689, 2021.

K. Ligett and A. Roth. Take it or leave it: Running a survey when privacy comes at a cost. In *International workshop on internet and network economics*, pages 378–391. Springer, 2012.

Y. Liu and Y. Chen. Learning to incentivize: Eliciting effort via output agreement. *arXiv preprint arXiv:1604.04928*, 2016.

Y. Liu and Y. Chen. Sequential peer prediction: Learning to elicit effort using posted prices. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.

R. Meir, A. D. Procaccia, and J. S. Rosenschein. Algorithms for strategyproof classification. *Artificial Intelligence*, 186:123–156, 2012.

I. Mironov. Rényi differential privacy. In *2017 IEEE 30th computer security foundations symposium (CSF)*, pages 263–275. IEEE, 2017.

R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

K. Nissim, C. Orlandi, and R. Smorodinsky. Privacy-aware mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 774–789, 2012.

K. Nissim, S. Vadhan, and D. Xiao. Redrawing the boundaries on purchasing data from privacy-sensitive individuals. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pages 411–422, 2014.

M. M. Pai and A. Roth. Privacy and mechanism design. *ACM SIGecom Exchanges*, 12(1):8–29, 2013.

J. Perote and J. Perote-Pena. The impossibility of strategy-proof clustering. *Economics Bulletin*, 4 (23):1–9, 2003.

E. A. Posner and E. G. Weyl. *Radical Markets*. Princeton University Press, 2019.

A. Roth and G. Schoenebeck. Conducting truthful surveys, cheaply. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 826–843, 2012.

WILL.I.AM. We need to own our data as a human right—and be compensated for it. 2019. URL `https://www.economist.com/open-future/2019/01/21/we-need-to-own-our-data-as-a-human-right-and-be-compensated-for-it`.

B. Yu. Assouad, fano, and le cam. In *Festschrift for Lucien Le Cam*, pages 423–435. Springer, 1997.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] Please check the conclusion section.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] Please check Sections 2,3, and 4.
   (b) Did you include complete proofs of all theoretical results? [Yes] Please check appendix
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

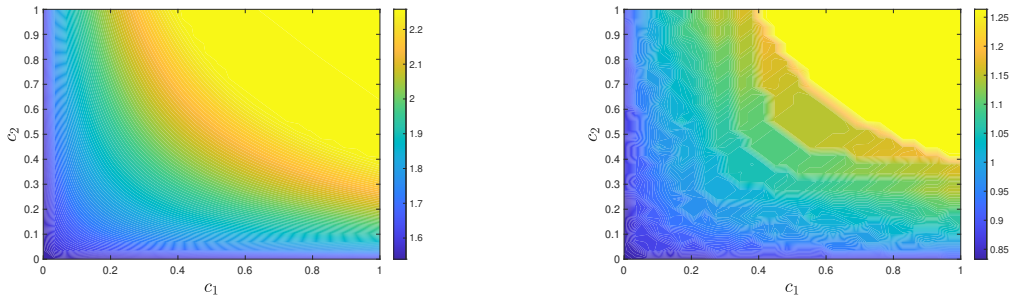

(a) (b)

Figure 2: (a) the platform's objective and (b) the platform's estimation error as a function of (c_1, c_2) for two users with $\gamma = 1$, $\alpha = 2$, and $\text{VAR} = 1/4$.

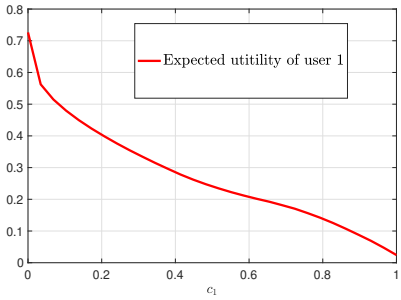

Figure 3: User 1's expected utility a function of c_1 for two users with $\gamma = 1$, $\alpha = 2$, and $\text{VAR} = 1/4$.

A Additional results and proofs

This appendix includes the additional results discussed in the text and the omitted proofs.

A.1 An illustrative example

Here, we illustrate the results in a simple setting with two users. For $i = 1, 2$, we let c_i be uniformly distributed over $[0, 1]$ so that the virtual value is $\psi_i(c_i) = 2c_i$. We also let $\gamma = 1$, $\alpha = 2$, and $\text{VAR} = 1/4$.

Figure 2 shows the platform's objective (i.e., the solution of Problem (7)) and its variance as a function of the privacy sensitivities. As we observe, higher privacy sensitivities (which means that users care more about the local privacy compared to central privacy) leads to a higher platform's cost and a higher estimation error. This is because guaranteeing local privacy is more demanding compared to central privacy. Figure 3 illustrates the expected utility of user 1 (similarly user 1) as a function of her privacy sensitivity. We observe that, unlike classical mechanism design settings, the utility is a continuous function of the user's type (as opposed to a threshold function). Again, we observe that higher privacy sensitivity implies better local privacy, which is more demanding and decreases the user's expected utility.

A.2 Proofs

Proof of Theorem 1

We establish the lower bound by using the Le Cam's method Yu [1997] which reduces the lower bound problem to a hypothesis testing problem between two distributions. To prove the lower bound, we need to show that for any $\hat{\theta} \in \mathcal{Q}((\varepsilon_i^{(l)})_{i=1}^n)$, we have

$$\sup_{P \in \mathcal{P}} \mathbb{E}_{(X_i \sim P)_{i=1}^n, \hat{\theta}} \left[\left| \hat{\theta}((\mathcal{C}_i(X_i))_{i=1}^n) - \theta(P) \right|^2 \right] \geq c \min \left(\frac{1}{\sum_{i=1}^n \varepsilon_i^{(l)}}, 1 \right). \tag{8}$$

