The Value of Privacy: Strategic Data Subjects, Incentive Mechanisms and Fundamental Limits

Weina Wang
School of Electrical, Computer
and Energy Engineering
Arizona State University
Tempe, AZ 85287
weina.wang@asu.edu

Lei Ying
School of Electrical, Computer
and Energy Engineering
Arizona State University
Tempe, AZ 85287
lei.ying.2@asu.edu

Junshan Zhang
School of Electrical, Computer
and Energy Engineering
Arizona State University
Tempe, AZ 85287
junshan.zhang@asu.edu

ABSTRACT

We study the value of data privacy in a game-theoretic model of trading private data, where a data collector purchases private data from strategic data subjects (individuals) through an incentive mechanism. The private data of each individual represents her knowledge about an underlying state, which is the information that the data collector desires to learn. Different from most of the existing work on privacy-aware surveys, our model does not assume the data collector to be trustworthy. Then, an individual takes full control of its own data privacy and reports only a privacy-preserving version of her data.

In this paper, the value of ϵ units of privacy is measured by the minimum payment of all nonnegative payment mechanisms, under which an individual's best response at a Nash equilibrium is to report the data with a privacy level of ϵ . The higher ϵ is, the less private the reported data is. We derive lower and upper bounds on the value of privacy which are asymptotically tight as the number of data subjects becomes large. Specifically, the lower bound assures that it is impossible to use less amount of payment to buy ϵ units of privacy, and the upper bound is given by an achievable payment mechanism that we designed. Based on these fundamental limits, we further derive lower and upper bounds on the minimum total payment for the data collector to achieve a given learning accuracy target, and show that the total payment of the designed mechanism is at most one individual's payment away from the minimum.

1. INTRODUCTION

From the monetary coupons offered for revealing opinions of a product to the large-scale trade of personal information by data brokers such as Acxiom [20], the commoditization of private data has been trending up when big data analytics is playing a more and more critical role in advertising, scientific research, etc. However, in the wake of a number of recent scandals, such as the Netflix data breach and the Vet-

erans Affairs data theft, data privacy is emerging as one of the most serious concerns of big data analytics. This raises a fundamental question "whether big-data and privacy can go hand-by-hand or giving up our privacy is inevitable in the big-data era." One common practice of collecting private data is called informed consent. With information on "who is collecting the data, what data is collected, and how the data will be used," data subjects decide upon whether to report data or not. The data collector is supposed to use the data only in the manner disclosed to data subjects. This practice, however, has two fundamental issues: (i) data subjects have no control of data privacy after transferring private data to the data collector; and (ii) the data collector has to take full responsibility of protecting users' private data, which not only costs significant investment on infrastructure and maintenance, but also may lead to reputation damage if data breach occurs. In some applications, such as collecting certain browsing history records to enhance the phishing and malware protection of web browsers [10, 11], the data collectors prefer to avoid holding individuals' raw data for subpoena concerns.

Taking a forward-looking view, we envisage a market model for private data analytics such that private data is treated as a commodity and traded in the market. In particular, the data collector will use an incentive mechanism to pay (or reward) individuals for reporting informative data, and individuals control their own data privacy by reporting noisy data with the appropriate level of privacy protection (or level of noise added) being strategically chosen to maximize their payoffs. A distinctive merit of this privacy protection approach is that data subjects take full control of their own privacy and the data collector gets informative data but does not need to bear the responsibility of protecting data privacy. This differentiates our approach from the existing work [15, 12, 22, 26, 13, 24, 14], where the data collector is assumed to be a trustworthy entity who is willing to and has the capability to protect users' privacy.

One significant challenge of the proposed paradigm is that the data collector has no direct control (perhaps no information either) over the quality of reported data. To tackle this challenge, we cast the problem into a game-theoretic setting, which allows us to quantify two fundamental tradeoffs: the tradeoff between cost and accuracy from the data collector's perspective, and the tradeoff between reward and privacy from an individual's perspective (the value of privacy for a data subject). In return, with the reward (incentive) as the bridge, it establishes the tradeoff of data privacy concerned

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGMETRICS '16, June 14–18, 2016, Antibes Juan-Les-Pins, France.

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4266-7/16/06...\$15.00

DOI: <http://dx.doi.org/10.1145/2896377.2901461>

by an individual versus data quality concerned by the data collector.

Specifically, we consider a game-theoretic model of collecting private data in hypothesis testing, where the data collector is interested in learning information from a population of N individuals. An illustration of our model is shown in Figure 1. The information is represented by a binary random variable W , which is called the *state*. Each individual i possesses a binary *signal* S_i , which is her private data, representing her knowledge about the state W . Conditional on the state W , the signals are independently generated such that the probability for each signal S_i to be the same as W is θ , where $0.5 < \theta < 1$. To protect her privacy, an individual reports only a privacy-preserving version of her signal, denoted by X_i , or chooses to not participate after considering both the payment from the data collector and the loss of privacy. The data collector needs to decide the amount of payment and the payment mechanism to get informative reports, i.e., not completely random data. Intuitively, the higher the payment is, the more informative the reported data should be. We will answer the following fundamental questions in this paper: *What is the minimum payment needed from the data collector to obtain reported data with a privacy level ϵ ? Which payment mechanism can be used to collect private data with minimum cost?* This setting without accounting for data privacy has garnered much attention in the literature (see, e.g., [23, 1, 21]), including the application of estimating the underlying value of a new technology by eliciting opinions from individuals.

Intuitively, the data collector can purchase more informative data (so higher privacy) by offering higher payment. However, the strategic behavior of the privacy-aware individuals makes this more complicated. Due to privacy concerns, an individual's action/strategy is the conditional distributions of the reported data given the realizations of the signal. But the actions of the individuals are not observable to the data collector. Instead, what the data collector receives is the reported data, generated randomly according to the individuals' strategies, so the payments can only be designed based on the reported data. This differs our problem from the conventional mechanism design.

Furthermore, the privacy-aware individuals weigh the privacy loss against the payment to choose the best quantity of privacy to trade. To make an individual willing to trade ϵ level of privacy, the data collector needs to make sure doing this benefits the individual most. We reiterate that the data collector has access only to the reported data instead of the individuals' actions. Note that only compensating the privacy cost incurred is not sufficient. The payment mechanism needs to ensure that ϵ is the best privacy level such that when an individual uses a less-private strategy, the decrease in her payment is faster than the decrease in her privacy cost, and similarly, when an individual uses a more-private strategy, the increase in her payment is slower than the increase in her privacy cost. In other words, with a game-theoretic approach, we consider an individual's best response in a Nash equilibrium, and the value of data privacy is measured by the minimum payment that makes this equilibrium strategy have a privacy level of ϵ , which represents the monetary value of data privacy in a market for private data.

Summary of Main Results

It is assumed that individuals use the celebrated notion of differential privacy [8, 7] to evaluate their data privacy. When an individual i uses an ϵ -differentially private randomization strategy to generate X_i , the privacy loss incurred is ϵ , and the individual's cost of privacy loss is a function of ϵ , whose form is assumed to be publicly known. The value of ϵ units of privacy, denoted by $V(\epsilon)$, is measured by the minimum payment of all nonnegative payment mechanisms under which an individual's best response in a Nash equilibrium is to report the data with privacy level ϵ , where nonnegativity ensures that individuals would not be *charged* for reporting data. We are interested in the range that $\epsilon > 0$, simply because when $\epsilon = 0$, the reported data is independent of the private data and thus would be of no use for data analysis. Our contributions are summarized as follows:

1. We establish a lower bound on $V(\epsilon)$. First we characterize the strategies of individuals at a Nash equilibrium to prove that from a payment perspective, it suffices to focus on nonnegative payment mechanisms at which the best response of an individual in a Nash equilibrium is a symmetric randomized response with a privacy level of ϵ . This strategy generates the reported data by flipping the signal with probability $\frac{1}{e^\epsilon + 1}$: for convenience, this is called the ϵ -strategy. Next we prove that the expected payments resulting from any Nash equilibrium of any payment mechanism can be "replicated" by a genie-aided payment mechanism, where the payments are determined with the aid of a genie who knows the underlying state W . This makes the analysis of the Nash equilibria more tractable by decoupling the individuals in the payments. The lower bound is then given by necessary conditions for ϵ to be the best privacy level in the genie-aided mechanism. We remark that although the genie-aided mechanism that achieves the lower bound is not implementable, it can be well-approximated, when the number of individuals is large, by the feasible payment mechanism that we design to establish the upper bound.
2. We observe that the equilibrium strategies exhibit some interesting characteristics: the strategy of an individual in a Nash equilibrium is either a symmetric randomized response, which treats the realizations of the private signal symmetrically, or a non-informative strategy, where the reported data is independent of the signal. This characterization holds regardless of the prior distribution of the state, and it also holds for more general probability models of the signals. This characterization advances our understanding of the behavior of privacy-aware individuals. It is worth pointing out that finding an equilibrium strategy of a privacy-aware individual under some payment mechanism involves non-convex optimization.
3. We prove an upper bound on $V(\epsilon)$ by designing a payment mechanism $\mathbf{R}^{(N, \epsilon)}$, in which the strategy profile consisting of ϵ -strategies constitutes a Nash equilibrium. The expected payment to each individual at this equilibrium gives an upper bound on $V(\epsilon)$. This upper bound converges to the lower bound exponentially fast as the number of individuals N becomes large, which indicates that the lower and upper bounds are asymptotically tight.
4. The above fundamental bounds on the value of privacy can be further used to study the *payment-accuracy prob-*

lem, where the data collector aims to minimize the total payment while achieving an accuracy target in learning the state W . Given an accuracy target τ , which can be regarded as the maximum allowable error, let $F(\tau)$ denote the minimum total payment for achieving τ . We obtain lower and upper bounds on $F(\tau)$ based on the lower and upper bounds on the value of privacy. The upper bound is given by the designed mechanism $\mathbf{R}^{(N,\epsilon)}$ with properly chosen parameters, which shows that the total payment of the designed mechanism is at most one individual's payment away from the minimum.

2. RELATED WORK

Most existing work on privacy-aware surveys [15, 12, 22, 26, 13, 24, 14] assumes that there is a trusted data curator or data collector. The private data is either already kept by the data collector, or is elicited using mechanisms that are designed with the aim of truthfulness. What the data collector purchases is the “right” of using individuals’ data in an announced way. Our work differs from the existing work by considering a data collector who is not trusted by individuals. In this scenario, the data collector directly purchases the private data, in which privacy is embedded.

