

Constrained Private Mechanisms for Count Data

Graham Cormode¹, Tejas Kulkarni², and Divesh Srivastava³

Abstract—Concern about how to aggregate sensitive user data without compromising individual privacy is a major barrier to greater availability of data. Differential privacy has emerged as an accepted model to release sensitive information while giving a statistical guarantee for privacy. Many different algorithms are possible to address different target functions. We focus on the core problem of count queries, and seek to design *mechanisms* to release data associated with a group of n individuals. Prior work has focused on designing mechanisms by raw optimization of a loss function, without regard to the consequences on the results. This can lead to mechanisms with undesirable properties, such as never reporting some outputs (gaps), and overreporting others (spikes). We tame these pathological behaviors by introducing a set of desirable properties that mechanisms can obey. Any combination of these can be satisfied by solving a linear program (LP) which minimizes a cost function, with constraints enforcing the properties. We focus on a particular cost function, and provide explicit constructions that are optimal for certain combinations of properties, and show a closed form for their cost. In the end, there are only a handful of distinct optimal mechanisms to choose between: one is the well-known (truncated) geometric mechanism; the second a novel mechanism that we introduce here, and the remainder are found as the solution to particular LPs. These all avoid the bad behaviors we identify. We demonstrate in a set of experiments on real and synthetic data which is preferable in practice, for different combinations of data distributions, constraints, and privacy parameters.

Index Terms—Differential privacy, mechanism design, constrained optimization, count queries

1 INTRODUCTION

THERE has been considerable progress on the problem of how to release sensitive information with privacy guarantees in recent years. Various formulations have been proposed, with the model of differential privacy emerging as the most popular and robust [1]. Differential privacy (DP) lays down rules on the likelihood of seeing particular outputs given related inputs. Many different algorithms have been proposed to meet this guarantee, based on different objectives and input types [2]. The resulting descriptions of output probabilities for different inputs are referred to as mechanisms, such as the Laplace mechanism, Geometric mechanism and Exponential mechanism described in more detail subsequently.

In this paper, we focus on count queries, a fundamental problem in private data release that underpins many applications, from basic statistics of a dataset to complex spatial and graphical distributions. Count queries are needed to materialize frequency distributions, instantiate statistical models, and as the basis of SQL COUNT * queries. Counts can be applied to arbitrary groups, and based on complex predicates; hence they represent a very general tool. Abstractly, we have a group of n individuals, who each hold a private bit (encoding, for example, whether or not they possess a particular sensitive characteristic). The aim is

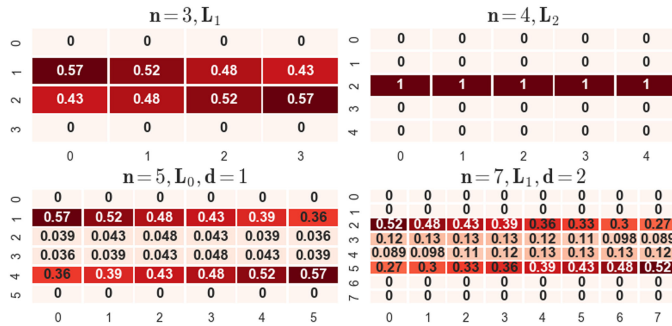
to release information about the sum of the bits, while meeting the stringent differential privacy guarantee. The usual model assumes the existence of a trusted aggregator, who receives the individual bits, and who aims to release a noisy representation of their sum. Since the value of the true answer is in $\{0 \dots n\}$, it is natural to restrict the output of the mechanism to this range also, to ensure downstream compatibility with subsequent data analysis expecting integer counts in this range. If we analyze how existing approaches to differential privacy handle this case, we find there are weaknesses. We consider the most relevant approach, mechanisms obtained via a linear programming framework.

Linear Programming Framework [3] (cf. Section 3). Ghosh et al. considered count queries and proved powerful theorems about utility-optimal mechanisms. They showed how to design mechanisms for count queries which minimize a loss function, via linear programming. The mechanisms obtained by solving linear programs specify, for each possible input, a probability distribution over allowable outputs. However, for common objectives, including to minimize the expected absolute error (denoted \mathbb{L}_1) and squared error (\mathbb{L}_2), we observed that the “optimal” mechanisms have some anomalous behavior, such as never reporting some values.

Fig. 1 gives some examples of this phenomenon in action. We show four optimal mechanisms generated by solving linear program described in Section 3 for different input sizes (n), under a privacy guarantee controlled by a parameter α (explained later, and set to a fixed value here). Each column gives the probability distribution over the outputs in the range 0 to n , for a given input count (also 0 to n). The case of optimizing the squared error (\mathbb{L}_2) is most striking: the “optimal” thing to do in this case is to ignore the input and always report ‘2’! But other cases are also problematic: all these optimal mechanisms never report some outputs

- G. Cormode and T. Kulkarni are with the University of Warwick, Coventry CV4 7AL, United Kingdom.
E-mail: {g.cormode, t.kulkarni.2}@warwick.ac.uk.
- D. Srivastava is with the AT&T Labs-Research, Bedminster, NJ 07921 USA.
E-mail: divesh@research.att.com.

Manuscript received 28 Sept. 2018; revised 20 Mar. 2019; accepted 11 Apr. 2019. Date of publication 16 May 2019; date of current version 11 Jan. 2021. (Corresponding author: Tejas Kulkarni.)
Recommended for acceptance by P. K. Chrysanthis, B. C. Ooi, and J. Dittrich.
Digital Object Identifier no. 10.1109/TKDE.2019.2912179

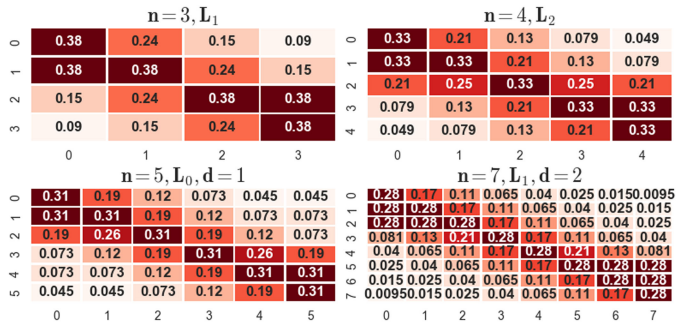
Fig. 1. Heatmaps of unconstrained mechanisms for $\alpha = 0.62$.

(gaps), and disproportionately report some others (spikes). For example, minimizing the absolute error for $n = 7$ has a chance of reporting the values 2 or 5 with at least 0.7 probability, regardless of the input value. Similarly, if we try to minimize the probability of reporting an answer that is more than 1 step away from the true input (denoted as \mathbb{L}_0 with $d = 1$), there is an over 90 percent chance of reporting 1 or 4.

Clearly, such results are counter-intuitive and show that blind optimization of simple objective functions can lead to unexpected and undesirable outcomes. To address this, we initiate the study of *constrained mechanism design*: requiring mechanisms to satisfy additional properties ensuring desired structure in the obtained mechanism and avoiding these pathologies. For example, we define the notion of *fairness*, which requires that the probability of reporting the true input is the same for all inputs; and *weak honesty*, where we require that the probability of reporting the true input is at least uniform (i.e., at least $\frac{1}{n+1}$). These both ultimately entail that every output is reported with a non-zero probability. We also consider various monotonicity properties, which preclude big spikes in probability for responses that are far from the truth. In total, we describe seven natural properties that one could demand of a mechanism. Our first result is to show how to extend the linear programming framework to incorporate these properties and eliminate pathological outcomes.

Fig. 2 shows the heatmap of constrained mechanisms satisfying all properties. The anomalies (spikes and gaps) seen in Fig. 1 are now eliminated. Recall that optimizing in the unconstrained \mathbb{L}_2 case returns a trivial solution that outputs 2 irrespective of input. Now, the probability mass in the corresponding constrained mechanism is more distributed and with probability at least $\frac{2}{3}$, the mechanism outputs a value differing from the true answer by at most 1 for all inputs. Similar observations can be made in the other instances. We go on to perform a detailed study of constrained mechanism design for count queries, and show some surprising outcomes:

- *Fully constrained mechanisms minimizing $\mathbb{L}_0, \mathbb{L}_1, \mathbb{L}_2$ are similar.* The mechanisms for satisfying all constraints irrespective of what objective function they are minimizing are similar. This means analyzing properties on just one of the loss functions should give us an approximate idea of utility offered on others. Hence, we focus most of our attention on \mathbb{L}_0 loss function.
- *No blow up in number of mechanisms for \mathbb{L}_0 .* Given 7 different properties, there are $2^7 = 128$ different combinations that could be requested. Does this mean that

Fig. 2. Heatmaps of constrained mechanisms for $\alpha = 0.62$.

there are over a hundred distinct constrained mechanisms? We show that this is not the case: there are at most four different behaviours that can be observed. Two behaviors correspond to explicit constructions of mechanisms: the (truncated) geometric mechanism (GM) proposed in [3], which corresponds to the unconstrained optimal solution; and a new “explicit fair mechanism” (EM) which simultaneously achieves all the properties that we introduce. In between are two mechanisms which achieve variations of the weak honesty property above, which are found by solving an optimization problem.

- *No significant loss in utility.* The Geometric mechanism obtains the minimal value of the \mathbb{L}_0 loss function, for which we give a closed form in terms of the privacy parameter α . However, our most constrained mechanism (the explicit fair mechanism, EM) is only incrementally more expensive: the loss function value is higher by a factor of approximately $1 + \frac{1}{n}$, which becomes negligible for even moderate n . The costs of the other constrained mechanisms are sandwiched in between.

Consequently, we conclude that the addition of constraints provides significant structure to the space of mechanism design, and comes at very low cost. Given these observations, one may wonder whether there is any material difference in behavior between the constrained and unconstrained mechanisms? This is indeed the case. For example, Fig. 7 shows a quantitative difference between GM and EM for $n = 4$ (chosen to make the results easy to view). The heatmap shows that GM concentrates the probability mass on the two extreme outputs, 0 and n , while EM achieves a more balanced distribution, closer to the leading diagonal (corresponding to a truthful mechanism). If we assume a uniform input distribution, EM reports the true input with probability 0.224, while GM (which maximizes this quantity) achieves 0.238, only marginally higher but with a high skew. A third mechanism with the weak honesty property, WM, sits between the two.

Our experiments further study the implications of using constrained mechanisms, and compare their empirical behavior on a mixture of real and synthetic data. Differences are most apparent for moderate values of n : as n becomes very large, these “end effects” become less significant, and off-the-shelf mechanisms do a good enough job. Thus, we spend most of our effort studying groups corresponding to a moderate number of individuals, up to tens. Arguably, such small groups are most in need of protection, since they

have only a few participants: there is reduced safety in numbers for them.

Outline. We discuss prior work, introduce our model and define useful notions for differential privacy in Section 2. After defining the Linear Programming framework (Section 3), in Section 4, we present constraints that can be added to avoid degeneracy. In Section 5, we revisit the one user/one bit case (Local Differential Privacy), and show that Randomized Response represents a natural convergence of multiple different approaches to privacy. In Section 6, we observe that some existing approaches yield seemingly undesirable results for small groups (with only a few members), which motivates our further study of differentially private mechanisms. Additional properties which constrain the output can be obtained efficiently via solving a constrained optimization problem. We also propose an explicit construction of a mechanism which provably achieves all our proposed properties, and analyze the additional “cost” in terms of various measures of accuracy. Section 7 reports on our experiments on accuracy with synthetic and real data.

2 PRELIMINARIES

2.1 Model and Definitions

Our model assumes a group of n participants, each of whom has some private information which is encoded as a single bit. They share their information with a trusted aggregator, whose aim is to release information about the sum of the values while protecting the privacy of each participant. Although simple, this question is at the heart of all complex analysis and modelling, and demands a comprehensive solution. We simplify the description of the input to just record the true sum of values j , so we have $0 \leq j \leq n$. This captures the case of a count-query over a table D . Our goal is to design a *randomized mechanism* that, given input j produces output i , subject to certain constraints.

