Optimal Privacy-Preserving Data Collection: A Prospect Theory Perspective

Guocheng Liao, Xu Chen, and Jianwei Huang

***Abstract*—We study a mechanism design problem of privacy-preserving data collection with privacy protection uncertainty. A data collector wants to collect enough data to perform a certain computation that benefits the individuals who contribute the data, with the possibility of individual privacy leakage. The data collector adopts a privacy-preserving mechanism by adding some random noise to the computation result, which reduces the accuracy of the computation. Individuals decide whether to contribute data based on the potential benefit and the possible privacy cost induced by the mechanism. Due to the intrinsic uncertainty involved in privacy protection, we model individuals' privacy-aware participation using the prospect theory, which more accurately models individuals' behavior under uncertainty than the traditional expected utility theory. We show that the data collector's utility maximization problem involves a polynomial of high and fractional order, which is difficult to solve analytically. We get around this issue by proposing an approximation method, which allows us to obtain a closed form unique solution of the data collector's decision problem. We numerically show that the approximation error is small when the number of individuals is large. By comparing with the results under the expected utility theory, we conclude that a data collector who considers the more realistic prospect theory modeling should adopt a stricter privacy-preserving mechanism to boost her utility.**

I. Introduction

A. Background and Motivation

An effective utilization of big data becomes increasingly important in various aspects of our daily life in terms of providing new inspirations and opportunities for knowledge generation. Corporations and academic institutions are eager to collect a large amount of data from individuals with an aim to develop new measures of data-driven analysis.

A privacy-aware individual, however, has two primary concerns regarding the data collection. First, he worries about the personal information leakage if the data collector does not provide enough privacy protection. Second, he is interested in knowing whether he can receive certain reward after contributing the data. Such a reward may be related to the computation performed by the data collector over the data. For example, when a patient reports his medical record to researchers, he can obtain an evaluation regarding the probabilities of getting various diseases and the possible ways of prevention. The reward could also be monetary. Individual needs to weigh the trade-off between the reward and the privacy loss to decide whether to participate in the data collection.

From the data collector's perspective, she exploits the collected data to conduct a certain kind of computation. In order to protect the individuals privacy, she needs to adopt a privacy-preserving mechanism that adds some random noise to the computation result, such that an adversary cannot easily infer participants' actual data. The added noise, however, will reduce the accuracy of the computation. The collector needs to carefully design the privacy-preserving mechanism to trade off individuals' satisfaction and computation accuracy.

A series of research on privacy-preserving data collection has emerged recently. Most of these studies used monetary payment as the reward. For instance, the studies in [1]–[3] assumed that individuals will truthfully report data, and the rewards are determined based on participants' declared privacy costs. Wang *et al.* in [4], [5] considered the case where individuals may randomly instead of truthfully report data. The objective of the mechanism designed in [4], [5] is to either minimize the total payment or guarantee the accuracy.

Our work in this paper differs from previous studies [1]–[5] in terms of both reward type and individual behavior modeling. First, we assume that the reward is related to the computation result instead of monetary based. Since the data collector is conducting data-driven computation beneficial to individuals, she can utilize this valuable result as a reward. For example, a data collector can offer a disease evaluation result to patients who contribute their medical records.

Regarding the modeling of the individual behavior, Acquisti and Grossklags in [6] suggested that privacy decision making is affected by both the external incomplete information and the internal bounded rationality. Complex information manipulations like the anticipation of other individuals' strategies and the prediction of reward can be intractable due to incomplete information and limited individual computational resources, so that actual decision-making process may be far from the one predicted under the fully rational agent model. A simple decision of whether to participate [7] considering the result-related reward is more practical for individuals in reality.

Another unique aspect of our model is the consideration of uncertainty. As [6] suggested, the uncertainty of outcomes plays a significant role in privacy decision making, including those related to privacy-preserving data collection. This motivates us to use prospect theory (PT) [8] [9] as the modeling tool, as it normatively and descriptively inter-

Guocheng Liao and Jianwei Huang are with Department of Information Engineering, the Chinese University of Hong Kong; Email: {lg016, jwhuang}@ie.cuhk.edu.hk. Xu Chen is with School of Data and Computer Science, Sun Yat-sen University, Guangzhou, China; Email: chenxu35@mail.sysu.edu.cn. This work is supported by the General Research Fund (Project Number CUHK 14219016) established under the University Grant Committee of the Hong Kong Special Administrative Region, China, and the start-up fund from Sun Yat-sen University and Thousand Talents Program of China.

prets how individuals make decisions by evaluating uncertain gains and losses. However, there does not exist any concrete theoretical analysis in the literature regarding how prospect theory can help us characterize and understand human privacy decision-making. Our paper represents a first step towards understanding this important and under-explored area.

