Auctioning Data for Learning

Iordanis Koutsopoulos
Athens University of Economics and Business
Email:jordan@aueb.gr

Aristides Gionis
Aalto University
Email: aristides.gionis@aalto.fi

Maria Halkidi
University of Piraeus
Email: mhalk@unipi.gr

***Abstract*—In this paper we advocate that market mechanisms inspired by economics in conjunction with intelligent data selection is the key to fulfilling learning tasks in the presence of big data subject to privacy concerns of users. We design a market of private data that are gathered towards building a classifier. Each data owner has a private cost that quantifies his discomfort for providing data to the learner. Also, at each stage of the learning process, each data owner is characterized by a utility score, which expresses the utility of his data for the learner. The learner initiates a call for buying data, and interested owners respond by declaring their costs. The learner computes the utility score of each data owner. It then selects one owner to buy data from and associated payment. For the goals of minimizing the expected privacy cost of users and of minimizing the expected payment by the learner, we propose a variation of a Vickrey-Clarke-Groves (VCG) auction and an optimal auction respectively. For the case where the data arrive in a streaming fashion, we formulate the multi-round sequential decision version of the problem of learning the classifier. At each round the learner decides whether it will stop the learning process given the current classifier accuracy or perform one more auction. The problem amounts to weighing the current cost of classifier inaccuracy against the expected (privacy or reimbursement) costs incurred through the auction, plus the expected cost of the new classifier accuracy; we cast this problem as an optimal stopping one. The complexity of our framework scales only linearly with the size of the data set and fits in a broad range of private-data scenarios and data attributes.**

I. Introduction

Internet users now spend significant amounts of their time in social media and social networking activities; they also use more and more mobile phone applications. Hence, their digital traces are left online, with or without their consent or awareness. A number of organizations such as census bureaus and hospitals maintain large collections of personal data. Online advertising has also provided a big boost in the practice of collecting and using personal information.

The massive volumes of data, oftentimes in streaming fashion and from various sources, render data processing a daunting task. Hence, lightweight and scalable approaches are needed for *judiciously selecting* the data items to insert to the data processing engine. The end-goal is to select the appropriate subset of most informative data items, i.e. those that would contribute most to the learning process. Another distinct feature is that data is a valuable asset whose acquisition incurs a benefit to the entity that collects it but also a privacy loss to the one that provides it. In addition, data owners are rational entities that may strategize over the data they own in order to get the most out of it. Private data has recently become a commodity–it is gathered, bought and sold, thus giving rise to personal-information monetization [18]. The rationale lies in the fact that data owners should be given the option to control the conveyance of data and be compensated for that.

Our framework focuses on the case that private data are gathered and used towards a classification task. The mechanisms bring together ideas from the areas of auction mechanism design and machine learning, in particular active learning [16]. We consider a large set of data owners, each with a data item in their possession for sale, and a learner who is interested in buying data so as to feed them into the inference process. Each data owner has a private cost that quantifies the amount of discomfort for providing their data item. The cost is modeled as a random variable with a given probability distribution, whose realization is known only to that owner, but not to the other data owners nor to the learner. Additionally, at any point of the learning process, each data owner can be characterized by a *utility score*, which expresses the expected utility of that data owner for the learner, i.e. the expected added value of the data item to the task performed by the learner. This utility score is computable by the learner regardless of the volume of the data set, and may also be computable by other data owners, depending on the setup.

The learner initiates a call for buying data, and interested owners respond by declaring their costs. The learner computes the utility score of each data owner. It then selects an owner to buy data from and the associated payment. Our specific contributions are summarized as follows: (i)We put forward an approach for a private-data market with emphasis on classification tasks. Our framework scales only linearly with the size of the data set, (ii) For the goals of minimizing the expected privacy cost of users or minimizing the expected payment by the learner, we propose a variation of a Vickrey-Clarke-Groves (VCG) auction and an optimal auction respectively (see section V, VI), (iii) For the case where data arrive in a streaming fashion, we address the multi-round sequential decision version of the problem of learning the classifier (see section VII). At each round the learner decides whether it will stop the learning process given the current classifier accuracy or perform one more auction. The problem amounts to weighing the current cost of classifier inaccuracy against the expected (privacy or reimbursement) cost incurred through the auction, plus the expected cost of the new classifier accuracy; we cast this problem as an optimal stopping one.

II. Related work

A large body of published work exists in area of private-data markets, studying different aspects of the domain. Kleinberg et al. [9] use the multi-player game notions of *core* and *Shapley value* to devise pricing schemes for private information. In [8] the authors consider a data analyst who wishes to buy sensitive information in order to estimate a population

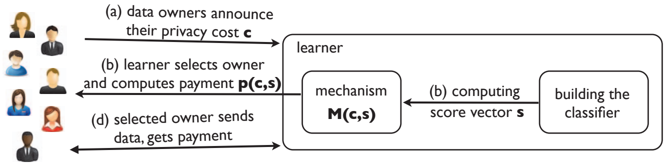
(a) data owners announce their privacy cost **c**
learner
(b) learner selects owner and computes payment **p(c,s)**
mechanism **M(c,s)**
(b) computing score vector **s**
building the classifier
(d) selected owner sends data, gets payment

Fig. 1. System overview: the data owners, the learner, and data owner selection process.

statistic. To quantify the notion of privacy preservation, they adopt the notion of differential privacy. Ideas from differential privacy and query pricing in data markets are also used in [12]. In [6], the authors design a truthful, individually rational, proportional-purchase mechanism that aims to maximize the estimation accuracy given that the analyst has a fixed budget. One issue with privacy cost modeling is that an individual's cost for privacy loss may be correlated with the private data. To deal with this issue, take-it-or-leave-it-offer mechanisms are proposed [7]. Riederer et al. [15] consider a setup where data aggregators compete to gain access to data through an auction that is run by a third party. In [4], the uncertainty about future consequences of current information disclosures is parallelized to financial options and borrows ideas from option pricing to account for valuation of privacy. Krause et al. [10] study the problem of attribute selection to maximize the utility of information gain, thus leading to efficient personalization while keeping the cost of user identifiability low. In [19] the impact of information disclosure on price setting is explored in a setup with a continuum of consumers with private preferences over a set of two goods, and two firms that produce the goods.