Definition 1 (Randomized Mechanism). *A randomized mechanism \mathcal{M} is a mapping $\mathcal{M} : D \Rightarrow R$, where $R = \{0, \dots, n\} = [n]$ is the range of the mechanism. We write $\Pr_{\mathcal{M}}[i|j]$ for the conditional probability that the output $\mathcal{M}(j)$ (on input $j \in D$) is $i \in R$. We will drop the subscript \mathcal{M} in context.*

Our mechanism maps inputs in the range 0 to n to outputs in the same range. While one could allow a different set of outputs, it is most natural to restrict to this range. Consider for example, a downstream analysis step which expects counts to be integers in the range $[n]$: we should ensure that this expectation is met by the result of applying mechanisms. Rather than attempt to map different outputs to this range, it is more direct to build mechanisms that cover this output set. It is therefore natural to represent \mathcal{M} as an $(n + 1) \times (n + 1)$ square matrix \mathcal{P} , where $\mathcal{P}_{i,j} = \Pr[\mathcal{M}(j) = i] = \Pr_{\mathcal{M}}[i|j]$. For brevity, we abbreviate this probability to $\Pr[i|j]$. Note that therefore \mathcal{P} is a *column stochastic matrix*: the entries in each column can be interpreted as probabilities, and sum to 1.

Privacy of a Mechanism. Differential privacy imposes constraints on the probabilities in our mechanism. Specifically, it bounds the ratio of probabilities of seeing the same output for *neighboring inputs* [1]. In our setting, the notion of neighboring is simply that they differ by (at most) one, which

happens when an individual changes their response. Hence, applying the definition, we obtain

Definition 2 (Differentially Private Mechanisms). *Mechanism \mathcal{M} is α -differentially private for $\alpha \in [0, 1]$ if*

$$\forall i, j : \alpha \leq \frac{\Pr[i|j]}{\Pr[i|j+1]} \leq \frac{1}{\alpha}.$$

Here α close to 1 provides a stronger notion of privacy and a tighter constraint on the probabilities, while α close to zero relaxes these constraints. It is common in differential privacy to write $\alpha = \exp(-\epsilon) \approx 1 - \epsilon$, for some $\epsilon > 0$. We adopt the α notation for conciseness, and translate results in terms of ϵ -differential privacy when appropriate. We say a DP constraint is *tight* if the relevant inequality is met with equality.

Utility of a Mechanism. The true test of the utility of a mechanism is the accuracy with which it allows queries to be answered over real data. However, we aim to design mechanisms prior to their application to data, and so we seek a suitable function to evaluate their quality. Since there are many column stochastic matrices that satisfy DP, the problem of finding a mechanism that provides the maximal utility can be framed as an optimization problem. Specifically, we can encode our notion of utility as a penalty function, where we seek to penalize the mechanism for reporting results that are far from the true answer.

Definition 3 (Objective Function Value). *We define the objective function $O_{t,\oplus}(\mathcal{P})$ of a mechanism \mathcal{P} as*

$$O_{t,\oplus}(\mathcal{P}) = \oplus_j \sum_i w_j \Pr[i|j] |i - j|^t,$$

where \oplus is an operator like \sum or \max , and $\sum_j w_j = 1$.

Observe that the weights w_j can be thought of as a prior distribution on the input values j . Then $O_{t,\oplus}(\mathcal{P})$ gives the expected error of the mechanism, when taking its output as the true answer, and $|i - j|^t$ penalizes the extent by which the output was incorrect. When not otherwise stated, we take $w_j = \frac{1}{n+1}$, i.e., a uniform prior over the inputs. Common choices for t in the definition would be $t = 2$, corresponding to a squared error (\mathbb{L}_2 norm), $t = 1$, corresponding to an absolute error (\mathbb{L}_1 norm), and $t = 0$, corresponding to the probability of any wrong answer (\mathbb{L}_0 norm). In what follows, we devote most of our attention to the case \mathbb{L}_0 . We argue that this is an important case: (i) maximizing the probability of reporting the truth is a natural objective in mechanism design; we aim to ensure that the reported answer is the maximum likelihood estimator (MLE) for the true answer, for use in downstream processing; (ii) due to the differential privacy constraints, maximizing the probability of the true answer has the additional effect of making nearby answers likely, as our experiments validate; (iii) our internal study shows that objectives like \mathbb{L}_1 and \mathbb{L}_2 often give pathological results, as seen in Fig. 1. Working with \mathbb{L}_0 gives more robust behavior. We therefore initiate the study of constrained mechanism design for \mathbb{L}_0 , and give some initial results for other objectives. It is convenient to apply a rescaling of the loss function by a factor of $\frac{n+1}{n}$: this sets the cost of a trivial mechanism to 1 (Definition 5). We refer to this rescaled cost as \mathbb{L}_0 , as this corresponds to a scaled

version of O_0 , \sum that sums the probabilities of a wrong answer, and so

$$\mathbb{L}_0(\mathcal{P}) = \frac{n+1}{n} - \frac{\text{trace } \mathcal{P}}{n}. \quad (1)$$

Abusing notation slightly, we also define the objective function, $\mathbb{L}_{0,d} = \frac{n+1}{n} \sum_{i,j:|i-j|\geq d} w_j \Pr[i|j]$ which computes a rescaled sum of probabilities more than d steps off the main diagonal, so that $\mathbb{L}_0 = \mathbb{L}_{0,0}$.

2.2 Prior Work and Existing Mechanisms

We now review the most relevant existing approaches that apply in our setting. The model of differential privacy [1], [4], [5] has received a lot of attention in the decade since it was christened, from a variety of communities including systems [6], machine learning and signal processing [7] and data management [8]. For a more thorough overview of the area, there are several detailed surveys [2], [9], [10].

The most relevant work to our interests is due to Ghosh et al. [3] who study the problem of designing mechanisms optimizing for expected utility. Their contributions are to introduce a linear programming formulation of the problem, and to show that a certain mechanism (denoted GM) emerges as the basis of other optimal mechanisms, discussed in more detail below. Gupte and Sundararajan proved a similar universality result for “minimax” loss functions and uniform weights w_j [11]. They provided a simple test for when a given mechanism can be obtained by first applying GM and then modifying the result (e.g., by randomly sampling from a distribution indexed by the observed output from GM). Subsequent work by Brenner and Nissim shows that such “universally optimal” mechanisms are not possible in general for other computations, such as computing histograms [12]. Other relevant work studies special cases of differential privacy. An important variant is the model of *local differential privacy* (LDP), where users first perturb their input before passing it to an (untrusted) aggregator. That is, each user applies a mechanism for a group of size $n = 1$. LDP is used in Google’s Chrome via the RAPPOR tool to collect browser and system statistics [13], and in Apple’s iOS 10 to collect app usage statistics [14]. In Section 5, we observe that all adaptations of standard DP methods converge on the same idea: a decades-old statistical sampling technique called Randomized Response [15].

The most relevant existing approaches to us are the following:

Mechanisms from Coin-Tossing: Randomized Response. There are many variations of Randomized Response [15]. A canonical form for the case $n = 1$ has the user report the true value of their input bit with probability $p > \frac{1}{2}$, but report the negation of their input with probability $1 - p$. Without loss of generality, we can assume that $p > \frac{1}{2}$. We can write this as a randomized mechanism

$$\mathcal{R} = \begin{bmatrix} p & 1-p \\ 1-p & p \end{bmatrix}. \quad (2)$$

It is immediate that this procedure achieves α -differential privacy for $\alpha = \frac{1-p}{p}$ (see Definition 2). Due to its simplicity and privacy guarantees, randomized response has recently found use in a number of systems, such as RAPPOR [13],

which applies randomized response in conjunction with a Bloom filter to accommodate many possible elements. Geng et al. in [16] give a natural extension of 1 bit randomized response to n -ary data, which reports its input with probability p , else another output is chosen uniformly. This gives low utility for count queries.

Defining Sampling Probabilities: Exponential Mechanism. McSherry and Talwar [17] proposed the Exponential Mechanism as a generic approach to designing mechanisms. Let \mathcal{D} be the domain of input dataset and \mathcal{R} the range of perturbed responses. The crux of the exponential mechanism is in designing a *quality function* $Q : \mathcal{D} \times \mathcal{R} \Rightarrow \mathbb{R}$ so that $Q(d, r)$ measures the desirability of providing output r for input d . The mechanism is then defined by setting

$$\Pr[r \in \mathcal{R}|d] = \exp\left(\frac{\epsilon Q(d,r)}{2s}\right) / \sum_{r' \in \mathcal{R}} \exp\left(\frac{\epsilon Q(d,r')}{2s}\right), \quad (3)$$

where s captures the amount by which changing an individual’s input can alter the output of Q in the worst case. It is proved that this mechanism obtains at least $\exp(-\epsilon)$ -differential privacy. However, although we can use Q to indicate that some outputs are more preferred, it is not possible to modify a given Q to directly enforce the properties that we desire, such as ensuring that the probability of returning the true output is at least as good as that of a uniform distribution (“weak honesty”, (14)).

Rounding Numeric Outputs: Laplace and Geometric Mechanisms. Perhaps the best known differentially private mechanism is the Laplace mechanism, which operates by adding random noise to the true answer from an appropriately scaled Laplace distribution (a continuous exponential distribution symmetric around zero). Note that in order to fit our definition of a mechanism (Definition 1), it will be necessary to round and truncate the output of the mechanism to the range $[n]$. Here, the Laplace mechanism does not easily fit the requirements. Instead, the appropriate method is the discrete analog of the Laplace mechanism, which is the (truncated) Geometric mechanism, introduced by Ghosh et al. [3], who showed that it is the basis for unconstrained mechanisms.

Definition 4. *Range Restricted Geometric Mechanism [3] (GM)*

Let q be the true (unperturbed) result of a count query. The GM responds with $\min(\max(0, q + \delta), n)$, where δ is a noise drawn from a random variable X with a double sided geometric distribution, $\Pr[X = \delta] = \frac{(1-\alpha)^{|\delta|}}{1+\alpha}$ for $\delta \in \mathbb{Z}$.

That is, GM adds noise from a two sided geometric distribution to the query result and remaps all outputs less than 0 onto 0 and greater than n to n . Though GM does not include any zero rows, we observe that each column distribution in GM has spikes at the extreme values, which tend to distort the true distribution quite dramatically, as the next example shows.

Example 1. Consider the case of $n = 2$, corresponding to a group of two individuals, with a moderate setting of the privacy parameter $\alpha = \frac{9}{10}$. For an input of 1 (i.e., one user has a 1, and the other has a 0), we obtain that the probability of seeing an output of 0 is ≈ 0.47 , and the same for an output of 2. Meanwhile, the probably of reporting the true output is ≈ 0.05 —in other words, the chance of seeing the true answer is eighteen times lower than seeing

an incorrect answer. Meanwhile, if the input is 0, then output 0 is returned with probability ≈ 0.53 : so the mechanism is much more likely to report the true answer when it is 0 than when it is 1. As we increase the privacy parameter α closer to 1 (more privacy), the probability of outputs other than 0 and n approaches 0.

As observed in Example 1, an apparent weakness of GM for interpretability is that it can give quite low probabilities for reporting accurate answers. In order to allow more sense to be made of the outputs of the designed mechanisms, we can specify additional constraints to guide the optimization to producing the best interpretable result. This prompts us to define a collection of plausible properties that a mechanism can obey. We will show analytically and empirically that these constraints do not significantly affect the obtained objective function values (i.e., the raw utility), but considerably improve the interpretability of the resulting mechanism. In particular, we demonstrate that it is possible to find a mechanism which achieves all the given properties with only marginal increase in objective function value, and improved interpretability.

