

Approximately Optimal Auctions for Selling Privacy when Costs are Correlated with Data

LISA K. FLEISCHER and YU-HAN LYU, Department of Computer Science, Dartmouth College

We consider a scenario in which a database stores sensitive data of users and an analyst wants to estimate statistics of the data. The users may suffer a cost when their data are used in which case they should be compensated. The analyst wishes to get an accurate estimate, while the users want to maximize their utility. We want to design a mechanism that can estimate statistics accurately without compromising users' privacy.

Since users' costs and sensitive data may be correlated, it is important to protect the privacy of both data *and* cost. We model this correlation by assuming that a user's unknown sensitive data determines a distribution from a set of publicly known distributions and a user's cost is drawn from that distribution. We propose a stronger model of privacy preserving mechanism where users are compensated whenever they reveal information about their data to the mechanism. In this model, we design a Bayesian incentive compatible and privacy preserving mechanism that guarantees accuracy and protects the privacy of both cost and data.

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Non-numerical Algorithms and Problems; J.4 [Social and Behavioral Sciences]: Economics

General Terms: Algorithms, Economics

Additional Key Words and Phrases: Mechanism Design, Differential Privacy

1. INTRODUCTION

Using the Internet, it is fairly easy to collect sensitive personal data. Online service providers implicitly compensate users who provide their personal data, by offering improved services based on their data. However, this implicit exchange may not be fair to the individual, since different people may have different costs — a loss in expected utility over future events — for use of their data. Moreover, companies rarely give well-defined guarantees concerning data privacy and compensation. When the compensation is less than the individual's perceived cost, the individual may choose not to participate. Here, we explore mechanisms to fairly compensate individuals for use of their personal data.

In order to motivate users to participate in a mechanism, the payment to a user should be at least the cost to the user. Thus, the mechanism should learn information about users' costs. Ghosh and Roth [Ghosh and Roth 2011] initiate a study of this problem. Their mechanism asks users to report their costs for the use of their data to estimate statistics, and then selects some of the users (based on their stated costs) to determine the statistics, and pays these users accordingly. This mechanism is problematic when costs and personal data are correlated, since users may be reluctant to

E-mail: {lkf,yuhanlyu}@cs.dartmouth.edu. Partially supported by NSF grants CCF-0728869 and CCF-1016778.

Author's address: L. Fleischer and Y. Lyu, Computer Science Department, Dartmouth College.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

EC'12, June 4–8, 2012, Valencia, Spain.

Copyright 2012 ACM 978-1-4503-1415-2/12/06...\$10.00.

reveal their costs if they are not guaranteed adequate compensation up front. For example, suppose that a database indicates whether a 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 the price of an undamaged car. Thus, his cost for revealing data is higher than the owner of an undamaged car. Revealing information about the costs may also reveal information about whether the car is damaged. Thus, it is important to also guarantee privacy of individual payments.

We study this problem where costs are correlated with data. We model this correlation by assuming that a user's unknown data determines a distribution from a set of accurate and publicly known distributions and the user's cost is drawn from that distribution. We propose a model of a privacy preserving mechanism where users are compensated whenever they reveal any information about their data to the mechanism, whether directly, or indirectly by revealing their costs. In this model, we design a Bayesian incentive compatible and individually rational mechanism, which produces accurate statistics and protects the privacy of data and costs.

Problem Setting. There are n users, which we call players, denoted by $[n]$. Each player has sensitive data $D_i \in [h]$, stored in a database $D \in [h]^n$. Initially D_i is the private information of player i . However, since D_i is also in the database, its value may be verified with player i 's permission. In addition, player i has a value for his loss of privacy of his data. This value v_i is private to player i , but it is correlated with D_i . This correlation is modeled as follows: If $D_i = t \in [h]$ then $v_i \sim F_t$, where F_t is a distribution of privacy costs for players of type t that is known to all players and the mechanism. F_t correctly represents the distributions of costs of type t players.

A query is a function $Q : [h]^n \rightarrow \mathbb{R}$, mapping a database to a response. An example of a query is "what is the number of people i in the database D with $D_i = j$?" A data analyst wants $Q(D)$. Since the data are sensitive, the data analyst accesses the database through a privacy preserving algorithm A . Therefore, the data analyst does not receive $Q(D)$ but receives an estimate $A(D)$. To ensure the estimate is accurate, the error $|Q(D) - A(D)|$ should be small with high probability.

Differential privacy, introduced in [Dwork et al. 2006], is an accepted way to measure privacy and privacy preserving algorithms. Two databases D and D' are *adjacent* if they differ in only one entry. An algorithm A satisfies ϵ -*differential privacy*, where $\epsilon > 0$, if for any pair of adjacent database D and D' and any set $I \subseteq \mathbb{R}$, $\Pr[A(D) \in I] \leq e^\epsilon \Pr[A(D') \in I]$. When $\epsilon = 0$, it implies that the algorithm does not depend on D . If the error $|Q(D) - A(D)|$ is small with high probability, then the algorithm should have large ϵ . Thus, privacy guarantees come at the expense of the accuracy.

Although an ϵ -differentially private algorithm can protect sensitive data, if a player allows his data to be used, he may incur a cost. We model this cost as linear in the privacy loss ϵ and his expected cost v_i .¹ Thus, for player i to agree to the use of his data, his expected payment should be at least ϵv_i .

A mechanism specifies a set of actions that players can take. The players take actions based on their data and private costs. Thus, the input of the mechanism is a database and a vector of actions. The outputs are an estimate \hat{s} and a payment vector $p = (p_1, \dots, p_n)$. Since player i has a linear cost ϵv_i , the utility of player i is $p_i - \epsilon v_i$ if

¹We can view this cost as due to the change in his utility from future events that depend on the answer he gives to the analyst. This cost is approximately linear in ϵ and his expected utility, denoted by v_i . Let $g(A(D))$ be the distribution of future events that depends on $A(D)$. Let w_i be the player i 's utility for future events. Since A is ϵ -differentially private, $g \circ A$ is also ϵ -differentially private. Thus, for random variables $y \sim g(A(D))$ and $y' \sim g(A(D'))$ and event b , $\Pr[y = b] \leq e^\epsilon \Pr[y' = b]$. Therefore, we have $E_{y \sim g(A(D))}[w_i(y)] - E_{y \sim g(A(D'))}[w_i(y)]$ is approximately $\epsilon E_{y \sim g(A(D'))}[w_i(y)]$ or $-\epsilon E_{y \sim g(A(D'))}[w_i(y)]$, when ϵ is small.

D_i is used in the mechanism, otherwise the utility is p_i . We assume that all players are rational and want to maximize their utilities. A mechanism is a *direct* mechanism if the action set equals the set of all real numbers. That is, a direct mechanism asks players to report their costs. A direct mechanism is truthful if every player reports his true cost in order to maximize his utility. Truth telling is a concept defined for direct mechanisms. In this paper, we propose an indirect mechanism. Thus, we want to extend the notion of truthfulness to indirect mechanisms. In our mechanism, there is a straightforward mapping, described in Section 3, from player's type set to player's action set. We say that a player *decides truthfully* if he picks the strategy corresponding to his type under this mapping.

In our paper, we will assume that the query/goal of the analyst is to estimate $n_j = |\{i : D_i = j\}|$. Without loss of generality, we assume throughout the paper that the data analyst wants to estimate n_1 . We seek to design a mechanism with the following properties.

- (1) **Accuracy:** A mechanism M is k -accurate, if for any database D , $\Pr[|\hat{s} - n_1| \geq k] \leq \frac{1}{3}$, when every player decides truthfully. Note that the accuracy guarantee is independent of the size of the database — the number k is fixed no matter how large the database is, or the sampled set is.
- (2) **Differential Privacy:** The estimate and payments satisfy ϵ -differential privacy.
- (3) **Truthfulness:** A mechanism is *dominant strategy truthful* if, for every player, deciding truthfully maximizes his utility. A mechanism is *Bayesian incentive compatible* (BIC) if, for every player, assuming that other players' costs are drawn from F according to their data and decide truthfully, deciding truthfully maximizes his utility.
- (4) **Individual Rationality:** If a player's utility is non-negative, then he should be willing to participate. A mechanism is *ex-post individually rational* (EPIR) if the utility is non-negative for every player when he decides truthfully. A mechanism is *ex-interim individually rational* (EIIR) if the expected utility is non-negative for every player when he decides truthfully, where the randomness comes from the mechanism and the costs of other players.
- (5) **Payment Minimization:** The summation of payments should be as little as possible.

To get permission to use a player's data, the mechanism must compensate the player by at least his perceived loss of privacy. But since costs are correlated with data, players may be reluctant to reveal their true costs, unless they will be compensated for this. To avoid this seeming chicken-and-egg problem, the mechanism designer cannot resort to the revelation principle, which states that any mechanism can be realized as a direct and truthful mechanism. In fact, [Ghosh and Roth 2011] prove that if costs and data can be arbitrarily correlated and player's cost of privacy can be unbounded, then for any $k < n/2$, no k -accurate, direct, dominant strategy truthful, EPIR, privacy preserving mechanism exists. On the other hand, we give a mechanism that provides k -accuracy for any input value k when costs are correlated with data, and there is no bound on players' cost of privacy. We get around the lower bound of [Ghosh and Roth 2011] by using an *indirect* mechanism, and modeling the correlation of values and data via publically known (and allowably unbounded) distributions.