A different, more empirical, line of research studies how users valuate their own privacy. Two different valuation models have been considered. The first one measures the amount of money or benefit a user considers sufficient to give away their personal data e.g., see [17]. The second one measures the prices or intangible costs a user is willing to pay to protect their privacy, e.g. see [2]. Another interesting user-driven approach is taken in [5] towards quantifying the privacy cost of different types of online information for users by constructing simple auctions for different types of online user data. Aperjis and Huberman [3] also argue that often data owners cannot determine a valuation of their private data and propose to overcome this problem by compensating owners based on their choice, among different pricing schemes. The problem of online data procurement for learning task is introduced in [1]. The authors present a scheme for learning and pricing data online. In this work we present a game theoretic approach to deal with the problem of purchasing data for learning tasks. The learner aims at buying private information from data owners to increase the accuracy of learning.

III. System model

We consider a learner in direct interaction with a set $\mathcal{N}$ data owners. The learner may represent an application provider that has launched a specific application and is interested in gathering data so as to improve the quality of the application. Mechanisms are played in rounds. At each round, the learner initiates a call for data contributions. Owners respond with an expression of interest and by declaring the price for which they are willing to sell their data. The learner selects one user to buy data and associated payment. The system is depicted in Figure 1.

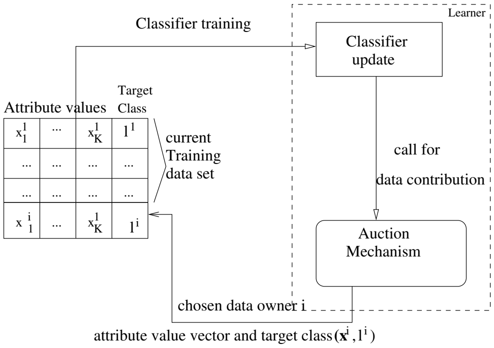
Learner
Classifier training
Classifier update
Attribute values
Target Class
x_1^1 ... x_K^1 l^1
... ...
... ...
x_1^i ... x_K^i l^i
current Training data set
call for data contribution
Auction Mechanism
chosen data owner i
attribute value vector and target class $(\mathbf{x}^i, l^i)$

Fig. 2. Owner selection for active learning.

A. Learner: The classification task

The learner performs a typical machine-learning task such as classification. We consider data with K attributes and a special one, the class label. Each attribute $j = 1, \dots, K$ takes values in a known set $\mathcal{A}_j$, which may be continuous- or discrete-valued, while the target class takes values in a known set $\mathcal{L}$. A data item i consists of a pair $(\mathbf{x}^i, \ell^i)$, where $\mathbf{x}^i = (x_1^i, \dots, x_K^i)$ is a vector of attribute values with respect to the K attributes, and $\ell^i \in \mathcal{L}$. A classifier is a mapping $f : \mathcal{A}_1 \times \cdots \times \mathcal{A}_K \to \mathcal{L}$ that takes an attribute value vector $\mathbf{x}^i$ and computes the predicted label $f(\mathbf{x}^i)$. The objective is to predict the true class ℓ^i, i.e., $f(\mathbf{x}^i) = \ell^i$.

We consider data owners for which part of their data is private (part of their attribute vector and/or the label). The learner aims at buying training data from data owners in order to learn the classifier. The objective of the learner is to select the subset of data items to use as training data. The limitation is that it cannot use as large a training data set as it wishes, since using data items incurs a cost, in particular, the privacy cost of a data owner who is selling the private information at hand. Thus the learner aims to minimize its cost while achieving a desirable level of accuracy for the classification task despite the large data volume. This situation is related to the problem of *active learning* [16], where a learning algorithm selects sequentially data items to use as training set, and selecting each data items incurs a cost. In active learning, many strategies have been proposed to guide the learning algorithm to select the next data item [16] (e.g., *uncertainty sampling*, *maximum expected change*, *least expected-error reduction*). Most typically, a score is assigned to each data item, and the learning algorithm selects the item with the best value on that score. Examples of selection strategies include *uncertainty sampling*, *maximum expected change*, *least expected-error reduction*, *least variance reduction*, and so on; for more details see the survey [16]. Any of these scores can be used as a *utility score* to assess the importance of buying a certain private data item from a data owner. The difference of our setting with active learning is that in addition to utility scores, data items have a cost that expresses the cost of privacy loss for data owners.

This cost is not uniform but is declared by each data owner. Furthermore, depending on the particular privacy settings, data owners may have access to utility scores of other data owners, and thus they can try to report their privacy costs in a way to maximize their expected profit. To address these issues, the learner selects the next data item for its training set via an auction mechanism, which takes account both the privacy costs and the utility scores of the data items. The details are discussed in Section III.

The learning process operates in rounds. At the beginning of round t, the learner has the training set from the previous round and it initiates a call for data contributions in order to select a new data item i and its label to insert into the training data set (Fig. 2).

B. Data owners: Privacy cost and utility score

Privacy cost. Part of the data item $(\mathbf{x}^i, \ell^i)$ of the data owner i is private and is brought for sale. Each data owner is characterized by their perceived privacy cost C_i for the data item in question. The cost quantifies the privacy discomfort of owner i if its data item is revealed to the learner. It may also be viewed as the minimum required compensation that owner i needs in order to reveal the data item. The cost may reflect the sensitivity of the owner about disclosing the private part of their data. It may depend only on the attribute set or on the attribute values as well. As an example for the former, revealing a "salary" attribute may be more sensitive than revealing a "residence location" attribute. As an example for the latter, declaring an attribute "drug user" as affirmative seems more privacy costly than declaring it as negative.

Cost C_i for owner i is a continuous random variable that takes values in interval $\mathcal{C}_i = [a_i, b_i]$. Costs C_i are independent from each other. We also assume that C_i is private information for each data owner i. That is, only i knows its realization c_i through evaluating the various factors that affect cost. The learner and data owners other than i have only partial information about C_i in the form of statistical knowledge. Let $f_i(\cdot)$ denote the probability density function (pdf) of C_i, and let $F_i(\cdot)$ be the corresponding cumulative density function (cdf). The learner and other data owners know only this pdf $f_i(\cdot)$. The pdf, and lower and upper limits a_i, b_i of its support set could emerge for example from the empirical distribution based on prior observations about the declared cost of the owner.

Utility score. Each data item i is characterized by a utility score s_i that denotes the importance of data point i for the learner. Utility scores depend on the exact learning task that is carried out, as well as on specific design choices. For learning the classifier, we employ active-learning strategies that quantify how important a data item is as the next training item. The importance of a data item is an increasing function of utility score.

