Impact of Social Learning on Privacy-Preserving Data Collection

Abdullah Basar Akbay, *Member, IEEE*, Weina Wang, *Member, IEEE*, and Junshan Zhang, *Fellow, IEEE*

***Abstract*—We study a game-theoretic model where a data collector purchases data from users through a payment mechanism. Each user has her personal signal which represents her knowledge about the underlying state the data collector desires to learn. Through social interactions, each user can also learn noisy versions of her friends' personal signals, which are called 'group signals'. We develop a Bayesian game theoretic framework to study the impact of social learning on users' data reporting strategies and devise the payment mechanism for the data collector accordingly. We show that the Bayesian-Nash equilibrium can be in the form of either a symmetric randomized response (SR) strategy or an informative non-disclosive (ND) strategy. Specifically, a generalized majority voting rule is applied by each user to her noisy group signals to determine which strategy to follow. Our findings reveal that both the data collector and the users can benefit from social learning which drives down the privacy costs and helps to improve the state estimation for a given total payment budget. Further, we derive bounds on the minimum total payment required to achieve a given level of state estimation accuracy.**

***Index Terms*—Data collection, differential privacy, social networks, crowdsourcing.**

I. Introduction

In THIS work, we study a game-theoretic market model in which users make strategic decisions to trade privacy-preserved versions of their personal data with a data collector. We assume that the users are rational, risk-neutral and self-interested. Our analysis generalizes the existing market models for private data collection [1] by incorporating the ubiquitous social interactions among users encountered in many settings in our everyday life. Specifically, we ask the question of what are the desired data reporting strategies (from an individual user perspective) and payment mechanisms (from a data collector perspective), when users can *learn* noisy versions of their friends' data through social interactions. Intuitively, social interactions among the users can help them to become better-informed, which in turn can impact their decision strategies by improving the quality of their data reporting. To the best of our knowledge, this article is the first to investigate the impact of social learning on privacy-preserving data collection which differentiates our study from the existing work [1]–[10], where the users' personal data are the only information at their disposal.

Manuscript received August 15, 2020; revised December 4, 2020; accepted January 14, 2021. Date of publication January 22, 2021; date of current version March 16, 2021. This work was supported in part by the National Science Foundation under Grant CNS-2003081, Grant CPS-1739344, and Grant SaTC-1618768. *(Corresponding author: Abdullah Basar Akbay.)*

Abdullah Basar Akbay and Junshan Zhang are with the School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ 85281 USA (e-mail: aakbay@asu.edu).

Weina Wang is with the Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213 USA.

This article has supplementary downloadable material available at https://doi.org/10.1109/JSAIT.2021.3053545, provided by the authors.

Digital Object Identifier 10.1109/JSAIT.2021.3053545

A. Data Collection and Social Learning

Social interactions can affect and introduce further complications into various data collection efforts including product ratings, political campaigns, smartphone applications or hotel and restaurant reviews. Consider a real-world setting where a group of friends, who are users on an online platform (*data collector*), such as *IMDB*, *Flixster* or *Netflix* which aims to collect audience ratings, watch a movie together. It is plausible that the rating or review from an individual is not formed solely based on their initial impression (*personal signal*) about the movie and it is also influenced by the opinions of her friends (*group signals*). Users are not bound to truthfully share their personal opinions with the data collector, and may opt out from data collection. The data collector can utilize a payment mechanism to incentivize participation and reward the users who report informative data. Nevertheless, the users are still not compelled to act truthfully and the data collector is not equipped with an instrument to directly authenticate their reported data.

The information is represented by a binary random variable, W, which is called the state. The users receive noisy individual copies of W. As a result of prevalent social interactions, the users can also obtain noisy observations of their friends' individual signals. The social learning among users, which can take place in many forms, including in face-to-face meetings and over multiple online social media (e.g., Facebook and Twitter), is captured by a *social learning graph* (or social graph for brevity). Each vertex of this undirected social graph corresponds to a user and each edge of this social graph points to information exchange between two friends.[1] The data collector attempts to learn the underlying state W based on data collection from the privacy-aware users, by using a payment mechanism to incentivize user participation. Based on her private signal S_i and noisy copies of her friends' signals, C_{ij}'s (when i and j are friends), each user i reports data X_i, which may incur privacy cost. As a result, each user can either choose to report data or not to participate. Figure 1 depicts an illustration of the information flow in this market model.

[1]The same analysis can be carried out for directed graph models. In Section V-C, simulations are provided for both directed and undirected real world networks.

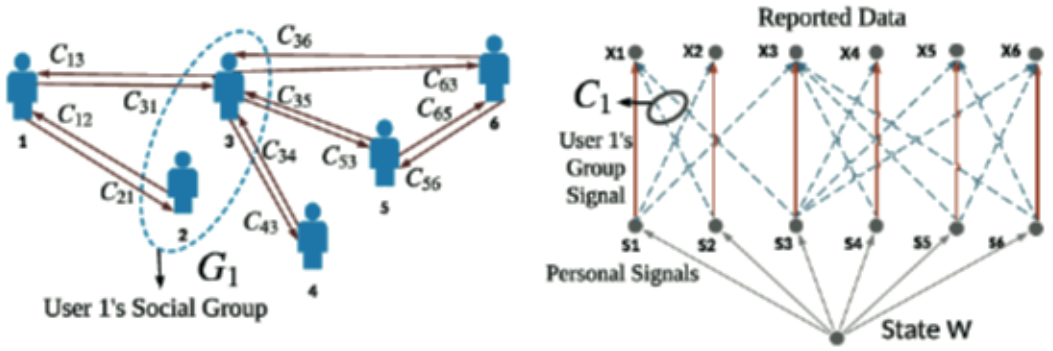

(a) Illustration of Social Learning Graph **(b) Data Reporting with Social Learning**

Fig. 1. The data collector is interested in learning the binary state W. Each user i possesses her personal data S_i and her group signals vector $\mathbf{C}_i$ whose components are noisy copies of her friends' personal signals. (a) If two users i and j are friends, each of them gets a noisy copy of the other's personal signal. (b) Conditioned on W, $S_1, S_2, \ldots, S_N$ are i.i.d. binary signals. The noisy copies of S_i, which are received by user i's friends, are i.i.d. binary signals conditioned on S_i. Taking S_i and $\mathbf{C}_i$ as inputs, user i generates her reported data X_i.

B. Challenges

We develop a Bayesian game theoretic framework to study the impact of social learning on users' data reporting strategies and devise the payment mechanism for the data collector. Using a variant of the peer-prediction method,[2] the data collector scores the reported data of user i by comparing it to his estimate of the underlying state W, which is computed based on the reported data of all other users. In the presence of social learning, the *quality* of the reports can vary across the users, as different users can have different numbers of friends and each user is capable of claiming the control of her privacy level against the data collector. As expected, social learning among users introduces coupling and heterogeneity in the reported data and significantly complicate the design. Under these challenges, we seek to answer the following key questions: When a user consents to publish her review, what is the best strategy for her to leverage her friends' noisy signals as opposed to her own personal signal in her reported data? Can the users benefit from social learning, and if yes what is the corresponding desired data reporting strategy? Can the data collector design incentive mechanisms to take advantage of social learning? Further, what payment mechanism enables the data collector to minimize the cost in the presence of social learning?

C. Relevant Work

Market models where privacy-aware strategic users treat their data as a commodity have received much interest [1]–[18]. Recently, privacy in edge computing paradigms have been also been studied from various perspectives [19]–[22]. In all these studies, the users are regarded as individual agents but social learning among them is not accounted for. The market model proposed in [1] can be differentiated in this stream of work where each user directly controls the privacy level of her reported data. Our proposed model can be regarded as a generalization of the model in [1] which assumes that the knowledge of each user is limited to their personal signals. As illustrated in Figure 1, the users have richer information about the underlying state beyond their private signals, thanks to the presence of social learning, and can therefore use this additional information to conceive their reporting strategy which can potentially have significant impact on the data collection.

User heterogeneity in peer prediction has recently gained attention, but there are very few results on handling its complications [2], [18], [23]. Furthermore, little attention has been paid to the cases where the reported data is correlated across users given the true state. We shall study both these two issues in the market model in this article. In particular, we will first consider the users as *local data curators*, then treat the data collector as a *fusion center*, and finally combine them and design the payment mechanism. This enable us to separate the convoluted dependence among the personal and group signals, and make the study tractable.

We also revisit the notion of truthful data reporting and informative user strategies in the framework of data privacy games where the users report their data using randomized response strategies to achieve privacy protection. We show that, in the presence of social learning, this conventional notion of "informative strategy" would not encompass some desired equilibria where each user reports informative data based on her friends' signals only. Building on this new insight, we introduce informative non-disclosive strategies which allows a user to formulate strategies based on only her learned group signals if there is strong concurrence. We caution that informative non-disclosive strategies does not create situations akin to herding [24] or information cascades [25]. In the proposed market model, the users take their actions in parallel, not sequentially, and the reported data in this study, are only revealed to the data collector not to the other participants.

D. Summary of Main Results

The primary objective of the data collector is the estimation of the underlying state W. The data collector does not have any observation about W and he relies on users' reported data $\mathbf{X}$. Since the users are strategic and they incur privacy costs, the data collector employs a payment mechanism to incentivize informative data reporting. In general, the payment each user receives also depend on the reported data of other users. Thus,

[2]See [4] for a recent and extensive survey of this field.

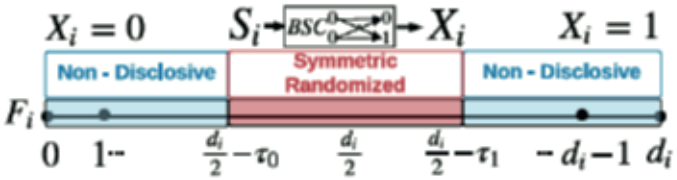

Fig. 2. Strategy profile at Bayesian Nash Equilibria (BNE): The data reporting strategy at BNE is in the form of either a non-disclosive strategy or a symmetric randomization strategy. Using F_i, defined as the sum of the social group signals, each user (with degree d_i) determines which strategy to follow.

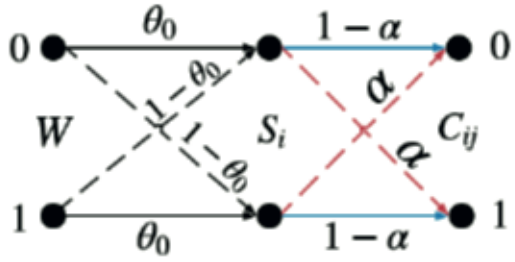

Fig. 3. The parameter θ_0 determines the quality of the personal signals. The noise level in social learning is modeled with crossover probability, α. The parameter θ_1, defined as $\theta_1 = \theta_0(1-\alpha) + (1-\theta_0)\alpha$, corresponds to the quality of the group signals.

the announcement of payment mechanism actuates a privacy-preserving data collection game among the users. Our main contributions can be summarized as follows.

- Our findings reveal that the data reporting strategies at the Bayesian-Nash equilibria (BNE) can be in the form of either symmetric randomized response (SR) strategy or non-disclosive (ND) strategy. When a user plays the ND strategy, her reported data is completely based on her learned group signals, independent of her private signal. Intuitively, albeit having a signal different from the majority of her friends' signals, a user might be better off to pretend in accordance with them and her reported data is still informative for the data collector.
- We design the data collector's payment mechanism using peer-prediction to incentivize the users for informative data reporting. Furthermore, our results demonstrate that, under the proposed mechanism, the data reporting strategy at the BNE is in the form of either an ND strategy or an SR strategy. As illustrated in Figure 2, each user follows a *majority voting* based data reporting rule to determine which strategy to employ. In general, as the noise level of the group signals increases, the ND strategy requires a 'higher' majority, and otherwise the user follows the SR strategy.
- To tackle the technical difficulty that the reported data is correlated across users given the underlying state, we use a Central Limit Theorem for dependence graphs to characterize the statistics of the reported data profile, based on which the data collector can evaluate the estimation error of W. The total expected payment is then characterized for a given accuracy target. Our analysis pinpoints to the positive impact of social learning on the privacy-preserving data collection game, in the sense that the data collector can lower the total payment significantly, compared to the case with no social learning.
- Our results demonstrates that both the data collector and the users can benefit from social learning which drives down the privacy costs and helps to improve the accuracy of the state estimation. In particular, some users' privacy cost can be driven to zero when ND strategies are used. This, in turn, benefits the data collector and drives down the overall cost, since his data resources are more informed and can report informative data at lower privacy costs.

The rest of the article is organized as follows. We introduce the models for the privacy-preserving data collection market and the social learning graph in Section II. We formalize the Bayesian game under this market model in Section III. We present the main results on the data reporting strategies and the payment mechanism design in Section IV. We discuss the impact of social learning on the payment and accuracy in Section V. Finally, we summarize and discuss possible extensions and open problems in Section VI. The key notation used in this article is shown in Table I.