Privacy Issues when Costs are Correlated with Data. The objective of a privacy preserving mechanism is that the increase in knowledge about a player's data due to output of the mechanism is small. Previous work on privacy in statistical databases assumes that the mechanism is associated with the database, such that the mechanism can access the whole database without compromising a player's privacy. How-

ever, if the mechanism is separated from the database, then a player might not trust the mechanism and might not want to reveal private information to the mechanism.

In our problem, in order to estimate n_1 , the mechanism should learn information about players' data. Suppose that the mechanism has a prior belief G about the data in D . That is, the mechanism believes that the probability of $D_i = j$ is $\Pr_G[D_i = j]$ according to the prior belief. The mechanism *learns* about D_i if the mechanism believes that $\Pr[D_i = j] \neq \Pr_G[D_i = j]$ after running the mechanism, for some j . There are two possible ways to learn about players' data. The first way is to read D_i explicitly. The second way is to read players' actions and deduce something about their D_i . For example, if the mechanism is direct and truthful, then the players report v_i truthfully. Suppose that the prior belief is that every player's data are drawn from a uniform distribution. That is, $\Pr_G[D_i = j]$ is the same for all i and j . If $F_j(v_i) < F_{j'}(v_i)$ for some j and j' , and player i truthfully reports v_i , then the mechanism's posterior belief is that $\Pr[D_i = j] < \Pr[D_i = j']$, which is different from the prior belief. Learning anything about a player's data may compromise a player's privacy and should be compensated. Thus, there are two kinds of cost to a player that should be compensated, one is for using the player's data and one is for learning about the player's data.

For the latter cost, we propose the concept of perfect data privacy, which is inspired by the concept of perfect objective privacy introduced in [Feigenbaum et al. 2010]. A mechanism satisfies *perfect data privacy* if whenever the mechanism's posterior belief about a player's data differs from its prior belief, the mechanism pays the player. Under perfect data privacy, mechanisms can learn about a player's cost, as long as that knowledge does not reveal anything about his data. However, for a perfectly data private mechanism, if the mechanism learns about a player's data, then the mechanism always compensates the player, even when the mechanism does not use the player's data to compute the estimate.

Our Main Contribution. We give a mechanism that is BIC, EIIR, $O(\epsilon^{-1})$ -accurate, perfectly data private, and ϵ -differentially private. To achieve our privacy guarantees, we propose a posted-price-like mechanism, described in Section 3. Given the set of types of players and the distributions of costs, the mechanism writes a contract that offers a different expected payment for each type. Each player is offered this contract. If a player accepts the contract, then his payment is determined by his verifiable type and the payment for his type in the contract. The player's action is either to accept the contract or reject the contract. A player *decides truthfully* if a player with type j accepts the contract when $\epsilon v_i \leq r_j$, where r_j is the payment for type j in the contract. We prove that this posted-price-like mechanism is BIC, EIIR, $O(\epsilon^{-1})$ -accurate, perfectly data private, and ϵ -differentially private.

We seek a mechanism with a small payment. In Section 4, we define a benchmark for the expected payment of a mechanism and compare the expected payment of our mechanism to this benchmark in two different settings. When costs are non-negative, we show that our mechanism is close to the benchmark.

We also prove a lower bound on the accuracy that a direct and data private mechanism can achieve in Section 2.

1.1. Related Work

Selling Privacy. Our paper is closely related to the privacy preserving mechanisms studied in [Ghosh and Roth 2011]. In [Ghosh and Roth 2011], they extend the definition of ϵ -differentially private algorithms to ϵ -differentially private mechanisms. Under their definition of an ϵ -differentially private mechanism, the randomness only comes from the mechanism. In our model, since we want to protect the privacy of the costs, which are drawn from distributions, our definition of an ϵ -differentially private mech-

anism relies both on the distributions of the costs and the randomness of the mechanism.

Differential Privacy. A comprehensive survey of differential privacy appears in [Dwork 2008]. Most of the previous results are based on random perturbations of the output, and assume that the mechanism has the ability to access the whole database. If the mechanism cannot access the whole database, Chaudhuri et al. [Chaudhuri and Mishra 2006] and Klonowski et al. [Klonowski et al. 2010] show that random sampling is enough to ensure differential privacy with high probability. That is, it is not necessary to add more noise to the output.

Differential Privacy and Mechanism Design. McSherry et al. [McSherry and Talwar 2007] use a privacy preserving algorithm as a tool to design an approximately dominant strategy truthful mechanism. Instead, we focus on treating sensitive data as a commodity that can be sold.

Privacy Concerns in Mechanisms. Traditional mechanism design theory focuses on drawing private information from players in order to compute a result. However, if players have privacy concerns, they may not want to reveal their information. Feigenbaum et al. [Feigenbaum et al. 2010] study how to quantify the information leakage to the mechanism based on communication complexity.

Xiao [Xiao 2011] quantifies the information leakage in a mechanism based on information theory. In his model, the outcome of a privacy preserving mechanism not only motivates the players to participate but also protects the private information of players. In independent work, Nissam et al. [Nissim et al. 2011] and Chen et al. [Chen et al. 2011] consider privacy issues in mechanism design in the context of elections and discrete facility location.

Posted-Price Mechanisms. In a posted-price mechanism, player i is offered a price r_i . If player i accepts that price, then i pays r_i to get the allocation. Goldberg et al. [Goldberg and Hartline 2005] show that the posted-price mechanism is collusion resistant. Moreover, the players do not need to know or report their private values precisely. They only decide to accept or reject the price. Chawla et al. [Chawla et al. 2010] point out that this could be useful in reducing the private information revealed to the mechanism.

Revenue Maximization in Bayesian Mechanism Design. In a classic paper, Myerson [Myerson 1981] characterizes the optimal BIC selling mechanism to maximization the expected revenue. In procurement mechanisms, each player is a supplier and each player's production cost is private information. The auctioneer is the buyer and wants to minimize the expected payment. In the computer science literature, an early paper in this area characterizes the minimum-cost dominant strategy truthful auction to buy an s - t path in a graph [Elkind et al. 2004]. Since then, there has been considerable interest in both *frugal* mechanism design (buying a feasible set at low cost), and budget-constrained mechanism design (buying as good a set as possible subject to a budget). Our work can be seen as a generalization of these questions to the setting of bidders who are reluctant to reveal their costs, and the feasibility of a set depends on the private costs (via the correlation with data).

2. MODEL AND LOWER BOUND

2.1. Model

There is a database $D \in [h]^n$ and n players, where each player has data D_i . Player i with $D_i = j$ has a private cost v_i drawn from a distribution with cumulative distribution function F_j . Note that this definition is different from the traditional definition

of a Bayesian setting. In the traditional definition, the distribution of v_i is known to every player and the mechanism. In our definition, the mechanism and players know that each player's v_i is drawn from one of a set of distributions, but the particular distribution depends on the individual player's data, which is unknown to everyone but that player.

The goal of our mechanism is to estimate n_1 based on D and determine the payment p_i for every player i . A mechanism first specifies the set of possible actions Y that players can take. Then, based on players' actions and the database, the mechanism determines the estimate and payment. Formally, a mechanism is a function $M : Y^n \times [h]^n \rightarrow \mathbb{R} \times \mathbb{R}^n$. The mechanism has an a priori belief G about the data in D . That is, the mechanism believes that the probability of $D_i = j$ is $\Pr_G[D_i = j]$. Recall that the mechanism learns about D_i if, after running the mechanism, the mechanism believes that $\Pr[D_i = j] \neq \Pr_G[D_i = j]$ for some j . We use a vector $x \in \{0, 1\}^n$ to indicate whether the mechanism learns something about each player's data. If the mechanism learns about D_i , then $x_i = 1$. A mechanism is *perfectly data private* if, when $x_i = 1$, player i 's expected payment from the mechanism is at least ϵv_i . We focus on randomized mechanisms in this paper, that is, x_i and payment p_i are random variables.

Next, we define the utility for a player. If $x_i = 1$, there is a cost ϵv_i to player i , since something about D_i is learned. For $y \in Y^n$ representing all players' actions, the *utility* for player i is $u_i(y, v_i) = p_i - \epsilon x_i v_i$, where $(\hat{s}, p) = M(y, D)$. In this paper, we assume that players are rational, so players want to maximize their expected utilities. The *strategy* of player i is a function $q_i : \mathbb{R} \times [h] \rightarrow Y$ mapping from v_i and D_i to an action. Since players want to maximize their expected utilities, they will take the action that is not worse than any other action.