To show this result, we replace $\sup_{P\in\mathcal{P}}$ by an average over two carefully chosen distributions in $\mathcal{P}$. More formally, let P_1 and P_2 be two distributions of choice in $\mathcal{P}$ with $\gamma := \frac{1}{2}|\theta(P_1) - \theta(P_2)|$. Note that

$$\sup_{P\in\mathcal{P}} \mathbb{E}_{(X_i\sim P)_{i=1}^n,\hat{\theta}} \left[\left|\hat{\theta}((\mathcal{C}_i(X_i))_{i=1}^n) - \theta(P)\right|^2\right] \geq \frac{1}{2}\sum_{j=1}^{2} \mathbb{E}_{(X_i\sim P_j)_{i=1}^n,\hat{\theta}} \left[\left|\hat{\theta}((\mathcal{C}_i(X_i))_{i=1}^n) - \theta(P_j)\right|^2\right]$$
$$= \frac{1}{2}\sum_{j=1}^{2} \mathbb{E}_{Y\sim Q_j}\left[|Y-\theta(P_j)|^2\right], \tag{9}$$

where, for any $j \in \{1,2\}$, Q_j denotes the distribution of $\hat{\theta}((\mathcal{C}_i(X_i))_{i=1}^n)$ when $X_1,\cdots,X_n$ are drawn from P_j. We next lower bound the right-hand side of (9) by Markov's inequality

$$\sup_{P\in\mathcal{P}} \mathbb{E}_{(X_i\sim P)_{i=1}^n,\hat{\theta}} \left[\left|\hat{\theta}((\mathcal{C}_i(X_i))_{i=1}^n) - \theta(P)\right|^2\right] \geq \gamma^2 \frac{1}{2}\sum_{j=1}^{2} \mathbb{P}\left(|Y-\theta(P_j)| \geq \gamma\right). \tag{10}$$

Now, consider a hypothesis testing problem with the goal of determining whether the underlying distribution is P_1 or P_2, given an observation of Y. One possible approach is to choose $j \in \{1,2\}$ for which $|Y-\theta(P_j)|$ is smaller. It can be shown that the probability of an incorrect estimate by this approach is upper bounded by $\frac{1}{2}\sum_{j=1}^{2}\mathbb{P}\left(|Y-\theta(P_j)| \geq \gamma\right)$ on the right-hand side of (10). Furthermore, a seminal result by Le Cam states that the infimum probability of incorrect decision among all possible mappings for the aforementioned hypothesis testing problem is given by $\frac{1}{2} - \frac{1}{2}\|Q_1 - Q_2\|_{\text{TV}}$. Therefore, we obtain the following lower bound

$$\mathcal{L}(\mathcal{P},\mathcal{Q},(\varepsilon_i^{(l)})_{i=1}^n) \geq \gamma^2\left(\frac{1}{2} - \frac{1}{2}\|Q_1 - Q_2\|_{\text{TV}}\right). \tag{11}$$

Next, we provide an upper bound on $\|Q_1 - Q_2\|_{\text{TV}}$. To do so, we use the connection between the total variation distance and the Hellinger distance. Recall that the Hellinger distance between two distributions μ and ν is given by

$$d_{\text{hel}}(\mu,\nu)^2 := \int (\sqrt{d\mu(x)} - \sqrt{d\nu(x)})^2.$$

Hellinger distance has a number of well-known desirable properties. In particular, we use the following two:

- For any two distributions μ_1 and μ_2, we have
$$\|\mu - \nu\|_{\text{TV}} \leq d_{\text{hel}}(\mu,\nu). \tag{12}$$
- Let $\mu := \mu_1 \times \cdots \times \mu_n$ and $\nu := \nu_1 \times \cdots \times \nu_n$. Then,
$$d_{\text{hel}}(\mu,\nu)^2 = 2 - 2\prod_{i=1}^{n}(1 - \frac{1}{2}d_{\text{hel}}(\mu_i,\nu_i)^2). \tag{13}$$

Let us go back to the problem of upper bounding $\|Q_1 - Q_2\|_{\text{TV}}$. The following lemma is the key result in our proof:

Lemma 2. *Let $\alpha \geq 2$ and suppose $\mathcal{C}(.) : \mathcal{X} - \to \mathbb{R}$ is an (ε,α)-RDP channel. For $j \in \{1,2\}$, let ν_j be the distribution of $\mathcal{C}(X)$ when $X \sim \mu_j$. Then,*

$$d_{hel}(\nu_1,\nu_2)^2 \leq 2(e^{\varepsilon}-1)\, d_{hel}(\mu_1,\mu_2)^2.$$

We defer the proof of Lemma 2 to the end of this section. Let us first complete the proof of lower bound using this lemma. Note that, by data processing inequality we have

$$\|Q_1 - Q_2\|_{\text{TV}} \leq \|(\mathcal{C}_i(X_i))_{i=1}^n - (\mathcal{C}_i(\tilde{X}_i))_{i=1}^n\|_{\text{TV}},$$

where $X_i \sim P_1$ and $\tilde{X}_i \sim P_2$. Next, using (12) and (13) implies

$$\|Q_1 - Q_2\|_{\text{TV}}^2 \leq \|(\mathcal{C}_i(X_i))_{i=1}^n - (\mathcal{C}_i(\tilde{X}_i))_{i=1}^n\|_{\text{TV}}^2$$
$$\leq d_{\text{hel}}((\mathcal{C}_i(X_i))_{i=1}^n, (\mathcal{C}_i(\tilde{X}_i))_{i=1}^n)^2$$
$$= 2 - 2\prod_{i=1}^{n}(1 - \frac{1}{2}d_{\text{hel}}(\mathcal{C}_i(X_i),\mathcal{C}_i(\tilde{X}_i))^2). \tag{14}$$

Next, note that, by Lemma 2, we have

$$d_{\text{hel}}(\mathcal{C}_i(X_i), \mathcal{C}_i(\tilde{X}_i))^2 \leq 2(e^{\varepsilon_i^{(l)}} - 1)\, d_{\text{hel}}(P_1, P_2)^2 \leq 4\varepsilon_i^{(l)}\, d_{\text{hel}}(P_1, P_2)^2,$$

where the last inequality follows from the fact that $\varepsilon_i^{(l)} \leq 1$. Plugging this back into (14), we obtain

$$\|Q_1 - Q_2\|_{\text{TV}}^2 \leq 2 - 2\prod_{i=1}^{n}(1 - 2\varepsilon_i^{(l)}\, d_{\text{hel}}(P_1, P_2)^2). \tag{15}$$

Next, note that, for nonnegative $y_1, \cdots, y_n$ we have

$$\prod_{i=1}^{n}(1 - y_i) \geq 1 - \sum_{i=1}^{n} y_i.$$

To show this, we can first verify it for $n = 2$ and then it is straightforward to show it for any n by induction. Using this inequality with $y_i = 2\varepsilon_i^{(l)}\, d_{\text{hel}}(P_1, P_2)^2$, we can further upper bound (15) by

$$\|Q_1 - Q_2\|_{\text{TV}}^2 \leq 4(\sum_{i=1}^{n} \varepsilon_i^{(l)})\, d_{\text{hel}}(P_1, P_2)^2. \tag{16}$$