In the seminal work by Ghosh and Roth [15], individuals’ data is already known to the data collector, and individuals bid their costs of privacy loss caused by data usage, where each individual’s privacy cost is modeled as a linear function of ϵ if her data is used in an ϵ -differentially private manner. The goal of the mechanism design is to elicit truthful bids of individuals’ cost functions, i.e., the coefficients. Subsequent work [12, 22, 26, 24] explores various models for individuals’ valuation of privacy, especially the correlation between the coefficients and the private bits.

This line of work has been extended to the scenario that the data is not available yet and needs to be reported by the individuals to the data collector, but the data collector is still trusted [13, 33, 3, 14]. Notably, Ghosh, Ligett and Roth [14] study the model in which the collected data is non-verifiable. The goal of the mechanism design there is to incentivize truthful data reporting (without adding any noise) from individuals. For more work on the interplay between differential privacy and mechanism design, Pai and Roth [25] give a comprehensive survey.

The local model of differential privacy, which is a generalization of randomized response [32] and is formalized in [19], has been studied in the literature [8, 7, 16, 6, 9, 18, 29, 30, 2, 27]. The hypothesis testing formulation in our paper is similar to a setting in [18], where the authors find an optimal mechanism that maximizes the statistical discrimination of the hypotheses subject to local differential privacy constraint. In practice, Google’s Chrome web browser has implemented the RAPPOR mechanism [10, 11] to collect users’ data, which guarantees that only limited privacy will be leaked by using randomized response in a novel manner. However, users may still not be willing to report data in the desired way due to the lack of an incentive mechanism.

3. SYSTEM MODEL

We consider a single-bit learning problem with privacy-aware individuals as shown in Figure 1. Recall that the data collector is interested in learning the state W , which is a binary random variable. For example, the state W can

Figure 1: Information structure of the model: The data collector is interested in the state W , which is a binary random variable. Each individual i possesses her private data, which is a binary signal S_i . Conditioned on W , $S_1, S_2, \dots, S_N$ are i.i.d. Individual i ’s reported data is X_i , which is generated based on S_i using a randomized strategy.

describe the underlying value of some new technology. Let P_W denote the prior PMF of W . We assume that $P_W(1) > 0$ and $P_W(0) > 0$.

Individuals and Strategies. Consider a population of N individuals and denote the set of individuals by $\mathcal{N} = \{1, 2, \dots, N\}$. Denote all individuals other than some given individual i by “ $-i$.” Each individual i possesses a binary signal S_i , which is her private data, reflecting her knowledge about the state W . For example, S_i can represent individual i ’s opinion towards the new technology. Let $\mathbf{S} = (S_1, S_2, \dots, S_N)$. Conditional on the state W , the signals $S_1, S_2, \dots, S_N$ are i.i.d. with the following conditional distributions:

$$\begin{aligned} \mathbb{P}(S_i = 1 \mid W = 1) &= \theta, & \mathbb{P}(S_i = 0 \mid W = 1) &= 1 - \theta, \\ \mathbb{P}(S_i = 0 \mid W = 0) &= \theta, & \mathbb{P}(S_i = 1 \mid W = 0) &= 1 - \theta, \end{aligned}$$

where the parameter θ with $0.5 < \theta < 1$ is called *the quality of signals*.

Let X_i denote the data reported by individual i and let $\mathbf{X} = (X_1, X_2, \dots, X_N)$. The acceptable values for reported data are 0, 1, and “nonparticipation.” So X_i takes values in the set $\mathcal{X} = \{0, 1, \perp\}$, where $\perp$ indicates that individual i declines to participate. A strategy of individual i for data reporting is a mapping $\sigma_i: \{0, 1\} \rightarrow \mathcal{D}(\mathcal{X})$, where $\mathcal{D}(\mathcal{X})$ is the set of probability distributions on $\mathcal{X}$. Let $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \dots, \sigma_N)$. The strategy σ_i prescribes a distribution to X_i for each possible value of S_i , which defines the conditional distribution of X_i given S_i . Since we will discuss different strategies of individual i , we let $\mathbb{P}_{\sigma_i}(X_i = x_i \mid S_i = s_i)$ with $x_i \in \mathcal{X}$ and $s_i \in \{0, 1\}$ denote the conditional probabilities defined by strategy σ_i . If a strategy σ_i satisfies that $\mathbb{P}_{\sigma_i}(X_i = 1 \mid S_i = 1) = \mathbb{P}_{\sigma_i}(X_i = 0 \mid S_i = 0)$ and $\mathbb{P}_{\sigma_i}(X_i = \perp \mid S_i = 1) = \mathbb{P}_{\sigma_i}(X_i = \perp \mid S_i = 0) = 0$, we say σ_i is a *symmetric randomized response*. If a strategy σ_i makes X_i and S_i independent, we say σ_i is *non-informative*; otherwise we say σ_i is *informative*.

Mechanism. The data collector uses a payment mechanism $\mathbf{R}: \mathcal{X}^N \rightarrow \mathbb{R}^N$ to determine the amount of payment to each individual, where $R_i(\mathbf{x})$ is the payment to individual i when the reported data is $\mathbf{X} = \mathbf{x}$. We are interested in payment mechanisms in which the payment to each individual is nonnegative, i.e., $R_i(\mathbf{x}) \geq 0$ for any individual i and any $\mathbf{x} \in \mathcal{X}^N$, which we call *nonnegative mechanisms*.

This constraint is motivated by the fact that in many practical applications such as surveys, the data collector has no means to charge users and can only use payments to incentivize user participation.

Privacy Cost. We quantify the privacy loss incurred when a strategy is in use by the level of (local) differential privacy [8, 7, 19, 9] of the strategy, defined as follows.

DEFINITION 1. *The level of (local) differential privacy, or simply the privacy level, of a strategy σ_i , denoted by $\zeta(\sigma_i)$, is defined to be*

$$\zeta(\sigma_i) = \max \left\{ \ln \left(\frac{\mathbb{P}_{\sigma_i}(X_i \in \mathcal{E} \mid S_i = s_i)}{\mathbb{P}_{\sigma_i}(X_i \in \mathcal{E} \mid S_i = 1 - s_i)} \right) : \mathcal{E} \subseteq \{0, 1, \perp\}, s_i \in \{0, 1\} \right\},$$

where we follow the convention that $0/0 = 1$, and the strategy σ_i is said to be $\zeta(\sigma_i)$ -differentially private.

The level of differential privacy quantifies the indistinguishability between the conditional distributions of the reported data given different values of the signal, therefore measuring how disclosive the strategy is. Note that the amount of privacy leakage quantified by differential privacy is “in addition” to what the adversaries already know. We refer the reader to [8] for more semantic implications of differential privacy.

The privacy loss causes a cost to an individual. We assume that when using strategies with the same privacy level, individuals experience the same cost of privacy. Thus, we model each individual’s cost of privacy by a function g of the privacy level. We call g the *cost function* and the cost the *privacy cost*. Our results can be extended to the case where the cost functions are heterogeneous (see the discussion in Section 4.3). We assume that the form of g is publicly known (Ghosh and Roth [15] and subsequent work study the scenario that cost functions are private and design truthful mechanisms to elicit them).

We say the cost function g is *proper* if it satisfies the following three conditions:

$$\begin{aligned} g(\xi) &\geq 0, \quad \forall \xi \geq 0, & (1) \\ g(0) &= 0, & (2) \\ g &\text{ is non-decreasing,} & (3) \end{aligned}$$

where (1) follows from the fact that a privacy cost is non-negative, (2) indicates that the privacy cost is 0 when the reported data is independent of the private data, and (3) means that the privacy cost will not decrease when the privacy loss becomes larger. In this paper, we will focus on a proper cost function that is convex, continuously differentiable, and $g(\xi) = 0$ only for $\xi = 0$. With a little abuse of notation, we also use $g(\sigma_i)$ to denote $g(\zeta(\sigma_i))$, which is the privacy cost to individual i when the strategy σ_i is used.

Game Formulation and Nash Equilibrium. In this market model, the data collector first announces a payment mechanism. Then this mechanism induces a strategic form game where the individuals are the players. The utility of each individual is the difference between her payment and her privacy cost. We assume that the individuals are risk neutral, i.e., they are interested in maximizing their expected utility. In this game, the prior distribution P_W , the signal quality parameter θ , the form of the payment mechanism $\mathbf{R}$ and the cost function g are common knowledge.

We focus on Nash equilibria of a payment mechanism, where each individual has no incentive to unilaterally change her strategy given other individuals’ strategies. Formally, a Nash equilibrium in our model is defined as follows.

DEFINITION 2. *A strategy profile σ is a Nash equilibrium in a payment mechanism $\mathbf{R}$ if for any individual i and any strategy σ'_i ,*

$$\mathbb{E}_{\sigma}[R_i(\mathbf{X}) - g(\sigma_i)] \geq \mathbb{E}_{(\sigma'_i, \sigma_{-i})}[R_i(\mathbf{X}) - g(\sigma'_i)],$$

where the expectation is over the reported data $\mathbf{X}$, and the subscripts σ and (σ'_i, σ_{-i}) indicate that $\mathbf{X}$ is generated by the strategy profile σ and (σ'_i, σ_{-i}) , respectively.

4. THE VALUE OF DATA PRIVACY

We say that the data collector *obtains* ϵ units of privacy from an individual i in a payment mechanism if individual i ’s best response in a Nash equilibrium of the mechanism is to report data with a privacy level of ϵ . Recall that we are interested in the regime $\epsilon > 0$ since the data collector wants the reported data to be useful for data analysis. Let $\mathcal{R}(i; \epsilon)$ denote the set of nonnegative payment mechanisms in which the data collector obtains ϵ units of privacy from individual i . Then we measure the value of ϵ units of privacy by the minimum payment to individual i of all mechanisms in $\mathcal{R}(i; \epsilon)$. Note that this measure does not depend on the specific identity of i due to the symmetry across individuals. For any mechanism $\mathbf{R} \in \mathcal{R}(i; \epsilon)$, let $\sigma^{(\mathbf{R}; \epsilon)}$ denote the corresponding Nash equilibrium. Then, formally, the value of ϵ units of privacy is measured by

$$V(\epsilon) = \inf_{\mathbf{R} \in \mathcal{R}(i; \epsilon)} \mathbb{E}_{\sigma^{(\mathbf{R}; \epsilon)}}[R_i(\mathbf{X})]. \quad (4)$$

In this section, we first derive a lower bound on $V(\epsilon)$ by characterizing the Nash equilibria and replicating mechanisms in $\mathcal{R}(i; \epsilon)$ by genie-aided mechanisms. We then design a payment mechanism in $\mathcal{R}(i; \epsilon)$, and consequently the equilibrium payment to individual i in this mechanism serves as an upper bound of $V(\epsilon)$. The gap between the lower and upper bounds diminishes to zero exponentially fast as the number of individuals becomes large, which indicates that the lower and upper bounds are asymptotically tight.