More specifically, PT better captures practical human behavioral characteristics with uncertainty compared with the traditionally widely-used expected utility theory (EUT). The framework of PT captures the impacts of three key aspects of decision making: reference point, S-shape asymmetrical valuation function, and weighting distortion. We will consider the impacts of reference point and S-shape valuation function in this paper. Recently, PT has been successfully applied to decision-making in engineering area such as cognitive radio networks [10], spectrum investment [11], and mobile data [12] trading to generate new engineering insights.

The overall objective of our work is to understand how the data collector should design her privacy-preserving mechanism and how PT analysis makes a difference in her optimal decision. To achieve this objective, we model the interaction between the data collector and individuals as a Stackelberg game. At the individuals' side, we use PT to capture individuals' subjective decision-making under the privacy protection uncertainty. At the data collector's side, a better privacy protection will attract more individuals, but the corresponding higher level of perturbation will degrade the accuracy of the computation. We compute the data collector's optimal strategy based on her prediction of individuals' participation decisions to various privacy protection levels. We compare the results under PT with the EUT benchmark, to understand how the modeling of PT can help the data collector design a better mechanism.

B. Key Contributions

The main contributions of this paper are as follows.

- *PT-aided individual behavior model under privacy protection uncertainty*: Since privacy-preserving mechanism innately involves uncertain outcomes due to the induced random noise, we adopt PT, which is more normatively and descriptively accurate compared with EUT when dealing with uncertainty, to capture how individuals subjectively respond to uncertainties.
- *Analysis of the data collector's utility maximization problem*. Since the data collector's utility maximization problem involves a polynomial of high and fractional order, it is difficult to obtain the analytical solution. We propose an approximation method that allows us to compute a unique optimal solution. The approximation is quite accurate under the realistic situation of a large number of individuals.
- *Practical insights based on the comparison between the PT model and the traditional EUT model*: We compare the results under the general PT model and that under the EUT model, and conclude that the data collector should adopt a stricter privacy-preserving mechanism based on PT. Properly considering individual' behavioral characteristics can increase the data collector's utility.

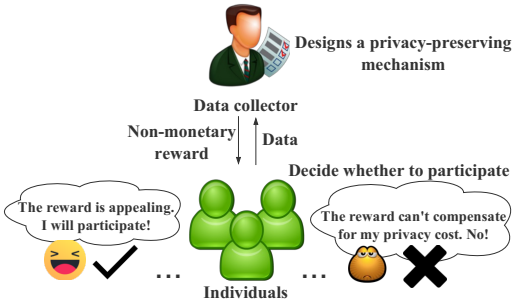

Fig. 1: System model: in Stage I, the data collector initiates a collection with a privacy-preserving mechanism. In Stage II, individual decides whether to participate in the data collection.

The rest of this paper is organized as follows. In Section II, we discuss the system model regarding the individual's participation problem and the data collector's utility maximization problem, respectively. In Section III, we solve these problems and analyze the impact of different PT parameters. In Section IV, we provide simulation results and discuss the corresponding insights. We conclude the paper in Section V.

II. System Model

Fig. 1 illustrates the system model of privacy-preserving data collection problem. In this model, a data collector wants to collect data from individuals, and provides a result-related reward as an incentive. Individuals decide whether to report data. The data collector designs a privacy-preserving mechanism to maximize her utility. We describe the individual's participation problem and the data collector's utility maximization problem in Sections II.A and II.B, respectively.

A. Individual's Participation Problem

In this subsection, we formulate individuals' participation problem under the privacy protection uncertainty. We use PT to model individuals' behavioral characteristics in this context.

The uncertain outcomes faced by individuals are induced by the privacy-preserving mechanism. We use the widely-adopted concept of differential privacy [13] to quantify these outcomes. The concept of ϵ-differential privacy given in Definition 1 serves as a basic framework for measuring privacy in the related literature [1]–[5]. In Definition 1, one entry in a database corresponds to one individual's reported data. A mechanism that is differential private ensures that the result of the computation will not change significantly when an individual's data is added to the database. This ensures that when the computation result is revealed, an adversary can hardly infer the information of a single individual's data.