3 UNCONSTRAINED MECHANISM DESIGN

A natural starting point is to use optimization tools to find optimal mechanisms. Following [3], the key observation is that the DP requirements can be written as linear constraints over variables which represent the entries of the mechanism. The objective function is also a linear function of these variables. Formally, we define variables $\rho_{i,j}$ for $\Pr[i|j]$, and write

$$\text{minimize: } \sum_{j=0}^n w_j \sum_{i=0}^n |i-j|^p \rho_{i,j} \quad (4)$$

$$\text{subject to: } 0 \leq \rho_{i,j} \leq 1 \quad \forall i, j \in [n] \quad (5)$$

$$\sum_{i=0}^n \rho_{i,j} = 1 \quad \forall j \in [n] \quad (6)$$

$$\rho_{i,j} \geq \alpha \rho_{i,j+1}, \text{ and } \rho_{i,j+1} \geq \alpha \rho_{i,j} \quad \forall i \in [n], j \in [n-1]. \quad (7)$$

The constraints can be understood as follows: (5), (6) ensure that the entries of the matrix are probabilities and each column encodes a probability distribution, i.e., sums to 1. Constraint (7) encodes the differential privacy constraints. Finally, (4) encodes a loss function of Definition 4 for the notion of utility we aim for. We refer to the set of constraints (5), (6) and (7) as BASICDP. The result is a linear program with a quadratic number of variables, and a quadratic number of constraints, each containing at most a linear number of variables. Therefore, solving the resulting LP obtains a mechanism minimizing the given objective function with the desired properties, in time polynomial in n .

Applying this approach yields results like those in Fig. 1. Our studies found that similar undesirable results were found across a range of choices of n , α and loss function. Simple attempts to prevent these outcomes are not effective. For example, we can ensure that no entry is zero by adding a constraint to the LP enforcing this. However, the consequence is that rows which were zero are now set to be the

smallest allowable value, which is unsatisfying. Instead, we propose an additional set of properties to eliminate degeneracy and provide more structure in our solutions.

4 CONSTRAINED MECHANISM DESIGN

We now propose a set of structural properties that help to control the objective function in addition to meeting differential privacy. We believe that these constraints are natural and intuitive and often observed in other mechanisms satisfying differential privacy. We present properties of three types: those which operate on rows of the matrix, those which apply to columns of the matrix, and those which apply to the diagonal.

Row Honesty (RH). A mechanism is *row honest* if

$$\forall i, j. \Pr[i|i] \geq \Pr[i|j]. \quad (8)$$

Row honesty means that a mechanism should have higher probability of reporting i when the input is i than for any other input.

Row Monotone (RM). A mechanism is *row monotone* if

$$\begin{aligned} \forall 1 \leq j \leq i : \Pr[i|j-1] &\leq \Pr[i|j] \\ \forall i \leq j < n : \Pr[i|j+1] &\leq \Pr[i|j]. \end{aligned} \quad (9)$$

This property generalizes row honesty: row monotonicity implies row honesty. It requires that entries in row i are monotone non-increasing as we move away from the diagonal element $\Pr[i|i]$. Note that row monotonicity is independent of differential privacy: we can find mechanisms that achieve DP but are not row monotone, and vice-versa.

Analogous to the row-wise properties, we define monotonicity and honesty along columns also.

Column Honesty (CH). A mechanism is *column honest* if

$$\forall i, j : \Pr[j|j] \geq \Pr[i|j]. \quad (10)$$

Column honesty requires that the mechanism be *honest* enough to report the true answer more often than any individual false answer. As demonstrated by Example 1, GM does not obey column honesty.

Column Monotone (CM). A mechanism is *column monotone* if

$$\begin{aligned} \forall 1 \leq i \leq j : \Pr[i-1|j] &\leq \Pr[i|j] \\ \forall j \leq i < n : \Pr[i+1|j] &\leq \Pr[i|j]. \end{aligned} \quad (11)$$

As in the row-wise case, column monotonicity implies column honesty (but not vice-versa). It captures the property that outputs closer to the true answer should be more likely than those further away.

Fairness (F). A mechanism is *fair* when the probability of reporting the true input is constant, i.e.,

$$\forall i, j : \Pr[i|i] = \Pr[j|j] := y. \quad (12)$$

Example 1 shows that GM is not a fair mechanism. If a mechanism is fair and has row honesty, then all off-diagonal elements are at most y , so the mechanism also satisfies column honesty. Symmetrically, a fair and column honest mechanism is row honest. While this may seem like a restrictive constraint, we observe that mechanisms

proposed in other contexts have this property, such as the staircase mechanism of [16].

Lemma 1. *If a mechanism is required to be fair, then any mechanism that minimizes the objective $O_{0,\Sigma}$ is simultaneously optimal for all settings of weights w_j .*

Proof. Let the diagonal element of the fair mechanism be y . The objective function value is

$$\sum_{j \in [n]} \sum_{i \in [n]} w_j \Pr[i|j](i-j)^0 = \sum_{j \in [n]} w_j(1-y) = 1-y. \quad (13)$$

That is, the value is independent of the w_j s. \square

Weak Honesty (WH). A mechanism satisfies *weak honesty* if

$$\forall i : \Pr[i|i] \geq \frac{1}{n+1}. \quad (14)$$

We can consider this property a weaker version of column honesty, as CH implies WH: for any column j , summing the column honesty property over all rows i we obtain

$$(n+1) \Pr[i|i] = \sum_{i=0}^n \Pr[j|j] \geq \sum_{i=0}^n \Pr[i|j] = 1,$$

so after rearranging, we have $\Pr[i|i] \geq \frac{1}{n+1}$.

Weak honesty ensures that a mechanism reports the true answer with probability at least that of uniform guessing (formalized as the uniform mechanism UM in Definition 5). It also ensures that the mechanism does not have any rows that are all zero (corresponding to outputs with no probability of being produced). GM does not always obey weak honesty, as is shown by Example 1.

The final property we consider is a natural symmetry property (formally, it is that the matrix \mathcal{P} is *centrosymmetric*): *Symmetry (S).* A mechanism is *symmetric* if

$$\forall i, j : \Pr[i|j] = \Pr[n-i|n-j]. \quad (15)$$

Since the input and output domains, and the objective functions are symmetric, it is natural to seek mechanisms which are also symmetric. Our next result shows that symmetry is always achievable without any loss in objective function.

Theorem 1. *Given a mechanism M which meets a subset of properties P from those defined above, we can construct a symmetric mechanism M^* which also satisfies all of P and achieves the same objective function value as M .*

Proof. Our construction to achieve symmetry is simple. Define a matrix M^S from M as $(M^S)_{i,j} = M_{n-i,n-j}$. Then set $M^* = \frac{1}{2}(M + M^S)$. We first observe that M^* is indeed symmetric, since it is equal to

$$\frac{1}{2}(M_{i,j} + M_{n-i,n-j}) = \frac{1}{2}(M_{n-i,n-j} + M_{n-(n-i),n-(n-j)}) = M_{n-i,n-j}^*,$$

as required by (15). The (L_0) objective function value is unchanged since (invoking (1))

$$\text{trace}(M^*) = \frac{1}{2}(\text{trace}(M) + \text{trace}(M^S)) = \text{trace}(M).$$

For the other diagonal properties (fairness and weak honesty), it is immediate that if either of these properties are satisfied by M , then they are also satisfied by M^* . We prove the claim for row properties; the case for column properties is symmetric.

- (i) *Differential privacy:* if we have $\alpha \leq M_{i,j}/M_{i,j+1} \leq 1/\alpha$ for all i, j , then this also holds for $M_{i,j}^S/M_{i,j+1}^S$. Summing both inequalities, and using that $\min(\frac{a}{b}, \frac{c}{d}) \leq \frac{a+c}{b+d} \leq \max(\frac{a}{b}, \frac{c}{d})$, this holds for M^* , hence M^* satisfies differential privacy.
- (ii) *Row monotonicity:* consider a pair i, j with $1 \leq i \leq j$. Then we have $M_{j,i-1} \leq M_{j,i}$ (from (9)). It is also the case that $n-j \leq n-i < n$, which means that $M_{n-j,n-i+1} \leq M_{n-j,n-i}$ (also from (9)). Then $M_{j,i-1}^S \leq M_{j,i}^S$. Combining these two inequalities, we have that $M_{j,i-1}^* \leq M_{j,i}^*$.
- (iii) *Row honesty:* if $\forall i, j. M_{i,i} \geq M_{i,j}$, then $M_{i,i}^S \geq M_{i,j}^S$ also. Summing both inequalities, we obtain $M_{i,i}^* \geq M_{i,j}^*$ as required. \square

Consequences of these Properties. We first argue that these properties all contribute to avoiding the degenerate mechanisms shown above. The (column, row) honesty and monotonicity properties work to prevent the ‘‘spikes’’ observed when a value far from the true input is made excessively likely. The (column) honesty properties do so by preventing a far output being more likely than the true input; the (column) monotonicity properties do so more strongly by ensuring that any further output is no more likely than one that is nearer to the true input. Fairness, column honesty and weak honesty prevent gaps (zero rows): they ensure that the diagonal entry in each row is non-zero, and then the DP requirement ensures that all other entries in the same row must also be non-zero. We next show that there is an efficient procedure to find an optimal constrained mechanism for any $n > 1$.

Theorem 2. *Given any subset of the structural constraints, we can find an optimal (constrained) mechanism which respects these constraints in time polynomial in n .*

Proof. We break the proof into two pieces. First, we argue that given any subset of structural constraints we can create a Linear Program describing it, and second we argue that there exists a mechanism satisfying them all. Observe that all seven properties listed above can be encoded as linear constraints. For example, symmetry is written as

$$\rho_{i,j} = \rho_{n-i,n-j} \quad \forall i, j \in [n],$$

while weak honesty is

$$\rho_{i,i} \geq 1/(n+1).$$

Row monotonicity becomes

$$\begin{aligned} \rho_{j,i-1} &\leq \rho_{j,i} & \forall j \in [n], i < j \\ \rho_{j,i+1} &\leq \rho_{j,i} & \forall i \in [n-1], j < i. \end{aligned}$$

\square

Consequently, we can create a linear program of size polynomial in n , by adding these to the BASICDP constraints (5), (6) and (7) established in Section 3. This shows

the first part of the proof. Next, we show that any such LP is feasible by defining a trivial baseline mechanism:

Definition 5 (Uniform Mechanism, UM). *The uniform mechanism of size n has $\Pr[i|j] = \frac{1}{n+1}$ for all $i, j \in [n]$.*

That is, UM ignores its input and picks an allowable output uniformly at random. It demonstrates that all our properties are (simultaneously) achievable, albeit trivially. By observation, the mechanism is symmetric and fair for any $\alpha' \leq 1$. It meets the inequalities specified for row monotonicity, column monotonicity and weak honesty with equality. UM also satisfies differential privacy for all $\alpha \leq 1$.

Clearly, UM is undesirable from the perspective of providing utility. We easily calculate that the objective function value $O_{0,\sum}$ achieved by UM is $\frac{n}{n+1}$, which is close to the maximum possible value of 1. Note that we chose our definition of the \mathbb{L}_0 function to assign this mechanism a (reweighted) score of 1.

5 CONSTRAINED MECHANISMS: $n = 1$

In this section, we consider an important special case of our problem: where a single user has a single private bit value. This is the limiting case of our setting, corresponding to $n = 1$. It turns out to be an important scenario that has been studied over many decades, as it asks each user to reveal a (noisy) version of their information for subsequent aggregation. We briefly revisit this case in the light of the objectives and properties defined above. The main conclusions we find are that for $n = 1$, all approaches to building DP mechanisms are essentially the same, and trivially obey all our constraints, making this a starting point for our subsequent study.

5.1 Randomized Response

We prove the following result in the Appendix, which can be found on the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TKDE.2019.2912179>:

Theorem 3. *In the one bit (binary) case, Randomized Response is the unique optimal non-trivial α -differentially private mechanism under any objective function $O_{t,\sum}$ when $\alpha \leq w_1/w_0 \leq 1/\alpha$.*

Fact 1. For $p \geq \frac{1}{2}$, Randomized Response satisfies all properties listed in Section 4.

The fact follows immediately by inspection: in the 2×2 case, the mechanism is symmetric (Equation (2)). This entails fairness. All other properties reduce to the condition that $p \geq 1 - p$, i.e., $p \geq \frac{1}{2}$.

