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

WEINA WANG, LEI YING, and JUNSHAN ZHANG, Arizona State University

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. One primary goal of the data collector is to learn some desired information from the elicited data. Specifically, this information is modeled by an underlying state, and the private data of each individual represents his of her knowledge about the state. Departing from most of the existing work on privacy-aware surveys, our model does not assume the data collector to be trustworthy. Further, an individual takes full control of his or her own data privacy and reports only a privacy-preserving version of his or her data.

In this article, the value of ϵ units of privacy is measured by the minimum payment among all nonnegative payment mechanisms, under which an individual's best response at a Nash equilibrium is to report his or her data in an ϵ-locally differentially private manner. The higher ϵ is, the less private the reported data is. We derive lower and upper bounds on the value of privacy that are asymptotically tight as the number of data subjects becomes large. Specifically, the lower bound assures that it is impossible to use a lower payment to buy ϵ units of privacy, and the upper bound is given by an achievable payment mechanism that we design. 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 accuracy target for learning the underlying state and show that the total payment of the designed mechanism is at most one individual's payment away from the minimum.

CCS Concepts: • **Security and privacy → Economics of security and privacy**; • **Theory of computation → Algorithmic game theory and mechanism design**;

Additional Key Words and Phrases: Data collection, differential privacy, randomized response

ACM Reference format:
Weina Wang, Lei Ying, and Junshan Zhang. 2018. The Value of Privacy: Strategic Data Subjects, Incentive Mechanisms, and Fundamental Limits. *ACM Trans. Econ. Comput.* 6, 2, Article 8 (August 2018), 26 pages.
https://doi.org/10.1145/3232863

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 [21], 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, and so on. However, in the wake of a number of recent scandals, such as the Netflix data breach and the Veterans Affairs data theft, data privacy is emerging as one of the most

This work was supported in part by the National Science Foundation under grants ECCS-1255425 and SaTC-1618768.
Authors' addresses: W. Wang, L. Ying, and J. Zhang, Goldwater Center, School of ECEE, Arizona State University, Tempe, AZ, 85281; emails: {weina.wang, lei.ying.2, junshan.zhang}@asu.edu.

https://doi.org/10.1145/3232863

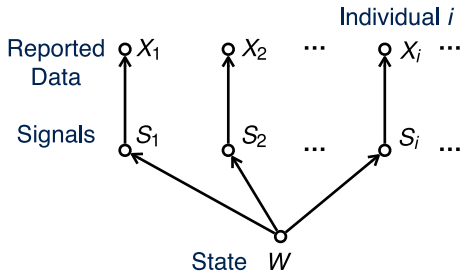

Fig. 1. Information structure of the model: The data collector is interested in the state W, which is a binary random variable. Each individual i possesses his or her private data, which is a binary signal S_i. Conditioned on W, $S_1, S_2, \ldots, S_N$ are i.i.d. Individual i's reported data is X_i,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

serious concerns of big data analytics. One common practice of collecting private data is called *informed consent*. With information on "who is collecting the data, what data are collected, and how the data will be used," data subjects decide 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 a data breach occurs. In some applications, such as collecting history records 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 where data subjects (individuals) are able to control their own data privacy by reporting perturbed data to the data collector. In particular, the data collector will use an incentive mechanism to pay (or reward) individuals for reporting informative data, and individuals report noisy data with the 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 [12–15, 22, 24, 26], 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 the 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 by an individual versus data quality concerned by the data collector.

Specifically, we consider a game-theoretic model of collecting private data for 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 model by a binary random variable W, which is called the *state*. Each individual i possesses a binary *signal* S_i, which is his or her private data, representing his or 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 his or 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 how to pay the individuals to get informative reports, i.e., not completely random data. Intuitively, the data collector should offer higher payments to purchase more informative data (which incurs more privacy leakage for individuals). More precisely, the payment mechanism should be carefully designed to provide right incentives for informative data reporting. We will answer the following fundamental questions in this article: *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?*

When reporting data to the data collector, a privacy-aware individual weighs 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. Note that only compensating the privacy cost incurred is not sufficient. The payment mechanism needs to ensure that ϵ is the best privacy level: If an individual uses a less-private strategy, then the decrease in his or her payment is larger than the decrease in her privacy cost; however, if an individual uses a more-private strategy, then the increase in her payment is smaller than the increase in her privacy cost. In other words, we focus on *Nash equilibria* (strategy profiles where individuals' strategies are best responses to others). We also note that Nash equilibrium is a more applicable solution concept than Bayesian Nash equilibrium for our model, due to the choice of privacy measure. This will become clear after we introduce the detailed model in Section 3. To quantify the monetary value of data privacy in a market for private data, we study the minimum payment that makes an equilibrium strategy have a privacy level of ϵ.

We remark that the problem of eliciting non-verifiable data has been studied in the peer prediction literature (see, e.g., the seminal article [23]). It is worth noting that individuals have no privacy concerns in a classical setting of peer prediction. Therefore, the relation between privacy and payment, which is the main focus of the current article, has not been addressed by the peer prediction literature. Nevertheless, this study has leveraged ideas from peer prediction to reward an individual when her reported data is less perturbed. We emphasize that this study takes into account individuals' behavior of striking the right balance between privacy loss and payments. This tradeoff makes truthfulness no longer a focal design goal, since truthful data reporting incurs a high privacy loss that is expensive to compensate. Instead, the data collector seeks to incentivize informative data reporting and achieve her learning goal using cost-effective mechanisms.

Summary of Main Results

We assume that individuals use the celebrated notion of differential privacy [7, 8] in the local model [5, 6, 19] 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 are 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 $V(\epsilon)$ through the following three steps. First, we prove that from a payment perspective, it suffices to focus on mechanisms under which an individual's equilibrium strategy is the following randomized response with a privacy level

of ϵ: The reported data are generated by flipping the signal with probability $\frac{1}{e^{\epsilon}+1}$. For convenience, we refer to this strategy as 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. Then the lower bound is given by necessary conditions for ϵ to be the best privacy level in a genie-aided mechanism. We remark that although the genie-aided mechanism that achieves this lower bound is not implementable, it can be well approximated by the feasible payment mechanism we design when the number of individuals is large.

(2) We observe that the strategy of an individual in a Nash equilibrium exhibits the following interesting characteristics: It is either a symmetric randomized response, where by symmetry we mean this strategy treats the two possible realizations of the private signal symmetrically, or a non-informative strategy, where the reported data are independent of the signal. This characterization holds regardless of the prior distribution over the state, and it also holds for more general probability models of binary 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 $\boldsymbol{R}^{(N,\epsilon)}$, 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 problem*, 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 $\boldsymbol{R}^{(N,\epsilon)}$ 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.