Finally, we introduce the solution concept. A profile of strategies q_1, \dots, q_n is a *Bayesian-Nash equilibrium* if for all i , v_i , and $y'_i \in Y$, $E[u_i(q(v_i, v_{-i}, D), v_i)] \geq E[u_i(y'_i, q_{-i}(v_{-i}, D_{-i}), v_i)]$, where the randomness is from the mechanism and the randomness of v_{-i} . A direct mechanism is *Bayesian incentive compatible* (BIC) if $q_i(v_i, D_i) = v_i$ is a Bayesian-Nash equilibrium for every player i .

2.2. Lower Bound

In order to ensure that players have incentive to participate the mechanism, we wish that the mechanism is individually rational. However, we can show that for any direct, BIC, and EIIR mechanism, there is a lower bound of accuracy. Since the condition of EIIR is weaker than EPIR, the lower bound for EIIR also implies a lower bound for EPIR mechanisms.

LEMMA 2.1. *If the functions F_i are arbitrary functions with unbounded range, then for any $k < n/2$, no k -accurate, direct, BIC, EIIR, and perfectly data private mechanism exists.*

PROOF. Suppose that M is a BIC, EIIR, perfectly data private, and k -accurate mechanism. First, we show that M must access at least one player's cost or data. Assume that M does not access any cost or data. Thus, M randomly output an estimate \hat{s} , which is independent of costs and data. For a database D^1 with all entries equal to one, since M is k -accurate, $\Pr[\hat{s} \in [n, n - k]] \geq \frac{2}{3}$. Similarly, if a database D^0 has no entries equal to one, then $\Pr[\hat{s} \in [0, k]] \geq \frac{2}{3}$. Because $k < n/2$, $[n, n - k]$ and $[0, k]$ do not overlap. But the summation of these two probabilities is greater than one, which is impossible. Hence, M must access at least one player's cost or data.

Suppose that $D_i \in \{1, 2\}$ and $F_1(v) \neq F_2(v)$ for all v . For any \hat{v} , if M access $v_i = \hat{v}$, then the mechanism must pay player i , since $F_1(\hat{v}) = \Pr[v_i = \hat{v} | D_i = j] \neq \Pr[v_i = \hat{v} | D_i = j'] = F_2(\hat{v})$ and M is perfectly data private. Let x_i be the indicator random

variable representing whether player i 's cost is accessed. Let p_j be the random variable representing player i 's payment. Since M is BIC, we suppose that players other than i report truthfully. Since the mechanism decides to access v_i based on v_{-i} , $\Pr[x_i = 1]$ is independent of v_i . Because M must access at least one player's cost, we can find a player i , such that $\Pr[x_i = 1] > 0$. For a fixed v_i , the expected utility of i is $E[p_i] - \epsilon v_i E[x_i]$. Since the range of F is unbounded, we can find another $v'_i > \frac{E[p_i]}{\epsilon E[x_i]}$. Since M is EIIR, we have $E[p'_i] \geq \epsilon v'_i E[x_i]$. Thus, for player i with cost v_i , if i overbids v'_i , the utility is $E[p'_i] - \epsilon v_i E[x_i] \geq \epsilon v'_i E[x_i] - \epsilon v_i E[x_i] > E[p_i] - \epsilon v_i E[x_i]$. Thus, player i can increase expected utility by overbidding. Hence, M is not BIC. \square

Our mechanism, which is explained in the next section, is an indirect mechanism since it does not ask for players' costs. The revelation principle, which states that if there exists an indirect mechanism implementing a function in Bayesian-Nash equilibrium, then there also exists a direct BIC mechanism implementing the same function, is irrelevant under the desire for perfect data privacy. It is easy to construct a direct mechanism from our indirect mechanism. However, this direct mechanism accesses all players' data without compensating all players. Thus, this direct mechanism is not perfectly data private.

ϵ -Differential Privacy. The traditional definition of ϵ -differential privacy compares the outcomes of the algorithm applied to adjacent databases. However, with a mechanism that offers payments, the mechanism may use both the database and the replies to the mechanism to compute an estimate and payments. Since replies depend on the individuals' costs, we compare the outcomes of the mechanism applied to two cost-data pairs (v, D) and (v', D') . A cost vector $v = (v_1, \dots, v_n)$ is *drawn according to a database* D , if v_i is drawn from F_j , where $D_i = j$. Two cost-data pairs (v, D) and (v', D') are *adjacent*, if D and D' differ only in the i -th entry and v and v' are independently drawn according to database D and D' . A BIC mechanism is ϵ -differentially private if, for any pair of adjacent cost-data pairs, the estimate and payments satisfy ϵ -differential privacy.

Bayesian Assumptions. Our definition of ϵ -differential privacy is based on the common belief F . That is, the player decides his strategy assuming that other players' costs are drawn from F and all players believe this assumption. If a player allows his data to be used, then he may incur a expected cost ϵv_i . The expected cost to the player depends on ϵ and thus also depends on the common belief F . Having a common belief is a traditional assumption in the Bayesian setting. Moreover, most BIC mechanisms become meaningless when the common belief is not true. Thus, we assume that the common belief F is correct.

3. MECHANISM

In this section, we give a perfectly data private, BIC, EIIR, ϵ -differentially private, and $O(\epsilon^{-1})$ -accurate mechanism. Every player i has data $D_i \in [h]$. To start, we assume that F_j is continuous for $j \in \{1, 2\}$.

The mechanism designs and offers contracts to players. The contract guarantees an expected payment to each player who accepts the contract. The players decide to accept or reject the contract. Thus, the possible actions for players are "accept" or "reject". The mechanism uses the data of players who accept the contract to estimate n_1 . The estimate is *unbiased* if the expected value of the estimate is n_1 . To obtain an unbiased estimate, the set of players who accept the contract should be unbiased, that is, the probability of a player accepting the contract should be equal for all players. Moreover, since the mechanism pays players, the costs of players in the accepting set should be bounded.

The mechanism first finds α_j for $j \in [h]$, such that $F_j(\alpha_j) = c$, where c will be determined later. Then, each player i is given a contract : “If $D_i = j$, your expected payment will be $\epsilon\alpha_j$.” A player i with $D_i = j$ decides truthfully if, when $v_i \leq \alpha_j$, player i accepts the contract and rejects otherwise. Let W be the set of players who accept the contract. If all players decide truthfully, the cost to each player in W is bounded by $\max_j \alpha_j$. Since for player i with $D_i = j$, $\Pr[v_i \leq \alpha_j] = c$, every player accepts the contract with probability c . Thus, W is an unbiased and cost-bounded sample set.

Since the probability that a player accepts the contract is c , the value $m := |\{i \in W : D_i = 1\}|$ is a random variable $\text{bin}(n_1, c)$ from a binomial distribution² $\text{Bin}(n_1, c)$. Since the expected value of m is cn_1 , $\frac{m}{c}$ is an unbiased estimate of n_1 . We say $\frac{m}{c}$ is a *naive* estimate of n_1 .

We explain how to produce an estimate that satisfies ϵ -differential privacy. Although the naive estimate is an unbiased estimate of n_1 , it does not satisfy differential privacy. Consider an adjacent pairs of cost-data pairs (v, D) and (v', D') , where D and D' differ in the i -th entry. Let n_1 be the number of player i with $D_i = 1$ and n'_1 be the number of players i with $D'_i = 1$. The naive estimate does not satisfy differential privacy, since if $D_i = 1$ and $v_i \leq \epsilon\alpha_1$, then an outsider can infer D_i easily by comparing the naive estimates of n_1 and n'_1 . Thus, we should introduce a random noise to the naive estimate to satisfy differential privacy.

The mechanism uses the Laplacian distribution as a source of the random noise. The Laplacian noise is commonly used to obtain differential privacy. A *Laplacian distribution* with mean 0 and parameter $b > 0$ is denoted by $\text{Lap}(b)$. The probability density function of $\text{Lap}(b)$ is

$$f(x) = \frac{1}{2b} \exp\left(-\frac{|x|}{b}\right).$$

Let $\text{lap}(b)$ denote a random variable drawn from $\text{Lap}(b)$.

In order to make estimate satisfy differential privacy, the mechanism adds random noise $\text{lap}(\frac{1}{\epsilon})$ to the naive estimate. Since the mean of the Laplacian noise is zero, $s = \frac{1}{c}(m + \text{lap}(\frac{1}{\epsilon}))$ is an unbiased estimate of n_1 . However, s might be larger than n or be negative, both of which are meaningless. We *truncate* s to get \hat{s} , that is when $s > n$, the mechanism outputs n and when $s < 0$, the mechanism outputs 0.