II. System Model

A. Personal Signals

Consider a market model where the data collector is interested in learning the underlying state W from a set $\mathcal{I} = \{1, 2, \ldots, N\}$ of $N \geq 2$ users. For ease of exposition, W is assumed to be a binary random variable (r.v.), for example, representing the product quality as *good* or *bad*.[3] We assume that the prior distribution $\mathrm{P}_W(\cdot)$ is common knowledge and both $\mathrm{P}_W(1), \mathrm{P}_W(0) > 0$. As illustrated in Figure 1, each user i possesses a binary signal S_i, which is her personal data, representing her knowledge about W. The personal signal profile of the entire population is denoted as $\mathbf{S} = [S_1\ S_2\ \cdots\ S_N]$. Given W, it is assumed that the binary signals S_i's are independent and identically distributed and the parameter θ_0 with $0.5 < \theta_0 < 1$ determines the *quality of the personal signals* for every user:

$$\begin{aligned} \mathrm{P}(S_i = 1 - w | W = w) &= 1 - \theta_0, \\ \mathrm{P}(S_i = w | W = w) &= \theta_0, \quad \text{for } w \in \{0, 1\}. \end{aligned} \tag{1}$$

B. Group Signals

The *social learning graph* $\mathcal{G} = \{\mathcal{I}, \mathcal{E}\}$ is used to model the social coupling among the users. The vertex set is the set of individuals $\mathcal{I}$ and the edge set is given as $\mathcal{E} = \{(i, j) \in \mathcal{I} \times \mathcal{I} : \mathcal{E}_{ij} = 1\}$ where $\mathcal{E}_{ij} = 1$ if and only if there is a social tie between i and j where $i \neq j$. User i's social group G_i is defined as the set of her friends: $G_i = \{j \in \mathcal{I} : \mathcal{E}_{ij} = 1\}$. The number of friends i has is called the degree d_i of that user. We assume

[3] Binary feedback and review systems are prevalent, e.g., Youtube [26] and Netflix [27] swapped out their five star rating systems for a binary system. On many platforms, it is observed that the vast majority of ratings are either the best or the worst option [28], [29].

TABLE I
NOMENCLATURE

Notation	Description	Notation	Description
W	Underlying state	T_i	The type of user i
N	Number of users	X_i	The reported data of user i
S_i	The personal signal of user i	σ_i	The strategy of user i
θ_0	Quality of the personal signals	$R_i(\mathbf{x})$	The payment user i receives
G_i	The social group of user i	$\zeta_i(\sigma_i, c_i)$	The privacy level of strategy σ_i
d_i	The degree of user i	$g(\zeta)$	Privacy cost function
ρ_d	Degree distribution	F_i	The sum of the group signals C_i
$\mathbf{C}_i$	The group signals of user i	σ_i^*	The majority voting data reporting
α	Crossover probability	$\tilde{\mathbf{R}}$	Proposed PP payment mechanism
θ_1	Quality of the group signals	$\mathbf{R}^g$	The genie-aided mechanism

that the social learning graph is a sparse[4] random[5] graph with node degrees following a distribution ρ_d with maximal degree $D_{\max}$. The degree distribution ρ_d is common knowledge for the data collector and the users. However, the data collector does not have any further knowledge about the social graph $\mathcal{G}$. The users know who their friends are, but they do not possess any further knowledge about their friends' social groups.

Each user i has noisy copies of her friends' personal signals. For user i with social group G_i, let vector $\mathbf{C}_i$ denote the vector of her *group signals*: $\mathbf{C}_i = [C_{ij_1} C_{ij_2} \dots C_{ij_{d_i}}]$, where C_{ij}'s are binary valued. To capture the noise in social learning of group signals, it is assumed that friends' personal signals are "flipped" with *crossover* probability α: $P(C_{ij} = 1|S_j = 0) = P(C_{ij} = 0|S_j = 1) = \alpha, 0 \le \alpha < 0.5$. Note that these "flips" are statistically independent, i.e., given j_1 and j_2 are friends with i, $\mathrm{P}(C_{j_1 i} = s_i|C_{j_2 i}, S_i = s_i) = \mathrm{P}(C_{j_1 i} = c_{j_1 i}|S_i = s_i) = 1-\alpha$. The parameter $\theta_1 = \theta_0(1-\alpha) + (1-\theta_0)\alpha$ points to the *quality of group signals*:

$$\begin{aligned} \mathrm{P}(C_{ij} = 1 - w|W = w) &= 1 - \theta_1, \\ \mathrm{P}(C_{ij} = w|W = w) &= \theta_1, \quad \text{for } w \in \{0, 1\}. \end{aligned} \tag{2}$$

C. Data Reporting Strategies

The type of a user is defined as $T_i = [S_i\ \mathbf{C}_i]$. Respectively, the type space $\mathcal{T}$ can be defined as $\mathcal{T} = \cup_{k=0}^{D_{\max}} \mathcal{T}_k$ where $\mathcal{T}_k = \{0, 1\}^{k+1}$. The users do not know the knowledge of other users' type vectors. The reported data of user i is denoted with X_i. It follows that $\mathbf{X} = [X_1\ X_2\ \dots\ X_N]$ is the reported data profile, where $X_i \in \mathcal{X} = \{0, 1, \perp\}$ and $\perp$ represents "non-participation". User i's strategy σ_i is a mapping from the type space $\mathcal{T}$ to $\Delta(\mathcal{X})$[6] and it specifies the probabilities $\mathrm{P}_{\sigma_i}(X_i \in \mathcal{F}|T_i = t_i)$ for all $\mathcal{F} \subseteq \mathcal{X}$. The data reporting strategies can be considered as a contingent plan of actions for different private type realizations.

[4]The average number of friends users have is much smaller than the total number of users.

[5]We follow the configuration model described in [30]. The degrees $\{d_i\}_{i=1}^N$ are independent and identically distributed random integers drawn from ρ_d. Pairs of users are chosen at random and edges are formed between them until complete pairing according to the drawn degree sequence. If complete pairing is not possible, one d_i can always be discarded and redrawn from ρ_d.

[6]$\Delta(\mathcal{X})$ being the set of all probability distributions over $\mathcal{X}$.

D. Data Privacy Model

Based on the celebrated notion of differential privacy [31], we define the privacy loss inflicted on the users as the level of local differential privacy when using the strategy σ_i. Given her group signal $\mathbf{C}_i = \mathbf{c}_i$, user i's privacy loss decreases as her data reporting makes her personal signal S_i more *indistinguishable*. The privacy level of strategy σ_i, given $\mathbf{C}_i = \mathbf{c}_i$, is defined as

$$\zeta_i(\sigma_i, \mathbf{c}_i) = \max_{\mathcal{F} \subseteq \{0,1,\perp\},\ s_i \in \{0,1\}} \times \ln\left(\frac{\mathrm{P}_{\sigma_i}(X_i \in \mathcal{F}|\mathbf{C}_i = \mathbf{c}_i, S_i = s_i)}{\mathrm{P}_{\sigma_i}(X_i \in \mathcal{F}|\mathbf{C}_i = \mathbf{c}_i, S_i = 1 - s_i)}\right), \tag{3}$$

where the convention $0/0 = 1$ is followed. The privacy cost of the user is determined by $g(\zeta_i(\sigma_i, \mathbf{c}_i))$. We assume that, $g(\cdot)$ is homogeneous across the users, convex, continuously differentiable, strictly increasing, nonnegative and $g(0) = 0$.

Intuitively, given her group signal $\mathbf{C}_i = \mathbf{c}_i$, user i's privacy cost decreases as her data reporting makes her private signal more indistinguishable. Notice that this function is defined for every possible group signal realization, offering a stronger privacy guarantee than an alternative definition

$$\zeta_i(\sigma_i) = \max_{\mathcal{F} \subseteq \{0,1,\perp\},\ s_i \in \{0,1\}} \ln\left(\frac{\mathrm{P}_{\sigma_i}(X_i \in \mathcal{F}|S_i = s_i)}{\mathrm{P}_{\sigma_i}(X_i \in \mathcal{F}|S_i = 1 - s_i)}\right),$$

since it assumed the worst case of the adversary already knowing $\mathbf{C}_i$. Figure 4 presents an example where the user i has two friends j_1 and j_2. In this example, $\zeta_i(\sigma_i, c_i)$ measures the indistinguishability of the private signal bit S_i for each $c_i \in \{0, 1\} \times \{0, 1\}$. Consider an extreme case where the user's reported data is her private signal, $X_i = S_i$ given $\mathbf{C}_i = \mathbf{c}_i$. In this case, $\zeta(\sigma_i, c_i)$ is equal to ∞, the maximum possible privacy leakage for her. Consider another case where the individual report X_i is independent from the personal signal S_i given $\mathbf{C}_i = \mathbf{c}_i$. In this case, $\zeta(\sigma_i, c_i)$ is equal to 0, the minimum possible privacy leakage for her.

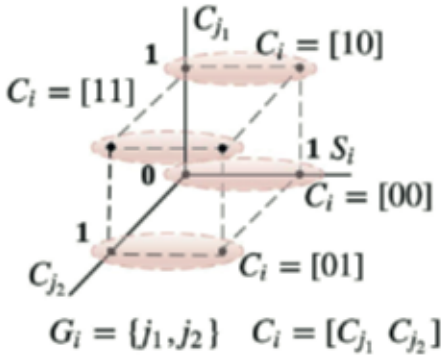

Fig. 4. $\zeta_i(\sigma_i, \mathbf{c}_i)$ measures the indistinguishability of personal signal S_i, given group signal $\mathbf{C}_i$, with respect to the user's strategy.

Fig. 5. A Market Model: a) Each user plays a Bayesian Game to report her privacy-preserved data to data collector. b) Data collector determines the amount of reward for each user based on the reported data profile $\mathbf{X} = [X_1 \dots X_N]$.

E. Payment Mechanism

The objective of the data collector is to estimate the underlying state W from users' reported data $\mathbf{X} = [X_1\ X_2\ \dots\ X_N]$ with minimum total payment subject to an accuracy constraint. The data collector does not have any observation about W and he relies on users' reported data $\mathbf{X}$. Since the users are strategic and they incur privacy costs, they may not provide informative reported data unless they are sufficiently incentivized. In this study, we assume that the data collector cannot impose penalties on the users. Therefore, positive rewards are the only options at his disposal to incentivize informative reporting as depicted on Fig. 5. We define the payment mechanism as $\mathbf{R} : \mathcal{X}^N \to \mathbb{R}^N$, where $R_i(\mathbf{x})$ specifies the amount of payment for user i given $\mathbf{X} = \mathbf{x}$ and $\mathbf{R}(\mathbf{x}) = [R_1(\mathbf{x})\ R_2(\mathbf{x})\ \dots\ R_N(\mathbf{x})]$. In general, the payment user i receives $R_i(\mathbf{x})$ also depend on the reported data of other users. Thus, the announcement of the payment mechanism instigates a strategic form game among the users where the utility of each user is the difference between her payment and her privacy cost. In the next section, we formalize the Bayesian game under this market model.

III. Characterization of Bayesian-Nash Equilibria

The Bayesian game under this market model is outlined as follows: The data collector announces a payment mechanism, which actuates a strategic form game where the users are the players aiming to maximize their expected utility, which is the difference between their rewards and their privacy costs. In this game, the common knowledge includes the prior state distribution $\mathbf{P}_W$, the signal quality parameter θ_0, the crossover probability α, the degree distribution[7] ρ_d, the privacy cost function g and the payment mechanism $\mathbf{R}$. Furthermore, it is assumed that the data collector and the users know how the personal signals (S_i's) and group signals ($\mathbf{C}_i$'s) are generated, according to (1) and (2). In this game with incomplete information, we focus on Bayesian-Nash equilibria where each user has no incentive to unilaterally change her strategy given other users' strategies. Formally, a Bayesian-Nash equilibrium (BNE) is defined in the following.[8]

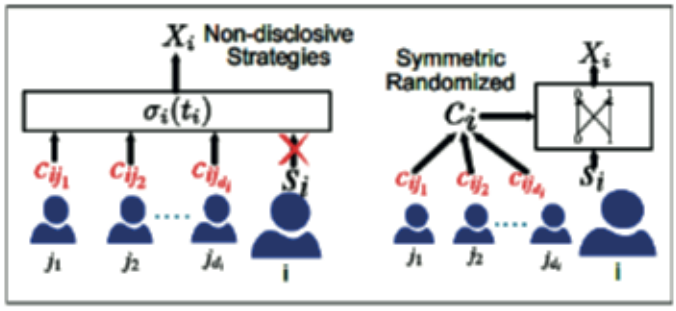

Fig. 6. Strategy profiles at BNE. (a) Non-Disclosive (ND) Strategies: X_i is independent from the personal signal S_i. In general, X_i is dependent on the group signal C_i. (b) Symmetric Randomized Response (SR) Strategies: S_i and X_i are the input and the output of a noisy binary symmetric channel. In general, the noise level depends on the group signals C_i.