(5) We also give a more in-depth analysis of the Nash equilibria of the designed mechanism $\boldsymbol{R}^{(N,\epsilon)}$. Besides the desired Nash equilibrium where every individual uses the ϵ-strategy, $\boldsymbol{R}^{(N,\epsilon)}$ also has other equilibria. When implementing this mechanism, the data collector should prime the individuals to use the ϵ-strategy, but the presence of multiple equilibria still puts obstacles to obtaining the desired outcome. We first show that any Nash equilibrium of $\boldsymbol{R}^{(N,\epsilon)}$ is homogeneous, i.e., all the individuals use the same strategy in a Nash equilibrium. By our characterization of equilibrium strategies, this strategy is either an informative symmetric randomized response or a non-informative strategy. For informative equilibrium strategies that report truthfully with a probability greater than $1/2$, we prove that any sequence of privacy levels of such equilibria converges to the desired level ϵ as the number of individuals N becomes large. Then with a large population, even if the designed mechanism $\boldsymbol{R}^{(N,\epsilon)}$ may have such equilibria with privacy levels different from the desired level ϵ, then these levels are close to ϵ. However, this convergence of privacy levels does not fully address the issues raised by the multiplicity of equilibria, since other equilibria, including the non-informative equilibria, still exist. We believe that it is possible

to leverage results that address the multiplicity in the peer prediction literature (e.g., Reference [20]) to mitigate this problem. Eliminating undesired equilibria for the designed mechanism would be an interesting direction for future work.

2 RELATED WORK

Most existing work on privacy-aware surveys [12–15, 22, 24, 26] assumes that there is a trusted data curator or data collector. The private data are either already kept by the data collectoror can be elicited truthfully. What the data collector purchases is the "right" of using individuals' data in an announced way. Our work differs from these articles by considering a data collector who is not trusted by individuals. In this scenario, individuals report noisy data for privacy protection.

In the seminal work by Ghosh and Roth [15], individuals' data are 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 his or 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, 24, 26] 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 are not available yet and need to be reported by the individuals to the data collector, but the data collector is still trusted [2, 13, 14, 32]. Notably, Ghosh et al. [14] study the model in which the collected data are 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 [31] and is formalized in Reference [19], has been studied in the literature [1, 5–9, 16, 18, 27, 29, 30]. The hypothesis testing formulation in our article is similar to a setting in Reference [18], where the authors find an optimal mechanism that maximizes the power to discriminate between data generated from different hypotheses subject to local differential privacy constraints. In practice, Google's Chrome web browser has implemented the RAPPOR mechanism [10, 11] to collect users' data, which guarantees that only limited privacy is 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.

Aside from privacy concerns, the problem we study in this article is closely related to the peer prediction mechanism proposed by Miller et al. [23]. In peer prediction, individuals do not have privacy concerns, and their utility only comes from the payments they receive. The peer prediction mechanism pays an individual based on how well his or her reported data predicts another randomly selected individual's data. It uses strictly proper scoring rules to evaluate the prediction, under which an individual's best way of predicting another individual's data is to report her own data truthfully. However, when an individual have privacy concerns, his or her best reporting strategy is not to obtain the highest possible payment, since his or her utility also includes a cost due to privacy loss. He or she will weigh this privacy cost against the payment to decide how to perturb his or her data. This tradeoff between privacy and payment is critical for the problem in this article and has not been studied in the peer prediction literature. The mechanism we design has a flavor of peer prediction in the sense that the payment to an individual also depends on how well his or her reported data predict others' reported data. But more importantly, we characterize how individuals behave when trading privacy for money, where the reported data are noisy rather than truthful, and then quantify the minimum payment needed to induce a desired privacy level.

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 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 $\mathcal{N} = \{1, 2, \ldots, N\}$. Denote all individuals other than some given individual i by "$-i$." Each individual i possesses a binary signal S_i, which is his or her private data, reflecting his or her knowledge about the state W. For example, S_i can represent individual i's opinion towards the new technology. Let $\boldsymbol{S} = (S_1, S_2, \ldots, S_N)$. Conditional on the state W, the signals $S_1, S_2, \ldots, S_N$ are i.i.d. with the following conditional distributions:

$$\mathbb{P}(S_i = 1 \mid W = 1) = \theta, \quad \mathbb{P}(S_i = 0 \mid W = 1) = 1 - \theta,$$
$$\mathbb{P}(S_i = 0 \mid W = 0) = \theta, \quad \mathbb{P}(S_i = 1 \mid W = 0) = 1 - \theta,$$

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 $\boldsymbol{X} = (X_1, X_2, \ldots, X_N)$. 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 $\sigma_i : \{0, 1\} \to \mathcal{D}(\mathcal{X})$, where $\mathcal{D}(\mathcal{X})$ is the set of probability distributions on $\mathcal{X}$. Let $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \ldots, \sigma_N)$. 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 $\mathbb{P}_{\sigma_i}(X_i = \perp \mid S_i = 1) = \mathbb{P}_{\sigma_i}(X_i = \perp \mid S_i = 0) = 0$, then we say σ_i is a *symmetric randomized response*. If a strategy σ_i makes X_i and S_i independent, then we say σ_i is *non-informative*; otherwise we say σ_i is *informative*.

Mechanism. The data collector uses a payment mechanism defined below to determine the amount of payment to each individual.

Definition 3.1. A payment mechanism $\boldsymbol{R} : \mathcal{X}^N \to \mathbb{R}^N$ is a function that maps individuals' reported data $\boldsymbol{X}$ to a vector $\boldsymbol{R}(\boldsymbol{X})$, whose ith entry $R_i(\boldsymbol{X})$ is the amount of payment to individual i.

We are interested in payment mechanisms in which the payment to each individual is nonnegative, i.e., $R_i(\boldsymbol{x}) \geq 0$ for any individual i and any $\boldsymbol{x} \in \mathcal{X}^N$, 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 [5–9, 19] of the strategy, defined as follows.

Definition 3.2. 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 indistinguishablity 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 Reference [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 studies 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:

$$g(\xi) \geq 0, \quad \forall \xi \geq 0, \tag{1}$$

$$g(0) = 0, \tag{2}$$

$$g \text{ is non-decreasing}, \tag{3}$$

where Equation (1) follows from the fact that a privacy cost is nonnegative, Equation (2) indicates that the privacy cost is 0 when the reported data is independent of the private data, and Equation (3) means that the privacy cost will not decrease when the privacy loss becomes larger. In this article, 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 his or her payment and his or 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 $\boldsymbol{R}$, and the cost function g are common knowledge.

We remark that there is no need to formulate this game as a Bayesian game for the following reasons. Each individual's privacy cost depends on not only how he or she reports data given the current realization of his or her signal, but also how he or she would report data given other possible realizations of his or her signal. So each individual's privacy cost is a function of the whole strategy regardless of the current realization of his or her signal. The form of this function is commonly known. The form of the payment mechanism is also commonly known. Therefore, the utility function of every individual is common knowledge in this game, and thus we do not need to formulate it as a Bayesian game.