Definition 1. *(ϵ-**differential privacy**) **[13]** A randomized mechanism $\mathcal{A}$ is ϵ-differential private if for any two neighboring databases D and D' that only differ in only one entry and for any set $\mathcal{S}$ of outputs:*

$$Pr[\mathcal{A}(D) \in \mathcal{S}] \leq \exp(\epsilon) \cdot Pr[\mathcal{A}(D') \in \mathcal{S}],$$

where $Pr[\cdot]$ denotes the probability of the event.

Consider the extreme case of $\epsilon = 0$, i.e., $exp(\epsilon) = 1$, Definition 1 implies that $\Pr[\mathcal{A}(D) \in \mathcal{S}] = \Pr[\mathcal{A}(D') \in \mathcal{S}]$. This means that any two neighboring databases will have the same output distribution regardless of the single entry difference, which means perfect privacy protection. When the value ϵ becomes larger, the privacy protection becomes worse.

Here we regard ϵ as the privacy level for a given mechanism. As we can see from Definition 1, the parameter ϵ measures the worst (or the highest) privacy level among all possible neighboring databases. More specifically, data from different participants (i.e., different entries) may have different effects on the output, and this definition measures the most significant effect among all possible single entries. So the actual level of privacy protection of a particular participant under this mechanism can be lower (better) than ϵ. This means that a participant's actual privacy level involves uncertainty. This motivates us to use the PT to model how a participant subjectively respond when he concerns about the potential risk of his actual privacy level.

One main characteristic of PT, the S-shape valuation function [9] as often given in (1), explains how individuals subjectively evaluate the outcome (the actual privacy level):

$$v(\epsilon) = \begin{cases} (\epsilon_{ref} - \epsilon)^{\beta}, & \text{if } \epsilon \leq \epsilon_{ref}, \\ -\lambda(\epsilon - \epsilon_{ref})^{\beta}, & \text{if } \epsilon > \epsilon_{ref}, \end{cases} \tag{1}$$

where $0 < \beta < 1$, $\lambda > 1$, and ϵ_{ref} is the reference point. Fig. 2 shows the valuation function under different parameters. An individual evaluates gain and loss based on the reference point. When actual privacy level ϵ is lower than the reference point ϵ_{ref}, the individual would perceive it as a gain. Otherwise, he would perceive it as a loss. The parameter β describes the concavity of gain and the convexity of loss, capturing the risk aversion toward gain and the risk-seeking toward loss. The loss penalty parameter λ captures the loss aversion, which means that the impact of loss is larger than that of gain of the same absolute value.

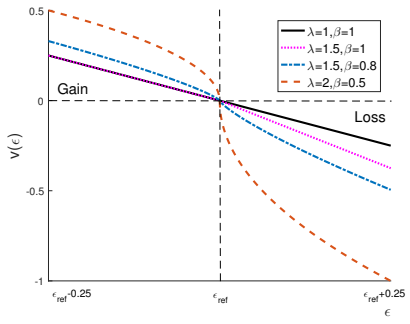

Fig. 2: (Inverse) S-shape asymmetrical valuation function.

In our setting, the privacy level (outcome) is continuous. We adopt the result from [14] to approximate infinite continuous outcomes with finite discrete outcomes. More specifically, we decompose the set of all possible continuous outcomes $[0, \epsilon]$ into m discrete outcomes $\frac{i\epsilon}{m}, i = 1, 2, ..., m$. Then a participant's prospect privacy level is the summation of weighted valuations of all discrete outcomes $p_i v(\frac{i\epsilon}{m}), i = 1, 2, ..., m$, where p_i is the weight (or probability) assigned to the corresponding outcome. For computational simplicity, we assume that the probability is evenly assigned to all the outcomes, then we obtain the prospect privacy level under an ϵ-differential private mechanism as follows:

$$\epsilon_p = \frac{1}{m} \sum_{i=1}^{m} v\left(\frac{i}{m}\epsilon\right). \tag{2}$$

When a participant's data is used in an ϵ-differential private mechanism, he will experience a privacy cost that is associated with the differential privacy level. Similar to [1], [2], [7], we model this privacy cost as a linear function of differential privacy level. Let $g(\epsilon_p)$ denote the linear function that maps the prospect differential privacy level to the privacy cost, i.e., $g(\epsilon_p) = c \cdot \epsilon_p$. The parameter c measures the privacy cost per privacy level. Since the data collector gathers the same type of data for a computation (i.e., income or moving rating), we assume that participants experience the same cost parameter c (similar as in [4], [5]) and PT parameters (ϵ_{ref}, λ, and β).