Definition 1 (BNE): A strategy profile $\sigma = \{\sigma_1, \sigma_2, \dots, \sigma_N\}$ is a Bayesian-Nash equilibrium (BNE) if, for each user $i \in \mathcal{I}$,

$$\sigma_i(\cdot) \in \arg\max_{\sigma_i'(\cdot) \in \Delta(\mathcal{X})} \mathbb{E}_{(\sigma_i', \sigma_{-i})}\left[R_i(\mathbf{X}) - g\left(\zeta\left(\sigma_i', \mathbf{C}_i\right)\right)\right] \quad (4a)$$

Because each type has positive probability, this *ex ante* formulation is equivalent to user i maximizing her expected utility conditional on $T_i = t_i$ for each t_i:

$$\begin{aligned} &\sum_{\mathbf{t}_{-i}} \mathrm{P}(\mathbf{T}_{-i} = \mathbf{t}_i | T_i = t_i) \mathbb{E}_{(\sigma_i, \sigma_{-i})} \\ &\quad \times \left[R_i(\mathbf{X}) - g(\zeta(\sigma_i, \mathbf{c}_i)) | T_i = t_i, \mathbf{T}_{-i} = \mathbf{t}_{-i}\right] \\ &\geq \sum_{\mathbf{t}_{-i}} \mathrm{P}(\mathbf{T}_{-i} = \mathbf{t}_i | T_i = t_i) \mathbb{E}_{(\sigma_i', \sigma_{-i})} \\ &\quad \times \left[R_i(\mathbf{X}) - g\left(\zeta\left(\sigma_i', \mathbf{c}_i\right)\right) | T_i = t_i, \mathbf{T}_{-i} = \mathbf{t}_{-i}\right]. \end{aligned} \quad (4b)$$

In equation (4), the randomized generation of the social graph is part of this probability space since the users' knowledge about the social graph is limited to their own social groups and they do not possess any knowledge about their friends' social groups.

We show that users' data reporting strategies at the Bayesian-Nash equilibria are in the form of either *symmetric randomized responses* or *non-disclosive strategies*. Firstly, we formally define the non-disclosive strategies as follows.

[7]We remark that many different types of real-world social networks and their statistical properties including their degree distributions are well-studied in the literature and available. In [30], [32]–[34], with a primary focus on their degree distributions, an extensive review of empirical studies of social networks is presented.

[8]Given a strategy profile $\sigma(\cdot)$, and a $\sigma_i' \in \Delta(\mathcal{X})$, $(\sigma_i'(\cdot), \sigma_{-i}(\cdot))$ denotes the profile where user i plays $\sigma_i'(\cdot)$ and the other users follow $\sigma(\cdot)$.

Definition 2 (ND Strategy): If S_i and X_i are independent conditioned on W and $\mathbf{C}_i$, then $\sigma_i([\cdot \;\; c_i])$ is a non-disclosive (ND) strategy; that is to say, for every $\mathcal{F} \subseteq \{0, 1, \perp\}$ and $s_i \in \{0, 1\}$ we have

$$\mathrm{P}_{\sigma_i}(X_i \in \mathcal{F} | S_i = s_i, \mathbf{C}_i = \mathbf{c}_i) = \mathrm{P}_{\sigma_i}(X_i \in \mathcal{F} | \mathbf{C}_i = \mathbf{c}_i).$$

When the user employs an ND strategy, the reported data X_i is independent from her personal signal S_i. This strategy does not disclose her private information and sets her privacy cost to 0, i.e., $\zeta_i(\sigma_i, c_i) = 0$. In general, her reported data still depends on her group signals C_i and is correlated with the underlying state W. Next we define the symmetric randomized response strategies as follows.

Definition 3 (SR Strategy): Given $\mathbf{C}_i = \mathbf{c}_i$, $\sigma_i([\cdot \;\; c_i])$ is called a symmetric randomized response (SR) strategy, if it satisfies the following conditions:

$$\begin{aligned}
&(1)\;\; \mathrm{P}_{\sigma_i}(X_i = \perp | S_i = 1, \mathbf{C}_i = \mathbf{c}_i) \\
&\qquad = \mathrm{P}_{\sigma_i}(X_i = \perp | S_i = 0, \mathbf{C}_i = \mathbf{c}_i) = 0, \\
&(2)\;\; \mathrm{P}_{\sigma_i}(X_i = 1 | S_i = 1, \mathbf{C}_i = \mathbf{c}_i) \\
&\qquad = \mathrm{P}_{\sigma_i}(X_i = 0 | S_i = 0, \mathbf{C}_i = \mathbf{c}_i).
\end{aligned}$$

When an SR strategy is played by the user, both the private signal S_i and the group signal C_i are correlated with the reported data X_i. We have the following lemma characterizing the Bayesian-Nash equilibria strategies (the proof can be found in Appendix A).

Lemma 1: For any non-negative payment mechanism, in every Bayesian-Nash equilibrium, the data reporting strategies are in the form of either a symmetric randomized response strategy or a non-disclosive strategy.

Note that if a user plays with a strategy where her reported data X_i is independent from both her personal data S_i and her group signal C_i, then X_i is pure noise. As a result, it is *uninformative* for the data collector, and it is a degenerated form of ND strategies. We remark that Lemma 1 is a generalization of [1, Lemma 1]. When a user i does not have any friends ($d_i = 0$), ND strategies reduce to be uninformative. On the other hand, in the presence of social learning ($d_i > 0$), ND strategies can be informative and positively contribute to the data collector's information elicitation.

IV. Data Collector's Payment Mechanism

The primary objective of the data collector is to estimate the state W from users' reported data $\mathbf{X}$. The hypothesis testing problem of the data collector can be stated as

$$\mathcal{H}_0 : W = 0, \quad \mathcal{H}_1 : W = 1. \tag{5}$$

The estimation of W from users' reported data $\mathbf{X}$ is viable only if there exists BNE strategies in which $\mathbf{X}$ is informative about W. The purpose of this section is to design the payment mechanism $\mathbf{R}$ in which each user can form her data reporting strategies, and the data collector can estimate W based on the reported data with minimum payment while achieving a given accuracy target.

In the presence of social learning, in addition to her personal signal, each user obtains noisy group signals through social interactions. Therefore, it is plausible to view each user as a *local data curator* who processes the data available to her and reports it to the data collector who acts as a *fusion center*. When users i and j are friends, C_{ij}, a noisy version of j's personal signal, is a component of T_i (and C_{ji} is a component of T_j). Further, if i and j are not friends but they have a common friend ℓ, then both i and j have noisy copies of S_ℓ in their type vectors. Consequently, given W, X_i and X_j can be correlated if i and j are friends or they have a common friend. To sum up, the optimal design hinges heavily upon the user types which are correlated across users given the true state and involves combinatorial optimization, and hence is very challenging to attain for the general case with finite N. To tackle this difficult task, firstly, we will focus on the problem at the local level by considering a hypothetical genie-aided payment mechanism and then study the optimal data reporting strategy in the asymptotic regime of N (for the sake of tractability). We present in Sections IV-C and IV-D the desired payment mechanism using peer prediction and the BNE strategies accordingly. Finally, in Section IV-E, we evaluate the total expected payment.

A. Users as Local Data Curators

We first study a hypothetical scenario where the data collector has access to the realization of the underlying state W and we discuss how we can design a *genie-aided* payment mechanism, to incentivize the users for informative data reporting. While it may sound vacuous as the data collector's main purpose of rewarding users for their reported data is to estimate the state W, it will become clear that this hypothetical scenario gives insight into the payment mechanism design where the data collector utilizes his estimation about the underlying state to determine the amount of reward for each user.

We consider a genie-aided payment mechanism $\mathbf{R}^{\mathrm{g}} : \mathcal{X}^N\{0, 1\} \to \mathbb{R}^N$, such that

$$R_i^{\mathrm{g}}((x_i, \mathbf{x}_{-i}), w) = R^{\mathrm{g}}(x_i, w).$$

Observe that the payments users receive do not depend on reported data of other users and they receive the same payment if their reported data is the same, i.e., $R_i(x_i, w) = R_j(x_j, w)$ if $x_i = x_j$. The expected payment of user i at strategy σ_i is given by

$$\begin{aligned}
\mathbb{E}_{\sigma_i}\big[R_i^{\mathrm{g}}(X_i, W) | T_i = t_i\big] = \sum_{w, x_i} & R^{\mathrm{g}}(x_i, w) \mathrm{P}(W = w | T_i = t_i) \\
& \times \mathrm{P}_{\sigma_i}(X_i = x_i | T_i = t_i),
\end{aligned}$$

where we have used the fact that X_i is independent from W given T_i. When $R^{\mathrm{g}}(0, 0) \geq R^{\mathrm{g}}(1, 0)$ and $R^{\mathrm{g}}(1, 1) \geq R^{\mathrm{g}}(0, 1)$, the expected payment of the user is maximized if X_i is selected as follows:

$$\frac{\mathrm{P}(T_i = t_i | W = 1)}{\mathrm{P}(T_i = t_i | W = 0)} \underset{X_i = 0}{\overset{X_i = 1}{\gtrless}} \frac{\mathrm{P}_W(0)\, R^{\mathrm{g}}(0, 0) - R^{\mathrm{g}}(1, 0)}{\mathrm{P}_W(1)\, R^{\mathrm{g}}(1, 1) - R^{\mathrm{g}}(0, 1)}.$$

Let $R^{\mathrm{g}}(1, 0) = R^{\mathrm{g}}(0, 1) = 0$. Further, setting the left-hand side of the above inequality to 1 renders the genie-aided payment

mechanism, where $Z^{\mathbf{g}} > 0$ is a design parameter:

$$R^{\mathbf{g}}(x_i, w) = \begin{cases} Z^{\mathbf{g}}/\mathrm{P}_W(0) & \text{if } x_i = w = 0, \\ Z^{\mathbf{g}}/\mathrm{P}_W(1) & \text{if } x_i = w = 1, \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

Consequently, in $\mathbf{R}^{\mathbf{g}}$, the expected payment of the user is maximized when their data reporting strategy is the maximum-likelihood (ML) decision rule:

$$\begin{aligned} \hat{W}_i^{\mathrm{ML}}(t_i) &= \arg \max_{w \in \{0,1\}} \mathrm{P}(T_i = t_i | W = w) \\ &= \arg \max_{w \in \{0,1\}} \mathrm{P}(S_i = s_i, C_i = c_i | W = w) \end{aligned}$$

For the data collector, the advantage of the ML rule is that the users do not use the prior distribution $\mathrm{P}_W(\cdot)$ information, which is already common knowledge. Thus, in this genie-aided scenario, the mechanism design problem reduces to a decentralized detection problem in which each user acts as a local decision maker.

Next, we explicitly state the ML rule as a function of S_i and $\mathbf{C}_i$. For convenience, denote the sum of the group signals by r.v. F_i and its realization by f_i:

$$F_i = \sum_{j \in G_i} C_{ij} \quad \text{and} \quad f_i = \sum_{j \in G_i} c_{ij}.$$

After some algebra, we have that

$$\frac{\theta_0^{s_i}(1-\theta_0)^{1-s_i}}{\theta_0^{1-s_i}(1-\theta_0)^{s_i}} \frac{\theta_1^{f_i}(1-\theta_1)^{d_i-f_i}}{\theta_1^{d_i-f_i}(1-\theta_1)^{f_i}} \underset{\hat{W}_i^{\mathrm{ML}}(t_i)=0}{\overset{\hat{W}_i^{\mathrm{ML}}(t_i)=1}{\gtrless}} 1$$

and

$$f_i \underset{\hat{W}_i^{\mathrm{ML}}(t_i)=0}{\overset{\hat{W}_i^{\mathrm{ML}}(t_i)=1}{\gtrless}} \frac{d_i}{2} - (2s_i - 1)\bar{A}, \text{ where } \bar{A} = \frac{1}{2} \frac{\log \frac{\theta_0}{1-\theta_0}}{\log \frac{\theta_1}{1-\theta_1}}. \tag{7}$$

It follows from (7) that

$$\hat{W}_i^{\mathrm{ML}}(s_i, f_i) = \begin{cases} 1, & \text{if } f_i > \frac{1}{2}d_i + \bar{A}, \\ 0, & \text{if } f_i < \frac{1}{2}d_i - \bar{A}, \\ s_i, & \text{otherwise.} \end{cases}$$

When $|f_i - \frac{d_i}{2}| > \bar{A}$, the ML rule reduces to reporting the *majority* bit of the group signals. This is an ND strategy and it incurs 0 privacy cost on the user. Therefore, the ML rule is a BNE strategy for the user when $|f_i - \frac{d_i}{2}| > \bar{A}$. When $|f_i - \frac{d_i}{2}| \leq \bar{A}$, the ML rule is directly reporting the personal signal which cannot be a BNE strategy because its privacy level is infinity. Consequently, the user has two options: 1. The user can send the majority bit of the group signals. 2. The user can employ an SR strategy and send a privacy-preserved version of S_i. Next, we define the majority voting-based data reporting strategies, denoted by σ_i^*.

Definition 4: The majority voting (MV)-based data reporting, σ_i^* has the following form:

$$\mathrm{P}_{\sigma_i^*}(X_i = 1 | T_i = t_i)$$

$$= \begin{cases} 1 & \text{if } f_i > \frac{d_i}{2} + \tau_1, \quad \textbf{(ND)} \\ 0 & \text{if } f_i < \frac{d_i}{2} - \tau_0, \quad \textbf{(ND)} \\ \frac{e^{\xi(f_i)}}{1+e^{\xi(f_i)}} & \text{if } f_i \in \left[\frac{d_i}{2} - \tau_0, \frac{d_i}{2} + \tau_1\right], \ s_i = 1, \quad \textbf{(SR)} \\ \frac{1}{1+e^{\xi(f_i)}} & \text{if } f_i \in \left[\frac{d_i}{2} - \tau_0, \frac{d_i}{2} + \tau_1\right], \ s_i = 0, \quad \textbf{(SR)} \end{cases}$$

$$\mathrm{P}_{\sigma_i^*}(X_i = 0 | T_i = t_i) = 1 - \mathrm{P}_{\sigma_i^*}(X_i = 1 | T_i = t_i),$$

where $0 \leq \tau_0, \tau_1 \leq \bar{A}$ and $\xi(f_i) \geq 0$.

