Journal of Artificial Intelligence Research 69 (2020) 1127-1164 Submitted 05/2020; published 12/2020

# A Differential Privacy Mechanism that Accounts for Network Effects for Crowdsourcing Systems

**Yuan Luo** Y.LUO@IMPERIAL.AC.UK
**Nicholas R. Jennings** N.JENNINGS@IMPERIAL.AC.UK
*Department of Computing*
*Imperial College London, London SW7 2BX UK*

## Abstract

In crowdsourcing systems, it is important for the crowdsource campaign initiator to incentivize users to share their data to produce results of the desired computational accuracy. This problem becomes especially challenging when users are concerned about the privacy of their data. To overcome this challenge, existing work often aims to provide users with differential privacy guarantees to incentivize privacy-sensitive users to share their data. However, this work neglects the network effect that a user enjoys greater privacy protection when he aligns his participation behaviour with that of other users. To explore this network effect, we formulate the interaction among users regarding their participation decisions as a population game, because a user's welfare from the interaction depends not only on his own participation decision but also the distribution of others' decisions. We show that the Nash equilibrium of this game consists of a threshold strategy, where all users whose privacy sensitivity is below a certain threshold will participate and the remaining users will not. We characterize the existence and uniqueness of this equilibrium, which depends on the privacy guarantee, the reward provided by the initiator and the population size. Based on this equilibria analysis, we design the PINE (Privacy Incentivization with Network Effects) mechanism and prove that it maximizes the initiator's payoff while providing participating users with a guaranteed degree of privacy protection. Numerical simulations, on both real and synthetic data, show that (i) PINE improves the initiator's expected payoff by up to 75%, compared to state of the art mechanisms that do not consider this effect; (ii) the performance gain by exploiting the network effect is particularly good when the majority of users are flexible over their privacy attitudes and when there are a large number of low quality task performers.

## 1. Introduction

A growing number of crowdsourcing applications aggregate and utilize users' personal data for decision-making purposes. Examples include movie rating systems such as Netflix and OK.com and online survey systems such as QuestionPro and SurveyMonkey. Users who share their personal data can certainly derive benefits from such information aggregation. For example, a user may receive good recommendations from a movie rating system or know whether he is paid fairly from an online survey. Furthermore, the benefit that a user derives typically increases with the number of participants in the system. Hence, crowdsourcing applications need to recruit and maintain large numbers of users to maintain the accuracy of their results.

However, the increasing use of individuals' data has been accompanied by growing concerns about their privacy. For example, Narayanan and Shmatikov (2008) were able to recover private information such as movie viewing history and political preferences from Netflix's published movie ratings. This risk of exposing private information deters privacy-sensitive users who do not want to suffer privacy losses when sharing their personal data (Jin et al., 2016). Given this, appropriate means for encouraging participation are central to the design of crowdsourcing systems (Lev et al., 2013; Chandra et al., 2017). In particular, it is important to develop methods for incentivizing such users to share their data with the crowdsource campaign initiator[1] or her agents who will process and publish the computational result[2] (Dandekar et al., 2014; Ghosh & Roth, 2015). To this end, recent work has started to explore the question by relating differential privacy guarantees to questions of incentives in mechanism design (*e.g.*, Ghosh & Roth, 2015; Chen et al., 2016b; Wang et al., 2016).

In more detail, a guarantee of *differential privacy* provides a degree of privacy protection to participating users when an initiator computes the data of a population and publishes the computational result (Dwork et al., 2006). Specifically, a differentially private mechanism involves the initiator publishing a noisy version of the computational result, where the noise level corresponds to a differential privacy guarantee (more noise means greater privacy). Using the added noise ensures that whether or not a user chooses to share his data only marginally changes the statistics inferred from the published computational result. As a consequence, any participant's information is harder to determine from the published results (Zhang et al., 2016). For example, if an average rating for a movie is 4, Netflix can publish the rating as 4.6 to provide a differential privacy guarantee. The difference between the published rating (*i.e.*, 4.6) and the true rating (*i.e.*, 4) is determined by the differentially private mechanism adopted by Netflix (McSherry & Mironov, 2009). Similarly, instead of publishing that the true percentage of people whose income is above 50K US dollars is 10%, QuestionPro could publish the percentage as 8%. Here, the differential privacy mechanism adopted by the survey system determines the difference between the true percentage and the announced percentage (Xu et al., 2013).

In addition to providing participants with differential privacy guarantees, one of the simplest ways to incentivize users to share their personal data is to pay them for using it. However, determining the correct price is challenging: low payments may not draw sufficient participants, causing insufficient data for an accurate computation, while high payments may make the system uneconomic. Existing incentive mechanisms with differential privacy guarantees aim to elicit each user's valuation for the various levels of guarantee and then tailor payments based on the elicited values to incentivize sufficient participants to produce accurate results and pay the minimum possible cost (see Section 2 for details). However, these mechanisms are often not suitable for crowdsourcing systems for a number of reasons. First, most real-world crowdsourcing applications that utilize personal data only mention how an individual's data might be used, instead of promising a differential privacy guarantee. For example, Google adds a single specified level of noise to users' web browsing data that is uploaded (Erlingsson et al., 2014) and Apple uses differential privacy to collect emoji usage data from users without identifying the added noise level (Apple, 2017a). In both of these cases, users cannot determine their privacy loss. Second, a crowdsourcing system

[1] For convenience, we term the crowdsource campaign initiator the "initiator" in the remainder of this paper.

[2] We refer to the initiator who collects and processes the data as "she" and individual users as "he".

can have thousands of participants and it is costly to elicit their valuations and determine payment individually with each of them based on his individual privacy loss. Finally, it is more difficult to identify a user's identity in a large population than in a smaller one. Thus, the more users contributing their personal data, the more difficult it is to identify a specific user's identity and therefore the more users should be willing to contribute their data. This *network effect* (Easley & Kleinberg, 2010) means a user's participation decision may be correlated with that of other users. To date, however, most existing work does not consider this network effect on users' participation decisions, meaning such mechanisms overpay to achieve a given level of privacy.

In this paper, we aim to bridge these gaps. Specifically, we design a mechanism that both exploits network effects to provide participants with a differential privacy guarantee and maximizes an initiator's payoff. The initiator's payoff is here defined as the difference between the payments made to the (privacy-sensitive) users and the computational accuracy of the population's data[3]. Designing such a mechanism poses interesting scientific challenges since it represents one of the first examples where privacy, game theory, social network theory and mechanism design are brought together to devise effective methods to incentivize a sufficient number of privacy-sensitive users with a minimum payment (Gatti et al., 2015). Furthermore, this mechanism design is complex as an initiator's payment strategy couples with the privacy guarantee level and all potential users' strategies. Such coupling exists for a number of reasons. First, an initiator's payment that incentivizes users' participation depends on the privacy guarantee level. A low privacy guarantee level requires the initiator to pay a high price to compensate for a user's privacy loss. A high privacy guarantee can reduce the payment made to a user. Second, the privacy guarantee level provided by the initiator can be increased by utilizing the network effect. Here, the benefit of utilizing the network effect depends on the number of participants which is determined by all potential users' strategies. Thus, the initiator's payment strategy, the privacy guarantee level and the users' participation strategies couple together. As existing mechanisms do not consider the interaction among these strategies, there is no way to design a mechanism by simply extending or combining existing techniques.

Against this background, we formulate the interactions between the users and the initiator (*i.e.*, the initiator first announces the payment to users, and the users respond by participating or not) as a two-stage Stackelberg game. Such games are often used to model a two-step process (*e.g.*, Shokri et al., 2012; Pawlick & Zhu, 2016; Liao et al., 2018). We then use backward induction (Von Stackelberg, 1934; Fudenberg & Tirole, 1991) to analyze users' participation strategies and the initiator's payment strategy. Based on the results of the analysis, we design a novel mechanism, which we call "PINE" (Privacy Incentivization with Network Effects). PINE utilizes the network effect to maximize the initiator's payoff, while providing privacy protection to participants. Moreover, PINE satisfies the *Individual Rationality* condition, *i.e.*, if a user's utility, the difference between the payment received from the initiator and the cost of sharing personal data, is non-negative, then he should be willing to participate. Thus, PINE can incentivize appropriate numbers of participants to maximize the initiator's payoff.

[3] Computational accuracy is defined as the difference between the initiator's published result and the true computation on the full set of data from the entire population (Liu & Chen, 2017).

This paper advances the state of the art in the following ways. First, we show how the network effect influences a user's participation decision. In particular, we show users' participation decisions couple together and form a population game among all the users. That is, a user's welfare from the interaction depends not only on his own participation decision, but also the distribution of others' decisions. We show that the Nash equilibrium of this game consists of a threshold strategy, where all users whose privacy sensitivities are lower than this threshold participate and the remainder do not. We characterize the existence and uniqueness of these equilibria, which depend on the privacy guarantee, the reward provided by the initiator and the population size. Second, we characterize the initiator's optimal payment strategy to maximize the initiator's expected payoff. Our analysis shows that if the initiator provides a high level privacy guarantee to participants, she can decrease her payment in line with the number of users. If the initiator provides a low level privacy guarantee to participants, her best payment strategy is to pay them nothing because the benefit brought by incentivizing users' participation is smaller than the payment made to the users. Thus, the initiator gives up incentivizing users in return for a low cost to maximize her payoff. If the initiator provides an in-between privacy guarantee level, the optimal payment only attracts a subset of the potential users to contribute their data. Third, our numerical simulations show that considering the network effect can significantly improve the initiator's expected payoff; the improvement is at least 50% and can be up to 75%, compared to the state of the art mechanisms that do not consider this effect. Finally, our empirical studies on both synthetic data and real-world data sets highlight the conditions under which utilizing the network effect can bring the greatest benefit. In particular, network effects are most apparent when the number of potential users is large. Moreover, compared to the setting where most users are either always willing or always reluctant to share their data, the benefit of utilizing network effects to maximize the initiator's expected payoff is most significant when the majority of users are willing to provide data only when the initiator provides sufficient incentivization and some degree of privacy guarantee. If the initiator can calculate an accurate computational result based on a small dataset, the benefit of the network effect is less important.

The rest of the paper is organized as follows. After reviewing the literature in Section 2, we describe a two-stage Stackelberg game in Section 3 to characterize the interaction between the users and the initiator. We apply backward induction to study this two-stage Stackelberg game and show the strategies of users and the initiator in Sections 4 and 5, respectively. Based on this analysis, we present PINE in Section 6. Then, we provide the simulation results in Section 7 to show the benefit of utilizing the network effect by comparing the performance of PINE to the state of the art mechanisms. Finally, we conclude in Section 8.

## 2. Related Work

A growing body of research has been carried out on a variety of topics at the intersection of privacy, game theory and mechanism design. One fruitful line of inquiry has considered market-based mechanism design, focusing on applying pricing or auctions to collect data from privacy-sensitive users. For example, Wang et al. (2016) study how a monopoly initiator engages in price discrimination during data collection to keep users' identities private. Li et al. (2014) propose a theoretical framework for assigning prices to users' answers based on their accuracy. Ho et al. (2016) design a monetary contract to collect high-quality data. Other work studies how to prevent unnecessary information about bids (such as users' actions, political preferences or social connections) from being leaked in an auction (*e.g.*, Brandt & Sandholm, 2008; Feigenbaum et al., 2010; Gao et al., 2016; Krause & Horvitz, 2010). There has also been work that incentivizes users to report their personal data truthfully by designing appropriate mechanisms (Kamar & Horvitz, 2012; Jain et al., 2018). All the above studies assume that revealing data to an initiator does not incur privacy costs to users. However, in practice, users may consider themselves suffering privacy losses when sharing their data (Gradwohl, 2017). Furthermore, an experimental study by Huberman, Adar, and Fine (2005) shows that users have diverse evaluations for their privacy losses and each user's evaluation is his private information which is unknown to the initiator.

An alternative approach to collecting data from users is to use differential privacy mechanisms. For example, the Chrome web browser uses RAPPOR (Erlingsson et al., 2014), a differentially private mechanism, to track the distribution of users' browser configuration behaviour while protecting users' privacy. Apple deploys local differential privacy in its iOS system for collecting their users' emoji usage data (Apple, 2017b). The basic ideas of these differentially private mechanisms are the same, *i.e.*, adding statistical noise to a user's individual data before this data is shared with Chrome or Apple. The differences come from the added noise level. Kifer and Machanavajjhala (2014) propose Pufferfish to protect users under correlated data evolution scenarios. There is also a related literature on providing various tasks with a certain level of differential privacy (see Dwork, 2008, for a detailed survey). However, what has been almost entirely missing is any normative guidance for balancing the tradeoff between incentivization of individuals' participation and an initiator's computational accuracy. Therefore, our paper proposes to answer this question of how the reward should be chosen to incentivize users' participation, while maximizing the initiator's payoff. Here, the initiator's payoff is defined as the difference between the payments made to the participants and the initiator's computational accuracy.

Several papers do study the conflict between users' participation that are affected by differential privacy and an initiator's computational accuracy by using Stackelberg games (see Zhu & Rass, 2018, for a detailed survey). For example, Shokri et al. (2012) formulate a game for preserving location privacy. In this game, after participating users choose the noise level added to their location, the initiator chooses the optimal way to reconstruct each user's location. In contrast, Pawlick and Zhu (2016) let the initiator choose a level of privacy protection first and users simultaneously make participation decision as well as choosing the added noise levels later. Dunyak and Zhu (2018) let the initiator decide privacy protection level and users react to privacy promises as Stackelberg followers. Liao et al. (2018) address the conflict issue by formulating a two-stage Stackelberg game where the initiator selects participants first and users respond by choosing appropriate levels of differential privacy. However, none of this work discusses the impact of payment and network effects on users' participating decisions and the initiator's payoff.

Closest to our work is the recent literature on modeling and designing mechanisms for computation in settings where a user's data is sensitive and sharing it incurs a privacy loss cost that should be compensated for (*e.g.*, Ghosh & Roth, 2015; Yang et al., 2018; Zhang et al., 2016, 2018; Jin et al., 2018). In such settings, users' evaluations for their privacy loss are independent of the actual data they have and whether they report this truthfully. The initiator typically aims to compensate users for their costs incurred by the privacy losses so that she can attract sufficient participants to improve her computational accuracy. As a user's cost for the privacy loss is his private information, the existing literature proposes a number of different mechanisms to elicit this information and designs appropriate rewards based on the elicited information to maximize the initiator's payoff. For example, Ghosh and Roth (2015) propose an auction-based mechanism, called MinCostAuction, to learn every user's privacy loss cost and then tailor payments based on the elicited values to incentivize sufficient participants to produce accurate results with a minimum possible cost. Zhang *et al.* (2018) apply contract theory to elicit every potential user's evaluation for his privacy loss and incentivize all potential users to participate. Based on the elicited result, the initiator negotiates with every user regarding the noise level added to the user's contributing data and the corresponding payment so that her payoff can be maximized. Their solution is called REAP. However, in crowdsourcing applications, an initiator usually recruits thousands of users (Nath & Narayanaswamy, 2014; Jain et al., 2016) so it would be very costly to negotiate compensation individually with each of them. Furthermore, in practice, a user may be unable to state a quantitative privacy loss but is usually able to say whether or not he is willing to share his data (Ghosh & Kleinberg, 2014). Thus, the above mechanisms are not easily implemented in practical crowdsourcing systems. Nevertheless, MinCostAuction and REAP provide a good basis for benchmarking our proposed mechanism (see Section 7.1 for more details), as they provide efficient solutions for the same problem as this paper.