5.2 Exponential Mechanism

Theorem 4. *In the one bit (binary) case, the Exponential Mechanism results in an instance of Randomized Response with $p = \frac{\exp(\epsilon/2)}{1+\exp(\epsilon/2)}$.*

Proof. In the binary case, we have $D = R = \{0, 1\}$. Without loss of generality, we can assume that $Q(0, 0) = Q(1, 1) := c$; $Q(1, 0) = Q(0, 1) := w$ (if not, this makes the privacy guarantee loose in one case). We also assume that $c \geq w$, since we should make the true response more likely than

the incorrect response. Then, by definition, $s = c - w$. The resulting mechanism has

$$\begin{aligned} \Pr[0|0] &= \frac{\exp(\epsilon c/2s)}{\exp(\epsilon w/2s) + \exp(\epsilon c/2s)} \\ &= \frac{\exp(-w\epsilon/2s)}{\exp(-w\epsilon/2s)} \frac{\exp(\epsilon c/2s)}{\exp(\epsilon w/2s) + \exp(\epsilon c/2s)} \\ &= \frac{\exp(\epsilon/2)}{1 + \exp(\epsilon/2)}. \end{aligned}$$

Meanwhile, $\Pr[1|0] = 1 - \Pr[0|0] = 1/(1 + \exp(\epsilon/2))$, $\Pr[1|1] = \Pr[0|0]$ and $\Pr[0|1] = \Pr[1|0]$. Consequently, the mechanism is equivalent to \mathcal{R} from (2), and the privacy guarantee is given by $\Pr[0|0]/\Pr[1|0] = \exp(-\epsilon/2)$. \square

Note that this direct application of the exponential mechanism construction actually yields $\exp(-\epsilon/2)$ privacy, stronger than specified, since it does not take full advantage of the additional simple structure of this scenario.

5.3 Geometric Mechanism

Lemma 2. *In the binary case, the Geometric mechanism results in an instance of Randomized Response with $p = \frac{1}{1+\alpha}$.*

Proof. When $n = 1$, we can consider each input separately. On input 0, the output is 0 if $\delta \leq 0$. From properties of the geometric distribution, we obtain

$$\Pr[0|0] = \Pr[X \leq 0] = \frac{1-\alpha}{1+\alpha} \cdot (1 + \alpha + \alpha^2 + \alpha^3 + \dots) = \frac{1}{1+\alpha}.$$

Then, $\Pr[1|0] = \Pr[X > 0] = \frac{\alpha}{1+\alpha}$.

The case for input 1 is symmetric. Hence the claim follows. \square

6 CONSTRAINED MECHANISMS: $n > 1$

For $n > 1$, it is not the case that all mechanisms automatically achieve all our enumerated properties. In this section, we consider mechanisms achieving various combinations of the structural properties.

6.1 The Geometric Mechanism

Next, we revisit the (range restricted) Geometric Mechanism, GM (Definition 4). In Fig. 3, we show the structure of the mechanism, which can be derived by simple calculation from Definition 4. Below, we show that it enjoys a number of special properties. In prior work, Ghosh et al. showed that GM plays an important role, as it can be transformed into an optimal mechanism for different objectives. Here, we argue (proof in the Appendix, available in the online supplemental material) a more direct result: that GM is directly optimal for a uniform objective function.¹

Theorem 5. *GM is the (unique) optimal mechanism satisfying BASICDP under the \mathbb{L}_0 objective function.*

Limitations of GM. Since GM is ‘optimal’ for \mathbb{L}_0 , should we conclude our study here? The answer is no, since GM fails to satisfy many of the desirable properties we identified in

1. Note that, compared to [3], we define mechanisms to enforce differential privacy along rows of \mathcal{P} rather than columns.

$$\begin{pmatrix} x & x\alpha & x\alpha^2 & x\alpha^3 & \cdots & x\alpha^n \\ y\alpha & y & y\alpha & y\alpha^2 & \cdots & y\alpha^{n-1} \\ y\alpha^2 & y\alpha & y & y\alpha & \cdots & y\alpha^{n-2} \\ y\alpha^3 & y\alpha^2 & y\alpha & y & \cdots & y\alpha^{n-3} \\ y\alpha^4 & y\alpha^3 & y\alpha^2 & y\alpha & \cdots & y\alpha^{n-4} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ x\alpha^n & x\alpha^{n-1} & x\alpha^{n-2} & x\alpha^{n-3} & \cdots & x \end{pmatrix}$$

Fig. 3. Structure of GM, where $x = \frac{1}{1+\alpha}$ and $y = \frac{1-\alpha}{1+\alpha}$.

Section 4, and as illustrated in Example 1. We have already observed that GM is not fair, and does not in general satisfy column honesty (or column monotonicity) or weak honesty. Next, we identify parameter settings for when they do hold.

Lemma 3. GM obeys weak honesty iff $n \geq \frac{2\alpha}{1-\alpha}$.

Proof. Weak honesty requires the diagonal elements to all exceed $\frac{1}{n+1}$. Since $y < x$, we focus on y . We require $y \geq \frac{1}{n+1}$ i.e. $\frac{1-\alpha}{1+\alpha} \geq \frac{1}{n+1}$. This reduces to $n+1 \geq \frac{1+\alpha}{1-\alpha}$ giving the requirement $n \geq \frac{2\alpha}{1-\alpha}$. \square

GM satisfies the column monotonicity condition for many i, j pairs. The critical place in the matrix where it can be violated is between the first and second rows (symmetrically, between penultimate and final rows). This corresponds to the problematic behavior of GM to report extreme outputs (0 or n) overly often in the increased privacy regime ($\alpha > \frac{1}{2}$).

Lemma 4. GM achieves column monotonicity iff $\alpha \leq \frac{1}{2}$.

Proof. We require $\Pr[1|1] \leq \Pr[0|1]$, i.e., $y \leq \alpha x$ or $\frac{1-\alpha}{1+\alpha} \leq \frac{\alpha}{1+\alpha}$. This gives the condition $\alpha \leq \frac{1}{2}$. It is straightforward to check that this ensures monotonicity in all other columns. \square

By inspection, GM is always symmetric, and row monotone. The (\mathbb{L}_0) objective function value achieved by GM is

$$\begin{aligned} & \frac{n+1}{n} \left(1 - \frac{(n-1)y + 2x}{n+1} \right) \\ &= \frac{n+1}{n} \left(1 - \frac{n-1}{n+1} \frac{1-\alpha}{1+\alpha} - \frac{2}{(1+\alpha)(n+1)} \right) = \frac{2\alpha}{1+\alpha}. \end{aligned}$$

We next design a different explicit mechanism which achieves more of the desired properties.

6.2 Explicit Fair Mechanism

Although we can achieve any desired combination of properties by solving an appropriate linear program, it is natural to ask whether there is any non-trivial explicit mechanism that achieves properties such as fairness with an objective function score comparable to that of GM. We answer this question in the positive. First, we consider the limits of what can be achieved under fairness. In the case of GM, all DP inequalities are tight. This is not possible when fairness is demanded. A fair mechanism M with all DP inequalities tight would be completely determined: $M_{i,j} = y\alpha^{|i-j|}$ for some y . It is easy to calculate for any such mechanism that there is no setting of y which ensures that all columns sum to 1, a contradiction. Hence, we cannot have a fair mechanism with all DP inequalities tight. Nevertheless, trying to

$$\begin{pmatrix} y & y\alpha & y\alpha^2 & y\alpha^3 & y\alpha^4 & y\alpha^4 & y\alpha^4 & y\alpha^4 \\ y\alpha & y & y\alpha & y\alpha^2 & y\alpha^3 & y\alpha^3 & y\alpha^3 & y\alpha^3 \\ y\alpha & y\alpha & y & y\alpha & y\alpha^2 & y\alpha^3 & y\alpha^3 & y\alpha^3 \\ y\alpha^2 & y\alpha^2 & y\alpha & y & y\alpha & y\alpha^2 & y\alpha^2 & y\alpha^2 \\ y\alpha^2 & y\alpha^2 & y\alpha^2 & y\alpha & y & y\alpha & y\alpha^2 & y\alpha^2 \\ y\alpha^3 & y\alpha^3 & y\alpha^3 & y\alpha^2 & y\alpha & y & y\alpha & y\alpha \\ y\alpha^3 & y\alpha^3 & y\alpha^3 & y\alpha^3 & y\alpha^2 & y\alpha & y & y\alpha \\ y\alpha^4 & y\alpha^4 & y\alpha^4 & y\alpha^4 & y\alpha^3 & y\alpha^2 & y\alpha & y \end{pmatrix}$$

Fig. 4. Explicit fair mechanism for $n = 7$.

achieve tightness provides us with a bound on what can be achieved.

Lemma 5. Let F be a fair mechanism of size $(n+1) \times (n+1)$ with y as the diagonal element. Then $y \leq \frac{1-\alpha}{1+\alpha-2\alpha^{\frac{n}{2}+1}}$.

Proof. There are some slight differences depending on whether we consider odd or even values of n . Without loss of generality, take n even. We will consider a fixed column j . For all i , we are required to have $\Pr[i|i] = y$ for some y . Repeatedly applying the DP inequality, we obtain an upper bound involving y as $\Pr[i|j] \geq y\alpha^{i-j}$ when $j < i$ and $\Pr[i|j] \geq y\alpha^{j-i}$ when $i > j$. Summing these for any given column j and equating to 1 provides an upper bound on y . We get the tightest bound by picking column $j = \frac{n}{2}$. Then $y + 2y \sum_{j=1}^{\frac{n}{2}} \alpha^j \leq 1$, so

$$y \leq \frac{1}{1 + 2 \sum_{j=1}^{\frac{n}{2}} \alpha^j} = \frac{1-\alpha}{1+\alpha-2\alpha^{\frac{n}{2}+1}}. \quad (16)$$

For n large enough, we can neglect the $\alpha^{n/2+1}$ term, and approximate this quantity by $\frac{1-\alpha}{1+\alpha}$. \square

Note that for optimality under an objective function $O_{t,\sum}$, we should make y as large as possible. Hence, any optimal mechanism will have y as close to this value as possible. Indeed, the above proof helps us to design an explicit mechanism EM that achieves fairness. The proof argues that in column $n/2$, the smallest values we can obtain above and below the y entry are αy , $\alpha^2 y$ and so on up to $\alpha^{n/2} y$. Then the sum of these terms is set to 1. All other columns must also sum to 1; a simple way to achieve this is to ensure all columns contain a permutation of the same set of terms. To ensure DP is satisfied, we should arrange these so that row-adjacent entries differ in their power of α by at most one.

Our explicit fair mechanism EM is then defined as follows:

$$\Pr[i|j] = \begin{cases} y\alpha^{|i-j|} & \text{if } |i-j| < \min(j, n-j) \\ y\alpha^{\lfloor \frac{|i-j| + \min(j, n-j)}{2} \rfloor} & \text{otherwise.} \end{cases} \quad (17)$$

Here, y is set to $\frac{1-\alpha}{1+\alpha-2\alpha^{n/2+1}}$, i.e., the value determined in Equation (16). From the proof of Lemmas 1 and (1), we have that the \mathbb{L}_0 score of this mechanism is $\frac{n+1}{n}(1-y)$, as it maximizes y subject to the bound of Lemma 1.