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

Definition 3.3. A strategy profile $\boldsymbol{\sigma}$ is a Nash equilibrium in a payment mechanism $\boldsymbol{R}$ if for any individual i and any strategy σ_i',

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

where the expectation is over the reported data $\boldsymbol{X}$, and the subscripts $\boldsymbol{\sigma}$ and $(\sigma_i', \boldsymbol{\sigma}_{-i})$, indicate that $\boldsymbol{X}$ is generated by the strategy profile $\boldsymbol{\sigma}$ and $(\sigma_i', \boldsymbol{\sigma}_{-i})$, respectively.

4 THE VALUE OF DATA PRIVACY

We consider nonnegative payment mechanism under which an individual i reports data with a privacy level of ϵ in a Nash equilibrium. We denote this set of mechanisms by $\mathcal{R}(i;\epsilon)$. Then we measure the value of ϵ units of privacy by the minimum payment to individual i among all mechanisms in $\mathcal{R}(i;\epsilon)$. This measure does not depend on the specific identity of i due to the symmetry across individuals. Note that a mechanism in $\mathcal{R}(i;\epsilon)$ may have multiple Nash equilibria. But since we are interested in the minimum payment possible, we include a mechanism in $\mathcal{R}(i;\epsilon)$ as long as one of its Nash equilibria satisfies that individual i reports data with privacy level ϵ. For any mechanism $\boldsymbol{R} \in \mathcal{R}(i;\epsilon)$, let $\boldsymbol{\sigma}^{(\boldsymbol{R};\epsilon)}$ denote the corresponding Nash equilibrium. Then, formally, the value of ϵ units of privacy is measured by

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

Recall that we are interested in the regime $\epsilon > 0$, since the data collector wants the reported data to be useful for data analysis.

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 4.1 below. For convenience, we define

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

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

Theorem 4.1. *The value of ϵ units of privacy measured in Equation (4) for any $\epsilon > 0$ is lower bounded as $V(\epsilon) \geq V_{\mathrm{LB}}(\epsilon)$. Specifically, for any nonnegative payment mechanism $\boldsymbol{R}$, 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 $V_{\mathrm{LB}}(\epsilon)$.*

We remark that the lower bound in Theorem 4.1 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. Note that such a genie-aided mechanism is not an actual mechanism in the sense of Definition 3.1, and thus it does not belong to the set $\mathcal{R}(i;\epsilon)$. 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 $\widehat{\boldsymbol{R}}^{(\epsilon)}$, determines the payment to each individual i based on his or her own reported data X_i and the state W as follows:

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

where

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

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 symmetric randomized response strategy as the *ϵ-strategy*. The expected payment to individual i at this strategy equals to the lower bound in Theorem 4.1.

Next we sketch the proof of Theorem 4.1. 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 are presented in Appendix A–C.

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 Reference [3] in a different setting.

Lemma 4.2. *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 4.2 holds for more general probability models of binary 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 4.2, if an individual's strategy in a Nash equilibrium has a privacy level of ϵ, where $\epsilon > 0$, then 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 4.3. *For any nonnegative payment mechanism $\boldsymbol{R}$ in which the strategy profile $(\sigma_i^{(-\epsilon)}, \boldsymbol{\sigma}_{-i})$ with some $\epsilon > 0$ is a Nash equilibrium, there exists another nonnegative payment mechanism $\boldsymbol{R}'$ in which $(\sigma_i^{(\epsilon)}, \boldsymbol{\sigma}_{-i})$ 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 $\boldsymbol{R}'$ that is constructed by applying $\boldsymbol{R}$ on the reported data after modifying X_i to $1 - X_i$.

4.1.2 Genie-Aided Payment Mechanism. A genie-aided payment mechanism $\widehat{\boldsymbol{R}} : \mathcal{X}^N \times \{0,1\} \to \mathbb{R}^N$ determines the payment to an individual based on not only the reported data $\boldsymbol{X}$ 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 he or she were aided by a genie. We consider nonnegative genie-aided payment mechanisms, where $\widehat{R}_i(\boldsymbol{X}, W)$, the payment to individual i, depends on only his or her own reported data X_i and the underlying state W. We write $\widehat{R}_i(X_i, W)$ to represent $\widehat{R}_i(\boldsymbol{X}, W)$ for conciseness. Therefore, for each individual i, a genie-aided mechanism makes use of the information of W but discards the information in $\boldsymbol{X}_{-i}$. 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 4.4. *For any nonnegative payment mechanism $\boldsymbol{R}$ and any Nash equilibrium $\boldsymbol{\sigma}$ of it, there exists a nonnegative genie-aided mechanism $\widehat{\boldsymbol{R}}$, such that $\boldsymbol{\sigma}$ is also a Nash equilibrium of $\widehat{\boldsymbol{R}}$ and the expected payment to each individual at this equilibrium is the same under $\boldsymbol{R}$ and $\widehat{\boldsymbol{R}}$.*

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

$$\widehat{R}_i(x_i, w) = \overline{R}_i(x_i; w) \coloneqq \mathbb{E}_{\boldsymbol{\sigma}}[R_i(\boldsymbol{X}) \mid X_i = x_i, W = w].$$

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 $\boldsymbol{X}_{-i}$ of others. The above genie-aided payment mechanism $\widehat{\boldsymbol{R}}$ is constructed such that no matter what strategy individual i uses, his or her expected utility is the same under $\boldsymbol{R}$ and $\widehat{\boldsymbol{R}}$. Since an individual calculates his or her best response according to the expected utility, his or her equilibrium behavior and expected payment are the same under $\widehat{\boldsymbol{R}}$ and $\boldsymbol{R}$. 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 $\widehat{\boldsymbol{\mathcal{R}}}(i;\epsilon)$ 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

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

which is a definition similar to the value of ϵ units of privacy, $V(\epsilon)$, measured in Equation (4). Then $\widehat{V}(\epsilon) \le V(\epsilon)$ for the following reasons. Consider any $\boldsymbol{R} \in \boldsymbol{\mathcal{R}}(i;\epsilon)$, i.e., any nonnegative payment mechanism $\boldsymbol{R}$ in which individual i's strategy in a Nash equilibrium has a privacy level of ϵ. With Lemma 4.2 and 4.3, we can assume without loss of generality that this equilibrium strategy is the ϵ-strategy. Then by Lemma 4.4, we can map $\boldsymbol{R}$ to a $\widehat{\boldsymbol{R}} \in \widehat{\boldsymbol{\mathcal{R}}}(i;\epsilon)$, such that

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

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

4.1.3 Proof of Theorem 4.1. With Lemmas 4.2, 4.3, and 4.4, we can prove the lower bound in Theorem 4.1 by focusing on the genie-aided mechanisms in $\widehat{\boldsymbol{\mathcal{R}}}(i;\epsilon)$. 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 4.1.