Finally, our model is also related to the literature that studies network effects. For example, Levit *et al.* (2018) study the impact of network effects on users' actions and social welfare. Ghosh and Ligett (2013) show how users make their participation decisions considering the network effect. Gradwohl (2017) analyzes the relationship between the network effect and users' privacy revealing risk. Tembine et al. (2009) formulate users' interactions as a population game and study the impact of the network effect on users' participation decisions. Manshaei et al. (2008) analyze a population game to show how users subscribe to a wireless access provider. However, all of these studies focus solely on interactions among different users without considering the initiator's payment scheme and how the network effect affects the payment strategy of the initiator when she needs to balance the computational accuracy and the cost. Thus, they cannot help the initiator to maximize her payoff. Some operations management and marketing science literature (*e.g.*, Candogan et al., 2010; Lin & Lu, 2011; Ajorlou et al., 2016; Luo et al., 2016; Chen et al., 2016a; Nair et al., 2015). does study the impact of network effects on an initiator's payment strategy. However, they focus on a public goods setting which does not require an initiator to make a computation on a collected dataset and they do not consider providing privacy protection to participants.

## 3. A Two-Stage Stackelberg Game for Crowdsourcing Systems

We consider a setting where an initiator wants to learn and publish some statistic about the population. Learning the average rating for a movie and knowing the average number of people whose income is above 50K US dollars in a certain area are two examples. The initiator will conduct a survey that asks a set $\mathcal{N} = \{1, 2, \ldots, N\}$ of $N$ users and gives a participating user a reward $\pi$. Users choose whether or not to contribute their data to the initiator. We let $n \leq N$ denote the number of actual participants in this survey.

### 3.1 The Initiator

Given that there are $n$ participating users, the initiator has a database $D \in \{0, 1\}^n$ with each entry of $d \in \{0, 1\}$ corresponding to some private information for a user[4], *e.g.*, his attitude towards a movie (good or bad) or his wealth (exceeding 50K dollars or not). The initiator computes the average of the dataset[5], *i.e.*, $g : \mathbb{R}^n \to \mathbb{R}$ with $g(D) = \frac{1}{n}\sum_{j=1}^{n} d_j$ and publishes the computational result in a *differentially private* way to preserve the privacy of all participants. In more detail, let $\mathcal{L} : \mathbb{R}^n \times \Omega \to \mathbb{R}$ denote the mechanism that the initiator adopts to compute and publish the mean of the population, where $\Omega$ denotes the sample space of noise added to the computational result $g(D)$, which is the set of all possible outcomes of the noise. Specifically,

$$\mathcal{L}(D) = g(D) + w, \tag{1}$$

where the added noise $w$ is drawn from the Laplace distribution $Lap(0, 1/\epsilon)$ with mean zero and standard deviation $\sqrt{2}/\epsilon$. We can view $1/\epsilon$ as the noise level added to the computational result[6]. Then we can check that $\mathcal{L}$ supports $\epsilon/n$-differential privacy (Dwork et al., 2006). Specifically, for any two neighbouring pairs of datasets $D$ and $D'$ where the dataset $D$ includes the private data of user $j \in \mathcal{N}$ and the dataset $D'$ does not, we have the following result given the number of actual participants $n$ and any possible computational result $\alpha$ based on the calculation of the dataset (*e.g.*, a movie rating calculated based on all participants' attitudes towards this movie or the average number of people whose income is above 50K US dollars in a country):

$$Pr[\mathcal{L}(D) = \alpha] \leq e^{\epsilon/n} Pr[\mathcal{L}(D') = \alpha]. \tag{2}$$

Based on (2), the differential privacy guarantee states that the probability that any output of mechanism $\mathcal{L}$ will be within an $e^{\epsilon/n}$ multiplicative factor whether or not a user's private data is included in the computational dataset. In particular, the parameter $\epsilon/n$ controls how much the distribution of the published computational result depends on data from a user[7]. We denote $\epsilon/n$ as a user's privacy revealing risk. This is because with an increase of

[4] We focus on binary private information for analytical convenience. We only need to revise the worst failure probability discussed in equation (3) to deal with the more general case.

[5] For analytical convenience, we assume the initiator wants to learn the mean of a population. Relaxing to learn other statistics of a population does not change our main insights (Zhang et al., 2016; Ghosh & Ligett, 2013).

[6] We add random noise with Laplacian distribution for analytical convenience. Adding random noise with other distributions such as Gaussian or binomial distribution does not change our main results (Dwork et al., 2006; Geng & Viswanath, 2014).

[7] We assume that the accuracy of the initiator's computational result depends only on the number, rather than the specific subset, of users that contribute their private data. This is a standard assumption (Ghosh & Roth, 2015; Zhang et al., 2018; Ghosh & Ligett, 2013; Jin et al., 2015) and holds in settings where no particular data is more or less incriminating than another, such as collecting data regarding watching movie behaviour, mobile applications for traffic monitoring, or polls in online communities.

$\epsilon/n$, the difference between two neighbouring databases increases. In such cases, it is easier to identify whether a user's private data is included in the computational dataset.

We assume that the value of $\epsilon$ is an exogenous variable[8] and the initiator will notify the users of this value. We can also see that a user's privacy revealing risk $\epsilon/n$ decreases as the value of $1/\epsilon$ (*i.e.*, the noise level added to the published computational result) increases.

The initiator's main concern is the accuracy of the published computational result with respect to the true mean of the entire population. We assume that the true mean of the target population is $\mu$ (*e.g.*, the average percentage of people in London who think the movie Dunkirk is good and the millionaire ratio in New York). Let function $A$ be a measure of accuracy. In general, there are several choices of what $A$ can measure, *e.g.*, average error (Chatzikokolakis et al., 2014), estimation distortion (Yang et al., 2018) and regularized empirical error (Ligett et al., 2017). As users' contributed data and added noise level are all random variables and the initiator has little information about the users' contributed data, we can only use a probability bounds analysis to quantify the uncertainties of various random variables. Such a quantitative calculation allows us to characterize the strategies of the initiator and users. In more detail, let $A : \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^+ \to \mathbb{R}^+$ measure the upper bound of the failure probability (*i.e.*, $Pr\big[\big|g(D) + w - \mu\big| \geq T\big]$ given the fixed error $T$). We call this the worst failure probability[9]. By applying the Chernoff bound, we have (Hoeffding, 1994)

$$A(n, T, \epsilon) = 2\exp\left(-\frac{nT^2}{12}\right) + \exp\left(-\frac{T\epsilon}{2}\right). \tag{3}$$

We assume that the initiator's utility, when she has perfect accuracy (*i.e.*, her published computational result is the same as the true mean), is zero. Let $v$ denote the monetary punishment to the initiator when she fails (*i.e.*, the difference between her published computational result and the true mean of the population is larger than the desired error $T$). Then, the initiator's expected utility is $-v \cdot A(n, T, \epsilon)$. Correspondingly, the expected payoff of the initiator $U^{\text{IN}} : \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^+ \to \mathbb{R}$ is defined as the difference between her utility obtained from the computation and the monetary payments to all participating users (Krause & Horvitz, 2010)

$$U^{\text{IN}}(v, n, T, \epsilon, \pi) = -v \cdot A(n, T, \epsilon) - \pi \cdot n. \tag{4}$$

### 3.2 The Users

After knowing the noise level $1/\epsilon$ added to the computational result and the payment $\pi$, each user decides whether to participate in a particular survey based on his payoff. We denote

[8] The value of $\epsilon$ depends on the circumstances of the data collection process. Experimental evaluations of differential privacy typically pick a value ranging from 0.01 to 10 (see Table 1 in Hsu et al., 2014). However, widespread practice assumes that the value of $\epsilon$ in computing the average of the dataset should not be larger than 7 (Machanavajjhala et al., 2008).

[9] Given a users' contributed dataset, the true mean of the target population and the mean value of the noise level, we can use numerical methods to calculate the empirical values of $Pr\big[\big|g(D) + w - \mu\big| \geq T\big]$. We refer to the value of $Pr\big[\big|g(D) + w - \mu\big| \geq T\big]$ as the failure probability. The detailed result is shown in Section 7.

the choice of user $j \in \mathcal{N}$ as $\ell_j \in \{0,1\}$, where $\ell_j = 1$ if he participates and $\ell_j = 0$ otherwise. Let $\theta_j$ denote user $j$'s privacy sensitivity. Recall that a user's privacy revealing risk is $\epsilon/n$ given there are $n$ participants. A type-$\theta_j$ user's evaluation for his privacy revealing risk (*i.e.*, privacy loss) is $\theta_j \cdot \frac{\epsilon}{n}$. We define the payoff of user $j$ as the difference between the payment he received from the initiator and his privacy loss (Jin et al., 2016; Ghosh & Kleinberg, 2014). In particular, given the same privacy revealing risk $\epsilon/n$, a highly privacy-sensitive user with a high value of $\theta_j$ would consider himself suffering a high privacy loss. For a type-$\theta_j$ user, his payoff $U^{\text{UE}} : \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^+ \to \mathbb{R}$ is

$$U^{\text{UE}}(\pi, \theta_j, \epsilon, n) = \begin{cases} \pi - \theta_j \cdot \dfrac{\epsilon}{n}, & \text{if } \ell_j = 1, \\ 0, & \text{if } \ell_j = 0, \end{cases} \tag{5}$$

Note that different users may have a different value of $\theta$ (Huberman et al., 2005), i.e., users are heterogeneous in terms of $\theta$. Thus, we assume that the users' privacy sensitivities follow a probability distribution whose probability density function (p.d.f.) $h$ is strictly positive and continuous on $[0, \bar{\theta}]$ for some $\bar{\theta} > 0$. Note that a user with $\theta = 0$ does not care about his privacy. For completeness of definition, we have $h(\theta) = 0$ for all $\theta \notin [0, \bar{\theta}]$. The cumulative distribution function (c.d.f.) is given by $H(\theta) = \int_{-\infty}^{\theta} h(y)\mathrm{d}y$ for all $\theta \in \mathbb{R}$.

As can be seen, different numbers of participating users result in a difference of a user's privacy revealing risk and a type-$\theta_j$ user's privacy revealing risk (*i.e.*, $\theta_j \cdot \frac{\epsilon}{n}$) decreases with the number of participants $n$. Thus, we can have the following observation: a user's participation decision is affected by the network effect. Specifically, the more users participating in a survey, the lower the privacy revealing risk faced by an individual participant and therefore the more users should be willing to participate in the survey.

Note that without considering the network effect, a user considers his privacy revealing risk is $\epsilon$ which is independent of the number of participants and is estimated based on the noise level announced by the initiator only. Thus, his payoff is given in the following equation:

$$U^{\text{UE}}(\pi, \theta_j, \epsilon) = \begin{cases} \pi - \theta_j \cdot \epsilon, & \text{if } \ell_j = 1, \\ 0, & \text{if } \ell_j = 0, \end{cases} \tag{6}$$

### 3.3 The Two-Stage Stackelberg Game

Based on the above discussion, a user's participation decision depends on the payment announced by the initiator. Thus, we can formulate the interactions between the initiator and the users as a two-stage Stackelberg game as discussed previously. Specifically, in Stage I, the initiator determines the payment $\pi$ given the noise level $1/\epsilon$, the number of potential users $N$, and the desired error $T$. In Stage II, users decide to participate in the survey or not based on the payment announced by the initiator at Stage I and the network effect brought by other users' participation decisions. We will study this two-stage game by backward induction. Namely, we first study the users' participation behaviours in Stage II (Section 4), then based on the participation analysis we study the initiator's optimal pricing decision that maximizes her expected payoff in Stage I (Section 5).

## 4. The Users' Best Strategies

In this section, given the payment specified by the initiator in Stage I, we study strategies of users who are self-interested and participate only when participation brings a non-negative payoff. As shown in (5), a user's payoff varies with the number of participants so we will formulate the interaction among users as a population game and study the existence and uniqueness of an equilibrium in this game. Before explaining this population game, we begin by defining one standard term that is used for defining population games.

**Definition 1.** *Consider a population of individuals that can use a set of pure strategies, $\boldsymbol{\ell}$. A social state is a vector $\boldsymbol{x}$ that gives a mass of users $x_\ell$ with which each strategy $\ell \in \boldsymbol{\ell}$ is played in the population (Sandholm, 2010).*

In our model, each user can choose to participate ($\ell = 1$) or not ($\ell = 0$), *i.e.*, $\boldsymbol{\ell} = \{1, 0\}$. Then, the social state is $\boldsymbol{x} = (x_1, x_0)$, where $x_1$ represents the mass of users choosing $\ell = 1$ and $x_0 = 1 - x_1$ denotes the mass of users choosing $\ell = 0$. Now we can define the Population Game formally as follows.

**Definition 2.** *[Population Game]*

- *Players: Potential users in the set $\mathcal{N}$ with different privacy sensitivities;*
- *Strategies: The user $j \in \mathcal{N}$'s strategy is his participation decision $\ell_j \in \boldsymbol{\ell}$;*
- *Social State: $\boldsymbol{x} = (x_1, x_0)$;*
- *Payoffs: The users' payoffs are defined in* (5).

In this game, a type-$\theta_j$ user's optimal strategy (*i.e.*, best response) $\ell_j^*$ is the strategy in $\boldsymbol{\ell}$ that is optimal at each social state $\boldsymbol{x}$ (Sandholm, 2010), *i.e.*, the strategy that maximizes the user's payoff defined in (5). Formally, we define the Nash equilibrium (NE) of the population game as follows.

**Definition 3.** *A social state $\boldsymbol{x}^* = (x_1^*, x_0^*)$ is a Nash equilibrium if each user chooses a best response to this state and no user can improve his payoff by deviating from his best response $\ell_j^*$ to a different response $\ell_j$.*

To find and study the Nash equilibrium, we first need to analyze users' best responses. As each user will form a belief on the number of participants when he makes a participation decision, we construct and analyze a dynamic model that specifies how users form their beliefs and make decisions. In more detail, we introduce a virtual discrete time system with time periods indexed $t = 1, 2, \ldots$, where user $j$ makes his participation decision at the beginning of slot $t$, based on his belief of the number of participants at the end of the previous time slot[10].

[10] Note that each user makes this participation decision only once (because a crowdsourcing system only allows a user to share his data once in order to avoid duplicate data) and it is independent of anyone else (because users are anonymous in crowdsourcing systems and do not know other users' identities). The main purpose of introducing the virtual discrete time system is to characterize the relation between the payment and the number of participants, and to facilitate the calculation of the initiator's optimal pricing strategy later. Such an analysis technique is commonplace in the existing literature, *e.g.*, Manshaei et al. (2008); Ren and Van der Schaar (2012); Luo et al. (2016).

Note that the population game assumes that all users have the same belief of the mass of participants and the number of participants at the same time slot. This is a standard assumption (Sandholm, 2010) and holds in many practical crowdsourcing settings as an initiator often updates the number of participants in real-time and publicizes this number in the systems (*e.g.*, Netflix, QuestionPro and SurveyMonkey).