Fig. 4 shows the instantiation of this mechanism for the case $n = 7$. Comparing to GM, we see that the diagonal elements are slightly increased, with the exception of the two corner diagonals, which are decreased. It is tempting to try to obtain the mechanism via the Exponential Mechanism,

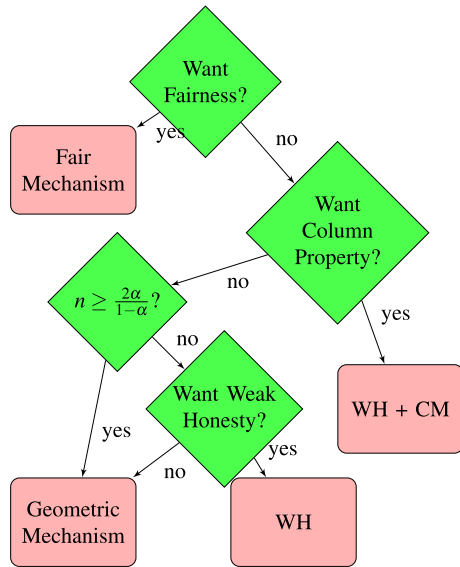


Fig. 5. Flowchart of properties for \mathbb{L}_0 objective ($\alpha > \frac{1}{2}$).

by using a quality function applied to $|i - j|$ similar in form to (17). Note however, that the constant factors of 2 in its definition (3) leads to a considerably weaker result than this explicit construction, equivalent to halving the privacy parameter ϵ . It is easy to check that in the $n = 7$ example, the mechanism is symmetric, and meets all of the properties defined in Section 4. In fact, this is the case for all values of n . The proof (presented in the Appendix, available in the online supplemental material) is rather lengthy and proceeds by considering a number of cases.

Theorem 6. *EM is an optimal mechanism under \mathbb{L}_0 that satisfies all properties listed in Section 4.*

6.3 Comparing Mechanisms

In Section 4, we define 7 different properties, denoted as RH, RM, CH, CM, WH, F, and S. We can seek a mechanism that satisfies any subset of these, suggesting that there are 128 combinations to explore. However, we are able to dramatically reduce this design space with the following analysis based on the \mathbb{L}_0 score function.

First, we have shown by Theorem 6 that EM has the optimal \mathbb{L}_0 score of any fair mechanism and has all other possible properties “for free”. Therefore, for any desired set of properties that include F, we can just use EM.

Second, we have shown by Theorem 5 that GM achieves symmetry and row monotonicity (and hence row honesty) at a cost which is optimal for any mechanism (i.e., BASICDP). Hence for any subset of {S, RM, RH}, it suffices to use GM.

In our experiments (Section 7.1), we show that there are only two remaining behaviors: either we solve the LP for the WH property alone, or we solve the LP for WH and CM properties. Both solutions come with symmetry (S) and row properties RH, RM at no additional cost. However, as noted in Lemma 3, GM satisfies WH when $n \geq \frac{2\alpha}{1-\alpha}$, so in this case, we can use GM. Last, from observations in Section 4, we have that $CM \Rightarrow CH \Rightarrow WH$, so any demand that requires any of these properties (and not F) can be satisfied by WM also. But in the weak privacy case that $\alpha \leq \frac{1}{2}$, GM has these properties, and so subsumes WM.

Property	GM	WM	EM	UM
Symmetry (S)	Y	Y	Y	Y
Row Monotone (RM)	Y	Y	Y	Y
Column Monotone (CM)	—	—	Y	Y
Fairness (F)	N	N	Y	Y
Weak Honesty (WH)	—	Y	Y	Y
\mathbb{L}_0	$\frac{2\alpha}{1+\alpha}$	$\geq \frac{2\alpha}{1+\alpha}$	$\approx \frac{2\alpha}{1+\alpha} \cdot \frac{n+1}{n}$	1

Fig. 6. Properties of named mechanisms.

To summarize this reasoning, in the case that $\alpha \leq \frac{1}{2}$, there are only two competitive mechanisms: EM if fairness is required, and GM for all other cases. When $\alpha > \frac{1}{2}$, things are a little more complicated, so we show a flowchart in Fig. 5: from 128 possibilities, there are only four distinct approaches to consider (two explicit mechanisms, and two solutions to an LP with different constraints), and the choice is determined primarily by whether the mechanism is required to satisfy fairness, column properties, weak honesty, or none. We also consider the baseline method UM for comparison. We present a summary of these four named mechanisms in Fig. 6: the explicit GM, UM and EM, and WM which is found by solving an LP. We write ‘—’ for a property when this depends on the setting of the parameters (discussed in the relevant section). We see that EM has a very similar objective function value \mathbb{L}_0 (recalling that we are trying to minimize this value), and all the properties considered so far. We do not have a closed form for the \mathbb{L}_0 score of WM, as it is found by solving the LP; however it is no less than that for GM (since GM satisfies a subset of the required properties of WM), and no more than that of EM (since EM satisfies all properties).

At this point, we might ask how different are these mechanisms in practice—perhaps they are all rather similar? Fig. 7 shows this is not the case for a small group size ($n = 4$). For a moderate value of the privacy parameter $\alpha = 0.9$, it presents the three non-trivial mechanisms using a heatmap to highlight where the large entries are. We immediately see that EM concentrates probability mass along a uniform diagonal (as required by fairness). Both GM and WM tend to favor extreme outputs (0 or 4 in this example) whatever the input, although GM is very skewed in this regard while WM is more uniform in allowing non-extreme outputs.

Last, we check that what we are doing is not a trivial modification of known mechanisms. Prior work [3], [11] showed how optimal unconstrained mechanisms can be derived from GM by transformations. Gupte and Sundarajan give a simple test: a mechanism \mathcal{P} can be derived from GM iff every set of three adjacent entries in the mechanism satisfy

$$(\Pr[i|j] - \alpha \Pr[i|j - 1]) \geq \alpha(\Pr[i|j + 1] - \alpha \Pr[i|j]).$$

We applied this test to mechanisms WM and verified that this condition is indeed violated for $n > 1$. For EM, this condition is automatically broken for all $n > 1$: we have $\Pr[2|0] = \Pr[2|1] = y\alpha$, while $\Pr[2|2] = y$. Then the condition is

$$y\alpha(1 - \alpha) \geq y\alpha(1 - \alpha^2) \equiv 1 \geq (1 + \alpha),$$

which is always false for $\alpha > 0$. Hence, these mechanisms are not derivable from GM.

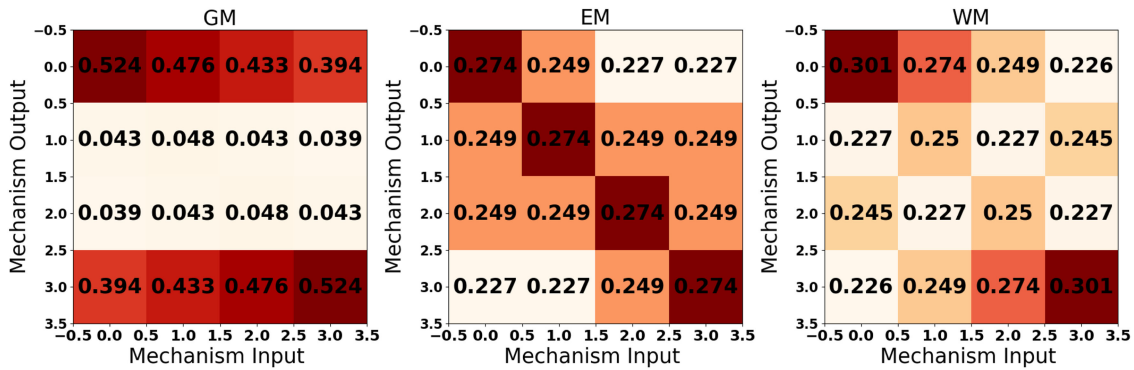
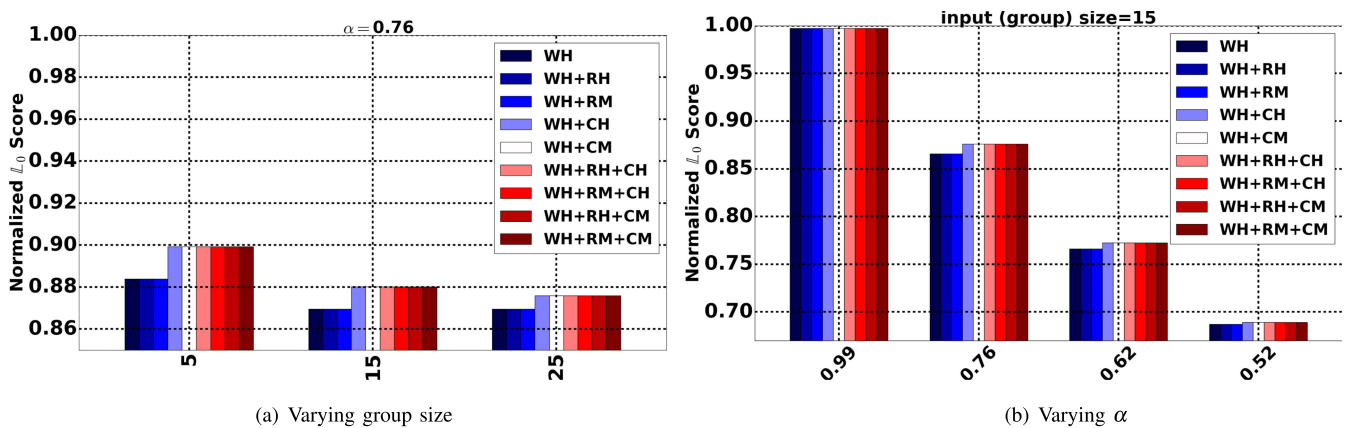
Fig. 7. Heatmaps for GM, EM, and WM with $n = 4$.

Fig. 8. Combinations of properties with weak honesty.

7 EXPERIMENTAL STUDY

The purpose of our experimental study is two-fold. In Section 7.1, we substantiate our earlier claims about properties of mechanisms satisfying weak honesty (but not fairness). In what follows, we look at other measures of utility of these mechanisms, to understand their robustness.

Default Experimental Settings. All experiments in this work were implemented in Python, making use of the standard library NumPy to handle the linear algebraic calculations, and PyLPSolve [18] to solve the generated LPs. Evaluation was made on a commodity machine running Linux. We omit detailed timing measurements, as the time to solve the LPs generated was negligible (sub-second).

Experimental Setting. We considered a variety of settings of parameter α (typical values chosen are $\{\frac{1}{2}, \frac{2}{3}, \frac{10}{11}, \frac{99}{100}\}$) and group size n (ranging from 2 up to hundreds).

7.1 \mathbb{L}_0 Objective Function

Our first experiment analyzes the effect of weak honesty combined with other properties drawn from $\{\text{CH}, \text{CM}, \text{RH}, \text{RM}\}$, including the empty set. There are 9 meaningful combinations of properties to ask for, which we write as $\{\emptyset, \text{RH}, \text{RM}, \text{CH}, \text{CM}, \text{RH}+\text{CH}, \text{RH}+\text{CM}, \text{RM}+\text{CH}, \text{RM}+\text{CM}\}$ —other combinations reduce to these, since RM implies RH, and CM implies CH.

As discussed in Section 6.3, there are cases when the solution found by solving the LP has cost $\frac{2\alpha}{1+\alpha}$ and is identical to GM: these are when $n \geq \frac{2\alpha}{1-\alpha}$ and only row-wise properties are requested, consistent with Fig. 5. This is borne out in

Fig. 8: we see that when WH alone is requested, or in combination with only row properties (RH or RM) we get a lower \mathbb{L}_0 value than when any column properties (CH or CM) are requested. Fig. 8a shows the case for different values of n . When $n > \frac{2\alpha}{1-\alpha}$, which is 6.33 in this example (where $\alpha = 0.76$), the cost of WH alone is $\frac{2\alpha}{1+\alpha} = 0.864$, the cost of GM. For large α (Fig. 8b), the cost of all combinations of WH are the same, and identical to the cost of EM; as α is decreased, we see two behaviors, where the lower \mathbb{L}_0 cost is that of GM. We confirmed this behavior for a wide range of n and α values. From now on, we use WM to refer to the mechanism with WH, RM and CM properties.

The relationship between the \mathbb{L}_0 scores for the three mechanisms is further clarified in Fig. 9. The plots show the \mathbb{L}_0 scores of GM, WM, EM and UM for different values of α . In Fig. 9a, $\alpha = \frac{2}{3}$ so the threshold $\frac{2\alpha}{1-\alpha} = 4$. Then GM satisfies WH for the whole range of n values shown, so WM converges on GM, while EM has a higher (but decreasing) cost. For Fig. 9b, $\alpha = \frac{10}{11}$ so the threshold is 20. Indeed, we see that the cost of WM converges with GM at $n = 20$. Last, in Fig. 9c, the threshold of 198 is far above the range of n values shown, so WM does not converge on GM here. Rather, for this high value of α , the y value for EM is above $\frac{1}{n+1}$ for all n : so in this case EM has weak honesty, and the cost of WM remains the same as that of the optimal fair EM.