With a little abuse of notation, $\zeta(\sigma_i^*, f_i)$ and $\zeta(\sigma_i^*, c_i)$ are used interchangeably for the privacy level of strategy σ_i^*. When $f_i \in [d_i/2 - \tau_0, d_i/2 + \tau_1]$, user i employs the SR strategy and we have $\zeta_i(\sigma_i^*, f_i) = \xi(f_i)$. When $f_i \notin [d_i/2 - \tau_0, d_i/2 + \tau_1]$, user i employs the ND strategy and we have $\zeta_i(\sigma_i^*, f_i) = 0$. The privacy level of the SR strategy, $\xi(f_i)$, and the thresholds τ_0 and τ_1, depend on $Z^{\mathbf{g}}$ from (6) and the system model parameters.

The majority voting-based data reporting strategy profile is denoted by $\sigma^* = \{\sigma_1^*, \ldots, \sigma_N^*\}$. Our next result states that σ^* is a BNE in the genie-aided payment mechanism $\mathbf{R}^{\mathbf{g}}$. Its proof is relegated to Appendix B.

Theorem 1: In the genie-aided payment mechanism $\mathbf{R}^{\mathbf{g}}$ (6), the majority voting-based data reporting strategy profile σ^* is a BNE.

In the next subsection, we analyze how the data collector can estimate the underlying state from users' reported data $\mathbf{X}$. In Section IV-C, building on the genie-aided mechanism $\mathbf{R}^{\mathbf{g}}$, we devise a peer-prediction-based payment mechanism $\bar{\mathbf{R}}$ where the data collector obtains the estimate of W from the users' reported data. In Section IV-D, we present the exact details of the BNE σ^*, and in particular, τ_0, τ_1 and $\xi(f_i)$, are determined accordingly.

B. Data Collector as Fusion Center

Recall that the objective of the data collector is to estimate the underlying state W from the users' reported data $\mathbf{X}$. The conditional distributions of the reported data profile $\mathbf{X}$, given the underlying state W, are dictated by the user data reporting strategies. For a given strategy profile σ, we can restate the binary hypothesis testing problem (5) as follows:

$$\mathcal{H}_0 : \mathbf{X} \sim \mathrm{P}_\sigma(\mathbf{X} = \mathbf{x} | W = 0), \ \mathcal{H}_1 : \mathbf{X} \sim \mathrm{P}_\sigma(\mathbf{X} = \mathbf{x} | W = 1).$$

The data collector employs the maximum a posteriori (MAP) decision rule, denoted by $\hat{W}_\sigma(\mathbf{x})$, in order to minimize the probability of error of the hypothesis testing problem:

$$\Lambda_\sigma(\mathbf{x}) := \frac{\mathrm{P}_\sigma(\mathbf{X} = \mathbf{x} | W = 1)}{\mathrm{P}_\sigma(\mathbf{X} = \mathbf{x} | W = 0)} \underset{\hat{W}_\sigma(\mathbf{x})=0}{\overset{\hat{W}_\sigma(\mathbf{x})=1}{\gtrless}} \frac{\mathrm{P}_W(0)}{\mathrm{P}_W(1)} \tag{8}$$

In general, X_i and X_j are correlated given W, if user i and j are friends or they have common friends owing to the social learning among the users. The closed-form evaluation of $\Lambda_\sigma(\mathbf{X})$ is often intractable for dependent observations. Therefore, in this study, we concentrate on symmetric data reporting strategies such that $\sigma_i(\cdot) = \sigma_j(\cdot)$ if the realizations of the type vectors of i and j are equal. For example, the majority voting-based data reporting strategy profile, $\sigma^*(\cdot)$, is a symmetric data reporting strategy profile. In what follows, we present two lemmas to study $\Lambda_\sigma(\mathbf{x})$ in the asymptotic regime of N. For convenience, let $\mathcal{I}_N$ be the collection of all permutations on the set indices $\mathcal{I} = \{1, 2, \ldots, N\}$. Then, for $\pi \in \mathcal{I}_N$, $\mathbf{x}_\pi = [x_1 \ldots x_N]_\pi$ denotes the permuted sequence $[x_{\pi(1)} \ldots x_{\pi(N)}]$. Our next result shows that the order of the reported data is irrelevant for the data collector's

binary hypothesis testing problem. Its proof is relegated to Appendix C.

Lemma 2: For every symmetric strategy profile σ, $w \in \{0,1\}$ and $\pi \in \mathcal{I}_N$, we have

$$\mathrm{P}_\sigma(\mathbf{X}=\mathbf{x}|W=w)=\mathrm{P}_\sigma(\mathbf{X}=\mathbf{x}_\pi|W=w).$$

From Lemma 2, it follows that $\hat{W}_\sigma(\mathbf{x})$ in (8) depends on $\sum_{i=1}^N x_i$ when σ is a symmetric strategy profile. To characterize the asymptotic statistics of $\sum_i X_i$, we employ a Central Limit Theorem (CLT) for dependence graphs [35]. For any symmetric strategy profile σ, we define $\mu_w(\sigma)$ as the conditional mean of X_i given $W=w$ with $w \in \{0,1\}$:

$$\begin{aligned}\mu_1(\sigma) &:= \mathrm{P}_\sigma(X_i=1|W=1),\\ \mu_0(\sigma) &:= \mathrm{P}_\sigma(X_i=1|W=0).\end{aligned} \tag{9}$$

Recall that $\mathcal{E}_{ij}=1$ if there is a social tie between i and j, otherwise $\mathcal{E}_{ij}=0$. Similarly, $B_{ij}=1$ if i and j have a common friend, otherwise $B_{ij}=0$. For convenience, we define ς_w and $\tilde{\varsigma}_w$ for $w \in \{0,1\}$, as follows:

$$\varsigma_w(\sigma) := \mathrm{P}_\sigma\big(X_i=w, X_j=w|W=w, B_{ij}=0, \mathcal{E}_{ij}=1\big), \tag{10a}$$

$$\tilde{\varsigma}_w(\sigma) := \mathrm{P}_\sigma\big(X_i=w, X_j=w|W=w, B_{ij}=1, \mathcal{E}_{ij}=0\big), \tag{10b}$$

In the rest of the article, for purposes of brevity, we drop the dependency of $\mu_w(\sigma)$, $\varsigma_w(\sigma)$, and $\tilde{\varsigma}_w(\sigma)$ on σ when it is clear from the context. To use the CLT for dependence graphs, the degree distribution ρ_d must meet the following sparsity criteria:

Assumption 1: Maximal degree $D_{\max} = o(N^{1/4})$ and $\mathbb{E}[D^{2+\Delta}]<\infty$ for some $\Delta>0$.

We have the following result on the asymptotics of $\sum_{i=1}^N X_i$ as $N \to \infty$.

Lemma 3: Under Assumption 1, conditioned on $W=w$, for a symmetric data reporting strategy profile σ, $\frac{\sum_{i=1}^N X_i - N\mu_w}{\sqrt{N\kappa_w}}$ converges in distribution to a standard normal random variable as $N \to \infty$, with

$$\kappa_1(\sigma) := \mu_1-\mu_1^2+\mathbb{E}[D](\varsigma_1-\tilde{\varsigma}_1)+\mathbb{E}\Big[D^2\Big]\Big(\tilde{\varsigma}_1-\mu_1^2\Big), \tag{11a}$$

$$\begin{aligned}\kappa_0(\sigma) := {}& \mu_0(1-\mu_0)+\mathbb{E}[D](\varsigma_0-\tilde{\varsigma}_0)\\ &+\mathbb{E}\Big[D^2\Big]\Big(\tilde{\varsigma}_0-(1-\mu_0)^2\Big).\end{aligned} \tag{11b}$$

The proof of this lemma is relegated to Appendix D.

Appealing to Lemmas 2 and 3, for large N, the MAP Decision rule $\hat{W}_{\sigma^*}(\mathbf{x})$ can be approximated as follows:

$$\frac{1}{\kappa_0}\left(\mu_0-\frac{\sum x_i}{N}\right)^2-\frac{1}{\kappa_1}\left(\mu_1-\frac{\sum x_i}{N}\right)^2 \underset{\hat{W}_{\sigma^*}(\mathbf{x})=0}{\overset{\hat{W}_{\sigma^*}(\mathbf{x})=1}{\gtrless}} \frac{2}{N}\ln\left(\sqrt{\frac{\kappa_1}{\kappa_0}}\frac{\mathrm{P}_W(0)}{\mathrm{P}_W(1)}\right). \tag{12}$$

The data collector aims to minimize the total payment under the budget of MAP detector error rate constraint. Let $\mathcal{R}(\sigma)$ denote the set of non-negative payment mechanisms in which σ is a BNE. Then, the mechanism design problem for the data collector can be formulated as follows:

$$\begin{aligned}\min_{\mathbf{R}\in\mathcal{R}(\sigma)} \ & \sum_{i=1}^N \mathbb{E}_\sigma[R_i(\mathbf{X})],\\ \text{s.t.} \ & \mathbb{E}_\sigma\Big[\mathrm{P}\Big(\hat{W}_\sigma(\mathbf{X})\neq W\Big)\Big]\le \mathcal{P}_e,\end{aligned} \tag{13}$$

where the maximum allowable error is represented by $\mathcal{P}_e$. Next, we design the peer-prediction-based payment mechanism. In Section IV-D, we determine the informative BNE strategies in the designed mechanism. In Section IV-E, we evaluate the total expected payment. In Section V, we will revisit the mechanism design problem and derive bounds on the minimum total payment required to achieve a given level of state estimation accuracy.

C. Payment Mechanism Design

Building on the genie-aided mechanism $\mathbf{R}^{\mathrm{g}}$, next we turn our attention to the design of a peer-prediction-based payment mechanism $\tilde{\mathbf{R}}$, where the data collector obtains the estimate of W from the users' reported data. In particular, we use majority voting as an effective aggregation method [1], [8], [36], [37] to obtain informative reported data from the users. By rewarding the users whose reported data is in agreement with the other users' reported data, this payment mechanism incentivizes the users to participate and report informatively using their personal and group signals. More specifically, we have the following payment mechanism $\tilde{\mathbf{R}}(\mathbf{X})$:

1) Each user reports her data, and the data collector counts the number of participants n excluding the users with "non-participation". For non-participating users, the payment is zero.
2) If $n=1$, the data collector pays zero to this participant. Otherwise, for each participating user i, the data collector computes the majority bit of the other participants' reported data:
$$M_{-i}=\begin{cases}1 & \text{if } \sum_{j:x_j\neq\perp, j\neq i} x_j \ge \left\lfloor\frac{n-1}{2}\right\rfloor+1;\\ 0 & \text{otherwise.}\end{cases}$$
3) Compute the payment for user i:
$$\tilde{R}_i(1,\mathbf{x}_{-i})=Z_1M_{-i}, \quad \tilde{R}_i(0,\mathbf{x}_{-i})=Z_0(1-M_{-i}), \tag{14}$$
where Z_0 and Z_1 are design parameters to be determined by the data collector.

In the genie-aided scenario, the payment mechanism $\mathbf{R}^{\mathrm{g}}$ is designed based on the hypothetical case where the underlying state W is given. The rationale behind the proposed payment mechanism $\tilde{\mathbf{R}}$ in (15) above is that the data collector obtains the estimate of W from the noisy user reports and utilizes it in the payment mechanism. Along the same line as in the genie-aided mechanism, each user first estimates the underlying state W based on her type t_i. The next key step lies in the computation of the probability of a user being consistent with the majority at the BNE strategy profile σ^*:

$$\beta_w=\mathrm{P}_{\sigma^*}(M_{-i}=w|W=w), \tag{15}$$

where $w \in \{0, 1\}$. Clearly, when the number of users is large, the asymptotic statistics of $\sum_j X_j$ is the same as the asymptotic statistics of $\sum_{j \in -i} X_j$. Therefore, β_0 and β_1 can be approximated using Lemma 3.

Based on the hypothetical genie-aided payment mechanism $\mathbf{R}^g$ defined in (6), we obtain the design parameters for the payment mechanism defined in (14) as follows:

$$Z_0 = Z\,\frac{\mathrm{P}_W(1)\beta_1 + \mathrm{P}_W(0)(1-\beta_0)}{(\beta_0+\beta_1-1)\mathrm{P}_W(1)\mathrm{P}_W(0)},$$
$$Z_1 = Z\,\frac{\mathrm{P}_W(1)(1-\beta_1) + \mathrm{P}_W(0)\beta_0}{(\beta_0+\beta_1-1)\mathrm{P}_W(1)\mathrm{P}_W(0)}. \tag{16}$$

where $Z > 0$ is a design parameter. For the degenerate case in which the data collector obtains the estimate of W with no error, we have that $\beta_1 = \beta_0 = 1$, indicating that $\tilde{\mathbf{R}}$ reduces to the genie-aided mechanism introduced in (6). Theorem 2 reveals that there exists a MV-based BNE when the data collector employs the payment mechanism $\tilde{\mathbf{R}}$.

Theorem 2: In the peer-prediction-based payment mechanism $\tilde{\mathbf{R}}$ (16), the majority voting-based data reporting strategy profile σ^* is a BNE.