Plugging this back into (11), we have

$$\mathcal{L}(\mathcal{P}, \mathcal{Q}, (\varepsilon_i^{(l)})_{i=1}^n) \geq \gamma^2 \left(\frac{1}{2} - \sqrt{\sum_{i=1}^{n} \varepsilon_i^{(l)}}\, d_{\text{hel}}(P_1, P_2) \right). \tag{17}$$

Next, we define P_1 and P_2 as

$$P_1(-1/2) = P_2(1/2) = \frac{1 + 2\gamma}{2}, \quad P_1(1/2) = P_2(-1/2) = \frac{1 - 2\gamma}{2}. \tag{18}$$

It is straightforward to verify that $|\theta(P_1) - \theta(P_2)| = 2\gamma$. Moreover, we have

$$d_{\text{hel}}(P_1, P_2)^2 = 2\left(\sqrt{\frac{1 + 2\gamma}{2}} - \sqrt{\frac{1 - 2\gamma}{2}} \right)^2 = 2(1 - \sqrt{1 - 4\gamma^2}) \leq 8\gamma^2,$$

where the last inequality follows from the fact that

$$1 - \sqrt{1 - 4\gamma^2} = \frac{4\gamma^2}{1 + \sqrt{1 - 4\gamma^2}} \leq 4\gamma^2.$$

Plugging this back into (17) implies

$$\mathcal{L}(\mathcal{P}, \mathcal{Q}, (\varepsilon_i^{(l)})_{i=1}^n) \geq \gamma^2 \left(\frac{1}{2} - \sqrt{8\sum_{i=1}^{n} \varepsilon_i^{(l)}}\, \gamma \right). \tag{19}$$

Finally, setting

$$\gamma = \min\left(\frac{1}{4\sqrt{8\sum_{i=1}^{n} \varepsilon_i^{(l)}}}, \frac{1}{2} \right)$$

completes the proof of lower bound.

To show the upper bound, first recall that a linear estimator with Gaussian mechanism is in the form of

$$\sum_{i=1}^{n} w_i \left(x_i + \mathcal{N}\left(0, \frac{\alpha}{2\varepsilon_i^{(l)}} \right) \right). \tag{20}$$

The mean square error of this estimator is given by

$$\sum_{i=1}^{n} w_i^2 (\text{VAR} + \frac{\alpha}{2\varepsilon_i^{(l)}}) \leq \alpha \sum_{i=1}^{n} \frac{w_i^2}{\varepsilon_i^{(l)}},$$

where the last inequality uses the fact that $\alpha \geq 2$ and $\varepsilon_i^{(l)} \leq 1$. Finally, setting

$$w_i = \frac{\varepsilon_i^{(l)}}{\sum_{j=1}^{n} \varepsilon_j^{(l)}}$$

gives us the desired upper bound.

Proof of Lemma 2

Note that

$$d_{\text{hel}}(\nu_1,\nu_2)^2 = \int (\sqrt{\nu_1(z)} - \sqrt{\nu_2(z)})^2 dz = \int \frac{(\nu_1(z)-\nu_2(z))^2}{(\sqrt{\nu_1(z)}+\sqrt{\nu_2(z)})^2} dz$$
$$\leq \int \frac{(\nu_1(z)-\nu_2(z))^2}{\nu_1(z)+\nu_2(z)} dz. \tag{21}$$

Note that, for any $j \in \{1,2\}$, we can cast $\nu_j(z)$ as

$$\nu_j(z) = \int_x \mathcal{C}(z|x) d\mu_j(x).$$

Moreover, for any $x' \in \mathcal{X}$, we have

$$\nu_1(z) - \nu_2(z) = \int_x \mathcal{C}(z|x)(d\mu_1(x) - d\mu_2(x)) = \int_x (\mathcal{C}(z|x) - \mathcal{C}(z|x'))(d\mu_1(x) - d\mu_2(x)).$$

Substituting these into (21), we have

$$d_{\text{hel}}(\nu_1,\nu_2)^2 \leq \int_z \frac{\left(\int_x (\mathcal{C}(z|x) - \mathcal{C}(z|x'))(d\mu_1(x) - d\mu_2(x))\right)^2}{\int_x \mathcal{C}(z|x)(d\mu_1(x) + d\mu_2(x))}\, dz. \tag{22}$$

Next, by Cauchy–Schwarz inequality, we obtain

$$\left(\int_x (\mathcal{C}(z|x) - \mathcal{C}(z|x'))(d\mu_1(x) - d\mu_2(x))\right)^2 \leq$$
$$\left(\int_x \mathcal{C}(z|x)(d\mu_1(x) + \mu_2(x))\right)\left(\int_x \frac{(\mathcal{C}(z|x) - \mathcal{C}(z|x'))^2}{\mathcal{C}(z|x)} \frac{(d\mu_1(x) - d\mu_2(x))^2}{d\mu_1(x) + d\mu_2(x)}\right).$$

Hence, using (22) and this inequality, we can further upper bound $d_{\text{hel}}(\nu_1,\nu_2)^2$ by

$$d_{\text{hel}}(\nu_1,\nu_2)^2 \leq \int_z \int_x \frac{(\mathcal{C}(z|x) - \mathcal{C}(z|x'))^2}{\mathcal{C}(z|x)} \frac{(d\mu_1(x) - d\mu_2(x))^2}{d\mu_1(x) + d\mu_2(x)} dz$$
$$= \int_x \left[\int_z \frac{(\mathcal{C}(z|x) - \mathcal{C}(z|x'))^2}{\mathcal{C}(z|x)}\, dz\right] \frac{(d\mu_1(x) - d\mu_2(x))^2}{d\mu_1(x) + d\mu_2(x)} \tag{23}$$

where the last equation follows from changing the order of integration using Fubini's theorem. Now, note that the first term on the right hand side of (23) can be cast as

$$\int_z \frac{(\mathcal{C}(z|x) - \mathcal{C}(z|x'))^2}{\mathcal{C}(z|x)}\, dz = \int_z \frac{\mathcal{C}(z|x')^2}{\mathcal{C}(z|x)}\, dz - 2\int_z \mathcal{C}(z|x')\, dz + \int_z \mathcal{C}(z|x)\, dz$$
$$= \exp(D_2(\mathcal{C}(x')||\mathcal{C}(x))) - 1. \tag{24}$$