We also use the Laplacian noises to produce payments that satisfy ϵ -differential privacy. By the construction of the contract, for any player i with $D_i = j$ who accepts the contract, the mechanism pays player i for $\epsilon\alpha_j$ in expectation. If the mechanism pays player i for $\epsilon\alpha_j$ deterministically, then an outsider can infer player i 's data easily. Thus, we should introduce noise to the payments. We add noise $\epsilon \text{lap}(\frac{\gamma}{\epsilon})$ to the payment, where $\gamma := |\max_j \alpha_j - \min_j \alpha_j|$. Thus, $p_i = \epsilon(\alpha_j + \text{lap}(\frac{\gamma}{\epsilon}))$. Since the expected value of $\text{lap}(\frac{\gamma}{\epsilon})$ is zero, the expected payment of player i is $\epsilon\alpha_j$, which satisfies the guarantee in the contract. Moreover, since $\epsilon\alpha_j$ is larger than ϵv_i , the mechanism is EIIR. The formal description of the mechanism is in Mechanism 1.

LEMMA 3.1. *Mechanism 1 is perfectly data private.*

PROOF. Let y_i be player i 's reply to the contract. By construction of the contract, if i decides truthfully, then $\Pr[y_i = \text{“accept”} \mid D_i = j] = c$ for all $j \in [h]$. That is, the probability of accepting the contract and D_i are independent. Thus, for any i , the mechanism cannot learn about D_i by reading y_i . Moreover, the mechanism only reads

²A *binomial distribution* with parameter n and p is denoted by $\text{Bin}(n, p)$. The probability density function of $\text{Bin}(n, p)$ is $f(k; n, p) = \binom{n}{k} p^k (1-p)^{n-k}$. Let $\text{bin}(n, p)$ denote a random variable drawn from $\text{Bin}(n, p)$. The expected value of $\text{bin}(n, p)$ is np and variance is $np(1-p)$.

Mechanism 1: ϵ -differentially private mechanism**input** : privacy parameter ϵ ; cost distributions $F_j, j \in [h]$ **output**: estimate \hat{s} ; payment p

- 1 Pick a real number $c \in (0, 1)$
 - 2 Find α_j for all $j \in [h]$, such that $F_j(\alpha_j) = c$.
 - 3 For each player i , offer a contract:
 - 4 If $D_i = j$, the expected payment will be $\epsilon\alpha_j$.
 - 5 Let $W = \{i : i \text{ accepts contract}\}$.
 - 6 Let $m = |\{i \in W : D_i = 1\}|$.
 - 7 Let $s = \frac{1}{\epsilon}(m + \text{lap}(\frac{1}{\epsilon}))$.
 - 8 $\hat{s} = s$ if $s \in [0, n]$, 0 if $s < 0$, n if $s > n$
 - 9 $p_i = \begin{cases} 0 & \text{if } i \notin W \\ \epsilon(\alpha_j + \text{lap}(\frac{\gamma}{\epsilon})), \text{ where } \gamma := |\max_j \alpha_j - \min_j \alpha_j| & \text{if } i \in W \text{ and } D_i = j \end{cases}$
 - 10 **return** (\hat{s}, p)
-

D_i , where $i \in W$. Since player $i \in W$ with $D_i = j$ is paid $\epsilon\alpha_j$ in expectation and $v_i \leq \alpha_j$, the mechanism satisfies the requirement. \square

LEMMA 3.2. *Mechanism 1 is BIC and EIIR.*

PROOF. (BIC) The payments for players who is not in W are always 0. For player i , there are two cases.

Case 1: $D_i = j$ and $v_i \leq \alpha_j$. Accepting the contract will get expected payment $\epsilon(\alpha_j - v_i) \geq 0$.

Case 2: $D_i = j$ and $v_i > \alpha_j$. Accepting the contract will get expected payment $\epsilon(\alpha_j - v_i) < 0$.

(EIIR) Suppose that every player decides truthfully. Then only players with $v_i \leq \alpha_j$ and $D_i = j$ for some j are in W . Since the expected payment for i with $D_i = j$ is $\epsilon\alpha_j$, the expected utility of the player is non-negative. \square

Two random variables x_1 and x_2 are ϵ -mutually bounded, if $\forall I \subseteq \mathbb{R}, \Pr[x_1 \in I] \leq e^\epsilon \Pr[x_2 \in I]$ and $\Pr[x_2 \in I] \leq e^\epsilon \Pr[x_1 \in I]$.

LEMMA 3.3 (FACT 2 IN [GHOSH AND ROTH 2011]). *If x_1 and x_2 are ϵ -mutually bounded and f is a function, then $f(x_1)$ and $f(x_2)$ are also ϵ -mutually bounded. \square*

LEMMA 3.4 ([DWORK ET AL. 2006]). *Let x_1 and x_2 be two random variables. If $|x_1 - x_2| \leq k$, then $x_1 + \text{lap}(\frac{k}{\epsilon})$ and $x_2 + \text{lap}(\frac{k}{\epsilon})$ are ϵ -mutually bounded. \square*

The next two lemmas address the ϵ -differential privacy of the payment and the estimate. Let (v, D) and (v', D') be adjacent cost-data pairs. Let (\hat{s}, p) and (\hat{s}', p') be the results for (v, D) and (v', D') respectively.

LEMMA 3.5. *For any $I \subseteq \mathbb{R}, \Pr[\hat{s} \in I] \leq e^\epsilon \Pr[\hat{s}' \in I]$.*

PROOF. Without loss of generality, we assume that $1 = D_i$ and $D'_i \neq 1$. First, $\Pr[\hat{s} \in I] = \int_{v_{-i} \in \mathbb{R}^{n-1}} \Pr[\hat{s} \in I \mid v_{-i}] \Pr[v_{-i}] dv_{-i}$. Similarly, $\Pr[\hat{s}' \in I] = \int_{v_{-i} \in \mathbb{R}^{n-1}} \Pr[\hat{s}' \in I \mid v_{-i}] \Pr[v_{-i}] dv_{-i}$. Let \hat{q}_w and \hat{q}'_w be two random variables, which are equal to \hat{s} and \hat{s}' when $v_{-i} = w$. If \hat{q}_w and \hat{q}'_w are ϵ -mutually bounded for all w , then \hat{s} and \hat{s}' are ϵ -mutually bounded, since then

$$\Pr[\hat{s} \in I] = \int_{w \in \mathbb{R}^{n-1}} \Pr[\hat{s} \in I \mid v_{-i} = w] \Pr[v_{-i} = w] dw$$

$$\begin{aligned}
&= \int_{w \in \mathbb{R}^{n-1}} \Pr[\hat{q}_w \in I] \Pr[v_{-i} = w] dw \\
&\leq \int_{w \in \mathbb{R}^{n-1}} e^\epsilon \Pr[\hat{q}'_w \in I] \Pr[v_{-i} = w] dw \\
&= \int_{w \in \mathbb{R}^{n-1}} e^\epsilon \Pr[\hat{s}' \in I \mid v_{-i} = w] \Pr[v_{-i} = w] dw = e^\epsilon \Pr[\hat{s}' \in I].
\end{aligned}$$

The case $\Pr[\hat{s}' \in I] \leq e^\epsilon \Pr[\hat{s} \in I]$ can be shown by a symmetric argument.

Here, we show that \hat{q}_w and \hat{q}'_w are ϵ -mutually bounded for all w . Fix $v_{-i} = w$. Let W_w and W'_w be the sets of players accepting the contract when applying the algorithm to inputs (v, D) and (v', D') respectively. Let $m_w := |\{i : D_i = 1, i \in W_w\}|$ and $m'_w := |\{i : D'_i = 1, i \in W'_w\}|$. When applying the mechanism to inputs (v, D) and (v', D') , the mechanism computes $s_w = \frac{1}{c}(m_w + \text{lap}(\frac{1}{\epsilon}))$ and $s'_w = \frac{1}{c}(m'_w + \text{lap}(\frac{1}{\epsilon}))$ respectively. Then, the mechanism truncates s_w and s'_w to get \hat{s}_w and \hat{s}'_w . By Lemma 3.3, since multiplication and truncation are functions, it suffices to show that $m_w + \text{lap}(\frac{1}{\epsilon})$ and $m'_w + \text{lap}(\frac{1}{\epsilon})$ are ϵ -mutually bounded when $v_{-i} = w$. Since $W \setminus W'$ is either the empty set or $\{i\}$, the difference between m_w and m'_w is at most one. Thus, Lemma 3.4 implies that $m_w + \text{lap}(\frac{1}{\epsilon})$ and $m'_w + \text{lap}(\frac{1}{\epsilon})$ are ϵ -mutually bounded. Thus, \hat{q}_w and \hat{q}'_w are ϵ -mutually bounded for all w , and hence \hat{s} and \hat{s}' are mutually bounded. \square

LEMMA 3.6. *For all $i \in [n]$ and for all $I \subseteq \mathbb{R}$, $\Pr[p_i \in I] \leq e^\epsilon \Pr[p'_i \in I]$.*

PROOF. Without loss of generality, we assume that $D_i = 1$ and $D'_i \neq 1$. For player $j \neq i$, if $j \notin W$, the payment is zero. If $j \in W$, the payment to j depends only on the data D_j and does not depend on the set of players receiving payments. Thus, p_j does not change and we only need to consider p_i . Note that $p_i \neq 0$ only happens if player i is in W . If $p_i \neq 0$, then p_i is a random variable $P^1 = \epsilon(\alpha_1 + \text{lap}(\frac{\gamma}{\epsilon}))$. Thus, for any $I \subseteq \mathbb{R} \setminus \{0\}$, the probability $\Pr[p_i \in I] = c \Pr[P^1 \in I]$, where c is the probability of that a player accepts the contract. The probability $\Pr[p_i = 0] = (1 - c) + c \Pr[P^1 = 0]$. Suppose that $D'_i = j'$. Symmetrically, let $P^2 = \epsilon(\alpha_{j'} + \text{lap}(\frac{\gamma}{\epsilon}))$, for any $I \subseteq \mathbb{R} \setminus \{0\}$, the probability $\Pr[p'_i \in I] = c \Pr[P^2 \in I]$ and $\Pr[p'_i = 0] = (1 - c) + c \Pr[P^2 = 0]$.