4.1 Lower Bound

We present a lower bound on $V(\epsilon)$ in Theorem 1 below. For convenience, we define

$$V_{\text{LB}}(\epsilon) = g'(\epsilon) \frac{e^\epsilon + 1}{e^\epsilon} \left(\frac{\theta}{2\theta - 1} (e^\epsilon + 1) - 1 \right), \quad (5)$$

where g' is the derivative of the privacy cost function of an individual and θ is the quality of signals.

THEOREM 1. *The value of ϵ units of privacy measured in (4) for any $\epsilon > 0$ is lower bounded as $V(\epsilon) \geq V_{\text{LB}}(\epsilon)$. Specifically, for any nonnegative payment mechanism $\mathbf{R}$, if the strategy of an individual i in a Nash equilibrium has a privacy level of ϵ with $\epsilon > 0$, then the expected payment to individual i at this equilibrium is lower bounded by $V_{\text{LB}}(\epsilon)$.*

We remark that the lower bound in Theorem 1 can be achieved by a hypothetical payment mechanism in which a genie who knows the realization of the underlying state W guides the data collector on how much to pay each individual. Intuitively, the knowledge of the state W provides

more information about the system, which helps the data collector to obtain privacy with less payment. While it may sound like a chicken-and-egg problem as the data collector's sole purpose of paying individuals for their private data is to learn the state W , it will become clear that the philosophy carries over and the data collector should utilize the best estimate of W in the payment mechanism to minimize the payment. The insight we gain from this mechanism sheds light on the asymptotically tight upper bound on the value of privacy in Section 4.2.

This genie-aided payment mechanism, denoted by $\hat{\mathbf{R}}^{(\epsilon)}$, determines the payment to each individual i based on her own reported data X_i and the state W as follows:

$$\hat{R}_i^{(\epsilon)}(X_i, W) = \frac{g'(\epsilon)(e^\epsilon + 1)^2}{2e^\epsilon} \hat{A}_{X_i, W}, \quad (6)$$

where

$$\hat{A}_{1,1} = \frac{1}{(2\theta - 1)P_W(1)}, \quad \hat{A}_{0,0} = \frac{1}{(2\theta - 1)P_W(0)}, \\ \hat{A}_{0,1} = \hat{A}_{1,0} = 0.$$

In this mechanism, it can be proved that the best response of individual i is the following symmetric randomized response, denoted by $\sigma_i^{(\epsilon)}$, which is ϵ -differentially private:

$$\mathbb{P}_{\sigma_i^{(\epsilon)}}(X_i = 1 \mid S_i = 1) = \mathbb{P}_{\sigma_i^{(\epsilon)}}(X_i = 0 \mid S_i = 0) = \frac{e^\epsilon}{e^\epsilon + 1}, \\ \mathbb{P}_{\sigma_i^{(\epsilon)}}(X_i = 1 \mid S_i = 0) = \mathbb{P}_{\sigma_i^{(\epsilon)}}(X_i = 0 \mid S_i = 1) = \frac{1}{e^\epsilon + 1}, \\ \mathbb{P}_{\sigma_i^{(\epsilon)}}(X_i = \perp \mid S_i = 1) = \mathbb{P}_{\sigma_i^{(\epsilon)}}(X_i = \perp \mid S_i = 0) = 0.$$

For convenience, we will refer to this strategy as the ϵ -strategy. The expected payment to individual i at this strategy equals to the lower bound in Theorem 1.

Next we sketch the proof of Theorem 1. We first give three lemmas that form the basis of the proof, and then present the proof based on that. The proofs of these lemmas can be found in our technical report [31].

4.1.1 Characterization of Nash Equilibria

We first characterize individuals' behavior in a Nash equilibrium. In general, an ϵ -differentially private strategy has uncountably many possible forms. However, provided that the strategy is part of a Nash equilibrium (i.e., a best response of an individual), the following lemma substantially reduces the space of possibilities. We remark that a similar phenomenon for privacy-aware individuals has been observed in [4] in a different setting.

LEMMA 1. *In any nonnegative payment mechanism, an individual's strategy in a Nash equilibrium is either a symmetric randomized response, or a non-informative strategy.*

We remark that Lemma 1 holds for more general probability models of the signals. The proof carries over as long as the support of the joint distribution of the signals is the entire domain $\{0, 1\}^N$.

By Lemma 1, if an individual's strategy in a Nash equilibrium has a privacy level of ϵ , where $\epsilon > 0$, this equilibrium strategy is either the ϵ -strategy or the $(-\epsilon)$ -strategy. The following lemma says that from the payment perspective, it suffices to further focus on the case that it is the ϵ -strategy.

LEMMA 2. *For any nonnegative payment mechanism $\mathbf{R}$ in which the strategy profile $(\sigma_i^{(-\epsilon)}, \sigma_{-i})$ with some $\epsilon > 0$ is a Nash equilibrium, there exists another nonnegative payment mechanism $\mathbf{R}'$ in which $(\sigma_i^{(\epsilon)}, \sigma_{-i})$ is a Nash equilibrium, and the expected payment to each individual at these two equilibria of the two mechanisms are the same.*

This lemma is proved by considering the payment mechanism $\mathbf{R}'$ that is constructed by applying $\mathbf{R}$ on the reported data after modifying X_i to $1 - X_i$.

4.1.2 Genie-Aided Payment Mechanism

A genie-aided payment mechanism $\hat{\mathbf{R}}: \mathcal{X}^N \times \{0, 1\} \rightarrow \mathbb{R}^N$ determines the payment to an individual based on not only the reported data $\mathbf{X}$ but also the underlying state W . Compared with a standard payment mechanism, a genie-aided mechanism is hypothetical since the data collector has access to the underlying state, as if she were aided by a genie. We consider nonnegative genie-aided payment mechanisms where $\hat{R}_i(\mathbf{X}, W)$, the payment to individual i , depends on only her own reported data X_i and the underlying state W . We write $\hat{R}_i(X_i, W)$ to represent $\hat{R}_i(\mathbf{X}, W)$ for conciseness. Therefore, for each individual i , a genie-aided mechanism makes use of the information of W but discards the information in $\mathbf{X}_{-i}$. The following lemma shows that the expected payments resulting from any Nash equilibrium of any payment mechanism can be replicated by a genie-aided payment mechanism with the same Nash equilibrium. Thus we can restrict our attention to genie-aided mechanisms to obtain a lower bound on the value of privacy.

LEMMA 3. *For any nonnegative payment mechanism $\mathbf{R}$ and any Nash equilibrium σ of it, there exists a nonnegative genie-aided mechanism $\hat{\mathbf{R}}$, such that σ is also a Nash equilibrium of $\hat{\mathbf{R}}$ and the expected payment to each individual at this equilibrium is the same under $\mathbf{R}$ and $\hat{\mathbf{R}}$.*

This lemma is proved by constructing the following genie-aided payment mechanism $\hat{\mathbf{R}}$ according to the desired equilibrium σ : for any individual i and any $x_i \in \mathcal{X}, w \in \{0, 1\}$,

$$\hat{R}_i(x_i, w) = \bar{R}_i(x_i; w) := \mathbb{E}_\sigma[R_i(\mathbf{X}) \mid X_i = x_i, W = w].$$

Our intuition is as follows. A genie-aided mechanism can use the state W to generate an incentive to individual i , which "mimics" the incentive provided by the reported data $\mathbf{X}_{-i}$ of others. The above genie-aided payment mechanism $\hat{\mathbf{R}}$ is constructed such that no matter what strategy individual i uses, her expected utility is the same under $\mathbf{R}$ and $\hat{\mathbf{R}}$. Since an individual calculates her best response according to the expected utility, her equilibrium behavior and expected payment are the same under $\hat{\mathbf{R}}$ and $\mathbf{R}$. We remark that the Nash equilibria of a genie-aided mechanism are much easier to analyze since the individuals are decoupled in the payments and thus an individual's strategy does not have an influence on other individuals' utility.

Let $\hat{\mathcal{R}}(i; \epsilon)$ denote the set of nonnegative genie-aided payment mechanisms in which the ϵ -strategy is an individual i 's strategy in a Nash equilibrium, and let $\sigma_i^{(\epsilon)}$ denote the ϵ -strategy. Consider

$$\hat{V}(\epsilon) = \inf_{\hat{\mathbf{R}} \in \hat{\mathcal{R}}(i; \epsilon)} \mathbb{E}_{\sigma_i^{(\epsilon)}}[\hat{R}_i(X_i, W)],$$

which is a definition similar to the value of ϵ units of privacy, $V(\epsilon)$, measured in (4). Then $\widehat{V}(\epsilon) \leq V(\epsilon)$ for the following reasons. Consider any $\mathbf{R} \in \mathcal{R}(i; \epsilon)$, i.e., any nonnegative payment mechanism $\mathbf{R}$ in which individual i 's strategy in a Nash equilibrium has a privacy level of ϵ . With Lemma 1 and 2, we can assume without loss of generality that this equilibrium strategy is the ϵ -strategy. Then by Lemma 3, we can map $\mathbf{R}$ to a $\widehat{\mathbf{R}} \in \widehat{\mathcal{R}}(i; \epsilon)$, such that

$$\mathbb{E}_{\sigma(\mathbf{R}; \epsilon)}[R_i(\mathbf{X})] = \mathbb{E}_{\sigma_i(\epsilon)}[\widehat{R}_i(X_i, W)].$$

Therefore, the infimum over $\widehat{\mathcal{R}}(i; \epsilon)$ is no greater than the infimum over $\mathcal{R}(i; \epsilon)$, i.e., $\widehat{V}(\epsilon) \leq V(\epsilon)$.

4.1.3 Proof of Theorem 1

With Lemma 1, 2 and 3, we can prove the lower bound in Theorem 1 by focusing on the genie-aided mechanisms in $\widehat{\mathcal{R}}(i; \epsilon)$. Then there is no need to consider the strategies of individuals other than individual i since a genie-aided mechanism pays individual i only according to X_i and W . A necessary condition for the ϵ -strategy to be a best response of individual i is that ϵ yields no worse expected payment than other privacy levels. We utilize this necessary condition to obtain a lower bound on the expected payment to individual i , which gives a lower bound on $\widehat{V}(\epsilon)$ and further proves the lower bound in Theorem 1.

