



6-5-2011

Selling Privacy at Auction

Aripta Ghosh
Yahoo Research

Aaron Roth
University of Pennsylvania, aaroth@cis.upenn.edu

Follow this and additional works at: http://repository.upenn.edu/cis_papers

 Part of the [Computer Sciences Commons](#)

Recommended Citation

Aripta Ghosh and Aaron Roth, "Selling Privacy at Auction", . June 2011.

© ACM, 2011. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in *Proceedings of the 12th ACM conference on Electronic commerce*, {(2011)} <http://doi.acm.org/10.1145/1993574.1993605> Email permissions@acm.org
doi>[10.1145/1993574.1993605](http://doi.acm.org/10.1145/1993574.1993605)

Arpita Ghosh and Aaron Roth. 2011. Selling privacy at auction. In *Proceedings of the 12th ACM conference on Electronic commerce (EC '11)*. ACM, New York, NY, USA, 199-208. DOI=10.1145/1993574.1993605 <http://doi.acm.org/10.1145/1993574.1993605>

This paper is posted at ScholarlyCommons. http://repository.upenn.edu/cis_papers/717
For more information, please contact libraryrepository@pobox.upenn.edu.

Selling Privacy at Auction

Abstract

We initiate the study of markets for private data, through the lens of differential privacy. Although the purchase and sale of private data has already begun on a large scale, a theory of privacy as a commodity is missing. In this paper, we propose to build such a theory. Specifically, we consider a setting in which a data analyst wishes to buy information from a population from which he can estimate some statistic. The analyst wishes to obtain an accurate estimate cheaply, while the owners of the private data experience some cost for their loss of privacy, and must be compensated for this loss. Agents are selfish, and wish to maximize their profit, so our goal is to design truthful mechanisms. Our main result is that such problems can naturally be viewed and optimally solved as variants of multi-unit procurement auctions. Based on this result, we derive auctions which are optimal up to small constant factors for two natural settings: 1. When the data analyst has a fixed accuracy goal, we show that an application of the classic Vickrey auction achieves the analyst's accuracy goal while minimizing his total payment. 2. When the data analyst has a fixed budget, we give a mechanism which maximizes the accuracy of the resulting estimate while guaranteeing that the resulting sum payments do not exceed the analyst's budget. In both cases, our comparison class is the set of envy-free mechanisms, which correspond to the natural class of fixed-price mechanisms in our setting. In both of these results, we ignore the privacy cost due to possible correlations between an individual's private data and his valuation for privacy itself. We then show that generically, no individually rational mechanism can compensate individuals for the privacy loss incurred due to their reported valuations for privacy. This is nevertheless an important issue, and modeling it correctly is one of the many exciting directions for future work.

Disciplines

Computer Sciences

Comments

© ACM, 2011. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in *Proceedings of the 12th ACM conference on Electronic commerce*, {(2011)} <http://doi.acm.org/10.1145/1993574.1993605> Email permissions@acm.org

doi>[10.1145/1993574.1993605](http://doi.acm.org/10.1145/1993574.1993605)

Arpita Ghosh and Aaron Roth. 2011. Selling privacy at auction. In *Proceedings of the 12th ACM conference on Electronic commerce (EC '11)*. ACM, New York, NY, USA, 199-208. DOI=10.1145/1993574.1993605 <http://doi.acm.org/10.1145/1993574.1993605>

Selling Privacy at Auction

Arpita Ghosh
Yahoo Research
Santa Clara, CA
arpita@yahoo-inc.com

Aaron Roth
Microsoft Research New England
Cambridge, MA
alroth@cs.cmu.edu

November 19, 2010

Abstract

We initiate the study of markets for private data, through the lens of differential privacy. Although the purchase and sale of private data has already begun on a large scale, a theory of privacy as a commodity is missing. In this paper, we propose to build such a theory. Specifically, we consider a setting in which a data analyst wishes to buy information from a population from which he can estimate some statistic. The analyst wishes to obtain an *accurate* estimate *cheaply*, while the owners of the private data experience some cost for their loss of privacy, and must be compensated for this loss. Agents are selfish, and wish to maximize their profit, so our goal is to design *truthful mechanisms*.

Our main result is that such problems can naturally be viewed and optimally solved as variants of multi-unit procurement auctions. Based on this result, we derive auctions which are optimal up to small constant factors for two natural settings:

1. When the data analyst has a fixed accuracy goal, we show that an application of the classic Vickrey auction achieves the analyst's accuracy goal while minimizing his total payment.
2. When the data analyst has a fixed budget, we give a mechanism which maximizes the accuracy of the resulting estimate while guaranteeing that the resulting sum payments do not exceed the analyst's budget.

In both cases, our comparison class is the set of envy-free mechanisms, which correspond to the natural class of fixed-price mechanisms in our setting.

In both of these results, we ignore the privacy cost due to possible correlations between an individual's private data and his valuation for privacy itself. We then show that generically, no individually rational mechanism can compensate individuals for the privacy loss incurred due to their reported valuations for privacy. This is nevertheless an important issue, and modeling it correctly is one of the many exciting directions for future work.

1 Introduction

Organizations such as the Census Bureau and hospitals have long maintained databases of personal information. However, with the advent of the Internet, many corporations are now able to aggregate enormous quantities of sensitive information, and use, buy, and sell it for financial gain. Up until recently, the purchase and sale of private information was the exclusive domain of aggregators – it was obtained for free from the actual owners of the data, for whom it was sensitive. However, recently, companies such as “mint.com” and “Bynamite” have started acting as brokers for private information at the consumer end, paying users for access to their sensitive information [Loh10, Cli10]. Many others, such as Yahoo, Microsoft, Google, and Facebook are also implicitly engaging in the purchase of private information in exchange for non-monetary compensation. In short, “privacy” has become a commodity that has already begun to be bought and sold, in a variety of ad-hoc ways.

Despite the commoditization of privacy in practice, markets for privacy lack a theoretical foundation. In this paper, we initiate the rigorous study of markets *for* private data. Our goal is not to provide a complete solution for the myriad problems involved in the sale of private data, but rather to introduce a crisp model with which to investigate some of the many questions unique to the sale of private data. The study of privacy as a commodity is of immediate relevance, and also a source of many interesting theoretical problems: we hope that this paper elicits more new questions than it answers.