Let $x_1^{t-1}$ denote the users' common belief of the mass of participants at the end of slot $t-1$. We also let function $f:[0,1]\times\mathbb{R}^+\to\mathbb{R}^+$ characterize the number of participants that the users believe at time $t$. For a user who knows the number of potential users except himself is $N-1$ and believes the mass of participants at the end of slot $t-1$ is $x_1^{t-1}$, the number of participants that this user believes at time $t$ is $f(x_1^{t-1},N)=x_1^{t-1}\cdot(N-1)$. Let a user's privacy sensitivity be $\theta$. Then this user participates at slot $t$ if and only if $\pi\geq\theta\cdot\frac{\epsilon}{f(x_1^{t-1},N)+1}$. That is, only those users with a privacy sensitivity smaller than or equal to $\frac{\pi\cdot[f(x_1^{t-1},N)+1]}{\epsilon}$ will participate at slot $t$. Let function $g:[0,1]\to[0,1]$ characterize the mass of participants in time slot $t$ given the participant fraction in the previous slot $t-1$. Recall that the c.d.f. of user's privacy sensitivity is $H$. Then, the mass of participants evolves following a sequence $\{x_1^t\}_{t=0}^{\infty}$ in $[0,1]$ generated by

$$x_1^t=g(x_1^{t-1})=H\left(\frac{\pi\cdot[f(x_1^{t-1},N)+1]}{\epsilon}\right)\tag{7}$$

for $t=1,2,\ldots$, starting from a given initial point $x_1^0\in[0,1]$. Note that the payment $\pi$ is fixed over time. Given the user participation dynamics (7), we are interested in whether the mass of participants will stabilize in the long run and, if so, to what value.

Let $\Delta x$ be the change of participant mass between two successive time slots, *e.g.*, $t-1$ and $t$, that is

$$\Delta x(x_1^{t-1})=x_1^t-x_1^{t-1}.\tag{8}$$

Obviously, if $\Delta x$ is zero in slot $t$, *i.e.*, $x_1^t=x_1^{t-1}$, then users will not change their participation strategies in the future and the mass of participants remains the same from that time slot on. Thus, this user participation dynamic system reaches an equilibrium. Formally,

**Proposition 1.** *A social state $\boldsymbol{x}^*=(x_1^*,x_0^*)$ where $x_1^*+x_0^*=1$ is a Nash equilibrium if and only if*

$$\Delta x(x_1^*)=H\left(\frac{\pi\cdot[f(x_1^*,N)+1]}{\epsilon}\right)-x_1^*=0.\tag{9}$$

We provide the detailed proof of Proposition 1 in Appendix A. Notice that under a fixed payment $\pi$, (9) can have multiple solutions, which results in different Nash equilibria. In such cases, there are a number of strategies that do equally well in the population and the population could drift in the direction of their other strategies and their corresponding Nash equilibrium. In more detail, we illustrate the dynamics of social state in Figure 1. The blue line denotes the isoline of $\Delta x=0$ and the red line denotes the line of $x_1+x_0=1$. By Proposition 1, the intersections between the blue curve and the red line are the Nash equilibria. The black arrows denote the social state changing directions at any given social state.

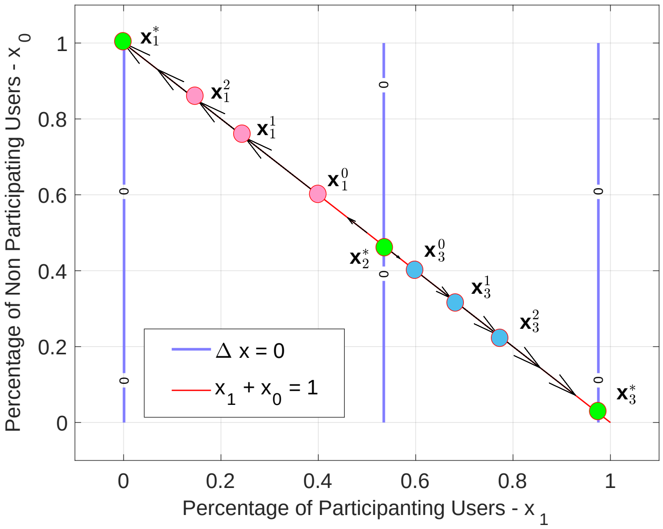

Figure 1: Dynamics of the social state $\boldsymbol{x} = (x_1, x_0)$. From the initial social state $\boldsymbol{x}_1^0 = (0.4, 0.6)$, the social state will gradually evolve to $\boldsymbol{x}_1^1$, $\boldsymbol{x}_1^2$, ..., and eventually achieve the equilibrium $\boldsymbol{x}_1^*$ located in the top-left corner. From the initial social state $\boldsymbol{x}_3^0 = (0.6, 0.4)$, the social state will evolve to the equilibrium $\boldsymbol{x}_3^*$ located in the bottom-right corner.

We let $\boldsymbol{x}_m^*, \forall m \in \{1, 2, \ldots, M\}$ denote the $m$-th Nash equilibrium in a population game, where $M$ denotes the number of Nash equilibrium in this game. Moreover, for $m \in \{1, 2, \ldots, M\}$, let $\boldsymbol{x}_m^0$ represent the initial social state before the social state eventually achieves the equilibrium $\boldsymbol{x}_m^*$. We use $\boldsymbol{x}_m^t$, $t = 1, 2, \ldots$ to show the social state at time slot $t$ given the initial social state $\boldsymbol{x}_m^0$. Thus, from the initial social state $\boldsymbol{x}_m^0$, the social state will evolve by following the path $\boldsymbol{x}_m^1, \boldsymbol{x}_m^2, \ldots$ and eventually achieve the equilibrium $\boldsymbol{x}_m^*$.

In this example, given $\pi = 0.16$ and $1/\epsilon = 0.1$, there are three equilibrium points $\boldsymbol{x}_1^*$, $\boldsymbol{x}_2^*$ and $\boldsymbol{x}_3^*$, and which will eventually emerge depends on the initial social state. For example, from the initial social state $\boldsymbol{x}_1^0 = \{0.4, 0.6\}$, the social state will change following the route $\boldsymbol{x}_1^0 \rightarrow \boldsymbol{x}_1^1 \rightarrow \boldsymbol{x}_1^2 \rightarrow \ldots \rightarrow \boldsymbol{x}_1^*$ as shown by the dashed arrow. From the initial social state $\boldsymbol{x}_3^0 = \{0.6, 0.4\}$, the social state will follow the route $\boldsymbol{x}_3^0 \rightarrow \boldsymbol{x}_3^1 \rightarrow \boldsymbol{x}_3^2 \rightarrow \ldots \rightarrow \boldsymbol{x}_3^*$. Note that no initial social state other than $\boldsymbol{x}_2^* = \{0.53, 0.42\}$ will converge to the Nash equilibrium $\boldsymbol{x}_2^*$. In fact, the equilibria $\boldsymbol{x}_1^*$ and $\boldsymbol{x}_3^*$ are stable in the sense that a small fluctuation around these equilibria will not drive the social state away from the equilibrium, while $\boldsymbol{x}_2^*$ is unstable as a tiny fluctuation on this equilibrium will drive the social state to a different value.

We formalize the stability of the Nash Equilibrium in the next two results.

**Definition 4.** *An equilibrium is stable (also called an evolutionarily stable state) if it is not affected by any small fluctuation around it (Sandholm, 2010).*

The stable equilibrium characterizes the robustness of an equilibrium to a small change in social state. Studying the stability of an equilibrium is important as it begs the questions of how equilibrium is established in the first place and whether this equilibrium will persist in the face of occasional small disturbances in behaviour. We characterize the uniqueness of the stable Nash equilibrium in the following proposition and theorem. The detailed proofs of all theorems and propositions are given in the appendices.

**Proposition 2.** *For the given number of potential users $N \geq 1$ and any payment $\pi$, there exists a unique stable Nash equilibrium if*

$$\max_{x_1 \in [0.1]} \frac{\mathrm{d}f(x_1, N)/\mathrm{d}x_1}{f(x_1, N) + 1} < \frac{1}{K}, \tag{10}$$

*where $K = \max_{\theta \in [0, \bar{\theta}]} h(\theta)\theta$.*

Notice that condition (10) is sufficient but not necessary for uniqueness. In particular, we observe through numerical simulations that in some cases (*e.g.*, $\theta \sim U(0,1)$ or $\theta \sim Beta(5,2)$), there still exists a unique stable Nash equilibrium for any payment $\pi$ although the condition (10) is violated. Nevertheless, the condition in (10) leads to the insight that if the number of participants that a user believes does not grow too rapidly (*i.e.*, $K \cdot \mathrm{d}f(x_1, N)/\mathrm{d}x_1$ is smaller than $f(x_1, N) + 1$ for any $x_1 \in [0, 1]$ ), there exists a unique stable Nash equilibrium.

Suppose the uniqueness condition in Proposition 2 is satisfied. Then based on the optimal participant mass $x_1^*$ derived from (9), we characterize the unique stable Nash equilibrium by the following theorem[11].

**Theorem 1.** *For any payment $\pi$ and noise level $1/\epsilon$, the unique stable Nash equilibrium is given by*

*(a) If $\pi N/\epsilon \geq \bar{\theta}$, there is a unique Nash equilibrium $\boldsymbol{x}^* = (1, 0)$, where the number of participants is*

$$n = N, \tag{11}$$

*and each user's best response is*

$$\ell_j^* = 1, \ \forall j \in \mathcal{N}. \tag{12}$$

*(b) If $\pi N/\epsilon < \bar{\theta}$, there is a unique Nash equilibrium $\boldsymbol{x}^* = (x_1^*, 1 - x_1^*)$, where the mass of participants $x_1$ is the solution of*

$$H\left(\frac{\pi \cdot [x_1 \cdot (N-1) + 1]}{\epsilon}\right) - x_1 = 0, \tag{13}$$

[11] Note that the uniqueness condition in Proposition 2 guarantees the existence of the unique stable Nash equilibrium and allows us to theoretically solve the optimal payment that maximizes the initiator's expected payoff as we will discuss in Section 5. If the assumption in Proposition 2 is violated, as long as a unique stable Nash equilibrium exists for any payment (*e.g.*, the p.d.f. of $\theta$ is uniformly distributed in [0,1] or $\theta \sim Beta(5, 2)$), we can still find the unique stable Nash equilibrium based on Theorem 1. As discussed in Section 1 and as we show in Section 5, it is practically impossible to find the optimal payment that maximizes the initiator's expected payoff if there exists multiple stable Nash equilibrium. This is because the payment strategy couples with the privacy guarantee level and all potential users' strategies.

*the number of participants is* $n = [x_1^* \cdot (N-1)+1]$ *and each user's best response is*

$$\ell_j^* = \begin{cases} 1, & \textit{if } \theta_j \leq \theta^*, \\ 0, & \textit{otherwise}, \end{cases} \tag{14}$$

*where* $\theta^* = \frac{\pi \cdot [x_1^* \cdot (N-1)+1]}{\epsilon}$ *is a unique privacy sensitivity threshold.*

The detailed proof of Theorem 1 is given in Appendix C. In practical terms, this means that if the initiator's payment is not less than the privacy loss of the user with the highest privacy sensitivity, a user should participate irrespective of his own privacy sensitivity. Otherwise, only users with low privacy sensitivities should join.

## 5. Maximizing the Initiator's Payoff

In this section, we study the optimal payment strategy for the initiator to maximize her expected payoff, based on the Nash equilibrium analysis in the previous section.

Based on Theorem 1, we can rewrite the initiator's expected payoff as

$$U^{\text{IN}}(\pi) = -2v \exp\left(-\frac{n^* T^2}{12}\right) - v \exp\left(-\frac{T\epsilon}{2}\right) - \pi \cdot n^*. \tag{15}$$

Recall that the number of participants depends on the payment $\pi$. Directly solving the optimal payment that maximizes (15) is challenging, due to the difficulty in analytically characterizing the impact of the initiator's payment on the number of participants and the initiator's expected payoff. To this end, we transform the original payment maximization problem into an equivalent threshold maximization problem. The key insight is to view the privacy sensitivity $\theta$ as the initiator's strategy and the payment as a function of the threshold. With such a transformation, the initiator's expected payoff is a function of the threshold, which allows us to characterize the property of the initiator's expected payoff analytically.

As Theorem 1 guarantees the uniqueness of the privacy sensitivity threshold $\theta$, there is a one-to-one correspondence between the optimal privacy sensitivity threshold $\theta$ and the payment $\pi$. In this sense, once the initiator chooses the payment $\pi$, she has equivalently determined the privacy sensitivity threshold $\theta$. And any user $j \in \mathcal{N}$ with privacy sensitivity $\theta_j$ not larger than this threshold $\theta$ will join the survey. Hence, we obtain the equivalent threshold maximization problem, where the strategy of the initiator is the privacy sensitivity threshold $\theta$, and the payment $\pi$ is the function of the threshold $\theta$.

Based on Theorem 1, we can derive the payment and the number of participants as the functions of the threshold $\theta$, *i.e.,*

$$\pi = \frac{\epsilon\theta}{(N-1)H(\theta)+1}, \tag{16}$$

$$n = (N-1)H(\theta)+1. \tag{17}$$

Accordingly, the expected payoff of the initiator can be written as

$$U^{\text{IN}}(\theta) = -2v \exp\left(-\frac{[(N-1)H(\theta)+1]T^2}{12}\right) - v\exp\left(-\frac{T\epsilon}{2}\right) - \epsilon\theta. \tag{18}$$

We first show the equivalence between the payment maximization problem and the threshold maximization problem.

**Proposition 3.** *If $\theta^*$ is an optimal solution of* (18)*, then $\pi^*$, calculated by substituting $\theta^*$ into* (16)*, is an optimal solution of* (15)*.*

As per Proposition 3, we can focus on finding the optimal solution of (18). It is clear that a solution $\theta^* \in [0, \bar{\theta}]$ that maximizes (18) exists as the constraint set (*i.e.*, $\theta \in [0,\bar{\theta}]$) is compact and the objective function (*i.e.*, $U^{\text{IN}}(\theta)$) is continuous.

By taking the second order derivative of (18) with respect to $\theta$, we have

$$\begin{aligned}\frac{\mathrm{d}^2\tilde{U}^{\text{IN}}(\theta)}{\mathrm{d}\theta^2} = & \frac{vT^2}{6}(N-1)\exp\left(-\frac{[(N-1)H(\theta)+1]T^2}{12}\right)\\ & \cdot\left[\frac{\mathrm{d}h(\theta)}{\mathrm{d}\theta} - [h(\theta)]^2\frac{(N-1)T^2}{12}\right].\end{aligned} \tag{19}$$

We can easily check that $\mathrm{d}^2\tilde{U}^{\text{IN}}(\theta)/\mathrm{d}\theta^2 < 0$ if $\mathrm{d}h(\theta)/\mathrm{d}\theta \le 0$ for any $\theta \in [0,\bar{\theta}]$. This means if the p.d.f. $h$ of a user's privacy sensitivity is a monotonic non-increasing function, such as a uniform or exponential distribution, the initiator's expected payoff in (18) is strictly concave in $\theta \in [0,\bar{\theta}]$ and there exists a unique solution that maximizes the initiator's expected payoff defined in (18). In such cases, we can obtain the optimal privacy sensitivity threshold $\theta^*$ (*i.e.*, the optimal solution of (18)) by using the following equation:

$$\frac{\mathrm{d}\tilde{U}^{\text{IN}}(\theta)}{\mathrm{d}\theta} = \frac{vT^2}{6}(N-1)h(\theta)\exp\left(-\frac{[(N-1)H(\theta)+1]T^2}{12}\right) - \epsilon. \tag{20}$$

Recall that the value of noise level $1/\epsilon$ is $[0,+\infty]$. By using (20), we obtain the optimal value of $\theta^*$ in Theorem 2 and $\pi^*$ in Theorem 3 given the p.d.f. $h$ is a monotonic non-increasing function for any $\theta \in [0,\bar{\theta}]$. The proofs are given in Appendix E and F, respectively.