PROOF OF THEOREM 1. By Lemma 1, 2 and 3, it suffices to focus on nonnegative genie-aided payment mechanisms in which the ϵ -strategy is an individual i 's strategy in a Nash equilibrium, i.e., mechanisms in $\widehat{\mathcal{R}}(i; \epsilon)$. Consider any $\widehat{\mathbf{R}} \in \widehat{\mathcal{R}}(i; \epsilon)$ and denote the ϵ -strategy by $\sigma_i^{(\epsilon)}$. Consider the ξ -strategy of individual i with any $\xi \geq 0$ and denote it by $\sigma_i^{(\xi)}$. Then the expected utility of individual i at the strategy $\sigma_i^{(\xi)}$ can be written as

$$\begin{aligned} & \mathbb{E}_{\sigma_i^{(\xi)}}[\widehat{R}_i(X_i, W)] - g(\sigma_i^{(\xi)}) \\ &= \sum_{x_i, s_i, w} \mathbb{P}_{\sigma_i^{(\xi)}}(X_i = x_i \mid S_i = s_i) \mathbb{P}(S_i = s_i, W = w) \widehat{R}_i(x_i, w) \\ & \quad - g(\xi), \\ &= \overline{K}_1 \frac{e^\xi}{e^\xi + 1} + \overline{K}_0 \frac{1}{e^\xi + 1} + \overline{K} - g(\xi), \end{aligned}$$

where

$$\begin{aligned} \overline{K}_1 &= \{\widehat{R}_i(1, 1)P_W(1)\theta + \widehat{R}_i(1, 0)P_W(0)(1 - \theta) \\ & \quad - \{\widehat{R}_i(0, 1)P_W(1)\theta + \widehat{R}_i(0, 0)P_W(0)(1 - \theta)\}, \\ \overline{K}_0 &= \{\widehat{R}_i(1, 1)P_W(1)(1 - \theta) + \widehat{R}_i(1, 0)P_W(0)\theta \\ & \quad - \{\widehat{R}_i(0, 1)P_W(1)(1 - \theta) + \widehat{R}_i(0, 0)P_W(0)\theta\}, \\ \overline{K} &= \widehat{R}_i(0, 1)P_W(1) + \widehat{R}_i(0, 0)P_W(0). \end{aligned}$$

It can be seen that $\overline{K}_1$, $\overline{K}_0$ and $\overline{K}$ do not depend on ξ . Let this expected utility define a function f of ξ ; i.e.,

$$f(\xi) = \overline{K}_1 \frac{e^\xi}{e^\xi + 1} + \overline{K}_0 \frac{1}{e^\xi + 1} - g(\xi) + \overline{K}.$$

Then a necessary condition for the ϵ -strategy to be an equilibrium strategy is that ϵ maximizes $f(\xi)$, which implies that $f'(\epsilon) = 0$ since $\epsilon > 0$. Since

$$f'(\xi) = (\overline{K}_1 - \overline{K}_0) \frac{e^\xi}{(e^\xi + 1)^2} - g'(\xi),$$

setting $f'(\epsilon) = 0$ yields that

$$\overline{K}_1 - \overline{K}_0 = g'(\epsilon) \frac{(e^\epsilon + 1)^2}{e^\epsilon}. \quad (7)$$

Now we calculate the expected payment to individual i at the ϵ -strategy:

$$\mathbb{E}_{\sigma_i^{(\epsilon)}}[\widehat{R}_i(X_i, W)] = -(\overline{K}_1 - \overline{K}_0) \frac{1}{e^\epsilon + 1} + (\overline{K}_1 + \overline{K}).$$

By definition,

$$\begin{aligned} \overline{K}_1 + \overline{K} &= \widehat{R}_i(1, 1)P_W(1)\theta + \widehat{R}_i(1, 0)P_W(0)(1 - \theta) \\ & \quad + \widehat{R}_i(0, 1)P_W(1)(1 - \theta) + \widehat{R}_i(0, 0)P_W(0)\theta, \end{aligned}$$

and

$$\begin{aligned} \overline{K}_1 - \overline{K}_0 &= (\widehat{R}_i(1, 1) - \widehat{R}_i(0, 1))P_W(1)(2\theta - 1) \\ & \quad + (\widehat{R}_i(0, 0) - \widehat{R}_i(1, 0))P_W(0)(2\theta - 1). \end{aligned}$$

Therefore,

$$\begin{aligned} \overline{K}_1 + \overline{K} &= \frac{\theta}{2\theta - 1}(\overline{K}_1 - \overline{K}_0) \\ & \quad + \widehat{R}_i(1, 0)P_W(0) + \widehat{R}_i(0, 1)P_W(1) \\ &\geq \frac{\theta}{2\theta - 1}(\overline{K}_1 - \overline{K}_0) \\ &= g'(\epsilon) \frac{(e^\epsilon + 1)^2}{e^\epsilon} \frac{\theta}{2\theta - 1}, \end{aligned}$$

where we have used the nonnegativity of $\widehat{\mathbf{R}}$. Then the expected payment to individual i is bounded as follows:

$$\begin{aligned} & \mathbb{E}_{\sigma_i^{(\epsilon)}}[\widehat{R}_i(X_i, W)] \\ &= -(\overline{K}_1 - \overline{K}_0) \frac{1}{e^\epsilon + 1} + (\overline{K}_1 + \overline{K}) \\ &\geq g'(\epsilon) \frac{e^\epsilon + 1}{e^\epsilon} \left(\frac{\theta}{2\theta - 1} (e^\epsilon + 1) - 1 \right), \end{aligned} \quad (8)$$

which proves the lower bound. $\square$

Now beyond the proof, we take a moment to check when this lower bound can be achieved. To achieve the lower bound, we need the equality in (8) to hold and the equation (7) to be satisfied, which is equivalent to the following conditions:

$$\widehat{R}_i(1, 0) = 0, \quad (9)$$

$$\widehat{R}_i(0, 1) = 0, \quad (10)$$

$$\begin{aligned} & (2\theta - 1) \left(\widehat{R}_i(1, 1)P_W(1) + \widehat{R}_i(0, 0)P_W(0) \right) \\ &= g'(\epsilon) \frac{(e^\epsilon + 1)^2}{e^\epsilon}. \end{aligned} \quad (11)$$

It is easy to check that the genie-aided payment mechanism $\widehat{\mathbf{R}}^{(\epsilon)}$ defined in (6) is in $\widehat{\mathcal{R}}(i; \epsilon)$ and satisfies (9)–(11), and therefore achieves the lower bound. Can this lower bound be achieved by a standard nonnegative payment mechanism? Consider any payment mechanism $\mathbf{R} \in \mathcal{R}(i; \epsilon)$. Following similar arguments, we can prove that to achieve the lower

bound, $\mathbf{R}$ needs to satisfy the following conditions:

$$\bar{R}_i(1; 0) = 0, \quad (12)$$

$$\bar{R}_i(0; 1) = 0, \quad (13)$$

$$(2\theta - 1)(\bar{R}_i(1; 1)P_W(1) + \bar{R}_i(0; 0)P_W(0)) = g'(\epsilon) \frac{(e^\epsilon + 1)^2}{e^\epsilon}, \quad (14)$$

where recall that $\bar{R}_i(x_i; w) = \mathbb{E}_{\sigma(\mathbf{R}, \epsilon)}[R_i(\mathbf{X}) \mid X_i = x_i, W = w]$ for $x_i, w \in \{0, 1\}$. It can be proved that if $\mathbf{R}$ satisfies (12) and (13), then $R_i(\mathbf{x}) = 0$ for any $\mathbf{x} \in \mathcal{X}^N$, which contradicts (14). Therefore, no standard nonnegative payment mechanism can achieve the lower bound. However, as will be shown in the next section, we can design a class of standard nonnegative payment mechanisms such that the expected payment approaches the lower bound as the number of individuals increases. The design follows the insights indicated by the genie-aided mechanism $\hat{\mathbf{R}}^{(\epsilon)}$: to minimize the payment, the data collector should utilize the best estimate of W in the payment mechanism based on the noisy reports.

4.2 Upper Bound

We present an upper bound on $V(\epsilon)$ in Theorem 2 below. For convenience, we define

$$d = \frac{1}{2} \ln \frac{(e^\epsilon + 1)^2}{4(\theta e^\epsilon + 1 - \theta)((1 - \theta)e^\epsilon + \theta)}, \quad (15)$$

where θ is the quality of signal. Note that $d > 0$ for any $\epsilon > 0$. Recall that $V_{\text{LB}}(\epsilon)$ is the lower bound in Theorem 1.

THEOREM 2. *The value of ϵ units of privacy measured in (4) is upper bounded as $V(\epsilon) \leq V_{\text{LB}}(\epsilon) + O(e^{-Nd})$, where the $O(\cdot)$ is for $N \rightarrow \infty$. Specifically, there exists a nonnegative payment mechanism $\mathbf{R}^{(N, \epsilon)}$ in which the strategy profile $\sigma^{(\epsilon)}$ consisting of ϵ -strategies is a Nash equilibrium, and the expected payment to each individual i at this equilibrium is upper bounded by $V_{\text{LB}}(\epsilon) + O(e^{-Nd})$.*

Comparing this upper bound with the lower bound $V_{\text{LB}}(\epsilon)$ in Theorem 1 we can see that the gap between the lower and upper bounds is just the term $O(e^{-Nd})$, which diminishes to zero exponentially fast as N goes to infinity.

We present the payment mechanism $\mathbf{R}^{(N, \epsilon)}$ in Section 4.2.1. We will show that under $\mathbf{R}^{(N, \epsilon)}$, the strategy profile $\sigma^{(\epsilon)}$ consisting of ϵ -strategies is a Nash equilibrium. Therefore, $\mathbf{R}^{(N, \epsilon)}$ is a member of $\mathcal{R}(i; \epsilon)$, and the payment to individual i at $\sigma^{(\epsilon)}$ gives an upper bound on the value of privacy.

The design of $\mathbf{R}^{(N, \epsilon)}$ is enlightened by the hypothetical payment mechanism $\hat{\mathbf{R}}^{(\epsilon)}$ defined in (6). But without direct access to the state W , the mechanism $\mathbf{R}^{(N, \epsilon)}$ relies on the reported data from an individual i 's peers, i.e., individuals other than individual i , to obtain an estimate of W . We borrow the idea of the peer-prediction method [23], which rewards more for the agreement between an individual and her peers to encourage truthful reporting. However, unlike the peer-prediction method, the individuals here have privacy concerns and they will weigh the privacy cost against the payment to choose the best privacy level. We modify the payments in $\hat{\mathbf{R}}^{(\epsilon)}$ to ensure that the ϵ -strategy is still a best response of each individual in $\mathbf{R}^{(N, \epsilon)}$, given that other individuals also follow the ϵ -strategy, which yields the desired Nash equilibrium $\sigma^{(\epsilon)}$.