First, let us briefly consider some of the issues that make *privacy* distinct from other commodities that we often deal with, and why this may complicate its sale:

1. First and foremost, in order to sell *privacy*, it is important to be able to define and quantify what we mean by privacy. In this regard, the commoditization of privacy has dovetailed nicely with the development of the theoretical underpinnings of privacy: recent work on *differential privacy* [DMNS06] provides a compelling definition and a precise way in which to quantify its sale. Importantly, as we will discuss, the guarantee of differential privacy has a natural utility-theoretic interpretation that makes it a natural quantity to buy and sell.
2. Private data is a good that exhibits intrinsic complementarities: a data analyst will typically not be interested in the private data of any particular individual, but rather in a representative sample from a large population. Nevertheless, he must purchase the data from particular individuals! Clearly, if there may be unknown correlations between individuals values for privacy and their private data, then the typical strategy of “buying from the cheapest sellers” is doomed to fail in this regard. How should an auction be structured by an analyst who wishes to calculate some value which is representative of an entire population?
3. An individual’s cost for privacy may itself be private information. Suppose that Alice visits an oncologist, and subsequently is observed to significantly increase her value for privacy: this is of course disclosive! Is it possible to run an auction for private data that compensates individuals for the privacy loss they incur, simply due to the effect that their bids have on the behavior of the mechanism?

1.1 Differential Privacy as a Commodity

Differential privacy, formally defined in Section 2, was introduced by Dwork et al. [DMNS06] as a technical definition for database privacy. Informally, an algorithm is ϵ -differentially private if changing the data of a single individual does not change the probability of any outcome of the mechanism by more than an $\exp(\epsilon) \approx (1 + \epsilon)$ multiplicative factor. Differential privacy also has a natural utility-theoretic interpretation that makes it a compelling measure with which to quantify privacy when buying or selling it¹.

An important property of an ϵ -differentially private algorithm A is that its composition with any other database-independent function f has the property that $f(A)$ remains ϵ -differentially private. This allows us to reason about events that might seem quite far removed from the actual output of the algorithm. Quite literally, a guarantee of ϵ -differential privacy is a guarantee that the probability of receiving phone calls during dinner, or of being denied health insurance will not increase by more than an $\exp(\epsilon)$ factor. This allows us to interpret differential privacy as a strong utility theoretic guarantee that holds simultaneously for arbitrary, unknown utility functions: for any individual, with any utility function u over (arbitrary) future events, an ϵ -differentially private computation will decrease his future expected utility by at most an $\exp(-\epsilon) \approx (1 - \epsilon)$ multiplicative factor, or equivalently, by an $\epsilon E[u(x)]$ additive factor, where the expectation is taken over all future events that the individual has preferences over. Therefore, there is a natural way for an individual to assign a cost to the use of his data in an ϵ -differentially

¹This utility theoretic interpretation has been used in another context: the work of McSherry and Talwar, and Nissim, Smorodinsky and Tennenholtz [MT07, NST10] using differential privacy as a tool for traditional mechanism design.

private manner: it should be worth to him an ϵ -fraction of his expected future utility. We expand on this in Appendix A.

1.2 Results

In this paper, we study the following stylized model. There are n individuals $[n]$, each of whom possesses a private bit b_i , which is already known by the administrator of the private database (for example, a hospital). Each individual also has a certain *cost* function $c_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, which determines what her cost $c_i(\epsilon)$ is for her private bit b_i to be used in an ϵ -differentially private manner. Any feasible mechanism must pay each individual enough to compensate him for the use of his private data. Moreover, individuals may mis-report their cost functions in an attempt to maximize their payment, and so we are interested in mechanisms which properly incentivize individuals to report their true cost for privacy. On the other side of the market, the *data analyst* wishes to estimate the quantity $s = \sum_{i=1}^n b_i$, and must compensate each individual through the mechanism's payments for this estimate. The data analyst may either have a fixed accuracy objective and wish to minimize his payments subject to obtaining the desired accuracy, or alternately have a fixed budget and wish to maximize the accuracy of his estimate within this budget.

We first consider the simpler model, in which individuals must be compensated for loss of privacy to their bits b_i , but not for any privacy-leakage due to implicit correlations between b_i and their cost function c_i (*i.e.*, if the mechanism does not use an individual's bit b_i at all in computing an estimate for the data analyst, the mechanism does not have to compensate individual i , even if changing her cost function would result in a different outcome for the mechanism). In trying to design an auction that guarantees the data analyst an accurate estimate of s , one might consider any number of complicated mechanisms that (for example) randomly sample individuals, and then attempt to buy from entire random samples – there are many variations therein, and indeed, this was the direction from which we first explored the problem. Our main result is that it is not necessary to consider such mechanisms. We show that we may abstract away the structure of the mechanism, and without loss of generality consider multi-unit procurement auctions. This has some immediate consequences: if we are interested in the setting for which the data analyst has a fixed accuracy goal, subject to which he wishes to minimize his payment, then we show that the standard VCG mechanism is optimal among the set of envy-free mechanisms. If we are instead interested in the setting for which the data analyst has a fixed budget subject to which he wishes to maximize his accuracy, then we are in a more unusual procurement-auction setting: the buyer wishes to maximize the number of sellers he can buy from, and the cost to the sellers is a function of who else sells their data! In this setting, we give a truthful mechanism that is instance-by-instance optimal among the set of all fixed-price (envy free) mechanisms. We remark that our choice of fixed-price mechanisms as a benchmark has become standard in prior-free mechanism design (see, e.g. [HK07, HR08]), but stands on firmer ground in auction settings for which Bayesian-optimal mechanisms are known also to charge fixed prices. We operate in a setting in which Bayesian-optimal mechanisms are not known, and so justifying (or improving) this choice of benchmark in our setting is an interesting open problem.

We then show a generic impossibility result: it is not, in general, possible for any mechanism to compensate individuals for their privacy loss due to unknown correlations between their private bits b_i and their cost functions c_i . If their costs are known to lie in some fixed range initially, it is possible to offer them some non-trivial privacy guarantee, but finding the correct model in which to study the issue of unknown correlations between data and valuation for privacy is another important direction in which to take this research agenda.

1.3 Related Work

1.3.1 Differential Privacy as a Tool in Mechanism Design

McSherry and Talwar proposed that differential privacy could itself be used as a *solution concept* in mechanism design [MT07]. They observed that a differentially private mechanism is approximately truthful, while simultaneously having some resilience to collusion. Using differential privacy as a solution concept as opposed to dominant strategy truthfulness, they gave some improved results in a variety of auction settings. Gupta et al. also used differential privacy as a solution concept in auction design [GLM⁺10].

In a beautiful follow-up paper, Nissim, Smorodinsky, and Tennenholtz [NST10] made the point that differential privacy may not be a compelling solution concept when beneficial deviations are easy to find (as indeed they are in the mechanism of [MT07]). Nevertheless, they demonstrated a generic methodology for using differentially private mechanisms as tools for designing exactly truthful mechanisms *that do not require payments*, and demonstrate the utility of this framework by designing new mechanisms for several problems.

In this paper, we consider an orthogonal problem: we do not try to use differential privacy as a tool in traditional mechanism design, but instead try to use the tools of traditional mechanism design to *sell* differential privacy as a commodity. Nevertheless, we also use the utility theoretic properties of differential privacy that allow McSherry and Talwar to prove that it implies approximate truthfulness to motivate why it is natural for individuals to have linear cost functions for differential privacy.