It is known that $D_\alpha(.||.)$ is nondecreasing in α. Thus, using $\alpha \geq 2$, we obtain

$$\int_z \frac{(\mathcal{C}(z|x) - \mathcal{C}(z|x'))^2}{\mathcal{C}(z|x)}\, dz \leq e^\varepsilon - 1. \tag{25}$$

Plugging this back into (23), we have

$$d_{\text{hel}}(\nu_1,\nu_2)^2 \leq (e^\varepsilon - 1) \int_x \frac{(d\mu_1(x) - d\mu_2(x))^2}{d\mu_1(x) + d\mu_2(x)}. \tag{26}$$

To complete the proof, we just need to show that

$$\int_x \frac{(d\mu_1(x) - d\mu_2(x))^2}{d\mu_1(x) + d\mu_2(x)} \leq 2 d_{\text{hel}}(\mu_1,\mu_2)^2.$$

To do so, note that

$$\int_x \frac{(d\mu_1(x) - d\mu_2(x))^2}{d\mu_1(x) + d\mu_2(x)} \leq 2 \int_x \frac{(d\mu_1(x) - d\mu_2(x))^2}{(\sqrt{d\mu_1(x)} + \sqrt{d\mu_2(x)})^2}$$
$$= 2\int_x \left(\sqrt{d\mu_1(x)} - \sqrt{d\mu_2(x)}\right)^2 = 2 d_{\text{hel}}(\mu_1,\mu_2)^2.$$

This concludes the proof of lemma 2 and hence the proof of Theorem 1. ■

Proof of Theorem 2

Recall the iterim quantities

$$\begin{aligned} t_i(c_i) =& \mathbb{E}_{\mathbf{c}_{-i}}\left[t(c_i, \mathbf{c}_{-i})\right], \\ \varepsilon_i^{(l)}(c_i) =& \mathbb{E}_{\mathbf{c}_{-i}}\left[\varepsilon_i^{(l)}(c_i, \mathbf{c}_{-i})\right], \text{ and} \\ \varepsilon_i^{(c)}(c_i) =& \mathbb{E}_{\mathbf{c}_{-i}}\left[\varepsilon_i^{(c)}(c_i, \mathbf{c}_{-i})\right] \text{ for all } i \in \mathcal{N}, c_i. \end{aligned}$$

Using these quantities, the incentive compatibility constraint becomes

$$t_i(c_i) - c_i\varepsilon_i^{(l)}(c_i) - (1-c_i)\varepsilon_i^{(c)}(c_i) \geq t_i(c_i') - c_i\varepsilon_i^{(l)}(c_i') - (1-c_i)\varepsilon_i^{(c)}(c_i').$$

By equating the derivative of the right-hand side with respect to c_i' at c_i to zero, we obtain

$$t_i'(c_i) - c_i\left(\varepsilon_i'^{(l)}(c_i) - \varepsilon_i'^{(c)}(c_i)\right) - \varepsilon_i'^{(c)}(c_i) = 0.$$

This equation gives us the derivative of the payment in terms of the privacy loss levels. By taking the integral of this expression we obtain

$$\begin{aligned} t_i(c_i) &= t_i(0) + \int_0^{c_i}\left(\varepsilon_i'^{(c)}(z) + z\left(\varepsilon_i'^{(l)}(z) - \varepsilon_i'^{(c)}(z)\right)\right)dz \\ &= t_i(0) + \varepsilon_i^{(c)}(c_i) - \varepsilon_i^{(c)}(0) + c_i\left(\varepsilon_i^{(l)}(c_i) - \varepsilon_i^{(c)}(c_i)\right) - \int_0^{c_i}\left(\varepsilon_i^{(l)}(z) - \varepsilon_i^{(c)}(z)\right)dz. \quad (27) \end{aligned}$$

We next show that the payment in (27) together with a weakly decreasing $\varepsilon_i^{(l)}(z) - \varepsilon_i^{(c)}(z)$ guarantees that the incentive compatibility constraint. To see this, we consider two possibilities depending on whether c_i' is larger or smaller than c_i:

- For $c_i' \geq c_i$, by using the payment in (27), the incentive compatibility constraint becomes equivalent to

  $$\left(\varepsilon_i^{(l)}(c_i') - \varepsilon_i^{(c)}(c_i')\right)(c_i - c_i') \geq \int_{c_i'}^{c_i}\left(\varepsilon_i^{(l)}(z) - \varepsilon_i^{(c)}(z)\right)dz,$$

  which holds because $\varepsilon_i^{(l)}(z) - \varepsilon_i^{(c)}(z)$ is weakly decreasing in z.
- For $c_i' \leq c_i$, by using the payment in (27), the incentive compatibility constraint becomes equivalent to

  $$\left(\varepsilon_i^{(l)}(c_i') - \varepsilon_i^{(c)}(c_i')\right)(c_i - c_i') \leq \int_{c_i}^{c_i'}\left(\varepsilon_i^{(l)}(z) - \varepsilon_i^{(c)}(z)\right)dz,$$

  which, again, holds because $\varepsilon_i^{(l)}(z) - \varepsilon_i^{(c)}(z)$ is weakly decreasing in z. This completes one direction of the proof.

To see the other direction, notice that using the first order condition for the incentive compatibility constraints, imply (27). To see the monotonicity, notice that the incentive compatibility implies

$$\varepsilon_i^{(c)}(c_i) + c_i\left(\varepsilon_i^{(l)}(c_i) - \varepsilon_i^{(c)}(c_i)\right) - t_i(c_i) \leq \varepsilon_i^{(c)}(c_i') + c_i\left(\varepsilon_i^{(l)}(c_i') - \varepsilon_i^{(c)}(c_i')\right) - t_i(c_i').$$

and

$$\varepsilon_i^{(c)}(c_i') + c_i'\left(\varepsilon_i^{(l)}(c_i') - \varepsilon_i^{(c)}(c_i')\right) - t_i(c_i') \leq \varepsilon_i^{(c)}(c_i) + c_i'\left(\varepsilon_i^{(l)}(c_i) - \varepsilon_i^{(c)}(c_i)\right) - t_i(c_i).$$

The summation of these two inequalities yields

$$\left(\left(\varepsilon_i^{(l)}(c_i) - \varepsilon_i^{(c)}(c_i)\right) - \left(\varepsilon_i^{(l)}(c_i') - \varepsilon_i^{(c)}(c_i')\right)\right)(c_i - c_i') \leq 0,$$

that proves $\varepsilon_i^{(l)}(\cdot) - \varepsilon_i^{(c)}(\cdot)$ is weakly decreasing.