Thus, it suffices to show that P^1 and P^2 are ϵ -mutually bounded. By Lemma 3.3, since multiplication is a function, it is sufficient to show that $\alpha_1 + \text{lap}(\frac{\gamma}{\epsilon})$ and $\alpha_{j'} + \text{lap}(\frac{\gamma}{\epsilon})$ are ϵ -mutually bounded. By Lemma 3.4, since $|\alpha_1 - \alpha_{j'}| \leq \gamma$, $\alpha_1 + \text{lap}(\frac{\gamma}{\epsilon})$ and $\alpha_{j'} + \text{lap}(\frac{\gamma}{\epsilon})$ are ϵ -mutually bounded. \square

LEMMA 3.7. *Mechanism 1 is $\sqrt{3(\frac{n_1(1-c)}{c} + \frac{2}{\epsilon^2 c^2})}$ -accurate.*

PROOF. Since the error term $|\hat{s} - n_1|$ is smaller than $|s - n_1|$, we can analyze $|s - n_1|$ to get a bound on the error. Since $E[m] = cn_1$, $E[s] = \frac{1}{c}(E[m] + E[\text{lap}(\frac{1}{\epsilon})]) = n_1$ by linearity of expectation.

$$|\hat{s} - n_1| \leq |s - n_1| = \frac{1}{c} |m + \text{lap}(\frac{1}{\epsilon}) - n_1 c| = \frac{1}{c} |\text{bin}(n_1, c) + \text{lap}(\frac{1}{\epsilon}) - E[\text{bin}(n_1, c) + \text{lap}(\frac{1}{\epsilon})]|.$$

In order to prove that accuracy with high probability, we use Chebyshev's inequality.

LEMMA 3.8 (CHEBYSHEV'S INEQUALITY). *Let X be a random variable with expected value μ and variance σ^2 . For any real number $k > 0$, $\Pr[|X - \mu| \geq k\sigma] \leq \frac{1}{k^2}$.*

We set $k = \sqrt{3}$ and let $X \sim \text{bin}(n_1, c) + \text{lap}(\frac{1}{\epsilon})$ with $\text{Var}[X] = n_1c(1-c) + \frac{2}{\epsilon^2}$ to get

$$\Pr \left[\left| \text{bin}(n_1, c) + \text{lap}(\frac{1}{\epsilon}) - E[\text{bin}(n_1, c) + \text{lap}(\frac{1}{\epsilon})] \right| \geq \sqrt{3(n_1c(1-c) + \frac{2}{\epsilon^2})} \right] \leq \frac{1}{3}.$$

This is equivalent to

$$\Pr \left[\frac{1}{c} \left| \text{bin}(n_1, c) + \text{lap}(\frac{1}{\epsilon}) - E[\text{bin}(n_1, c) + \text{lap}(\frac{1}{\epsilon})] \right| \geq \sqrt{3\left(\frac{n_1(1-c)}{c} + \frac{2}{\epsilon^2c^2}\right)} \right] \leq \frac{1}{3}.$$

$$\text{Thus, } \Pr \left[|\hat{s} - n_1| \geq \sqrt{3\left(\frac{n_1(1-c)}{c} + \frac{2}{\epsilon^2c^2}\right)} \right] \leq \frac{1}{3}. \quad \square$$

The mechanism can pick c freely. If the mechanism picks a constant c such that $\frac{n(1-c)}{c} \leq \frac{2}{\epsilon^2c^2}$, the mechanism is $O(\epsilon^{-1})$ accurate.

We will extend this result to general data entry and discrete cost distributions in Section 3.1. Thus, we have the main theorem.

THEOREM 3.9. *Mechanism 1 is BIC, EIIR, $O(\epsilon^{-1})$ -accurate, perfectly data private, and ϵ -differentially private. \square*

3.1. Extensions and Computational Issues

General Database Entries. Suppose that the entry of database has d attributes, that is, $D_i \in [h]^d$. Given a sequence a_1, \dots, a_d , where $a_j \in [h]$, the data analyst wants to estimate $|\{i : \forall_j D_{ij} = a_j\}|$. For any D_i , we can transform D_i to a single attribute data $D'_i = 1 + \sum_{i=0}^{d-1} D_{ij} \times d^i$, such that $D'_i \in [h^d]$. Then, we can apply the mechanism to estimate the number of players with $D'_i = 1 + \sum_{i=0}^{d-1} a_j \times d^i$.

Discrete Cost Distributions. When F_j is a discrete probability function, the major difficulty is that for a given c and j , we may not find a suitable α_j , such that $F_j(\alpha_j) = c$, because the cumulative probability function of a discrete distribution is a step function. However, the mechanism can provide different contracts to different players and this ability allows us to design a mechanism for discrete case.

The basic idea is that the mechanism uses randomness to pick α_j such that every player has equal probability c to accept the contract. For a given c and for each j , if there is no α_j such that $F_j(\alpha_j) = c$, then the mechanism finds the largest α_j^- and the smallest α_j^+ such that $F_j(\alpha_j^-) = c_j^- < c$ and $F_j(\alpha_j^+) = c_j^+ > c$. Note that a player i with $D_i = j$ accepts the contract if his cost is smaller than the expected payment. If the expected payment is α_j^+ , then the player accepts the contract with probability $c_j^+ > c$. On the other hand, if the expected payment is α_j^- , then the player accepts the contract with probability $c_j^- < c$. Let $\beta_j = \frac{c - c_j^-}{c_j^+ - c_j^-}$. Player i is given a contract “If $D_i = j$, your expected payment is $\epsilon\alpha_j$ in expectation,” where $\Pr[\alpha_j = \alpha_j^-] = 1 - \beta_j$ and $\Pr[\alpha_j = \alpha_j^+] = \beta_j$. Thus, $\Pr[v_i \leq \alpha_j] = c_j^- + \beta_j(c_j^+ - c_j^-) = c$, where the randomness is over the distribution of costs and the random choice of α_j . We can prove that the mechanism is perfectly data private, BIC, and EIIR by arguments similar to those in the proofs of Lemmas 3.1 and 3.2. Since every player has equal probability c to accept the contract, we can show that the mechanism satisfies ϵ -differential privacy of estimate and is $O(\epsilon^{-1})$ -accurate by arguments similar to those in the proofs of Lemmas 3.5 and 3.7. In order to satisfy differential privacy of payments, we let $\gamma := \max_j \alpha_j^+ - \min_j \alpha_j^-$.

Then, the payments satisfy ϵ -differential privacy by an argument similar to the proof of Lemma 3.6.

Cost of Mechanism. For a fixed ϵ , when c increases, the accuracy of Mechanism 1 is improved, since Mechanism 1 uses more players' data. However, Mechanism 1's expected total payment also increases. Since Mechanism 1 is $\sqrt{3(\frac{n_1(1-c)}{c} + \frac{2}{\epsilon^2 c^2})}$ -accurate, there is a trade-off between the accuracy and the expected total payment. Since the mechanism can pick c freely, for a given $\epsilon > \sqrt{\frac{8}{n}}$, the mechanism can pick $c = \frac{1 + \sqrt{1 - 8\epsilon^{-2}/n}}{2}$.

Let $\alpha = \max_j \alpha_j$. The expected total payment is $\epsilon \alpha c n = \alpha n (\frac{\epsilon + \sqrt{\epsilon^2 - 8/n}}{2})$. Then, Mechanism 1 picks a suitable ϵ , such that the expected total payment $\epsilon \alpha c n = B$. Hence, the mechanism is budget-feasible in expectation and is $O(\frac{1}{\epsilon c}) = O(\frac{\alpha n}{B})$ accurate.

Fixed Accuracy. If the data analyst wants a k -accurate mechanism, we can pick $c = \frac{1}{1+k^2/6n}$ and $\epsilon = \frac{2\sqrt{3}(1+k^2/6n)}{k}$, such that the mechanism is k -accurate. The expected total payment is $\epsilon \alpha c n = \frac{2\sqrt{3}\alpha n}{k}$.