The equilibrium payment to each individual in $\mathbf{R}^{(N, \epsilon)}$ converges to the lower bound in Theorem 1 as the number of individuals N goes to infinity. The intuition behind is that as the number of individuals N goes to infinity, the majority of the reported data from other individuals converges to the underlying state W , and thus $\mathbf{R}^{(N, \epsilon)}$ works similar as the genie-aided mechanism $\hat{\mathbf{R}}^{(\epsilon)}$, whose equilibrium payment to each individual equals to the lower bound in Theorem 1.

4.2.1 A Payment Mechanism $\mathbf{R}^{(N, \epsilon)}$

The payment mechanism $\mathbf{R}^{(N, \epsilon)}$ is designed for purchasing private data from N privacy-aware individuals, parameterized by a privacy parameter ϵ , where $N \geq 2$ and $\epsilon > 0$.

1. Each individual reports her data (which can be the decision of not participating).
2. The data collector counts the number of participants, which is denoted by n .
3. For non-participating individuals, the payment is zero.
4. If there is only one participant, pay zero to this participant. Otherwise, for each participating individual i , the data collector computes the variable

$$M_{-i} = \begin{cases} 1 & \text{if } \sum_{j: X_j \neq 1, j \neq i} X_j \geq \lfloor \frac{n-1}{2} \rfloor + 1, \\ 0 & \text{otherwise,} \end{cases}$$

which is the majority of the other participants' reported data. Then the data collector pays individual i the following amount of payment according to her reported data X_i and M_{-i} :

$$R_i^{(N, \epsilon)}(\mathbf{X}) = \frac{g'(\epsilon)(e^\epsilon + 1)^2}{2e^\epsilon} A_{X_i, M_{-i}},$$

where the parameters $A_{1,1}, A_{0,0}, A_{0,1}, A_{1,0}$ are defined in Section 4.2.2.

4.2.2 Payment Parameterization

Let

$$\alpha = \theta \frac{e^\epsilon}{e^\epsilon + 1} + (1 - \theta) \frac{1}{e^\epsilon + 1}.$$

The physical meaning of α can be seen by considering the strategy profile $\sigma^{(\epsilon)}$, where given the state W , the reported data $X_1, X_2, \dots, X_N$ are i.i.d. with

$$\mathbb{P}_{\sigma^{(\epsilon)}}(X_i = 1 \mid W = 1) = \mathbb{P}_{\sigma^{(\epsilon)}}(X_i = 0 \mid W = 0) = \alpha.$$

Given that the number of participants is n with $n \geq 2$, define the following quantities. Consider a random variable that follows the binomial distribution with parameters $n-1$ and α . Let $\beta^{(n)}$ denote the probability that this random variable is greater than or equal to $\lfloor \frac{n-1}{2} \rfloor + 1$. Let

$$\gamma^{(n)} = \begin{cases} 1 - \binom{n-1}{\frac{n-1}{2}} \alpha^{\frac{n-1}{2}} (1 - \alpha)^{\frac{n-1}{2}} & \text{if } n-1 \text{ is even,} \\ 1 & \text{if } n-1 \text{ is odd.} \end{cases} \quad (16)$$

To see the physical meaning of $\beta^{(n)}$ and $\gamma^{(n)}$, still consider $\sigma^{(\epsilon)}$, where the number of participants is $n = N$. Then for an individual i ,

$$\begin{aligned} \mathbb{P}_{\sigma^{(\epsilon)}}(M_{-i} = 1 \mid W = 1) &= \beta^{(N)}, \\ \mathbb{P}_{\sigma^{(\epsilon)}}(M_{-i} = 1 \mid W = 0) &= \gamma^{(N)} - \beta^{(N)}. \end{aligned}$$

With the introduced notation, the parameters $A_{1,1}$, $A_{0,0}$, $A_{0,1}$, $A_{1,0}$ used in the payment mechanism $\mathbf{R}^{(N,\epsilon)}$ are defined as follows:

$$\begin{aligned} A_{1,1} &= \frac{P_W(1)(1 - \beta^{(n)}) + P_W(0)(1 - (\gamma^{(n)} - \beta^{(n)}))}{(2\beta^{(n)} - \gamma^{(n)})(2\theta - 1)P_W(1)P_W(0)}, \\ A_{0,0} &= \frac{P_W(1)\beta^{(n)} + P_W(0)(\gamma^{(n)} - \beta^{(n)})}{(2\beta^{(n)} - \gamma^{(n)})(2\theta - 1)P_W(1)P_W(0)}, \\ A_{0,1} &= 0, \\ A_{1,0} &= 0. \end{aligned}$$

It is easy to verify that these parameters are nonnegative. Thus $\mathbf{R}^{(N,\epsilon)}$ is a nonnegative payment mechanism. The proof of the equilibrium properties of $\mathbf{R}^{(N,\epsilon)}$ in Theorem 2 is given below.

4.2.3 Proof of Theorem 2

PROOF. It suffices to prove that the strategy profile $\sigma^{(\epsilon)}$ is a Nash equilibrium in $\mathbf{R}^{(N,\epsilon)}$ and the expected payment to each individual i at this equilibrium satisfies that $\mathbb{E}_{\sigma^{(\epsilon)}}[R_i^{(N,\epsilon)}(\mathbf{X})] \leq V_{LB}(\epsilon) + O(e^{-Nd})$, where recall that $V_{LB}(\epsilon)$ is defined in (5). For conciseness, in the remainder of this proof, we suppress the explicit dependence on N and ϵ , and write $\mathbf{R}$ and σ to represent $\mathbf{R}^{(N,\epsilon)}$ and $\sigma^{(\epsilon)}$, respectively.

We first prove that the strategy profile σ is a Nash equilibrium in $\mathbf{R}$; i.e., for any individual i , the ϵ -strategy is a best response of individual i when other individuals follow σ_{-i} . Following the notation in the proof of Lemma 1, for any individual i we consider any strategy σ'_i of individual i and let

$$\begin{aligned} p_1 &= \mathbb{P}_{\sigma'_i}(X_i = 1 \mid S_i = 1), \quad q_1 = \mathbb{P}_{\sigma'_i}(X_i = 0 \mid S_i = 1), \\ p_0 &= \mathbb{P}_{\sigma'_i}(X_i = 1 \mid S_i = 0), \quad q_0 = \mathbb{P}_{\sigma'_i}(X_i = 0 \mid S_i = 0). \end{aligned}$$

Then by the proof of Lemma 1, the best response satisfies either $p_1 = p_0, q_1 = q_0$, or $p_1 = q_0, p_0 = q_1, p_1 + q_1 = 1$, depending on the form of the utility function $U_i(p_1, p_0, q_1, q_0)$, which is the expected utility of individual i at the strategy σ'_i when other individuals follow σ_{-i} . Thus, we derive the form of $U_i(p_1, p_0, q_1, q_0)$ next. Recall that we let $\bar{R}_i(x_i; w)$ denote $\mathbb{E}_{(\sigma'_i, \sigma_{-i})}[R_i(\mathbf{X}) \mid X_i = x_i, W = w]$ for $x_i, w \in \{0, 1\}$. Then

$$\begin{aligned} U_i(p_1, p_0, q_1, q_0) &= \mathbb{E}_{(\sigma'_i, \sigma_{-i})}[R_i(\mathbf{X}) - g(\zeta(\sigma'_i))] \\ &= K_1 p_1 + K_0 p_0 + L_1 q_1 + L_0 q_0 - g(\zeta(p_1, p_0, q_1, q_0)), \end{aligned}$$

with

$$\begin{aligned} K_1 &= \{\bar{R}_i(1; 1)P_W(1)\theta + \bar{R}_i(1; 0)P_W(0)(1 - \theta)\}, \\ K_0 &= \{\bar{R}_i(1; 1)P_W(1)(1 - \theta) + \bar{R}_i(1; 0)P_W(0)\theta\}, \\ L_1 &= \{\bar{R}_i(0; 1)P_W(1)\theta + \bar{R}_i(0; 0)P_W(0)(1 - \theta)\}, \\ L_0 &= \{\bar{R}_i(0; 1)P_W(1)(1 - \theta) + \bar{R}_i(0; 0)P_W(0)\theta\}. \end{aligned}$$

In the designed mechanism $\mathbf{R}$, the payment to individual i only depends on X_i and M_{-i} . Thus we write $R_i(X_i; M_{-i}) = R_i(\mathbf{X})$. Then the value of $\bar{R}_i(x_i; w)$ is calculated as follows:

$$\begin{aligned} \bar{R}_i(1; 1) &= \mathbb{E}_{(\sigma'_i, \sigma_{-i})}[R_i(\mathbf{X}) \mid X_i = 1, W = 1] \\ &= \beta^{(N)} R_i(1; 1), \\ \bar{R}_i(1; 0) &= \mathbb{E}_{(\sigma'_i, \sigma_{-i})}[R_i(\mathbf{X}) \mid X_i = 1, W = 0] \end{aligned}$$

$$\begin{aligned} &= (\gamma^{(N)} - \beta^{(N)}) R_i(1; 1), \\ \bar{R}_i(0; 1) &= \mathbb{E}_{(\sigma'_i, \sigma_{-i})}[R_i(\mathbf{X}) \mid X_i = 0, W = 1] \\ &= (1 - \beta^{(N)}) R_i(0; 0), \\ \bar{R}_i(0; 0) &= \mathbb{E}_{(\sigma'_i, \sigma_{-i})}[R_i(\mathbf{X}) \mid X_i = 0, W = 0] \\ &= (1 - (\gamma^{(N)} - \beta^{(N)})) R_i(0; 0), \end{aligned}$$

and it can be verified that K_1, K_0, L_1 and L_0 are all positive. Therefore, by the proof of Lemma 1, the possibility for the best response to be $p_1 = p_0, q_1 = q_0, 0 < p_1 + q_1 < 1$ can be eliminated and the best response strategy must be in one of the following three forms:

$$p_1 = p_0 = q_1 = q_0 = 0, \quad (17)$$

$$p_1 = p_0, \quad q_1 = q_0, \quad p_1 + q_1 = 1, \quad (18)$$

$$p_1 = q_0, \quad p_0 = q_1, \quad p_1 + q_1 = 1. \quad (19)$$

The strategy in (17) is to always not participate, which yields an utility of zero. For strategies in the form of (18) or (19), we can write the expected utility as a function of p_1 and p_0 as follows:

$$\bar{U}_i(p_1, p_0) = \bar{K}_1 p_1 + \bar{K}_0 p_0 + \bar{K} - g(\zeta(p_1, p_0)),$$

where $\bar{K}_1 = K_1 - L_1$, $\bar{K}_0 = K_0 - L_0$, $\bar{K} = L_1 + L_0$, and with a little abuse of notation, $\zeta(p_1, p_0) = \max\left\{\left|\ln \frac{p_1}{p_0}\right|, \left|\ln \frac{1-p_1}{1-p_0}\right|\right\}$.