We next evaluate the individual rationality constraint. Using (27), we can rewrite this constraint as

$$t_i(0) \geq \varepsilon_i^{(c)}(0) + \int_0^{c_i} \left(\varepsilon_i^{(l)}(z) - \varepsilon_i^{(c)}(z)\right) dz \quad \text{for all } c_i. \tag{28}$$

Using $\varepsilon_i^{(l)}(z) \geq \varepsilon_i^{(c)}(z)$ for all z, this inequality means that it only needs to hold for $c_i = \infty$. Hence, we could cast $t_i(0)$ as

$$t_i(0) = \varepsilon_i^{(c)}(0) + \int_0^{\infty} \left(\varepsilon_i^{(l)}(z) - \varepsilon_i^{(l)}(z)\right) dz.$$

Plugging this back in (27) yields

$$t_i(c_i) = \varepsilon_i^{(c)}(c_i) + c_i \left(\varepsilon_i^{(l)}(c_i) - \varepsilon_i^{(c)}(c_i)\right) + \int_{c_i} \left(\varepsilon_i^{(l)}(z) - \varepsilon_i^{(c)}(z)\right) dz.$$

which is the optimal payment when $\varepsilon^{(l)}(\cdot)$ is decreasing.

With this optimal payment, the expected payment becomes

$$\begin{aligned}
&\mathbb{E}_{c_i}\left[t_i(c_i)\right] = \mathbb{E}_{c_i}[\varepsilon_i^{(c)}(c_i)] + \mathbb{E}_{c_i}\left[c_i \left(\varepsilon_i^{(l)}(c_i) - \varepsilon_i^{(c)}(c_i)\right) + \int_{z=c_i} \left(\varepsilon_i^{(l)}(z) - \varepsilon_i^{(c)}(z)\right) dz\right] \\
&= \mathbb{E}_{c_i}[\varepsilon_i^{(c)}(c_i)] \\
&+ \int_{\mathbf{z}_{-i}} \int_{z_i} \left(z_i \left(\varepsilon_i^{(l)}(z_i, \mathbf{z}_{-i}) - \varepsilon_i^{(c)}(z_i, \mathbf{z}_{-i})\right) + \int_{y_i = z_i} \left(\varepsilon_i^{(l)}(y, \mathbf{z}_{-i}) - \varepsilon_i^{(c)}(y, \mathbf{z}_{-i})\right) dy_i\right) f_i(z_i) dz_i f_{-i}(\mathbf{z}_{-i}) d\mathbf{z}_{-i} \\
&\overset{(a)}{=} \mathbb{E}_{c_i}[\varepsilon_i^{(c)}(c_i)] \\
&+ \int_{\mathbf{z}_{-i}} \int_{z_i} \left(z_i \left(\varepsilon_i^{(l)}(z_i, \mathbf{z}_{-i}) - \varepsilon_i^{(c)}(z_i, \mathbf{z}_{-i})\right) + \left(\varepsilon_i^{(l)}(z_i, \mathbf{z}_{-i}) - \varepsilon_i^{(c)}(z_i, \mathbf{z}_{-i})\right) \frac{F_i(z_i)}{f_i(z_i)}\right) f_i(z_i) dz_i f_{-i}(\mathbf{z}_{-i}) d\mathbf{z}_{-i} \\
&= \mathbb{E}_{c_i}[\varepsilon_i^{(c)}(c_i)] + \int_{\mathbf{z}} \left(z_i + \frac{F_i(z_i)}{f_i(z_i)}\right) \left(\varepsilon_i^{(l)}(\mathbf{z}) - \varepsilon_i^{(c)}(\mathbf{z})\right) f(\mathbf{z}) d\mathbf{z},
\end{aligned} \tag{29}$$

where (a) follows from changing the order of the integrals. Substituting equation (29) in the platform's objective function yields

$$\begin{aligned}
&\mathbb{E}_{\mathbf{c}}\left[\gamma \text{MSE}(\boldsymbol{\varepsilon}^{(l)}(\mathbf{c}), \boldsymbol{\varepsilon}^{(c)}(\mathbf{c}), \hat{\theta}) + \sum_{i=1}^n t_i(\mathbf{c})\right] \\
&= \mathbb{E}_{\mathbf{c}}\left[\gamma \text{MSE}(\boldsymbol{\varepsilon}^{(l)}(\mathbf{c}), \boldsymbol{\varepsilon}^{(c)}(\mathbf{c}), \hat{\theta}) + \sum_{i=1}^n \psi_i(c_i) \varepsilon_i^{(l)}(\mathbf{c}) + \sum_{i=1}^n \left(1 - \psi_i(c_i)\right) \varepsilon_i^{(c)}(\mathbf{c})\right].
\end{aligned}$$

Notice that the maximizer of the above objective is the optimal local and central privacy losses, provided that $\varepsilon_i^{(l)}(\cdot) - \varepsilon_i^{(c)}(\cdot)$ is decreasing. For a given privacy sensitivity vector $\mathbf{c}$, let us consider the point-wise minimization given by

$$\min_{\{\boldsymbol{\varepsilon}^{(l)}\}_{i=1}^n, \{\boldsymbol{\varepsilon}^{(c)}\}_{i=1}^n} \quad \gamma \text{MSE}(\boldsymbol{\varepsilon}^{(l)}, \boldsymbol{\varepsilon}^{(c)}, \hat{\theta}) + \sum_{i=1}^n \varepsilon_i^{(l)} \psi_i(c_i) + \sum_{i=1}^n \varepsilon_i^{(c)} \left(1 - \psi_i(c_i)\right). \tag{30}$$

This point-wise optimization clearly finds the optimal $\varepsilon^{(l)}(\cdot)$ and $\varepsilon^{(c)}(\cdot)$, but the issue is that the corresponding $\varepsilon_i^{(l)}(\cdot) - \varepsilon_i^{(c)}(\cdot)$ may not be decreasing. We next show that under Assumption 1 this is always the case.