Computing $F^{-1}(c)$. In an ideal model, when F_j is a continuous distribution, we assume that mechanism can access the closed form of F_j , such that the mechanism can compute $\alpha = F_j^{-1}(c)$ accurately. However, when the mechanism cannot access the closed form of F_j , $F_j^{-1}(c)$ may not be computable. When it is impossible to access the closed form of F_j , we assume that there is an oracle, which returns $F_j(v)$ for any given value v . In the oracle model, the mechanism finds α_j^-, α_j^+ for all j , such that $F_j(\alpha_j^-) < c$, $F_j(\alpha_j^+) > c$, and $\alpha_j^+ - \alpha_j^- < \delta$ for $\delta < 1/n$ using binary search. Then, the mechanism uses the method that we use for discrete cost distributions to construct the contract. That is, let $\beta_j = \frac{c - c_j^-}{c_j^+ - c_j^-}$. Player i is given a contract "If $D_i = j$, your expected payment is $\epsilon \alpha_j$ in expectation," where $\Pr[\alpha_j = \alpha_j^-] = 1 - \beta_j$ and $\Pr[\alpha_j = \alpha_j^+] = \beta_j$. Thus, $\Pr[v_i \leq \alpha_j] = c_j^- + \beta_j(c_j^+ - c_j^-) = c$, where the randomness is over the distribution of costs and the random choice of α_j . Hence, the mechanism is still perfectly data private, BIC, EIIR, ϵ -differential private, and $O(1/\epsilon)$ -accurate. In the oracle model, the expected payment for player i with $D_i = j$ who accepts the contract is at most α_j^+ . In the ideal model, the expected payment for player i with $D_i = j$ who accepts the contract is $F_j^{-1}(c)$, which is smaller than α_j^+ . Since $\alpha_j^+ - F_j^{-1}(c)$ is at most $\delta < 1/n$, the difference between the expected payments in the ideal model and in the oracle model is at most $1/n$ for each player. Thus, the difference between the expected total payment in the ideal model and in the oracle model is at most 1.

4. OPTIMALITY

In this section, we define a benchmark for the expected payment of a mechanism and compare the expected payment of Mechanism 1 to this benchmark in two different settings. The benchmark mechanism is not only truthful but also knows D_i for all i and has no privacy requirements. We show that when all costs are non-negative, Mechanism 1 is provably close to the benchmark.

The benchmark is the minimum expected payment among all truthful mechanisms M^* that satisfy the following properties. In order to get a meaningful estimate, for any $k < n/2$, a k -accurate mechanism learns a subset of players' data. We call this subset a *sample set*. Since obtaining an estimate based on an unbiased sample is a common approach in statistics, we assume that M^* uses an unbiased sample. Suppose

that there are n_j players with $D_i = j$ for $j \in [h]$. Since the sample set is unbiased, there exists c such that M^* buys $w_j = cn_j$ data from players with $D_i = j$. After getting an unbiased sample, M^* uses w_1/c as the straightforward estimate of n_1 . Since the choices of c may affect the accuracy guarantee, we compare the payment of Mechanism 1 to the payment of M^* , where Mechanism 1 and M^* have the same size of sample sets. Thus, M^* is a truthful mechanism that gets an unbiased sample with size cn for a fixed number c .

Since there is no competition between players with data j and players with data $j' \neq j$, M^* can run auctions for players with $D_i = j$ for all $j \in [h]$ independently and buy w_j data from players with $D_i = j$. The mechanism that guarantees buying w units is called *w-unit procurement mechanism*. Thus, M^* is a mechanism that runs a truthful, w_j -unit procurement mechanism for each $j \in [h]$.

Mechanism 1 buys in expectation w_j data from players with $D_i = j$ for $j \in [h]$. We compare the expected payment of Mechanism 1 for buying in expectation w_j data from players with $D_i = j$ with the expected payment of M^* for buying w_j data from players with $D_i = j$ for each j . If the expected payment of Mechanism 1 is at most r times the expected payment of M^* for each j , then the total expected payment of Mechanism 1 is at most r times the total expected payment of M^* . Thus, we focus on a single auction that all players have the same D_i and both Mechanism 1 and the M^* want to buy w data from n players.

For multi-unit procurement mechanisms, let x_i be the indicator random variable denoting whether the mechanism buys from player i . Let v_i be the cost to the player i , if $x_i = 1$. Let p_i be the payment of player i . The utility for player i is $p_i - x_i v_i$. Note that when we consider privacy preserving mechanisms, the utility of player i is $p_i - \epsilon x_i v_i$. However, since ϵ is the same for all players, we can ignore ϵ without loss of generality, that is, scaling every player's cost by ϵ . Without loss of generality, we suppose that players report costs $v_1 \leq v_2 \leq \dots \leq v_n$.

4.1. Envy-free Benchmark

A mechanism is *envy-free* if for all v and for all i, j , $p_i - v_i x_i \geq p_j - v_i x_j$. We show that for any envy-free, multi-unit procurement mechanism, every data that is bought by the mechanism is purchased at the same price. Suppose that a multi-unit procurement mechanism buys data from two players at two different prices. Since the player with the lower price wants to have the higher price, the mechanism is not envy-free. We compare the expected payment of Mechanism 1 with the expected payment of the optimal, envy-free, dominant strategy truthful, multi-unit procurement mechanism. We use envy-free mechanisms as a benchmark, because for procurement mechanisms in a Bayesian setting, the optimal mechanisms are known to charge a fixed price.³

We introduce another commonly used solution concept as follows. A profile of strategies q_1, \dots, q_n is a *dominant strategy equilibrium* if for all i , v_i, v_{-i} , and $y'_i \in Y$, $E[u_i(q(v_i, v_{-i}, D), v_i)] \geq E[u_i((y'_i, q_{-i}(v_{-i}, D_{-i})), v_i)]$, where the randomness is from the mechanism. A direct mechanism is *dominant strategy truthful* if $q_i(v_i, D_i) = v_i$ is a dominant strategy equilibrium for every player i . The following lemma characterizes the total payment for any dominant strategy truthful, EPIR, and envy-free mechanisms.

LEMMA 4.1 (THEOREM 4.6 IN [GHOSH AND ROTH 2011]). *No dominant strategy truthful, EPIR, and envy-free w-unit procurement mechanism can have total payment less than wv_{w+1} . □*

³Envy-free benchmarks are also common in prior-free mechanism design [Hartline and Roughgarden 2008].

Let F be the cumulative distribution function of players' costs, that is, $F(a) = \Pr[v \leq a]$. By Lemma 4.1, the total expected payment of any dominant strategy truthful, EPIR, and envy-free w -unit procurement mechanism is at least $wE_{v \sim F}[v_{w+1}]$. Thus, our benchmark is $wE_{v \sim F}[v_{w+1}]$.

Now, we compare the benchmark with the expected payment of Mechanism 1. There are two cases. First, when there exists α such that $F(\alpha) = \frac{w}{n}$, Mechanism 1 offers a posted price α for each player in order to buy w players' data in expectation. If player i accepts the price, the mechanism buys from player i with expected payment α . Since each player has probability $\frac{w}{n}$ to accept the contract, the total expected payment of Mechanism 1 is $w\alpha$.

Second, when there is no α such that $F(\alpha) = \frac{w}{n}$, we give an extension to Mechanism 1 in Section 3.1. The extension finds the largest α^- and the smallest α^+ , such that $F(\alpha^-) < \frac{w}{n}$ and $F(\alpha^+) > \frac{w}{n}$. Let $c^- := F(\alpha^-)$, $c^+ := F(\alpha^+)$, and $\beta := \frac{\frac{w}{n} - c^-}{c^+ - c^-}$. Then, the mechanism offers a price α^+ with probability β and price α^- with probability $1 - \beta$. For a player with cost at most α^- , since the player always accepts the offer, the expected payment is $(\alpha^-(1 - \beta) + \alpha^+\beta)$. For a player with cost equal to α^+ , since the player accepts the offer only when the offered price is α^+ , the expected payment is $\alpha^+\beta$. For a player with cost larger than α^+ , since the player always rejects the offer, the expected payment is 0. Since each player has a cost at most α^- with probability c^- and has a cost equal to α^+ with probability $c^+ - c^-$, each player's expected payment is $c^-(\alpha^-(1 - \beta) + \alpha^+\beta) + (c^+ - c^-)\alpha^+\beta$. Thus, the total expected payment is $n(c^-(\alpha^-(1 - \beta) + \alpha^+\beta) + (c^+ - c^-)\alpha^+\beta)$ by the linearity of expectation. Moreover,

$$\begin{aligned} n(c^-(\alpha^-(1 - \beta) + \alpha^+\beta) + (c^+ - c^-)\alpha^+\beta) &= n(c^-(\alpha^-(1 - \beta) + \alpha^+\beta) + (\frac{w}{n} - c^-)\alpha^+) \\ &= n(\frac{w}{n}\alpha^+ + c^-(\alpha^-(1 - \beta) + \alpha^+\beta - \alpha^+)) \\ &= n(\frac{w}{n}\alpha^+ + c^-((1 - \beta)(\alpha^- - \alpha^+))) \\ &= n(\frac{w}{n}\alpha^+ - c^-((1 - \beta)(\alpha^+ - \alpha^-))) \\ &= w(\alpha^+ - \frac{nc^-}{w}(1 - \beta)(\alpha^+ - \alpha^-)). \end{aligned}$$

When there exists α , such that $F(\alpha) = \frac{w}{n}$, the expected payment of Mechanism 1 is $w\alpha$. When α does not exist, the expected payment is $w(\alpha^+ - \frac{nc^-}{w}(1 - \beta)(\alpha^+ - \alpha^-))$. Thus, we should compare both $w\alpha$ and $w(\alpha^+ - \frac{nc^-}{w}(1 - \beta)(\alpha^+ - \alpha^-))$ with $wE_{v \sim F}[v_{w+1}]$. It suffices to compare α and $\alpha^+ - \frac{nc^-}{w}(1 - \beta)(\alpha^+ - \alpha^-)$ with $E_{v \sim F}[v_{w+1}]$.