**Theorem 2.** *Suppose the p.d.f. $h$ is monotonic and non-increasing in $[0,\bar{\theta}]$. Then, there exists a unique optimal solution $\theta^*$ for the initiator, where for*

*(a) low noise level: $1/\epsilon \le 1/[\frac{vT^2}{6}(N-1)h(0)\exp(-\frac{T^2}{12})]$ the optimal threshold is $\theta^* = 0$ and only users with privacy sensitivity $\theta = 0$ will participate.*

*(b) medium noise level: $1/[\frac{vT^2}{6}(N-1)h(0)\exp(-\frac{T^2}{12})] < 1/\epsilon < 1/[\frac{vT^2}{6}(N-1)h(\bar{\theta})\exp(-\frac{NT^2}{12})]$, the optimal threshold $\theta^*$ satisfies*

$$\frac{vT^2}{6}(N-1)h(\theta)\cdot\exp\left(-\frac{[(N-1)H(\theta)+1]T^2}{12}\right) - \epsilon = 0. \tag{21}$$

*and users with privacy sensitivity $\theta \le \theta^*$ will participate.*

(c) *high noise level:* $1/\epsilon \geq 1/[\frac{vT^2}{6}(N-1)h(\bar{\theta})\exp(-\frac{NT^2}{12})]$*, the optimal threshold* $\theta^* = \bar{\theta}$ *and all the potential users will participate.*

Based on Proposition 3 and Theorem 2, we obtain the following theorem to characterize the optimal payment.

**Theorem 3.** *Suppose the p.d.f.* $h$ *is monotonic and non-increasing in* $[0, \bar{\theta}]$*. Then, there exists a unique optimal solution* $\pi^*$ *for the initiator, where for*

(a) *low noise level:* $\frac{1}{\epsilon} \leq 1/[\frac{vT^2}{6}(N-1)h(0)\exp(-\frac{T^2}{12})]$*, the optimal payment is*

$$\pi^* = 0; \tag{22}$$

(b) *medium noise level:* $1/[\frac{vT^2}{6}(N-1)h(0)\exp(-\frac{T^2}{12})] < \frac{1}{\epsilon} < 1/[\frac{vT^2}{6}(N-1)h(\bar{\theta})\exp(-\frac{NT^2}{12})]$*, the optimal payment is*

$$\pi^* = \frac{\epsilon\theta^*}{(N-1)H(\theta^*)+1}, \tag{23}$$

*where* $\theta^*$ *is the optimal solution of* (21)*.*

(c) *high noise level:* $\frac{1}{\epsilon} \geq 1/[\frac{vT^2}{6}(N-1)h(\bar{\theta})\exp(-\frac{NT^2}{12})]$*, the optimal payment is*

$$\pi^* = \epsilon\bar{\theta}/N. \tag{24}$$

Theorem 3 implies that the initiator's optimal payment increases with the noise level, given the p.d.f. $h$ is a monotonic non-increasing function for any $\theta \in [0, \bar{\theta}]$. As different added noise levels affect the initiator's computational accuracy, Theorem 3 helps the initiator to achieve a good tradeoff between the computational accuracy and her total payment. In more detail, we have the following propositions to characterize the relationship between the initiator's payment $\pi^*$ and the added noise level, given the p.d.f. $h$ is a monotonic non-increasing function for any $\theta \in [0, \bar{\theta}]$.

**Proposition 4.** *Suppose the p.d.f.* $h$ *is monotonic and non-increasing in* $[0, \bar{\theta}]$*. Then with a low noise level,* i.e., $1/\epsilon \leq 1/[\frac{vT^2}{6}(N-1)h(0)\exp(-\frac{T^2}{12})]$*, the optimal payment* $\pi^*$ *is zero and only users with no privacy sensitivity will join the survey.*

Intuitively, it is easier to infer the true computational result $g(D)$ from the published computational result $\mathcal{L}(D)$ in the low noise setting than in the high noise one. With a decrease in the added noise level $1/\epsilon$, a user's privacy revealing risk $\epsilon/n$ increases, which prevents users from participating in the survey. In such cases, the initiator needs to pay a large amount of money to incentivize users to participate. However, when the noise level is too low, *i.e.,* $1/\epsilon \leq 1/[\frac{vT^2}{6}(N-1)h(0)\exp(-\frac{T^2}{12})]$, the benefit regarding the accuracy brought by incentivizing a large number of users to participate is smaller than the payment made to the users. As the initiator's goal is to maximize her expected payoff, she is willing to give up the accuracy in return for a low cost.

**Proposition 5.** *Suppose the p.d.f.* $h$ *is monotonic and non-increasing in* $[0, \bar{\theta}]$*. Then with a high noise level,* i.e., $1/\epsilon \geq 1/[\frac{vT^2}{6}(N-1)h(\bar{\theta})\exp(-\frac{NT^2}{12})]$*, the optimal payment* $\pi^*$ *decreases as the number of potential users* $N$ *increases.*

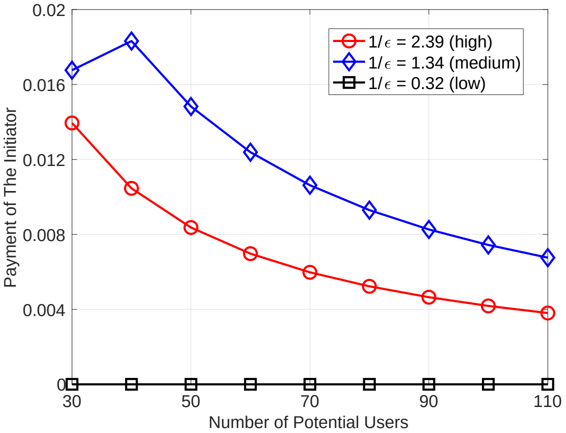

Figure 2: Optimal payment with different noise levels $1/\epsilon$ v.s. different numbers of potential users when $\theta$ is uniformly distributed in $[0, 1]$.

When the noise level is high, a user's privacy revealing risk decreases and the initiator can incentivize users to participate without paying too much. When the number of potential users increases, the number of participants also increases given the same noise level and payment. Due to the network effect, the more users participating in this survey, the more difficult it is to identify whether a given user's private data is included in the computational dataset. This means more users are willing to participate and the initiator can decrease the payment made to the users.

**Proposition 6.** *Suppose the p.d.f. $h$ is monotonic and non-increasing in $[0, \bar{\theta}]$. Then, with a medium noise level, $1/[\frac{vT^2}{6}(N-1)h(0)\exp(-\frac{T^2}{12})] < 1/\epsilon < 1/[\frac{vT^2}{6}(N-1)h(\bar{\theta})\exp(-\frac{NT^2}{12})]$, the relationship between the optimal payment $\pi^*$ and the noise level $1/\epsilon$ depends on the value of $T$ and the functions of $H$ and $h$. For example, when all users' privacy sensitivities are uniformly distributed in $[0, 1]$, $T = 0.4$ and $1/\epsilon = 1.33$, the optimal payment $\pi^*$ increases with the number of potential users when $N \leq 40$, and decreases with the number of potential users $N$ when $N > 40$.*

Proposition 6 reflects the impact of the network effect under a medium noise level. When the number of potential users is not so large, *e.g.*, $N \leq 40$, the network effect's benefit is less obvious as the number of participants is small. Under a medium noise level, the initiator has to increase the payment to incentivize as many participants as possible from the pool of potential users to guarantee the accuracy. When the potential users' number is large enough, the network effect's impact allows the initiator to decrease the payment without jeopardizing the computational accuracy.

In more detail, Figure 2 shows the value of the optimal payment with different numbers of potential users when $T = 0.4$ and a user's privacy sensitivity $\theta$ is uniformly distributed in $[0, 1]$. Note that the performance for other values of $T$ are similar. We can see that the payment is zero and independent of the number of potential users when the noise level is low (*i.e.*, $1/\epsilon = 0.37$) and the payment decreases with the number of potential users when the noise level is high (*i.e.*, $1/\epsilon = 2.39$), as discussed in Propositions 4 and 5. As per Proposition 6, the payment first increases with the number of potential users and then decreases with it when the noise level is medium (*i.e.*, $1/\epsilon = 1.35$). The number of potential users that determines the monotonicity of the payment is $N = 40$ as shown in Figure 2.

Note that Theorem 2, Theorem 3 and Propositions 4-6 only hold in a scenario where the p.d.f. $h$ of a user's privacy sensitivity is monotonic and non-increasing on the interval $[0, \bar{\theta}]$, as such a function guarantees (18) has a unique solution. However, in practice, the privacy sensitivities of users may follow other distributions, like log-normal or beta, which result in multiple solutions that can maximize the initiator's expected payoff defined in (18). As the initiator's expected payoff achieved in all these solutions is the same and the purpose of the initiator is to use minimum cost to obtain high computational accuracy so that her expected payoff is maximized, we have the following theorem.

**Theorem 4** (Optimal Payment under Other Functions of $h$)**.** *Suppose the p.d.f. $h$ is not monotonic or non-increasing on the interval* $[0, \bar{\theta}]$*. Then, if there still exists a unique solution that maximizes the initiator's payoff defined in* (18)*, the initiator selects that solution as the optimal $\theta^*$ and the payment calculated by substituting $\theta^*$ into* (16) *is the optimal payment. Otherwise, the initiator selects the solution that achieves the minimum payment as the optimal solution $\theta^*$ and the payment calculated by substituting $\theta^*$ into* (16) *is the optimal payment.*

Theorem 4 implies that we can still find the optimal payment that maximizes the initiator's payoff under a general distribution. Thus, we can design a mechanism that works under a general distribution of users' privacy sensitivities.

## 6. The PINE Mechanism

Based on the analysis of our two-stage Stackelberg game in Sections 4 and 5, we present the higher-level structure of the PINE mechanism in Algorithm 1, where $\mathcal{N}^{\mathrm{P}}$ denotes the set of participating users.

As shown by the input in Algorithm 1, the payment announced by the initiator is calculated based on Theorem 3 or 4 (depending on the p.d.f. $h$ of a user's privacy sensitivity). Here, the value of the payment depends on the noise level $1/\epsilon$ that is predetermined by the initiator, the initiator's targeted error $T$ and the number of potential users $N$ that the initiator can ask to contribute data. Steps 2-8 show a type-$\theta_j$ user's behaviour which is characterized by Theorem 1. After collecting data from the participating users, the initiator computes and publishes the average of the dataset as described in Section 3.1 and shown by Steps 9-11.

We also characterize the properties of PINE in the following proposition.

**Proposition 7.** PINE *is Individually Rationality (IR), exhibits $\epsilon/n$-differential privacy and maximizes the initiator's payoff defined in* (15)*.*

Proposition 7 shows that PINE incentivizes users to participate as it satisfies the IR condition. Meanwhile, PINE maximizes the initiator's expected payoff (as shown by Theorems 3 and 4), while providing all participants with a differential privacy guarantee (as discussed in Section 3.2).

**Algorithm 1** PINE

```
Input: the noise level: $1/\epsilon$
       the number of potential users: $N$
       the target error: $T$
       the payment: $\pi$   %% if p.d.f. $h$ is monotonic and non-increasing, then
                            %% calculate $\pi$ based on Theorem 3, otherwise use
                            %% Theorem 4.

procedure
    for user $j \in \mathcal{N}$ whose privacy sensitivity is $\theta_j \in [0, \bar{\theta}]$ do
        if $\pi N/\epsilon \geq \bar{\theta}$ then
            $\ell_j = 1$, i.e., reporting data $d_j \in \{0, 1\}$ to the initiator;
        else if $\pi N/\epsilon < \bar{\theta}$ and $\theta_j \leq \theta^*$ (the value of $\theta^*$ is determined by Theorem 1) then
            $\ell_j = 1$, i.e., reporting data $d_j \in \{0, 1\}$ to the initiator;
        else
            $\ell_j = 0$, i.e., not to participates.
    Initiator computes the average of the dataset: $g(D) = \frac{1}{|\mathcal{N}^{\mathrm{P}}|} \sum_{j \in \mathcal{N}^{\mathrm{P}}} d_j$.
    Initiator adds noise $w$ to the computational result: $\mathcal{L}(D) = g(D) + w$, where $w \sim Lap(0, 1/\epsilon)$
    Initiator publishes the differentially private result $\mathcal{L}(D)$.
```

## 7. Empirical Evaluation

In this section, we conduct an empirical study to evaluate the performance of PINE. This complements our theoretical analysis. In particular, the latter concentrates on how to utilize the network effect to maximize the initiator's expected payoff, while in this section we are interested in quantifying the benefit of utilizing the network effect in terms of the initiator's expected payoff and how the number of potential users affects this benefit. We first focus on synthetic data. This allows us to systematically control the properties of data such as the distribution of the users' privacy sensitivities, the accuracy of the users' data and the number of potential users so that the performance of PINE is solely affected by the system parameters that we are interested in. However, as we are also interested in real crowdsourcing settings, we then evaluate the performance of PINE on two real-world data sets. Because movie rating and survey investigation are two of the most common applications in crowdsourcing systems (Howe, 2008), we select standard data sets from these domains. Specifically, we use the Netflix Prize Data Set (Narayanan & Shmatikov, 2008) and the census income data set (Dheeru & Karra Taniskidou, 2017).

Note that given a dataset (either synthetic or real-world), the true mean of the dataset and the mean value of the added noise level, we can calculate the empirical values of $Pr\big[|g(D)+w-\mu| \geq T\big]$ which we refer to the failure probability. As the worst failure probability defined in (3) is the upper bound of failure probability, the value of failure probability which depends on the users' contributed data and the realized added noise level is a better reflection of the performance of mechanisms in practical settings. Thus, we calculate the expected payoff of the initiator in this section based on the following equation

$$\bar{U}^{\text{IN}}(v, n, T, w, \mu, \pi) = -v \cdot Pr\Big[\big|g(D) + w - \mu\big| \geq T\Big] - \pi \cdot n. \quad (25)$$

Compared with the initiator's expected payoff defined in (4), the difference comes from the initiator's utility obtained from the computation which is calculated based on her failure probability instead of the worst failure probability.

In the following, we first describe the benchmarks (Section 7.1), then detail our results using synthetic data (Section 7.2), before moving onto real-world data sets in Section 7.3.

### 7.1 Benchmarks

To evaluate PINE, we compare its performance to a number of benchmark methods:

1. Simple: Similar to PINE, except this mechanism does not consider the network effect and results in an $\epsilon$ privacy revealing risk to a user that is independent of the number of participants (see Appendix L).

2. REAP: As detailed in Section 2, REAP incentivizes all users and adds different noise levels to users' contributing data. The payment to a user is determined by all users' privacy sensitivities that are elicited by the initiator from users. Although REAP minimizes the initiator's worst failure probability given a fixed budget, we revise it to maximize the initiator's expected payoff defined in (4) for a fair comparison[12].