1.3.2 Auctions Which Preserve Privacy

Recently, Feigenbaum, Jaggard, and Schapira considered (using a different notion of privacy) how the implementation of an auction can affect how many bits of information are leaked about individuals bids [FJS10]. Specifically, they study to what extent information must be leaked in second price auctions and in the millionaires problem. Protecting the privacy of bids is an important problem, and although it is not the main focus of this paper, we consider it in the context of differential privacy in Section 5. We consider somewhat orthogonal notions of privacy and implementation that make our results incomparable to those of [FJS10].

1.3.3 Privacy in the Economics Literature

Privacy and its relation to mechanism design has also been studied in the economics literature, although primarily in the context of how preferences for privacy by agents may affect mechanisms, rather than in the context of markets *for* privacy. For example, Calzolari and Pavan study the optimal disclosure policy when designing contracts for buyers who are in the position of repeatedly choosing between multiple sellers [CP06], and the recent work of Taylor, Conitzer, and Wagman [TCW10] studies the relationship between the ability of consumers to keep their identity private, and the ability of a monopolist to engage in price discrimination.

One exception is the essay of Laudon [Lau96], which proposes the idea of a market for personal information—a ‘National Information Market’—where individuals can choose to sell or lease their information (possibly to be used in aggregation with other individuals’ information) in exchange for a share of the revenue generated from its use; he argues that only individuals whose cost from the ‘annoyance’ caused by releasing their information is lower than the payment they receive will participate in this market. Our individually rational auctions for privacy are conceptually similar to this, but are investigated within the formal framework of differential privacy, and from the perspective of auction design.

1.3.4 Relationship to the Privacy Literature

The now large literature on differential privacy (see [Dwo08] for an excellent overview) has almost exclusively focused on techniques for guaranteeing ϵ -differential privacy for various tasks, where ϵ has been taken as a given parameter. What has been almost entirely missing is any normative guidance for how to pick ϵ . There is a natural tradeoff between the privacy parameter ϵ and the accuracy of privacy-preserving estimates (which is well-understood in the case of single statistics, see [GRS09, BN10]). Therefore, this paper proposes to answer the question of how ϵ should be chosen: it should be the smallest value that the data analyst is able to afford, given the individuals’ valuations for privacy (or equivalently, the smallest value that the owners of the data are willing to accept in exchange for their payment).

We also highlight in this work the explicit tradeoff between compensating individuals for the use of their private information, and the accuracy of our resulting estimates. Implicit in previous works on privacy has been the idea that for fixed values of ϵ , individuals should be willing to participate in private databases given only some small positive incentive. However, this incentive may be different for different individuals, and without running an auction, a data collector is engaging in selection bias: he is only collecting data from those individuals who value their privacy at a low enough level to make participation in a given database worth while. Such individuals might not be representative of the general population, and resulting estimates may therefore be inaccurate. This source of inaccuracy is hidden in previous works, but we point out that it should be a real concern, and we explicitly address it in this paper.

2 Preliminaries

We consider a database consisting of the data of n individuals $\{1, \dots, n\}$ whom we denote by $[n]$. Each individual i is associated with a private bit $b_i \in \{0, 1\}$. (We may think of this bit as representing the answer to some arbitrary yes or no question). Each bit b_i is already known to a trusted database administrator (for example, a hospital), and so throughout our discussion, we will not endow individuals with the ability to lie about their private bit.

2.1 Differential Privacy

We say that the collection of all individuals' private bits is a *database* $D \in \{0, 1\}^n$. Two databases $D, D^{(i)} \in \{0, 1\}^n$ are *neighbors* if they differ only in the private bit of a single individual, *i.e.*, if $D_j = D_j^{(i)}$ for all $j \neq i$. The quantification of privacy we employ is that of *differential privacy*, due to Dwork et al. [DMNS06]:

Definition 2.1. An algorithm $A : \{0, 1\}^n \rightarrow \mathbb{R}$ satisfies ϵ_i -differential privacy with respect to individual i if for any pair of neighboring databases $D, D^{(i)} \in \{0, 1\}^n$ differing only in their i 'th bit, and for any $S \subset \mathbb{R}$:

$$\frac{\Pr[A(D) \in S]}{\Pr[A(D^{(i)}) \in S]} \leq e^{\epsilon_i}$$

An algorithm A is ϵ_i -*minimally private* with respect to individual i if $\epsilon_i = \inf \epsilon$ such that A is ϵ -differentially private with respect to individual i . Throughout this paper, whenever we say that an algorithm is ϵ_i -differentially private, we mean that it is ϵ_i -minimally differentially private.

Remark 2.2. ϵ -differential privacy becomes less meaningful for large values of ϵ . In this paper, we will restrict our attention to values of $\epsilon < 1$. Note that in this case, $\exp(\epsilon) \approx 1 + \epsilon$.

The following easy fact follows from a hybrid argument:

Fact 1. Consider an algorithm $A : \{0, 1\}^n \rightarrow \mathbb{R}$ that satisfies ϵ_i -differential privacy with respect to each individual i , and let $T \subset [n]$ denote a set of indices. Consider two databases $D, D^T \in \{0, 1\}^n$ at Hamming distance $|T|$ that differ exactly on the indices in T . Then:

$$\frac{\Pr[A(D) \in S]}{\Pr[A(D^T) \in S]} \leq e^{\sum_{i \in T} \epsilon_i}$$

A useful primitive for differential privacy is the Laplacian distribution:

Definition 2.3. Denote by $\text{Lap}(\sigma)$ the symmetric Laplacian distribution with mean 0 and scaling σ . This distribution has probability density function:

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

We will sometimes abuse notation and write $\text{Lap}(\sigma)$ to denote the realization of a random variable drawn from the Laplacian distribution with parameter σ .

2.2 Mechanism Design

Every individual has some (unknown to the mechanism) cost function $c_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, where $c_i(\epsilon)$ represents player i 's cost for having his bit b_i used in an ϵ -differentially private manner. Because we consider small values of ϵ for which $\exp(\epsilon) \approx (1 + \epsilon)$, it will be most natural to consider *linear* cost functions for which $c_i(\epsilon) = v_i \epsilon$ for some unknown $v_i \in \mathbb{R}_+$. For clarity, we will assume throughout that individuals have linear cost functions, but our results will hold for any single-parameter family of cost functions that admit a total ordering independent of ϵ . That is, the property that our results will require is that for any $i \neq j$, and for any $\epsilon, \epsilon' \in \mathbb{R}_+$, it should hold that $c_i(\epsilon) = c_i(\epsilon, v_i) \leq c_j(\epsilon, v_j)$ if and only if $v_i \leq v_j$. Linear cost functions of course obey this property, but so do many other natural choices, such as exponential cost functions of the form $c_i(\epsilon) = \exp(\epsilon v_i)$.