The proof of Theorem 2 is given in Appendix E.[9] For brevity, we define $p(s_i, f_i)$ and $q(s_i, f_i)$ as follows:

$$p(s_i, f_i) = \mathrm{P}_{\sigma_i^*}(X_i = 1 | S_i = s_i, F_i = f_i),$$
$$q(s_i, f_i) = \mathrm{P}_{\sigma_i^*}(X_i = 0 | S_i = s_i, F_i = f_i). \tag{17}$$

Theorem 2 establishes the existence of a MV-based BNE under the payment mechanism $\tilde{\mathbf{R}}$. To complete the design of the payment mechanism, it remains to characterize $p(s_i, f_i)$ and $q(s_i, f_i)$ with respect to the mechanism design parameter Z. For this purpose, we express Z in terms of an auxiliary parameter $\epsilon \geq 0$ as follows:

$$Z = g'(\epsilon)\frac{(e^\epsilon+1)^2}{2e^\epsilon(2\theta_0-1)}. \tag{18}$$

It is clear that Z is continuously differentiable, increasing and nonnegative in ϵ which corresponds to the privacy level of strategy σ^* when there is a tie in the group signal, i.e., $\xi(d_i/2) = \epsilon$:

$$p\left(s_i, \frac{d_i}{2}\right) = \frac{e^\epsilon}{1+e^\epsilon} \quad \text{and} \quad q\left(s_i, \frac{d_i}{2}\right) = \frac{1}{1+e^\epsilon}.$$

By increasing ϵ, the data collector can increase the accuracy of the reported data at σ^* albeit the higher total payment.

[9]It is worth noting that in the payment mechanism $\tilde{\mathbf{R}}$, σ^* is not the only equilibrium. At σ^*, no user can gain by playing uninformative when other users employ the MV rule. However, in $\tilde{\mathbf{R}}$, uninformative equilibria also exist, as it is the case in many peer-prediction and information elicitation mechanisms [11], [12], [15], [38]. Indeed, the equilibria of output agreement based peer-prediction mechanisms cannot avoid having an uninformative equilibria [12], [17]. Interested readers can find detailed discussions for different colluding scenarios in the information elicitation mechanism in [13]. One set of such equilibria is that the users form lying coalitions and collude to report the same uninformative data, e.g., uninformative pure-strategy equilibria such that all agents coordinate to simply report $X_i = 1$ or $X_i = 0$. One set of such equilibria is that the users form lying coalitions and collude to report the same uninformative data. However, we caution that the social learning model does not imply any cooperation and communication among friends.

D. Data Reporting Strategies at BNE

In this subsection, we put forward an algorithm to find $p(s_i, f_i)$ and $q(s_i, f_i)$. Recall that, $\xi(f_i)$ corresponds to the privacy level of σ_i^* when $f_i \in [\frac{d_i}{2} - \tau_0, \frac{d_i}{2} + \tau_1]$. Given ϵ and f_i, user i can determine $\xi(f_i)$ as follows:

$$\xi(f_i) = \epsilon + (\mathrm{P}_W(1) - \mathrm{P}_W(0))\left(\frac{e^\epsilon - 1}{e^\epsilon + 1}\right. \tag{19}$$
$$\left. + 2\frac{g''(\epsilon)}{g'(\epsilon)}\frac{\mathrm{P}_W(0)\theta_1^{d_i-2f_i} + \mathrm{P}_W(1)(1-\theta_1)^{d_i-2f_i}}{\theta_1^{d_i-2f_i}(1-\theta_1)^{d_i-2f_i}}\right)^{-1}.$$

Next, user i needs to compare the expected utilities of playing the ND and the SR strategies in order to decide upon between them. For this purpose, user i evaluates $\Upsilon_0[\xi(f_i)]$ and $\Upsilon_1[\xi(f_i)]$:

$$\Upsilon_\ell\big[\xi(f_i)\big] = \begin{cases} 0, & \text{if } A_\ell\big[\xi(f_i)\big] \leq 0, \\ A_\ell\big[\xi(f_i)\big], & \text{if } A_\ell\big[\xi(f_i)\big] \in (0, \bar{A}), \\ \bar{A}, & \text{if } A_\ell\big[\xi(f_i)\big] \geq \bar{A}, \end{cases}$$
$$\text{with } \bar{A} = \frac{1}{2}\frac{\log\frac{\theta_0}{1-\theta_0}}{\log\frac{\theta_1}{1-\theta_1}}, \tag{20}$$

and

$$A_\ell\big[\xi(f_i)\big] = \frac{2\ell-1}{2\ln\left(\frac{\theta_1}{1-\theta_1}\right)}\ln$$
$$\times\left(\frac{e^{\xi(f_i)}(1-\theta_0)+\theta_0+\mathrm{P}_W(\ell)B\big[\xi(f_i)\big]}{e^{\xi(f_i)}\theta_0+1-\theta_0-\mathrm{P}_W(1-\ell)B\big[\xi(f_i)\big]}\right),$$
$$\text{where } B\big[\xi(f_i)\big] = (2\theta_0-1)2e^\epsilon\frac{g(\xi(f_i))}{g'(\epsilon)}\frac{\left(e^{\xi(f_i)}+1\right)}{(e^\epsilon+1)^2}.$$

If $f_i \leq d_i/2 - \Upsilon_0[\xi(f_i)]$, then the user plays the ND strategy and the reported data $X_i = 0$. Similarly, if $f_i \geq d_i/2 - \Upsilon_0[\xi(f_i)]$. then the user plays the ND strategy and $X_i = 1$. Otherwise, the user plays the SR strategy with privacy level $\xi(f_i)$. We detail the procedure to find the MV-based BNE strategy σ^* in Algorithm 1. The following result formalizes this argument. The proof of Proposition 1 can be found in Appendix F.

Proposition 1: The BNE strategy profile σ^* can be found by using Algorithm 1.

E. A Closer Look at Data Reporting Strategies: Two Special Cases

To get a more concrete sense, in what follows we study two special cases where Algorithm 1 can be further simplified.

1) The Case With Noiseless Group Signals: In general, there is a discrepancy between the quality of S_i and C_i: $\theta_1 < \theta_0$. We first consider the extreme case in which the group signals are noiseless, $\alpha = 0$ and $\theta_0 = \theta_1$. In this case, $\bar{A} = 0.5$ and Algorithm 1 reduces to a *simple majority* rule: A simple majority in C_i suffices to determine whether to play the ND strategy or the SR strategy.

Corollary 1: For the case with noiseless group signals, the BNE strategy profile σ^* has the following form:

Algorithm 1: MV-Based BNE Strategy σ^*

Input: Type t_i, number of friends d_i.
Output: $p(s_i, f_i)$ and $q(s_i, f_i)$
Determine $\xi(f_i)$, $\Upsilon_0[\xi(f_i)]$ and $\Upsilon_1[\xi(f_i)]$ using Equations 19 and 20.
if $f_i \leq d_i/2 - \Upsilon_0[\xi(f_i)]$ **then** $p(0, f_i) = p(1, f_i) = 0$ and $q(0, f_i) = q(1, f_i) = 1$ **(ND)**
else if $f_i \geq d_i/2 + \Upsilon_1[\xi(f_i)]$ **then** $p(0, f_i) = p(1, f_i) = 1$ and $q(1, f_i) = q(0, f_i) = 0$ **(ND)**
else $p(1, f_i) = q(0, f_i) = e^{\xi(f_i)}/(1 + e^{\xi(f_i)})$ and $p(0, f_i) = q(1, f_i) = 1/(1 + e^{\xi(f_i)})$ **(SR)**

$$p(s_i, f_i) = \mathrm{P}_{\sigma_i^*}(X_i = 1 | S_i = s_i, F_i = f_i) = \begin{cases} 1 & \text{if } f_i > \frac{d_i}{2}, \\ 0 & \text{if } f_i < \frac{d_i}{2}, \\ \frac{s_i e^\epsilon + 1 - s_i}{1 + e^\epsilon} & \text{if } f_i = \frac{d_i}{2}, \end{cases} \quad \text{and } q(s_i, f_i) = 1 - p(s_i, f_i).$$

According to Corollary 1, if d_i is even, the user plays the SR strategy with privacy level ϵ if $f_i = d_i/2$. Note that, by increasing ϵ, the data collector raises the payment per user in order to collect more accurate reported from the users who play the SR strategy. If $f_i \neq d_i/2$ the user plays the ND strategy and the reported data X_i is equal to the majority bit within the group signal C_i. If d_i is odd, the user never plays the SR strategy.

2) The Case With Equal Priors: When $\mathrm{P}_W(1) = \mathrm{P}_W(0) = 0.5$, it is clear from (19) that $\xi(f_i) = \epsilon$, for every $f_i \in \{0, 1, \ldots, d_i\}$. It directly follows from Algorithm 1 that the thresholds of σ^* have the same value and can be found as follows:

$$\tau(\epsilon) = \begin{cases} 0, & \text{if } A \leq 0 \\ A, & \text{if } A \in (0, \bar{A}), \\ \bar{A}, & \text{if } A \geq \bar{A} \end{cases}$$

$$A = \frac{\ln\left(\frac{e^\epsilon\left[e^\epsilon\theta_0 + 1 - (2\theta_0 - 1)g(\epsilon)/g'(\epsilon)\right] + 1 - \theta_0}{e^\epsilon\left[e^\epsilon(1-\theta_0) + 1 + (2\theta_0 - 1)g(\epsilon)/g'(\epsilon)\right] + \theta_0}\right)}{2\ln\left(\frac{\theta_1}{1-\theta_1}\right)}. \quad (21)$$

Consequently, our next result determines the MV-based data reporting strategies.

Corollary 2: For the case with equal priors, the BNE strategy profile σ^* has the following form:

$$p(s_i, f_i) = \mathrm{P}_{\sigma_i^*}(X_i = 1 | S_i = s_i, F_i = f_i) = \begin{cases} 1 & \text{if } f_i > \frac{d_i}{2} + \tau(\epsilon), \\ 0 & \text{if } f_i < \frac{d_i}{2} - \tau(\epsilon), \\ \frac{s_i e^\epsilon + 1 - s_i}{1 + e^\epsilon} & \text{if } f_i \in \left[\frac{d_i}{2} - \tau(\epsilon), \frac{d_i}{2} + \tau(\epsilon)\right], \end{cases}$$
$$q(s_i, f_i) = \mathrm{P}_{\sigma_i^*}(X_i = 0 | S_i = s_i, F_i = f_i) = 1 - p(s_i, f_i).$$

The MV-based data reporting strategies depend heavily on the crossover probability α and the payment mechanism parameter ϵ. Recall that the quality of group signals is defined as $\theta_1 = \theta_0 - \alpha(2\theta_0 - 1)$. It is clear from (21) that τ increases with α, the noise level in the group signals: The user plays the SR strategy more often at the BNE as α increases since the group signals become less informative. Note that α is a system model parameter and it is not under the control of the users and the data collector. By choosing the payment mechanism parameter ϵ, the data collector can control the privacy level of the users who play the SR strategy. As ϵ increases, the users play the SR strategy more often and inject less noise on the reported data at the BNE. For conciseness, in the remainder of this section, we suppress the explicit dependence of τ on ϵ.

Next we study the computation of β_0 and β_1 defined in (15), under equal priors assumption. We first need to define several terms. Define $\gamma(k; d, p)$ and $\Gamma(k, \ell; d, p)$ corresponding to a Binomial distribution with parameters m (number of trials) and n (probability of success) as follows:

$$\gamma(k; m, n) = \begin{cases} \binom{d}{k} n^k (1-n)^{m-k} & \text{if } k \in \{0, 1, \ldots, n\}, \\ 0 & \text{otherwise;} \end{cases},$$

$$\Gamma(k, \ell; m, n) = \sum_{i=\lceil k \rceil}^{\lfloor \ell \rfloor} \gamma(i; m, n), \quad (22)$$

where $\lceil k \rceil := \min\{m \in \mathbb{Z} : m \geq k\}$ and $\lfloor \ell \rfloor := \max\{n \in \mathbb{Z} : n \leq \ell\}$. Recall that, ρ_d is the degree distribution of the social learning graph. We define $\tilde{\rho}$ as follows:

$$\tilde{\rho}_d := \mathrm{P}(D_i = d | D_i > 0) = \begin{cases} 0, & \text{if } d = 0; \\ \rho_d/(1 - \rho_0), & \text{else.} \end{cases} \quad (23)$$

Note that, $\tilde{\rho}$ is well defined unless $\rho_0 = 1$ which corresponds to the case there is no social learning among the users. In the rest of the article, we use the subscript notation $\mathbb{E}_{\tilde{\rho}}$ when we use $\tilde{\rho}$ for the expectation of the user degrees. The following results determines μ_w and κ_w, which are defined in (9) and (11), at the MV BNE σ^*.