The data collector will perform computation on the collected data, and return the computation result to the participants as the reward. We assume that individuals have heterogeneous evaluations of the reward. Let W_i be individual's valuation of the reward, which is measured in the same unit as privacy cost. We use $f(W)$ to denote the distribution of W among individuals. The utility of a participating individual is the summation of reward valuation and privacy cost, i.e., $W_i + g(\epsilon_p)$.

For a non-participating individual, his actual privacy level is zero, i.e., perfect privacy protection. If his reference point ϵ_{ref} is positive, then not participating indicates a "gain" of privacy. If his reference point ϵ_{ref} is zero, he will not gain or lose any privacy. Hence his utility is $g(v(0)) = g(\epsilon_{ref}^{\beta})$.

Each individual i needs to solve the following optimization to decide his optimal action a_i,

$$\begin{aligned} \max_{a_i} \quad & U_i(a_i) = a_i(W_i + g(\epsilon_p)) + (1 - a_i)g(v(0)) \\ \text{s.t.} \quad & a_i \in \{0, 1\}. \end{aligned} \tag{3}$$

Action $a_i = 1$ means participation, and $a_i = 0$ otherwise. Similar to [1]–[3], for the ease of exposition, in this study we assume that participants will truthfully report their data due to the trusted data collector.

B. Data Collector's Utility Maximization Problem

In this subsection, we model how the data collector designs the differential privacy mechanism to maximize her utility.

The data collector benefits more if she manages to collect more data, as it enables a more convincing computation result [15]. Meanwhile, the data collector adopts a differential private mechanism that adds some random noise to the computation result, which leads to an accuracy penalty. So the data

collector's utility function depends on two factors: the amount of data collected and the accuracy penalty encountered.

We use $R(n)$ to denote the data collector's benefit of collecting data from n participants. We assume that $R(n)$ is non-negative, monotonic increasing, strictly concave, and upper bounded. As n grows large, the marginal benefit of collecting data from one more participant reduces, hence the concave shape of the function. Furthermore, the data amount is not the only factor that affects the computation result. Other factors such as methods of representation and optimization [15] all influence the computation. Hence function $R(n)$ would be bounded even when n goes to infinity. For the ease of exposition, we follow [16] and use the following benefit function with the parameters $[k, l]$ in our analysis,

$$R(n, [k, l]) = 1 - \frac{k}{1 + l \cdot n}, \text{where } k > 0 \text{ and } l > 0. \quad (4)$$

We use $L(\epsilon)$ to denote the accuracy penalty encountered by the data collector. We adopt a widely used representation [17]. Let $l(e)$ be the penalty if the noise is e. We assume that the penalty depends on the multitude of the noise and $l(e)$ is non-negative and non-decreasing in $|e|$. We consider one of the possible representations, $l(e) = e^2$ [17], which emphasizes the variance in the error. Let $f_\epsilon(e)$ be the probability density function of noise e under an ϵ-differential privacy mechanism. Then we get the expected accuracy penalty as follows:

$$L(\epsilon) = \int_{-\infty}^{+\infty} l(e) f_\epsilon(e) de = 2 \int_0^{+\infty} l(e) f_\epsilon(e) de. \quad (5)$$

Differential private mechanisms typically use Laplace distribution to generate the random noise [1], [2], [7], [13]. Let $Lap(b)$ with parameter b denote the Laplace distribution, which has the probability density function $f(x) = \frac{1}{2b} \exp(-\frac{|x|}{b})$ with zero mean and variance $2b^2$. By adding the noise to the computation result following the distribution $Lap(\frac{S(f)}{\epsilon})$, the data collector can achieve the ϵ-privacy protection. Here $S(f)$ is the sensitivity of a function f.

The sensitivity of a function measures the maximum variation that any single variable can cause to the computation result. For example, the mean function with the representation $f(X) = \frac{1}{n} \sum_i x_i$, where $X = [x_1, ..., x_n]$ is collected data, has the sensitivity of $S(f) = x_{\max}/n$. Here, $x_{\max} = \max_i |x_i|$.

For our problem, by substituting $Lap(\frac{S(f)}{\epsilon})$ for $f_\epsilon(e)$ in (5), we obtain $L(\epsilon) = 2\frac{S(f)^2}{\epsilon^2}$. The data collector needs to choose the privacy level ϵ to maximize her utility, i.e.,

$$\max_{\epsilon > 0} \; U_c(\epsilon) = R(n(\epsilon)) - L(\epsilon). \quad (6)$$

Here $n(\epsilon)$ is the number of participants under the ϵ-differential private mechanism, which will be derived based on the individuals' response to the mechanism (as in Section III.A).