3. MinCostAuction: As detailed in Section 2, MinCostAuction determines the optimal number of participants to minimize the total payment of the initiator given a target error. Thus, for a fair comparison, we apply MinCostAuction by revising its payment to an incentive for a subset of users' participation so that the initiator's expected payoff defined in (4) is maximized. Furthermore, MinCostAuction requires the added noise level to be $1/\epsilon = T/(2 \cdot \ln(3))$ to guarantee $T$-accuracy (*i.e.*, the difference between the initiator's published computational result and the true value of the population is less than $T$). Thus, we also let the noise be $1/\epsilon = T/(2 \cdot \ln(3))$ in all our experiments.

### 7.2 Results in the Synthetic Data Settings

In this section, we study how much the initiator's expected payoff (calculated based on (25)) is affected by utilizing the network effect. We also study the impact of the number of potential users $N$ on the value of the network effect.

Throughout our simulations, let the number of potential users $N$ range from[13] 50 to 250 and let these users' data take the value 1 with probability 0.25 and the value 0 with probability 0.75 (*i.e.*, the true mean of the users' contributed data is $\mu = 0.75$). We run the simulation 5000 times and randomly generate the users' contributed data in each round. We calculate the initiator's failure probability and expected payoff based on the average

[12] With such a revision, REAP aims to choose the optimal payment that maximizes the initiator's expected payoff, which is the same as the goal of PINE.

[13] Recall that this is the number of users for an individual task, not all the people in the crowdsourcing system. Even for real crowdsourcing systems, the number of users who perform each given task is typically in this range (Hara et al., 2018; Difallah et al., 2018).

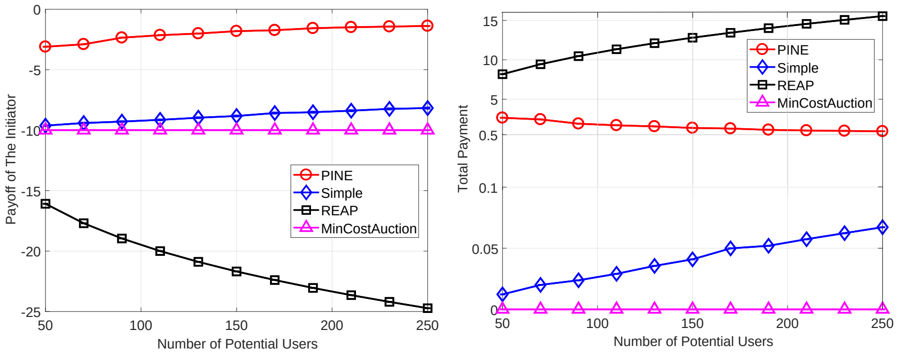

Figure 3: (a) The expected payoff of the initiator and (b) the total payment of the initiator vs. different numbers of potential users.

value in these 5000 runs. We also fix the monetary punishment $v = 10$, the initiator's target error $T = 0.8$ and the corresponding noise level $1/\epsilon = T/(2 \cdot \ln(3)) = 0.36$. To verify the assumption in Theorem 3, we assume each user's privacy sensitivity $\theta$ is uniformly distributed in $[0, 1]$. Note that the p.d.f. of the uniform distribution is monotonic and non-increasing so that we can calculate the optimal payment based on Theorem 3. As we discussed in Section 4, there still exists a unique stable participation equilibrium for any payment $\pi$ when $\theta \sim U(0, 1)$. Thus, we can obtain the number of participants based on Theorem 1. In all of the graphics, the error bars are too small to be visible.

Figure 3.a illustrates the expected payoff of the initiator with different numbers of users. Here, it can be seen that PINE outperforms the other three mechanisms by utilizing the network effect. The improvement regarding the initiator's expected payoff increases with the number of potential users and is at least 50% when the number of potential users is small (*i.e.*, $N = 50$) and 75% when the number of potential users is large (*i.e.*, $N = 250$). This is because with the increase in the number of potential users, the benefit of the network effect also increases. Hence, PINE can decrease the total payment (see Figure 3.b) to attract more participants (see Figure 4.a). Meanwhile, the increased number of participants can lower the failure probability (see Figure 4.b). The decreased payment and the failure probability increase the initiator's expected payoff. Figure 3.a also shows the expected payoff of the initiator achieved by PINE is a concave function of the number of potential users due to the diminishing marginal effect brought by the network effect.

In Figure 3.a, the Simple mechanism increase with the number of potential users $N$. This is because two things happen when $N$ increases: (i) more users are recruited to contribute data, which reduces the initiator's failure probability (see Figure 4.b) and (ii) the initiator recruits more users with litter increasing in her payment (see Figures 3.a and 4.a). MinCostAuction is independent of the number of potential users $N$ as it does not incentivize any user to participate. This is because the benefit regarding the accuracy brought by incentivizing users to participate is smaller than the payment made to the users. As the initiator's goal is to maximize her expected payoff, she is willing to give up the accuracy in return for a low cost.

Figure 3.a also shows that the performance of REAP decreases with the number of potential users. This is because REAP incentivizes all users to participate so its total

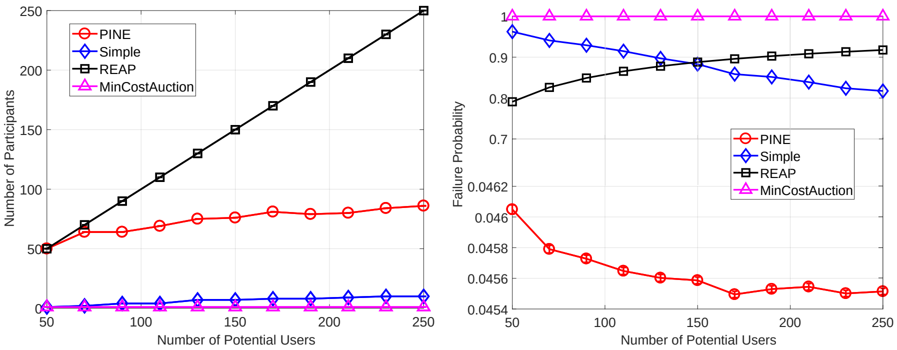

Figure 4: (a) The number of participants and (b) the failure probability vs. different numbers of potential users.

payment increases with the number of potential users (see Figure 3.b). Meanwhile, REAP adds different noise levels to every user's contributing data. With the increasing number of participants, the data collected by the initiator contains more noise which makes it more difficult for the initiator to infer the true mean. Thus, the failure probability of REAP increases with the number of potential users (see Figure 4.b). With the increasing payment and failure probability, the initiator's expected payoff achieved by REAP decreases.

To conclude this section, we note that PINE outperforms other benchmarks by effectively utilizing network effects. The largest performance gain is achieved with a high number of potential users, as the network effect increases with the number of participants (see discussion in Section 3.2) and the latter depends on the number of potential users (see Theorem 1 and 3). Thus, PINE is particularly beneficial for settings with a large number of potential users.

### 7.3 Results in the Real-World Data Settings

While we have so far used synthetic data to show the impact of the number of potential users on the performance of PINE, we now explore PINE's performance in realistic settings. This complements the synthetic data setting, because the mechanism has to cope with actual behaviour rather than mathematically convenient assumptions.

Considering the Netflix data set first. We pick the movie with the most ratings from this data set so that we can have the largest number of potential users (*i.e.*, 581) from this data set. Based on an experimental study by Huberman et al. (2005) on users' general attitudes towards privacy[14], we assume 22% of participants are privacy fundamentalists who are extremely concerned about any use of their data (*i.e.*, their privacy sensitivities $0.78 < \theta_j \leq 1$), 48% are the pragmatic majority who are concerned about their data, but less so than the fundamentalists (*i.e.*, their privacy sensitivities are $0.3 < \theta_j \leq 0.78$), 27% are marginally concerned and generally willing to provide data (*i.e.*, their privacy sensitivities are $0.03 <$

[14] In the experiment of Huberman et al. (2005), they collected 128 answers from users to the question "How important to you is your personal privacy information" and clustered these users' attitudes into four types: Critical, Very Important, Somewhat Important and Not Important. As shown by Figure 3 in their paper, the price requested by users for sharing their information increases with their attitude types. Thus, we need to pay more to users who value their privacy more.

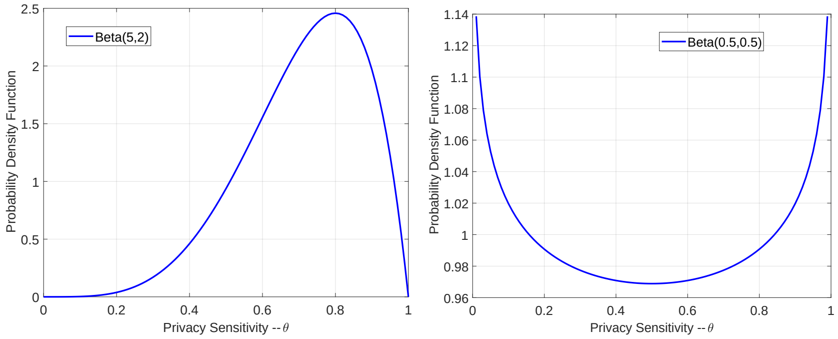

Figure 5: Density plot for different Beta distributions for generating $\theta$: (a) $\theta \sim Beta(5, 2)$ and (b) $\theta \sim Beta(0.5, 0.5)$.

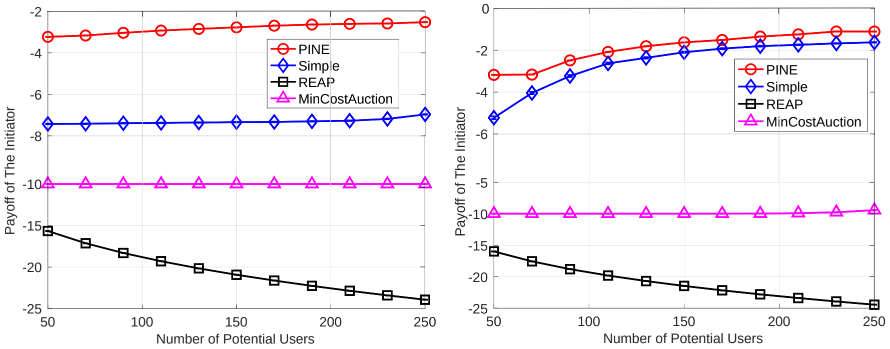

Figure 6: The expected payoff of the initiator vs. different numbers of potential users when (a) $\theta \sim Beta(5, 2)$ and (b) $\theta \sim Beta(0.5, 0.5)$.

$\theta_j \leq 0.3$), and 3% do not really value their privacy (*i.e.*, their privacy sensitivities $0 \leq \theta_j \leq 0.03$). Thus, we put a $Beta(5, 2)$ distribution for each user's privacy sensitivity $\theta$ as shown in Figure 5.a. To characterize the impact of $\theta$'s distribution on the mechanisms' performance, we also evaluate the performance when $\theta$ follows a $Beta(0.5, 0.5)$ distribution (shown in Figure 5.b) where most users are either always willing (*i.e.*, their privacy sensitivities are $0 \leq \theta_j \leq 0.03$) or reluctant (*i.e.*, their privacy sensitivities are $0.9 < \theta_j \leq 1$) to share their data. By comparing Figure 5.a with Figure 5.b, we can see that these two distributions show highly opposing distributions of users' attitude towards privacy.

We can check that the sufficient condition in Proposition 2 is violated in the case where $\theta \sim Beta(5, 2)$ and $\theta \sim Beta(0..5, 0.5)$ for any $N > 1$. However, our numerical simulations show that there still exists a unique stable participation equilibrium. Thus, the number of participants when $\theta$ follows these two distributions can be obtained based on Theorem 1. As the p.d.f. of Beta distribution is not monotonic and non-increasing, we obtain the payment of the initiator based on Theorem 4.

In this setting, the movie's true rating is the mean value of all the users' ratings provided by this data set (*i.e.*, 581 ratings in total). The initiator estimates the movie's rating by averaging all participating users' ratings. We run 5000 times in this simulation and randomly

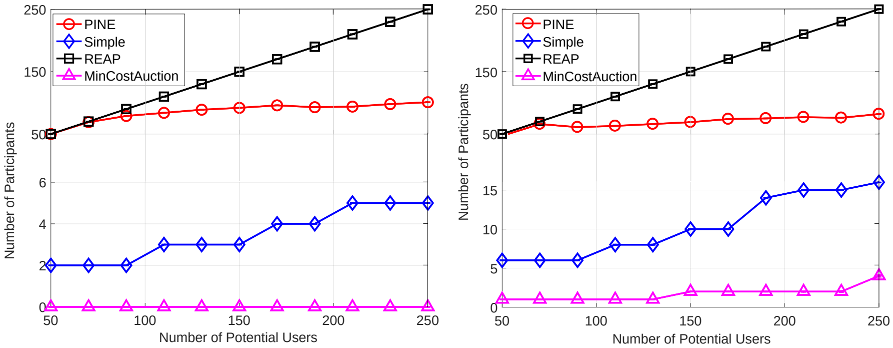

Figure 7: The number of participants vs. different numbers of potential users when (a) $\theta \sim Beta(5,2)$ and (b) $\theta \sim Beta(0.5,0.5)$.

select $N$ users out of the 581 users as potential users in each round and let each potential user's privacy sensitivity be randomly drawn from a beta distribution based on the experimental study by Huberman et al. (2005). We calculate the initiator's failure probability based on the average value in these 5000 runs. We fix the monetary punishment $v = 10$, the initiator's target error $T = 0.8$ and the corresponding noise level $1/\epsilon = T/(2 \cdot \ln(3)) = 0.36$ for illustrative purpose (the results for other distributions and other values of $T$ are broadly similar).

Figure 6 illustrates the expected payoff of the initiator with different numbers of potential users ($N$ from 50 to 250) when (a) $\theta \sim Beta(5,2)$ and (b) $\theta \sim Beta(0.5,0.5)$. The error bars in Figure 6 are too small to be visible. From both Figure 6.a and Figure 6.b, we can see that PINE outperforms the three benchmark mechanisms and similar observations can be drawn with those from the synthetic data in Section 7.2. By comparing Figure 6.a with Figure 6.b, we can see the network effect has a higher impact on the initiator's expected payoff when $\theta \sim Beta(5,2)$ than when $\theta \sim Beta(0.5,0.5)$. This is because when $\theta \sim Beta(5,2)$, most of the users are pragmatists whose concerns about privacy can be significantly reduced by the presence of privacy protection measures. As PINE both adds noise to protect participants' privacy, as do the other three mechanisms, and utilizes the network effect to provide users with further privacy protection, the mechanism can attract a large number of users without increasing the payment significantly (see Figures 7.a and 8.a). However, when $\theta \sim Beta(0.5,0.5)$, most users either do not care about their privacy (*i.e.*, $\theta = 0$) or value their data extremely importantly (*i.e.*, $\theta = 1$). This means incentivizing users with high privacy sensitivities is expensive, so PINE focuses on incentivizing users with a small value of $\theta$ with a small amount of payment (see Figures 7.b and 8.b). However, among these users, most of them have $\theta = 0$ and so their participation strategies are independent of their privacy revealing risk. Thus, the impact of the network effect is less obvious.

