Impossibility of Differentially Private Universally Optimal Mechanisms*

Hai Brenner[†] Kobbi Nissim[‡]

August 3, 2010

Abstract

The notion of *a universally utility-maximizing privacy mechanism* was recently introduced by Ghosh, Roughgarden, and Sundararajan [STOC 2009]. These are mechanisms that guarantee optimal utility to a large class of information consumers, *simultaneously*, while preserving *Differential Privacy* [Dwork, McSherry, Nissim, and Smith, TCC 2006]. Ghosh et al. have demonstrated, quite surprisingly, a case where such a universally-optimal differentially-private mechanisms exists, when the information consumers are Bayesian. This result was recently extended by Gupte and Sundararajan [PODS 2010] to risk-averse consumers.

Both positive results deal with mechanisms (approximately) computing a *single count query* (i.e., the number of individuals satisfying a specific property in a given population), and the starting point of our work is a trial at extending these results to similar settings, such as sum queries with non-binary individual values, histograms, and two (or more) count queries. We show, however, that universally-optimal mechanisms do not exist for all these queries, both for Bayesian and risk-averse consumers.

For the Bayesian case, we go further, and give a characterization of those functions that admit universally-optimal mechanisms, showing that a universally-optimal mechanism exists, essentially, only for a (single) count query. At the heart of our proof is a representation of a query function f by its *privacy constraint graph* G_f whose edges correspond to values resulting by applying f to neighboring databases.

[*] Research partly supported by the Israel Science Foundation (grant No. 860/06), and by the Frankel Center for Computer Science, Dept. of Computer Science, Ben-Gurion University.

[†] Dept. of Mathematics, Ben-Gurion University. `haib@bgu.ac.il`.

[‡] Microsoft AI, Israel, and Dept. of Computer Science, Ben-Gurion University. `kobbi@cs.bgu.ac.il`.

Contents

1 Introduction **1**
1.1 Our Results and Directions for Future Progress . . . 2
1.2 Related Work . . . 3

2 Preliminaries **3**
2.1 Differential Privacy [6] . . . 3
2.2 Oblivious Mechanisms . . . 4
2.2.1 Oblivious Differentially Private Mechanisms for a Count Query . . . 4
2.3 Utility Models . . . 5
2.4 Remapping and Generality . . . 6
2.5 Universal Mechanisms . . . 7

3 Impossibility of Universally Optimal Mechanisms for Generalizations of Count Queries **7**
3.1 The Basic Impossibility Result for Sum Queries . . . 8
3.2 Generalizing the Impossibility Result for Sum Queries . . . 10
3.2.1 Generalizing the Sum Query Impossibility Result to $m > 2$ . . . 10
3.2.2 Generalizing the Sum Query Impossibility Result to $n > 1$ . . . 10
3.3 Impossibility of Universally Optimal Mechanisms for Histogram Queries . . . 11
3.4 Impossibility of Universally Optimal Mechanisms for Bundles of Count Queries . . . 13

4 A Characterization of Universal Optimality in the Bayesian Setting **13**
4.1 The Basic Case: G_f is a Cycle . . . 15
4.2 Impossibility of Universal Optimality When G_f Contains a Cycle . . . 17
4.3 Impossibility of Universal Optimality When $\Delta(G_f) \geq 3$ . . . 19

1 Introduction

Differential Privacy [6] is a rigorous notion of privacy that allows learning global ('holistic') information about a collection of individuals while preserving each individual's information private. The literature of differential privacy is now rich in techniques for constructing differentially privacy mechanisms, including some generic techniques such as the addition of Laplace noise with magnitude calibrated to global sensitivity [6], addition of instance based noise calibrated to smooth sensitivity [13], and the exponential mechanism [12]. These and other techniques allow performing a wide scope of analyses in a differentially private manner, including conducting surveys over sensitive information, computing statistics, datamining, and sanitization. The reader is referred to [3] for a recent survey.

An immediate consequence of differential privacy is that (unless computing a constant function) a mechanism cannot compute a deterministic function. In other words, a differentially private version of an analysis would be a randomized approximation to the analysis, and furthermore, it would generally be possible to choose from a host of implementations for a task (e.g., the three generic techniques mentioned about may result with different mechanisms). Naturally, the designer of the analysis should choose one that is *useful*. Usefulness, however, depends on how the outcome of the analysis would be used, i.e., on the preferences of its *consumer*, that we henceforth refer to as an *information consumer*. Such a trade-off between uncertainty and utility, while taking consumer's preferences into account, is the subject of rational-choice theory and decision theory, as noted in [9, 10].

We discuss the two models of utility which were previously discussed in [9, 10]. In both, the information consumer has *side information* (her own world-view or previous knowledge), and a *loss-function* which quantifies the consumer's preferences and the quality of the solution for her problem. Intuitively, it describes how bad is a deviation from the exact answer for the consumer, a measure of her intolerance towards the inaccuracy imposed by differentially private mechanisms. Finally, the models assume that the consumers are *rational* - they combine the structure of the mechanism, their side information and their personal loss-function (preferences) with the goal of minimizing their loss, or, equivalently maximizing their utility. The two models differ in the way side information is formulated and respectively how utility function is defined. Subject to the requirements of differential privacy, one usually has a choice from a collection of implementations. As discussed in decision-theory and assuming rational information consumers, each consumer will choose a mechanism which maximizes her utility. This is an *optimal* mechanism for this consumer.

Information consumers' accuracy requirements vary: for some consumers only an exact answer would be of value, whereas others may aim at minimizing the estimate bias (ℓ_1 error), or its variance (ℓ_2 error), and, clearly, many other criteria exist. It seems that a discussion of the utility of differentially private mechanisms should take this rich variety into account. The recent work of Ghosh, Roughgarden, and Sundararajan [9] has put forward a serious attempt at doing exactly that with respect to (oblivious) Bayesian information consumers. In this utility model, the consumer's side information is described as an a priori distribution on the exact result of the analysis. The recent work of Gupte and Sundararajan [10] considers a related model where the information consumers are *risk-averse.* Here, the information consumer's knowledge is a set of possible values the exact analysis can take, and an optimal mechanism minimizes the consumer's worst-case expected loss.

Composition theorems for differential privacy only guarantee that the degradation in privacy is not more than exponential in the number invocations. Hence, while different consumers may exhibit different optimal mechanisms, a very important goal is to avoid invoking that multiplicity of mechanisms. This degradation is part of the motivation for the work on *sanitization* where a family of queries are answered at once [5, 1, 8, 7], the work on *privacy under continual observation* [4], and the construction of the *Median Mechanism* [14]. A surprising result of Ghosh, Roughdarden, and Sundararajan [9] is that invoking a multiplicity of optimal mechanisms may not be necessary. They consider a database that is a collection of Binary inputs (e.g., pertaining to having some disease) and Bayesian information consumers that wish to count the number of *one* entries in the database (equivalently, compute the sum of the entries). They show the existence of a single mechanism that enables optimality for *all* Bayesian information consumers (the mechanism needs to be invoked only once). The mechanism itself is not optimal for all Bayesian information consumers, however, each consumer can perform a deterministic remapping on the outcome of the common mechanism, where the remapping is chosen according to her notion of utility, and locally output a result that is effectively according to one of her optimal mechanism. Such a common mechanism is referred to as *universally optimal*. An analogous result for risk-averse information consumers was shown in [10].

Are these results of [9] and [10] that deal with the simple case of a single count query "accidental", or can they be extended to other queries? to multiple queries? One would anticipate that universally-optimal mechanisms should exist (at least) for those queries that are closely related to counting, such as sum queries where the inputs are non-binary, histograms, and bundles of two or more count queries.

1.1 Our Results and Directions for Future Progress

In contrast with the anticipation expressed in the previous paragraph, we show that settings in which universally optimal mechanisms exist are extremely rare, and, in particular, in both the setting of Bayesian and of risk-averse information consumers, universally optimal mechanisms do not exist even for sum queries where the inputs are non-binary, histograms, and bundles of two or more count queries.

Moreover, in the case of Bayesian information consumers, we give a characterization of those functions of the data that admit universally optimal mechanisms. The characterization makes use of a combinatorial structure of the query function $f : \mathcal{D}^n \to \mathcal{R}_f$, where $\mathcal{D}$ is the domain of the database records and $\mathcal{R}_f$ is the output space of the query function. We define this combinatorial structure of the query G_f and call it a *privacy constraint graph*. The vertices of G_f correspond to values in $\mathcal{R}_f$, and edges correspond to pairs of values resulting by applying f to neighboring databases. (This graph was examined in some proofs in [11] as well). We show:

Theorem 4.2 (Informal)**.** *If G_f contains a cycle then no universally optimal mechanism exists for f.*

Theorem 4.3 (Informal)**.** *If G_f is a tree that contains a vertex of degree 3 or more, then no universally optimal mechanism exists for f for better values of the privacy parameter.*

Facing the impossibility of universal optimality, an alternative may be found in an approximate notion, which enables (approximate) optimality to (approximately) all of the information consumers. A good notion of approximate optimality should allow constructing such mechanisms for sum queries, histograms, and more. Furthermore, it should allow performing several queries and satisfy a composition requirement, in a sense that when applying two such mechanisms to two different queries, the resulting composed mechanism should be somewhat approximately optimal for the two queries together.

Finally, we note that, following prior work we focus on *oblivious* mechanisms (see Section 2.2 for the technical definition). In Section 3, we show that for the intuitive generalizations of count queries, enabling *non*-oblivious universal mechanisms from which optimal oblivious mechanisms are derived, still leaves the construction of universally optimal mechanisms impossible. The question whether non-oblivious universally-optimal mechanisms exist for some other natural abstract queries, from which all oblivious universally-optimal mechanisms may be derived is left open.