C. Problem Formulation

Based on the data collector and an individual's problem derived in (6) and (3), respectively, we formulate the overall system as a two-stage game, which is illustrated in Fig. 1. In Stage I, the data collector launches a data collection with an ϵ-differential private mechanism. In Stage II, each individual decides whether to participate in the data collection. We use backward induction to solve this two-stage optimization problem. We will first derive the solution to individual's participation problem (3) in Stage II given an ϵ-differential private mechanism. Then we analyze the data collector's utility maximization problem (6) in Stage I given the solution to the individual's participation problem.

III. Solving the Two-stage Problem

A. Individual's Decision-Making

We first consider the case where individuals' reference point is zero, which means that individuals are intolerant and any privacy level induced in the data collection process will be considered as a loss. We will consider the non-zero reference point through simulations in Section IV.

In an individual's participation problem (3), the individual will decide to participate if and only if $U_i(1) \geq U_i(0)$, i.e., $W_i \geq g(v(0)) - g(\epsilon_p) = -g(\epsilon_p)$. This is because under the zero reference point case, utility of non-participation is zero, i.e., $g(v(0)) = g(\epsilon_{ref}^\beta) = 0$.

We consider a group of individuals with a population size N. Recall that the valuation of reward W_i among the individuals follows a distribution of $f(W)$. Then the number of participants n is given by:

$$n(\epsilon) = N \cdot Pr(W > -g(\epsilon_p)) = N \int_{-g(\epsilon_p)}^{\infty} f(W) dW. \quad (7)$$

The number of participants $n(\epsilon)$ is non-increasing with ϵ.

B. Data Collector's Optimal Differential Private Mechanism

Then we focus on the data collector's utility maximization problem in (6). We assume that she possesses adequate information of the target individuals, i.e., the privacy loss coefficient c [4], [5] and the distribution of valuation W and PT parameters [10], by abundant previous data-related investigations. Then she can decide the optimal ϵ to maximize her utility. As an example, we consider that the data ranging from zero to one representing individuals' ratings of a new movie or income levels. The computation performed by the collector is to calculate the mean of the collected data. The sensitivity of this computation is $S(f) = \frac{1}{n}$. Then the accuracy penalty is given as follows:

$$L(\epsilon) = 2\frac{S(f)^2}{\epsilon^2} = \frac{2}{n^2 \epsilon^2}. \quad (8)$$

For the sake of simplicity, we let the individual valuation of reward W follow a uniform distribution in $[m_W - \sigma_W, m_W + \sigma_W]$ in our theoretical analysis, i.e.,

$$f(W) = \begin{cases} \frac{1}{2\sigma_W}, & \text{if } W \in [m_W - \sigma_W, m_W + \sigma_W], \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

In the simulation study in Section IV, we will use a more general truncated normal distribution (which includes the uniform distribution as a special case).

We apply (7), (8), and (9) to the data collector's utility maximization problem (6), which leads to a one-variable optimization problem. We can apply many one-dimensional search methods [18] to numerically solve this problem, but it can be time-consuming. To obtain a computationally efficient solution, we would like to derive an analytical solution. The key challenge is that the first-order derivative of the objective function involves a six-order polynomial, and it is difficult to theoretically compute its root. Next, we describe how we approximate the first-order derivative to derive an (approximate) optimal solution of (6).

More specifically, when $-g(\epsilon_p) \in (m_W - \sigma_W, m_W + \sigma_W)$, we can compute the data collector's objective function in (6) together with its first-order derivative as follows:

$$U_c(\epsilon) = 1 - \frac{k}{1 + lN\frac{m_W+\sigma_W+g(\epsilon_p)}{2\sigma_W}} - \frac{2}{N^2\epsilon^2\left(\frac{m_W+\sigma_W+g(\epsilon_p)}{2\sigma_W}\right)^2}, \tag{10}$$

$$U_c'(\epsilon) = \frac{klN\frac{g_p'(\epsilon)}{2\sigma_W}}{\left[1 + lN\frac{m_W+\sigma_W+g(\epsilon_p)}{2\sigma_W}\right]^2} + \frac{4}{N^2}\frac{\frac{m_W+\sigma_W+g(\epsilon_p)+g'(\epsilon_p)\epsilon}{2\sigma_W}}{\left[\frac{m_W+\sigma_W+g(\epsilon_p)}{2\sigma_W}\right]^3\epsilon^3}. \tag{11}$$