IV. Definition and properties of the mechanism

Upon the call for data contribution, data owners report the privacy cost for their data items. The learner computes the utility score for each data item. Let $\mathbf{c} = (c_i : i \in \mathcal{N})$ be the vector of declared privacy costs and $\mathbf{s} = (s_i : i \in \mathcal{N})$ be the vector of utility scores of data items.

A mechanism $M(\mathbf{c}, \mathbf{s})$ maps a vector $(\mathbf{c}, \mathbf{s})$ to a data owner *selection vector* $\mathbf{x}(\mathbf{c}, \mathbf{s}) = (x_i(\mathbf{c}, \mathbf{s}) : i \in \mathcal{N})$, and a *compensation* vector $\mathbf{p}(\mathbf{c}, \mathbf{s}) = (p_i(\mathbf{c}, \mathbf{s}) : i \in \mathcal{N})$. The selection variable $x_i(\mathbf{c}, \mathbf{s})$ denotes the probability that data owner i will be selected, and $p_i(\mathbf{c}, \mathbf{s})$ is the payment to selected data owner i. For each data owner i, we define $t_i(\mathbf{c}, \mathbf{s}) = p_i(\mathbf{c}, \mathbf{s}) x_i(\mathbf{c}, \mathbf{s})$ as the expected compensation to the owner i, with the understanding that $t_i(\mathbf{c}, \mathbf{s}) = p_i(\mathbf{c}, \mathbf{s})$ when $x_i(\mathbf{c}, \mathbf{s}) = 1$. The *net utility* of data owner i is [1]

$$u_i(\mathbf{c}, \mathbf{s}) = x_i(p_i - c_i) = t_i - c_i x_i \,. \tag{1}$$

A. Bayesian game

Once the mechanism is announced to users, a Bayesian game is played. Recall that each data owner i knows only his own actual privacy cost c_i and has only statistical knowledge about costs of others. The latter are collectively denoted as $\mathbf{c}_{-i} = (c_j : j \in \mathcal{N}, j \neq i)$. Each data owner i strategically tries to determine the privacy cost that he declares in order to maximize his expected utility $\mathbb{E}_{\mathbf{c}_{-i}}[u_i(\mathbf{c}, \mathbf{s})] = \mathbb{E}_{\mathbf{c}_{-i}}[t_i(\mathbf{c}, \mathbf{s}) - c_i x_i(\mathbf{c}, \mathbf{s})]$, where expectation is taken with respect to the privacy cost of other data owners. A declared cost vector $\mathbf{y}^* = (y_1^*, \ldots, y_N^*)$ is *Bayesian Nash equilibrium* if for each data owner $i \in \mathcal{N}$, it is $\mathbb{E}_{\mathbf{y}^*_{-i}}[u_i(y_i^*, \mathbf{y}^*_{-i}, \mathbf{s})] \geq \mathbb{E}_{\mathbf{y}^*_{-i}}[u_i(y_i, \mathbf{y}^*_{-i}, \mathbf{s})]$, for all $y_i \in \mathcal{C}_i$, with $y_i \neq y_i^*$. Namely, at a Bayesian Nash equilibrium, no data owner can benefit from uni-laterally changing its declared privacy cost, provided that others do not change theirs.

Incentive compatibility (IC). A mechanism is called incentive compatible (IC), if the strategy in which each data owner reports his *true* privacy cost is a Bayesian Nash equilibrium. Namely, for each $i \in \mathcal{N}$,

$$\mathbb{E}_{\mathbf{c}_{-i}}[t_i(\mathbf{c}, \mathbf{s}) - c_i x_i(\mathbf{c}, \mathbf{s})] \geq \mathbb{E}_{\mathbf{c}_{-i}}[t_i(c_i', \mathbf{c}_{-i}, \mathbf{s}) - c_i x_i(c_i', \mathbf{c}_{-i}, \mathbf{s})]$$

for all $c_i' \in \mathcal{C}_i$, $c_i' \neq c_i$, with $\mathbf{c} = (c_i, \mathbf{c}_{-i})$. That is, each data owner always has larger net utility if he truthfully reports his privacy cost c_i than if he misreports it to be $c_i' \neq c_i$, given that all other owners are truthful.

Individual rationality (IR). A mechanism is called individually rational(IR), if $\forall i \in \mathcal{N}$ and $c_i \in \mathcal{C}_i$, it is

$$\mathbb{E}_{\mathbf{c}_{-i}}[u_i(\mathbf{c}, \mathbf{s})] \geq 0, \;\; \text{i.e.,} \;\; \mathbb{E}_{\mathbf{c}_{-i}}[t_i(\mathbf{c}, \mathbf{s}) - c_i x_i(\mathbf{c}, \mathbf{s})] \geq 0.$$

Individual rationality implies that at equilibrium each data owner has net utility at least as much as the one obtained when it does not participate in the market, which is zero.

B. Problem statement

The first challenge to confront is user rationality and selfishness. Driven by their selfish behavior, data owners strategically try to misreport their true cost in an effort to attract higher payment and maximize their own net utility. The second obstacle is that the learner is *unaware* of actual privacy costs of data owners. Hence, we focus to the class of mechanisms for which the Bayesian Nash equilibrium has two desirable properties: (i) it consists of strategies where data owners honestly declare their true privacy cost, because they have no motivation to do otherwise, (ii) data owners are urged to

[1] For simplicity, we use the briefer notation x_i, p_i and t_i to refer to the variables $x_i(\mathbf{c}, \mathbf{s})$, $p_i(\mathbf{c}, \mathbf{s})$ and $t_i(\mathbf{c}, \mathbf{s})$ respectively.

participate in the market, i.e., their net utility at equilibrium is always greater than zero (the net utility under no participation). Third, the learner needs to consider the different utility scores of data items provided by owners, and to adhere to a given minimum specified expected utility score e for selected data items so as to ensure good performance in its tasks. This constraint can be expressed as: $\sum_{i=1}^{N} s_i x_i(\mathbf{c}, \mathbf{s}) \geq e$. Let $\mathcal{M}(\mathbf{c}, \mathbf{s})$ be the space of all mechanisms $M(\mathbf{c}, \mathbf{s})$ that satisfy the following properties:

- **P1**: Selection vector $\mathbf{x}(\mathbf{c}, \mathbf{s})$ satisfies $\sum_{i=1}^{N} s_i x_i(\mathbf{c}, \mathbf{s}) \geq e$.
- **P2**: Selection vector $\mathbf{x}(\mathbf{c}, \mathbf{s})$ satisfies: $\sum_{i=1}^{N} x_i(\mathbf{c}, \mathbf{s}) = 1$.
- **P3**: $M(\mathbf{c}, \mathbf{s})$ satisfies incentive compatibility (IC).
- **P4**: $M(\mathbf{c}, \mathbf{s})$ satisfies individual rationality (IR).