1.2 Related Work

Most relevant to our work are the papers by Ghosh, Roughgarden, and Sundararajan [9] and by Gupte and Sundararajan [10]. Ghosh et al. show that the geometric mechanism (a discrete version of the Laplace mechanism of [6]) yields optimal utility for all Bayesian information consumers for a count query. Their proof begins by observing that all differentially private mechanisms correspond to the feasible region of a Linear Program (a polytope), and that minimizing disutility can be expressed as minimizing a linear functional. Hence, every Bayesian information consumer has an optimal mechanism corresponding to a vertex of the polytope, which in turn corresponds to a subset of the constraints of the Linear Program which are tight (optimal mechanisms, not corresponding to the polytope vertices, may also exist). They introduce a *constraint matrix* that uniquely corresponds to a vertex of the polytope, and indicates which constraints are tight, and which are slack on that vertex. Those constraint matrices that correspond to optimal mechanisms, are shown to have some special structure that allows to derive mechanisms with the same signature (and thus equal) from the geometric mechanism using some deterministic remapping on its output.

We are also interested in observing the tight constraints in some mechanisms. We will not need the full description of the structure of such a constraint matrix. Instead we only use the observation that tight privacy constraints can be derived only from mechanisms that also obey similar tight constraints.

Gupte and Sundararajan show similar results for the risk-averse utility model, where consumers try to minimizes their maximal worst-case disutility. They provide a full characterization of the mechanisms which are derivable (by random remapping) from the geometric mechanism and use this characterization to construct a universally-optimal mechanism for a count query. An interesting feature of the construction is that it releases noisy answers of the query at different privacy levels, thus keeping more privacy against specific consumers, and enabling more utility to others.

Also related to our work is the recent work of Kifer and Lin [11] that studies privacy and utility, in a very general setting, from an axiomatic point of view. They introduce a partial order on mechanism where mechanism Y is at least as *general* as mechanism X if X can be derived from Y by post processing. They also introduce the concept of maximal generality, which turns to be useful in our proofs.

2 Preliminaries

2.1 Differential Privacy [6]

Simply speaking, a mechanism which preserves differential-privacy will output for any two databases which 'look alike' the same result, with similar probabilities. More formally, consider databases $D_1, D_2 \in \mathcal{D}^n$ which consist of n records out of some domain $\mathcal{D}$. The Hamming Distance between

D_1 and D_2 is the number of records on which they differ. We will call databases at distance one *neighboring.*

Definition 2.1 (Differential Privacy [6])**.** Let $\mathcal{M} : \mathcal{D}^n \to \mathcal{R}$ be a probabilistic mechanism. $\mathcal{M}$ preserves α*-differential-privacy* for $\alpha \in (0, 1)$ if for any two neighboring databases $D_1, D_2 \in \mathcal{D}^n$ and any (measurable) subset of the mechanism's range $S \subseteq \mathcal{R}$,

$$Pr\left[\mathcal{M}(D_1) \in S\right] \geq \alpha \cdot Pr\left[\mathcal{M}(D_2) \in S\right]. \tag{1}$$

The probability is taken over the coin tosses of the mechanism $\mathcal{M}$.

Notice that the greater α is the less the mechanism's output depends on the exact query result, and so better privacy is attained.

2.2 Oblivious Mechanisms

We consider a setting where several information consumers are interested in estimating the value of some query $f(\cdot)$ applied to a database $D \in \mathcal{D}^n$, and answered by a differentially private mechanism $\mathcal{M}$. Ghosh et al. [9] show that if no restriction is put on the mechanism, then no universally optimal mechanism exists for count queries (intuitively, universal optimality, defined below, means that all potential consumers minimize their loss simultaneously). On the other hand, universally optimal mechanisms sometimes do exist if we restrict our mechanisms such that their output distribution depends only on the the exact query result (a.k.a. *oblivious mechanisms*). This is why in [9] (and later in [10]) only oblivious mechanisms are considered[1]. We follow suit and only consider oblivious mechanisms. We show in Subsection 3.2.1 that this restriction does not weaken the basic results presented in Section 3.

Definition 2.2 (Oblivious Mechanism)**.** Let $f : \mathcal{D}^n \to \mathcal{R}_f$ be a query. A mechanism $\mathcal{M} : \mathcal{D}^n \to \mathcal{R}$ is f*-oblivious* (or simply *oblivious*) if there exists a randomized function $\tilde{\mathcal{M}} : \mathcal{R}_f \to \mathcal{R}$ such that, for all $D \in \mathcal{D}^n$, the distributions induced by $\mathcal{M}(D)$ and $\tilde{\mathcal{M}}(f(D))$ are identical.

Combining α-differential privacy with obliviousness, we get that for every $i, i' \in \mathcal{R}_f$ which are outputs of neighboring databases D, D' (i.e., $f(D) = i$ and $f(D') = i'$), then $\Pr[\tilde{\mathcal{M}}(i) \in S] \geq \alpha \cdot \Pr[\tilde{\mathcal{M}}(i') \in S]$ for all $S \subseteq \mathcal{R}$.

2.2.1 Oblivious Differentially Private Mechanisms for a Count Query

An oblivious finite-range mechanism $\mathcal{M} : \mathcal{D}^n \to \mathcal{R}$ estimating $f : \mathcal{D}^n \to \mathcal{R}_f$ can be described by a row-stochastic matrix $X = (x_{i,j})$ of the underlying randomized mapping $\tilde{\mathcal{M}}$, whose rows are indexed by elements of $\mathcal{R}_f$, and whose columns are indexed by elements of $\mathcal{R}$, where $x_{i,j}$ equals the probability of outputting $j \in \mathcal{R}$ when $f(D) = i$. Since $\mathcal{R}$ is finite, and information consumers anyway remap the outcome of $\mathcal{M}$, we can assume, wlog, that $\mathcal{R} = \{0, 1, 2, \ldots, |\mathcal{R}| - 1\}$.

We now consider the case where $\mathcal{D} = \{0, 1\}$ and $f(D)$ counts the number of one entries in D. Hence, $\mathcal{R}_f = \{0, \ldots, n\}$ and the matrix X is of dimensions $(n+1) \times |\mathcal{R}|$. Preserving α-differential privacy poses constraints on the transition matrix X beyond row-stochasticity. Note that for the

[1] Impossibility of universal optimality when the mechanisms are not restricted to being oblivious is proved in [9] for Bayesian information consumers. For risk-averse consumers, [10] show that non-oblivious mechanisms may be replaced with oblivious ones without affecting the consumers' utility for the worse.

count query, the query results of two neighboring databases may differ by at most one. Differential privacy hence imposes the constrains $x_{i,j} \geq \alpha \cdot x_{i+1,j}$ and $x_{i+1,j} \geq \alpha \cdot x_{i,j}$ where $i \in \mathcal{R}_f = \{0 \ldots n-1\}$ and $j \in \mathcal{R}$. Adding row-stochasticity and differential privacy, we get that an oblivious differentially private mechanism for the count query should satisfy the following linear constraints:

$$x_{i,r} \geq \alpha x_{i+1,r} \qquad \forall i \in \{0, \ldots, n-1\}, \forall r \in \mathcal{R} \tag{2}$$

$$\alpha x_{i,r} \leq x_{i+1,r} \qquad \forall i \in \{0, \ldots, n-1\}, \forall r \in \mathcal{R} \tag{3}$$

$$\sum_{r \in \mathcal{R}} x_{i,r} = 1 \qquad \forall i \in \{0, \ldots, n\} \tag{4}$$

$$x_{i,r} \geq 0 \qquad \forall i \in \{0, \ldots, n\}, \forall r \in \mathcal{R} \tag{5}$$

2.3 Utility Models

We use the utility models defined in [9] and [10]. In both, a *loss function* $\ell(i, r)$ quantifies an information consumer's disutility when she chooses to use answer r while the correct answer is i. Given a loss function $\ell(\cdot, \cdot)$ of an information consumer, if the exact answer is i then her expected loss is $\sum_{r \in \mathcal{R}} x_{i,r} \cdot \ell(i, r)$.[2] Loss functions vary between consumers, and the only assumptions made in [9, 10] is that $\ell(i, r)$ depends on i and $|i - r|$ and is monotonically non-decreasing in $|i - r|$ for all i. (This is a reasonable requirement that turns to be crucial for the existence of a universally optimal mechanism [9].) Examples of loss functions include $\ell_1(i, r) = |i - r|$ (consumers who care to minimize expected mean error); $\ell_2(i, r) = (i - r)^2$ (minimize error variance); and $\ell_{bin}(i, r)$ that evaluates to 0 if $i = r$ and to 1 otherwise (minimize number of errors).

Information consumers differ in their knowledge about the exact $f(D)$. References [9] and [10] model this knowledge differently as we now describe.

Bayesian Model [9] In the Bayesian utility model, an information consumer's knowledge is represented by a vector $\bar{p}$ where p_i is the consumer's a priori probability that $f(D) = i$. Having a vector of prior probabilities $\bar{p}$ and loss function $\ell(\cdot, \cdot)$, the consumer's expected loss can be expressed as $\sum_i p_i \cdot \sum_r x_{i,r} \cdot \ell(i, r)$. The *optimal mechanisms* for this information consumer hence are the solutions of the linear program in the variables $x_{i,r}$ consisting the constraints in Equations (2)-(5) and the objective

$$\text{minimize} \sum_{i \in \mathcal{R}_f} p_i \cdot \sum_{r \in \mathcal{R}} x_{i,r} \cdot \ell(i, r). \tag{6}$$

Risk-Averse Model [10] In the risk-averse utility model an information consumer's knowledge restricts the possible values for the exact $f(D)$. This is expressed by a set $S \subseteq \mathcal{R}_f$ of the possible values $f(D)$ can take. The consumer is interested in minimizing her maximal expected loss conditioned on $f(D) \in S$, i.e., $\max_{i \in S} \sum_r x_{i,r} \cdot \ell(i, r)$. Similarly to the above, the optimal mechanism for an information consumer is a solution to a linear program consisting the constraints in Equations (2)-(5) and the objective

$$\text{minimize} \max_{i \in S} \sum_{r \in \mathcal{R}} x_{i,r} \cdot \ell(i, r). \tag{7}$$

[2] This is only true if the consumer uses the mechanism X *directly*, i.e., the consumer leaves the mechanism's output as is, and does not apply a post-processing step. The ability to apply such a post-processing step on the mechanism's output will be discussed in the next sub-section.

