Congenial Differential Privacy under Mandated Disclosure

Ruobin Gong
Rutgers University
Piscataway, New Jersey
ruobin.gong@rutgers.edu

Xiao-Li Meng
Harvard University
Cambridge, Massachusetts
meng@stat.harvard.edu

ABSTRACT

Differentially private data releases are often required to satisfy a set of external constraints that reflect the legal, ethical, and logical mandates to which the data curator is obligated. The enforcement of constraints, when treated as post-processing, adds an extra phase in the production of privatized data. It is well understood in the theory of multi-phase processing that congeniality, a form of procedural compatibility between phases, is a prerequisite for the end users to straightforwardly obtain statistically valid results. Congenial differential privacy is theoretically principled, which facilitates transparency and intelligibility of the mechanism that would otherwise be undermined by ad-hoc post-processing procedures. We advocate for the systematic integration of mandated disclosure into the design of the privacy mechanism via standard probabilistic conditioning on the invariant margins. Conditioning automatically renders congeniality because any extra post-processing phase becomes unnecessary. We provide both initial theoretical guarantees and a Markov chain algorithm for our proposal. We also discuss intriguing theoretical issues that arise in comparing congenital differential privacy and optimization-based post-processing, as well as directions for further research.

KEYWORDS

Belief Function, conditioning, invariants, post-processing, statistical intelligibility, uncongeniality, Monte Carlo

ACM Reference Format:
Ruobin Gong and Xiao-Li Meng. 2020. Congenial Differential Privacy under Mandated Disclosure. In *Proceedings of the 2020 ACM-IMS Foundations of Data Science Conference (FODS '20), October 19–20, 2020, Virtual Event, USA.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3412815.3416892

1 PRIVACY AS DATA PROCESSING

1.1 A blurry yet essential picture

The curation and dissemination of large-scale datasets benefits science and society by supplying factual knowledge to assist discoveries, policy decisions, and promote transparency of information. As more data become accessible to more entities, however, the unobstructed access to information collected from individuals poses the risk of infringing on their privacy. Differential privacy is a mathematical concept that quantifies the extent of disclosure of confidential information in a database. It enjoys several advantages over its previous counterparts in statistical disclosure limitation. Most important of all, especially to data analysts who wish to conduct statistical inference on privatized data releases, is the transparency of the algorithm. The mechanism through which the privatized data release is generated can be spelled out in full, with its statistical properties fully understood. This enables analysts to incorporate it as part of a model, hence permitting the statistical validity of the resulting inference [12].

The protection of confidential data with differential privacy relies on the careful design of a probabilistic mechanism, one that can veil the microscopic identities of individual respondents while preserving the macroscopic aspects of the data with high fidelity. The probabilistic nature of the mechanism is necessary, and enables a tradeoff as such to be made [4]. Typically, a differentially private mechanism injects a random perturbation into an otherwise deterministic query to be applied to the confidential database. One can say, then, that differentially private releases are "blurry" versions of the confidential data, just the same way a skilled Impressionist painter captures the essence of a pond of waterlilies without sketching out the contour of every petal and leaf.

When randomness is involved, however, certain truthful aspects of the data is invariably compromised, no matter how well-designed the privacy algorithm may be. Imagine if a picture of waterlilies was commissioned, not by an art collector, but by a botanist whose sole purpose is to study the structural formation of the plant, such as the exact length and width of its petals. She would be terribly disappointed at the Impressionist painting, even if it was the work of Claude Monet himself!

Circumstances in practice dictates that aspects of the data release may be deemed as unfit to be tampered with. These are usually key statistics reflecting the fundamental purpose of data collection, as required by law, policy, or other external constraints as put forth by the stakeholders. The data curator is mandated to disclose these statistics accurately at any expense, while at the same time shielding the remainder part of the data release with a veil of privacy. This poses a challenge to the design of the privacy mechanism subject to mandated disclosure. The central question is, how to integrate data privatization, an inherently random act, with the mandated disclosure of partial yet deterministic information, while maintaining both the logical consistency of the overall data release and the quality of the privacy protection.