Property **P1** refers to a minimum expected utility score e that reflects on a minimum level of learning quality, while **P2** is a feasibility constraint, since x_i the probability of selection. Next, we drop the dependence of x_i and t_i on $\mathbf{c}, \mathbf{s}$ for notational simplicity.

V. VCG auctions for private data

We consider the case where the learner wishes to minimize the expected privacy cost of data owners while maintaining a given minimum expected utility score e. This objective can be viewed as one that optimizes social cost if that refers to total expected privacy cost.V This implies that among the set of eligible data items that guarantee a minimum utility score, the learner selects the one that is least invasive for owner privacy. The motivation for the learner to realize such an objective is that it indirectly alleviates privacy concerns of data owners, and thus it reduces their hesitance to participate in the market. Consider the following linear optimization problem:

$$\min_{\mathbf{x}} \sum_{i=1}^{N} c_i x_i, \tag{2}$$

subject to:

$$\sum_{i=1}^{N} s_i x_i \geq e, \quad \text{and} \quad \sum_{i=1}^{N} x_i = 1 \tag{3}$$

where $x_i \in [0,1]$ is the probability that data item i is selected. The following cases exist: (i) If $e < \min_i s_i$, the solution is trivial: $x_j = 1$, for $j = \arg\min_i c_i$, and $x_i = 0$, for $i \neq j$, $i = 1, \ldots, N$, (ii) If $e > \max_i s_i$, the problem is infeasible, and there is no solution that satisfies **P1**. (iii) If $\min_i s_i < e < \max_i s_i$, then the solution emerges from Problem (2).

Depending on the values of $\mathbf{c}$, $\mathbf{s}$ and e, the solution may be an integer one (in which one variable has solution value 1 and the others are 0), or a fractional one, (there exist more than one variables with a solution value in $(0,1)$). In the former case, the continuous solution of Problem (2) coincides with the solution of the 0–1 problem which is given as follows. We define the subset of items $\mathcal{N}(\mathbf{s}) = \{i \in \mathcal{N} : s_i \geq e\}$. We rank costs in increasing order, and we set $j = \arg\min_{i \in \mathcal{N}(\mathbf{s})} c_i$. The optimal solution is $x_j^* = 1$ and $x_i = 0$, for $i \neq j$, $i = 1, \ldots, N$. On the other hand, if the optimal solution is fractional, it is not applicable in our case, since one data item needs to be selected. In this case, we adopt the integral optimal solution that we just described. Overall, the following greedy selection rule selects the winner owner to be the one with minimum privacy cost out of the set of ones whose scores are sufficiently high:

$$x_i(\mathbf{c}, \mathbf{s}) = \begin{cases} 1, & \text{if } i = i^* = \arg\min_{j \in \mathcal{N}(\mathbf{s})} c_j \\ 0, & \text{else}\,. \end{cases} \tag{4}$$

The payment rule is as in VCG auction: the winner is compensated with an amount equal to the second smallest reported privacy cost out of those owners whose score exceeds e.

$$t_i(\mathbf{c}, \mathbf{s}) = \begin{cases} c_j, & \text{if } i = i^* \text{ and } j = \arg\min_{k \in \mathcal{N}(\mathbf{s}), k \neq i^*} c_k \\ 0, & \text{else}\,. \end{cases} \tag{5}$$

A different point of view is defined by the optimization problem,

$$\min_{\mathbf{x}} \sum_{i=1}^{N} \left(c_i x_i + \lambda (e - s_i x_i) \right) \tag{6}$$

subject to $\sum_{i=1}^{N} x_i = 1$, with $x_i \in [0,1]$, where λ is an a priori defined weight factor. Here, the first term denotes the concern about privacy cost minimization, whereas the second term reflects quality of the utility score compared to the minimum required one, e. This auction could also be considered as one that minimizes the social cost, if the latter is considered as the sum of privacy cost and the negative of the difference of the score from e.

The rule is to select the owner $i^* = \arg\min_i (c_i - \lambda s_i)$. Clearly, among data items with the same utility score, the rule selects the one with the smallest privacy cost for its owner. Also, among data items with the same privacy cost, the one with the highest score is selected. The payment rule is:

$$t_i(\mathbf{c}, \mathbf{s}) = \begin{cases} c_j, & \text{if } i = i^* \text{ and } j = \arg\min_{k \neq i^*} (c_k - \lambda s_k) \\ 0, & \text{else}\,. \end{cases} \tag{7}$$

VCG auction has the IC and IR properties [11].

VI. Optimal auction design

A second objective for the learner is to minimize the expected compensation to data owners, while adhering to the utility score constraint **P1**. This objective is a pragmatic one for the learner who aims at making the mechanism viable and sustainable, rather than being benign to the data owners with respect to their privacy cost. This is expressed as:

$$\min_{M(\mathbf{c}, \mathbf{s}) \in \mathcal{M}(\mathbf{c}, \mathbf{s})} \mathbb{E}_{\mathbf{c}} \{ \sum_{i=1}^{N} t_i(\mathbf{c}, \mathbf{s}) \}\,. \tag{8}$$

We rely on the framework of Myerson [13], [14, Chap.3], which we modify to adapt to a procurement auction.

A. Conditions for IC and IR.

For each owner i, let c_i and c_i' be its true and declared privacy cost. Define as $X_i(c_i')$ the probability that i is selected, if i declares cost c_i' while all other owners declare their true costs,

$$X_i(c_i') = \mathbb{E}_{\mathbf{c}_{-i}} [x_i(c_i', \mathbf{c}_{-i}, \mathbf{s})]\,. \tag{9}$$

Also, define as $T_i(c_i')$ the expected payment to i, if he declares privacy cost c_i' and others declare true costs, i.e., $T_i(c_i') = \mathbb{E}_{\mathbf{c}_{-i}}[t_i(c_i', \mathbf{c}_{-i}, \mathbf{s})]$. Finally, denote by $U_i(c_i')$ the expected net utility for owner i if he declares privacy cost c_i' instead of c_i. Clearly, $U_i(c_i') = T_i(c_i') - c_i X_i(c_i')$. Then, the condition for incentive compatibility expressed in terms of the quantities above, is

$$U_i(c_i) \geq U_i(c_i') \Leftrightarrow T_i(c_i) - c_i X_i(c_i) \geq T_i(c_i') - c_i X_i(c_i') \tag{10}$$

and for individual rationality, it is, $U_i(c_i) \geq 0 \Leftrightarrow T_i(c_i) - c_i X_i(c_i) \geq 0$. We can prove the following theorem.