7.2 Experiments on Real Data

We make use of the UCI Adult dataset, a workhorse for privacy experiments [19]. Our instance of the dataset contains demographic information on 32 K adults with 15 columns

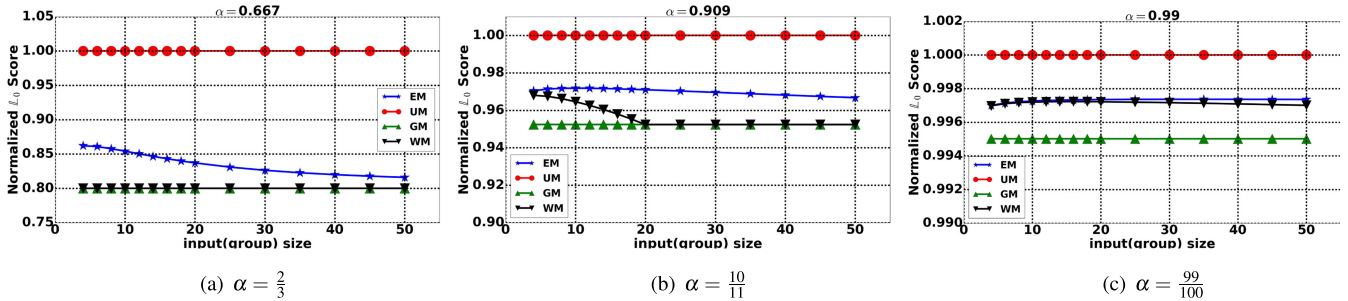


Fig. 9. Final groups of mechanisms with distinct behaviors.

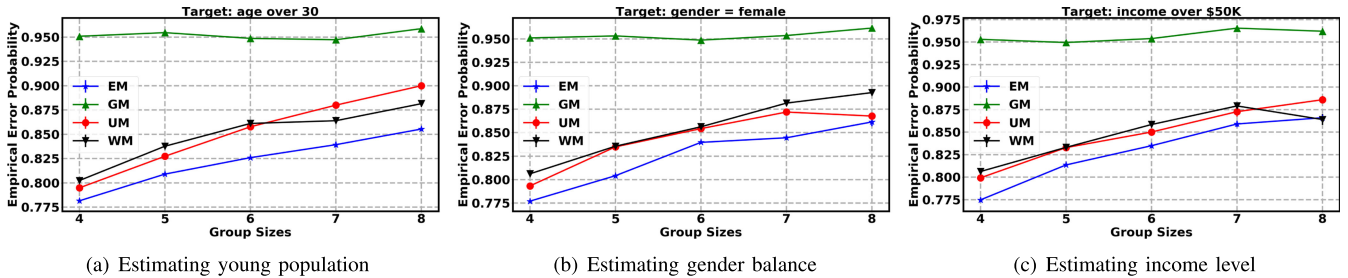


Fig. 10. Empirical error probability on adult dataset for $\alpha = 0.9$.

listing age, job type, education, relationship status, gender, and (binary) income level. We created three binary targets, treated as sensitive: income level (high/low), gender (male/female), and young (age over/under 30). To form small groups, we gathered the rows (corresponding to individuals) arbitrarily into groups of a desired size.

Fig. 10 shows results for the \mathbb{L}_0 objective, that is, where we focus on the fraction of times the mechanism reports an incorrect answer, as a function of group size. Specifically, we count the number of groups whose noisy count for each target attribute is not equal to their true count. We expect this quantity to be fairly high, as it measures how often our mechanism is honest, i.e., returns the true input. Other experiments (not shown) computed the corresponding probability for returning an answer that is close to the true one, e.g., off by at most one, and showed similar patterns. The plot includes error bars from 50 repetitions of this process to show 1 standard error.

Observe first that the performance of UM is essentially independent of the input data: the chance of it picking the correct answer is always $1/(n + 1)$ for a group of size n , and indeed we see this behavior (up to random variation). We would hope that our optimized mechanisms can outperform this trivial method. Perhaps surprisingly, on this data GM does appreciably worse. This highlights the limitations of GM. In this data, the common inputs are around the middle of the group size (i.e., typically close to $n/2$). It is on these inputs that GM does poorly, and only does well for inputs that are 0 or n , which happen to be rare in this dataset (in other words, the data distribution does not match the prior for which GM is optimal). The condition of weak honesty is not sufficient to improve significantly over random guessing: for this data, we see that WM tracks UM quite closely. It is only the most constrained mechanism that fares better on this evaluation metric for this data: EM which achieves fairness gives the best probability of returning the unperturbed input. In corresponding experiments with higher values of α in the range 0.9 to 0.99, corresponding to

the strongest privacy guarantees adopted in prior work on differential privacy, there is not much to choose between EM and WM, and it gets even harder to show substantial improvement over uniform guessing. In order to understand the behaviors of the mechanisms further, we next consider synthetic data, where we can directly control the data skewness within groups.

7.3 Experiments on Synthetic Data

In our experiments with synthetic data we generate a population of 10,000 individuals each with a private bit and divide them into small groups of the same size, n . Each individual has the same probability p of having their bit be one, so the distribution within each group is Binomial. Hence, the expected count for each group is pn . Our experiments vary the parameters p , n and α .

$\mathbb{L}_{0,1}$ Error. Our experiments so far have used the target objective function \mathbb{L}_0 to evaluate the quality of the mechanism. This is sufficient to distinguish the different mechanisms, but all mechanisms achieve a score which is still quite close to 1, obtained by uniform guessing. To better demonstrate the usefulness of the obtained mechanisms, we use other functions to evaluate their accuracy. Fig. 11 uses the related measure of $\mathbb{L}_{0,1}$ i.e., the fraction of groups which output a value differing from their true answer by more than 1, as we vary data distribution (determined by p), group size n , and privacy parameter α . We stress that though we use $\mathbb{L}_{0,1}$ for evaluation, we continue to use mechanisms designed for minimizing the \mathbb{L}_0 error. Each subplot in the figure represents a configuration of $\langle \alpha, n \rangle$, describing how $\mathbb{L}_{0,1}$ error changes with input distribution parameter p . Each experiment is repeated 30 times and we observe that the results have very small variance.

It is apparent that the shape of the input distribution has a pronounced effect on the quality of the output. We confirm that GM can do well when the input is very biased (p close to 0 or 1), which generates more instances with extreme input values. However, when the input is more spread across the

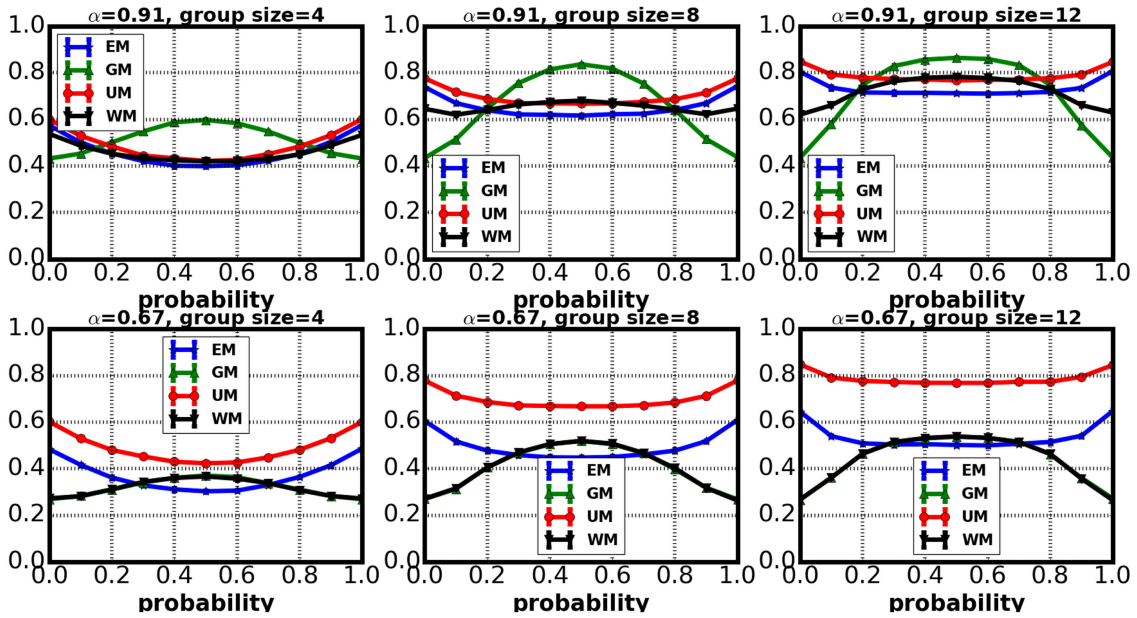


Fig. 11. $L_{0,1}$ score for Binomial data, for $n = \{4, 8, 12\}$ and $\alpha = \{0.91, 0.67\}$.

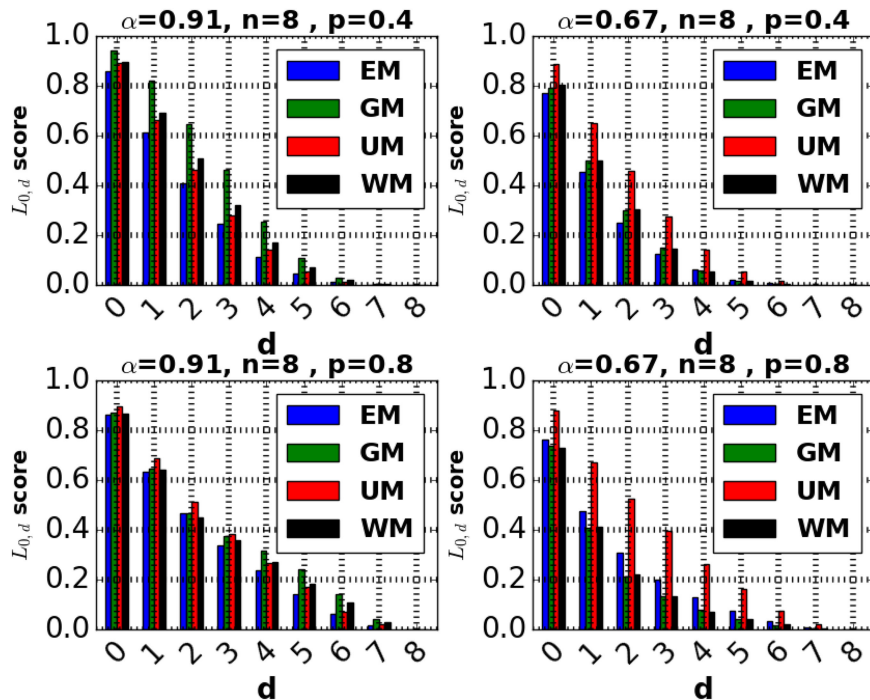


Fig. 12. Histograms of $L_{0,d}$ scores for Binomial data.

input space, the more constrained mechanisms consistently give better results. For higher α , the constrained methods have similar behavior, and improve only slightly over UM (while GM is often worse than uniform). Enforcing fairness tends to make EM less sensitive to the input distribution, except when the input is an extreme value (0 or n). When α is lower (second row), the overall scale of error decreases and WM and GM converge, as noted previously.

$L_{0,d}$ Error. In the previous experiment, we fixed $d = 1$ and evaluated our mechanisms for variety of input distributions. Next we vary d while holding input size and input distributions steady, and compute $L_{0,d}$ error. Fig. 12 plots the fraction of population reporting a value that is more than d steps away

from the true answer for various d values with $n = 8$. This captures the probability mass in the tail of each mechanism.