A mechanism $M : \mathbb{R}_+^n \times \{0, 1\}^n \rightarrow \mathbb{R} \times \mathbb{R}_+^n$ takes as input a vector of cost functions $v = (v_1, \dots, v_n) \in \mathbb{R}_+^n$ and a database $D \in \{0, 1\}^n$, and outputs the evaluation of some algorithm $A(D)$ that is $\epsilon_i(v)$ -differentially private with respect to D to the data analyst, as well as a vector of payments $p(v) \in \mathbb{R}_+^n$ to each individual in D . For any $v'_i \in \mathbb{R}_+$ we let (v_{-i}, v'_i) denote the vector that results from changing entry v_i in v to v'_i . A player i derives utility $u_i = p_i(v) - v_i \epsilon_i(v)$ from such an outcome. Since any individual may opt against participating in our mechanism, we require first that our mechanisms be *individually rational*:

Definition 2.4. A mechanism $M : \mathbb{R}_+^n \times \{0, 1\}^n \rightarrow \mathbb{R} \times \mathbb{R}_+^n$ is *individually rational* if for all $v \in \mathbb{R}_+^n$:

$$p_i(v) \geq v_i \epsilon_i(v)$$

That is, each player must be guaranteed non-negative utility by participating and truthfully reporting his value to the mechanism.

Since individuals may misreport their costs so as to maximize their gain, we also require our mechanisms to be *truthful*:

Definition 2.5. A mechanism $M : \mathbb{R}_+^n \times \{0, 1\}^n \rightarrow \mathbb{R} \times \mathbb{R}_+^n$ is *dominant-strategy truthful* if for all $v \in \mathbb{R}_+^n$, for all $i \in [n]$, and for all $v'_i \in \mathbb{R}_+$:

$$p_i(v) - v_i \epsilon_i(v) \geq p_i(v_{-i}, v'_i) - v_i \epsilon_i(v_{-i}, v'_i),$$

that is, no player can ever increase his utility by misreporting his value for privacy.

The mechanism is run on behalf of some data analyst, who wishes to know an estimate of the statistic $s \equiv \sum_{i=1}^n b_i$. The mechanism outputs some randomized estimate of this quantity $\hat{s} = A(D)$, and the analyst prefers more accurate answers.

Definition 2.6. A mechanism M satisfies k -accuracy if for any $D \in \{0, 1\}^n$, it outputs an estimate $\hat{s} = A(D)$ such that:

$$\Pr[|\hat{s} - s| \geq k] \leq \frac{2}{3}$$

where the probability is taken over the internal coins of the mechanism.

The constant $2/3$ is of course inconsequential, and can be boosted to any desired value by repeated application of the mechanism.

We may consider two dual objectives for our mechanism. Our data analyst may have a fixed goal of k -accuracy for some k in which case we want to design mechanisms which deliver k -accurate estimates of s so as to minimize the sum of the payments. Alternately, our data analyst may have a fixed budget $B \in \mathbb{R}_+$ (say an NSF grant that can be used for data procurement). In this case, our goal is to design a mechanism which is k -accurate for the smallest possible value of k , while under the constraint that the sum of the payments never exceeds B .

3 Characterizing Accurate Mechanisms

In this section, we show necessary and sufficient conditions on the amount of privacy that a mechanism must purchase from each player in order to guarantee a fixed level of accuracy—to obtain a given level of accuracy, we show that a mechanism must purchase *at least* ϵ -privacy, from *at least* $|H|$ people, where the values of ϵ and $|H|$ depend on the desired accuracy. We emphasize that these necessary conditions are independent of any truthfulness requirements on the mechanism, and arise purely because of the need to achieve accuracy. This greatly simplifies the mechanism-design process for auctions for private data, because it allows us to without loss of generality restrict our attention to multi-unit procurement auctions.

Theorem 3.1. *Let $0 < \alpha < 1$. Any differentially private mechanism that is $\alpha \cdot n/4$ -accurate must select a set of users $H \subseteq [n]$ such that:*

1. $\epsilon_i \geq \frac{1}{\alpha n}$ for all $i \in H$.
2. $|H| \geq (1 - \alpha)n$.

Proof. Let M be a mechanism that is $\alpha \cdot n/4$ -accurate, and let $H \subset [n]$ be the set of individuals i such that $\epsilon_i \geq 1/\alpha n$. For point of contradiction, suppose that $|H| < (1 - \alpha)n$. Let $\bar{H} = [n] \setminus H$. We have that $|\bar{H}| > \alpha n$. Let $S = \{x \in \mathbb{R} : |x - s| < \frac{\alpha n}{4}\}$, where $s = \sum_{i=1}^n b_i$. By the accuracy of the mechanism, we have that the estimate \hat{s} output by the mechanism $M(v, D)$ satisfies:

$$\Pr[\hat{s} \in S] \geq \frac{2}{3}$$

Let $\bar{H}^1 = \{i \in \bar{H} : b_i = 1\}$ and let $\bar{H}^0 = \{i \in \bar{H} : b_i = 0\}$. Since \bar{H}^0 and \bar{H}^1 form a partition of \bar{H} , it must be that $\max(|\bar{H}^0|, |\bar{H}^1|) > \alpha n/2$. Without loss of generality, assume that $|\bar{H}^0| > \alpha n/2$ (the other case is identical). Let $T \subset \bar{H}^0$ such that $|T| = \alpha n/2$. Let D' be the database that results in setting each bit $b'_i = b_i$ if $i \notin T$, and $b'_i = 1$ otherwise. Note that D' and D have hamming distance $|T| = \alpha n/2$, and differ exactly on the indices of

T. Let \hat{s}' be the estimate generated by $M(v, D')$. By differential privacy together with Fact 1 we have:

$$\begin{aligned} \Pr[\hat{s}' \in S] &\geq \exp\left(-\sum_{i \in T} \epsilon_i\right) \cdot \Pr[\hat{s} \in S] \\ &\geq \exp\left(-\frac{\alpha n}{2} \cdot \frac{1}{\alpha n}\right) \cdot \frac{2}{3} \\ &= \frac{2}{3\sqrt{e}} \\ &> \frac{1}{3} \end{aligned}$$

Let $s' = \sum_{i=1}^n b'_i$. Note that $s' = s + \alpha n/2$. If $\hat{s}' \in S$, then by definition: $|\hat{s}' - s| < \alpha n/4$. By the triangle inequality, we must therefore have that: $|\hat{s}' - s'| > \alpha n/4$, and so it must have been that M was not $\alpha \cdot n/4$ accurate. \square

Corollary 3.2. Any αn -accurate individually rational mechanism must pay out a total payment of at least:

$$\sum_{i=1}^n p_i \geq \sum_{i=1}^{(1-4\alpha)n} c_i \left(\frac{1}{4\alpha n}\right)$$

where bidders are ordered such that $c_1(\cdot) \leq c_2(\cdot) \leq \dots \leq c_n(\cdot)$.