OF THEOREM 4.1. By Lemmas 4.2, 4.3, and 4.4, 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 $\widehat{\boldsymbol{\mathcal{R}}}(i;\epsilon)$. Consider any $\widehat{\boldsymbol{R}} \in \widehat{\boldsymbol{\mathcal{R}}}(i;\epsilon)$ 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

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

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

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

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

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

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

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

Therefore,

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

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

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

which proves the lower bound. □

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 Equation (8) to hold and Equation (7) to be satisfied, which is equivalent to the following conditions:

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

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

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

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

$$\overline{R}_i(1;0) = 0, \tag{12}$$

$$\overline{R}_i(0;1) = 0, \tag{13}$$

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

where recall that $\overline{R}_i(x_i;w) = \mathbb{E}_{\boldsymbol{\sigma}^{(\boldsymbol{R},\epsilon)}}[R_i(\boldsymbol{X}) \mid X_i = x_i, W = w]$ for $x_i, w \in \{0,1\}$. It can be proved that if $\boldsymbol{R}$ satisfies Equations (12) and (13), then $R_i(\boldsymbol{x}) = 0$ for any $\boldsymbol{x} \in \mathcal{X}^N$, which contradicts Equation (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 $\widehat{\boldsymbol{R}}^{(\epsilon)}$: 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 4.5 below. For convenience, we define

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

where θ is the quality of signal. Note that $d > 0$ for any $\epsilon > 0$. Recall that $V_{\mathrm{LB}}(\epsilon)$ is the lower bound in Theorem 4.1.

Theorem 4.5. *The value of ϵ units of privacy measured in Reference [4] is upper bounded as $V(\epsilon) \le V_{\mathrm{LB}}(\epsilon) + O(e^{-Nd})$, where the $O(\cdot)$ is for $N \to \infty$. Specifically, there exists a nonnegative payment mechanism $\boldsymbol{R}^{(N,\epsilon)}$ in which the strategy profile $\boldsymbol{\sigma}^{(\epsilon)}$ consisting of ϵ-strategies is a Nash equilibrium, and the expected payment to each individual i at this equilibrium is upper bounded by $V_{\mathrm{LB}}(\epsilon) + O(e^{-Nd})$.*

Comparing this upper bound with the lower bound $V_{\mathrm{LB}}(\epsilon)$ in Theorem 4.1, 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 $\boldsymbol{R}^{(N,\epsilon)}$ in Section 4.2.1. We will show that under $\boldsymbol{R}^{(N,\epsilon)}$, the strategy profile $\boldsymbol{\sigma}^{(\epsilon)}$ consisting of ϵ-strategies is a Nash equilibrium. Therefore, $\boldsymbol{R}^{(N,\epsilon)}$ is a member of $\mathcal{R}(i;\epsilon)$, and the payment to individual i at $\boldsymbol{\sigma}^{(\epsilon)}$ gives an upper bound on the value of privacy.

The design of $\boldsymbol{R}^{(N,\epsilon)}$ is enlightened by the hypothetical payment mechanism $\widehat{\boldsymbol{R}}^{(\epsilon)}$ defined in Equation (6). But without direct access to the state W, the mechanism $\boldsymbol{R}^{(N,\epsilon)}$ 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 his or 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 $\widehat{\boldsymbol{R}}^{(\epsilon)}$ to ensure that the ϵ-strategy is still a best response of each individual in $\boldsymbol{R}^{(N,\epsilon)}$, given that other individuals also follow the ϵ-strategy, which yields the desired Nash equilibrium $\boldsymbol{\sigma}^{(\epsilon)}$.

The equilibrium payment to each individual in $\boldsymbol{R}^{(N,\epsilon)}$ converges to the lower bound in Theorem 4.1 as the number of individuals N goes to infinity. The intuition behind this 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 $\boldsymbol{R}^{(N,\epsilon)}$ works similarly to the genie-aided mechanism $\widehat{\boldsymbol{R}}^{(\epsilon)}$, whose equilibrium payment to each individual equals to the lower bound in Theorem 4.1.

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

(1) Each individual reports his or 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, then pay zero to this participant. Otherwise, for each participating individual i, the data collector computes the variable

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

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

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

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

$$\alpha = \theta \frac{e^\epsilon}{e^\epsilon + 1} + (1-\theta)\frac{1}{e^\epsilon + 1}. \tag{16}$$

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

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

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

$$\gamma^{(n)} = \begin{cases} 1 - \binom{n-1}{\frac{n-1}{2}} \alpha^{\frac{n-1}{2}} (1-\alpha)^{\frac{n-1}{2}} & \text{if } n-1 \text{ is even,} \\ 1 & \text{if } n-1 \text{ is odd.} \end{cases} \tag{17}$$

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

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

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

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

It is easy to verify that these parameters are nonnegative. Thus $\boldsymbol{R}^{(N,\epsilon)}$ is a nonnegative payment mechanism. The proof of the equilibrium properties of $\boldsymbol{R}^{(N,\epsilon)}$ in Theorem 4.5 is given below.

4.2.3 *Proof of Theorem 4.5.*

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

We first prove that the strategy profile $\boldsymbol{\sigma}$ is a Nash equilibrium in $\boldsymbol{R}$; i.e., for any individual i, the ϵ-strategy is a best response of individual i when other individuals follow $\boldsymbol{\sigma}_{-i}$. Following the notation in the proof of Lemma 4.2, for any individual i we consider any strategy σ_i' of individual i and let

$$p_1 = \mathbb{P}_{\sigma_i'}(X_i = 1 \mid S_i = 1), \quad q_1 = \mathbb{P}_{\sigma_i'}(X_i = 0 \mid S_i = 1),$$
$$p_0 = \mathbb{P}_{\sigma_i'}(X_i = 1 \mid S_i = 0), \quad q_0 = \mathbb{P}_{\sigma_i'}(X_i = 0 \mid S_i = 0).$$

Then, by the proof of Lemma 4.2, 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 $\boldsymbol{\sigma}_{-i}$. Thus, we derive the form of $U_i(p_1, p_0, q_1, q_0)$ next. Recall that we let $\overline{R}_i(x_i; w)$ denote $\mathbb{E}_{(\sigma_i', \boldsymbol{\sigma}_{-i})}[R_i(\boldsymbol{X}) \mid X_i = x_i, W = w]$ for $x_i, w \in \{0, 1\}$. Then