In the top row, we use a more proportionate input distribution. Here, EM outperforms all other mechanisms, sometimes by a substantial fraction. Interestingly, the margin between EM and GM only increases with larger d . Once again we see that for higher α values, use of GM can yield accuracy worse than mere random guessing. For lower α 's GM's accuracy increases dramatically but still remains worse than EM's.

In the bottom row, the input distribution is more skewed, which tends to favor GM. However, EM does not do substantially worse than GM even for this biased input

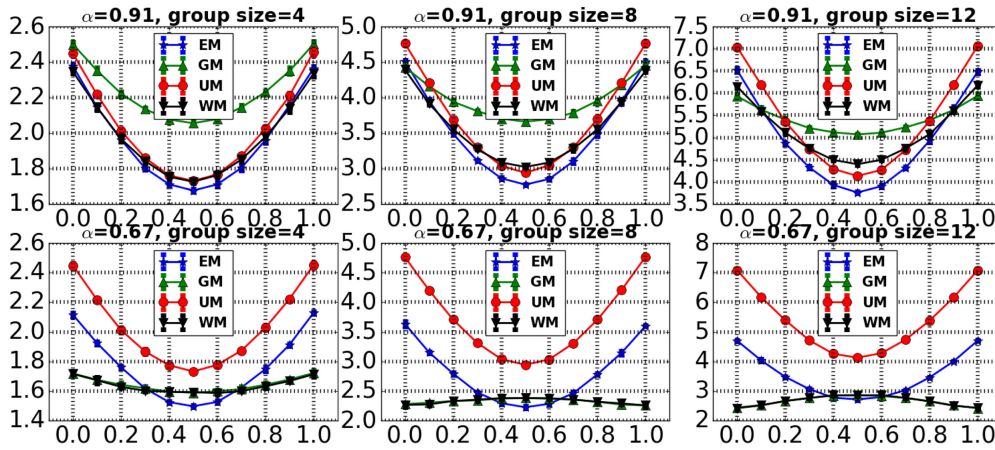


Fig. 13. Root mean square error plots for Binomial data.

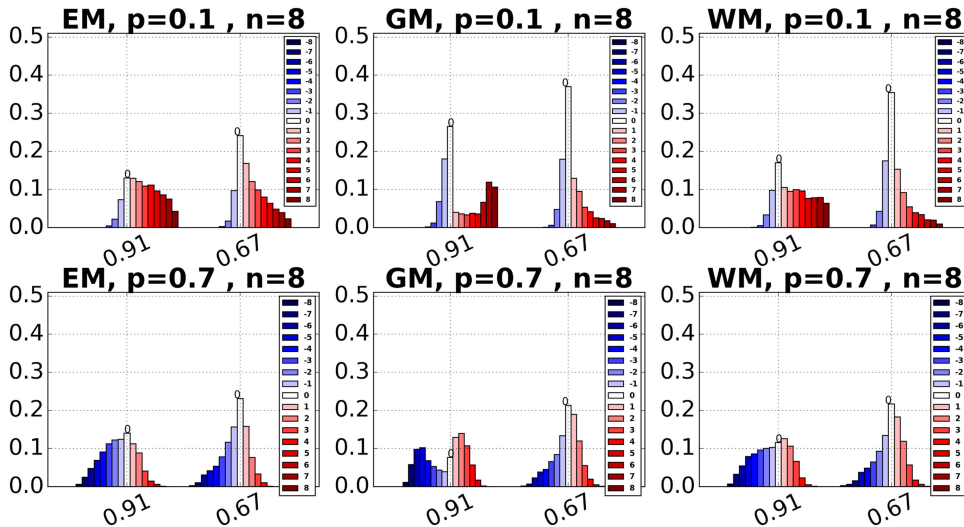


Fig. 14. Error histograms on group size 8 for $p = 0.1$ and $p = 0.7$, with $\alpha = \{0.91, 0.67\}$.

distribution. The intermediate mechanism found by WM tends to fall between GM and EM. We observed similar behavior for other values of n .

Root Mean Square Error (RMSE). Our next set of experiments compute the RMSE error (a measure of variance and bias of a mechanism) of estimates from small groups. Note that none of our mechanisms are designed to optimize this metric, but we can nevertheless use it as a measure of the overall spread of error. Fig. 13 shows plots with error bars showing one standard deviation from 30 repetitions.

As seen in previous experiments, a more symmetric input distribution (p closer to 0.5) tends to be easier for most mechanisms—although we see cases where GM finds this more difficult. Increasing the group size increases the RMSE, as there is a wider range of possible outputs, and the constraints ensure that there is some probability of producing each possible output. Yet again, we see that increasing α tends to make GM less competitive and find many cases where GM is worse than random guessing (UM). The interesting case may be for fairly high privacy requirements ($\alpha = 0.91$), where we observe that EM tends to give lower error across all group sizes and input distributions.

Error Histograms. As previously discussed, the measure of error probability gives some insight into the difference

between mechanisms, but holds them to a high standard. This probability is high for all mechanisms, as we do not expect them to give the exact correct answer. To see the spread of error from another perspective, we plot error histograms for our mechanisms—for a given input distribution, how often is the response correct, how often is it an overestimate by one, and so on. For example, when an input of 1 is reported as 0, the error is -1 . Fig. 14 shows the error histograms for a representative group size $\{8\}$ with $p = 0.1$ and $p = 0.7$ for two extreme α values. For each case, we show error histograms for the three mechanisms EM, GM, and WM.

In the $p = 0.1^2$ case, the input does not permit significant underestimation (most true answers are small). All mechanisms are more likely to give zero error. This is enforced by the fairness (for EM) or weak honesty (WM) properties. For GM, we see that for $\alpha = 0.91$, there is a second peak corresponding to an output of n . So it tends to have a larger error when it does not output the true answer. We observe that the column monotonicity properties of EM and WM tend to force a smoother error distribution.

$2. p$ is a parameter for producing synthetic input data introduced in Section 7.3.

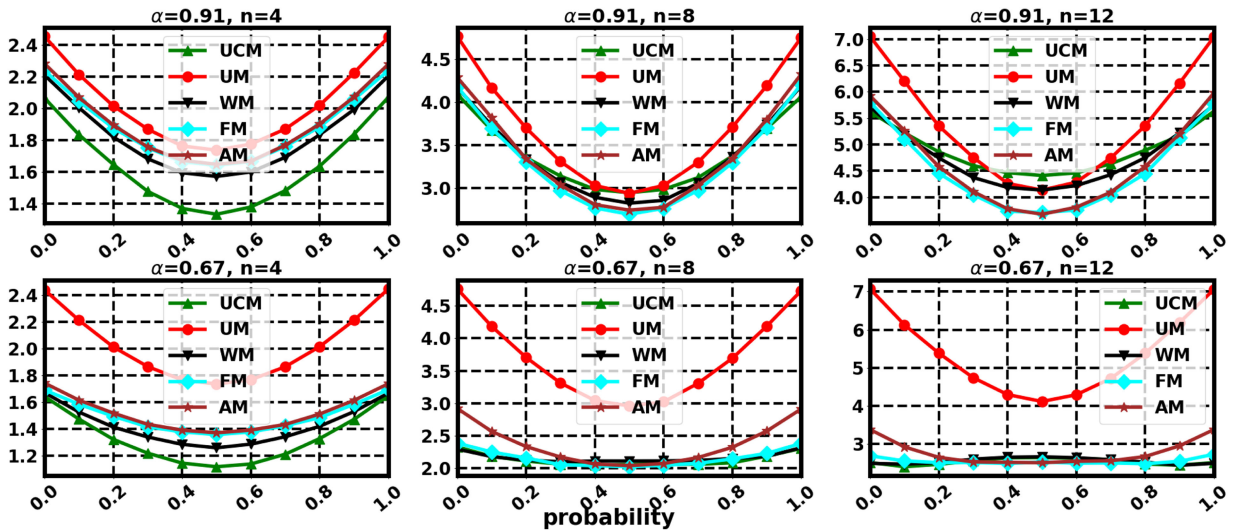


Fig. 15. Root mean square error plots on Binomial data for \mathbb{L}_1 objective function mechanisms.

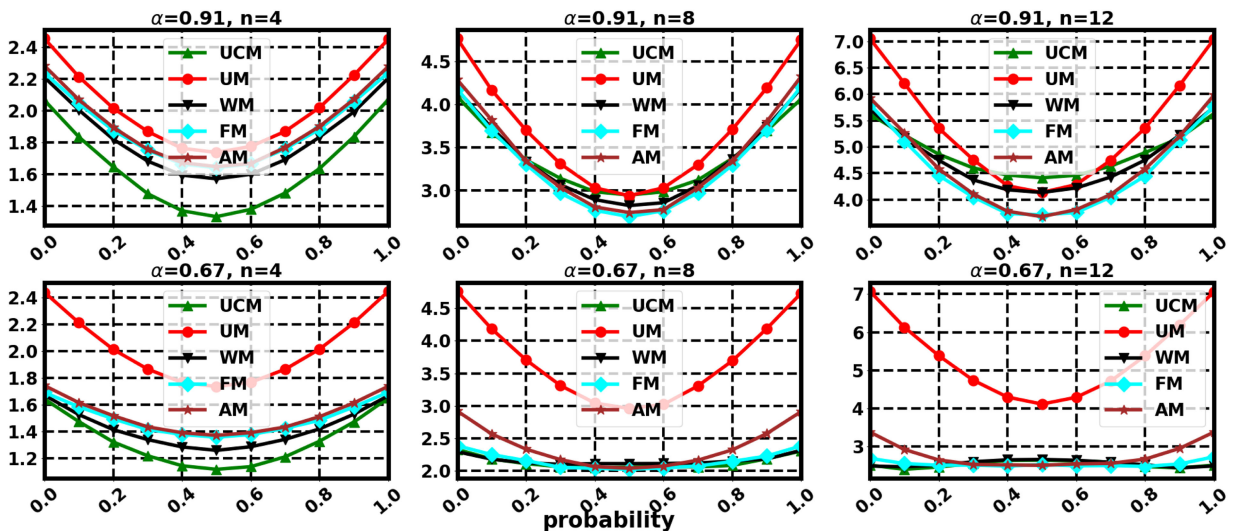


Fig. 16. Root mean square error plots for Binomial data for \mathbb{L}_2 objective function mechanisms.

Some similar behavior is observed for $p = 0.7$. Here, the support of the input distribution is broader, and hence so is the support of the error distribution. We still observe that GM tends to have a bimodal error distribution for high α , with a dip around zero error. As α decreases, the mechanisms become more similar, in particular WM tends to look more like GM (as we have seen, they converge to the same mechanism for $n \geq \frac{2\alpha}{1-\alpha}$). We have observed similar trends for other n values.

7.4 \mathbb{L}_1 and \mathbb{L}_2 Objective Functions

We have already seen that unconstrained mechanisms for \mathbb{L}_1 and \mathbb{L}_2 can have pathological outcomes. In this section, we return to these objective functions, and study their behavior under the imposition of conditions. In contrast to the \mathbb{L}_0 case, we observed that the number of distinct mechanisms obtained under selection of different subsets of conditions was quite large. In order to constrain the number of mechanisms under consideration, we restrict our attention to a small number of options: enforcing Weak Honesty (denoted WM) or Fairness (denoted FM) only; or requiring

either no properties at all (the unconstrained mechanism, UCM), or all properties simultaneously (the all properties mechanism, AM). We also compare to the trivial uniform mechanism (UM) for calibration. Among these four options, we expect UCM to obtain the lowest error since it can directly optimize the target function, with the comensurate disadvantages discussed previously.

Root Mean Square Error (MSE). Figs. 15 and 16 show plots for the root mean square error on binomially distributed input data, similar to Fig. 13. For this measure of accuracy, UCM provides among the best results. However, for small groups, we would tend to prefer WM, since it provides a similar level of accuracy while avoiding the degerate behaviors. AM performs well when the input distribution is close to the symmetric case ($p = 0.5$), but has weaker results when the input is more skewed (smaller or larger p values). FM is observed to do better as the group size increases.