We remark that this corollary assumes only individual rationality, and is in general achievable only by an omniscient mechanism that knows all players' cost functions. No *truthful* αn -accurate mechanism is able to pay as little as this benchmark in general.

Theorem 3.1 gave necessary conditions on the privacy costs of an accurate mechanism. Next, we show that up to small constant factors, they are also *sufficient* conditions for an accurate mechanism:

Theorem 3.3. Let $0 < \alpha < 1$. There exists a differentially private mechanism that is $(\frac{1}{2} + \ln 3)\alpha \cdot n$ -accurate and selects a set of individuals $H \subseteq [n]$ such that:

1. $\epsilon_i = \begin{cases} \frac{1}{\alpha n}, & \text{for } i \in H; \\ 0, & \text{for } i \notin H. \end{cases}$
2. $|H| = (1 - \alpha)n$.

Proof. Let $H \subset [n]$ be any collection of individuals of size $|H| = (1 - \alpha)n$, selected independently of their private bits b_i , and let $t = \sum_{i \in H} b_i + \alpha n/2$. Observe that for any database D , $|t - s| \leq \alpha n/2$. Consider the mechanism that outputs $\hat{s} = t + \text{Lap}(\alpha n)$. First, we claim that this mechanism is $(1/2 + \ln 3)\alpha n$ -accurate. This follows by the triangle inequality conditioned on the event that $\text{Lap}(\alpha n) \leq (\ln 3)\alpha n$. It remains to verify that this holds with probability at least $2/3$. This is in fact the case:

$$\Pr[|\text{Lap}(\alpha n)| \geq (\ln 3)\alpha n] = \frac{1}{2\alpha n} \left(\int_{-\infty}^{-(\ln 3)\alpha n} \exp\left(-\frac{|x|}{\alpha n}\right) dx + \int_{(\ln 3)\alpha n}^{\infty} \exp\left(-\frac{|x|}{\alpha n}\right) dx \right) = \frac{1}{3}$$

We now verify the differential privacy guarantee, which follows from the analysis given in [DMNS06] of the Laplace mechanism. Let \hat{s} be the estimate calculated on database D (via sum t) and let \hat{s}' be the estimate calculated on neighboring database $D^{(i)}$ (via sum t'). Clearly, for any $i \notin H$ and for any $S \subset \mathbb{R}$, $\Pr[\hat{s} \in S] = \Pr[\hat{s}' \in S]$ and so $\epsilon_i = 0$. Now consider some $i \in H$ and $S \subset \mathbb{R}$. For any $S \subset \mathbb{R}$ and $r \in \mathbb{R}$, let $S - r$ denote $\{x - r : x \in S\}$.

$$\begin{aligned} \Pr[\hat{s} \in S] &= \Pr[\text{Lap}(\alpha n) \in S - t] \\ &= \int_{x \in S - t} \frac{1}{2\alpha n} \exp\left(-\frac{|x|}{\alpha n}\right) dx \\ &\leq \exp\left(\frac{1}{\alpha n}\right) \cdot \int_{x \in S - t'} \frac{1}{2\alpha n} \exp\left(-\frac{|x|}{\alpha n}\right) dx \\ &= \exp\left(\frac{1}{\alpha n}\right) \cdot \Pr[\hat{s}' \in S] \end{aligned}$$

where the inequality follows from the fact that $|t - t'| \leq 1$. \square

Theorems 3.3 and 3.1 taken together have the effect of greatly simplifying the space of possible mechanisms for private data that we need to consider. They imply that without loss of generality (up to small constant factors in their error term), when searching for αn -accurate mechanisms, we may restrict our attention to a special class of multi-unit procurement auctions, where we seek to purchase exactly $1/\alpha n$ units of some good (in this case, differential privacy) from exactly $(1 - \alpha)n$ individuals. Once we do this, we have purchased a sufficient quantity of privacy to run the Laplace mechanism employed in Theorem 3.3, which guarantees the desired accuracy! In the next section, we consider such mechanisms.

4 Deriving Truthful Mechanisms

4.1 Maximizing Accuracy Subject to a Budget Constraint

In this section, following the characterization of accurate mechanisms in Section 3, we restrict our attention to algorithms that guarantee $O(\alpha n)$ -accuracy by purchasing $1/\alpha n$ units of privacy from exactly $(1 - \alpha)n$ individuals. We consider the problem of obtaining an estimate \hat{s} of maximum accuracy, subject to a hard budget constraint²: $\sum_{i=1}^n p_i \leq B$. This is a natural objective, for example, in the case of a data analyst who has B dollars of grant money with which to buy data for a study, and wishes to buy the most accurate data that he can afford. We give a truthful and individually rational mechanism for this problem, and show that it is instance-by-instance optimal among the class of envy-free mechanisms. For clarity, we assume in this section that all cost functions are linear, i.e., $c_i(\epsilon) = v_i \epsilon$.

FairQuery(v, D, B) :

Sort v such that $v_1 \leq v_2 \leq \dots \leq v_n$.

Let k be the largest integer such that $\frac{v_k}{n-k} \leq \frac{B}{k}$.

Output $\hat{s} = \sum_{i=1}^k b_i + \frac{n-k}{2} + \text{Lap}(n-k)$

Pay each $i > k$ $p_i = 0$ and each $i \leq k$ $p_i = \min(\frac{B}{k}, \frac{v_{k+1}}{n-k})$.

We first prove that FairQuery is truthful and individually rational.

Theorem 4.1. *FairQuery is truthful and individually rational, and never exceeds the data analyst's budget B .*

Proof. First note that by the analysis from Theorem 3.3, for any $i \leq k$, $\epsilon_i = \frac{1}{n-k}$, and for any $i > k$, $\epsilon_i = 0$. For $i > k$ therefore, $p_i = v_i \cdot 0 = 0$. For $i \leq k$, $p_i = \min(\frac{B}{k}, \frac{v_{k+1}}{n-k}) \geq v_i/(n-k)$ because $v_i/(n-k) \leq B/k$ by construction and $v_i \leq v_{k+1}$ by definition. Hence, individual rationality is satisfied. Note also that $\sum_{i=1}^n p_i = k \cdot \min(\frac{B}{k}, \frac{v_{k+1}}{n-k}) \leq B$, and so the budget constraint is also satisfied. It remains to verify truthfulness:

Fix any v, i, v'_i and consider $k = k(v)$, $k' = k(v_{-i}, v'_i)$, $p_i = p_i(v)$, $p'_i = p'_i(v_{-i}, v'_i)$, $\epsilon_i = \epsilon_i(v)$, and $\epsilon'_i = \epsilon'_i(v_{-i}, v'_i)$. There are four cases:

1. Case 1: $v'_i < v_i$ and $p_i > 0$. In this case, v'_i moves earlier in the ordering and $\epsilon_i = \epsilon'_i$, and $p_i = p'_i$.
2. Case 2: $v'_i > v_i$ and $p_i = 0$. In this case, v'_i moves later in the ordering and $\epsilon_i = \epsilon'_i = p_i = p'_i = 0$.
3. Case 3: $v'_i < v_i$ and $p_i = 0$. In this case, v'_i moves earlier in the ordering, but if $p'_i > 0$ then by construction $p'_i = \min(\frac{B}{k'}, \frac{v_{k'+1}}{n-k'}) \leq v_i/(n-k')$. This follows because k' is such that $v_{k'+1} \leq v_i$ for all $i > k$ such that $p'_i > 0$.
4. Case 4: $v'_i > v_i$ and $p_i > 0$. In this case, v'_i moves later in the ordering, and either $p'_i = p_i$ and $\epsilon'_i = \epsilon_i$, or $p'_i = 0$ and $\epsilon_i = 0$. In the second case, by individual rationality, $p_i - v_i \epsilon_i \geq 0 = p'_i - v_i \epsilon'_i$.

Thus in all four cases, deviations are not beneficial, and the mechanism is truthful. \square

The next natural question to ask is: does FairQuery guarantee the data analyst a good level of accuracy, given his budget? As is always the case in prior-free mechanism design, it is important to specify what our benchmark is – good compared to what? Because mechanisms of the kind that we are considering always buy the same amount of privacy from an individual from whom they buy any privacy at all, a natural benchmark to consider is the set of all “envy-free” mechanisms which guarantee that no individual would prefer the outcome granted to any other.

²This question is related to the problem of designing budget feasible mechanisms in [Sin10, CGL11], but differs in that our privacy auction has *externalities*: a seller's cost for her good is a function of how many other sellers are chosen as winners by the mechanism.

Definition 4.2. A mechanism for private data is *envy-free* if for all possible valuation vectors v , and for all individuals i, j , $p_i - \epsilon_i v_i \geq p_j - \epsilon_j v_j$. That is, after the mechanism has determined the privacy costs and payments to each individual, there are no individuals who would prefer to have the payment and privacy cost granted to any other individual.

Observation 4.3. Any truthful envy-free mechanism which buys either no privacy or ϵ -privacy from each individual (i.e., if $\epsilon_i > 0, \epsilon_j > 0$ then $\epsilon_i = \epsilon_j$) must have the property that for all i, j with $\epsilon_i > \epsilon_j > 0$, $p_i = p_j$. That is, such mechanisms must pay each individual from whom privacy is purchased the same fixed price.

Note that by the characterization in Section 3, we may restrict ourselves to considering such mechanisms. Therefore we have:

Proposition 4.4. For any set of valuations $v \in \mathbb{R}_+^n$ (i.e., on an instance-by-instance basis) FairQuery achieves the optimal accuracy given budget B , among the set of all truthful, individually rational envy-free mechanisms.

Proof. First, observe that easy fact that FairQuery is indeed envy-free. We then merely observe that for any vector of valuations v , if FairQuery is $n - k$ -accurate, then by the definition of k , it must be that $\frac{v_{k+1}}{(n-k-1)} > \frac{B}{k+1}$, and so any mechanism that was $n - k'$ accurate for $k' > k$ must have $p_{k+1} > \frac{B}{(k+1)}$ by individual rationality. But by envy-freeness, it must have $p_i = p_{k+1} > \frac{B}{(k+1)}$ for all $i \leq k$. But in this case, we would have

$$\sum_{i=1}^n p_i \geq k' \cdot p_{k+1} > (k+1) \cdot \frac{B}{k+1} > B$$

which would violate the budget constraint. \square

4.2 Minimizing Payment Subject to an Accuracy Constraint

In this section, we consider mechanisms for the dual goal of truthfully obtaining a k -accurate estimate for some fixed accuracy constraint k while minimizing the payment required. Again, we restrict ourselves to the model of multi-unit procurement auctions justified in Section 3. In this setting, we show that the VCG mechanism is in fact optimal.

Recall that for a fixed accuracy goal αn , by Theorem 3.3, it is sufficient to buy $\frac{(1/2+\ln 3)}{\alpha n}$ units of privacy from $(1 - \frac{\alpha}{(1/2+\ln 3)})n$ people. We may therefore view our setting as a multi-unit procurement auction in which every individual is selling a single good ($\frac{(1/2+\ln 3)}{\alpha n}$ units of privacy), for which they have valuation $v_i = c_i(\frac{(1/2+\ln 3)}{\alpha n})$ (note that v_i is now the *total* cost for the $\frac{(1/2+\ln 3)}{\alpha n}$ units of privacy). The constraint on accuracy simply states that we must buy $(1 - \frac{\alpha}{(1/2+\ln 3)})n$ units of the good. In this case, we can analyze a simple application of the standard VCG mechanism:

MinCostAuction(v, D, α):

Let $\alpha' = \frac{\alpha}{1/2+\ln 3}$ and $k = \lceil (1 - \alpha')n \rceil$.

Sort $v = c_i(\frac{1}{n-k})$ such that $v_1 \leq v_2 \leq \dots \leq v_n$.

Output $\hat{s} = \sum_{i=1}^k b_i + \frac{n-k}{2} + \text{Lap}(\alpha'n)$

Pay each $i > k$ $p_i = 0$ and each $i \leq k$ $p_i = v_{k+1}$.

We first show that MinCostAuction does indeed satisfy the constraints of truthfulness and individual rationality, while obtaining sufficient accuracy.

Proposition 4.5. MinCostAuction is truthful, individually rational and αn -accurate.

Proof. That MinCostAuction is αn -accurate follows immediately from Theorem 3.3. Moreover, by Theorem 3.3, for each $i \leq k$, $\epsilon_i = 1/(\alpha'n)$ and for $i > k$, $\epsilon_i = 0$. Truthfulness and individual rationality then follow immediately from the fact that each $v_i = c_i(1/(\alpha'n))$ and MinCostAuction is an instantiation of the classical VCG mechanism. \square

MinCostAuction achieves its target utility at a cost of $\sum_{i=1}^n p_i = k \cdot v_{k+1}$. We now show that no other envy-free multi-unit procurement auction with the same accuracy guarantees (i.e. one that guarantees buying k units) makes smaller payments than MinCostAuction.