Theorem 1: A mechanism $M(\mathbf{c}, \mathbf{s}) = (\mathbf{x}(\mathbf{c}, \mathbf{s}), \mathbf{t}(\mathbf{c}, \mathbf{s}))$ is IC and IR if and only if for all i, (a) $X_i(c_i')$ is non-increasing in c_i', and (b) the following holds:

$$T_i(c_i') = \Gamma_i + c_i' X_i(c_i') + \int_{c_i'}^{b_i} X_i(s)\, ds\,. \tag{11}$$

where $\Gamma_i = U_i(b_i) = T_i(b_i) - b_i X_i(b_i)$, and b_i is the upper limit of the support set of the cost pdf.

B. Optimal auction formulation.

The objective of the learner is written as:

$$\sum_{i=1}^{N} \mathbb{E}_{\mathbf{c}}[t_i(\mathbf{c}, \mathbf{s})] = \sum_{i=1}^{N} \mathbb{E}_{c_i}\{\mathbb{E}_{\mathbf{c}_{-i}}[t_i(c_i, \mathbf{c}_{-i}, \mathbf{s})]\} = \sum_{i=1}^{N} \mathbb{E}_{c_i}\{T_i(c_i)\} \tag{12}$$

To impose IC and IR, we substitute $T_i(c_i)$ from (11). For any IC and IR mechanism, we write each term in (12) as

$$\mathbb{E}_{c_i}\{T_i(c_i)\} = \Gamma_i + \int_{\mathcal{C}_i} c_i X_i(c_i) f_i(c_i)\, dc_i + \int_{\mathcal{C}_i} \int_{c_i}^{b_i} X_i(r)\, dr\, f_i(c_i)\, dc_i\,.$$

We consider the definition for $X_i(\cdot)$ from (9). Assuming $\mathcal{C} = \times_i \mathcal{C}_i$, $\mathcal{C}_{-i} = \times_{j \neq i} \mathcal{C}_j$, and $f(\mathbf{c}) = \prod_i f_i(c_i)$ due to cost independence, we can show that a mechanism $M(\mathbf{c}, \mathbf{s})$ with $\Gamma_i = 0$, that minimizes

$$\int_{\mathcal{C}} \sum_{i=1}^{N} \big[X_i(\mathbf{c})\big(c_i + \frac{F_i(c_i)}{f_i(c_i)}\big)\big] f(\mathbf{c})\, d\mathbf{c} = \sum_{i=1}^{N} \mathbb{E}_{\mathbf{c}}\big[X_i(\mathbf{c})\big(c_i + \frac{F_i(c_i)}{f_i(c_i)}\big)\big]\,, \tag{13}$$

and satisfies properties **P1–P4** solves optimally problem (8) of the learner and is IC and IR. In the sequel, we can substitute c_i' with c_i due to IC.

C. The mechanism and its interpretation.

Data owner selection: For reported privacy cost vector $\mathbf{c}$ and utility score vector $\mathbf{s}$, consider mechanism $M(\mathbf{c}, \mathbf{s}) = (\mathbf{x}(\mathbf{c}, \mathbf{s}), \mathbf{t}(\mathbf{c}, \mathbf{s}))$. Let the data owner selection vector $\mathbf{x}(\mathbf{c}, \mathbf{s}) = (x_i(\mathbf{c}, \mathbf{s}) : i = 1, \ldots, N)$ be the solution to problem

$$\mathbf{x}(\mathbf{c}, \mathbf{s}) = \arg\min_{\mathbf{x}} \sum_{i=1}^{N} x_i\big(c_i + \frac{F_i(c_i)}{f_i(c_i)}\big) \tag{14}$$

subject to: $\sum_{i=1}^{N} s_i x_i \geq e$, and $\sum_{i=1}^{N} x_i = 1$

The integer solution to this problem is:

$$x_i(\mathbf{c}, \mathbf{s}) = \begin{cases} 1, & \text{if } i = i^* = \arg\min_{j \in \mathcal{N}(\mathbf{s})} \omega_j(c_j) \\ 0, & \text{else}\,, \end{cases} \tag{15}$$

where $\mathcal{N}(\mathbf{s}) = \{i \in \mathcal{N} : s_i \geq e\}$ and $\omega_j(c_j) = c_j + \frac{F_j(c_j)}{f_j(c_j)}$. The owner selection rule says that: (*i*) the declared privacy costs should be low, (*ii*) the probability density to the declared cost value due to prior knowledge (i.e., the pdf value at that point) should be high, and (*iii*) the declared cost value should be close to the lower limit of cost values, so that the cumulative probability mass at values lower than the declared one is low (ideally the declared cost should be the lower limit of the cost interval).

Recall from Theorem 1 in order that the mechanism is IC and IR, $X_i(c_i)$ should be non-increasing. This is satisfied only if $\omega_i(c_i)$ is non-decreasing in c_i. Indeed, we can write $x_i(c_i, \mathbf{c}_{-i}) = 1$, if $\omega_i(c_i) = \min_{j \in \mathcal{N}(\mathbf{s})} \omega_j(c_j)$, and 0 otherwise. Therefore, if $\omega_i(c_i)$ is non-decreasing in c_i, then $x_i(c_i, \mathbf{c}_{-i})$ is non-increasing in c_i, and so is $X_i(c_i)$. The requirement that $\omega_i(c_i)$ is non-decreasing is met if the pdf $f_i(\cdot)$ is non-increasing. This is satisfied by a wide range of pdfs.