One of the key challenges of computing the root of $U_c'(\epsilon) = 0$ is due to the denominator term $\left(1 + lN\frac{m_W+\sigma_W+g(\epsilon_p)}{2\sigma_W}\right)^2$, which leads to a six-order polynomial at the numerator after combining the terms. However, we observe that $lN\frac{m_W+\sigma_W+g(\epsilon_p)}{2\sigma_W} \gg 1$ when the number of individuals N is large, which motivates us to approximate $1 + lN\frac{m_W+\sigma_W+g(\epsilon_p)}{2\sigma_W}$ with $lN\frac{m_W+\sigma_W+g(\epsilon_p)}{2\sigma_W}$. We further set $\beta = 1$ to avoid a fractional order so that $g(\epsilon_p) = -M\epsilon$, where $M = \frac{c\lambda(m+1)}{2m}$ based on (2) and $\epsilon \in \left(\frac{m_W-\sigma_W}{M}, \frac{m_W+\sigma_W}{M}\right)$. Hence we can obtain an approximated (denoted by the superscript a) version of (11):

$$U_c'^a(\epsilon) = \frac{4}{\left(\frac{m_W+\sigma_W-M\epsilon}{2\sigma_W}\right)^3\epsilon^3N^2}f_1(\epsilon). \tag{12}$$

Here

$$f_1(\epsilon) = \left(\frac{m_W+\sigma_W}{2\sigma_W} - \frac{M}{2\sigma_W}\epsilon\right)\left(1 - C\frac{M}{2\sigma_W}\epsilon^3\right) - \frac{M}{2\sigma_W}\epsilon, \tag{13}$$

where $C = \frac{kN}{4l}$. As $\epsilon < \frac{m_W+\sigma_W}{M}$, so $\frac{m_W+\sigma_W-M\epsilon}{2\sigma_W} > 0$ in (12). This means that computing the root of $U_c'^a(\epsilon) = 0$ is equivalent to computing the root of $f_1(\epsilon) = 0$. The equation $f_1(\epsilon) = 0$ has two real roots and two imaginary roots in the whole feasible set. Notice that function $f_1(\epsilon)$ is continuous, and $f_1(0)\cdot f_1(\frac{m_W+\sigma_W}{M}) < 0$ and $f_1(\frac{m_W+\sigma_W}{M})\cdot f_1(+\infty) < 0$. This implies that the equation $f_1(\epsilon) = 0$ has at least one root in $(0, \frac{m_W+\sigma_W}{M})$ and at least one root in $(\frac{m_W+\sigma_W}{M}, +\infty)$.

Combining the above discussions, we know that there is only one root in $(0, \frac{m_W+\sigma_W}{M})$, and we denote this root as ϵ_1. More specifically, in $(0, \epsilon_1)$ we have $f_1(\epsilon) > 0$ and $U_c'^a(\epsilon) > 0$, and in $(\epsilon_1, \frac{m_W+\sigma_W}{M})$ we have $f_1(\epsilon) < 0$ and $U_c'^a(\epsilon) < 0$. Then the objective function $U_c(\epsilon)$ of (10) achieves its maximum value (approximately) at $\epsilon = \epsilon_1$.

Furthermore, we need to verify whether ϵ_1 is within the interval $(\frac{m_W-\sigma_W}{M}, \frac{m_W+\sigma_W}{M})$. This requires us to consider two cases, as summarized in Theorem 1. The proof of Theorem 1 is provided in our online appendix [19].

Theorem 1. *When approximating the first order derivative (11) with (12), we obtain the unique optimal solution ϵ^* of problem (6) as follows:*

$$\epsilon^* = \begin{cases} \frac{m_W-\sigma_W}{M}, & \text{if } \frac{m_W-\sigma_W}{2\sigma_W}\left[C\left(\frac{m_W-\sigma_W}{M}\right)^2 + 1\right] \geq 1, \\ \epsilon_1, & \text{otherwise,} \end{cases} \tag{14}$$

where $C = \frac{kN}{4l}$ and $M = \frac{c\lambda(m+1)}{2m}$.