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

with

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

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

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

and it can be verified that K_1, K_0, L_1, and L_0 are all positive. Therefore, by the proof of Lemma 4.2, 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:

$$p_1 = p_0 = q_1 = q_0 = 0, \tag{18}$$

$$p_1 = p_0, \quad q_1 = q_0, \quad p_1 + q_1 = 1, \tag{19}$$

$$p_1 = q_0, \quad p_0 = q_1, \quad p_1 + q_1 = 1. \tag{20}$$

The strategy in Equation (18) is to always not participate, which yields a utility of zero. For strategies in the form of Equation (19) or Equation (20), we can write the expected utility as a function of p_1 and p_0 as follows:

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

where $\overline{K}_1 = K_1 - L_1$, $\overline{K}_0 = K_0 - L_0$, $\overline{K} = L_1 + L_0$, and with a little abuse of notation, $\zeta(p_1, p_0) = \max\{|\ln \frac{p_1}{p_0}|, |\ln \frac{1-p_1}{1-p_0}|\}$. Inserting the value of $R_i(X_i; M_{-i})$ gives

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

Then a strategy in the form of Equation (19) yields a utility of $\overline{K} > 0$. A strategy in the form of Equation (20) 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:

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

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

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

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

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

$$f''(\xi) = -(\overline{K}_1 - \overline{K}_0) \frac{e^\xi (e^\xi - 1)}{(e^\xi + 1)^3} - g''(\xi) \leq 0,$$

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

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

By the convexity of g, $g(\epsilon) \leq g'(\epsilon)\epsilon \leq g'(\epsilon)\frac{e^\epsilon - e^{-\epsilon}}{2}$. Thus $f(\epsilon) \geq \overline{K}$, 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 $\boldsymbol{\sigma}$, which can be written as

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

By definition,

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

Then

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

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

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

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

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

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

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

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

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

By the monotonicity of h,

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

Notice that

$$1-\gamma^{(N)} = \begin{cases} \binom{N-1}{\frac{N-1}{2}}\alpha^{\frac{N-1}{2}}(1-\alpha)^{\frac{N-1}{2}} & \text{if } N-1 \text{ is even,} \\ 0 & \text{if } N-1 \text{ is odd.} \end{cases}$$

Then when $N-1$ is odd, $\gamma^{(N)} = 1$, and when $N-1$ is even,

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

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

Therefore,

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

as $N\to\infty$, which completes the proof. □

4.3 Extensions to Heterogeneous Cost Functions and Heterogeneous Privacy Levels

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 4.1 and 4.5 except for the heterogeneous cost function $g_i(\epsilon)$, hold,

$$g_i'(\epsilon)\frac{e^\epsilon+1}{e^\epsilon}\left(\frac{\theta}{2\theta-1}(e^\epsilon+1)-1\right) \le V_i(\epsilon) \le g_i'(\epsilon)\frac{e^\epsilon+1}{e^\epsilon}\left(\frac{\theta}{2\theta-1}(e^\epsilon+1)-1\right)+O(e^{-Nd}).$$

The lower bound above can be derived directly from the proof of Theorem 4.1, 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 similarly to $\boldsymbol{R}^{(N,\epsilon)}$, with the g' in $R_i^{(N,\epsilon)}$ replaced by g_i'. In this mechanism, the strategy profile $\boldsymbol{\sigma}^{(\epsilon)}$ is still a Nash equilibrium, and the expected payment to individual i at this equilibrium can still be upper bounded as in Theorem 4.5 but again with g' replaced by g_i'.

We are also able to induce data reporting with different privacy levels from different individualsby modifying the mechanism $R^{(N,\epsilon)}$ properly. Suppose that we want a Nash equilibrium where individual i uses the ϵ_i-strategy. Then we modify the payments as follows. Recall that in the mechanism $R^{(N,\epsilon)}$, the payment to individual i is $\frac{g'(\epsilon)(e^\epsilon+1)^2}{2e^\epsilon}A_{X_i,M_{-i}}$. We first replace the ϵ in the term $\frac{g'(\epsilon)(e^\epsilon+1)^2}{2e^\epsilon}$ with ϵ_i. We then modify the term $A_{X_i,M_{-i}}$. Recall that to calculate $A_{X_i,M_{-i}}$, we have defined $\beta^{(n)}$ and $\gamma^{(n)}$. Here we still define

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

where M_{-i} is the majority of the other $n-1$ participants' reported data. Note that the probability is calculated assuming that each participant uses his or her own ϵ_i. Using arguments similar to those used in the proof of Theorem 4.5, we can show that this modified mechanism has a Nash equilibrium where individual i uses the ϵ_i-strategy.

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 $X_1, X_2, \ldots, X_N$, 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 4.2, we can write an equilibrium strategy profile in the form of $(\sigma_i^{(\epsilon_i)}) = (\sigma_1^{(\epsilon_1)}, \sigma_2^{(\epsilon_2)}, \ldots, \sigma_N^{(\epsilon_N)})$ with $\epsilon_i \in \mathbb{R}\setminus\{0\}\cup\{\perp\!\!\!\perp\}$, 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 $\epsilon = \perp\!\!\!\perp$. Let $\boldsymbol{\mathcal{R}}(\epsilon_1, \epsilon_2, \ldots, \epsilon_N)$ 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 [4] 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 $\epsilon_1, \epsilon_2, \ldots, \epsilon_N$, and the payment mechanism $\boldsymbol{R}$ 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*:

$$\min_{\substack{N\in\mathbb{N},\ \epsilon_i\in\mathbb{R}\setminus\{0\}\cup\{\perp\!\!\!\perp\},\forall i \\ \boldsymbol{R}\in\boldsymbol{\mathcal{R}}(\epsilon_1,\epsilon_2,\ldots,\epsilon_N)}} \quad \sum_{i=1}^{N} \mathbb{E}_{(\sigma_i^{(\epsilon_i)})}[R_i(\boldsymbol{X})]$$
$$\text{subject to} \quad e^{-\sum_{i=1}^{N} D(\epsilon_i)} \le \tau,$$

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 5.1 below. For convenience, we define

$$\widetilde{\epsilon} = \inf\left\{\arg\max\left\{\frac{D(\epsilon)}{V_{\mathrm{LB}}(\epsilon)} : \epsilon > 0\right\}\right\}, \quad \widetilde{N} = \left\lceil \frac{\ln(1/\tau)}{D(\widetilde{\epsilon})} \right\rceil, \tag{21}$$

where recall that $V_{\mathrm{LB}}(\epsilon)$ is the lower bound in Theorem 4.1.