Figure 6 also shows that the Simple mechanism has a relatively better performance when $\theta \sim Beta(0.5,0.5)$. This is because when $\theta \sim Beta(0.5,0.5)$, the number of users with a small value of $\theta$ increases. Thus, the Simple mechanism can attract some participants with quite a small amount of money (see Figures 7.b and 8.b). However, when $\theta \sim Beta(5,2)$, as shown by Figure 5.a, most users have a median value of privacy sensitivity (*i.e.*, $\theta_j \in$

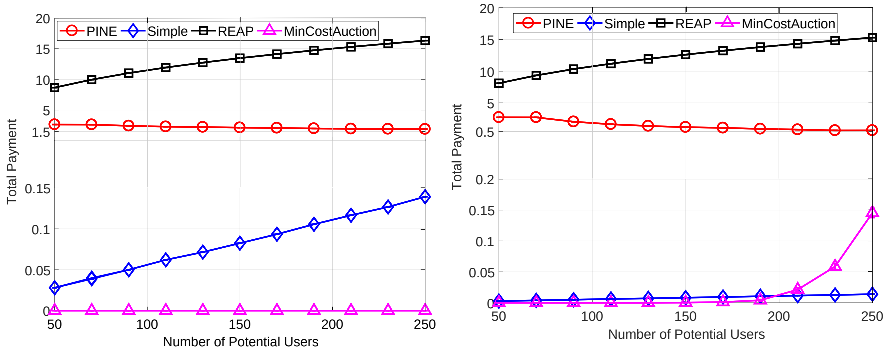

Figure 8: The total payment of the initiator vs. different numbers of potential users when (a) $\theta \sim Beta(5,2)$ and (b) $\theta \sim Beta(0.5,0.5)$.

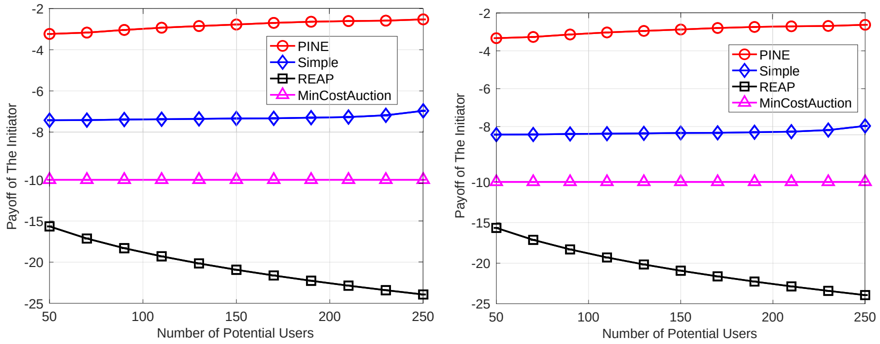

Figure 9: The expected payoff of the initiator vs. different numbers of potential users based on (a) a Nexflix data set and (b) a census income data set when $\theta \sim Beta(5,2)$.

$[0.5, 0.9]$). In such cases, the Simple mechanism needs to pay more to incentivize users to participate (see Figure 8.a).

As the movie rating data set contains a medium number of ratings, we also want to see if we obtain the same results on a larger data set and also from a different domain. Thus, we turn to the census income data set (Dheeru & Karra Taniskidou, 2017) that has 32561 instances. By comparing different mechanisms' performance on these two data sets, we can explore the impact of the data set on the initiator's expected payoff. The standard census income data set contains census income data extracted from the 1994 and 1995 current population surveys conducted by the U.S. Census Bureau. In particular, each instance shows whether a participant's income is below or above the 50K US dollars level. The purpose of the initiator is to compute the percentage of people whose income is above 50K US dollars based on this data set. Similar to our Netflix studies, the true percentage is the mean value of all users' responses provided by this data set (*i.e.*, 32561 responses in total) and the initiator estimates the percentage of people by averaging all participating users' responses. We also run this simulation 5000 times and calculate the initiator's failure probability based on the average value in these 5000 runs. We put a $Beta(5,2)$ distribution for each user's privacy sensitivity $\theta$ to be consistent with our Netflix studies and keep the other parameters the same.

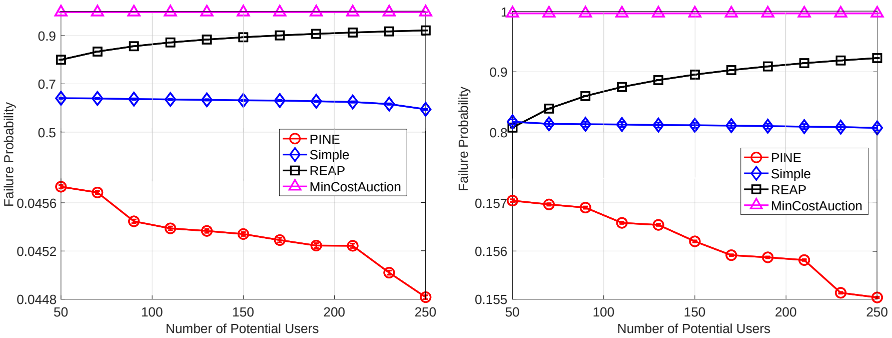

Figure 10: The failure probability vs. different numbers of potential users based on (a) a Nexflix data set and (b) a census income data set when $\theta \sim Beta(5, 2)$.

Figure 9 illustrates the expected payoff of the initiator with different numbers of potential users ($N$ from 50 to 250) based on (a) the Netflix data set and (b) the census income data set. Comparing with these two figures, we can see that the impact of the number of potential users on the initiator's payoff (calculated based on (25)) are broadly similar. However, both figures show that the impact of the network effect in the census data set is higher than that in the Netflix income data set. This is because in the Netflix income data set, 90.8% of the participants rate the movie as good. Thus, even the Simple mechanism which only attracts a few participants does not result in a high failure probability (see Figure 10.a). Thus, even though PINE can utilize the network effect to attract more participants without paying too much, the benefit of reducing the failure probability is small. This means there is only a small difference between PINE and the Simple mechanism.

Based on the real-world data, we can see that PINE consistently outperforms all three benchmarks. However, the empirical evaluation on real-world data sets provides several new insights. First, the performance gain by exploiting the network effect is particularly strong when most potential users are pragmatists who are willing to provide data when the initiator provides sufficient incentivization and some degree of privacy guarantee. Second, although the benefit of utilizing the network effect degrades when the users' contributed data has high accuracy, PINE still outperforms existing benchmarks by decreasing the total payment.

## 8. Conclusions

We have presented a new model that can be used to underpin the design of incentives for crowdsourcing applications with privacy-sensitive users. Our mechanism utilizes network effects to optimize the initiator's payoff and provides all participating users with differential privacy guarantees. Our analysis shows that with the network effect, a user's privacy revealing risk depends on the number of participants. Thus all potential users' participation decisions couple together and form a Population Game. By analyzing this population game, we prove the existence and uniqueness of the equilibrium and show that it depends on the privacy guarantee provided by the initiator, the reward provided by the initiator, and the population size. Based on the equilibrium result, we design a novel mechanism, PINE, that incentivizes appropriate numbers of participants to maximize the initiator's payoff.

Based on our theoretical results, we use synthetic data to systematically explore the benefit of utilizing the network effect and study the impact of system parameters on the performance of PINE. We show that utilizing the network effect can significantly improve the initiator's payoff compared to the state of the art mechanisms by up to 75%. Furthermore, the performance is particularly strong when the population size is large. To further characterize the performance of PINE, we also apply it to two standard real-world data sets. These results show that when most potential users are pragmatists whose concerns about privacy can be significantly reduced by the presence of privacy protection measures and users' contributed data are less accurate, PINE can effectively exploit the network effect to improve the initiator's payoff.

More generally, this work is an important initial step towards a sound theoretical basis for crowdsourcing systems which aim to improve their computational accuracy and provide effective privacy protection for participants with minimum payment. Such guarantees are likely to become increasingly important as people become ever more aware of the value of their data. Although couched in terms of crowdsourcing, this work is applicable to a wide range of applications that require users to send information to data aggregators performing monitoring or control tasks (e.g., database-assisted TV white space networks, cloud-computing systems and smart grid systems). In all of these systems, a data aggregator (*i.e.*, a database in a TV white space network, a data center in a cloud-computing system and a concentrator in a smart grid system) acts like an initiator in our paper. They all collect users' data such as spectrum occupation, server usage and electricity consumption and announce the computational result based on the collected dataset to help users make their resource usage decisions such as accessing new spectrum, selecting a new server and increasing the electricity consumption. As the collected data contains users' private information, privacy-sensitive users will be reluctant to share their data and the data aggregator needs to incentivize them to contribute data with appropriate monetary reward. Due to the similarity between these systems and the crowdsourcing systems we considered in this paper, PINE could be used with minimal adaptation.

There are several possible directions to extend this work. First, as in common in the literature, we assume that there is no correlation between a user's evaluation for his privacy loss and his private data. While this is justifiable in several online settings such as collecting data regarding watching movie behaviour, mobile applications for traffic monitoring, or polls in online communities, there are also several other settings where this assumption is unlikely to hold. For example, consider a database of HIV status: value 1 corresponds to the data owner having HIV, while value 0 means that he does not. Typically, users who have HIV will set a much higher value on privacy of their value 1 than those who do not have HIV. Another example is to consider a database indicating whether a motor vehicle has been damaged. When the database can be publicly accessed, the owner of a damaged car cannot sell the car for the same price as an equivalent undamaged car. Thus, his privacy evaluation for revealing data is higher than the owner of an undamaged car. Exploring these settings where a user's evaluation for his privacy loss depends on his private data will lead to a different formulation and result. Second, as is also common in the literature, we assume that the noise level is an exogenous system parameter. Again, however, in some cases this may not be appropriate. For example, the initiator may be able to adjust the value of the noise level to improve her computational accuracy (Machanavajjhala et al., 2011).

Hsu et al. (2014) have some preliminary results on how to choose $\epsilon$ without considering the network effect. The detailed analysis for finding the optimal value of $\epsilon$ by considering the network effect, however, is much more challenging. It involves the negotiation process between the initiator and all users and requires the initiator to know private information of users such as their data and their evaluations for their privacy loss. Third, we assume that the initiator is a trusted entity to whom all the users send their private data once they have decided to participate. However, PINE fails to guarantee that the initiator will not collude with a third party for her own personal benefit (*e.g.*, the initiator can sell this data to advertisement companies to make profit) or due to legal commitments (*e.g.*, a government may force the initiator to reveal the private data). To date, there is no effective method preventing the initiator from collusion. When an initiator is considered to be untrustworthy, the existing study lets users contribute their randomized data instead of true data to the initiator to protect users' privacy (Jin et al., 2016). However, such a method will decrease the computational accuracy of the initiator due to the randomness added to users' data. Thus, designing a mechanism to prevent this collusion without jeopardizing the benefits of the initiator and users deserves careful study in future work.

## Appendix A. Proof of Proposition 1

Recall that in (8), $\Delta x(x_1^t) = x_1^t - x_1^{t-1}$. Once the value of $\Delta x(x_1^t)$ is zero in slot $t$, we have $x_1^t = x_1^{t-1}$ which means that users will no longer change their participation strategies in the future and the social state remains the same from that time slot onwards. This satisfies the definition of Nash Equilibrium. Note that (7) shows the value of $x_1^t$ depends on the value of $x_1^{t-1}$. Thus, by substituting (7) into (8) and setting (8) equal to zero, we get the result.

## Appendix B. Proof of Proposition 2

Recall that $x_1 + x_0 = 1$ and the value of $x_1$ evolves following a sequence that is generated by (7). Thus, to prove the Nash equilibrium $(x_1^*, x_0^*)$ is unique and stable, we only need to show the sequence always converges to a unique fixed point irrespective of the starting point. As shown by Bertsekas and Tsitsiklis (1989), given a contraction mapping $\boldsymbol{T}$, the update sequence generated by $y^t = \boldsymbol{T}(y^{t-1})$, $t = 1, 2, \ldots$ can converge to a fixed point $y^*$ satisfying $\boldsymbol{T}(y^*) = y^*$ starting from any initial value $y^0 \in \chi$. Here, $\chi$ is a closed subset of $\mathbb{R}^n$. Thus, to prove Proposition 2, we will show that the function $g$ defined in (7) is a contraction mapping on $[0, 1]$ with respect to the absolute value norm if condition (10) is satisfied. A contraction mapping function is defined as follows.

**Definition 5.** *A mapping $\boldsymbol{T} : \chi \to \chi$, where $\chi$ is a closed subset of $\mathbb{R}^n$ is called a contraction if there is a real number $\kappa \in [0, 1)$ such that*

$$\|\boldsymbol{T}(y_1) - \boldsymbol{T}(y_2)\| \leq \kappa \cdot \|y_1 - y_2\|, \; \forall y_1, y_2 \in \chi, \tag{26}$$

*where $\| \cdot \|$ is some norm defined on $\chi$.*

Suppose $\pi = 0$. Then $g(x_1) = 0$ for all $x_1 \in [0, 1]$, and thus $g$ is a contraction with $\kappa = 0$.

Suppose $\pi > 0$. Let $x_{1,a}$ and $x_{1,b}$ be two different real numbers arbitrarily chosen from the interval $[0,1]$. Without loss of generality, we let $x_{1,a} > x_{1,b}$. We will show that

$$|g(x_{1,a}) - g(x_{1,b})| \leq \kappa \cdot |x_{1,a} - x_{1,b}|, \tag{27}$$

where $\kappa = K \cdot \max_{x_1 \in [0.1]} \frac{\mathrm{d}f(x_1,N)/\mathrm{d}x_1}{f(x_1,N)+1}$ and $K = \max_{\theta \in [0,\bar{\theta}]} h(\theta)\theta$. Based on (10), we can easily check that $\kappa \in [0,1)$. Hence, once we prove (27) holds, $g$ is a contraction. As $0 < \frac{\pi \cdot [f(x_{1,b},N)+1]}{\epsilon} < \frac{\pi \cdot [f(x_{1,a},N)+1]}{\epsilon}$, we have the following three settings.