From Theorem 1, we can see that the data collector has a unique (approximately) optimal strategy of differential private mechanism to maximize the utility. We also see from (14) that the following condition plays an important role,

$$\frac{m_W-\sigma_W}{2\sigma_W}\left[C\left(\frac{m_W-\sigma_W}{M}\right)^2 + 1\right] \geq 1. \tag{15}$$

Intuitively, condition (15) holds when the minimum valuation of reward $m_W - \sigma_W$ is high, or the number of individuals N is large, or the individual' privacy cost parameter c is small or the loss penalty parameter λ is small.

When condition (15) holds, the absolute value of prospect privacy cost $-g(\epsilon_p)$ equals to the minimum valuation of reward $m_W - \sigma_W$, i.e., $-g(\epsilon_p) = M\epsilon^* = m_W - \sigma_W$.

C. Comparison with EUT

In this section, we consider the traditional EUT case [4]. EUT is a special case of PT where $\lambda = 1$ and $\beta = 1$ in (2). The general PT case considered corresponds to the case where at least one of these two parameters is not equal to one.

Corollary 1. *The data collector's optimal ϵ^* under general PT is lower than that under EUT.*

We provide the proof of Corollary 1 in the online appendix [19]. We conclude that compared with traditional EUT modeling, the data collector should adopt a stricter privacy-preserving mechanism when considering the individuals' loss aversion and risk aversion predicted by the PT.

IV. Simulation results and insights

In this section, we evaluate the accuracy of the approximation and the impact of reference point.

A. The Accuracy of Approximation in (12)

We compare the optimal ϵ^* with and without approximation. The result without approximation is obtained by the exhaustive search. We set $m_W = \sigma_W = 0.5$, i.e., the minimum valuation of reward is zero, so the approximated solution is ϵ_1 in (14). We set $c = 1$ and $\lambda = 2.25$. We choose $k = 0.989$ and $l = 0.0039$ for the data amount benefit in (4). We change the number of individuals N and compare the optimal ϵ^* with approximation under $\beta = 1$ and that without the approximation under different values of β.

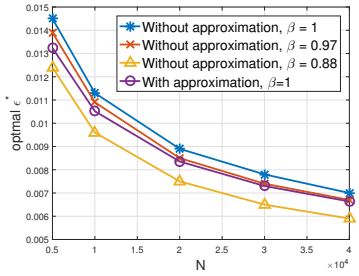

Fig. 3: Comparison between optimal ϵ^* with and without approximation vs. N.

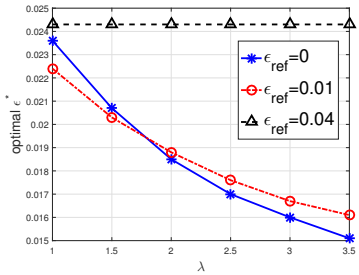

Fig. 4: Optimal ϵ^* vs. λ under different reference point ϵ_{ref}.

From Fig. 3 we see that the optimal ϵ^* decreases in N, as a larger number of individual population can potentially provide a higher accuracy, which dramatically reduces the accuracy penalty. So a stricter privacy-preserving mechanism can attract more individuals to improve accuracy. Furthermore, we find that for the case of $\beta = 1$, the inaccuracy due to the approximation is relatively small, and it decreases as N increases. As N changes from 5000 to 40000, the relative difference between ϵ^* with approximation and that without approximation changes from 8.7% to 5.2%. This is because the condition $lN\frac{m_W+\sigma_W+g_p(\epsilon)}{2\sigma_W} \gg 1$ becomes more accurate. However, for the case where $\beta \neq 1$, the approximation error can be larger. When $\beta = 0.88$ and $N = 5000$, the approximation error is 6.8%, which is still quite reasonable.

B. The Impact of Reference Point

We numerically study the impact of a non-zero reference point on the data collector's optimal strategy. In this case, a non-participating individual will enjoy a gain from better privacy protection than his reference point, i.e., $g(v(0)) = c\epsilon_{ref}^{\beta} > 0$ in (3). We set $c = 1$, $\beta = 0.88$, and $N = 5000$, and adopt the truncated normal distribution in $[0, 2]$ with mean one and variance one for the valuation of reward.