Theorem 5.1. *The optimal payment $F(\tau)$ in the payment–accuracy problem for a given accuracy target $\tau \in (0, 1)$ is bounded as: $(\widetilde{N} - 1)V_{\mathrm{LB}}(\widetilde{\epsilon}) \le F(\tau) \le \widetilde{N}V_{\mathrm{LB}}(\widetilde{\epsilon}) + O(\tau \ln(1/\tau))$, where the $O(\cdot)$ is for $\tau \to 0$.*

The upper bound in Theorem 5.1 is given by the designed mechanism $\boldsymbol{R}^{(N,\epsilon)}$ with parameters chosen as $\epsilon = \widetilde{\epsilon}$ and $N = \widetilde{N}$. Note that $\widetilde{\epsilon}$ 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 $\boldsymbol{R}^{(\widetilde{N},\widetilde{\epsilon})}$ is roughly equal to the lower bound $V_{\mathrm{LB}}(\widetilde{\epsilon})$. Then Theorem 5.1 indicates that the total payment of the designed mechanism $\boldsymbol{R}^{(\widetilde{N},\widetilde{\epsilon})}$ 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 5.1 is proved by Lemma 5.2 and Lemma 5.3 below, where the lower bound is given by the lower bound on the value of privacy, and the upper bound is given by $\boldsymbol{R}^{(\widetilde{N},\widetilde{\epsilon})}$.

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 $\epsilon_i = \perp\!\!\!\perp$, $D(\epsilon_i) = 0$, 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 4.3.

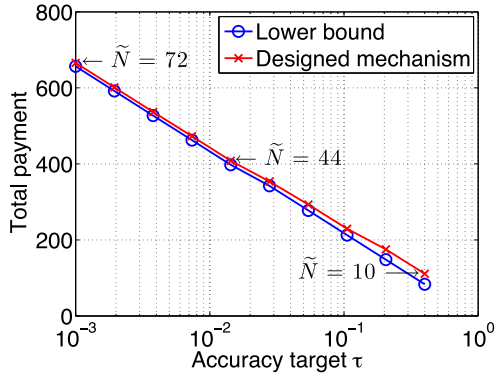

Fig. 2. Illustration of the lower and upper bounds in Theorem 5.1 on the minimum total payment for achieving an accuracy target τ, where the upper bound is given by the designed mechanism $\boldsymbol{R}^{(\widetilde{N},\widetilde{\epsilon})}$. In this example, the state has the following prior PMF: $P_W(1) = 0.7$, $P_W(0) = 0.3$. 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.

Now we use the lower bound on the value of privacy to prove the lower bound on $F(\tau)$. By Theorem 4.1,

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

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

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

Lemma 5.2. *Any feasible solution $(N, \epsilon_1, \epsilon_2, \ldots, \epsilon_N)$ of (P1) satisfies*

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

where $\widetilde{\epsilon}$ and $\widetilde{N}$ are defined in Equation (21).

Lemma 5.2 states that the total expected payment of the data collector is at least $(\widetilde{N}-1)V_{\mathrm{LB}}(\widetilde{\epsilon})$. Note that the value given by the genie-aided payment mechanism $\widehat{\boldsymbol{R}}^{(\widetilde{\epsilon})}$ for $\widetilde{N}$ individuals is $\widetilde{N}V_{\mathrm{LB}}(\widetilde{\epsilon})$, which is at most one $V_{\mathrm{LB}}(\widetilde{\epsilon})$ away from the optimal value of (P1). We can think of $V_{\mathrm{LB}}(\epsilon)$ 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 $(\widetilde{N}, \widetilde{\epsilon}, \ldots, \widetilde{\epsilon})$ to be a near-optimal choice is that the privacy level $\widetilde{\epsilon}$ gives the best quality/price ratio and $\widetilde{N}$ is the fewest number of individuals to meet the accuracy target. The proof of Lemma 5.2 is presented is Appendix D. With this lemma, the lower bound on $F(\tau)$ in Theorem 5.1 is straightforward.

5.2.2 *Upper Bound.*

Lemma 5.3. *Choose the parameters in the payment mechanism $\boldsymbol{R}^{(N,\epsilon)}$ defined in Section 4.2.1 to be $\epsilon = \widetilde{\epsilon}$ and $N = \widetilde{N}$, where $\widetilde{\epsilon}$ and $\widetilde{N}$ are defined in Equation (21). Then in the Nash equilibrium $\boldsymbol{\sigma}^{(\widetilde{\epsilon})}$*

of $\boldsymbol{R}^{(\widetilde{N},\widetilde{\epsilon})}$*, the accuracy target* τ *can be achieved, and the total expected payment is upper bounded as*

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

This lemma follows from Theorem 4.5 and we omit the proof here. Since the payment mechanism $\boldsymbol{R}^{(N,\epsilon)}$ together with $\epsilon = \widetilde{\epsilon}$ and $N = \widetilde{N}$ is a feasible solution of the payment–accuracy problem, the upper bound in this lemma gives the upper bound on $F(\tau)$ in Theorem 5.1.

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

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

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

$$p_e \leq e^{-\sum_{i=1}^{N} D(\epsilon_i)}. \tag{22}$$

Therefore, if we want to guarantee that $p_e \leq p_e^{\max}$ for some maximum allowable probability of error $p_e^{\max}$, then 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 (22) 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 EQUILIBRIUM ANALYSIS OF MECHANISM $R^{(N,\epsilon)}$

We have showed that our designed mechanism $\boldsymbol{R}^{(N,\epsilon)}$ given in Section 4.2.1 has a desired Nash equilibrium $\boldsymbol{\sigma}^{(\epsilon)}$, where every individual uses the ϵ-strategy. Besides this equilibrium, the mechanism $\boldsymbol{R}^{(N,\epsilon)}$ also has other equilibria. For example, the strategy profile where every individual reports the same data disregard of their signals is also a Nash equilibrium. In this section, we further analyze $\boldsymbol{R}^{(N,\epsilon)}$ to give a more in-depth characterization of its Nash equilibria. We have the following main results:

- *Homogeneity.* Any Nash equilibrium of the mechanism $\boldsymbol{R}^{(N,\epsilon)}$ is homogeneous, i.e., all the individuals use the same strategy in a Nash equilibrium.
- *Convergence.* For the mechanism $\boldsymbol{R}^{(N,\epsilon)}$, consider a Nash equilibrium where each individual uses the ϵ_N-strategy with $\epsilon_N > 0$. Then the privacy level ϵ_N converges to the desired level ϵ as the number of individuals N goes to infinity for non-uniform priors over W.

The multiplicity of Nash equilibria makes it difficult to obtain the outcome of the desired Nash equilibrium when implementing the mechanism. In practice, the data collector would benefit from priming the individuals to use the ϵ-strategy. The above results partially address this concern by guaranteeing that even if the designed mechanism $\boldsymbol{R}^{(N,\epsilon)}$ may have equilibria with positive privacy levels different from the desired level ϵ, these levels are close to ϵ when the population size is large. However, other equilibria still exist and may give every individual a higher payment than the desired equilibrium $\boldsymbol{\sigma}^{(\epsilon)}$. For example, when the prior PMF of W has $P_W(0) < P_W(1)$, the payment to every individual at the non-informative equilibrium where every individual reports 0

is larger than the payment at $\boldsymbol{\sigma}^{(\epsilon)}$. Further, individuals have no privacy loss at a non-informative equilibrium. This issue has not been resolved in the current article, and we will further investigate it in future work, possibly by utilizing techniques that mitigate the multiplicity problem in the peer prediction literature. For example, Kong et al. [20] add punishment when all the individuals report the same data, which makes the non-informative equilibria give lower payments to every individual than the truthful equilibrium.