1. $\frac{\pi \cdot [f(x_{1,a},N)+1]}{\epsilon} < \bar{\theta}$ and $\frac{\pi \cdot [f(x_{1,b},N)+1]}{\epsilon} < \bar{\theta}$. Note that $g$ is continuous in $[0,1]$ and differentiable in $[x_{1,b}, x_{1,a}]$. Thus, by the mean value theorem, there exists $x_{1,c} \in (x_{1,b}, x_{1,a})$ such that

$$g'(x_{1,c}) = \frac{g(x_{1,a}) - g(x_{1,b})}{x_{1,a} - x_{1,b}}. \tag{28}$$

Then we have

$$\begin{aligned}
&|g(x_{1,a}) - g(x_{1,b})| \\
&= |g'(x_{1,c})| \cdot |x_{1,a} - x_{1,b}| \\
&= \left| h\left(\frac{\pi \cdot [f(x_{1,c},N)+1]}{\epsilon}\right) \cdot \frac{\pi[f(x_{1,c},N)+1]}{\epsilon} \cdot \frac{f'(x_{1,c},N)}{f(x_{1,c},N)+1} \right| \cdot |x_{1,a} - x_{1,b}| \\
&\leq \kappa \cdot |x_{1,a} - x_{1,b}|.
\end{aligned} \tag{29}$$

2. $\frac{\pi \cdot [f(x_{1,b},N)+1]}{\epsilon} < \bar{\theta} \leq \frac{\pi \cdot [f(x_{1,a},N)+1]}{\epsilon}$. We first let $\bar{x} = \max\{x_1 \in [0,1] \mid \frac{\bar{\theta} \cdot \epsilon}{f(x_1,N)+1} \geq \pi\}$. Note that $x_{1,b} < \bar{x} \leq x_{1,a}$. By applying the mean value theorem to $g$ on the interval $[x_{1,b}, \bar{x}]$, we have

$$|g(\bar{x}) - g(x_{1,b})| \leq \kappa \cdot |\bar{x} - x_{1,b}|. \tag{30}$$

As $g(\bar{x}) = g(x_{1,a}) = 1$ and $\kappa \geq 0$, we have

$$\begin{aligned}
|g(x_{1,a}) - g(x_{1,b})| &= |g(\bar{x}) - g(x_{1,b})| \\
&\leq \kappa \cdot |\bar{x} - x_{1,b}| \\
&\leq \kappa \cdot |x_{1,a} - x_{1,b}|.
\end{aligned} \tag{31}$$

3. $\frac{\pi \cdot [f(x_{1,a},N)+1]}{\epsilon} \geq \bar{\theta}$ and $\frac{\pi \cdot [f(x_{1,b},N)+1]}{\epsilon} \geq \bar{\theta}$. In such a case, $g(x_{1,a}) = g(x_{1,b}) = 1$, and (27) is trivially satisfied.

## Appendix C. Proof of Theorem 1

We first prove that in an equilibrium $\boldsymbol{x}^* = (x_1^*, 1 - x_1^*)$, all users whose privacy sensitivities are below or equal to a certain threshold $\theta^* = \frac{\pi \cdot [x_1^* \cdot (N-1)+1]}{\epsilon}$ participate and all users with stricter privacy concerns, *i.e.*, with privacy sensitivities above $\theta^*$, do not participate, and no user can benefit by deviating from this strategy.

Recall that the payoff of user $j \in \mathcal{N}$ is

$$U^{\text{UE}}(\pi, \theta_j, \epsilon) = \begin{cases} \pi - \theta_j \cdot \dfrac{\epsilon}{n^*}, & \text{if } \ell_j = 1, \\ 0, & \text{if } \ell_j = 0, \end{cases} \tag{32}$$

where $n^* = [x_1^* \cdot (N-1) + 1]$. By checking (32), we can see that every user with $\theta_j \leq \theta^*$ must obtain a non-negative payoff, and every user with $\theta_j > \theta^*$ must receive a negative expected payoff from participating. Hence, $\theta^*$ is an equilibrium threshold that determines the participation decisions of users.

Recall that the value of $\theta$ lies in $[0, \bar{\theta}]$ for all users. Then when $\pi N/\epsilon \geq \bar{\theta}$ and $n^* = N$, we have $\theta^* \geq \bar{\theta}$ and all potential users' privacy sensitivities are below or equal to the optimal threshold $\theta^*$. Thus, all users participate. Only when $\pi N/\epsilon < \bar{\theta}$, we have

$$n^* = [x_1^* \cdot (N-1) + 1] < N, \tag{33}$$

where $x_1^*$ is the solution of $H\left(\frac{\pi \cdot [x_1 \cdot (N-1)+1]}{\epsilon}\right) - x_1 = 0$. Based on the above equation, we have $\theta^* < \bar{\theta}$, and only users with privacy sensitivity $\theta_j \leq \theta^*$ participate.

Based on Proposition 2 and its proof provided in Appendix B, we can prove the uniqueness and stability of the Nash equilibrium.

## Appendix D. Proof of Proposition 3

Based on the proof of Theorem 1, we can see that the number of participants has a unique value given a payment $\pi$. As shown in (17), there exists a one-to-one correspondence between $n$ and $\theta$ due to the monotonic property of c.d.f. $H$. This means that given a payment $\pi$, the threshold $\theta$ is uniquely determined. Thus we can rewrite the initiator's expected payoff, which is a function of payment $\pi$ defined in (15), as a function of the optimal threshold $\theta$. Due to the one-to-one correspondence between the optimal privacy sensitivity threshold $\theta$ and the payment $\pi$, (18) and (15) have the same value given a value $\theta^*$ and $\pi^*$. Here, the payment $\pi^*$ is calculated by substituting $\theta^*$ into (16). This proves the result.

## Appendix E. Proof of Theorem 2

By taking the second order derivative of (18) with respect to $\theta$, we have

$$\begin{aligned} \frac{\mathrm{d}^2 \tilde{U}^{\text{IN}}(\theta)}{\mathrm{d}\theta^2} = {} & \frac{vT^2}{6}(N-1) \exp\left(-\frac{[(N-1)H(\theta)+1]T^2}{12}\right) \\ & \cdot \left[\frac{\mathrm{d}h(\theta)}{\mathrm{d}\theta} - [h(\theta)]^2 \frac{(N-1)T^2}{12}\right]. \end{aligned} \tag{34}$$

We can easily check that $\frac{\mathrm{d}^2 \tilde{U}^{\text{IN}}(\theta)}{\mathrm{d}\theta^2} < 0$ if $\frac{\mathrm{d}h(\theta)}{\mathrm{d}\theta} \leq 0$ for any $\theta \in [0, \bar{\theta}]$, which means, if the p.d.f. $h$ of a user's privacy sensitivity is a monotonic non-increasing function such as a uniform distribution or an exponential distribution, the initiator's expected payoff in (18) is strictly concave in $\theta \in [0, \bar{\theta}]$ and there exists a unique solution that maximizes the initiator's expected payoff defined in (18). In such cases, we can obtain the optimal privacy sensitivity threshold $\theta^*$ (*i.e.*, the optimal solution of (18)) by checking the following equation:

$$\frac{\mathrm{d}\tilde{U}^{\text{IN}}(\theta)}{\mathrm{d}\theta} = \frac{vT^2}{6} \cdot (N-1) \cdot h(\theta) \cdot \exp\left(-\frac{[(N-1)H(\theta)+1]T^2}{12}\right) - \epsilon. \tag{35}$$

Recall that the value of noise level $\frac{1}{\epsilon}$ is $[0, +\infty]$, $h$ is strictly positive and non-increasing on $[0, \bar{\theta}]$ and $h(\theta) = 0$ for all $\theta \notin [0, \bar{\theta}]$. When the noise level satisfies $\frac{1}{\epsilon} \leq 1/[\frac{vT^2}{6} \cdot (N-1) \cdot h(0) \cdot \exp(-\frac{T^2}{12})]$, we have

$$\frac{\mathrm{d}\tilde{U}^{\text{IN}}(\theta)}{\mathrm{d}\theta}\Big|_{\theta=0} = \frac{vT^2}{6} \cdot (N-1) \cdot h(0) \cdot \exp\left(-\frac{T^2}{12}\right) - \epsilon \leq 0. \tag{36}$$

This means that the value of $\frac{\mathrm{d}\tilde{U}^{\text{IN}}(\theta)}{\mathrm{d}\theta}$ is always negative given $\theta \in [0, \bar{\theta}]$. Hence, the initiator's expected payoff defined in (18) decreases with the value of $\theta \in [0, \bar{\theta}]$. The optimal threshold that maximizes the expected payoff of the initiator is $\theta^* = 0$.

When the noise level satisfies $\frac{1}{\epsilon} \geq 1/[\frac{vT^2}{6} \cdot (N-1) \cdot h(\bar{\theta}) \cdot \exp(-\frac{NT^2}{12})]$, we have

$$\frac{\mathrm{d}\tilde{U}^{\text{IN}}(\theta)}{\mathrm{d}\theta}\Big|_{\theta=\bar{\theta}} = \frac{vT^2}{6} \cdot (N-1) \cdot h(\bar{\theta}) \cdot \exp\left(-\frac{NT^2}{12}\right) - \epsilon \geq 0. \tag{37}$$

This means that the value of $\frac{\mathrm{d}\tilde{U}^{\text{IN}}(\theta)}{\mathrm{d}\theta}$ is always positive given $\theta \in [0, \bar{\theta}]$. Hence, the initiator's expected payoff defined in (18) increases with the value of $\theta \in [0, \bar{\theta}]$. The optimal threshold that maximizes the expected payoff of the initiator is $\theta^* = \bar{\theta}$.

When the noise level satisfies $1/[\frac{vT^2}{6} \cdot (N-1) \cdot h(0) \cdot \exp(-\frac{T^2}{12})] < \frac{1}{\epsilon}$ and $\frac{1}{\epsilon} < 1/[\frac{vT^2}{6} \cdot (N-1) \cdot h(\bar{\theta}) \cdot \exp(-\frac{NT^2}{12})]$, we have

$$\frac{\mathrm{d}\tilde{U}^{\text{IN}}(\theta)}{\mathrm{d}\theta}\Big|_{\theta=0} = \frac{vT^2}{6} \cdot (N-1) \cdot h(0) \cdot \exp\left(-\frac{T^2}{12}\right) - \epsilon > 0. \tag{38}$$

$$\frac{\mathrm{d}\tilde{U}^{\text{IN}}(\theta)}{\mathrm{d}\theta}\Big|_{\theta=\bar{\theta}} = \frac{vT^2}{6} \cdot (N-1) \cdot h(\bar{\theta}) \cdot \exp\left(-\frac{NT^2}{12}\right) - \epsilon < 0. \tag{39}$$

Due to the monotonicity of $\frac{\mathrm{d}\tilde{U}^{\text{IN}}(\theta)}{\mathrm{d}\theta}$, there exists a unique value $\theta^* \in [0, \bar{\theta}]$ such that

$$\frac{\mathrm{d}\tilde{U}^{\text{IN}}(\theta)}{\mathrm{d}\theta}\Big|_{\theta=\theta^*} = \frac{vT^2}{6} \cdot (N-1) \cdot h(\theta^*) \cdot \exp\left(-\frac{[(N-1)H(\theta^*)+1]T^2}{12}\right) - \epsilon = 0, \tag{40}$$

and this value $\theta^*$ maximizes the expected payoff of the initiator.

## Appendix F. Proof of Theorem 3

Based on Proposition 3, the optimal payment is obtained by substituting the results of Theorem 2 into (16). In more detail, when the noise level satisfies $\frac{1}{\epsilon} \leq 1/[\frac{vT^2}{6} \cdot (N-1) \cdot h(0) \cdot \exp(-\frac{T^2}{12})]$, the optimal payment is

$$\pi\left(\theta = 0\right) = 0; \tag{41}$$

when the noise level satisfies $\frac{1}{\epsilon} \geq 1/[\frac{vT^2}{6}(N-1) \cdot h(\bar{\theta}) \cdot \exp(-\frac{NT^2}{12})]$, the optimal payment is

$$\pi\left(\theta = 1\right) = \frac{\epsilon \cdot \bar{\theta}}{N}; \tag{42}$$

when the noise level satisfies $1/[\frac{vT^2}{6} \cdot (N-1) \cdot h(0) \cdot \exp(-\frac{T^2}{12})] < \frac{1}{\epsilon}$ and $\frac{1}{\epsilon} < 1/[\frac{vT^2}{6} \cdot (N-1) \cdot h(\bar{\theta}) \cdot \exp(-\frac{NT^2}{12})]$, the optimal payment is

$$\pi\left(\theta = \theta^*\right) = \frac{\epsilon \cdot \theta^*}{(N-1)H(\theta^*)+1}, \tag{43}$$

where $\theta^*$ is the optimal solution of (21).

## Appendix G. Proof of Proposition 4

As shown by Theorem 3, the optimal payment is $\pi^* = 0$ when the added noise level satisfies $\frac{1}{\epsilon} \leq 1/[\frac{vT^2}{6} \cdot (N-1) \cdot h(0) \cdot \exp(-\frac{T^2}{12})]$. Recall that only users with non-negative expected payoff will participate in the survey. By checking (6), we can see that only the users with no privacy sensitivity satisfy this condition. This proves the result.

## Appendix H. Proof of Proposition 5

Based on Theorem 3, when $\frac{1}{\epsilon} \geq 1/[\frac{vT^2}{6} \cdot (N-1) \cdot h(\bar{\theta}) \cdot \exp(-\frac{NT^2}{12})]$, the optimal payment is $\pi^* = \frac{\epsilon \cdot \bar{\theta}}{N}$, and this payment decreases as the number of potential users $N$ increases.

## Appendix I. Proof of Proposition 6

When the added noise level $1/[\frac{vT^2}{6} \cdot (N-1) \cdot h(0) \cdot \exp(-\frac{T^2}{12})] < 1/\epsilon$ and $1/\epsilon < 1/[\frac{vT^2}{6} \cdot (N-1) \cdot h(\bar{\theta}) \cdot \exp(-\frac{NT^2}{12})]$, the optimal payment is $\pi^* = \frac{\epsilon \cdot \theta^*}{(N-1)H(\theta^*)+1}$, where $\theta^*$ is the optimal solution of (21). We can see that the relationship between $\pi$ and $\theta$ depends on function $H$, $\epsilon$ and $T$.

## Appendix J. Proof of Proposition 7

We first prove that PINE satisfies the Individual Rationality (IR) condition. Theorem 1 shows that in the Nash equilibrium, every user $j \in \mathcal{N}$ with $\theta_j \in [0, \theta^*]$ participates in the survey. As shown in (32), a type-$\theta$ user's payoff decreases with his privacy sensitivity when he participates in the survey. Hence, we only need to prove that given the payment characterized by PINE, the payoff of a type-$\theta^*$ user must be zero. Otherwise, the initiator can reduce the payment by a small value of $\delta > 0$, which does not violate the IR condition but raises the initiator's expected payoff. By substituting (17) into (32), we have

$$U^{\mathrm{UE}}(\pi, \theta_j, \epsilon) = \pi - \theta^* \cdot \frac{\epsilon}{(N-1)H(\theta^*)+1} = 0. \tag{44}$$

We then prove that PINE supports differential privacy. The sensitivity of a function is the maximum by which the value of that function can change under any two neighbouring datasets. Recall that $g(D) = \frac{1}{n}\sum_{j=1}^{n} d_j$. For any two neighbouring pairs of datasets $D$ and $D'$ where the dataset $D$ includes private data of user $j \in \mathcal{N}$ and the dataset $D'$ does not, the difference between $g(D)$ and $g(D')$ is $1/n$. Hence, the sensitivity of $g(D)$ is $1/n$. Recall that PINE requires the initiator to add noise drawn from the Laplace distribution $Lap(0, 1/\epsilon)$ to her computational result $g(D)$ whose sensitivity is $1/n$. As shown by (Dwork et al., 2006), we have

$$Pr[\mathcal{L}(D) = a] \leq e^{\epsilon/n} \cdot Pr[\mathcal{L}(D') = a], \tag{45}$$

which shows that PINE supports $\epsilon/n$-differential privacy.

Based on the proof of Theorem 3, we can prove that PINE maximizes the initiator's expected payoff.