Proposition 2: For the case with equal priors, $\kappa_w(\sigma^*)$ and $\mu_w(\sigma^*)$ for $w \in \{0, 1\}$ are found as follows:

$$\mu_1(\sigma^*) = 1 - \mu_0(\sigma^*) = \mathbb{E}\left[\Gamma\left(\left\lfloor \tfrac{D}{2} + \tau + 1 \right\rfloor, D; D, \theta_1\right)\right] + \lambda(\epsilon)\Gamma\left(\tfrac{D}{2} - \tau, \tfrac{D}{2} + \tau; d, \theta_1\right), \quad (24)$$

$$\kappa_1(\sigma^*) = \kappa_0(\sigma^*) = \mu_1 - \mu_1^2 + \tilde{\Delta}\mathbb{E}\left[D^2\right] + \Delta\left(\mathbb{E}\left[D^2\right] - \mathbb{E}[D]\right), \quad (25)$$

where λ, Δ and $\tilde{\Delta}$ are defined as

$$\Delta := \frac{\theta_0(1-\theta_0)(1-2\alpha)}{e^\epsilon + 1} \times \mathbb{E}_{\tilde{\rho}}\left[\left(e^\epsilon(1-\theta_0) + \theta_0\right)\gamma\left(\left\lfloor \tfrac{D}{2} + \tau \right\rfloor; D - 1, \theta_1\right) + \left(\theta_0 e^\epsilon + 1 - \theta_0\right)\gamma\left(\left\lceil \tfrac{D}{2} - \tau - 1 \right\rceil; D - 1, \theta_1\right)\right]$$

$$\lambda(\epsilon) := \frac{\theta_0 e^\epsilon + 1 - \theta_0}{e^\epsilon + 1},$$

$$\tilde{\Delta} := \frac{\rho_0}{(1-\rho_0)^2}\left(\mu_1^2(2 - \rho_0) - 2\mu_1\lambda(\epsilon) + \rho_0\lambda^2(\epsilon)\right). \quad (26)$$

The proof of Proposition 2 is relegated to Appendix G.

Appealing to Lemma 3 and Proposition 2, we can compute $\beta_w = \mathrm{P}_{\sigma^*}(M_{-i} = w|W = w)$, for $w \in \{0, 1\}$ as follows:

$$\beta := \beta_0 = \beta_1 = \Phi\left(\sqrt{\frac{N-1}{\kappa_1(\sigma^*)}}\left(\mu_1(\sigma^*) - \frac{1}{2}\right)\right), \quad (27)$$

where Φ denotes the cumulative distribution function of the standard normal distribution. Consequently, for the case with equal priors, the payment mechanism parameters Z_0 and Z_1 can be found as follows:

$$Z_0 = Z_1 = \frac{g'(\epsilon)(e^\epsilon + 1)^2}{e^\epsilon(2\theta_0 - 1)(2\beta - 1)}.$$

It is clear that $\beta \to 1$ as N grows and the proposed payment mechanism $\tilde{\mathbf{R}}$ (14) boils down to the genie-aided payment mechanism $\mathbf{R}^{\mathrm{g}}$ (6).

Our next result determines the expected payment of the proposed payment mechanism $\tilde{\mathbf{R}}$:

Theorem 3: For the case with equal priors, the total expected payment at the BNE is the following:

$$\sum_{i=1}^{N} \mathbb{E}_{\sigma^*}\left[\tilde{R}_i(\mathbf{X})\right] = \frac{g'(\epsilon)(e^\epsilon + 1)^2 N}{e^\epsilon(2\theta_0 - 1)(2\beta - 1)}\left(1 - \beta + \frac{\mu_1(\sigma^*)}{2\beta - 1}\right).$$

The proof of Theorem 3 is relegated to Appendix H. In Section V, we will revisit the mechanism design problem and the payment vs. accuracy trade-off is further analyzed based on Theorem 3.

F. The Privacy of the Group Signals

In majority voting-based data reporting strategies, each user locally estimates the underlying state from the sum of the group signal. Any attack attempt which targets to learn a user's personal signal from her friends' reported data would require the exact knowledge of social learning graph, which is assumed to be not available to the data collector and users. This is a sensible assumption, because the social learning among privacy-aware users can take place in many different forms, including face-to-face meetings and over multiple online social media, and hence it is difficult for the adversary to obtain the social learning graph.

To get a more concrete sense, we consider a worst case scenario where there is an attacker who has the exact knowledge of the entire social learning graph and the attacker can observe the group signals the users receive from their friends. For this worst case scenario, we formally quantify the privacy leakage of the personal signal S_j, when the data reporting strategy σ_i is used by user i and $i \neq j$, as follows:

$$\eta_{ij}(\sigma_i) = \max_{\mathcal{F} \subseteq \{0,1,\perp\}, s_j \in \{0,1\}, w \in \{0,1\}} \times \ln\left(\frac{\mathrm{P}_{\sigma_i}(X_i \in \mathcal{F}|S_j = s_j, W = w)}{\mathrm{P}_{\sigma_i}(X_i \in \mathcal{F}|S_j = 1 - s_j, W = w)}\right), \quad (28)$$

where the convention $0/0 = 1$ is followed.[10] Intuitively, the smaller $\eta_{ij}(\sigma_i)$ is, the more indistinguishable x_i from s_j, and hence the less privacy leakage is. For any data reporting strategy σ_i, if users i and j are not friends, X_i is conditionally independent from S_j given W and hence $\eta_{ij}(\sigma_i) = 0$. Next, we focus on the privacy leakage of group signals at the BNE strategy profile σ^*. For ease of exposition, we consider the case with equal priors. If users i and j are friends, it follows that

$$\eta_{ij}(\sigma_i^*) = \ln \times \frac{(1-\alpha)\mathrm{P}_{\sigma_i^*}(X_i = 1|C_{ij} = 1, W = 1) + \alpha\mathrm{P}_{\sigma_i^*}(X_i = 1|C_{ij} = 0, W = 1)}{(1-\alpha)\mathrm{P}_{\sigma_i^*}(X_i = 1|C_{ij} = 0, W = 1) + \alpha\mathrm{P}_{\sigma_i^*}(X_i = 1|C_{ij} = 1, W = 1)}.$$

For convenience, define $F_{i,-j} := F_i - C_{ij}$. In general, $\mathrm{P}_{\sigma_i^*}(X_i = 1|C_{ij} = 1, W = 1, D_i = d_i) = \mathrm{P}_{\sigma_i^*}(X_i = 1|C_{ij} = 0, W = 1, D_i = d_i)$ unless $F_{i,-j} = \lfloor \frac{d_i}{2} - \tau \rfloor$ or $F_{i,-j} = \lfloor \frac{d_i}{2} + \tau \rfloor$. Only in these two cases, the group signal C_{ij} can be considered as the "tie-breaking vote" and the attacker can gain some information about user j's personal signal. It is clear that the probability of this event decreases as the social learning among the users strengthens, i.e., d_i increases. The following result determines $\eta_{ij}(\sigma_i^*)$ as

$$\eta_{ij}(\sigma_i^*) = \ln \times \frac{\mathbb{E}_{\tilde{\rho}}\left[(1-\alpha)\left(\lambda\gamma\left(\left\lceil\frac{D+1}{2} - \tau\right\rceil + 1; D, \theta_1\right) + (1-\lambda)\gamma\left(\left\lfloor\frac{D+1}{2} + \tau\right\rfloor; D, \theta_1\right)\right) + \lambda\Gamma\left(\frac{D+1}{2} - \tau, \frac{D+1}{2} + \tau; D, \theta_1\right) + \Gamma\left(\left\lfloor\frac{D+1}{2} + \tau\right\rfloor, D; D, \theta_1\right)\right]}{\mathbb{E}_{\tilde{\rho}}\left[\alpha\left(\lambda\gamma\left(\left\lceil\frac{D+1}{2} - \tau\right\rceil + 1; D, \theta_1\right) + (1-\lambda)\gamma\left(\left\lfloor\frac{D+1}{2} + \tau\right\rfloor; D, \theta_1\right)\right) + \lambda\Gamma\left(\frac{D+1}{2} - \tau, \frac{D+1}{2} + \tau; D, \theta_1\right) + \Gamma\left(\left\lfloor\frac{D+1}{2} + \tau\right\rfloor, D; D, \theta_1\right)\right]}, \quad (29)$$

where $\Gamma(\cdot)$ and $\gamma(\cdot)$, $\mathbb{E}_{\tilde{\rho}}[\cdot] = \mathbb{E}_{\rho}[\cdot|D > 0]$ and λ are defined in (22), (23), and (26), respectively.

Note that $\eta_{ij}(\sigma_i^*)$ is a complicated function. In Section V-C3, we shall use numerical examples to illustrate the dependency of $\eta_{ij}(\sigma_i^*)$ on the payment mechanism and the system model parameters, which corroborate that *the privacy leakage is insignificant in most cases and it can get severe only if the following three conditions hold: 1) The noise level of the group signals is very low, 2) the users have very few friends and 3) the payment is very low.* It is clear from (29) that $\eta_{ij}(\sigma_i^*)$ decreases when the noise level of the group signals increases. It is natural for a user to share only a randomized version (with moderate crossover probability α) of her signal to protect her personal signals against privacy leakage. In the next section, we discuss the performance of the MV based data reporting strategies with a focus on the cases in which α is moderate or reasonably high.

V. The Impact of Social Learning

In this section, we analyze the impact of social learning on the trade-off between payment and accuracy and that between payment and privacy cost. We also present examples, using social learning graph models based on synthetic data and/or real-world data, to evaluate the performance of our proposed mechanisms.

A. Payment vs. Accuracy

The data collector aims to minimize the total payment while achieving a given accuracy target in estimating W. In particular, the accuracy is measured by the error rate of the MAP

[10] In this definition, we quantify how much the change of only user j's personal signal alters the probability of any reported data output of user i. If we do not condition on the underlying state W, the flip of s_j changes the posterior probabilities of other group signals since the personal signals are correlated through the underlying state W.

detector (8). Recall that the mechanism design problem for the data collector is defined in (13) as

$$\min_{\mathbf{R}\in\mathcal{R}(\sigma)} \sum_{i=1}^{N} \mathbb{E}_\sigma[R_i(\mathbf{X})], \quad \text{s.t.} \quad \mathbb{E}_\sigma\Big[\mathrm{P}\Big(\hat{W}_\sigma(\mathbf{X}) \neq W\Big)\Big] \leq \mathcal{P}_e,$$

where $\mathcal{R}(\sigma)$ denotes the set of non-negative payment mechanisms in which σ is a BNE and the maximum allowable error is represented by $\mathcal{P}_e$. It is known that in general it is difficult to characterize the error rate in closed-form at a given BNE strategy profile σ. Therefore, we measure the accuracy based on an information-theoretic metric which is closely associated to the error rate of the MAP decision rule as follows [39]:

$$\mathbb{E}_\sigma\Big[\mathrm{P}\Big(\hat{W}_\sigma(\mathbf{X}) \neq W\Big)\Big] \leq e^{-\mathcal{B}(\sigma)},$$

where $\mathcal{B}(\sigma)$ denotes the Bhattacharyya distance [40]

$$\mathcal{B}(\sigma) = -\ln \sum_{\mathbf{x}\in\mathcal{X}^N} \sqrt{\mathrm{P}_\sigma(\mathbf{X}=\mathbf{x}|W=1)\mathrm{P}_\sigma(\mathbf{X}=\mathbf{x}|W=0)}.$$

Thus, the mechanism design problem can be restated as follows:

$$\min_{\mathbf{R}\in\mathcal{R}(\sigma)} \sum_{i=1}^{N} \mathbb{E}_\sigma[R_i(\mathbf{X})], \quad \text{s.t.} \quad \mathcal{B}(\sigma) \geq -\ln \mathcal{P}_e. \tag{30}$$

Define $\mathcal{L}(\mathcal{P}_e)$ as the minimum total payment while satisfying the error rate constraint $\mathcal{P}_e$. Appealing to Lemma 2 and 3, we can simplify the expression for $\mathcal{B}(\sigma)$, for symmetric strategy profiles by approximating $\mathrm{P}_\sigma(\mathbf{X}=\mathbf{x}|W=w)$ as a Gaussian distribution for large N. Thus, $\mathcal{B}(\sigma)$ can be calculated explicitly as follows [39]:

$$\mathcal{B}(\sigma) = \frac{N}{4}\frac{(\mu_1(\sigma)-\mu_0(\sigma))^2}{\kappa_1(\sigma)+\kappa_0(\sigma)}.$$

B. Bounds on Payment

Our next result shows that, if the required estimation accuracy, in terms of $\mathcal{P}_e$, is loose, the total payment can be driven to be arbitrarily small by using the following non-disclosive strategy, denoted as σ^{nd}:

$$\mathrm{P}_{\sigma^{\mathrm{nd}}}(X_i=1|S_i=s_i,F_i=f_i) = \begin{cases} 1 & \text{if } f_i > d_i/2, \\ 0 & \text{if } f_i < d_i/2, \\ 0.5 & \text{else;} \end{cases}$$

$$\mathrm{P}_{\sigma^{\mathrm{nd}}}(X_i=0|S_i=s_i,F_i=f_i) = 1-\mathrm{P}_{\sigma^{\mathrm{nd}}} \times (X_i=0|S_i=s_i,F_i=f_i).$$

If there is a tie within her group signals, the user tosses a fair coin. It is clear that the above σ^{nd} is a specific form of MV-based data reporting strategies with $\tau_1=\tau_0=0$ and $\xi(f_i)=0$. From (9), it directly follows that

$$\mu_1\Big(\sigma^{\mathrm{nd}}\Big) = \mathbb{E}_\rho\big[\Gamma(\lfloor D/2+1\rfloor, D; D, \theta_1) + 0.5\gamma(D/2; D, \theta_1)\big],$$

$$\mu_0\Big(\sigma^{\mathrm{nd}}\Big) = \mathbb{E}_\rho\big[\Gamma(\lfloor D/2+1\rfloor, D; D, 1-\theta_1 + 0.5\gamma(D/2; D, 1-\theta_1)\big]. \tag{31}$$