Data owner compensation: The expected payment that guarantees IC and IR is $T_i(c_i) = c_i X_i(c_i) + \int_{c_i}^{b_i} X_i(r)\, dr$ due to Theorem 1. The payment is

$$t_i(c_i, \mathbf{c}_{-i}, \mathbf{s}) = c_i x_i(c_i, \mathbf{c}_{-i}, \mathbf{s}) + \int_{c_i}^{b_i} x_i(r, \mathbf{c}_{-i}, \mathbf{s})\, dr. \tag{16}$$

If the learner knew privacy cost c_i of an owner i, it would have compensated him with an amount equal to c_i per unit of probability of being selected; hence the total cost for user i would be $c_i x_i(\mathbf{c}, \mathbf{s})$. However, due to incomplete cost information, the learner gives owners an extra compensation $\int_{c_i}^{b_i} x_i(r, \mathbf{c}_{-i}, \mathbf{s})\, dr$ in order to motivate them to reveal their true cost. The payment rule (16) can be made more intuitive if for each data owner $i \in \mathcal{N}(\mathbf{s})$ we define

$$z_i(\mathbf{c}_{-i}) = \sup\{c : \omega_i(c) \leq \min_{k \in \mathcal{N}(\mathbf{s}), k \neq i} \omega_k(c_k)\}\,. \tag{17}$$

Thus, $z_i(\mathbf{c}_{-i})$ is the maximum declared cost of owner i that can make him win the auction against declared costs of others. The selection rule is written as:

$$x_i(\mathbf{c}, \mathbf{s}) = \begin{cases} 1, & \text{if } i \in \mathcal{N}(\mathbf{s}) \text{ and } c_i \leq z_i(\mathbf{c}_{-i}) \\ 0, & \text{else}\,. \end{cases} \tag{18}$$

If $x_i(c_i, \mathbf{c}_{-i}, \mathbf{s}) = 1$, i.e., data owner i is selected, then it is $\int_{c_i}^{b_i} x_i(r, \mathbf{c}_{-i}, \mathbf{s})\, dr = z_i(\mathbf{c}_{-i}) - c_i$, and his compensation is

$$t_i(c_i, \mathbf{c}_{-i}, \mathbf{s}) = c_i \cdot 1 + z_i(\mathbf{c}_{-i}) - c_i = z_i(\mathbf{c}_{-i})\,. \tag{19}$$

Thus, the payment to the selected owner is the maximum value of his declared cost so as he wins the auction. In general, owner selection and payment rules in the optimal auction differ from those in VCG auction. However, if data owner cost probability distributions and their supports are the same, the rules of VCG and the optimal auction coincide.

VII. Streaming Data: The multiple-round problem

We study the case where data are available through their owners in a *streaming* fashion in consecutive time epochs or rounds. At each round, there exists a set of data owners and corresponding data items available for sale.

The multi-round sequential decision version of the problem of learning the classifier is as follows. Each data owner declares their privacy cost for his data item at that epoch. At each round, the learner needs to decide whether it will perform one more auction for buying a data item, or whether it will stop the learning process for the current classifier accuracy.

The problem amounts to reaching a satisfactory tradeoff between classifier accuracy and item acquisition cost. The learner weighs the current attained accuracy versus the expected incurred privacy cost plus expected accuracy in the future. The expectation is with respect to randomness in utility scores of data items, as explained in the sequel. If the decision is to continue with an auction, the learner selects an item and computes the payment to its owner. The item should be the one that minimizes expected privacy cost plus the expected cost due to classifier inaccuracy, as projected the future. This item is obtained and added to the classifier training set, the classifier is updated and its accuracy changes. Then the next round begins, with new data items and perhaps new data owners.

The procedure continues until a decision is taken to stop the process or a maximum number M of rounds is reached, where M is dictated by practical constraints, such as latency or computation ones. The objective of the learner is to minimize a linear combination of total expected cost for data item acquisition and of classifier accuracy when the learning process stops. The total expected cost may refer to either (i) total expected privacy cost incurred to data owners, or (ii) total expected compensation to data owners. We are interested in the class of policies that are IC and IR at each round. These properties hold if the item selection is of the form (4) for VCG and (15) for optimal auctions respectively, and if payments are of the form (5), (19). Note that threshold e plays a critical role at each round of item selection; as long as the item selection and payment rules abide to threshold rules above, the mechanism is IC and IR at each round. The presentation in the sequel focuses on VCG auctions. The case of optimal auctions is similar.

A. *Optimal stopping formulation.*

The learning process may stop at or before the M-th round. At each round k, $k = 0, 1, \ldots, M-1$, each data owner brings a data item for sale. Denote with $\mathcal{S}_k$ the set of items that participate in auction round k. Let y_k for $k = 0, 1, \ldots, M-1$ denote the system state after auction round $k-1$ (i.e. before auction round k). State $y_k = \mathcal{T}_{k-1}$, where $\mathcal{T}_{k-1}$ is the set of obtained data items by the learner at previous rounds up to round $k-1$; this set is the current training set for the classifier. Denote by $\alpha(\mathcal{T}_{k-1})$ the accuracy of the classifier after round $k-1$ with the corresponding training data set. We define an additional state T, the termination state. If $y_k = T$ for some $k \leq M-1$, it is meant that the learning process has stopped at or before round $k-1$. If $y_k \neq T$, the learning process has not stopped and the current training data set is $\mathcal{T}_{k-1}$. The state space is $\mathcal{Y} = \mathcal{P} \cup \{T\}$, where $\mathcal{P}$ is the power-set of the set of data items.

Let u_k denote the control before round k, $k = 0, 1, \ldots, M-1$. The control consists of two elements, u^0 (stop the learning process) and u^1 (continue the learning process). In the latter case, a data item must be selected out of set $\mathcal{S}_k$. Thus, the control space at round k is $\mathcal{S}_k \cup \{u^0\}$. The state evolution is as follows. If $y_k = T$, or if $y_k \neq T$ and $u_k = u^0$, then $y_{k+1} = T$. That is, the stopping action drives the system to termination state T. If $y_k \neq T$ and $u_k = u^0$, a cost $g(y_k) = 1 - \alpha(\mathcal{T}_{k-1})$ is incurred, equal to the classifier inaccuracy, and the system subsequently stays in state T at no cost. Otherwise, a new data item $i_k \in \mathcal{S}_k$ is obtained, and $y_{k+1} = \mathcal{T}_k$, with $\mathcal{T}_k = \mathcal{T}_{k-1} \cup \{i_k\}$. Note that $\alpha(\mathcal{T}_k) = s_{i_k}$, where s_{i_k} is the utility score of the new data item.

The objective is to minimize the following cost,

$$\mathbb{E}\big\{\lambda g_M(y_M) + \sum_{k=0}^{M-1} c(y_k, u_k, w_k) + \lambda g(y_k, u_k, w_k)\big\}\,, \quad (20)$$

where expectation is with respect to the randomness of data item utility scores; we denote this randomness at round k with term w_k. For $k = 0, 1, \ldots, M-1$, the term $c(y_k, u_k, w_k)$ denotes the privacy cost for data owners at auction round k, and it will be defined later. If $y_k = T$, the privacy cost is 0, otherwise it is equal to the privacy cost incurred to the owner from which the data was obtained. Further, $g(y_k, u_k, w_k)$ is the cost of classifier inaccuracy. This is equal to $1 - \alpha(\mathcal{T}_{k-1})$, if $y_k \neq T$ and $u_k = u^0$, otherwise it is 0. Thus the classifier inaccuracy cost is imposed only when the learning process is terminated.