Let $\{\boldsymbol{\varepsilon}^{(l)}\}_{i=1}^n$ and $\{\boldsymbol{\varepsilon}^{(c)}\}_{i=1}^n$ be the solution of optimization problem (30) for $c_1, \ldots, c_n$. Now, suppose we increase one of the c_i's, which, without loss of generality, we assume is the first one. Let $c_1' > c_1$ and $c_i' = c_i$ for $i = 2, \ldots, n$ and suppose $\{\boldsymbol{\varepsilon}'^{(l)}\}_{i=1}^n, \{\boldsymbol{\varepsilon}'^{(c)}\}_{i=1}^n$ is the corresponding optimal solution of optimization problem (30). The optimality condition implies that

$$\begin{aligned}
&\gamma \text{MSE}(\boldsymbol{\varepsilon}^{(l)}, \boldsymbol{\varepsilon}^{(c)}, \hat{\theta}) + \sum_{i=1}^n \varepsilon_i^{(l)} \psi_i(c_i) + \sum_{i=1}^n \varepsilon_i^{(c)} \left(1 - \psi_i(c_i)\right) \\
&\leq \gamma \text{MSE}(\boldsymbol{\varepsilon}'^{(l)}, \boldsymbol{\varepsilon}'^{(c)}, \hat{\theta}) + \sum_{i=1}^n \varepsilon_i'^{(l)} \psi_i(c_i) + \sum_{i=1}^n \varepsilon_i'^{(c)} \left(1 - \psi_i(c_i)\right)
\end{aligned}$$

and

$$\gamma \text{MSE}(\boldsymbol{\varepsilon}'^{(l)}, \boldsymbol{\varepsilon}'^{(c)}, \hat{\theta}) + \sum_{i=1}^{n} \varepsilon'^{(l)}_i \psi_i(c'_i) + \sum_{i=1}^{n} \varepsilon'^{(c)}_i \left(1 - \psi_i(c_i)\right)$$
$$\leq \gamma \text{MSE}(\boldsymbol{\varepsilon}^{(l)}, \boldsymbol{\varepsilon}^{(c)}, \hat{\theta}) + \sum_{i=1}^{n} \varepsilon^{(l)}_i \psi_i(c'_i) + \sum_{i=1}^{n} \varepsilon^{(c)}_i \left(1 - \psi_i(c_i)\right)$$

The summation of both sides of these inequalities, together with $c_i = c'_i$ for $i = 2, \ldots, n$, results in

$$\left(\left(\varepsilon^{(l)}_1 - \varepsilon^{(c)}_1\right) - \left(\varepsilon'^{(l)}_1 - \varepsilon'^{(c)}_1\right)\right)\left(\psi_1(c_1) - \psi_1(c'_1)\right) \leq 0.$$

Assumption 1 and the above inequality establishes that the solution of problem (30) is weakly decreasing in the privacy sensitivity. ■

Proof of Corollary 1

The proof follows by invoking Theorem 2 and noting that with

$$\hat{\theta} = \sum_{i=1}^{n} w_i \left(x_i + \mathcal{N}\left(0, \frac{\alpha}{2\varepsilon^{(l)}_i}\right)\right)$$

we have

$$\varepsilon^{(c)}_i = \frac{w_i^2}{\sum_{j=1}^{n} \frac{w_j^2}{\varepsilon^{(l)2}_j}}$$

and

$$\text{MSE}(\boldsymbol{\varepsilon}^{(l)}, \boldsymbol{\varepsilon}^{(c)}, \hat{\theta}) = \text{VAR} \sum_{i=1}^{n} w_i^2 + \sum_{i=1}^{n} w_i^2 \frac{\alpha}{2\varepsilon^{(l)}_i}.$$

This completes the proof. ■

Proof of Proposition 1

With a Gaussian mechanism that adopts both local and central noises, using a similar argument to that of Theorem 2 and Corollary 1, the optimal central privacy loss levels are

$$\varepsilon^{(c)}_i = \frac{w^{*2}_i}{\sum_{j=1}^{n} \frac{w^{*2}_j}{y^*_j} + \frac{1}{\varepsilon}},$$

where $(w^*_1, \ldots, w^*_n)$, $(y^*_1, \ldots, y^*_n)$, and ε are the optimal solution of

$$\min_{\mathbf{w}, \mathbf{y}, \varepsilon} \ \text{VAR}\gamma \sum_{i=1}^{n} w_i^2 + \frac{\gamma\alpha}{2} \sum_{i=1}^{n} \frac{w_i^2}{y_i} + \frac{\alpha\gamma}{2\varepsilon} + \sum_{i=1}^{n} \left(1 - \psi_i(c_i)\right) \frac{w_i^2}{\sum_{j=1}^{n} \frac{w_j^2}{y_j} + \frac{1}{\varepsilon}} + \sum_{i=1}^{n} \psi_i(c_i) y_i$$
$$\text{s.t. } w_i, y_i \geq 0, \text{ for all } i \in \mathcal{N}$$
$$\sum_{i=1}^{n} w_i = 1.$$

For any solution of the above optimization problem we define the following alternative solution:

$$y'_i = \frac{w_i^2}{\frac{w_i^2}{y_i} + \frac{w_i}{\varepsilon}}, \varepsilon' = \infty, \text{ and } w'_i = w_i \quad \text{for all } i \in \mathcal{N}.$$

We have

$$\begin{aligned}
&\text{VAR}\gamma \sum_{i=1}^{n} {w'_i}^2 + \frac{\gamma\alpha}{2} \sum_{i=1}^{n} \frac{{w'_i}^2}{y'_i} + \sum_{i=1}^{n} (1-\psi_i(c_i)) \frac{{w'_i}^2}{\sum_{j=1}^{n} \frac{{w'_j}^2}{y'_j}} + \sum_{i=1}^{n} \psi_i(c_i) y'_i \\
&\overset{(a)}{=} \text{VAR}\gamma \sum_{i=1}^{n} {w_i}^2 + \frac{\gamma\alpha}{2} \sum_{i=1}^{n} \frac{{w_i}^2}{y_i} + \frac{\alpha\gamma}{2\varepsilon} + \sum_{i=1}^{n} (1-\psi_i(c_i)) \frac{{w_i}^2}{\sum_{j=1}^{n} \frac{{w_j}^2}{y_j} + \frac{1}{\varepsilon}} + \sum_{i=1}^{n} \psi_i(c_i) y'_i \\
&\overset{(b)}{\leq} \text{VAR}\gamma \sum_{i=1}^{n} {w_i}^2 + \frac{\gamma\alpha}{2} \sum_{i=1}^{n} \frac{{w_i}^2}{y_i} + \frac{\alpha\gamma}{2\varepsilon} + \sum_{i=1}^{n} (1-\psi_i(c_i)) \frac{{w_i}^2}{\sum_{j=1}^{n} \frac{{w_j}^2}{y_j} + \frac{1}{\varepsilon}} + \sum_{i=1}^{n} \psi_i(c_i) y_i,
\end{aligned}$$

where (a) follows the construction of the new solution and (b) follows from

$$y'_i = \frac{w_i^2}{\frac{w_i^2}{y_i} + \frac{w_i}{\varepsilon}} \leq y_i \quad \text{for all } i \in \mathcal{N}.$$

This completes the proof.