LEMMA 4.2. *1. If there exists α such that $F(\alpha) = \frac{w}{n}$, then $E_{v \sim F}[v_{w+1}] \geq \frac{1}{2}\alpha$.
2. If there is no α such that $F(\alpha) = \frac{w}{n}$, then $E_{v \sim F}[v_{w+1}] \geq \frac{1}{2}(\alpha^+ - \frac{nc^-}{w}(1 - \beta)(\alpha^+ - \alpha^-))$.*

PROOF. We show the second statement. The first statement follows by setting $\alpha^- = \alpha^+ = \alpha$.

Let $\eta = \alpha^+ - \frac{nc^-}{w}(1 - \beta)(\alpha^+ - \alpha^-)$. By conditional probability,

$$\begin{aligned} E[v_{w+1}] &= \Pr[v_{w+1} \leq \eta] \times E[v_{w+1} \mid v_{w+1} \leq \eta] + \Pr[v_{w+1} > \eta] \times E[v_{w+1} \mid v_{w+1} > \eta] \\ &\geq \Pr[v_{w+1} > \eta] \times \eta \quad (\text{costs are non-negative}). \end{aligned}$$

It suffices to show that $\Pr[v_{w+1} > \eta] \geq \frac{1}{2}$. Since $c^- < \frac{w}{n}$, $\frac{nc^-}{w} < 1$. Since $\beta < 1$ and $\frac{nc^-}{w} < 1$, $\alpha^- < \eta < \alpha^+$. If $v_{w+1} > \eta$, then $v_{w+1} \geq \alpha^+$, since α^+ is the smallest num-

ber larger than α^- with non-zero probability. Let $v_{(i)}$ denote the cost of player i . If $v_{w+1} \geq \alpha^+$, then at most w players' $v_{(i)}$ are no larger than α^- . Since each $v_{(i)}$ is independently drawn from F , $\Pr[v_{(i)} \leq \alpha^-] = c^-$. Let X_i be the indicator random variable such that $X_i = 1$ if $v_{(i)} \leq \alpha^-$, otherwise $X_i = 0$. Let $X = \sum_{i=1}^n X_i$. The probability that at most w players have $v_{(i)}$ no larger than α^- is $\Pr[X \leq w]$. Since the X_i 's are independent, identical, indicator random variables, X is a random variable from a binomial distribution $\text{Bin}(n, c^-)$. Thus, $\Pr[v_{w+1} \geq \alpha^+] = \Pr[\text{bin}(n, c^-) \leq w]$.

Now, we show that $\Pr[\text{bin}(n, c^-) \leq w] \geq \frac{1}{2}$. We say m is the *median* of a distribution D over real numbers if, $\Pr[Z \leq m] \geq \frac{1}{2}$ and $\Pr[Z \geq m] \geq \frac{1}{2}$, where Z is a random variable drawn from D . For a binomial distribution $\text{Bin}(n, p)$, the expected value np and the median m satisfy $\lfloor np \rfloor \leq m \leq \lceil np \rceil$ [Kaas and Buhrman 1980]. Since $c^- < \frac{w}{n}$, the expected value of $\text{bin}(n, c^-)$ is smaller than w . Since $\lceil nc^- \rceil \leq w$, the median m of $\text{Bin}(n, c^-)$ is at most w . Thus, $\Pr[\text{bin}(n, c^-) \leq w] \geq \frac{1}{2}$. \square

Lemmas 4.1 and 4.2 imply the following theorem.

THEOREM 4.3. *Mechanism 1's expected payment is 2-approximate to the benchmark.* \square

4.2. Anti-regular Distributions

In this section, we compare the expected payment of Mechanism 1 with the expected payment of the optimal BIC, multi-unit procurement mechanism. We first characterize randomized BIC procurement mechanisms. For a randomized mechanism and a given bid v_i , let $\bar{x}_i(v_i)$ be the probability that the mechanism buys from player i and let $p_i(v_i)$ be the random variable denoting the payment for player i , where both \bar{x}_i and p_i 's randomness come from the mechanism and v_{-i} . Suppose that when $v_i = \infty$, the mechanism will not buy from player i . That is, $\bar{x}_i(\infty) = 0$ and $E[p_i(\infty)] = 0$. The characterization for the BIC, procurement mechanisms is analogous to the characterization of BIC selling mechanisms, which is a well-known result in auction theory.

LEMMA 4.4. *A randomized procurement mechanism is BIC if and only if for every i the procurement probability \bar{x} and payment p satisfies*

- (i) $\bar{x}_i(v_i)$ is decreasing in v_i ;
- (ii) $E[p_i(v_i)] = v_i \bar{x}_i(v_i) + \int_{v_i}^{\infty} \bar{x}_i(t) dt$. \square

To prove the optimality of selling mechanisms, Myerson [Myerson 1981] introduces a virtual value function. The analogous function for procurement mechanisms is a *virtual cost function*, which is $\phi(z) := z + \frac{F(z)}{f(z)}$. Thus, to ensure that $\phi(z)$ is well-defined and the integral of f is well-defined (used in the proof of Lemma 4.5 and Lemma 4.7), we assume

ASSUMPTION 1. *Let f be the density probability function of distribution F with range $[a, b] \subseteq [0, \infty)$. f is piecewise continuous and $f(z)$ is positive for all $z \in [a, b]$.*

A distribution F is *anti-regular* if F satisfies Assumption 1 and $\phi(z)$ is increasing in z .⁴

When the distribution F is anti-regular, [Elkind et al. 2004] characterize the optimal dominant strategy truthful mechanism to minimize the expected payment for path auctions. Although their problem is not exactly the same as w -unit procurement mechanisms, their result can be extended to procurement mechanisms easily.

⁴For selling mechanisms, a distribution is *regular* if the *virtual value* $\phi'(z) = z - \frac{1-F(z)}{f(z)}$ is increasing in z .

LEMMA 4.5. *When the distribution F is anti-regular, the optimal BIC w -unit procurement buys from the w players with the smallest virtual cost. \square*

Since $\phi(z)$ is increasing in z , the optimal mechanism buys from the first w players. By Lemma 4.4, the expected payment for player $i \leq w$ is v_{w+1} . Thus, the total expected payment of the optimal BIC mechanism is $wE_{v \sim F}[v_{w+1}]$. Thus, our benchmark is $wE_{v \sim F}[v_{w+1}]$. We compare the expected payment of Mechanism 1 with $wE_{v \sim F}[v_{w+1}]$, when F is anti-regular.

THEOREM 4.6. *When F is anti-regular, Mechanism 1's expected payment is 2-approximate to the benchmark.*

PROOF. Since F satisfies Assumption 1 by definition of anti-regular, F^{-1} is well-defined. The total expected payment of Mechanism 1 is $wF^{-1}(\frac{w}{n})$. When F is anti-regular, the benchmark is $wE_{v \sim F}[v_{w+1}]$. By Lemma 4.2, Mechanism 1 is 2-approximate. \square

4.3. General Distributions

When the distribution satisfies Assumption 1 but $\phi(z)$ is not increasing in z , buying from the w players with smallest virtual cost may result in a non-truthful mechanism. We can use the ironing procedure, which is designed by Myerson [Myerson 1981], to resolve this issue. For a fixed cost vector v , ironing procedure irons on interval $[a, b)$, if $v_i \in [a, b)$, then v_i is replaced by a random number v'_i , which is drawn from the distribution F on $[a, b)$. By a way similar to Myerson's method, we can identify a set S of intervals, such that the *ironed virtual cost function* $\bar{\phi}(z) = E[\phi(z)]$ is increasing in z . Moreover, for an ironed interval $[a, b)$, $\bar{\phi}(z)$ is the same for all $z \in [a, b)$.