It is clear that $\mu_0(\sigma^{\mathrm{nd}}) = 1-\mu_1(\sigma^{\mathrm{nd}})$ and $\mu_1(\sigma^{\mathrm{nd}}) > 1/2$. Thus, appealing to Proposition 2, we can find $\kappa_w(\sigma^{\mathrm{nd}})$ for $w \in \{0,1\}$ as follows:

$$\kappa_w\Big(\sigma^{\mathrm{nd}}\Big) = \mu_1\Big(\sigma^{\mathrm{nd}}\Big) - \mu_1^2\Big(\sigma^{\mathrm{nd}}\Big) + \tilde{\Delta}^{\mathrm{nd}}\mathbb{E}\Big[D^2\Big] + \Delta^{\mathrm{nd}}\Big(\mathbb{E}\Big[D^2\Big] - \mathbb{E}[D]\Big), \tag{32}$$

where Δ^{nd} and $\tilde{\Delta}^{\mathrm{nd}}$ are found as

$$\tilde{\Delta}^{\mathrm{nd}} = \Big(\mu_1^2\Big(\sigma^{\mathrm{nd}}\Big)(2-\rho_0) - \mu_1\Big(\sigma^{\mathrm{nd}}\Big) + 0.25\Big)\rho_0/(1-\rho_0)^2,$$

$$\Delta^{\mathrm{nd}} = \Big(\theta_0-\theta_0^2\Big)(0.5-\alpha) \times \mathbb{E}_{\tilde{\rho}}\big[\gamma(\lfloor D/2\rfloor; D-1, \theta_1) + \gamma(\lceil D/2-1\rceil; D-1, \theta_1)\big].$$

After some algebra, we can find $\mathcal{B}(\sigma^{\mathrm{nd}})$ from (31) and (32) as follows:

$$\mathcal{B}\Big(\sigma^{\mathrm{nd}}\Big) = \frac{N}{8}\left(\frac{1+\mathbb{E}[D^2]\tilde{\Delta}^{\mathrm{nd}} + (\mathbb{E}[D^2]-\mathbb{E}[D])\Delta^{\mathrm{nd}}}{\big(2\mu_1(\sigma^{\mathrm{nd}})-1\big)^2} - \frac{1}{4}\right)^{-1}$$

Based on the above, we have the next result that the data collector can drive the total payment to be arbitrarily small for a given N, provided that $\mathcal{P}_e \geq e^{-\mathcal{B}(\sigma^{\mathrm{ND}})}$. The proof is relegated to Appendix I.

Proposition 3: For the case with equal priors, if $\mathcal{P}_e \geq e^{-\mathcal{B}(\sigma^{\mathrm{ND}})}$, then we have that $\mathcal{L}(\mathcal{P}_e) = \delta N$ for any $\delta > 0$, indicating that the total payment can be driven to be arbitrarily small.

Remarks: Theorem 3 pinpoints to the positive impact of social learning for all participants of the privacy-preserving data collection game. For the data collector, it implies that he can lower the payment significantly when there are sufficiently many users. From the perspective of the users, each of them incurs zero privacy cost.

If the error constraint is tighter, then the data collector can employ the designed payment mechanism, $\tilde{\mathbf{R}}$. Under the equal priors assumption, we can find $\mathcal{B}(\sigma^*)$ from (24) and (25) as follows:

$$\mathcal{B}(\sigma^*) = \frac{N}{8}\left(\frac{1+\mathbb{E}[D^2]\tilde{\Delta} + (\mathbb{E}[D^2]-\mathbb{E}[D])\Delta}{(2\mu_1(\sigma^*)-1)^2} - \frac{1}{4}\right)^{-1}.$$

Note that $\mathcal{B}(\sigma^*) \geq \mathcal{B}(\sigma^{\mathrm{nd}})$ and the data collector reduces the error rate of the MAP detector $\hat{W}_{\sigma^*}(\mathbf{X})$ by gathering informative reported data from the users who play with the SR strategies. Based on Theorem 3, our next result reveals that, when $\mathcal{P}_e < e^{-\mathcal{B}(\sigma^{\mathrm{nd}})}$, the expected payment of the payment mechanism $\tilde{\mathbf{R}}$ constitutes an upper bound on $\mathcal{L}(\mathcal{P}_e)$.

Proposition 4: For the case with equal priors, when $\mathcal{P}_e < e^{-\mathcal{B}(\sigma^{\mathrm{nd}})}$, we have that

$$\mathcal{L}(\mathcal{P}_e) \leq Z\left(1-\beta+\frac{\mu_1(\sigma^*)}{2\beta-1}\right)N.$$

In the next section, we discuss the impact of social learning on the payment and accuracy with numerical examples.

C. Numerical Examples

In this section, we use examples to examine the impact of social learning on the trade-off between payment and accuracy and that between payment and privacy cost, using social learning graph models based on synthetic data and/or real-world data.

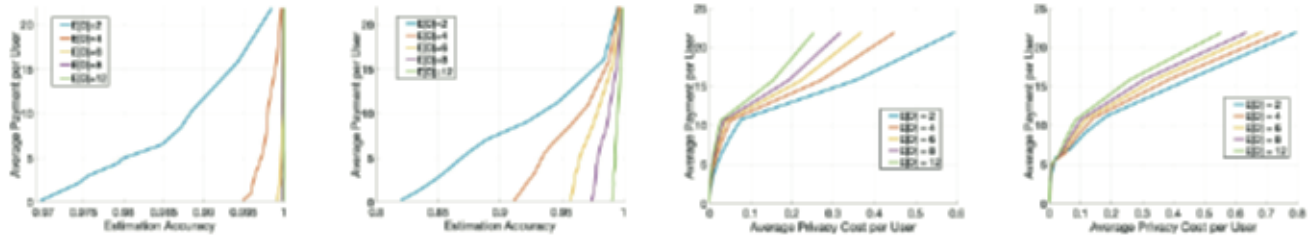

(a) $\mathbb{E}_{\sigma^\star}[\hat{R}_i(\mathbf{X})]$ vs. $\mathrm{P}_{\sigma^\star}(\hat{W}_{\sigma^\star}(\mathbf{X})=W)$ (α=0.35) (b) $\mathbb{E}_{\sigma^\star}[\hat{R}_i(\mathbf{X})]$ vs. $\mathrm{P}_{\sigma^\star}(\hat{W}_{\sigma^\star}(\mathbf{X})=W)$ (α=0.425) (c) $\mathbb{E}_{\sigma^\star}[\hat{R}_i(\mathbf{X})]$ vs. $\mathbb{E}[g(\{\sigma_i^\star, C_i\})]$ ($\alpha = 0.35$) (d) $\mathbb{E}_{\sigma^\star}[\hat{R}_i(\mathbf{X})]$ vs. $\mathbb{E}[g(\{\sigma_i^\star, C_i\})]$ ($\alpha = 0.425$)

Fig. 7. The impact of social learning: synthetic social learning graphs using the Erdös-Renyi random graph model with $N = 250$, $\theta_0 = 0.7$, $\mathrm{P}_W(0) = \mathrm{P}_W(1) = 0.5$ and $g(\zeta) = \zeta^2$. (a-b): Average payment per user vs. average degree. (c-d) Average payment per user vs. average privacy cost per user.

1) Synthetic Social Learning Graphs: To illustrate the impact of different parameters of the social learning graph and the payment mechanism on the performance, we first consider two synthetic models for the social learning graph. In the simulations, we use the Erdös-Renyi and Watts&Strogatz Model [41] random graph models. In the Erdös-Renyi model, each social tie is considered to be present with independent probability $\mathbb{E}[D]/(N-1)$. For large N, the degree distribution can be approximated by the Poisson distribution. Watts&Strogatz model starts from a ring lattice with N users and $\mathbb{E}[D]$ edges per user, and rewire each edge at random with probability p. In the simulations, we set $N = 250$, $\theta_0 = 0.7$ and $\mathrm{P}_W(0) = \mathrm{P}_W(1) = 0.5$ and consider the quadratic cost function, $g(\zeta) = \zeta^2$. For the Watts&Strogatz model, the rewiring probability is set as $p = 0.1$.

For the Erdös-Renyi model, Fig. 7a depicts the average payment each user receives in the payment mechanism $\tilde{\mathbf{R}}$ with respect to the state estimation accuracy, $\mathrm{P}_{\sigma^\star}(\hat{W}_{\sigma^\star}(\mathbf{X}) = W)$. It corroborates that the data collector can get an accurate estimate of the underlying state, with a much smaller payment compared to the case with no social learning. When the social learning among users strengthens (equivalently, the average degree $\mathbb{E}[D]$ increases), the privacy cost decreases because they receive informative social group signals C_i more often and hence they play the SR strategy less often. Fig. 7b demonstrates the payment vs. accuracy trade-off when the group signal noise level α is very high ($\alpha = 0.425$). In this case, to achieve a given accuracy level for the state estimator, the data collector needs to gather informative reported data from the users who play the SR strategy. The higher the payment is, the less noise the reported data would have (albeit the higher privacy cost), and hence the more accurate the state estimator is. As illustrated in Fig. 7c and Fig. 7d when the degree of a user increases, it is more likely for this user to play the ND strategy and hence her privacy cost drops. Accordingly, the total payment decreases. For the Watts&Strogatz model, Fig. 8 also verifies that the total payment required to achieve a given estimation accuracy decreases when the social learning among users strengthens.

2) Real World Social Learning Graphs: To evaluate the impact of social learning in practice, we also use two social learning graph models based on real-world data. Firstly, we study Arxiv GR-QC (General Relativity and Quantum Cosmology) collaboration network [42]. The graph contains an edge between authors i and j if they co-authored at least one paper. The graph has 5242 nodes and 14496 edges. Secondly, we use the Gnutella peer-to-peer file sharing network from August 2002 [42]. Nodes represent hosts in the file sharing network and edges represent connections between the Gnutella hosts. It has 6301 nodes and 20777 edges. Fig. 9 depicts the state estimation accuracy with respect to the payment per user under the proposed payment mechanism $\tilde{\mathbf{R}}$. These simulation studies also corroborate that the data collector can obtain an accurate estimate of W, with small amounts of payments despite the fact that very high noise is injected into group signals and private signals.

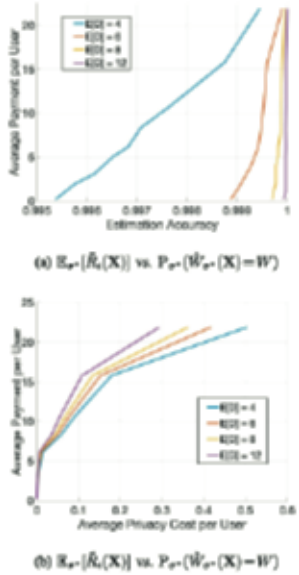

(a) $\mathbb{E}_{\sigma^\star}[\hat{R}_i(\mathbf{X})]$ vs. $\mathrm{P}_{\sigma^\star}(\hat{W}_{\sigma^\star}(\mathbf{X})=W)$

(b) $\mathbb{E}_{\sigma^\star}[\hat{R}_i(\mathbf{X})]$ vs. $\mathrm{P}_{\sigma^\star}(\hat{W}_{\sigma^\star}(\mathbf{X})=W)$

Fig. 8. The impact of social learning: synthetic social learning graphs using the Watts&Strogatz random graph model with $N = 250$, $\theta_0 = 0.7$, $\alpha = 0.4$, $\mathrm{P}_W(0) = \mathrm{P}_W(1) = 0.5$ and $g(\zeta) = \zeta^2$. (a): Average payment per user vs. average degree. (b) Average payment per user vs. average privacy cost per user.

3) Privacy Leakage of Group Signals: Next, we use numerical examples to illustrate the dependency of $\eta_{ij}(\sigma_i^\star)$ on the payment mechanism and the system model parameters. Recall that the peer-prediction-based payment mechanism $\tilde{\mathbf{R}}$ is determined in terms of $\epsilon \geq 0$ in (18) and the payment is increasing in ϵ. Fig. 10a depicts the privacy leakage, given in (29), with respect to α for different ϵ values. It reveals that as α increases, $\eta_{ij}(\sigma_i^\star)$ declines in general, with jumps which occur when $\lfloor \tau \rfloor$ changes. When $\lfloor \tau \rfloor$ increases, the user plays

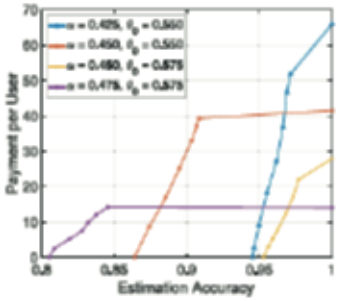

(a) Arxiv GR-QC Collaboration Network

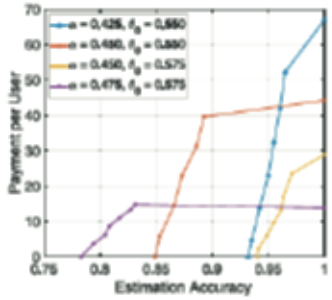

(b) Gnutella Peer-to-Peer Network

Fig. 9. $\mathbb{E}[\bar{R}_i(\mathbf{X})]$ vs. $\mathrm{P}_{\sigma^*}(\hat{W}_{\sigma^*}(\mathbf{X}) = W)$. The impact of social learning: payment per user vs. accuracy.