6.1 Homogeneity of Nash Equilibrium

Theorem 6.1. *In any Nash equilibrium of the mechanism $\boldsymbol{R}^{(N,\epsilon)}$, the equilibrium strategies of the individuals have the same form.*

This homogeneity result is established by the properties of $\boldsymbol{R}^{(N,\epsilon)}$'s Nash equilibria presented in Lemma 6.2 below, the proof of which is presented in Appendix E. In a Nash equilibrium, we only need to focus on symmetric randomized response strategies, i.e., ϵ'-strategies with $\epsilon' \neq 0$, and non-informative strategies due to Lemma 4.2 in Section 4.1.1. Lemma 6.2 gives three properties of a Nash equilibrium of $\boldsymbol{R}^{(N,\epsilon)}$. The first property indicates that the randomized response strategies of different individuals in a Nash equilibrium all have the same privacy level. The second property implies that the non-informative strategies of different individuals in a Nash equilibrium all have the same form. Then, finally, the third property shows that a Nash equilibrium cannot consist of both informative and non-informative strategies. Combining these three properties leads to the homogeneity result in Theorem 6.1.

Lemma 6.2. *Any Nash equilibrium of the mechanism $\boldsymbol{R}^{(N,\epsilon)}$ has the following properties:*

(1) Suppose that the equilibrium strategies of two individuals i and j are the ϵ_i-strategy and ϵ_j-strategy with $\epsilon_i \neq 0$ and $\epsilon_j \neq 0$, respectively. Then $\epsilon_i = \epsilon_j$.
(2) Suppose that the equilibrium strategies of two individuals are non-informative strategies. Then these two equilibrium strategies are the same.
(3) Suppose that the equilibrium strategies of two individuals i and j are σ_i and σ_j. Then it is impossible to have that σ_i is an ϵ'-strategy with some $\epsilon' \neq 0$ and σ_j is a non-informative strategy.

6.2 Convergence of Positive Privacy Levels

In this section, we focus on the Nash equilibria where individuals report data informatively and with a probability greater than 1/2 the reported data is truthful, i.e., the equilibrium strategies are ϵ'-strategies with some $\epsilon' > 0$. We consider a general case where $\mathbb{P}_W(1) \neq \mathbb{P}_W(0)$. Consider the sequence of mechanisms $\{R^{(N,\epsilon)}, N = 2, 3, \dots\}$. Let $\boldsymbol{\sigma}^{(N,\epsilon_N)}$ denote the strategy profile where the population has N individuals and every individual uses the ϵ_N-strategy. For each N, pick any $\epsilon_N > 0$ such that $\boldsymbol{\sigma}^{(N,\epsilon_N)}$ is a Nash equilibrium of the mechanism $\boldsymbol{R}^{(N,\epsilon)}$ (at least we can pick ϵ). Then the following theorem guarantees that any such sequence $\{\epsilon_N, N = 2, 3, \dots\}$ converges to the desired privacy level ϵ.

Theorem 6.3. *For any sequence of strategy profiles $\{\boldsymbol{\sigma}^{(N,\epsilon_N)}, N = 2, 3, \dots\}$, where each $\epsilon_N > 0$ and $\boldsymbol{\sigma}^{(N,\epsilon_N)}$ is a Nash equilibrium of the mechanism $\boldsymbol{R}^{(N,\epsilon)}$, we have*

$$\lim_{N\to\infty} \epsilon_N = \epsilon.$$

The intuition behind Theorem 6.3 is that when the population size N becomes large, at any Nash equilibrium with non-diminishing data quality, the majority of the reported data from individuals except one individual converges to the underlying state W. Then the mechanism $\boldsymbol{R}^{(N,\epsilon)}$ works

similar as the genie-aided mechanism $\widehat{\boldsymbol{R}}^{(\epsilon)}$, which has a unique Nash equilibrium that consists of ϵ-strategies. So the main proof difficulty is to show that the data quality does not diminish as N becomes large, i.e., ϵ_N does not go to zero as $N \to \infty$. A detailed proof of Theorem 6.3 is presented in Appendix F.

7 DISCUSSIONS ON FURTHER GENERALIZATIONS

The model we study in this article considers binary signals $S_1, \ldots, S_N$ that come from a binary underlying state W, where the prior is symmetric in the sense that $\mathbb{P}(S_i = 1 \mid W = 1) = \mathbb{P}(S_i = 0 \mid W = 0)$. A relatively straightforward extension would be to consider asymmetric priors where $\mathbb{P}(S_i = 1 \mid W = 1) \neq \mathbb{P}(S_i = 0 \mid W = 0)$. For this setting, our characterization of individuals' equilibrium strategies in Lemma 4.2 and our techniques of establishing a lower bound on the payment through analyzing genie-aided mechanisms still apply. The mechanism $R^{(N,\epsilon)}$ we design can also be adapted according to this prior to make sure that it has a desired Nash equilibrium and thus gives an upper bound on the payment. However, naively following the steps we have yields a gap between the lower and upper bounds. It is not immediately clear how we should obtain asymptotically tight lower and upper bounds.

The characterization of individuals' equilibrium strategies in Lemma 4.2 further applies to more general probability models of binary signals. For example, we can just assume that $S_1, \ldots, S_N$ come from a general joint probability distribution over $\{0,1\}^N$, without any underlying state in the model. For this setting, we need to properly model the goal of the data collector, since learning the state no longer makes sense.

When the signals are beyond binary, e.g., they can be k-ary data with $k > 2$, we expect it to require substantially new techniques to study the relation between privacy and payment. To see this, we consider the data reporting strategy of an individual, which prescribes a distribution to the reported data for each possible value of the signal. So a strategy can be represented by k^2 probabilities. For example, for binary data, these are the p_1, p_0, q_1, q_0 in the proof of Theorem 4.5 in Section 4.2.3. Then the privacy loss, measured by the level of local differential privacy, depends on these k^2 probabilities in a complicated, non-convex way. This dependence becomes difficult to track when k is larger than 2. Therefore, it is hard to characterize an individual's best response strategy for $k > 2$, let alone establish lower and upper bounds on the payment to induce certain privacy level. This difficulty also illustrates the difference between the problem we study and the classical setting for peer prediction: We focus on the relation between privacy and payment, whereas peer prediction aims at eliciting honest report.