LEMMA 4.7. *The w -unit procurement mechanism that buys from the w players with smallest ironed virtual cost and breaks ties uniformly at random is the optimal BIC mechanism when the distribution satisfies Assumption 1. \square*

Thus, our benchmark is the expected payment of the optimal BIC mechanism, M , when the distribution satisfies Assumption 1. In order to calculate the expected payment of M , we specify the payment rule as follows. Let $\bar{x}_i(v_i, v_{-i})$ be the probability that M buys from player i , where the randomness comes from the mechanism. Since M buys from the w players with smallest ironed virtual cost, $\bar{x}_i(v_i, v_{-i})$ is decreasing in v_i for any fixed v_{-i} . Let $p_i(v_i, v_{-i})$ be the random variable denoting the payment for player i , where $E[p_i(v_i, v_{-i})] = v_i \bar{x}_i(v_i, v_{-i}) + \int_{v_i}^{\infty} \bar{x}_i(t, v_{-i}) dt$ and the randomness comes from the mechanism. It is easy to see that this payment rule satisfies Lemma 4.4.

We compare the expected payment of Mechanism 1 with the benchmark.

THEOREM 4.8. *Let F satisfy Assumption 1. Let S be the set of ironed intervals for F . If every interval $[a, b) \in S$ satisfies $a \geq b/r$ for some $r > 1$, then the expected payment of Mechanism 1 is $2r$ -approximate to the benchmark.*

PROOF. Since F satisfies Assumption 1, F^{-1} is well-defined. The expected payment of Mechanism 1 is $wF^{-1}(\frac{w}{n})$. We compare the expected payment of the optimal BIC mechanism, M , with $wF^{-1}(\frac{w}{n})$. Let $p_i(v)$ be the random variable representing the payment for player i in M when the cost vector is v . We show that $E_{v \sim F, M}[\sum_{i \in [n]} p_i(v)] \geq E_{v \sim F}[v_{w+1}]/r$, which implies $E_{v \sim F, M}[\sum_{i \in [n]} p_i(v)] \geq F^{-1}(\frac{w}{n})/2r$ by Lemma 4.2 and hence Mechanism 1 is $2r$ -approximate.

There are two sources of randomness in mechanism M . One is from the cost vector v since v is drawn from a distribution F . Another one is M itself since M is a randomized mechanism. For a fixed cost vector v , let p_i^v be the random variable representing the

payment for player i , where the randomness only comes from M . We show that for any fixed v , $E_M[\sum_{i \in [n]} p_i^v] \geq wv_{w+1}/r$. This implies $E_{v \sim F, M}[\sum_{i \in [n]} p_i(v)] \geq E_{v \sim F}[v_{w+1}]/r$. There are three cases.

Case 1: v_{w+1} is not in any ironed interval. Since M chooses the w players with smallest ironed virtual costs and the ironed virtual cost is increasing, M buys from the first w players. For player $i \leq w$, if v_i increases to $t < v_{w+1}$, by the monotonicity of $\bar{\phi}$, the mechanism still buys from player i . That is $\bar{x}_i(t, v_{-i}) = 1$ for all $t < v_{w+1}$. When $t > v_{w+1}$, the mechanism will not buy from player i . Thus, by definition of the expected payment, the expected payment for each player $i \leq w$ is $v_i + \int_{v_i}^{\infty} \bar{x}_i(t, v_{-i}) dt = v_i + \int_{v_i}^{v_{w+1}} \bar{x}_i(t, v_{-i}) dt = v_{w+1}$. Since expected payment for player $i > w$ is 0, $E_M[\sum_{i \in [n]} p_i^v] = wv_{w+1}$.

Case 2: v_{w+1} is in an ironed interval $[a, b)$ but $v_w \notin [a, b)$. Since for all player $i \leq w$, $v_i \notin [a, b)$, M buys from the first w players. For player $i \leq w$, $\bar{x}_i(t, v_{-i}) = 1$ for all $t < a$. By definition of the expected payment, the expected payment for player $i \leq w$ is $v_i + \int_{v_i}^{\infty} \bar{x}_i(t, v_{-i}) dt \geq v_i + \int_{v_i}^a \bar{x}_i(t, v_{-i}) dt = a$. Thus, $E_M[\sum_{i \in [n]} p_i^v] \geq wa \geq wb/r \geq wv_{w+1}/r$.

Case 3: v_{w+1} and v_w are in the same ironed interval $[a, b)$. Let $l_1 = |\{i : v_i < a\}|$ and $l_2 = |\{i : v_i \in [a, b)\}|$. Thus, $l_1 < w$ and $l_1 + l_2 > w$. The mechanism always buys from the first l_1 players. Since $\bar{\phi}(t)$ is the same for all $t \in [a, b)$ and the mechanism breaks ties uniformly at random, the mechanism buys from player i , $l_1 + 1 \leq i \leq l_1 + l_2$, with probability $\frac{w-l_1}{l_2}$. For player $i \leq l_1$, $\bar{x}_i(t, v_{-i}) = 1$ if $t < a$. By definition of the expected payment, the expected payment for player $i \leq l_1$ is $v_i + \int_{v_i}^{\infty} \bar{x}_i(t, v_{-i}) dt \geq v_i + \int_{v_i}^a \bar{x}_i(t, v_{-i}) dt = a$. For player i , $l_1 < i \leq l_1 + l_2$, when v_i increases to $t < b$, since $\bar{\phi}(t)$ is the same for all $t \in [a, b)$, the probability that the mechanism buys from player i does not change. That is, $\bar{x}_i(t, v_{-i}) = \frac{w-l_1}{l_2}$ if $t \in [a, b)$. By definition of the expected payment, the expected payment for player i , $l_1 < i \leq l_1 + l_2$, is $v_i \bar{x}_i(v_i, v_{-i}) + \int_{v_i}^{\infty} \bar{x}_i(t, v_{-i}) dt = \frac{b(w-l_1)}{l_2}$. Therefore, $E_M[\sum_{i \in [n]} p_i^v] \geq al_1 + l_2 \frac{b(w-l_1)}{l_2} \geq wa \geq wb/r \geq wv_{w+1}/r$. \square

REFERENCES

- CHAUDHURI, K. AND MISHRA, N. 2006. When random sampling preserves privacy. In *CRYPTO*. Lecture Notes in Computer Science Series, vol. 4117. Springer, 198–213.
- CHAWLA, S., HARTLINE, J. D., MALEC, D. L., AND SIVAN, B. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*. STOC '10. ACM, New York, NY, USA, 311–320.
- CHEN, Y., CHONG, S., KASH, I., MORAN, T., AND VADHAN, S. 2011. Truthful mechanisms for agents that value privacy. *arXiv:1111.5472*.
- DWORK, C. 2008. Differential privacy: A survey of results. In *Proceedings of Theory and Applications of Models of Computation, 5th International Conference*. Lecture Notes in Computer Science Series, vol. 4978. Springer, 1–19.
- DWORK, C., MCSHERRY, F., NISSIM, K., AND SMITH, A. 2006. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Third Theory of Cryptography Conference*. Lecture Notes in Computer Science Series, vol. 3876. Springer, 265–284.
- ELKIND, E., SAHAI, A., AND STEIGLITZ, K. 2004. Frugality in path auctions. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '04. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 701–709.
- FEIGENBAUM, J., JAGGARD, A. D., AND SCHAPIRA, M. 2010. Approximate privacy: foundations and quantification (extended abstract). In *Proceedings of the 11th ACM Conference on Electronic Commerce*. EC '10. ACM, New York, NY, USA, 167–178.
- GHOSH, A. AND ROTH, A. 2011. Selling privacy at auction. In *Proceedings of the 12th ACM Conference on Electronic Commerce*. EC '11. ACM, New York, NY, USA, 199–208.

- GOLDBERG, A. V. AND HARTLINE, J. D. 2005. Collusion-resistant mechanisms for single-parameter agents. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '05. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 620–629.
- HARTLINE, J. D. AND ROUGHGARDEN, T. 2008. Optimal mechanism design and money burning. In *Proceedings of the 40th annual ACM Symposium on Theory of Computing*. STOC '08. New York, NY, USA, 75–84.
- KAAS, R. AND BUHRMAN, J. 1980. Mean, median and mode in binomial distributions. *Statistica Neerlandica* 34, 1, 13–18.
- KLONOWSKI, M., PRZYKUCKI, M., STRUMINSKI, T., AND SULKOWSKA, M. 2010. Practical universal random sampling. In *Proceedings of Advances in Information and Computer Security - 5th International Workshop on Security*. Lecture Notes in Computer Science Series, vol. 6434. Springer, 84–100.
- MCSHERRY, F. AND TALWAR, K. 2007. Mechanism design via differential privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society, Washington, DC, USA, 94–103.
- MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1, 58–73.
- NISAN, N., ROUGHGARDEN, T., TARDOS, E., AND VAZIRANI, V. V. 2007. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, Chapter 9.
- NISSIM, K., ORLANDI, C., AND SMORODINSKY, R. 2011. Privacy-aware mechanism design. *arXiv:1111.3350*.
- ROCKAFELLAR, R. 1997. *Convex Analysis*. Princeton University Press.
- XIAO, D. 2011. Is privacy compatible with truthfulness? Cryptology ePrint Archive, Report 2011/005.