In the definition above, $\lambda \geq 0$ is a weight factor, measured in appropriate units, which reflects relative significance of classifier inaccuracy and data owner privacy cost. Finally, $g(y_M)$ is the terminal cost which is equal to $1 - \alpha(\mathcal{T}_{M-1})$, if $y_M \neq T$, otherwise it is 0.

Let $J_k(y_k)$ denote the optimal expected cost-to-go at round k, $k = 0, 1, \ldots, M-1$. After auction round $M-1$, we have $J_M(y_M) = \lambda(1 - \alpha(y_M)) = \lambda(1 - \alpha(\mathcal{T}_{M-1}))\,.$

For $k = M-1$, namely after auction round $M-2$ and before round $M-1$, the optimal cost-to-go based on the dynamic programming equation is

$$\begin{aligned} J_{M-1}(y_{M-1}) = \min\big\{&\lambda(1-\alpha(y_{M-1})), \\ &\min_e \mathbb{E}\big[\min_{i \in \mathcal{S}_{M-1}: s_i \geq e} c_i + \lambda(1 - \alpha(y_M))\big]\big\} \end{aligned} \quad (21)$$

with $y_M = s_{i_{M-1}}$, namely the score of the selected data item. Before the last auction, $M-1$, the learner decides whether to stop with current classifier accuracy and incur cost $\lambda(1 - \alpha(y_{M-1}))$, or continue with the last auction. Intuitively, the learner chooses to stop if the current inaccuracy is less than the expected privacy cost due to the auction, plus the expected cost due to the new inaccuracy.

If the learner decides to employ the auction, it needs to select the data item to buy. Recall that item selection must be of threshold form (4) in order to guarantee IC and IR at that round. However, the learner can influence the selection of the item through threshold e. Specifically, it selects the threshold that minimizes expected privacy cost to data owners plus the new accuracy, $\lambda(1-\alpha(y_M))$ of the classifier. The following tradeoff exists. By increasing e, the minimum privacy cost is increased, since the set of items from which the min-cost one is selected becomes smaller. On the other hand, high e results in selection of a data item with large utility score, which in turn leads to larger accuracy at the next stage, and therefore to reduced cost due to inaccuracy. Without loss of generality, we set the threshold e to be one of the scores of candidate data items at that round.

For $k < M-1$, the optimal cost-to-go is written as

$$J_k(y_k) = \min\left\{\lambda(1-\alpha(y_k)), \min_e \mathbb{E}\Big[\min_{i\in\mathcal{S}_k: s_i\ge e} c_i + J_{k+1}(y_{k+1})\Big]\right\} \tag{22}$$

with $y_{k+1} = s_{i_k}$ the score of selected item at round k.

Expectation $\mathbb{E}[\cdot]$ is with respect to uncertainty in the utility score of data items, which in turn stems from uncertainty about hidden, private labels of data items. If there exist L possible values for the label, then the label may take each value with probability $1/L$ (assuming no prior knowledge). Expectation $\mathbb{E}[\cdot]$ is then computed with respect to the uniform distribution of all data item scores over the L possible values. Conditioned on the value of the label, the utility score can be computed. The above formulation is an optimal stopping problem.

VIII. Experimental evaluation

A. Data sets.

To evaluate the performance of our approach we use two sizable data sets in terms of number of entries and attributes from the UCI Repository. It should be clear by now that the approach is applicable to much larger data sets. The first data set is the “Adult” data set (also known as “Census Income” data set). The data items consist of 14 attributes and are classified into two classes based on income, i.e. income $> 50K$, $\le 50K$. The training data set contains 32561 instances and the testing data set 16281 instances. We used the naive Bayes learning algorithm for training the classifier. The accuracy of the classifier using the entire training data set is 0.7851. The second data set (further referred to as H.A.R) has been built from recordings of 30 volunteers performing activities of daily living while carrying a waist-mounted smartphone with embedded inertial sensors. Each person performed six activities. The data set has been randomly partitioned into two sets, where 70% of the volunteers were selected for generating the training data and 30% the test data. The observations in this data set consists of 561 attributes. For the purposes of our study we used only the instances with label WALKING, SITTING and LAYING. After the filtering, the test data set contains 1524 observations. For the H.A.R data set, we used decision tree as classifier. The classification accuracy using the entire training data set is 1.

B. Discussion on experimental results.

We consider reverse VCG auction and reverse optimal auction for data item selection. Experiments were carried out 100 times and we present the average value of the various evaluation metrics. We evaluate the classification accuracy at different rounds of the auction considering different values of accuracy threshold for the selection of data items. Also we present the cumulative payment at each round and required total budget.

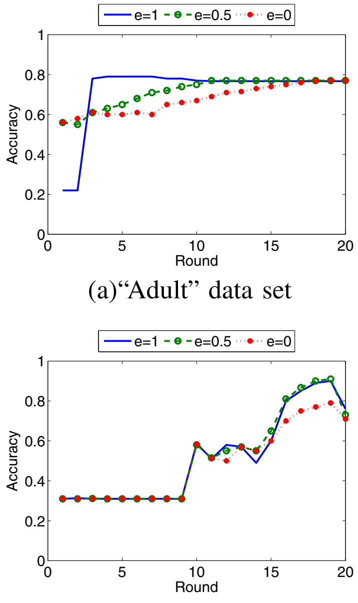

(a)“Adult” data set

(b) H.A.R. data set

Fig. 3. Accuracy of classifier at each round of reverse VCG auction.