$\mathbb{L}_{1,d}$ and $\mathbb{L}_{2,d}$ Functions. Figs. 17 and 18 show plots of the errors for the functions $\mathbb{L}_{1,d}$ and $\mathbb{L}_{2,d}$ respectively, similar to Fig. 12 for $\mathbb{L}_{0,d}$. When $p = 0.1$, most groups have sums close to 0. For larger α 's and smaller d 's, UCM performs worst.

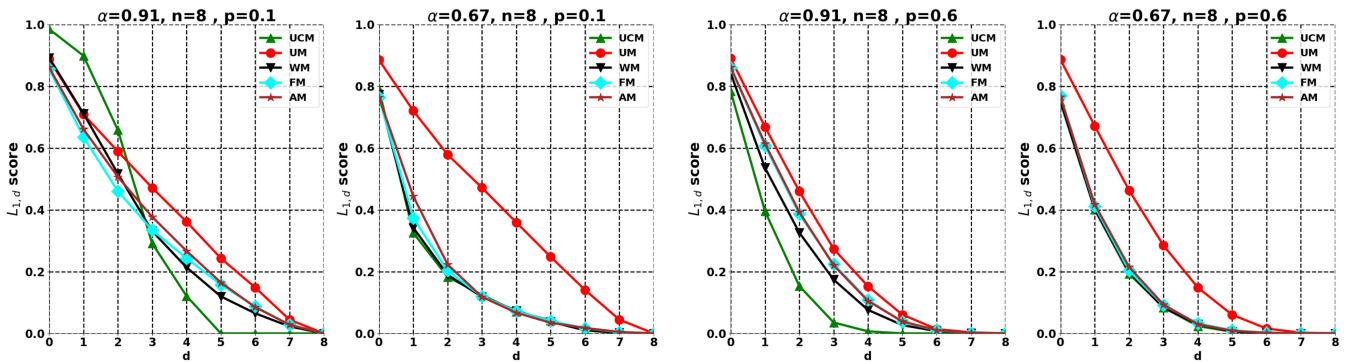


Fig. 17. Line plots for $\mathbb{L}_{1,d}$ scores for Binomial data ($p = 0.1$ and $p = 0.6$).

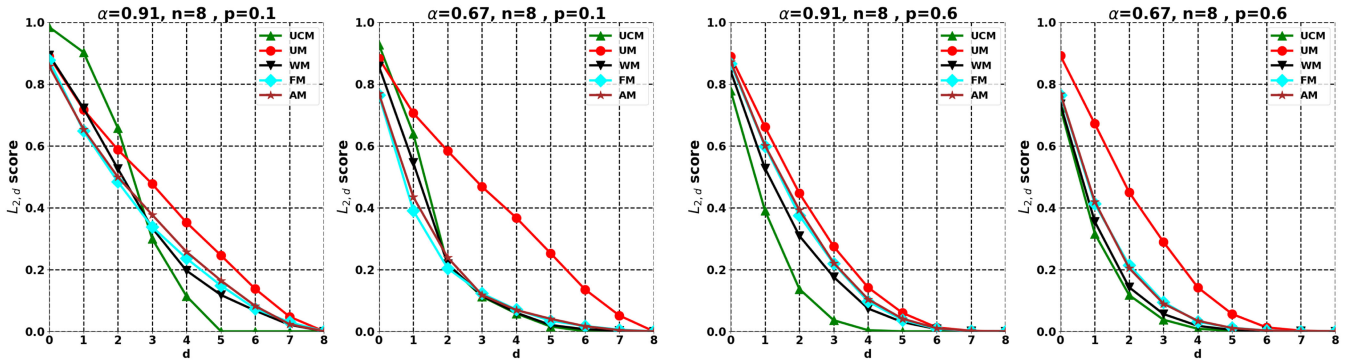


Fig. 18. Line plots for $\mathbb{L}_{2,d}$ scores for Binomial data ($p = 0.1$ and $p = 0.6$).

This situation is reversed as d increases. That is, the mechanism reduces the probability mass that is far from the true answer, at the expense of increasing the mass close to the true answer but not equal to it. It tends to map most inputs to outputs close to $\lfloor \frac{n}{2} \rfloor$. UCM is then the preferred mechanism for more balanced distributions ($p = 0.6$). AM, FM and WM all behave quite similarly to each other, and their lines almost overlap for smaller α . In summary, AM, FM are slightly preferable for skewed input distributions and strong privacy requirements, whereas WM is suitable in general for distributions with less bias.

8 CONCLUDING REMARKS

Further Applications. While materializing basic count statistics directly supports a number of standard queries, it is natural to consider more complex analyses of data. For example, we may wish to materialize machine learning models of the data. The count-based mechanisms we have described can naturally be used to instantiate simple predictive Bayesian models, such as the Naive Bayes classifier or low-degree Bayesian networks, via the Chow-Liu tree [20]. In more detail, these models require the creation of marginal distributions on combinations of attributes in the data. However, these marginal distributions can be obtained from counts: the (empirical) probability $\Pr[Y|X]$ can be obtained by estimating the count of individuals in a group who have the properties of both Y and X , and dividing by the count who have property X . The confidence in these estimates can be increased by averaging over a large number of groups. The procedure meets differential privacy, since each user locates in exactly one cell of a joint distribution, and so we can apply the parallel composition theorem

of differential privacy to reason about the overall privacy guarantee.

Initial experiments indicate that this direction is feasible. We observed that the results from GM and WM were similar in terms of quality, while those for EM were much weaker. This indicates the weak honesty property is a reasonable one to request, but adding fairness tended to distort the results too much. The quality of the results were still much weaker than without privacy, unless the parameter α was set very low, corresponding to a low level of privacy. This suggests that for applications like this, novel mechanisms more directly tuned to the desired application are needed, since the generic count-based mechanisms we consider do not give good results under tighter privacy requirements.

Summary and Future Work. We have proposed and studied several structural properties for privacy preserving mechanisms for count queries. We show how any combination of desired properties can be provided optimally under \mathbb{L}_0 by one of a few distinct mechanisms. Our experiments show that the “optimal” GM often displays the undesirable property of tending to output extreme values (0 or n). In practice, this means it is often not the mechanism of choice, particularly when α is large (above 0.7), but can be acceptable for smaller privacy parameters. EM and WM are quite different in structure, but are often similar in performance.

It is natural to consider other possible properties—for example, one could imagine taking a version of the DP constraint applied to columns of the mechanism (in addition to the rows): this would enforce that the ratio of probabilities between neighboring *outputs* is bounded, as well as that of neighboring *inputs*. The next logical direction is to provide a deeper study of mechanisms with various properties using \mathbb{L}_1 or \mathbb{L}_2 as objective function, building on our

empirical observations. It will be interesting to study tailor-made linear programming mechanisms that aim to optimize other queries such as range queries.

ACKNOWLEDGMENTS

This work is supported in part by The Alan Turing Institute under EPSRC grant EP/N510129/1, Marie Curie Career Integration Grant 618202, European Research Council Grant 647557, an AT&T Labs VURI award, and a Warwick Collaborative Postgraduate Research Scholarship.

REFERENCES

- [1] C. Dwork, "Differential privacy," in *Proc. Int. Colloq. Automata Languages Program.*, 2006. [Online]. Available: http://dx.doi.org/10.1007/11787006_1
- [2] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Found. Trends Theoretical Comput. Sci.*, vol. 9, no. 3/4, pp. 211–407, 2014. [Online]. Available: <http://dx.doi.org/10.1561/04000000042>
- [3] A. Ghosh, T. Roughgarden, and M. Sundararajan, "Universally utility-maximizing privacy mechanisms," in *Proc. Annu. ACM Symp. Theory Comput.*, 2009. [Online]. Available: <http://doi.acm.org/10.1145/1536414.1536464>
- [4] C. Dwork, F. McSherry, K. Nissim, and A. D. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proc. 3rd Conf. Theory Cryptography*, 2006, pp. 265–284.
- [5] C. Dwork and K. Nissim, "Privacy-preserving datamining on vertically partitioned databases," in *Proc. Annu. Int. Cryptology Conf.*, 2004, pp. 528–544.
- [6] E. Shi, H. Chan, E. Rieffel, R. Chow, and D. Song, "Privacy-preserving aggregation of time-series data," in *Proc. Netw. Distrib. Syst. Security Symp.*, 2011.
- [7] A. D. Sarwate and K. Chaudhuri, "Signal processing and machine learning with differential privacy: Algorithms and challenges for continuous data," *IEEE Signal Process. Mag.*, vol. 30, no. 5, pp. 86–94, Sep. 2013.
- [8] Y. Yang, Z. Zhang, G. Miklau, M. Winslett, and X. Xiao, "Differential privacy in data publication and analysis," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2012, pp. 601–606.
- [9] C. Dwork, "The promise of differential privacy. A tutorial on algorithmic techniques," in *Proc. IEEE Symp. Found. Comput. Sci.*, 2011. [Online]. Available: <http://research.microsoft.com/apps/pubs/default.aspx?id=155617>
- [10] A. Machanavajhala, X. He, and M. Hay, "Differential privacy in the wild (tutorial)," in *VLDB*, 2016. [Online]. Available: http://vldb2016.persistent.com/differential_privacy_in_the_wild.php
- [11] M. Gupte and M. Sundararajan, "Universally optimal privacy mechanisms for minimax agents," in *Proc. ACM SIGMOD-SIGACT-SIGART Symp. Principles Database Syst.*, 2010, pp. 135–146. [Online]. Available: <http://doi.acm.org/10.1145/1807085.1807105>
- [12] H. Brenner and K. Nissim, "Impossibility of differentially private universally optimal mechanisms," in *Proc. 51th Annu. IEEE Symp. Found. Comput. Sci.*, 2010, pp. 71–80.
- [13] U. Erlingsson, V. Pihur, and A. Korolova, "RAPPOR: Randomized aggregatable privacy-preserving ordinal response," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2014. [Online]. Available: <http://doi.acm.org/10.1145/2660267.2660348>
- [14] C. Zibreg, 2016. [Online]. Available: www.idownloadblog.com/2016/06/25/differential-privacy-overview/
- [15] A. Chaudhuri and R. Mukerjee, *Randomized Response: Theory and Techniques*. New York, NY, USA: Marcel Dekker, 1988.
- [16] Q. Geng, P. Kairouz, S. Oh, and P. Viswanath, "The staircase mechanism in differential privacy," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 7, pp. 1176–1184, Oct. 2015.
- [17] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *Proc. IEEE Symp. Found. Comput. Sci.*, 2007. [Online]. Available: <https://www.microsoft.com/en-us/research/publication/mechanism-design-via-differential-privacy/>
- [18] PyLPsolve—An object oriented wrapper for the Ipsolve, 2010. [Online]. Available: www.stat.washington.edu/~hoytak/code/pylpsolve/

- [19] C. Blake and C. Merz, "UCI machine learning repository," 1998. [Online]. Available: <https://archive.ics.uci.edu/ml/datasets/Adult>
- [20] C. Chow and C. Liu, "Approximating discrete probability distributions with dependence trees," *IEEE Trans. Inf. Theory*, vol. 14, no. 3, pp. 462–467, May 1968.



Graham Cormode is a professor in computer science with the University of Warwick, United Kingdom, where he works on research topics in data management, privacy, and big data analysis. Previously, he was a principal member of technical staff with AT&T Labs-Research. He is a University Liaison director with the Alan Turing Institute, and in 2017 he was the co-recipient of the Adams Prize for mathematics for his work on statistical analysis of Big Data.



Tejas Kulkarni received the master's degree from the Indian Institute of Technology, Madras. He is working toward the PhD degree at the University of Warwick, with a focus on differential privacy.



Divesh Srivastava received the BTech degree from the Indian Institute of Technology, Bombay, and the PhD degree from the University of Wisconsin, Madison. He is head of the Database Research Department, AT&T Labs-Research. He has served as associate editor-in-chief of the *IEEE Transactions on Knowledge and Data Engineering*, and the PC co-chair of many conferences, including ICDE 2019. His research interests span a variety of topics in data management. He is a fellow of the ACM, the vice president of the VLDB Endowment, and an associate editor of the *ACM Transactions on Data Science*.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.