2.4 Remapping and Generality

An information consumer might have access to a private mechanism U which is not tailored specifically for her needs (i.e., to her prior knowledge and loss function). Yet, she may be able to recover a better mechanism for her needs by means of post-processing, which we will denote *remapping*. To intuit remapping, consider a consumer that knows that for the specific database the count query cannot yield the answer 0. If that consumer receives a 0, it may be beneficial for her to remap it to 1. (Recall that the loss function is monotone in $|i-r|$.) Denoting the given mechanism by U and the remapping by T (a row-stochastic linear transformation, T has no access to the information of the database other then the output of U), the actual mechanism that is used by the information consumer is denoted $T \circ U$ (in matrix form: UT).

Notice that given a mechanism U with a finite range, an information consumer can find the optimal remapping T for her (such that $T \circ U$ has optimal utility), by constructing a linear program in which $T = (t_{i,j})$ are the program variables [10].

Definition 2.3 (Derivable Mechanisms, Generality Partial Order [11])**.** Let X, Y be private mechanisms. We say that a mechanism X is *derivable* from a mechanism Y if there exists a random remapping T of the results of mechanism Y, such that $X = T \circ Y$. We also say that Y is *at least as general as* X, and denote this relation by $X \preceq_G Y$. If $X \preceq_G Y$ and $Y \preceq_G X$ we say that X, Y are *equivalent*.

Definition 2.4 (Maximal Generality [11])**.** Let X be an α-differentially private mechanism. X is *maximally general* if for every α-differentially private mechanism Y, if $X \preceq_G Y$ then $Y \preceq_G X$.

After introducing the notion of maximally general mechanisms (for any definition of privacy), Kifer et al. fully characterize all maximally general private mechanisms with a finite input space in the differential privacy setting. First they introduce the concept of *column-graphs*[3] of a private mechanism, which mark the tight privacy constraints in one column of the mechanism X.

Definition 2.5 (Column graph [11])**.** Let X be an α-differentially private mechanism with a finite input space. Let r be some possible output of X, and x_r be its corresponding column in X. Let I be the input space of X (corresponding to X's rows). The graph associated with this column has I as the set of nodes, and for any $i_1, i_2 \in I$, there is a directed edge (i_1, i_2) if i_1 and i_2 match neighboring databases and $x_{i_1,r} = \alpha x_{i_2,r}$, and a directed edge (i_2, i_1) if $x_{i_2,r} = \alpha x_{i_1,r}$. The direction of the edges is only necessary to distinguish between maximally general mechanisms which have similar undirected column-graphs, but it will not be essential to the rest of this article.

Kifer and Lin characterize the maximally general differentially private mechanisms with a finite input space:

Theorem 2.6 ([11])**.** *Fix a privacy parameter α and a database query f with a finite range for databases of a specific size. Let X be an α-differentially private mechanism with a finite range. Then X is maximally general iff each column graph of X's columns (according to the privacy constraints implied by f) is connected.*

This theorem shows that we wish to maximize the set of tight privacy constraints in order to make a private mechanism as general as possible. Notice that having just one entry of a column in X and the spanning tree of this column's graph (we need to know the direction of the edges as well), determines all the entries of this column.

[3] Kifer et al. actually define *row graphs* and not *column graphs*. We follow the matrix structure of [9, 10] which is simply the transposed matrix of the one used by Kifer et al., hence the difference in terminology.

2.5 Universal Mechanisms

Consider a collection of Bayesian information consumers, and suppose we wish to enable each of the information consumers to sample a result from a differentially private mechanism optimizing her utility. Ghosh et al. [9] showed that this does not necessarily require executing multiple mechanisms: if the query is a count query, then it is possible to construct one *universally optimal* mechanism U, from which all information consumers can *simultaneously* recover an optimal mechanism for their needs by *remapping*. I.e., every information consumer has an optimal private mechanism which is derivable from U. This result is repeated for risk-averse information consumers by Gupte et al. [10]. More formally:

Theorem 2.7 (Universal optimality, Bayesian consumers [9])**.** *Fix a privacy parameter $\alpha \in (0, 1)$. There exists an α-differentially private mechanism U for a single count query, such that for every prior $\bar{p}$ and every monotone loss function $\ell(\cdot, \cdot)$ there exists a (deterministic) remapping T such that $T \circ U$ implements an optimal oblivious mechanism for $\bar{p}, \ell(\cdot, \cdot)$.*

Theorem 2.8 (Universal optimality, risk-averse consumers [10])**.** *Fix a privacy parameter $\alpha \in (0, 1)$. There exists an α-differentially private mechanism U for a single count query, such that for every set S of possible outcomes and every monotone loss function $\ell(\cdot, \cdot)$ there exists a (probabilistic) remapping T such that $T \circ U$ implements an optimal oblivious mechanism for $S, \ell(\cdot, \cdot)$.*

It turns out that in both theorems U is realized by the geometric mechanism – a variant of the mechanism adding Laplace noise of [6]. Note that there may be optimal mechanisms which cannot be derived from the geometric mechanism, but for every information consumer there is at least one private mechanism that is derivable from the geometric mechanism and is optimal for her.

3 Impossibility of Universally Optimal Mechanisms for Generalizations of Count Queries

When the domain of the database records is $\{0, 1\}$, a count query is equivalent to a sum query. Theorems 2.7 and 2.8 can hence be thought of as applying to a *sum query* over the integers, where the domain of the database is Binary. It is natural to ask whether the results of these theorems can be extended to showing that universally optimal mechanisms exist for sum queries when the underlying data is taken from a larger domain such as $\mathcal{D} = \{0, 1, \dots, m\}$ where $m \geq 2$. We answer this question negatively.

Consider the case $m = 2$. Recall that an oblivious differentially private mechanism can be described by a row-stochastic matrix $X = (x_{i,j})$, such that $x_{i,j}$ is the probability of the mechanism to return j when the exact result is i. A difference of the case $m = 2$ from count queries ($m = 1$) is that applying a sum query to two neighboring databases may yield results which differ by 0, 1, or 2 (instead of 0 or 1). Therefore, in the linear program describing mechanism X equations (2) and (3), should be replaced by the following four constraints (the range for i in the other equations should be modified to $0, \dots, 2n$):

$$x_{i,r} \geq \alpha x_{i+1,r}, \quad \alpha x_{i,r} \leq x_{i+1,r} \qquad \forall i \in \{0, \dots, 2n-1\}, \forall r \in \mathcal{R}$$
$$x_{i,r} \geq \alpha x_{i+2,r}, \quad \alpha x_{i,r} \leq x_{i+2,r} \qquad \forall i \in \{0, \dots, 2n-2\}, \forall r \in \mathcal{R}$$

Once again, a consumer's optimal mechanism can be found by solving a linear program with all the constraints and the appropriate target function.

3.1 The Basic Impossibility Result for Sum Queries

We first consider the case where the database contains $n = 1$ record, taking values in {0,1,2} (i.e., $m = 2$). Later, we generalize to $n \geq 1$ and $m \geq 2$. Note that in the case of $n = 1$, the non-oblivious mechanisms are identical to oblivious mechanisms. We consider non-oblivious universal mechanisms as well when generalizing this result lo larger values of n.

Observation 3.1. *In the Bayesian model there exists an information consumer whose only optimal mechanism is* $X = \frac{1}{1+2\alpha} \cdot \begin{bmatrix} 1 & \alpha & \alpha \\ \alpha & 1 & \alpha \\ \alpha & \alpha & 1 \end{bmatrix}$ *and an information consumer whose optimal mechanisms are all of the form* $Y = \frac{1}{1+\alpha} \cdot \begin{bmatrix} 1 & \alpha & 0 \\ \alpha & 1 & 0 \\ q & 1+\alpha-q & 0 \end{bmatrix}$, *where* $q \in [\alpha, 1]$.

Proof. Consider an information consumer with a prior $\bar{p} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ and a loss function ℓ_{bin} (i.e., a penalty of 1 whenever she chooses an answer different from the exact result, and no penalty otherwise). It is easy to see that no optimal mechanism for this consumer outputs a value not in $\{0, 1, 2\}$.

The information consumer wishes to minimize

$$\sum_{i=0}^{2} p_i \sum_{r=0}^{2} x_{i,r} \cdot \ell(i, r) = \frac{1}{3} \sum_{i=0}^{2} \sum_{r \neq i} x_{i,r} = \frac{1}{3} \sum_{i=0}^{2} (1 - x_{i,i}) = 1 - \frac{1}{3} \sum_{i=0}^{2} x_{i,i}.$$

And so, the consumer's goal is to maximize $\sum_{i=0}^{2} x_{i,i}$ subject to maintaining α-differential privacy.

For $i \in \{0, 1, 2\}$, having α-differential privacy implies

$$\alpha x_{i,i} \leq x_{j,i} \quad \forall j \in \{0, 1, 2\} \setminus \{i\}, \tag{8}$$

and hence (by summing up Equation (8) for $j \neq i$), we get

$$2\alpha x_{i,i} = \sum_{\substack{j=0 \\ j \neq i}}^{2} \alpha x_{i,i} \leq \sum_{\substack{j=0 \\ j \neq i}}^{2} x_{j,i}. \tag{9}$$

Summing up Equation (9) for $i \in \{0, 1, 2\}$ we get

$$\sum_{i=0}^{2} 2\alpha x_{i,i} \leq \sum_{i=0}^{2} \sum_{\substack{j=0 \\ j \neq i}}^{2} x_{j,i} = \sum_{i=0}^{2} (1 - x_{i,i}) = 3 - \sum_{i=0}^{2} x_{i,i},$$

and we can now conclude that $\sum_{i=0}^{2} x_{i,i} \leq \frac{3}{2\alpha+1}$. This inequality is tight iff Equation (8) is tight (i.e., $x_{j,i} = \alpha x_{i,i}$) for every $i \neq j$. In that case, we get the following system of linear equations:

$$\begin{aligned} x_{11} + \alpha x_{22} + \alpha x_{33} &= 1 \\ \alpha x_{11} + x_{22} + \alpha x_{33} &= 1 \\ \alpha x_{11} + \alpha x_{22} + x_{33} &= 1 \end{aligned}$$

Since the three equations are linearly independent, we get a *unique* solution: $x_{1,1} = x_{2,2} = x_{3,3} = \frac{1}{1+2\alpha}$.