Theorem 4.6. *No truthful, individually rational, envy-free multi-unit procurement auction that guarantees purchasing k units can have total payment less than $k \cdot v_{k+1}$.*

Proof. For the sake of contradiction, suppose we have such a mechanism M . Fix some vector of valuations v that yields payments $p(v)$ such that $\sum_{i=1}^n p_i(v) < k \cdot v_{k+1}$ (again, note that v_i now denotes the total cost for purchasing data, not the per-unit privacy cost). First, if it is not already the case, we will construct a bid profile such that an item is purchased from some seller who is not among the k lowest sellers. It must be that there exists some i such that an item is purchased from i at a price of p_i , such that $v_i \leq p_i < v_{k+1}$ (otherwise $\sum_{i=1}^n p_i(v) \geq k \cdot v_{k+1}$). Let $v' = (v_{-i}, (v_i + v_{k+1})/2)$ be a bid profile in which bidder i raises his bid to be above p_i while remaining below v_{k+1} . Let $p' = p'(v')$ be the new payment vector. By individual rationality and truthfulness, it must be that in this new bid profile v' , player i is no longer allocated an item: by individual rationality, he would have to be paid $p'_i > p_i$ if he were allocated an item, but if his true valuation were v_i , then this would be a beneficial deviation, contradicting truthfulness. Because the mechanism is constrained to always buy at least k items, it must be that in v' , an item is now purchased from some seller j such that $j \geq k + 1$. By individual rationality, $p'_j \geq v_j \geq v_{k+1}$. But by envy-freeness, it must be that for every seller i from whom an item was purchased, $p'_i = p'_j \geq v_{k+1}$. Because at least k items are purchased, we therefore have $\sum_{i=1}^n p'_i \geq k \cdot v_{k+1}$, which contradicts the purported payment guarantee of mechanism M . \square

5 Preserving the Privacy of the Bid

In Section 4, we considered truthful, individually rational mechanisms that compensated users for the privacy loss due to the mechanisms' use of the individual's private bits b_i , but *not* due to the mechanisms' use of their valuations for privacy, v_i . Nevertheless, as we observed in the introduction, it is quite reasonable to assume that individual's valuations for privacy are correlated with their private bits. Can we design mechanisms that treat individuals' valuations for privacy as private data as well, and compensate individuals for the privacy loss due to the use of their valuations v_i ? In this section, we show that the answer is generically 'no' if we allow individuals to have arbitrarily high valuations for privacy. Moreover, we note that if we try to impose an a-priori bound on individual's valuations for privacy, then we re-introduce the same source of sampling bias that we had hoped to solve by running an auction.

A mechanism has two outputs: the estimate \hat{s} , and the payment P that the data analyst must make. Note that if the bids are private data as well, then a mechanism which is ϵ_i -differentially private with respect to bidder i must satisfy, for every set of estimate/payment tuples $S \subset \mathbb{R}_+^2$ and for each $(v, D) \in \mathbb{R}_+^n \times \{0, 1\}^n$, $\Pr[M(v, D) \in S] \leq \exp(\epsilon_i) \Pr[M(v^{(i)}, D^{(i)}) \in S]$, where $v^{(i)}$ and $D^{(i)}$ are arbitrary vectors that are identical to v and D everywhere except possibly on their i th index.

Theorem 5.1. *If bidder valuations for privacy may be arbitrarily large (i.e., $v \in \mathbb{R}_+^n$) then no individually rational mechanism M can protect the privacy of the bidder valuations and promise k -accuracy for any $k < n/2$ (i.e., any nontrivial value).*

Proof. Assume that M is k -accurate for some $k < n/2$. Run the mechanism $M(v, D)$ and obtain an estimate \hat{s} and privacy costs ϵ_i for each $i \in [n]$. Let $P = \sum_{i=1}^n p_i$ be the payment that the data analyst makes. By individual rationality, $P \geq \sum_{i=1}^n \epsilon_i v_i \geq \min_i v_i \cdot \sum_{i=1}^n \epsilon_i$. We trivially have that either $\Pr[\hat{s} \in [0, n/2]] \geq 1/2$ or $\Pr[\hat{s} \in [n/2, n]] \geq 1/2$. Without loss of generality, assume $\Pr[\hat{s} \in [0, n/2]] \geq 1/2$. Let $D' = 1^n$, and let \hat{s}' be the estimate obtained by running $M(v, D')$. By accuracy, we have that: $\Pr[\hat{s}' \in (n/2, n]] \geq \frac{2}{3}$. However, by differential privacy, together with Fact 1 we have:

$$\frac{2}{3} \leq \Pr[\hat{s}' \in (n/2, n]] \leq \exp\left(\sum_{i=1}^n \epsilon_i\right) \Pr[\hat{s} \in (n/2, n]] \leq \frac{\exp(\sum_{i=1}^n \epsilon_i)}{2}$$

Solving, we find that $\sum_{i=1}^n \epsilon_i \geq \ln(4/3)$, independent of v . We therefore have by individual rationality that $\Pr[P \in [0, \ln(4/3) \min_i v_i]] = 0$. By differential privacy, this must hold simultaneously for all inputs to the mechanism (v, D) : that is, such a mechanism can not charge a finite price P for *any* input, which completes the proof. \square

Remark 5.2. A natural (partial) way around the impossibility result of Theorem 5.1 is to restrict bidder valuations to lie in a bounded range (e.g. $[0, 1]$). This is unsatisfying, however, because it re-introduces the very source of sampling bias that we wanted to solve by running an auction. That is, bidders who happen to value their privacy at

a higher rate than allowed by the mechanism will simply not participate in the auction, which might systematically skew the resulting estimate in a way that we cannot measure.

6 Future Directions

The main contribution of this paper is to formalize the notion of auctions *for* private data, and to show that the design space of such auctions can without loss of generality be taken to be the simple setting of multi-unit procurement auctions. This initiates an intriguing new area of study that raises many questions. Among these are:

1. What is the proper benchmark for auctions in our setting? In this paper, we used the class of fixed-price (or envy free) mechanisms, which has become standard in the field of prior-free mechanism design [HR08, HK07]. However, this approach is better motivated in settings in which Bayesian optimal mechanisms are well understood and indeed charge fixed prices to winners. Bayesian optimal mechanisms are not known for our settings (e.g. budget constrained auctions with the objective function of buying as many units as possible). Studying Bayesian optimal mechanism design for these auctions, which correspond to natural markets for privacy would help identify and justify appropriate benchmarks.
2. We have shown that generically, no mechanism can compensate individuals for the loss of privacy which results from correlations between their private data and their reported costs for privacy. Nevertheless, such correlations exist! It is unsatisfying to restrict individual valuations for privacy to lie in a bounded range, because this reintroduces the very source of bias that we hoped to overcome by designing auctions. However, is there some restricted sense in which we can protect (and compensate users for) the privacy of their *valuations for privacy*? This requires the development of new models.
3. We have assumed throughout this paper that the private bits of the users, b_i are already known to some database administrator, such as a hospital. Although this is a natural assumption in some settings, what if it does not hold? Is there any way to mediate the purchase of private data directly from individuals who have the power to lie about their private data?
4. In this paper we considered an extremely simple market, in which there was a single data analyst wanting to buy data from a population. How about a two sided market, in which there are multiple data analysts, competing for access to the private data from multiple populations? Can we privately compute the market clearing prices for access to data in this way?
5. In this paper we considered a one-shot mechanism. In reality, the administrator of a private database will face multiple requests for access to his data as time goes on. How should the data analyst reason about these online requests and his value for the marginal privacy loss that he will incur after answering each request?