Proof of Theorem 3

Consider the optimization problem

$$\begin{aligned}
\min_{\mathbf{w},\mathbf{y}} \ & \text{VAR}\gamma \sum_{i=1}^{n} w_i^2 + \frac{\gamma\alpha}{2} \sum_{i=1}^{n} \frac{w_i^2}{y_i} + \sum_{i=1}^{n} \frac{w_i^2}{\sum_{j=1}^{n} \frac{w_j^2}{y_j}} (1-\psi_i(c_i)) + \sum_{i=1}^{n} \psi_i(c_i) y_i \qquad (31) \\
\text{s.t. } & w_i, y_i \geq 0, \text{ for all } i \in \mathcal{N} \\
& \sum_{i=1}^{n} w_i = 1.
\end{aligned}$$

We can rewrite this optimization problem as

$$\begin{aligned}
\min_{\mathbf{w},\mathbf{y},S} \ & \text{VAR}\gamma \sum_{i=1}^{n} w_i^2 + \frac{\gamma\alpha}{2} S + \frac{1}{S} \sum_{i=1}^{n} w_i^2 (1-\psi_i(c_i)) + \sum_{i=1}^{n} \psi_i(c_i) y_i \\
\text{s.t. } & w_i, y_i \geq 0, \text{ for all } i \in \mathcal{N} \\
& \sum_{i=1}^{n} w_i = 1, \\
& \sum_{i=1}^{n} \frac{w_i^2}{y_i} = S.
\end{aligned}$$

Let us fix S. The Lagrangian of this optimization problem becomes

$$\begin{aligned}
&\sum_{i=1}^{n} w_i^2 \left(\text{VAR}\gamma + \frac{1-\psi_i(c_i)}{S} \right) + \sum_{i=1}^{n} \psi_i(c_i) y_i + p \left(\sum_{i=1}^{n} \frac{w_i^2}{y_i} - S \right) - q \left(\sum_{i=1}^{n} w_i - 1 \right) \\
&- \sum_{i=1}^{n} u_i w_i - \sum_{i=1}^{n} v_i y_i,
\end{aligned}$$

where $u_i, v_i \geq 0$ and $u_i w_i = v_i y_i = 0$, for all i.

Equating the derivative with respect to w_i to zero, yields

$$2w_i \left(\text{VAR}\gamma + \frac{1-\psi_i(c_i)}{S} \right) + \frac{2pw_i}{y_i} - q = u_i. \qquad (32)$$

Hence, if $w_i^* > 0$, then $u_i = 0$ which implies

$$2 \left(\text{VAR}\gamma + \frac{1-\psi_i(c_i)}{S} + \frac{p}{y_i} \right) w_i = q. \qquad (33)$$

On the other hand, if $w_i^* = 0$, then $u_i = -q \geq 0$.

Equating the derivative with respect to y_i to zero implies

$$\psi_i(c_i) - \frac{pw_i^2}{y_i^2} - v_i.$$

Hence, if $y_i^* = 0$, then $w_i^* = 0$ and $v_i = \psi_i(c_i)$. On ther other hand, if $y_i > 0$, then $v_i = 0$ and we have

$$w_i = \sqrt{\frac{\psi_i(c_i)}{p}} y_i. \tag{34}$$

Now, we claim there is no i for which $w_i^* = 0$ (and hence there is no i for which $y_i^* = 0$). Assume this is not the case, and hence there exists some i_0 for which $w_{i_0}^* = 0$. Therefore, as we established earlier, we have $q = -u_{i_0} \leq 0$. On the other hand, note that there exists j for which $w_j^* > 0$. Hence, for that $y_j^* > 0$ as well. Therefore, using (33), along with the fact that $q \leq 0$, we should have

$$\frac{1 - \psi_j(c_j)}{S} + \frac{p}{y_j^*} \leq 0,$$

which implies

$$\frac{\psi_j(c_j) - 1}{S} \geq \frac{p}{y_j^*}.$$

Hence, using $S \geq (w_j^*)^2/y_j^*$,we have

$$\psi_j(c_j) - 1 \geq S\frac{p}{y_j^*} \geq \frac{(w_j^*)^2 p}{y_j^*}.$$

However, since $y_j^* > 0$, by (34), the right hand side is equal to $\psi_j(c_j)$, which implies $\psi_j(c_j) - 1 \geq \psi_j(c_j)$ which is a contradiction! As a result, (33) and (34) hold for all i.

By invoking (34) in (33), we obtain

$$w_i = \frac{1}{\text{VAR}\gamma + \frac{1-\psi_i(c_i)}{S}} \left(\frac{q}{2} - \sqrt{\psi_i(c_i)p} \right).$$

To simplify the analysis, we define the interim variable

$$\nu_i = \frac{1}{\gamma \text{VAR} + (1 - \psi_i(c_i))/S}. \tag{35}$$

Taking summation of the above equation for $i = 1, \ldots, n$ and using $\sum_{i=1}^n w_i = 1$, we obtain

$$\frac{q}{2} = \frac{1 + \sum_j \nu_j \sqrt{\psi_j(c_j)p}}{\sum_j \nu_j},$$

which together with (32) results in

$$w_i = \frac{\nu_i}{\sum_j \nu_j} + \frac{\nu_i}{\sum_j \nu_j} \left(\sum_{j=1}^n \nu_j (\sqrt{\psi_j(c_j)p} - \sqrt{\psi_i(c_i)p}) \right). \tag{36}$$

Therefore, by using (34) and (36), once we have S and p, we can find y_i and w_i for all $i \in \mathcal{N}$.