A similar proof shows that mechanisms of the form Y are the only mechanisms optimal for information consumers with a prior $p_0 = p_1 = \frac{1}{2}, p_2 = 0$ and loss function ℓ_{bin}.

It may seem like we restrict ourselves only to information consumers with the ℓ_{bin} loss function. Note that, according to Theorem 2.6, there are not so many maximally general mechanisms whose range is a subset of $\{0, 1, 2\}$, and some of them are not optimal for any consumer. Therefore, the mechanisms described are also the only optimal mechanisms for a variety of other information consumers, such as whose prior is $p_0 = p_1 = \frac{1}{2}$, $p_2 = 0$ and loss function is ℓ_1. Also, even more such consumers can be found easily in any sequence of consumers which converge to consumers with such unique optimal mechanisms (i.e., their priors and loss functions converge to the prior and loss function of the consumer we chose). Such information consumers with close priors and close loss functions to the ones described above will have the same unique optimal mechanisms. □

Observation 3.2. *In the risk-averse model there exists an information consumer whose only optimal mechanism is* $X = \frac{1}{1+2\alpha} \cdot \begin{bmatrix} 1 & \alpha & \alpha \\ \alpha & 1 & \alpha \\ \alpha & \alpha & 1 \end{bmatrix}$ *and an information consumer whose optimal mechanisms are all of the form* $Y = \frac{1}{1+\alpha} \cdot \begin{bmatrix} 1 & \alpha & 0 \\ \alpha & 1 & 0 \\ q & 1+\alpha-q & 0 \end{bmatrix}$*, where* $q \in [\alpha, 1]$.

Proof. Consider an information consumer whose loss function is ℓ_{bin} who knows the support of the query is $S = \{0, 1, 2\}$. As in the previous observation, the support of any optimal mechanism for this consumer must be a subset of $\{0, 1, 2\}$. Notice that if the consumer uses the mechanism described by X then her maximal expected loss is $\frac{2\alpha}{1+2\alpha}$.

Assume for a contradiction that the consumer has another mechanism X' with maximal expected loss at most $\frac{2\alpha}{1+2\alpha}$. I.e.,

$$\max\{x'_{0,1} + x'_{0,2}, x'_{1,0} + x'_{1,2}, x'_{2,0} + x'_{2,1}\} \le \frac{2\alpha}{(1 + 2\alpha)}. \tag{10}$$

Since $X' \neq X$, Equation (10) implies that $x'_{i,j} < \frac{\alpha}{1+2\alpha}$ for some $i \neq j$. Taking into account that X' is α-differentially private we get $x'_{j,j} \le \frac{1}{\alpha} \cdot x'_{i,j} < \frac{1}{1+2\alpha}$, and hence the maximal expected loss is at least $\sum_{i \neq j} x'_{i,j} = 1 - x'_{j,j} > 1 - \frac{1}{1+2\alpha} = \frac{2\alpha}{1+2\alpha}$, in contradiction to the assumption that this mechanism is at least as good as X for this information consumer.

A similar proof shows that mechanisms of the form Y are the only mechanisms optimal for an information consumer with auxiliary knowledge of the support $S = \{0, 1\}$ and loss function ℓ_{bin}. As in the previous observation, the mechanisms described are also the only optimal mechanisms for a variety of other information consumers. □

We will now use these two observations to show that in both models no universally optimal mechanism U exists. (This is true even if we allow U to have a non-discrete range.)

Claim 3.3. *No α-differentially private mechanism can derive both X and an instance of Y.*

Proof. Assume for a contradiction that such a mechanism U exists, so X and some instance of Y are both derivable from U. For simplicity we refer to this instance as Y. By Theorem 2.6, X is a maximally general mechanism. Therefore $U \preceq_G X$, and hence $Y \preceq_G X$, i.e., there exists a random remapping T such that $Y = XT$. Denote by x_j the j^{th} column of X, and by y_k the k^{th} column of Y. We get that

$$y_k = t_{0,k} \cdot x_0 + t_{1,k} \cdot x_1 + t_{2,k} \cdot x_2, \qquad \forall k \in \{0, 1, 2\}$$

Note that some α-differentially privacy constraints in Y are tight. Specifically, $y_{1,0} = \alpha y_{0,0}$ and $y_{0,1} = \alpha y_{1,1}$. As Y's columns are non-negative linear combinations of X's columns, such a tight constraint in a column of Y appears only if this column is a linear combination of columns of X in which the same privacy constraints are also tight. Note that the first two entries of every column in Y correspond to a tight constraint. But since $x_{0,2} = x_{1,2} > 0$, mapping this column of X by T to any column of Y (even with just a positive probability), yields a mechanism with a column in which the first two entries do not correspond to a tight constraint. Therefore, a contradiction. □

3.2 Generalizing the Impossibility Result for Sum Queries

So far we have shown the following: if n, the number of records in the database, is 1, and the range of values is $0, \ldots, m$ where $m = 2$, then no universal private mechanism for sum queries yields optimal utility for all consumers. Next, we generalize these impossibility results to the case $m \geq 2$ (and $n = 1$), and later present also the case where $n > 1$. Hence, we will conclude the following theorem:

Theorem 3.4. *No universally optimal mechanism exists for sum queries for databases whose records take values in the set $\{0, 1, \ldots, m\}$ where $m \geq 2$. This holds both for the Bayesian and the risk-averse utility models.*

3.2.1 Generalizing the Sum Query Impossibility Result to $m > 2$

Consider the case where the database consists of one record, and the possible values in this record are 0 to m. Let

$$X = \frac{1}{1+m\alpha} \cdot \begin{bmatrix} 1 & \alpha & \alpha & \cdots & \alpha \\ \alpha & 1 & \alpha & \cdots & \alpha \\ \alpha & \alpha & 1 & \cdots & \alpha \\ \vdots & \vdots & & \ddots & \vdots \\ \alpha & \alpha & \alpha & \cdots & 1 \end{bmatrix} ; \; Y = \frac{1}{1+(m-1)\alpha} \cdot \begin{bmatrix} 1 & \alpha & \alpha & \cdots & \alpha & 0 \\ \alpha & 1 & \alpha & \cdots & \alpha & 0 \\ \alpha & \alpha & 1 & \cdots & \alpha & 0 \\ \vdots & \vdots & & \ddots & \vdots & \\ \alpha & \alpha & \alpha & \cdots & 1 & 0 \\ q_1 & q_2 & q_3 & \cdots & q_m & 0 \end{bmatrix}, \tag{11}$$

where $\alpha \leq q_i \leq 1$ and $\sum_{i=1}^{m} q_i = 1 + (m-1)\alpha$. Similar arguments to those used for the case $m = 2$ show that X is the unique optimal mechanism for an information consumer with loss function ℓ_{bin} and prior $p_0 = p_1 = \cdots = p_m = \frac{1}{m+1}$ in the Bayesian utility model and for an information consumer with support $S = \{0, 1, \ldots, m\}$ in the risk-averse utility model. Also, mechanisms of the form Y are the only optimal mechanisms for the information consumers with loss function ℓ_{bin} and prior $p_0 = p_1 = \cdots = p_{m-1} = \frac{1}{m}$, $p_m = 0$ in the Bayesian model, and for an information consumer with support $S = \{0, 1, \ldots, m-1\}$ in the risk-averse model. Once again, these mechanisms are also the only optimal private mechanisms for a variety of other consumers as well. Using the same arguments to those in the proof of Claim 3.3, it follows that X and Y are not derivable from one single mechanism.

3.2.2 Generalizing the Sum Query Impossibility Result to $n > 1$

Now consider the case where the number of records in the database is larger than 1. We first prove the impossibility of an *oblivious* universally optimal mechanism. Consider two consumers with loss function ℓ_{bin}. The first consumer believes that the result of the sum query is bounded by m (in the Bayesian case, the consumer holds a uniform prior over $\{0,\dots,m\}$). No optimal mechanism for this consumer returns values larger than m, so in the mechanism matrix the columns corresponding to values greater than m contain zeros. Refer to some optimal mechanism for this consumer as X'. Ignoring rows and columns of X' that correspond to values greater than m, the remaining entries exactly form the mechanism X of Equation (11). (Observe that such an extension of mechanism X is indeed feasible, as any row which pertains to a value greater than m can be identical to the row which pertains to the value m, and so the privacy constraints hold. Such a mechanism is also optimal, as the utility is a function of only the rows $\{0,1,\dots,m\}$, due to the consumer's prior, so we cannot achieve a better utility than the utility gained by mechanism X). The second consumer believes that the query result cannot be larger than $m-1$ (in the Bayesian case, the consumer holds a uniform prior over $\{0,\dots,m-1\}$). Refer to some optimal mechanism for this consumer as Y'. A similar argument shows, that ignoring rows and columns that pertains to values greater than m, the remaining entries match the mechanism Y of Equation (11). Assume for a contradiction that X' and Y' are both derivable from some mechanism U'. Therefore there exist remappings T, S such that $X' = U'T$ and $Y' = U'S$. Let U be the mechanism U' reduced to only the inputs $\{0,1,\dots,m\}$. Reducing U' to U, we get that $X = UT$ and $Y = US$. According to the previous subsection these two mechanisms cannot be derived from a single oblivious mechanism, due to the same arguments in the proof of Claim 3.3. Thus, a contradiction.