7 Acknowledgements

This paper has benefitted from discussions with many people. In particular, we would like to thank Frank McSherry for suggesting the problem of treating differential privacy as a currency to be bought and sold at the 2010 IPAM workshop on differential privacy, Tim Roughgarden for helpful early discussions on benchmarks for prior-free mechanisms for private data, and Sham Kakade, Ian Kash, Michael Kearns, Katrina Ligett, Malleesh Pai, Ariel Procaccia, David Parkes, Kunal Talwar, Salil Vadhan, Jon Ullman, and Bumin Yenmez for helpful comments and discussions.

References

- [BN10] H. Brenner and K. Nissim. Impossibility of Differentially Private Universally Optimal Mechanisms. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2010.
- [CGL11] N. Chen, N. Gravin, and P. Lu. On the Approximability of Budget Feasible Mechanisms. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2011.
- [Cli10] Stephanie Clifford. Web startups offer bargains for users data. *The New York Times*, May 2010. <http://www.nytimes.com/2010/05/31/business/media/31privacy.html>.
- [CP06] G. Calzolari and A. Pavan. On the optimality of privacy in sequential contracting. *Journal of Economic Theory*, 130(1):168–204, 2006.

- [DMNS06] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Third Theory of Cryptography Conference TCC*, volume 3876 of *Lecture Notes in Computer Science*, page 265. Springer, 2006.
- [Dwo08] C. Dwork. Differential privacy: A survey of results. In *Proceedings of Theory and Applications of Models of Computation, 5th International Conference, TAMC 2008*, volume 4978 of *Lecture Notes in Computer Science*, page 1. Springer, 2008.
- [FJS10] J. Feigenbaum, A.D. Jaggard, and M. Schapira. Approximate privacy: foundations and quantification (extended abstract). In *Proceedings of the 11th ACM conference on Electronic commerce*, pages 167–178. ACM, 2010.
- [GLM⁺10] A. Gupta, K. Ligett, F. McSherry, A. Roth, and K. Talwar. Differentially Private Approximation Algorithms. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, 2010.
- [GRS09] A. Ghosh, T. Roughgarden, and M. Sundararajan. Universally utility-maximizing privacy mechanisms. In *Proceedings of the 41st annual ACM symposium on Symposium on theory of computing*, pages 351–360. ACM New York, NY, USA, 2009.
- [HK07] J. Hartline and A. Karlin. Profit maximization in mechanism design. *Algorithmic Game Theory*, page 331, 2007.
- [HR08] J.D. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*. Citeseer, 2008.
- [Lau96] K. Laudon. Markets and privacy. *Communications of the ACM*, 39(9):-104, 1996.
- [Loh10] Steve Lohr. You want my personal data? reward me for it. *The New York Times*, July 2010. <http://www.nytimes.com/2010/07/18/business/18unboxed.html>.
- [MT07] F. McSherry and K. Talwar. Mechanism design via differential privacy. In *Proceedings of the 48th Annual Symposium on Foundations of Computer Science*, 2007.
- [NST10] K. Nissim, R. Smorodinsky, and M. Tennenholtz. Approximately Optimal Mechanism Design via Differential Privacy. *Arxiv preprint arXiv:1004.2888*, 2010.
- [Sin10] Y. Singer. Budget Feasible Mechanisms. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2010.
- [TCW10] C.R. Taylor, V. Conitzer, and L. Wagman. Online Privacy and Price Discrimination. *Economic Research*, 2010.

A A Utility Theoretic View of Valuing Differential Privacy

In this section, we provide a brief justification for why individuals should be able to quantify their cost for experiencing an ϵ -differentially private use of their private data. Say that \mathcal{A} denotes the set of all future events for which an individual i has preferences over outcomes, and $u_i : \mathcal{A} \rightarrow \mathbb{R}$ is a function mapping events to i 's utility for that event. Suppose that $D \in \mathcal{D}$ is a data-set containing individual i 's private data, and that $M : \mathcal{D} \rightarrow T$ is a mechanism operating on D promising ϵ_i -differential privacy to individual i . Let D' be a data-set that is identical to D except that it does not include the data of individual i (equivalently, it includes the data of individual i , but it is used in a 0-differentially private manner), and let $f : T \rightarrow \Delta\mathcal{A}$ be the (arbitrary) function that determines the distribution over all future events, conditioned on the output of mechanism M .

A basic consequence of differential privacy is the following:

Fact 2. *If $M : \mathcal{D} \rightarrow T$ is ϵ_i -differentially private with respect to individual i , and $f : T \rightarrow U$ is any arbitrary (randomized) function independent of D , then the composition $f \circ M : \mathcal{D} \rightarrow U$ is also ϵ_i -differentially private with respect to individual i .*

By the guarantee of differential privacy together with Fact 2, we have:

$$\begin{aligned}
\mathbb{E}_{x \sim f(M(D))}[u_i(x)] &= \sum_{x \in \mathcal{A}} u_i(x) \cdot \Pr_{f(M(D))}[x] \\
&\leq \sum_{x \in \mathcal{A}} u_i(x) \cdot \exp(\epsilon_i) \Pr_{f(M(D'))}[x] \\
&= \exp(\epsilon_i) \mathbb{E}_{x \sim f(M(D'))}[u_i(x)]
\end{aligned}$$

Similarly,

$$\mathbb{E}_{x \sim f(M(D))}[u_i(x)] \geq \exp(-\epsilon_i) \mathbb{E}_{x \sim f(M(D'))}[u_i(x)]$$

Therefore, when individual i is deciding whether or not to allow his data to be used in an ϵ_i -differentially private way, he is facing the decision about whether he would like his data to be used in such a way that could change his future utility by at most an additive factor of

$$\Delta u_i \equiv (\exp(\epsilon_i) - 1) \mathbb{E}_{x \sim f(M(D'))}[u_i(x)]$$

and so this is a natural quantity for i to value his privacy at. Note that for small values of ϵ_i , this is approximately $\epsilon_i \cdot \mathbb{E}_{x \sim f(M(D'))}[u_i(x)]$, which (by setting $v_i = \mathbb{E}_{x \sim f(M(D'))}[u_i(x)]$) conveniently leads to the form of linear utility functions that we explore in this paper.