Inserting the value of $R_i(X_i; M_{-i})$ gives

$$\bar{K}_1 = \frac{g'(\epsilon)(e^\epsilon + 1)^2}{2e^\epsilon}, \quad \bar{K}_0 = -\frac{g'(\epsilon)(e^\epsilon + 1)^2}{2e^\epsilon}.$$

Then a strategy in the form of (18) yields an utility of $\bar{K} > 0$. A strategy in the form of (19) can be written as

$$p_1 = q_0 = \frac{e^\xi}{e^\xi + 1}, \quad p_0 = q_1 = \frac{1}{e^\xi + 1}.$$

Then the expected utility can be further written as a function f of ξ as follows:

$$f(\xi) = \bar{K}_1 \frac{e^\xi}{e^\xi + 1} + \bar{K}_0 \frac{1}{e^\xi + 1} - g(|\xi|) + \bar{K}.$$

Therefore, to prove that the ϵ -strategy is a best response of individual i , it suffices to prove that ϵ maximizes $f(\xi)$ and $f(\epsilon) \geq \bar{K}$. For any $\xi < 0$, it is easy to see that

$$\bar{K}_1 \frac{e^\xi}{e^\xi + 1} + \bar{K}_0 \frac{1}{e^\xi + 1} < 0 < \bar{K}_1 \frac{e^{-\xi}}{e^{-\xi} + 1} + \bar{K}_0 \frac{1}{e^{-\xi} + 1}.$$

Thus $f(\xi)$ achieves its maximum value at some $\xi \geq 0$. For any $\xi \geq 0$,

$$\begin{aligned} f'(\xi) &= (\bar{K}_1 - \bar{K}_0) \frac{e^\xi}{(e^\xi + 1)^2} - g'(\xi), \\ f''(\xi) &= -(\bar{K}_1 - \bar{K}_0) \frac{e^\xi(e^\xi - 1)}{(e^\xi + 1)^3} - g''(\xi) \leq 0, \end{aligned}$$

where the second inequality is due to the convexity of the cost function g . Therefore, f is concave. Since $f'(\epsilon) = 0$, ϵ maximizes $f(\xi)$. The optimal value is

$$f(\epsilon) = g'(\epsilon) \frac{e^\epsilon - e^{-\epsilon}}{2} - g(\epsilon) + \bar{K}.$$

By the convexity of g , $g(\epsilon) \leq g'(\epsilon)\epsilon \leq g'(\epsilon) \frac{e^\epsilon - e^{-\epsilon}}{2}$. Thus $f(\epsilon) \geq \bar{K}$, which completes the proof for the ϵ -strategy to be a best response of individual i .

Next we calculate the expected payment to individual i at σ , which can be written as

$$\mathbb{E}_\sigma[R_i(\mathbf{X})] = -(\bar{K}_1 - \bar{K}_0) \frac{1}{e^\epsilon + 1} + \bar{K}_1 + \bar{K}.$$

By definitions,

$$\begin{aligned} & \bar{K}_1 + \bar{K} \\ &= \frac{g'(\epsilon)(e^\epsilon + 1)^2}{2e^\epsilon} \frac{1}{(2\beta^{(N)} - \gamma^{(N)})(2\theta - 1)} \\ & \quad \cdot \left(2(\beta^{(N)})^2 + (4\theta - 2 - 2\gamma^{(N)})\beta^{(N)} \right. \\ & \quad \left. + 2(1 - \theta)\gamma^{(N)} + \beta^{(N)}(1 - \beta^{(N)}) \frac{P_W(1)}{P_W(0)} \right. \\ & \quad \left. + (\gamma^{(N)} - \beta^{(N)})(1 - (\gamma^{(N)} - \beta^{(N)})) \frac{P_W(0)}{P_W(1)} \right) \\ &=: \frac{g'(\epsilon)(e^\epsilon + 1)^2}{2e^\epsilon} h(\beta^{(N)}). \end{aligned}$$

Then

$$\begin{aligned} \mathbb{E}_\sigma[R_i(\mathbf{X})] &= \frac{g'(\epsilon)(e^\epsilon + 1)}{e^\epsilon} \left(\frac{1}{2} h(\beta^{(N)})(e^\epsilon + 1) - 1 \right) \\ &= V_{LB}(\epsilon) + \frac{g'(\epsilon)(e^\epsilon + 1)^2}{2e^\epsilon} \left(h(\beta^{(N)}) - \frac{2\theta}{2\theta - 1} \right). \end{aligned}$$

To derive an upper bound on the expected payment, we first analyze the function h . Rearranging terms gives

$$\begin{aligned} h(\beta^{(N)}) &= \frac{1}{2\theta - 1} \frac{1}{2\beta^{(N)} - \gamma^{(N)}} \\ & \quad \cdot \left((2 - t)(\beta^{(N)})^2 + \left(4\theta - 2 - 2\gamma^{(N)} + \frac{P_W(1)}{P_W(0)} \right. \right. \\ & \quad \left. \left. + (2\gamma^{(N)} - 1) \frac{P_W(0)}{P_W(1)} \right) \beta^{(N)} \right. \\ & \quad \left. + 2(1 - \theta)\gamma^{(N)} + \gamma^{(N)}(1 - \gamma^{(N)}) \frac{P_W(0)}{P_W(1)} \right), \end{aligned}$$

where $t = \frac{(P_W(1))^2 + (P_W(0))^2}{P_W(1)P_W(0)} \geq 2$. Taking derivative yields

$$\begin{aligned} h'(\beta^{(N)}) &= \frac{1}{2\theta - 1} \frac{1}{(2\beta^{(N)} - \gamma^{(N)})^2} \\ & \quad \cdot \left(2(2 - t) \left(\beta^{(N)} - \frac{\gamma^{(N)}}{2} \right)^2 - (\gamma^{(N)})^2 \right. \\ & \quad \left. - \frac{\gamma^{(N)}t}{2} (2 - \gamma^{(N)}) - 2\gamma^{(N)}(1 - \gamma^{(N)}) \right). \end{aligned}$$

Therefore, $h'(\beta^{(N)}) \leq 0$ and h is a non-increasing function.

Next we derive a lower bound on $\beta^{(N)}$. Let $Y_1, Y_2, \dots, Y_{N-1}$ be i.i.d. Bernoulli random variables with parameter α . Then by the definition of $\beta^{(N)}$:

$$\begin{aligned} \beta^{(N)} &= \mathbb{P} \left(\sum_{i=1}^{N-1} Y_i \geq \left\lfloor \frac{N-1}{2} \right\rfloor + 1 \right) \\ &= \gamma^{(N)} - \mathbb{P} \left(\sum_{i=1}^{N-1} (1 - Y_i) \geq N - 1 - \left\lfloor \frac{N-1}{2} \right\rfloor + 1 \right) \\ &\geq \gamma^{(N)} - \mathbb{P} \left(\sum_{i=1}^{N-1} (1 - Y_i) \geq \frac{N-1}{2} \right) \end{aligned}$$

$$\geq \gamma^{(N)} - e^{-(N-1)d},$$

where $d = \frac{1}{2} \ln \frac{1}{4\alpha(1-\alpha)} > 0$ is the parameter defined in (15) and the last inequality follow from the Chernoff bound [28].

By the monotonicity of h ,

$$\begin{aligned} & h(\beta^{(N)}) - \frac{2\theta}{2\theta - 1} \\ &\leq h(\gamma^{(N)} - e^{-(N-1)d}) - \frac{2\theta}{2\theta - 1} \\ &= \frac{1}{2\theta - 1} \frac{1}{\gamma^{(N)} - 2e^{-(N-1)d}} \\ & \quad \cdot \left((2 - t)e^{-2(N-1)d} + \left(2(1 - \gamma^{(N)}) + 2\gamma^{(N)}t \right. \right. \\ & \quad \left. \left. - \frac{P_W(1)}{P_W(0)} - (2\gamma^{(N)} - 1) \frac{P_W(0)}{P_W(1)} \right) e^{-(N-1)d} \right. \\ & \quad \left. + \gamma^{(N)} \frac{P_W(1)}{P_W(0)} + (\gamma^{(N)})^2 \frac{P_W(0)}{P_W(1)} - (\gamma^{(N)})^2 t \right) \\ &\leq \frac{1}{2\theta - 1} \frac{1}{\gamma^{(N)} - 2e^{-(N-1)d}} \\ & \quad \cdot \left((2 - t)e^{-2(N-1)d} + (2(1 - \gamma^{(N)}) + t)e^{-(N-1)d} \right. \\ & \quad \left. + \gamma^{(N)}(1 - \gamma^{(N)}) \frac{P_W(1)}{P_W(0)} \right). \end{aligned}$$

Recall the definition of $\gamma^{(N)}$ in (16). Then when $N - 1$ is odd, $\gamma^{(N)} = 1$, and when $N - 1$ is even,

$$\begin{aligned} 1 - \gamma^{(N)} &= \left(\frac{N-1}{2} \right) \alpha^{\frac{N-1}{2}} (1 - \alpha)^{\frac{N-1}{2}} \\ &= e^{-(N-1)d} \cdot \left(\frac{N-1}{2} \right) 2^{-(N-1)}, \end{aligned}$$

where $\lim_{N \rightarrow \infty} \left(\frac{N-1}{2} \right) 2^{-(N-1)} = 0$. Thus $1 - \gamma^{(N)} = O(e^{-Nd})$ as $N \rightarrow \infty$.