Now suppose for a contradiction that both the mechanisms are derived from a single *non*-oblivious mechanism U^*. This means that U^*'s input space corresponds to databases rather than to query results. Suppose there is a remapping T such that $X^* = U^*T$. This means that the rows of X^* correspond to databases as well. We assume that X^* is oblivious (as universal optimality was shown not to exist even for count queries when consumers choose non-oblivious optimal mechanisms [9]). Therefore, applying U^* on two databases with the same query result and then applying T on U^*'s output, yields identical rows in X^* (which is described as a single row in the oblivious matrix X above). Note that although X^*'s input and output spaces are discrete (and so we can refer to X^* as a matrix), we assume nothing on U^*'s outputs and T's inputs. Reducing U^* to an input space of only $m+1$ databases with different query results and applying the remapping T on this reduced mechanism's output, yields mechanism X completely. Similarly, applying some remapping S on the same reduced mechanism yields mechanism Y. Now reduce U^* to inputs which are the databases $(0,0,\dots,0,q)$ where q is any possible record value. Refer to this mechanism as U. According to the assumptions, we get that $X = UT, Y = US$. Also note that every two possible inputs of U are neighboring databases, and so U must satisfy privacy constraints as any oblivious mechanism. Therefore, we get a simple reduction to the case of an oblivious mechanism U, and the same impossibility result applies also to the case of non-oblivious universal mechanism[4]. Thus, we conclude Theorem 3.4

3.3 Impossibility of Universally Optimal Mechanisms for Histogram Queries

The previous subsection shows that no universally optimal mechanisms exist for sum queries. In this and the following sections we consider other generalizations of count queries. One natural generalization is to histogram queries, and another is to bundles of simultaneous count queries. We

[4] Actually, this also shows that enabling universal non-oblivious mechanisms cannot resolve such impossibility for every query whenever there are 3 (or more) values which are the exact query results of 3 different neighboring databases.

begin with histogram queries. Note that a count query may be thought of as a histogram query where the database records are partitioned into two categories: those which satisfy a predicate, and those which do not. Consider now a histogram query which partitions the database records into three categories or more.

Theorem 3.5. *No universally optimal mechanism exists for histogram queries, except for histograms for one predicate and its complement or trivial predicates. This holds both for the Bayesian and the risk-averse utility models.*

Proof. Once again, consider first the case where there is only one record in the database, and the query is for a histogram which partitions the possible records into three categories. The only possible results for such a query are $(1, 0, 0)$, $(0, 1, 0)$ and $(0, 0, 1)$. Notice that all these histograms result from neighboring databases. Now consider information consumers whose loss function is either ℓ_{bin} or ℓ_1 (in the case of a histogram over one record they both result with 0 if the output matches the exact result and a constant otherwise). Refer to the first possible result as 0, the second possible result as 1, and the third possible result as 2. Notice now that we have exactly the same constraints for valid mechanisms as we had for the sum query with just one record. Also, the utility expression for each of the consumers is the same. The problem of universally optimizing the utility for all ℓ_{bin} information consumers (or ℓ_1 consumers) is now reduced to the same problem for sum queries. According to Subsection 3.1, universally optimizing the utility for all such consumers is impossible, and so it is impossible to construct a mechanism for this specific case as well.

We now generalize this result for histograms over larger databases and partitions of any number of categories larger than 2. First, consider the case of querying one record for a histogram of $c \geq 3$ categories. This can easily be reduced to problems we have already answered negatively. One way is to notice that as in the case where $c = 3$ (in which we reduced this problem to the problem of sum queries where the records' values bound is $m = 2$), larger values of c can easily be reduced to sum queries with larger bounds on the records' values m. For every number of c partitions, there are exactly c possible results for the histogram over one record. They are all the results of neighboring databases. Refer to these results as $0, 1, 2, \ldots, c - 1$. Again, this is exactly like constructing a universally optimal mechanism for sum queries over one record, in which the bound on its values is $m = c - 1$. This is impossible as was shown in Subsection 3.2.1. Another way to be convinced is to refer to the partitions as $A_1, A_2, \ldots, A_c$. Now consider only consumers whose loss functions are depend on the number of records in $A_1, A_2, (A_3 \cup A_4 \cup \ldots \cup A_c)$. These loss functions are monotone. This reduces the current problem to the problem of histograms over a partition of only 3 categories, to which we already proved negative results.

We now generalize this result further to any size of the database. The same argument that was applied in Subsection 3.2.2 for sum queries, applies here as well (even for the case of non-oblivious universal mechanisms). Consider only consumers with a prior such that all records except perhaps one fall into one specific category of the histogram. Querying for a histogram on such a database reduces to the result of the same histogram over one record only. Even if we consider only these consumers, we know that no one single mechanism can optimize their utilities over all possible mechanisms. Since there is no such mechanism that optimizes these consumers' utilities, there is obviously no mechanism that yields optimized results for all possible consumers. Therefore, even for larger databases, there is no universally optimal private mechanism for histogram queries. □

3.4 Impossibility of Universally Optimal Mechanisms for Bundles of Count Queries

We now consider the generalization of single count queries to a bundle of count queries, where a bundle contains several simple (non trivial) count queries that need to be answered simultaneously. Note that a consumer's disutility for a bundle query need not be the sum of the losses for the separate basic queries – it may be a more involved function of the bundle outputs. For instance, a consumer with the ℓ_{bin} loss function has no loss if all the results he uses are correct, and has one unit of loss if one or more of the results he uses are incorrect, no matter how many. Furthermore, information consumers may have auxiliary knowledge about the dependency between bundle outputs.

Theorem 3.6. *No universally optimal mechanism exists for bundles of more than one simultaneous non-trivial count queries. This holds both for the Bayesian and the risk-averse utility models.*

Proof. Such a generalization of count queries proves to be no different than the other intuitive generalizations we have already discussed. Note that two simultaneous non-trivial different predicates actually partition the records domain into 4 categories: those which satisfy both predicates, those which satisfy none of them, and those which satisfy just the first or just the second. If the predicates are somehow related, then the predicates might partition the domain into only 3 categories. This may happen in various cases, namely if one of the predicates is a subset of the other, if no record can possibly satisfy both of the predicates, or if any possible record must satisfy at least one of the predicates. Either way, there are always three different outputs for such bundles which result from three neighboring databases. (This is of course true also if the bundle consists of more than two simultaneous count queries). Once more, consider two different information consumers. The first has the ℓ_{bin} loss function and a uniform prior over these three outputs (Resp. in the risk-averse model, her support is the set of these three outputs). The second consumer also uses the ℓ_{bin} loss function and has a uniform prior over two of these outputs. (Resp. in the risk-averse model, her support is a set of two of these three outputs). Name these different outputs 0, 1 and 2. As in the previous subsection, the problem of universally optimizing the utility for all ℓ_{bin} information consumers is now reduced to the same problem presented in Subsection 3.1 . (The constraints for valid mechanisms are the same, and the utility expression for each of the consumers is the same). The only optimal mechanisms for the chosen information consumers are the same as those in Observations 3.1 and 3.2. According to Claim 3.3, such mechanisms are not derivable from any single private mechanism, and so universally optimizing the utility for all such consumers is impossible in the queries bundles as well.

□

4 A Characterization of Universal Optimality in the Bayesian Setting

We now discuss a more general setting, where a query (not necessarily related to sum or count) is answered by a differentially private mechanism in the Bayesian utility model. We follow other works on this subject and only consider oblivious private mechanisms. Note that although our results do not exclude the possibility of non-oblivious differentially private mechanisms, our techniques yield that no such non-oblivious universally optimal mechanisms exist for many natural functions. Specifically, enabling universal non-oblivious mechanisms cannot resolve such impossibilities for a query whenever there are 3 (or more) values which are the exact query results of 3 different neighboring databases. This is due to the same argument that was used in Subsection 3.2.2.

Let the database records be taken from a discrete domain $\mathcal{D}$ and let the query be $f : \mathcal{D}^n \to \mathcal{R}_f$ (wlog, we will assume that f is a surjective function, in which case $\mathcal{R}_f = \{f(D) : D \in \mathcal{D}^n\}$ is also a discrete set). Define the following graph where edges correspond to answers f may give on neighboring databases (and hence to restrictions on output distributions implied by differential privacy):

Definition 4.1 (Privacy Constraint Graph)**.** Fix a query $f : \mathcal{D}^n \to \mathcal{R}_f$. The *Privacy Constraint Graph* for f is the undirected graph $G_f = (V, E)$ where $V = \mathcal{R}_f$ is the set of all possible query results and $E = \{(f(D_1), f(D_2)) : D_1, D_2 \in \mathcal{D}^n \text{ are neighboring}\}$. The *degree* of the constraint graph, $\Delta(G_f)$, is the maximum over its vertices' degrees. For $i_2, i_2 \in \mathcal{R}_f$, G_f induces a distance metric $d_{G_f}(i_1, i_2)$ that equals the length of the shortest path in G_f from i_1 to i_2.

Observe that the constraint graph is connected for any query f: If $i_1 = f(D_1)$ and $i_2 = f(D_2)$ then there is a sequence of neighboring databases starting with D_1 and ending in D_2, and hence a path from i_1 to i_2 in G_f.

Recall that the results of [9, 10] are restricted to loss functions $\ell(i, r)$ that are monotonically non decreasing in the metric $|i - r|$. In our more general setting, we avoid interpreting outcome of f as points of a specific metric space, and hence we only consider the ℓ_{bin} loss function, which would remain monotone under any imposed metric.

Outline of this Section. We are now ready to describe the results of this section. Let f be a query, and G_f its constraint graph. We first show that if G_f is a single cycle, then no universally optimal mechanism exists for f. This impossibility result is then extended to the case where G_f contains a cycle.

Theorem 4.2. *Fix a query $f : \mathcal{D}^n \to \mathcal{R}_f$, and let G_f be its constraint graph. Consider Bayesian information consumers with loss function ℓ_{bin}. If G_f contains a cycle then no universally optimal mechanism exists for these consumers.*

Constraint graphs of sum queries (for $m \geq 2$), histograms and bundles of queries all have cycles of length 3, so, in the Bayesian utility model, Theorem 4.2 generalizes all our previous results.