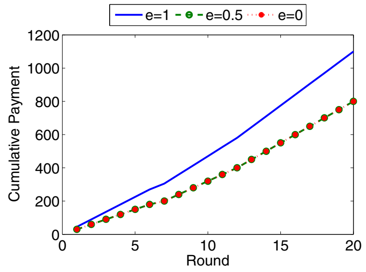

(a)“Adult” data set

Fig. 4. Cumulative payment at each round of reverse VCG auction for different threshold accuracy values

Selection with reverse VCG auction. Figure 3 shows the classification accuracy at different rounds of the VCG auction. The classification accuracy at the last round of the auction takes the same value independently of the value of threshold e for both data sets (“Adult”, H.A.R). This is because selected data items during the auction are the same, so the training data set is the same for each value of e. For the “Adult” data set, for a high value of e, the accuracy of the classifier reaches its highest value from the 3rd round. When the threshold is set to its low or medium value, the accuracy of the classifier reaches the highest value after the 10th or 16th round. For the H.A.R data set, which requires more data items for better accuracy due to the three class classification problem, the accuracy of classifier takes similar values for the first 13 rounds. In subsequent rounds, the accuracy increases for the majority of the auction rounds when we set e to its medium or high value. In both data sets the algorithm achieves the accuracy that we obtain when all the instances of the data set are used for

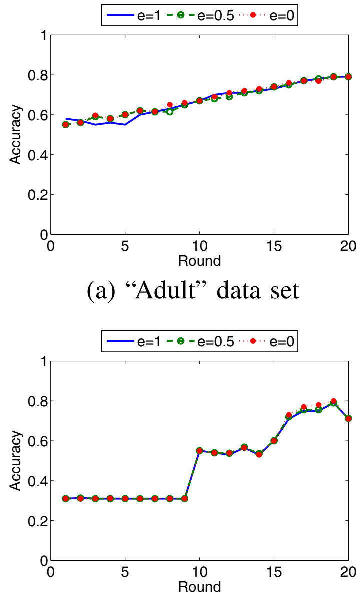

(a) “Adult” data set

(b) H.A.R. data set

Fig. 5. Accuracy of the classifier at each round of reverse optimal auction.

training. Further, when the threshold increases, the accuracy increases and reaches its highest value with few auctions.

Figure 4 shows that the cumulative payment for the “Adult” data set increases linearly. This implies that privacy costs in each round of the reverse VCG auction are of the same order of magnitude. Furthermore, we can infer that the higher the accuracy threshold of the learner, the more total budget is needed. The graph for the H.A.R data set is similar and it is omitted due to limited space. Based on these observations for both data sets, we conclude that for high values for the accuracy threshold, the learner selects data items that increase classification accuracy but they are expensive. Also, the higher the threshold, the more data items are filtered, and thus they are possibly more costly.

Selection with reverse optimal auction. The learner chooses the data item with minimum privacy cost after the filtering process. Figure 5 shows that the accuracy of the classifier takes similar values for the three threshold values used. The data selection has minor impact on the accuracy of the classifier. Figure 6 shows that the cumulative payment per round increases linearly for the “Adult” data set. A high threshold needs a total budget that is more than two times that of a learner with low threshold. The graph of cumulative payment for the H.A.R. data set is similar and thus is omitted.

IX. Conclusion

We proposed a framework for private data markets geared towards the task of learning a classifier. Our approach is scalable with the size of data in terms of data entries and attributes. Furthermore, the sequential decision-making version of the problem addresses the core of the problem when data arrive in streaming fashion. Again, the procedures behind decision-making at each round are of low complexity and particularly suitable for large data sets.

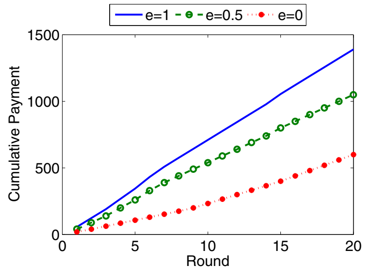

Fig. 6. “Adult” data set: The cumulative payment at each round of *reverse optimal auction*.

Acknowledgment

I. Koutsopoulos acknowledges the support of ERC08-RECITAL project, co-financed by Greece and the European Social Fund through the Operational Program Education and Lifelong Learning- NCRF 2007-2013

References

[1] J. Abernethy, Y. Chen, C. Ho and B. Waggoner’ “Actively Purchasing Data for Learning”, arXiv:1502.05774, 2015.

[2] A. Acquisti and J. Grossklags, “Privacy and rationality in decision making”, In *IEEE Security and Privacy*, 2005.

[3] C. Aperjis and B. A. Huberman, “A market for unbiased private data: Paying individuals according to their privacy attitudes”, *arXiv: 1205.0030*, 2012.

[4] S. Berthold and R. Bohme, “Valuating privacy with option pricing theory”, In *Proc. of Workshop on Economics of Inf. Security*, 2009.

[5] J. P. Carrascal, C. Riederer, V. Erramilli, M. Cherubini, and R. de Oliveira, “Your browsing behavior for a big mac: Economics of personal information online”, *arXiv: 1112.6098*, 2011.

[6] P. Dandekar, N. Fawaz, and S. Ioannidis, “Privacy Auctions for Inner Product Disclosures”, *arXiv:1111.2885*, 2011.

[7] L. Fleischer and Y.-H. Lyu, “Approximately optimal auctions for selling privacy when costs are correlated with data”, In *Proc. of ACM Conf. on Electronic Commerce (EC)*, 2012.

[8] A. Ghosh and A. Roth, “Selling Privacy at Auction”, In *Proc. of the Electronic Commerce Conf.*, 2011.

[9] J. Kleinberg, C. H. Papadimitriou, and P. Raghavan, “On the value of private information”, In *Proc. of the TARK Conf.*, 2001.

[10] A. Krause and E. Horvitz, “A utility theoretic approach to privacy and personalization”, *Proc. of Artificial Intelligence Conf.*, 2008.

[11] V. Krishna *Auction Theory*, Academic Press, 2002.

[12] C. Li, D. Y. Li, G. Miklau, and D. Suciu, “A theory of pricing private data”, *arXiv: 1208.5258*, 2012.

[13] R, Myerson, *Optimal auction design*, Mathematics of Operations Research, vol.6, pp.58-73, 1981.

[14] Y. Narahari, D. Garg, R. Narayanam and H. Prakash, *Game theoretic problems in Network Economics and Mechanism Design solution* Springer-Verlag, 2009.

[15] C. Riederer, V. Erramilli, A. Chaintreau, B. Krishnamurthy, and P. Rodriguez, “For sale: Your Data. By: You”, In *Proc. of Hotnets*, 2011.

[16] B. Settles. “Active learning literature survey”, Tech. report 1648, University of Winsconsin-Madison, 2010.

[17] S. Spiekermann, J. Grossklags, and B. Berendt, “E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior”, In *Proc. of the Electronic Commerce Conf.*, 2001.

[18] “Start-ups seek to help users put a price on their personal data”, NY Times, Feb. 2012.

[19] C. R. Taylor, “Consumer privacy and the market for customer information”, *RAND Journal of Economics*, 35(4):631–650, 2004.