Next, we derive a relation between S and p. Note that (36) implies that w_i can be cast as

$$\zeta_i(S) + \xi_i(S)\sqrt{p}$$

with

$$\zeta_i(S) = \frac{\nu_i}{\sum_j \nu_j} \text{ and } \xi_i(S) = \frac{\nu_i}{\sum_j \nu_j} \left(\sum_{j=1}^n \nu_j (\sqrt{\psi_j(c_j)} - \sqrt{\psi_i(c_i)}) \right).$$

Using (34), we have

$$S = \sum_{i=1}^{n} \frac{w_i^2}{y_i} = \sum_{i=1}^{n} w_i \frac{\sqrt{\psi_i(c_i)}}{\sqrt{p}} = \sum_{i=1}^{n} \frac{\zeta_i(S)\sqrt{\psi_i(c_i)}}{\sqrt{p}} + \sum_{i=1}^{n} \sqrt{\psi_i(c_i)}\xi_i(S).$$

This implies

$$p = \left(\frac{\sum_{i=1}^{n} \zeta_i(S)\sqrt{\psi_i(c_i)}}{S - \sum_{i=1}^{n} \sqrt{\psi_i(c_i)}\xi_i(S)} \right)^2. \tag{37}$$

We next show that we can search over a grid to find the approximately optimal S. In this regard, we derive a lower and upper bound on the optimal S.

To do so, first note that the objective function (31) is given by

$$\begin{aligned} \text{OBJ} &= \text{VAR}\gamma \sum_{i=1}^{n} w_i^2 + \frac{\gamma\alpha}{2} S + \frac{1}{S} \sum_{i=1}^{n} w_i^2 \left(1 - \psi_i(c_i)\right) + \sum_{i=1}^{n} \psi_i(c_i) y_i \\ &\geq \frac{\gamma\alpha}{2} S + \frac{1}{S} \sum_{i=1}^{n} w_i^2 + \sum_{i=1}^{n} \psi_i(c_i)(y_i - \frac{w_i^2}{S}). \end{aligned}$$

It is straightforward to see $y_i \geq \frac{w_i^2}{S}$ for all i, and thus, we have

$$\text{OBJ} \geq \frac{\gamma\alpha}{2} S + \frac{1}{S} \sum_{i=1}^{n} w_i^2. \tag{38}$$

Using (38) along with the fact that Cauchy–Schwarz inequality implies $\sum_{i=1}^{n} w_i^2 \geq 1/n$, we have

$$\text{Optimal objective (OPT)} \geq \frac{\gamma\alpha}{2} S^* + \frac{1}{nS^*}.$$

As a result, we have

$$\frac{\text{OPT}}{\gamma\alpha/2} \geq S^* \geq \frac{1}{\text{OPT}n}. \tag{39}$$

Letting $y_i = 1$, $w_i = \frac{1}{n}$, and $\varepsilon = 1$ in the objective of Problem (31) gives us an upper bound on the optimal objective OPT. Let us denote this upper bound by M. We have

$$\frac{M}{\gamma\alpha/2} \geq S^* \geq \frac{1}{Mn}. \tag{40}$$

Therefore, we obtain an approximate optimal solution by grid search. This provides an $O(\delta)$ optimal solution for the platform's problem because the objective of Problem (31) is Lipschitz continuous. ■

A.3 Revelation principle

Consider the strategy of user i is a function of its relative privacy sensitivity shown by $\beta_i(c_i)$. For a given estimator $\hat{\theta}$ and mechanism $(\boldsymbol{\varepsilon}^{(l)}, \boldsymbol{\varepsilon}^{(c)}, \mathbf{t})$, the action profile $\{\beta_i(\cdot)\}_{i=1}^n$ is an equilibrium if

$$\begin{aligned} &\mathbb{E}_{\mathbf{c}_{-i}} \left[t_i(\beta_{-\mathbf{i}}(\mathbf{c}_{-\mathbf{i}}), \beta_{\mathbf{i}}(\mathbf{c}_{\mathbf{i}})) - \mathbf{c_i}\varepsilon_{\mathbf{i}}^{(\mathbf{l})}(\beta_{-\mathbf{i}}(\mathbf{c}_{-\mathbf{i}}), \beta_{\mathbf{i}}(\mathbf{c}_{\mathbf{i}})) - \varepsilon_{\mathbf{i}}^{(\mathbf{c})}(\beta_{-\mathbf{i}}(\mathbf{c}_{-\mathbf{i}}), \beta_{\mathbf{i}}(\mathbf{c}_{\mathbf{i}})) \right] \\ &\geq \mathbb{E}_{\mathbf{c}_{-i}} \left[t_i(\beta_{-\mathbf{i}}(\mathbf{c}_{-\mathbf{i}}), \beta'_{\mathbf{i}}(\mathbf{c}_{\mathbf{i}})) - \mathbf{c_i}\varepsilon_{\mathbf{i}}^{(\mathbf{l})}(\beta_{-\mathbf{i}}(\mathbf{c}_{-\mathbf{i}}), \beta'_{\mathbf{i}}(\mathbf{c}_{\mathbf{i}})) - \varepsilon_{\mathbf{i}}^{(\mathbf{c})}(\beta_{-\mathbf{i}}(\mathbf{c}_{-\mathbf{i}}), \beta'_{\mathbf{i}}(\mathbf{c}_{\mathbf{i}})) \right] \end{aligned}$$

for all $i \in \mathcal{N}, c_i, \beta'_i(\cdot)$. By letting $(\tilde{\boldsymbol{\varepsilon}}^{(l)}, \tilde{\boldsymbol{\varepsilon}}^{(c)}, \tilde{\mathbf{t}})$ be such that $\tilde{\varepsilon}_i^{(l)}(c_1, \ldots, c_n) = \varepsilon_i^{(l)}(\beta_1(c_1), \ldots, \beta_n(c_n))$, $\tilde{\varepsilon}_i^{(c)}(c_1, \ldots, c_n) = \varepsilon_i^{(c)}(\beta_1(c_1), \ldots, \beta_n(c_n))$, and $\tilde{t}_i(c_1, \ldots, c_n) = t_i(\beta_1(c_1), \ldots, \beta_n(c_n))$, the users will report truthfully and that the platform's objective is the same as the original mechanism. This establishes the revelation principle. ■