Next, we consider the case where G_f is a tree and show that if G_f contains a vertex of degree 3 or higher, then no α-differentially private universally optimal mechanism exists for f for $\alpha > 1/(\Delta(G_f) - 1)$. (Recall that the closer α is to one, the better privacy we get.)

Theorem 4.3. *Fix a query $f : \mathcal{D}^n \to \mathcal{R}_f$, and let G_f be its constraint graph. Consider Bayesian information consumers with loss function ℓ_{bin}. If the privacy parameter $\alpha > 1/(\Delta(G_f) - 1)$ then no universally optimal mechanism exists for these consumers.*

We can conclude from theorems 4.2 and 4.3 that for $\alpha > 0.5$, the only functions f for which universally optimal mechanisms exist are those where G_f is a simple chain, as is the case for the count query.

The proof structure is similar to the one presented in the previous section for sum queries. We begin with the case where G_f is a simple cycle. We consider two consumers with different priors and loss function ℓ_{bin}, and show that the optimal mechanisms for these consumers must have specific structures (in the sense that some privacy constraints are satisfied tightly). Once again, we show that for two mechanisms with such structures, there is no mechanism which is at least as general as these two (i.e., there is no single mechanism which derives both of them).

Next, we extend the proof to the case where G_f contains a cycle. We focus on a cycle in G_f of smallest size m, and consider two information consumers. The consumers are similar to those for the case where G_f is a cycle, and so are the optimal mechanisms for them, except that we need to prove that these optimal mechanisms can be extended in a differentially private manner to the entire range of f. For that we introduce a labeling of G_f in which the labels of adjacent vertices differ by at most one modulo m.

Last, we discuss the case where G_f is a tree containing a vertex of degree at least 3. Focusing on that vertex and three of its adjacent vertices, we present three consumers with different priors. Again, we focus on the corresponding entries in the matrices of their optimal mechanisms, and find which constraints must be tight. Assuming all three mechanisms are derived from a single mechanism U, we present three different partitions of U's range according to which constraints are tight for every measurable subset of U's range. Combining the attributes from these partitions, we get one elaborated partition of U's range. We can then assume U's range is finite and reveal the structure of its matrix columns. Such a structure of U's columns (for the consumers we chose) is feasible iff we compromise for a privacy parameter $\alpha \leq 0.5$. Finally, we generalize this claim to any degree of one vertex.

4.1 The Basic Case: G_f is a Cycle

We begin with the simple case where G_f is a single cycle of $m > 2$ vertices[5].

Claim 4.4. *If the constraint graph G_f of $f : \mathcal{D}^n \to \mathcal{R}_f$ is a single cycle, then no universally optimal mechanism for Bayesian information consumers exists for f.*

Proof. Assume G_f is the cycle $C_m = (v_0, v_1, \ldots, v_{m-1}, v_0)$. We already proved impossibility of universal optimality for the case $m = 3$ in Claim 3.3. We now deal with the case $m > 3$. As in the proof of Claim 3.3, we will present two information consumers, and their corresponding optimal mechanisms, and prove that these cannot be derived from a single mechanism.

We first consider an information consumer with loss function ℓ_{bin} and prior $p_{v_0} = p_{v_1} = \cdots = p_{v_{m-1}} = 1/m$, and construct the unique optimal mechanism X for this consumer. (X is represented by an $m \times m$ matrix since with the ℓ_{bin} loss function the support of the optimal mechanism's range must match the support of the consumer's prior.) An optimal mechanism minimizes

$$\sum_{v_i \in C_m} p_{v_i} \sum_{r \in C_m} x_{v_i,r} \cdot \ell_{bin}(v_i, r) = \sum_{v_i \in C_m} p_{v_i} \cdot (1 - x_{v_i,v_i}) = 1 - \frac{1}{m} \sum_{v_i \in C_m} x_{v_i,v_i},$$

and hence, the consumer's goal is to maximize $\sum_{v_i \in C_m} x_{v_i,v_i}$ subject to maintaining α-differential privacy. Maintaining α-differential privacy implies

$$\alpha^{d_{G_f}(v_i,v_j)} x_{v_i,v_i} \leq x_{v_j,v_i} \quad \forall v_i, v_j \in C_m, \tag{12}$$

and hence, by summing up the inequalities for all v_i, v_j, we get

$$\sum_{v_i \in C_m} \sum_{\substack{v_j \in C_m \\ v_j \neq v_i}} \alpha^{d_{G_f}(v_i,v_j)} x_{v_i,v_i} \leq \sum_{v_i \in C_m} \sum_{\substack{v_j \in C_m \\ v_j \neq v_i}} x_{v_j,v_i} = \sum_{v_i \in C_m} 1 - x_{v_i,v_i} = m - \sum_{v_i \in C_m} x_{v_i,v_i},$$

[5] An example query that yields such a graph is $f : \{0,1\}^n \to [m]$ defined as $f(d_1, \ldots, d_n) = \sum_{i=1} d_i \mod m$. If $n \geq m > 2$ then G_f is a cycle of size m.

and we conclude that

$$\sum_{v_i \in C_m} x_{v_i,v_i} \leq \frac{m}{1 + \sum_{\substack{v_j \in C_m \\ v_i \neq v_j}} \alpha^{d_{G_f}(v_i,v_j)}}.$$

This inequality is tight iff Equation (12) is tight (i.e., $x_{v_j,v_i} = \alpha^{d_{G_f}(v_i,v_j)} x_{v_i,v_i}$) for every $v_i \neq v_j \in C_m$. In such a case, we can find the mechanism's entries by solving a system of m linear equations (the sum of each row in the mechanism must be 1), in a similar argument to the one presented in the proof of Observation 3.1. Since these are m independent linear equations in m variables, our optimal solution for $x_{v_1,v_1}, \ldots, x_{v_m,v_m}$ is unique.

Utilizing the symmetry of the equations, we get that every row of X is a cyclic shift of:

$$\begin{aligned} &\delta \cdot (1, \alpha^1, \alpha^2, \ldots, \alpha^{(m-1)/2}, \alpha^{(m-1)/2}, \alpha^{(m-1)/2-1}, \ldots, \alpha^2, \alpha^1) && \text{if } m \text{ is odd,} \\ &\delta \cdot (1, \alpha^1, \alpha^2, \ldots, \alpha^{m/2-1}, \alpha^{m/2}, \alpha^{m/2-1}, \ldots, \alpha^2, \alpha^1) && \text{if } m \text{ is even.} \end{aligned} \tag{13}$$

where δ is chosen such that X is row-stochastic. The mechanism X satisfies α-differential privacy, it is optimal for our information consumer, and it is unique.

Our second information consumer uses ℓ_{bin} as her loss function, and prior $p_{v_0} = p_{v_1} = p_{v_2} = 1/3$ and $p_{v_3} = \cdots = p_{v_{m-1}} = 0$. Note that since $m > 3$ the vertices v_0, v_2 are not adjacent in G_f (so $d_{G_f}(v_0, v_2) = 2$). In constructing an optimal mechanism Y for the information consumer we will only consider the rows and columns pertaining to vertices v_0, v_1, v_2, noting that the columns for all other vertices contain only zeros, and there is some freedom with respect to the rows for the other vertices. Applying similar arguments as for mechanism X, we get that the columns of Y are of the forms $(1, \alpha^1, \alpha^2)^T$, $(\alpha^1, 1, \alpha^1)^T$, $(\alpha^2, \alpha^1, 1)^T$ (each of the columns may be multiplied by a different coefficient). By forcing row stochasticity, we can solve the following equations to get the coefficients:

$$\begin{bmatrix} 1 & \alpha^1 & \alpha^2 \\ \alpha^1 & 1 & \alpha^1 \\ \alpha^2 & \alpha^1 & 1 \end{bmatrix} \times \begin{pmatrix} c_1 \\ c_2 \\ c_3 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$$

and we get a unique structure on the entries of these rows and columns of Y. This mechanism is of no surprise, as these entries are merely the finite-range version of the geometric mechanism (as shown in [9]).

Summarizing our findings, we get that

$$X = \delta \cdot \begin{bmatrix} 1 & \alpha & \alpha^2 & \cdots & \alpha^2 & \alpha \\ \alpha & 1 & \alpha & \cdots & \alpha^3 & \alpha^2 \\ \alpha^2 & \alpha & 1 & \cdots & \alpha^4 & \alpha^3 \\ \vdots & \vdots & & \ddots & & \vdots \\ \alpha & \alpha^2 & \alpha^3 & \cdots & \alpha & 1 \end{bmatrix}; \; Y = \begin{bmatrix} c_1 & c_2 \cdot \alpha^1 & c_3 \cdot \alpha^2 & 0 & \cdots & 0 \\ c_1 \cdot \alpha^1 & c_2 & c_3 \cdot \alpha^1 & 0 & \cdots & 0 \\ c_1 \cdot \alpha^2 & c_2 \cdot \alpha^1 & c_3 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \vdots & \vdots & \vdots & 0 & \cdots & 0 \end{bmatrix}. \tag{14}$$

We now show that instances of such mechanisms X and Y are not derivable from a single mechanism. Since the conditions stated for these mechanisms are necessary for them to be optimal for the two consumers we chose, this will prove that there is no universally optimal mechanism in such a scenario.