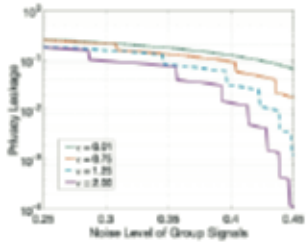

(a) $\eta_{ij}(\sigma_i^*)$ vs. α ($|\mathbb{S}_i[D]| = 6$).

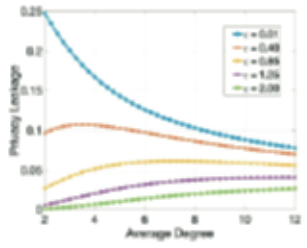

(b) $\eta_{ij}(\sigma_i^*)$ vs. $|\mathbb{S}_i[D]|$ ($\alpha = 0.40$).

Fig. 10. The impact of system parameters and payment mechanism on privacy leakage of group signals: (a) privacy leakage vs noise level of group signals (b) privacy leakage vs average degree ($\rho \sim$ Poisson, $\theta_0 = 0.7$ and $g(\zeta) = \zeta^2$.).

the SR strategy more often and hence the privacy leakage of the group signals, η_{ij}, decreases. Further, Fig. 10b depicts the privacy leakage with respect to the average degree of the social graph. If the payment is near zero with $\epsilon = 0.01$, then $\sigma_i^\star$ is effectively reduced to the simple majority rule in the group signals. In this case, as the average degree of the social learning graph increases, $\eta_{ij}(\sigma_i)$ monotonically decreases, simply because the privacy leakage occurs only through the sum signal, F_i. By contrast, if the payment is sufficiently larger than 0 with $\epsilon \geq 0.4$, the privacy leakage is always minimal.

VI. Conclusion and Future Work

In this article, we study a market model in which users can learn noisy versions of their social friends' data and make strategic decisions to report privacy-preserved versions of their personal data to a data collector. Thanks to the existence of social learning, the users have richer information about the underlying state beyond their personal signals. We develop a Bayesian game theoretic framework to study the impact of social learning on users' data reporting strategies and devise the payment mechanism for the data collector. Our findings reveal that, in general, the desired data reporting strategy at the Bayesian-Nash equilibria can be in the form of either symmetric randomized response or informative non-disclosive strategy. In particular, when a user plays the non-disclosive strategy, she reports her data completely based on her social group signals, independent of her personal signal, which drives her privacy cost to 0. As a result, both the data collector and the users benefit from social learning which lowers the privacy costs and helps to improve the state estimation at a given payment budget.

More specifically, our findings reveal that the desired data reporting strategy at the BNE is in the form of either a non-disclosive strategy or a symmetric randomized strategy. We show that the desired data reporting strategy is a *majority voting*-based data reporting rule which is applied by each user to her group signals to determine which strategy to follow. It is worth noting that the payment mechanism is designed to achieve *informative* equilibria, because no user can gain by playing uninformative when other users follow informative data reporting strategies. We caution that the social learning model does not imply any collusion among friends. We use a Central Limit Theorem for dependence graphs to evaluate the estimation error of the underlying state. The total expected payment is characterized subject to a constraint on the estimation error. Our analysis reveals both the data collector and users benefit from social learning: The data collector can get an accurate estimate of the underlying state, with a much smaller payment (compared to the case with no learning), thanks to social learning.

We are currently generalizing this study to account for the social learning costs of the users for sending out the noisy personal signals to their neighbors. To compensate this, the users might offer rewards to their friends before obtaining their noisy signals. In this setup, the Bayesian game becomes significantly more complicated where the users first employ "data acquisition strategies" and then data reporting strategies. In this study, we focus on the case where the data collector and users interact only once. Designing mechanisms for iterated games is also a very promising and important direction for further work. We also work on the opinion formation dynamics which is based on the fusion of private signals and group signals across heterogeneous users, e.g., diffusion models with influential and stubborn users. It is also of great interest to investigate the impact of "fake signals" (from fake news), and

our effort along this line is underway. In this market model, the utility of each user is designed to protect her own privacy and the "social" privacy cost of group signals is not accounted for yet. As elaborated in Section IV-B, this model makes sense when the noise level of the group signal is moderate or reasonably high (which is often the case in practical scenarios) or the average degree in the social learning graph is high and hence the privacy leakage of their friends' signals is insignificant. Alternatively, it is of interest to investigate the market model in which the users are "socially-aware" and they provide privacy guarantee to her friends' signals; our work along this line is underway.

REFERENCES

[1] W. Wang, L. Ying, and J. Zhang, "The value of privacy: Strategic data subjects, incentive mechanisms and fundamental limits," in *Proc. ACM Int. Conf. Meas. Model. Comput. Sci.*, 2016, pp. 249–260.

[2] A. Agarwal, D. Mandal, D. C. Parkes, and N. Shah, "Peer prediction with heterogeneous users," in *Proc. ACM Conf. Econ. Comput.*, 2017, pp. 81–98.

[3] Y. Chen, S. Chong, I. A. Kash, T. Moran, and S. P. Vadhan, "Truthful mechanisms for agents that value privacy," in *Proc. ACM Conf. Electron. Commerce*, 2013, pp. 215–232.

[4] B. Faltings and G. Radanovic, "Game theory for data science: Eliciting truthful information," in *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 11. San Rafael, CA, USA: Morgan Claypool, Sep. 2017, pp. 1–151.

[5] L. K. Fleischer and Y.-H. Lyu, "Approximately optimal auctions for selling privacy when costs are correlated with data," in *Proc. ACM Conf. Electron. Commer.*, 2012, pp. 568–585.

[6] A. Ghosh and A. Roth, "Selling privacy at auction," in *Proc. ACM Conf. Electron. Commer.*, 2011, pp. 199–208.

[7] A. Ghosh, K. Ligett, A. Roth, and G. Schoenebeck, "Buying private data without verification," in *Proc. ACM Conf. Econ. Comput.*, 2014, pp. 931–948.

[8] Y. Liu and Y. Chen, "Learning to incentivize: Eliciting effort via output agreement," in *Proc. Int. Joint Conf. Artif. Intell.*, 2016, pp. 3782–3788.

[9] B. Waggoner, R. Frongillo, and J. D. Abernethy, "A market framework for eliciting private data," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 3492–3500.

[10] Y. Kong and G. Schoenebeck, "An information theoretic framework for designing information elicitation mechanisms that reward truth-telling," *ACM Trans. Econ. Comput.*, vol. 7, no. 1, p. 2, Jan. 2019.

[11] N. Miller, P. Resnick, and R. Zeckhauser, "Eliciting informative feedback: The peer-prediction method," *Manag. Sci.*, vol. 51, no. 9, pp. 1359–1373, Sep. 2005.

[12] R. Jurca and B. Faltings, "Enforcing truthful strategies in incentive compatible reputation mechanisms," in *Proc. 1st Int. Conf. Internet Netw. Econ. (WINE)*, 2005, pp. 268–277.

[13] R. Jurca and B. Faltings, "Mechanisms for making crowds truthful," *J. Artif. Int. Res.*, vol. 34, no. 1, pp. 209–253, Mar. 2009.

[14] A. Dasgupta and A. Ghosh, "Crowdsourced judgement elicitation with endogenous proficiency," in *Proc. 22nd Int. Conf. World Wide Web*, 2013, pp. 319–330.

[15] V. Shnayder, R. M. Frongillo, and D. C. Parke, "Measuring performance of peer prediction mechanisms using replicator dynamics," in *Proc. Int. Joint Conf. Artif. Intell.*, 2016, pp. 2611–2617.

[16] Y. Kong, K. Ligett, and G. Schoenebeck, "Putting peer prediction under the micro(economic)scope and making truth-telling focal," in *Web and Internet Economics*. Berlin, Germany: Springer, 2016, pp. 251–264.

[17] B. Waggoner and Y. Chen, "Output agreement mechanisms and common knowledge," in *Proc. 2nd AAAI Conf. Human Comput. Crowdsourcing*, vol. 2, no. 1, 2014. [Online]. Available: https://ojs.aaai.org/index.php/HCOMP/article/view/13151

[18] V. Shnayder, A. Agarwal, R. Frongillo, and D. C. Parkes, "Informed truthfulness in multi-task peer prediction," in *Proc. ACM Conf. Econ. Comput.*, 2016, pp. 179–196.

[19] X. He, J. Liu, R. Jin, and H. Dai, "Privacy-aware offloading in mobile-edge computing," in *Proc. IEEE Global Comm. Conf. (GLOBECOM)*, 2017, pp. 1–6.

[20] L. Xiao, X. Wan, C. Dai, X. Du, X. Chen, and M. Guizani, "Security in mobile edge caching with reinforcement learning," *IEEE Wireless Commun.*, vol. 25, no. 3, pp. 116–122, Jun. 2018.

[21] M. Min *et al.*, "Learning-based privacy-aware offloading for healthcare IoT with energy harvesting," *IEEE Internet Things J.*, vol. 6, no. 3, pp. 4307–4316, Jun. 2019.

[22] M. Du, K. Wang, Z. Xia, and Y. Zhang, "Differential privacy preserving of training model in wireless big data with edge computing," *IEEE Trans. Big Data*, vol. 6, no. 2, pp. 283–295, Jun. 2020.

[23] C. Huang, H. Yu, J. Huang, and R. A. Berry, "Crowdsourcing with heterogeneous workers in social networks," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, 2019, pp. 1–6.

[24] A. V. Banerjee, "A simple model of herd behavior," *Quart. J. Econ.*, vol. 107, no. 3, pp. 797–818, Aug. 1992.

[25] S. Bikhchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *J. Polit. Econ.*, vol. 100, no. 5, pp. 992–1026, Oct. 1992.

[26] S. Rajaraman. (Sep. 2009). *Five Stars Dominate Ratings*. [Online]. Available: https://web.archive.org/web/20201203190556/https://blog.youtube/news-and-events/five-stars-dominate-ratings

[27] D. Sims. (Jul. 2017). *Netflix Believes in the Power of Thumbs*. [Online]. Available: https://web.archive.org/web/20190710211653/https://www.theatlantic.com/entertainment/archive/2017/03/netflix-believes-in-the-power-of-thumbs/520242

[28] N. Hu, P. A. Pavlou, and J. Zhang, "Can online reviews reveal a product's true quality? Empirical findings and analytical modeling of online word-of-mouth communication," in *Proc. 7th ACM Conf. Electron. Commer.*, 2006, pp. 324–330.

[29] T. Zhou, H. A. T. Kiet, B. J. Kim, B.-H. Wang, and P. Holme, "Role of activity in human dynamics," *EPL Europhys. Lett.*, vol. 82, no. 2, 2008, Art. no. 28002.

[30] M. E. J. Newman, "The structure and function of complex networks," *SIAM Rev.*, vol. 45, no. 2, pp. 167–256, May 2003.

[31] C. Dwork, "Differential privacy," in *Proc. Int. Conf. Automata Lang. Program. Vol. II*, 2006, pp. 1–12.

[32] M. S. Handcock, "Statistical models for social networks: Inference and degeneracy," in *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers*. Washington, DC, USA: Nat. Acad. Press, 2003, pp. 229–240.

[33] D. J. Watts, "The 'new' science of networks," *Annu. Rev. Sociol.*, vol. 30, no. 1, pp. 243–270, 2004.

[34] R. Toivonen, L. Kovanen, M. Kivelä, J.-P. Onnela, J. Saramäki, and K. Kaski, "A comparative study of social network models: Network evolution models and nodal attribute models," *Social Netw.*, vol. 31, no. 4, pp. 240–254, 2009.

[35] S. Janson, "Normal convergence by higher semiinvariants with applications to sums of dependent random variables and random graphs," *Ann. Probab.*, vol. 16, no. 1, pp. 305–312, Jan. 1988.

[36] L. von Ahn and L. Dabbish, "Designing games with a purpose," *Commun. ACM*, vol. 51, no. 8, pp. 58–67, Aug. 2008.

[37] V. S. Sheng, F. Provost, and P. G. Ipeirotis, "Get another label? Improving data quality and data mining using multiple, noisy labelers," in *Proc. ACM Int. Conf. Knowl. Disc. Data Min.*, 2008, pp. 614–622.

[38] D. Prelec, "A bayesian truth serum for subjective data," *Science*, vol. 306, no. 5695, pp. 462–466, Oct. 2004.

[39] T. Kailath, "The divergence and Bhattacharyya distance measures in signal selection," *IEEE Trans. Commun. Technol.*, vol. 15, no. 1, pp. 52–60, Feb. 1967.

[40] A. Bhattacharyya, "On a measure of divergence between two multinomial populations," *Sankhyā, Indian J. Stat.*, vol. 7, no. 4, pp. 401–406, 1946.

[41] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, pp. 440–442, Jun. 1998.

[42] J. Leskovec and A. Krevl. (Jun. 2014). *SNAP Datasets: Stanford Large Network Dataset Collection*. [Online]. Available: http://snap.stanford.edu/data

[43] B. Bollobás, "A probabilistic proof of an asymptotic formula for the number of labelled regular graphs," *Eur. J. Comb.*, vol. 1, no. 4, pp. 311–316, Dec. 1980.

[44] P. Billingsley, *Probability and Measure* (Probability and Statistics), 4th ed. New York, NY, USA: Wiley, 2012.

[45] M. E. J. Newman, S. H. Strogatz, and D. J. Watts, "Random graphs with arbitrary degree distributions and their applications," *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 64, Jul. 2001, Art. no. 026118.