Fig. 4 shows how optimal ϵ^* without approximation changes with reference points under a fixed λ. A higher ϵ_{ref} means that individuals are more tolerant about their privacy loss. We first focus on the case where λ is large, i.e., $\lambda \geq 2$. Under a large λ, the optimal ϵ^* increases with the reference point ϵ_{ref}. In this case, individuals are sensitive to loss due to the large λ (as is shown in (2)). So when ϵ_{ref} increases, the absolute value of privacy cost from participation significantly decreases. Then the data collector can adopt a less strict privacy-preserving mechanism (which corresponds to a larger ϵ^*). Next we consider the case where λ is small, i.e., $\lambda \leq 1.5$. When ϵ_{ref} increases from zero, at first (e.g., $\epsilon_{ref} = 0.01$) the privacy protection gain from non-participation is more significant. The data collector needs to enforce a stricter privacy protection, so the optimal ϵ^* is lower than that under $\epsilon_{ref} = 0$. However, when ϵ_{ref} is relatively large (i.e., $\epsilon_{ref} = 0.04$), individuals perceive less privacy loss, hence a less strict mechanism is still acceptable.

V. Conclusion

In this paper, we utilize prospect theory to analyze privacy-preserving data collection problem with the privacy protection uncertainty. We conclude that a data collector should adopt a stricter privacy-preserving mechanism owing to individuals' loss aversion, compared with the modeling under the expected utility theory. For the future work, we will consider the possibility of untruthful reporting of data after participation.

References

[1] A. Ghosh and A. Roth, "Selling privacy at auction," *Games and Economic Behavior*, vol. 91, pp. 334–346, 2015.

[2] L. K. Fleischer and Y.-H. Lyu, "Approximately optimal auctions for selling privacy when costs are correlated with data," in *Proceedings of ACM Conference on Electronic Commerce*, 2012.

[3] A. Roth and G. Schoenebeck, "Conducting truthful surveys, cheaply," in *Proceedings of ACM Conference on Electronic Commerce*, 2012.

[4] W. Wang, L. Ying, and J. Zhang, "The value of privacy: Strategic data subjects, incentive mechanisms and fundamental limits," *arXiv preprint arXiv:1603.06870*, 2016.

[5] ——, "A game-theoretic approach to quality control for collecting privacy-preserving data," in *Annual Allerton Conference on Communication, Control, and Computing*, Sept 2015.

[6] A. Acquisti and J. Grossklags, "What can behavioral economics teach us about privacy," *Digital Privacy: Theory, Technologies and Practices*, vol. 18, pp. 363–377, 2007.

[7] A. Ghosh and K. Ligett, "Privacy and coordination: Computing on databases with endogenous participation," in *Proceedings of ACM Conference on Electronic Commerce*, 2013.

[8] D. Kahneman and A. Tversky, "Prospect theory: An analysis of decision under risk," *Econometrica*, vol. 47, no. 2, pp. 263–291, 1979.

[9] A. Tversky and D. Kahneman, "Advances in prospect theory: Cumulative representation of uncertainty," *Journal of Risk and Uncertainty*, vol. 5, no. 4, pp. 297–323, 1992.

[10] Y. Yang, L. T. Park, and N. B. Mandayam, "Prospect pricing in cognitive radio networks," *IEEE Transactions on Cognitive Communications and Networking*, vol. 1, no. 1, pp. 56–70, 2015.

[11] J. Yu, M. H. Cheung, and J. Huang, "Spectrum investment under uncertainty: A behavioral economics perspective," *IEEE Journal on Selected Area in Communication*, vol. 34, no. 10, pp. 2667–2677, 2016.

[12] J. Yu, M. H. Cheung, J. Huang, and H. V. Poor, "Mobile data trading: Behavioral economics analysis and algorithm design," *IEEE Journal on Selected Areas in Communications*, vol. 35, no. 4, pp. 994–1005, 2017.

[13] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography Conference*, 2006.

[14] M. O. Rieger and M. Wang, "Prospect theory for continuous distributions," *Journal of Risk and Uncertainty*, vol. 36, no. 1, pp. 83–102, 2008.

[15] P. Domingos, "A few useful things to know about machine learning," *Communications of the ACM*, vol. 55, no. 10, pp. 78–87, 2012.

[16] D. Niyato, M. A. Alsheikh, P. Wang, D. I. Kim, and Z. Han, "Market model and optimal pricing scheme of big data and internet of things (iot)," in *IEEE ICC*, May 2016.

[17] A. Ghosh, T. Roughgarden, and M. Sundararajan, "Universally utility-maximizing privacy mechanisms," *SIAM Journal on Computing*, vol. 41, no. 6, pp. 1673–1693, 2012.

[18] K. E. Atkinson, *An introduction to numerical analysis*. John Wiley & Sons, 2008.

[19] G. Liao, X. Chen, and J. Huang, *online appendix, available at http://jianwei.ie.cuhk.edu.hk/publication/AppendixGC17Privacy.pdf*.