Suppose, towards a contradiction, that there exists a mechanism U which derives both X and some instance of Y. According to the characterization of generally maximal differentially private mechanisms (Theorem 2.6), X is maximally general. Therefore, we get that U is derivable from X and so Y is derivable from X as well. Therefore, there exists a remapping matrix T such that $Y = XT$. Remember that Y's columns are linear combinations of X's columns with non-negative coefficients, as described in the proof of Claim 3.3. Any tight constraint met in one of Y's columns must match the same tight constraints in all of X's columns which appear in the linear combination of that column. Once again, any specific column of X must appear in at least one linear combination of one of Y's columns with a positive coefficient (as any possible output of X must be remapped to the values $\{v_0, v_1, v_2\}$ by T). Notice that one of X's columns is

$$\begin{aligned} &\delta \cdot (\alpha^{(m-1)/2}, \alpha^{(m-1)/2}, \alpha^{(m-1)/2-1}, \ldots, 1, \ldots, \alpha^{(m-1)/2-1})^T && \text{if } m \text{ is odd,} \\ &\delta \cdot (\alpha^{m/2-1}, \alpha^{m/2}, \alpha^{m/2-1}, \ldots, 1, \ldots, \alpha^{m/2-2})^T && \text{if } m \text{ is even,} \end{aligned}$$

Mapping this column into any of Y's first three columns (with any positive probability) cannot yield the tight constraints which appear in the first three entries of the chosen column in Y. Therefore, no such remapping T is feasible and we get a contradiction.

□

4.2 Impossibility of Universal Optimality When G_f Contains a Cycle

We now give a proof for Theorem 4.2 which deals with the case where G_f *contains* a cycle.

Proof. Let $C_m = (v_0, v_1, \ldots, v_{m-1}, v_0)$ be a cycle of smallest size in G_f. Based on C_m, we will consider two consumers whose optimal mechanisms contain as sub-matrices the matrices X, Y from the proof of Claim 4.4, and hence they cannot be derived from a single mechanism.

The First Consumer: uniform prior over C_m Consider an information consumer with loss function ℓ_{bin} and prior $p_{v_0} = p_{v_1} = \cdots = p_{v_{m-1}} = 1/m$ and $p_u = 0$ for every $u \notin C_m$. We will construct an optimal mechanism X' for this consumer, and will prove that (in some sense) it is unique. We begin with a labeling algorithm of the vertices in G:

1. Given $C_m = (v_0, v_1, \ldots, v_{m-1}, v_0)$, set $l(v_i) = i$ for $i \in \{0, \ldots, m-1\}$.
2. For s from 1 to $m - 1$:
   (a) Let V_s be the set of unlabeled vertices that are adjacent to vertices labeled $s - 1$.
   (b) Let $l(u) = s$ for all $u \in V_s$.
3. Let $l(u) = m - 1$ for all remaining vertices u.

Claim 4.5. *After applying the above algorithm, the labels for every two adjacent vertices differ by at most 1 (modulo m).*

Proof. We show that at any stage of the labeling, any two adjacent vertices satisfy the requirement that their labels differ by at most 1 (modulo m).

Note first that this holds for all labeled vertices after Step 1. Consider a vertex $u \in V_s$ (i.e., $l(u) = s$ is set in iteration s), and an adjacent vertex u' that is labeled $l(u') = s'$ prior to or on iteration s. Clearly, if $u' \in V_s \cup V_{s-1}$ then $s' \in \{s - 1, s\}$ and the statement holds for (u, u'). Otherwise, we consider two sub-cases. In the first, $l(u') = s' < s - 1$, and we are led to a contradiction since u remains unlabeled after iteration $s'+1$ whereas by definition $u \in V_{s'+1}$. In the second sub-case $l(u') = s' > s+1$ (if $s' = s+1$ the claim holds) and hence it must have been that u' was labeled in Step 1, i.e., $u' = v_{s'}$ for $s' \in \{s+1, \ldots, m-1\}$. Following the path of labels which led to the label of u we can get to the vertex v_0 via a path of length s. Noting that this path is disjoint from the length $m-s'$ path $v_{s'} \rightsquigarrow v_0 = v_{s'}, v_{s'+1}, \ldots, v_{m-1}, v_0$, we get that G contains the cycle $v_{s'} \rightsquigarrow v_0 \rightsquigarrow u \rightsquigarrow u'$ that is of length $m - s' + s + 1 < m$, in contradiction to C_m being the smallest cycle in G. To conclude the proof, note that every vertex $u \in G$ adjacent to some $u' \in G$ such that $l(u') \in \{0, 1, \ldots, m-2\}$ has been labeled in iteration $l(u')+1$ or earlier. Therefore in Step 3, the vertices which are not labeled yet are adjacent only to unlabeled vertices and to vertices with label $m-1$. Labeling the remaining vertices with $m-1$ satisfies the requirement. □

We now use the graph labels to construct an optimal mechanism X', represented by a matrix of dimensions $|\mathcal{R}_f| \times |\mathcal{R}_f|$. The entries of rows $u \notin C_m$ have no effect on the expected loss of this consumer, as $p_u = 0$. There are, however, restrictions on these rows, as the mechanism X' must be differentially private. We construct X' as follows:

1. For all $u \notin C_m$, set column u of X' to be a column of zeros.

2. For all $u \in C_m$, set row u of X' as in the optimal mechanism X described in the proof of Claim 4.4 (i.e., Equation (14)).

3. For all $u \notin C_m$, set row u of X' to be identical to the row corresponding with the vertex identically labeled in C_m.

Clearly, the resulting mechanism is row-stochastic. The privacy constraints also hold: suppose $u, u' \in \mathcal{R}_f$ are query results of neighboring databases. Therefore, they are adjacent in the constraint graph, and their labels differ by at most 1 (modulo m). And so, their matching rows in mechanism X' are either identical or they are the same as rows of two adjacent vertices $v_i, v_j \in C_m$ in mechanism X. Since the construction of rows in X hold to the privacy constraints, so do the rows of X'. In other words, we just showed that mechanism X can be extended to any query f whose constraint graph G_f contains C_m but no smaller cycles.

Notice that only rows of C_m affect the expected loss in X', which is hence identical to that of X. Since any mechanism in this scenario has to satisfy all the restrictions for just the vertices of the cycle C_m, and more, the expected loss for any optimal mechanism in the current scenario is lower bounded with that of X. Hence, we can conclude that X' is optimal for the information consumer, and furthermore, X' restricted to the rows corresponding to C_m is unique.

The Second Consumer: uniform prior over v_0, v_1, v_2 Consider an information consumer with loss function ℓ_{bin} and prior $p_{v_0} = p_{v_1} = p_{v_2} = 1/3$ and $p_u = 0$ for every other $u \in \mathcal{R}_f$. We argue that every optimal mechanism Y' for this consumer has the same structure on rows v_0, v_1, v_2 as mechanism Y in Equation (14). As the impossibility of universal optimality for the case of $m = 3$ was already covered, and we assumed $m > 3$, v_0 and v_2 are not adjacent in G_f. This enables us to label the vertices in such a way: $l(v_0) = 0$, $l(v_2) = 2$ and $l(u) = 1$ for any other vertex in G_f. Again, it is clear that every two adjacent vertices have labels which differ by 1 at most. Similar arguments as the ones presented for the first consumer, show that the first three rows of every optimal mechanism for this consumer (i.e. the rows for v_0, v_1, v_2) have the same structure as the first three rows of mechanism Y in Equation (14).

Assume towards a contradiction that both X' and Y' are derivable from a single mechanism U'. Therefore there exist remappings T, S such that $X' = U'T$ and $Y' = U'S$. Let U be the mechanism U' reduced to only the inputs of the cycle $C_m = \{v_0, v_1, \ldots, v_{m-1}\}$. Reducing U' to U, we get that $X = UT$ and $Y = US$. According to the previous subsection these two mechanisms cannot be derived from a single oblivious mechanism, due to the same arguments in the proof of Claim 3.3. Thus, we get a contradiction.

□

4.3 Impossibility of Universal Optimality When $\Delta(G_f) \geq 3$

We now focus on acyclic constraint graphs and prove Theorem 4.3 and its conclusion that for $\alpha > 0.5$ no universally optimal mechanisms exists unless the constraint graph is a simple chain.

Proof. For simplicity of this proof, we first focus only on 3 neighbors of a specific vertex, and prove that no universally optimal mechanism exists for $\alpha > 1/(3-1) = 0.5$. Later, we generalize this result for a vertex of any degree by taking into account all of the vertex's neighbors. The generalization is done using the same methods we use to prove the simpler case.

Let v_0 be a vertex in G_f with a degree greater than 2. Let v_1, v_2, v_3 be 3 of its neighbors. We choose some consumers with loss function ℓ_{bin} and zero a priori probability for all values other than v_0, v_1, v_2, v_3. We define some necessary conditions on the optimal mechanisms of these consumers and show it is impossible to simultaneously derive optimal mechanisms for these consumers from a single mechanism U (when $\alpha > 0.5$).

Note that by the tree structure of G_f, every mechanism that satisfies the requirements of differential privacy on query results v_0, v_1, v_2, v_3 can be easily extended to a differentially private mechanism on all results of $\mathcal{R}_f$[6]. Furthermore, since our consumers have zero a priori probability for all other values, the entries in rows corresponding to values other than v_0, v_1, v_2, v_3 do not affect the consumers' expected loss. Hence, it suffices to show the impossibility result for the case where G_f is restricted to v_0, v_1, v_2, v_3.[7]

Consider first an information consumer with prior $p_{v_0} = p_{v_1} = p_{v_2} = 1/3, p_{v_3} = 0$. note that v_1, v_0, v_2 is a simple path of length 3 in G_f for which the optimal mechanism was described in Section 4.1. Any optimal mechanism for this consumer is of the form

$$Y = \begin{bmatrix} c_0 & c_1 \cdot \alpha & c_2 \cdot \alpha & 0 \\ c_0 \cdot \alpha & c_1 & c_2 \cdot \alpha^2 & 0 \\ c_0 \cdot \alpha & c_1 \cdot \alpha^2 & c_2 & 0 \\ q_0 & q_1 & q_2 & 0 \end{bmatrix}$$

where $q_0 + q_1 + q_2 = 1$ and they are subject to some privacy constraints. The restrictions on the optimal mechanism for this consumer are the same as those on mechanism Y in Equation (14), only now v_0 is the vertex in the middle, so the first two rows were swapped, as were the first two columns.