Therefore,

$$\begin{aligned} & \mathbb{E}_\sigma[R_i(\mathbf{X})] \\ &\leq V_{LB}(\epsilon) + \frac{g'(\epsilon)(e^\epsilon + 1)^2}{2e^\epsilon} \left(h(\gamma^{(N)} - e^{-(N-1)d}) - \frac{2\theta}{2\theta - 1} \right) \\ &\leq V_{LB}(\epsilon) + \frac{g'(\epsilon)(e^\epsilon + 1)^2}{2e^\epsilon} \frac{1}{2\theta - 1} \frac{1}{\gamma^{(N)} - 2e^{-(N-1)d}} \\ & \quad \cdot \left((2 - t)e^{-2(N-1)d} + (2(1 - \gamma^{(N)}) + t)e^{-(N-1)d} + O(e^{-Nd}) \right) \\ &= V_{LB}(\epsilon) + O(e^{-Nd}), \end{aligned}$$

as $N \rightarrow \infty$, which completes the proof. $\square$

4.3 Extension to Heterogeneous Cost Functions

Our results on the value of privacy are also valid in the scenario where individuals' privacy cost functions are heterogeneous and known. In this case, the value of ϵ units of privacy is still measured by the minimum payment of all nonnegative payment mechanisms under which an individual's best response in a Nash equilibrium is to report the data with a privacy level of ϵ . However, with heterogeneous cost functions, this value differs from individual to individual. Following similar notation, we let $V_i(\epsilon)$ denote the value of ϵ units of privacy to individual i , and let g_i denote the

cost function of individual i . Then the following lower and upper bounds, which are almost identical to those in Theorem 1 and 2 except the heterogeneous cost function $g_i(\epsilon)$, hold

$$\begin{aligned} g'_i(\epsilon) \frac{e^\epsilon + 1}{e^\epsilon} \left(\frac{\theta}{2\theta - 1} (e^\epsilon + 1) - 1 \right) &\leq V_i(\epsilon) \\ &\leq g'_i(\epsilon) \frac{e^\epsilon + 1}{e^\epsilon} \left(\frac{\theta}{2\theta - 1} (e^\epsilon + 1) - 1 \right) + O(e^{-Nd}). \end{aligned}$$

The lower bound above can be derived directly from the proof of Theorem 1, since the proof does not depend on whether the cost functions are homogeneous or not. The upper bound above is given by a payment mechanism that works similar to $\mathbf{R}^{(N, \epsilon)}$, with the g' in $R_i^{(N, \epsilon)}$ replaced by g'_i . In this mechanism, the strategy profile $\sigma^{(\epsilon)}$ is still a Nash equilibrium, and the expected payment to individual i at this equilibrium can still be upper bounded as in Theorem 2 but again with g' replaced by g'_i .

5. PAYMENT VS. ACCURACY

In this section, we apply the fundamental bounds on the value of privacy to the payment-accuracy problem, where the data collector aims to minimize the total payment while achieving an accuracy target in learning the state. The solution of this problem can be used to guide the design of review systems. For example, to evaluate the underlying value of a new product, a review system can utilize the results in this section to design a payment mechanism for eliciting informative feedback from testers.

5.1 Payment-Accuracy Problem

The data collector learns the state W from the reported data $X_1, X_2, \dots, X_N$, which is collected through some payment mechanism, by performing hypothesis testing between the following two hypotheses:

$$\begin{aligned} H_0: W &= 0, \\ H_1: W &= 1. \end{aligned}$$

The conditional distributions of the reported data given the hypotheses are specified by the strategy profile in a Nash equilibrium of the payment mechanism. According to Lemma 1, we can write an equilibrium strategy profile in the form of $(\sigma_i^{(\epsilon_i)}) = (\sigma_1^{(\epsilon_1)}, \sigma_2^{(\epsilon_2)}, \dots, \sigma_N^{(\epsilon_N)})$ with $\epsilon_i \in \mathbb{R} \setminus \{0\} \cup \{\perp\}$, where recall that $\sigma_i^{(\epsilon_i)}$ is the ϵ_i -strategy. For ease of notation, a non-informative strategy is also called an ϵ -strategy but with $\epsilon = \perp$. Let $\mathcal{R}(\epsilon_1, \epsilon_2, \dots, \epsilon_N)$ denote the set of nonnegative payment mechanisms in which $(\sigma_i^{(\epsilon_i)})$ is a Nash equilibrium.

We consider an information-theoretic approach based on the Chernoff information [5] to measure the accuracy that can be achieved in hypothesis testing. For each individual i , let $D(\epsilon_i)$ denote the Chernoff information between the conditional distributions of X_i given $W = 1$ and $W = 0$. The larger $D(\epsilon_i)$ is, the more possible that the two hypotheses can be distinguished. In later discussions we will see that the Chernoff information is closely related to the best achievable probability of error.

The data collector aims to minimize the total payment while achieving an accuracy target. The design choices include the number of individuals N , the parameters $\epsilon_1, \epsilon_2, \dots, \epsilon_N$, and the payment mechanism $\mathbf{R}$ in which the strategy profile

$(\sigma_i^{(\epsilon_i)})$ is a Nash equilibrium. Then we formulate the mechanism design problem as the following optimization problem, which we call the *payment-accuracy problem*:

$$\begin{aligned} \min_{\substack{N \in \mathbb{N}, \epsilon_i \in \mathbb{R} \setminus \{0\} \cup \{\perp\}, \forall i \\ \mathbf{R} \in \mathcal{R}(\epsilon_1, \epsilon_2, \dots, \epsilon_N)}} & \sum_{i=1}^N \mathbb{E}_{(\sigma_i^{(\epsilon_i)})} [R_i(\mathbf{X})] \\ \text{subject to} & e^{-\sum_{i=1}^N D(\epsilon_i)} \leq \tau, \end{aligned}$$

where the accuracy target is represented by τ , which is related to the maximum allowable error. We focus on the range $\tau \in (0, 1)$ for nontriviality. Let $F(\tau)$ denote the optimal payment in this problem, i.e., the infimum of the total payment while satisfying the accuracy target τ .

5.2 Bounds on the Payment-Accuracy Problem

We present bounds on $F(\tau)$ in Theorem 3 below. For convenience, we define

$$\tilde{\epsilon} = \inf \left\{ \arg \max \left\{ \frac{D(\epsilon)}{V_{LB}(\epsilon)} : \epsilon > 0 \right\} \right\}, \quad \tilde{N} = \left\lceil \frac{\ln(1/\tau)}{D(\tilde{\epsilon})} \right\rceil, \quad (20)$$

where recall that $V_{LB}(\epsilon)$ is the lower bound in Theorem 1.

THEOREM 3. *The optimal payment $F(\tau)$ in the payment-accuracy problem for a given accuracy target $\tau \in (0, 1)$ is bounded as: $(\tilde{N} - 1)V_{LB}(\tilde{\epsilon}) \leq F(\tau) \leq \tilde{N}V_{LB}(\tilde{\epsilon}) + O(\tau \ln(1/\tau))$, where the $O(\cdot)$ is for $\tau \rightarrow 0$.*

The upper bound in Theorem 3 is given by the designed mechanism $\mathbf{R}^{(N, \epsilon)}$ with parameters chosen as $\epsilon = \tilde{\epsilon}$ and $N = \tilde{N}$. Note that $\tilde{\epsilon}$ can be proved to have a well-defined finite value independent of τ . By the lower and upper bounds on the value of privacy, the payment to each individual in $\mathbf{R}^{(\tilde{N}, \tilde{\epsilon})}$ is roughly equal to the lower bound $V_{LB}(\tilde{\epsilon})$. Then Theorem 3 indicates that the total payment of the designed mechanism $\mathbf{R}^{(\tilde{N}, \tilde{\epsilon})}$ is at most one individual's payment away from the minimum, with the diminishing term $O(\tau \ln(1/\tau))$ omitted. Figure 2 shows an illustration of the lower and upper bounds.

Theorem 3 is proved by Lemma 4 and Lemma 5 below, where the lower bound is given by the lower bound on the value of privacy, and the upper bound is given by $\mathbf{R}^{(\tilde{N}, \tilde{\epsilon})}$.

5.2.1 Lower Bound

First, notice that it suffices to limit the choice of each ϵ_i to $(0, +\infty)$ in the payment-accuracy problem, since when $\epsilon_i = \perp$, $D(\epsilon_i) = 0$, and when $\epsilon_i < 0$, $D(\epsilon_i) = D(|\epsilon_i|)$ and there exists another nonnegative payment mechanism with the same payment property and a Nash equilibrium at $(\sigma_i^{(|\epsilon_i|)})$ by Lemma 2.

Now we use the lower bound on the value of privacy to prove the lower bound on $F(\tau)$. By Theorem 1,

$$\inf_{\mathbf{R} \in \mathcal{R}(\epsilon_1, \epsilon_2, \dots, \epsilon_N)} \sum_{i=1}^N \mathbb{E}_{(\sigma_i^{(\epsilon_i)})} [R_i(\mathbf{X})] \geq \sum_{i=1}^N V_{LB}(\epsilon_i).$$

Therefore, the optimal payment $F(\tau)$ is lower bounded by the optimal value of the following optimization problem (P1):

$$\begin{aligned} \min_{N \in \mathbb{N}, \epsilon_i \in (0, +\infty), \forall i} & \sum_{i=1}^N V_{LB}(\epsilon_i) \\ \text{subject to} & e^{-\sum_{i=1}^N D(\epsilon_i)} \leq \tau. \end{aligned} \quad (\text{P1})$$

Figure 2: Illustration of the lower and upper bounds in Theorem 3 on the minimum total payment for achieving an accuracy target τ , where the upper bound is given by the designed mechanism $\mathbf{R}^{(\tilde{N}, \tilde{\epsilon})}$. In this example, the prior PMF of the state is $P_W(1) = 0.7$, $P_W(0) = 0.3$. The quality of signals is $\theta = 0.8$. The cost function is $g(\epsilon) = \epsilon$. The range of τ shown in the figure is 0.001–0.4.

LEMMA 4. Any feasible solution $(N, \epsilon_1, \epsilon_2, \dots, \epsilon_N)$ of (P1) satisfies

$$\sum_{i=1}^N V_{\text{LB}}(\epsilon_i) \geq (\tilde{N} - 1)V_{\text{LB}}(\tilde{\epsilon}),$$

where $\tilde{\epsilon}$ and $\tilde{N}$ are defined in (20).

Lemma 4 states that the total expected payment of the data collector is at least $(\tilde{N} - 1)V_{\text{LB}}(\tilde{\epsilon})$. Note that the value given by the genie-aided payment mechanism $\hat{\mathbf{R}}^{(\epsilon)}$ for $\tilde{N}$ individuals is $\tilde{N}V_{\text{LB}}(\tilde{\epsilon})$, which is at most one $V_{\text{LB}}(\tilde{\epsilon})$ away from the optimal value of (P1). We can think of $V_{\text{LB}}(\epsilon)$ as the price for ϵ units of privacy and $D(\epsilon)$ as the quality that the data collector gets from ϵ units of privacy due to its contribution to the accuracy. Then the intuition for $(\tilde{N}, \tilde{\epsilon}, \dots, \tilde{\epsilon})$ to be a near-optimal choice is that the privacy level $\tilde{\epsilon}$ gives the best quality/price ratio and $\tilde{N}$ is the fewest number of individuals to meet the accuracy target. The proof of Lemma 4 is presented in our technical report [31]. With this lemma, the lower bound on $F(\tau)$ in Theorem 3 is straightforward.