## Appendix K. Proof of Theorem 4

It is clear that a solution $\theta^* \in [0, \bar{\theta}]$ that maximizes (18) exists as the constraint set (*i.e.*, $\theta \in [0, \bar{\theta}]$) is compact and the objective function (*i.e.*, $U^{\text{IN}}(\theta)$) is continuous. If there is still a unique solution that maximizes the initiator's payoff defined in (18), the optimal payment can be obtained by substituting $\theta^*$ into (16) based on Proposition 3. If there are multiple solutions that maximize the initiator's payoff defined in (18), we can check that the initiator's expected payoffs achieved for all these solutions are the same. As the purpose of the initiator is to use minimum cost to obtain high computational accuracy so that her expected payoff is maximized, the initiator should select the solution that brings the minimum payment. Based on Proposition 3, the corresponding reward calculated based on (16) by substituting the selected optimal solution $\theta^*$ is the optimal solution that maximizes the initiator's expected payoff. Thus, we get the result.

## Appendix L. The Simple Benchmark Mechanism

In order to evaluate the performance of PINE and show the benefit of utilizing the network effect, we need a benchmark mechanism that does not consider the network effect. Thus our Simple mechanism is based on the underlying ideas of Jin et al. (2018); Zhang et al. (2016); and Yang et al. (2018) and determines the initiator's optimal payment strategy and the way of publishing her computational result without including the network effect. In more detail, without knowing the benefit of the network effect, a type-$\theta_j$ user considers his privacy revealing risk is $\epsilon$ which is independent of the number of participants and is estimated based on the added noise level (*i.e.*, $1/\epsilon$) announced by the initiator. Compared to (5), his payoff is given in the following equation:

$$U^{\text{UE}}(\pi, \theta_j, \epsilon) = \begin{cases} \pi - \theta_j \cdot \epsilon, & \text{if } \ell_j = 1, \\ 0, & \text{if } \ell_j = 0, \end{cases} \tag{46}$$

As the user's payoff is decreasing with $\theta$, given the initiator's payment $\pi$, only users with $\theta_j \leq \pi/\epsilon$ will participate. And the number of participants will be $n = (N-1) \cdot H(\pi/\epsilon) + 1$ instead of (13).

As the initiator adds noise that is drawn from the Laplace distribution with mean zero and standard deviation $\sqrt{2}/\epsilon$ to her computational result, the initiator's optimal payment

$\pi^*$ is chosen to maximize the following equation instead of that defined in (15):

$$U^{\text{IN,SIM}}(\pi) = -v \cdot \left[ 2 \exp\left( - \frac{[(N-1) \cdot H(\pi/\epsilon) + 1] T^2}{12} \right) + \exp\left( - \frac{T\epsilon}{2} \right) \right] \\ - \pi \cdot [(N-1) \cdot H(\pi/\epsilon) + 1], \tag{47}$$

Compared to PINE that uses network effects to decrease a user's privacy revealing risk (*i.e.*, a user's privacy revealing risk decreases with the number of participants as shown in (5)), the Simple mechanism does not consider the network effect and a user's privacy revealing risk is a constant value (*i.e.*, $\epsilon$ as shown in (46)).

## References

Ajorlou, A., Jadbabaie, A., & Kakhbod, A. (2016). Dynamic pricing in social networks: The word-of-mouth effect. *Management Science*.

Apple (2017a). Differential privacy overview. `https://images.apple.com/privacy/docs/Differential_Privacy_Overview.pdf`.

Apple, D. P. T. (2017b). Learning with privacy at scale. `https://machinelearning.apple.com/2017/12/06/learning-with-privacy-at-scale.html`.

Bertsekas, D. P., & Tsitsiklis, J. N. (1989). *Parallel and distributed computation: numerical methods*, Vol. 23. Prentice hall Englewood Cliffs, NJ.

Brandt, F., & Sandholm, T. (2008). On the existence of unconditionally privacy-preserving auction protocols. *ACM Transactions on Information and System Security*, *11*(2), 6.

Candogan, O., Bimpikis, K., & Ozdaglar, A. (2010). Optimal pricing in the presence of local network effects. In *International Workshop on Internet and Network Economics*, pp. 118–132. Springer.

Chandra, P., Gujar, S., & Narahari, Y. (2017). Referral-embedded provision point mechanisms for crowdfunding of public projects. In *International Conference on Autonomous Agents and Multiagent Systems*, pp. 642–650.

Chatzikokolakis, K., Palamidessi, C., & Stronati, M. (2014). A predictive differentially-private mechanism for mobility traces. In *International Symposium on Privacy Enhancing Technologies Symposium*, pp. 21–41. Springer.

Chen, Y., Li, B., & Zhang, Q. (2016a). Incentivizing crowdsourcing systems with network effects. In *International Conference on Computer Communications*, pp. 1–9.

Chen, Y., Chong, S., Kash, I. A., Moran, T., & Vadhan, S. (2016b). Truthful mechanisms for agents that value privacy. *ACM Transactions on Economics and Computation*, *4*(3), 13.

Dandekar, P., Fawaz, N., & Ioannidis, S. (2014). Privacy auctions for recommender systems. *ACM Transactions on Economics and Computation*, *2*(3), 12.

Dheeru, D., & Karra Taniskidou, E. (2017). UCI machine learning repository..

Difallah, D., Filatova, E., & Ipeirotis, P. (2018). Demographics and dynamics of mechanical turk workers. In *International Conference on Web Search and Data Mining*, pp. 135–143.

Dunyak, A., & Zhu, Q. (2018). Understanding mean-field effects of large-population user data obfuscation in machine learning. In *Annual Conference on Information Sciences and Systems*, pp. 1–6.

Dwork, C. (2008). Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation*, pp. 1–19. Springer.

Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pp. 265–284. Springer.

Easley, D., & Kleinberg, J. (2010). *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge University Press.

Erlingsson, Ú., Pihur, V., & Korolova, A. (2014). Rappor: Randomized aggregatable privacy-preserving ordinal response. In *ACM Conference on Computer and Communications Security*, pp. 1054–1067.

Feigenbaum, J., Jaggard, A. D., & Schapira, M. (2010). Approximate privacy: foundations and quantification. In *ACM Conference on Electronic Commerce*, pp. 167–178.

Fudenberg, D., & Tirole, J. (1991). Game theory. Tech. rep., The MIT press.

Gao, Y., Toni, F., Wang, H., & Xu, F. (2016). Argumentation-based multi-agent decision making with privacy preserved. In *International Conference on Autonomous Agents and Multiagent Systems*, pp. 1153–1161.

Gatti, N., Lazaric, A., Rocco, M., & Trovò, F. (2015). Truthful learning mechanisms for multi-slot sponsored search auctions with externalities. *Artificial Intelligence*, *227*, 93–139.

Geng, Q., & Viswanath, P. (2014). The optimal mechanism in differential privacy. In *IEEE International Symposium on Information Theory*, pp. 2371–2375.

Ghosh, A., & Kleinberg, R. (2014). Optimal contest design for simple agents. In *ACM Conference on Economics and Computation*, pp. 913–930.

Ghosh, A., & Ligett, K. (2013). Privacy and coordination: Computing on databases with endogenous participation. In *ACM Conference on Electronic Commerce*, pp. 543–560.

Ghosh, A., & Roth, A. (2015). Selling privacy at auction. *Games and Economic Behavior*, *91*, 334–346.

Gradwohl, R. (2017). Information sharing and privacy in networks. In *ACM Conference on Economics and Computation*, pp. 349–350.

Hara, K., Adams, A., Milland, K., Savage, S., Callison-Burch, C., & Bigham, J. P. (2018). A data-driven analysis of workers' earnings on amazon mechanical turk. In *Conference on Human Factors in Computing Systems*, p. 449.

Ho, C.-J., Slivkins, A., & Vaughan, J. W. (2016). Adaptive contract design for crowdsourcing markets: Bandit algorithms for repeated principal-agent problems. *Journal of Artificial Intelligence Research*, *55*, 317–359.

Hoeffding, W. (1994). Probability inequalities for sums of bounded random variables. In *The Collected Works of Wassily Hoeffding*, pp. 409–426. Springer.

Howe, J. (2008). *Crowdsourcing: How the power of the crowd is driving the future of business*. Random House.

Hsu, J., Gaboardi, M., Haeberlen, A., Khanna, S., Narayan, A., Pierce, B. C., & Roth, A. (2014). Differential privacy: An economic method for choosing epsilon. In *IEEE Computer Security Foundations Symposium*, pp. 398–410.

Huberman, B. A., Adar, E., & Fine, L. R. (2005). Valuating privacy. *IEEE Security & Privacy*, *3*(5), 22–25.

Jain, S., Ghalme, G., Bhat, S., Gujar, S., & Narahari, Y. (2016). A deterministic mab mechanism for crowdsourcing with logarithmic regret and immediate payments. In *International Conference on Autonomous Agents and Multiagent Systems*, pp. 86–94.

Jain, S., Gujar, S., Bhat, S., Zoeter, O., & Narahari, Y. (2018). A quality assuring, cost optimal multi-armed bandit mechanism for expertsourcing. *Artificial Intelligence*, *254*, 44–63.

Jin, H., Su, L., Chen, D., Nahrstedt, K., & Xu, J. (2015). Quality of information aware incentive mechanisms for mobile crowd sensing systems. In *ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pp. 167–176.

Jin, H., Su, L., Xiao, H., & Nahrstedt, K. (2016). Inception: Incentivizing privacy-preserving data aggregation for mobile crowd sensing systems. In *ACM Mobile Ad Hoc Networking and Computing*, Vol. 16, pp. 341–350.

Jin, H., Su, L., Xiao, H., & Nahrstedt, K. (2018). Incentive mechanism for privacy-aware data aggregation in mobile crowd sensing systems. *IEEE/ACM Transactions on Networking (TON)*, *26*(5), 2019–2032.

Kamar, E., & Horvitz, E. (2012). Incentives for truthful reporting in crowdsourcing. In *International Conference on Autonomous Agents and Multiagent Systems*, pp. 1329–1330.

Kifer, D., & Machanavajjhala, A. (2014). Pufferfish: A framework for mathematical privacy definitions. *ACM Transactions on Database Systems*, *39*(1), 3.

Krause, A., & Horvitz, E. (2010). A utility-theoretic approach to privacy in online services. *Journal of Artificial Intelligence Research*, *39*, 633–662.

Lev, O., Polukarov, M., Bachrach, Y., & Rosenschein, J. S. (2013). Mergers and collusion in all-pay auctions and crowdsourcing contests. In *International Conference on Autonomous Agents and Multiagent Systems*, pp. 675–682.

Levit, V., Komarovsky, Z., Grinshpoun, T., & Meisels, A. (2018). Incentive-based search for efficient equilibria of the public goods game. *Artificial Intelligence*, *262*, 142–162.

Li, C., Li, D. Y., Miklau, G., & Suciu, D. (2014). A theory of pricing private data. *ACM Transactions on Database Systems*, *39*(4), 34.

Liao, G., Chen, X., & Huang, J. (2018). Social-aware privacy-preserving correlated data collection. In *ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pp. 11–20.

Ligett, K., Neel, S., Roth, A., Waggoner, B., & Wu, S. Z. (2017). Accuracy first: Selecting a differential privacy level for accuracy constrained erm. In *Advances in Neural Information Processing Systems*, pp. 2566–2576.

Lin, K.-Y., & Lu, H.-P. (2011). Why people use social networking sites: An empirical study integrating network externalities and motivation theory. *Computers in Human Behavior*, *27*(3), 1152–1161.

Liu, Y., & Chen, Y. (2017). Sequential peer prediction: Learning to elicit effort using posted prices.. In *AAAI Conference on Artificial Intelligence*, pp. 607–613.

Luo, Y., Gao, L., & Huang, J. (2016). An integrated spectrum and information market for green cognitive communications. *IEEE Journal on Selected Areas in Communications*, *34*(12), 3326–3338.

Machanavajjhala, A., Kifer, D., Abowd, J., Gehrke, J., & Vilhuber, L. (2008). Privacy: Theory meets practice on the map. In *IEEE International Conference on Data Engineering*, pp. 277–286.

Machanavajjhala, A., Korolova, A., & Sarma, A. D. (2011). Personalized social recommendations: accurate or private. *Proceedings of the Very Large Data Base Endowment*, *4*(7), 440–450.

Manshaei, M. H., Freudiger, J., Félegyházi, M., Marbach, P., & Hubaux, J.-P. (2008). On wireless social community networks. In *International Conference on Computer Communications*, pp. 1552–1560.

McSherry, F., & Mironov, I. (2009). Differentially private recommender systems: building privacy into the net. In *International Conference on Knowledge Discovery and Data Mining*, pp. 627–636.

Nair, J., Wierman, A., & Zwart, B. (2015). Provisioning of large-scale systems: The interplay between network effects and strategic behavior in the user base. *Management Science*, *62*(6), 1830–1841.

Narayanan, A., & Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets. In *IEEE Symposium on Security and Privacy.*, pp. 111–125.

Nath, S., & Narayanaswamy, B. M. (2014). Productive output in hierarchical crowdsourcing. In *International Conference on Autonomous Agents and Multiagent Systems*, pp. 469–476.

Pawlick, J., & Zhu, Q. (2016). A stackelberg game perspective on the conflict between machine learning and data obfuscation. In *IEEE International Workshop on Information Forensics and Security*, pp. 1–6.

Ren, S., & Van der Schaar, M. (2012). Data demand dynamics in wireless communications markets. *IEEE Transactions on Signal Processing*, *60*(4), 1986–2000.

Sandholm, W. H. (2010). *Population games and evolutionary dynamics*. MIT press.

Shokri, R., Theodorakopoulos, G., Troncoso, C., Hubaux, J.-P., & Le Boudec, J.-Y. (2012). Protecting location privacy: optimal strategy against localization attacks. In *ACM Conference on Computer and Communications Security*, pp. 617–627.

Tembine, H., Altman, E., El-Azouzi, R., & Hayel, Y. (2009). Evolutionary games in wireless networks. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, *40*(3), 634–646.

Von Stackelberg, H. (1934). *Marktform und gleichgewicht*. J. springer.

Wang, W., Ying, L., & Zhang, J. (2016). The value of privacy: Strategic data subjects, incentive mechanisms and fundamental limits. In *ACM International Conference on Measurement and Modeling of Computer Science*, pp. 249–260.

Xu, J., Zhang, Z., Xiao, X., Yang, Y., Yu, G., & Winslett, M. (2013). Differentially private histogram publication. *The Very Large Data Base Journal*, *22*(6), 797–822.

Yang, L., Zhang, M., He, S., Li, M., & Zhang, J. (2018). Crowd-empowered privacy-preserving data aggregation for mobile crowdsensing. In *ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pp. 151–160.

Zhang, M., Yang, L., Gong, X., & Zhang, J. (2016). Privacy-preserving crowdsensing: Privacy valuation, network effect, and profit maximization. In *IEEE Global Communications Conference*, pp. 1–6.

Zhang, Z., He, S., Chen, J., & Zhang, J. (2018). Reap: An efficient incentive mechanism for reconciling aggregation accuracy and individual privacy in crowdsensing. *IEEE Transactions on Information Forensics and Security*, *13*(12), 2995–3007.

Zhang, Z., Rubinstein, B. I., Dimitrakakis, C., et al. (2016). On the differential privacy of bayesian inference.. In *AAAI Conference on Artificial Intelligence*, pp. 2365–2371.

Zhu, Q., & Rass, S. (2018). Game theory meets network security: A tutorial. In *ACM Conference on Computer and Communications Security*, pp. 2163–2165.