[6] One possible extension is as follows: Suppose X is a mechanism from $\{v_0, v_1, v_2, v_3\}$ to $\{v_0, v_1, v_2, v_3\}$. Label each of the vertices v_0, v_1, v_2, v_3 by $l(v_i) = i$, then label every other vertex in the graph with the same label as its nearest labeled vertex. Construct a mechanism X' from X like this: Set $x'_{v_i,v_j} = x_{v_i,v_j}$ for every $i, j \in \{0, 1, 2, 3\}$. Set $x'_{v_i,u} = 0$ for every $u \notin \{v_0, v_1, v_2, v_3\}$. For every $u \notin \{v_0, v_1, v_2, v_3\}$, set the row of u to be the same as the row of $v_{l(u)}$.

[7] An example query that yields such a graph is $f : \{1, 2, 3\}^n \to \{0, 1, 2, 3\}$ defined as $f(D) = i$ if all records in D equal i, 0 otherwise.

Suppose that such a mechanism was derived from a universal mechanism U by some remapping T. Suppose for now that U's range is discrete and so it can be expressed in matrix form. (We abuse a little the notion of a matrix and allow U to have infinitely many columns, and T to have infinitely many rows, if needed). As noted before, this means that Y's columns are linear combinations with positive coefficients of columns in U. Also, remember that since the coefficients are non-negative, linearly combining columns which do not hold tight privacy constraints, cannot yield a column with tight constraints. Since T is row-stochastic, every row of T has at least one positive entry. This means that every column in U is remapped (with some positive probability) to a column in Y. Assume U does not have zero columns (otherwise we could just ignore them as they pertain to results which are not in U's range). From the reasons above and the the structure of constraints in Y which are tight, we conclude that all of U's columns can be partitioned into columns of the forms: $\delta_1 \cdot (1, \alpha, \alpha, *)^T$, $\delta_2 \cdot (\alpha, 1, \alpha^2, *)^T$, $\delta_3 \cdot (\alpha, \alpha^2, 1, *)^T$. The first set of columns is summed by T into the first column of Y, The second set is summed by T into the second column of Y, and the third set is summed to the third column of Y. The $*$ can take infinitely many values as it does not necessarily match to a tight constraint in Y.

Considering now an information consumer with a prior $p_{v_0} = p_{v_1} = p_{v_3} = 1/3, p_{v_2} = 0$, and applying the same arguments, we have that the non-zero columns of the universal mechanism U are partitioned into columns of the forms $\delta_1 \cdot (1, \alpha, *, \alpha)^T, \delta_2 \cdot (\alpha, 1, *, \alpha^2)^T, \delta_3 \cdot (\alpha, \alpha^2, *, 1)^T$. Similarly, considering a consumer with a prior $p_{v_0} = p_{v_2} = p_{v_3} = 1/3, p_{v_1} = 0$, we have that the non-zero columns of the universal mechanism U are partitioned into columns of the forms $\delta_1 \cdot (1, *, \alpha, \alpha)^T, \delta_2 \cdot (\alpha, *, 1, \alpha^2)^T, \delta_3 \cdot (\alpha, *, \alpha^2, 1)^T$.

Notice that every non-zero column in U must match one category in each of the partitions described above. Combining these conditions together, we have that the non-zero columns of U are partitioned into columns of the forms $\gamma_1 \cdot (1, \alpha, \alpha, \alpha)^T, \gamma_2 \cdot (\alpha, 1, \alpha^2, \alpha^2)^T, \gamma_3 \cdot (\alpha, \alpha^2, 1, \alpha^2)^T, \gamma_4 \cdot (\alpha, \alpha^2, \alpha^2, 1)^T$. As the columns in every category are proportional to one another, we assume that the mechanism U has exactly one column in each of these categories. We can assume that, since if U' is a mechanism with two non-zero columns which are proportional to one another, we can produce a mechanism U by replacing these columns with a single column containing their sum. Then U is derivable from U', and vice versa. Therefore these mechanisms are equivalent.

Note that we assumed U's range is discrete only for convenience. U's range can be continuous as well, as explained by Kifer and Lin [11]. Define T's inverse to be for every vertex v, $T^{-}(v) = \{o' \in Rng(U) : Pr[T(o') = v] > 0\}$. The same arguments from before hold, and we get that for every measurable $O' \subseteq T^{-}(v)$, and any adjacent vertices v_i, v_j, $\frac{Pr[T \circ U(v_i)=v]}{Pr[T \circ U(v_j)=v]} = \frac{Pr[U(v_i) \in O']}{Pr[U(v_j) \in O']}$, unless one of the probabilities is zero in which case all the probabilities are zero due to differential privacy constraints. This is because we assume $\frac{Pr[T \circ U(v_i)=v]}{Pr[T \circ U(v_j)=v]}$ is tight by differential privacy constraints, and it can be expressed as a positive combination over measurable sets in $T^{-}(v)$ which, therefore, must be tight as well. And so, the same structure of tight constraints as they appear in the derived mechanism, must appear also for every measurable subset of $T^{-}(v)$ for every v in the derived mechanism's output.

We conclude that a universal mechanism must be of the form:

$$U = \begin{bmatrix} c_0 & c_1 \cdot \alpha & c_2 \cdot \alpha & c_3 \cdot \alpha \\ c_0 \cdot \alpha & c_1 & c_2 \cdot \alpha^2 & c_3 \cdot \alpha^2 \\ c_0 \cdot \alpha & c_1 \cdot \alpha^2 & c_2 & c_3 \cdot \alpha^2 \\ c_0 \cdot \alpha & c_1 \cdot \alpha^2 & c_2 \cdot \alpha^2 & c_3 \end{bmatrix}.$$

The privacy and non-negativity constraints hold if $c_i \geq 0$ for every i. Imposing row-stochasticity, we can solve for the coefficients and get the unique solution: $c_1 = c_2 = c_3 = 1/(\alpha+1), c_0 = (1-2\alpha)/(\alpha+1)$. Mechanism U is only feasible if $c_0 \geq 0$, or equivalently $\alpha \leq 0.5$.

Note that, so far, we used only three of the vertices adjacent to v_0. Suppose v_0 has $k > 3$ neighbors. We actually can achieve stronger results by treating more consumers, each with a prior of uniform probability over only three vertices (one of which is v_0). Using the same arguments, and combining the partitions imposed by each of the consumers on U's columns, we get that U's positive columns are partitioned into columns of the forms: $\gamma_0 \cdot (1, \alpha, \alpha, \ldots, \alpha)^T$, $\gamma_1 \cdot (\alpha, 1, \alpha^2, \ldots, \alpha^2)^T$, $\gamma_2 \cdot (\alpha, \alpha^2, 1, \ldots, \alpha^2)^T$, $\ldots$, $\gamma_k \cdot (\alpha, \alpha^2, \alpha^2, \ldots, 1)^T$. Thus, a universal mechanism for all these consumers must have the structure:

$$U = \begin{bmatrix} c_0 & c_1 \cdot \alpha & c_2 \cdot \alpha & \ldots & c_k \cdot \alpha \\ c_0 \cdot \alpha & c_1 & c_2 \cdot \alpha^2 & \ldots & c_k \cdot \alpha^2 \\ c_0 \cdot \alpha & c_1 \cdot \alpha^2 & c_2 & \ldots & c_k \cdot \alpha^2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ c_0 \cdot \alpha & c_1 \cdot \alpha^2 & c_2 \cdot \alpha^2 & \ldots & c_k \end{bmatrix}.$$

Imposing row-stochasticity, we can solve for the coefficients and get the unique solution: $c_i = 1/(\alpha+1)$ for every $i > 0$ and $c_0 = (1-(k-1)\alpha)/(\alpha+1)$. Mechanism U is only feasible if $c_0 \geq 0$, or equivalently $\alpha \leq 1/(k-1)$.

□

References

[1] Avrim Blum, Katrina Ligett, Aaron Roth. A learning theory approach to non-interactive database privacy. In *STOC 2008*, pages 609–618.

[2] Irit Dinur, Kobbi Nissim. Revealing information while preserving privacy. In *PODS 2003*, pages 202–210.

[3] Cynthia Dwork. The Differential Privacy Frontier (Extended Abstract). In *TCC 2009*, pages 496–502.

[4] Cynthia Dwork, Moni Naor, Toniann Pitassi, Guy N. Rothblum. Differential privacy under continual observation. In *STOC 2010*, pages 715–724.

[5] Cynthia Dwork, Kobbi Nissim. Privacy-preserving datamining on vertically partitioned databases. In *CRYPTO 2004*, pages 528–544.

[6] Cynthia Dwork, Frank McSherry, Kobbi Nissim, Adam Smith. Calibrating noise to sensitivity in private data analysis. In *TCC 2006*, pages 265–284.

[7] Cynthia Dwork, Moni Naor, Omer Reingold, Guy N. Rothblum, Salil P. Vadhan. On the complexity of differentially private data release: efficient algorithms and hardness results. In *STOC 2009*, pages 381–390.

[8] Dan Feldman, Amos Fiat, Haim Kaplan, Kobbi Nissim. Private coresets. In *STOC 2009*, pages 361–370.

[9] Arpita Ghosh, Tim Roughgarden, Mukund Sundararajan. Universally utility-maximizing privacy mechanisms. In *STOC 2009*, pages 351–360.

[10] Mangesh Gupte, Mukund Sundararajan. Universally optimal privacy mechanisms for minimax agents. In *PODS 2010*, pages 135–146.

[11] Daniel Kifer, Bing-Rong Lin. Towards an axiomatization of statistical privacy and utility. In *PODS 2010*, pages 147–158.

[12] Frank McSherry, Kunal Talwar. Mechanism Design via Differential Privacy. In *FOCS 2007*, pages 94–103.

[13] Kobbi Nissim, Sofya Raskhodnikova, Adam Smith. Smooth sensitivity and sampling in private data analysis. In *STOC 2007*, pages 75–84.

[14] Aaron Roth, Tim Roughgarden. Interactive privacy via the median mechanism. In *STOC 2010*, pages 765–774.