5.2.2 Upper Bound

LEMMA 5. Choose the parameters in the payment mechanism $\mathbf{R}^{(N, \epsilon)}$ defined in Section 4.2.1 to be $\epsilon = \tilde{\epsilon}$ and $N = \tilde{N}$, where $\tilde{\epsilon}$ and $\tilde{N}$ are defined in (20). Then in the Nash equilibrium $\sigma^{(\tilde{\epsilon})}$ of $\mathbf{R}^{(\tilde{N}, \tilde{\epsilon})}$, the accuracy target τ can be achieved, and the total expected payment is upper bounded as

$$\mathbb{E}_{\sigma^{(\tilde{\epsilon})}} \left[\sum_{i=1}^{\tilde{N}} R_i^{(\tilde{N}, \tilde{\epsilon})}(\mathbf{X}) \right] \leq \tilde{N}V_{\text{LB}}(\tilde{\epsilon}) + O(\tau \ln(1/\tau)).$$

This lemma follows from Theorem 2 and we omit the proof here. Since the payment mechanism $\mathbf{R}^{(N, \epsilon)}$ together with $\epsilon = \tilde{\epsilon}$ and $N = \tilde{N}$ is a feasible solution of the payment–accuracy problem, the upper bound in this lemma gives the upper bound on $F(\tau)$ in Theorem 3.

5.3 Discussions on the Accuracy Metric

When we study the relation between payment and accuracy, the accuracy can also be measured by the best achievable probability of error, defined as

$$p_e = \inf_{\psi} \mathbb{P}_{(\sigma_i^{(\epsilon_i)})}(\psi(\mathbf{X}) \neq W),$$

where $\psi(\mathbf{x})$ is a decision function, with $\psi(\mathbf{x}) = 0$ implying that H_0 is accepted and $\psi(\mathbf{x}) = 1$ implying that H_1 is accepted. However, p_e is difficult to deal with analytically since its exact form in terms of $\epsilon_1, \epsilon_2, \dots, \epsilon_N$ is intractable.

We measure the accuracy based on the Chernoff information, which is an information-theoretic metric closely related to p_e . It can be proved by the Bhattacharyya bound [17] that at the strategy profile $(\sigma_i^{(\epsilon_i)})$,

$$p_e \leq e^{-\sum_{i=1}^N D(\epsilon_i)}. \quad (21)$$

Therefore, if we want to guarantee that $p_e \leq p_e^{\max}$ for some maximum allowable probability of error $p_e^{\max}$, we can choose $\tau = p_e^{\max}$ in the payment–accuracy problem. In fact, the metric based on the Chernoff information is very close to the metric p_e , since the upper bound (21) is tight in exponent when all the ϵ_i are the same, i.e., when the reported data is i.i.d. given the hypothesis.

6. CONCLUSIONS

In this paper, we studied “the value of privacy” under a game-theoretic model, where a data collector pays strategic individuals to buy their private data for a learning purpose. The individuals do not consider the data collector to be trustworthy, and thus experience a cost of privacy loss during data reporting. The value of ϵ units of privacy is measured by the minimum payment of all nonnegative payment mechanisms under which an individual’s best response in a Nash equilibrium is to report the data with a privacy level of ϵ . We derived asymptotically tight lower and upper bounds on the value of privacy as the number of individuals becomes large, where the upper bound was given by a designed payment mechanism $\mathbf{R}^{(N, \epsilon)}$. We further applied these fundamental limits to find the minimum total payment for the data collector to achieve certain learning accuracy target, and derived lower and upper bounds on the minimum payment. The total payment of the designed mechanism $\mathbf{R}^{(N, \epsilon)}$ with properly chosen parameters is at most one individual’s payment away from the minimum. It would be of great interest to study the value of privacy under (ϵ, δ) -differential privacy, and to extend our results to more general models of private and reported data, e.g., models with larger alphabets for the state, the signals and the reported data.

7. ACKNOWLEDGEMENT

This work was supported in part by the NSF under Grant ECCS-1255425.

8. REFERENCES

- [1] ACEMOGLU, D., DAHLEH, M. A., LOBEL, I., AND OZDAGLAR, A. Bayesian learning in social networks. *Review of Econ. Stud.* 78, 4 (Oct. 2011), 1201–1236.
- [2] BASSILY, R., AND SMITH, A. Local, private, efficient protocols for succinct histograms. In *Proc. Ann. ACM Symp. Theory of Computing (STOC)* (Portland, OR, 2015), pp. 127–135.

- [3] CHEN, Y., CHONG, S., KASH, I. A., MORAN, T., AND VADHAN, S. Truthful mechanisms for agents that value privacy. In *Proc. ACM Conf. Electronic Commerce (EC)* (Philadelphia, PA, 2013), pp. 215–232.
- [4] CHEN, Y., SHEFFET, O., AND VADHAN, S. Privacy games. In *Int. Conf. Web and Internet Economics (WINE)* (2014), vol. 8877, pp. 371–385.
- [5] COVER, T. M., AND THOMAS, J. A. *Elements of Information Theory*, 2nd ed. John Wiley & Sons, Hoboken, NJ, 2006.
- [6] DUCHI, J. C., JORDAN, M. I., AND WAINWRIGHT, M. J. Local privacy and minimax bounds: Sharp rates for probability estimation. In *Advances Neural Information Processing Systems (NIPS)* (Lake Tahoe, NV, Dec. 2013), pp. 1529–1537.
- [7] DWORK, C. Differential privacy. In *Proc. Int. Conf. Automata, Languages and Programming (ICALP)* (Venice, Italy, 2006), pp. 1–12.
- [8] DWORK, C., MCSHERRY, F., NISSIM, K., AND SMITH, A. Calibrating noise to sensitivity in private data analysis. In *Proc. Conf. Theory of Cryptography (TCC)* (New York, NY, 2006), pp. 265–284.
- [9] DWORK, C., AND ROTH, A. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.* 9, 3–4 (Aug. 2014), 211–407.
- [10] ERLINGSSON, Ú., PIHUR, V., AND KOROLOVA, A. RAPPOR: Randomized aggregatable privacy-preserving ordinal response. In *Proc. ACM SIGSAC Conf. Computer and Communication Security (CCS)* (Scottsdale, AZ, 2014), pp. 1054–1067.
- [11] FANTI, G. C., PIHUR, V., AND ERLINGSSON, Ú. Building a RAPPOR with the unknown: Privacy-preserving learning of associations and data dictionaries. *arXiv:1503.01214 [cs.CR]* (2015).
- [12] FLEISCHER, L. K., AND LYU, Y. Approximately optimal auctions for selling privacy when costs are correlated with data. In *Proc. ACM Conf. Electronic Commerce (EC)* (Valencia, Spain, 2012), pp. 568–585.
- [13] GHOSH, A., AND LIGETT, K. Privacy and coordination: Computing on databases with endogenous participation. In *Proc. ACM Conf. Electronic Commerce (EC)* (Philadelphia, PA, 2013), pp. 543–560.
- [14] GHOSH, A., LIGETT, K., ROTH, A., AND SCHOENEBECK, G. Buying private data without verification. In *Proc. ACM Conf. Economics and Computation (EC)* (Palo Alto, CA, 2014), pp. 931–948.
- [15] GHOSH, A., AND ROTH, A. Selling privacy at auction. In *Proc. ACM Conf. Electronic Commerce (EC)* (San Jose, CA, 2011), pp. 199–208.
- [16] HSU, J., KHANNA, S., AND ROTH, A. Distributed private heavy hitters. In *Proc. Int. Conf. Automata, Languages and Programming (ICALP)* (Warwick, UK, 2012), pp. 461–472.
- [17] KAILATH, T. The divergence and Bhattacharyya distance measures in signal selection. *IEEE Trans. Commun. Technol.* 15, 1 (Feb. 1967), 52–60.
- [18] KAIROUZ, P., OH, S., AND VISWANATH, P. Extremal mechanisms for local differential privacy. In *Advances Neural Information Processing Systems (NIPS)* (Montreal, Canada, Dec. 2014), pp. 2879–2887.
- [19] KASIVISWANATHAN, S. P., LEE, H. K., NISSIM, K., RASKHODNIKOVA, S., AND SMITH, A. What can we learn privately? *SIAM J. Comput.* 40, 3 (May 2011), 793–826.
- [20] KROFT, S. The data brokers: selling your personal information. *CBS News* (Mar. 2014).
- [21] LE, T. N., SUBRAMANIAN, V. G., AND BERRY, R. A. The value of noise for informational cascades. In *Proc. IEEE Int. Symp. Information Theory (ISIT)* (Honolulu, HI, 2014), pp. 1101–1105.
- [22] LIGETT, K., AND ROTH, A. Take it or leave it: Running a survey when privacy comes at a cost. In *Proc. Int. Workshop Internet and Network Economics (WINE)* (Liverpool, UK, 2012), pp. 378–391.
- [23] MILLER, N., RESNICK, P., AND ZECKHAUSER, R. Eliciting informative feedback: The peer-prediction method. In *Computing with Social Trust*, Human–Computer Interaction Series. Springer London, 2009, pp. 185–212.
- [24] NISSIM, K., VADHAN, S., AND XIAO, D. Redrawing the boundaries on purchasing data from privacy-sensitive individuals. In *Proc. Conf. Innovations in Theoretical Computer Science (ITCS)* (Princeton, NJ, 2014), pp. 411–422.
- [25] PAI, M. M., AND ROTH, A. Privacy and mechanism design. *SIGecom Exch.* 12, 1 (June 2013), 8–29.
- [26] ROTH, A., AND SCHOENEBECK, G. Conducting truthful surveys, cheaply. In *Proc. ACM Conf. Electronic Commerce (EC)* (Valencia, Spain, 2012), pp. 826–843.
- [27] SHOKRI, R. Privacy games: Optimal user-centric data obfuscation. In *Proc. Privacy Enhancing Technologies (PETS)* (Philadelphia, PA, 2015), pp. 299–315.
- [28] SRIKANT, R., AND YING, L. *Communication Networks: An Optimization, Control and Stochastic Networks Perspective*. Cambridge Univ. Press, New York, 2014.
- [29] WANG, W., YING, L., AND ZHANG, J. On the relation between identifiability, differential privacy, and mutual-information privacy. In *Proc. Ann. Allerton Conf. Communication, Control and Computing* (Monticello, IL, Sept. 2014), pp. 1086–1092.
- [30] WANG, W., YING, L., AND ZHANG, J. A minimax distortion view of differentially private query release. In *Proc. Asilomar Conf. Signals, Systems, and Computers* (Pacific Grove, CA, Nov. 2015).
- [31] WANG, W., YING, L., AND ZHANG, J. The value of privacy: Strategic data subjects, incentive mechanisms and fundamental limits. *arXiv:1603.06870 [cs.GT]* (2016).
- [32] WARNER, S. L. Randomized response: A survey technique for eliminating evasive answer bias. *J. Amer. Stat. Assoc.* 60, 309 (Mar. 1965), 63–69.
- [33] XIAO, D. Is privacy compatible with truthfulness? In *Proc. Conf. Innovations in Theoretical Computer Science (ITCS)* (Berkeley, CA, 2013), pp. 67–86.