8 CONCLUSIONS

In this article, 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 non-negative 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 $\boldsymbol{R}^{(N,\epsilon)}$. 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 $\boldsymbol{R}^{(N,\epsilon)}$ with properly chosen parameters is at most one individual's payment away from the minimum. A more in-depth analysis of the Nash equilibria of $\boldsymbol{R}^{(N,\epsilon)}$ was also given,

which indicates that any Nash equilibrium of $\boldsymbol{R}^{(N,\epsilon)}$ is homogeneous, and if we focus on the Nash equilibria where individuals report data informatively and with a probability greater than 1/2 that the reported data are truthful, then the privacy levels of such equilibria converge to the desired level ϵ as the number of individuals N becomes large.

REFERENCES

[1] Raef Bassily and Adam Smith. 2015. Local, private, efficient protocols for succinct histograms. In *Proceedings of the Annual ACM Symposium on the Theory of Computing (STOC'15)*. 127–135.

[2] Yiling Chen, Stephen Chong, Ian A. Kash, Tal Moran, and Salil Vadhan. 2013. Truthful mechanisms for agents that value privacy. In *Proceedings of the ACM Conference on the Electronic Commerce (EC'13)*. 215–232.

[3] Yiling Chen, Or Sheffet, and Salil Vadhan. 2014. Privacy games. In *International Conference on the Web and Internet Economics (WINE'14)*, Vol. 8877. 371–385.

[4] Thomas M. Cover and Joy A. Thomas. 2006. *Elements of Information Theory* (2nd ed.). John Wiley & Sons, Hoboken, NJ.

[5] John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. 2013. Local privacy and minimax bounds: Sharp rates for probability estimation. In *Advances Neural Information Processing Systems (NIPS'13)*. 1529–1537.

[6] John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. 2013. Local privacy and statistical minimax rates. In *Proceedings of the Annual IEEE Symposium on the Foundations of Computer Science (FOCS'13)*. 429–438.

[7] Cynthia Dwork. 2006. Differential privacy. In *Proceedings of the International Conference on Automata, Languages and Programming (ICALP'06)*. 1–12.

[8] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Conference on the Theory of Cryptography (TCC'06)*. 265–284.

[9] Cynthia Dwork and Aaron Roth. 2014. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.* 9, 3–4 (Aug. 2014), 211–407.

[10] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. 2014. RAPPOR: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the ACM SIGSAC Conference on the Computer and Communication Security (CCS'14)*. 1054–1067.

[11] Giulia C. Fanti, Vasyl Pihur, and Úlfar Erlingsson. 2016. Building a RAPPOR with the unknown: Privacy-preserving learning of associations and data dictionaries. In *Proceedings on Privacy Enhancing Technologies (PoPETs'16)*. 41–61.

[12] Lisa K. Fleischer and Yu-Han Lyu. 2012. Approximately optimal auctions for selling privacy when costs are correlated with data. In *Proceedings of the ACM Conference on the Electronic Commerce (EC'12)*. 568–585.

[13] Arpita Ghosh and Katrina Ligett. 2013. Privacy and coordination: Computing on databases with endogenous participation. In *Proceedings of the ACM Conference on the Electronic Commerce (EC'13)*. 543–560.

[14] Arpita Ghosh, Katrina Ligett, Aaron Roth, and Grant Schoenebeck. 2014. Buying private data without verification. In *Proceedings of the ACM Conference on the Economics and Computation (EC'14)*. 931–948.

[15] Arpita Ghosh and Aaron Roth. 2011. Selling privacy at auction. In *Proceedings of the ACM Conference on the Electronic Commerce (EC'11)*. 199–208.

[16] Justin Hsu, Sanjeev Khanna, and Aaron Roth. 2012. Distributed private heavy hitters. In *Proceedings of the International Conference on the Automata, Languages and Programming (ICALP'12)*. 461–472.

[17] Thomas Kailath. 1967. The divergence and Bhattacharyya distance measures in signal selection. *IEEE Trans. Commun. Technol.* 15, 1 (Feb. 1967), 52–60.

[18] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. 2014. Extremal mechanisms for local differential privacy. In *Proceedings of the Conference on Advances Neural Information Processing Systems (NIPS'14)*. 2879–2887.

[19] Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. 2011. What can we learn privately? *SIAM J. Comput.* 40, 3 (May 2011), 793–826.

[20] Yuqing Kong, Katrina Ligett, and Grant Schoenebeck. 2016. Putting peer prediction under the micro(economic)scope and making truth-telling focal. In *International Conference on the Web and Internet Economics (WINE'16)*. 251–264.

[21] Steve Kroft. 2014. The data brokers: Selling your personal information. *CBS News* (March 2014).

[22] Katrina Ligett and Aaron Roth. 2012. Take it or leave it: Running a survey when privacy comes at a cost. In *Proceedings of the International Workshop Internet and Network Economics (WINE'12)*. 378–391.

[23] Nolan Miller, Paul Resnick, and Richard Zeckhauser. 2005. Eliciting informative feedback: The peer-prediction method. *Manage. Sci.* 51 (Sep. 2005), 1359–1373.

[24] Kobbi Nissim, Salil Vadhan, and David Xiao. 2014. Redrawing the boundaries on purchasing data from privacy-sensitive individuals. In *Proceedings of the Conference on the Innovations in Theoretical Computer Science (ITCS'14)*. 411–422.

[25] Mallesh M. Pai and Aaron Roth. 2013. Privacy and mechanism design. *SIGecom Exch.* 12, 1 (Jun. 2013), 8–29.
[26] Aaron Roth and Grant Schoenebeck. 2012. Conducting truthful surveys, cheaply. In *Proceedings of the ACM Conference on the Electronic Commerce (EC'12)*. 826–843.
[27] Reza Shokri. 2015. Privacy games: Optimal user-centric data obfuscation. In *Proceedings of the Conference on Privacy Enhancing Technologies (PETS'15)*. 299–315.
[28] R. Srikant and Lei Ying. 2014. *Communication Networks: An Optimization, Control and Stochastic Networks Perspective*. Cambridge University Press, New York, NY.
[29] Weina Wang, Lei Ying, and Junshan Zhang. 2014. On the relation between identifiability, differential privacy, and mutual-information privacy. In *Proceedings of the Annual Allerton Conference on the Communication, Control and Computing*. 1086–1092.
[30] Weina Wang, Lei Ying, and Junshan Zhang. 2015. A minimax distortion view of differentially private query release. In *Proceedings of the Asilomar Conference on the Signals, Systems, and Computers*. 1046–1050.
[31] Stanley L. Warner. 1965. Randomized response: A survey technique for eliminating evasive answer bias. *J. Am. Stat. Assoc.* 60, 309 (Mar. 1965), 63–69.
[32] David Xiao. 2013. Is privacy compatible with truthfulness?. In *Proceedings of the Conference on Innovations in Theoretical Computer Science (ITCS'13)*. 67–86.

Received June 2016; revised June 2017; accepted June 2018