1.2 Congenial privacy

In this work we conceptualize the privacy mechanism as one of the many phases of data processing. The concept of congeniality, or

FODS '20, October 19–20, 2020, Virtual Event, USA

ACM ISBN 978-1-4503-8103-1/20/10…$15.00
https://doi.org/10.1145/3412815.3416892

rather *uncongeniality* [19], is then relevant. The theory of uncongeniality was developed for investigating a seemingly paradoxical phenomenon discovered by researchers dealing with imputations for the U.S. Decennial Census and similar public-use data files. Fay [7] and Kott [17] found that one can have inconsistent variance estimation for multiple imputation inference [20], even when both the imputation model and analysis procedure are valid. The issue turned out to be a mathematical incompatibility between the imputation model and the analysis procedure, even if neither is incompatible with the underlying model that generates the data. In other words, there is no probabilistically coherent model that can simultaneously imply the imputation model and the analysis procedure, a situation termed uncongeniality by Meng [19]. To make matters worse, the imputation model, such as adopted by the Census Bureau, is typically not disclosed to the analyst of the imputed data, or at least not fully (e.g., due to confidential information used to help better predict the missing values). The lack of transparency makes it impossible for the analyst to correct for the uncongeniality, and worse, even to realize the problem.

The framework of uncongeniality was later generalized to the multivariate setting [22] and to general multi-phase processing [3], which covers the current application by the same overarching principles. Two properties are critical for good privacy practice: transparency and congeniality. When the protection of privacy must observe mandated constraints, our proposal is to use conditional distributions as derived from the original unconstrained differential privacy mechanism conditioning on invariant margins. This approach achieves both automatically. Transparency is automatic, because the conditional distribution is determined by the original unconstrained distribution and the invariant constraints, both fully disclosed by design. Congeniality stipulates the use of a *single coherent probabilistic model* to ensure both differential privacy and mandated disclosure. This requirement is automatically satisfied when we use the conditional distribution derived from the original differential privacy mechanism, restricted solely to obey the mandated disclosure. A third advantage is that our proposal does not need any additional choice of procedural ingredients, such as projection distance, which is required by optimization-based post-processing such as adopted by the Census TopDown algorithm [1].

There is, however, no free lunch. The first price we pay for congenial privacy is computational. Sampling from a distribution truncated on some space, as determined by the invariant margins, is generally not trivial, especially when the truncated region is of irregular shape. The second price is that we may pay more privacy loss budget than necessary, since the budget designed for the original mechanism depends on the sensitivity of the query measure on the unconstrained space, which is larger than that for the constrained space. When deriving the appropriate class of conditional distributions for the constrained mechanism, the new privacy loss budget should ideally be calibrated directly according to the query behavior on the constrained space, as opposed to be inherited from the unconstrained mechanism, which ensures a likely overly conservative level of privacy protection for the entire space. When the analytical complexity and computational requirements for the two approaches of budget calibration are similar, we certainly recommend the former.

1.3 The mechanism of differential privacy

Let $x = (x_1, \ldots, x_n)$ denote a database consisting of n individuals, and $\mathcal{X}$ the space on which it is defined. A query function $s : \mathcal{X} \to \mathbb{R}^d$ embodies the knowledge contained in the database that stakeholders – scientists, policy makers and the general public – would like to learn. What determines the value of $s(x)$ is of course x, or equivalently all its component x_i values, corresponding to individual respondents included in the database. It is precisely these individuals records, or the x_i's, that are the subject of privacy protection. How can the data curator say useful things about $s(x)$, while saying barely anything about each of the x_i's?

The mechanism that can instill privacy into the curator's release appeals to randomness. A random query function $M : \mathcal{X} \to \mathbb{R}^d$ is said to satisfy ϵ-differential privacy [5], if for all pairs of neighboring datasets $(x, x') \in \mathcal{X} \times \mathcal{X}$, we have that

$$P(M(x) \in B) \leq \exp(\epsilon) P\left(M\left(x'\right) \in B\right) \tag{1.1}$$

for all Borel-measurable sets $B \in \mathscr{B}\left(\mathbb{R}^d\right)$. In this work, the term *neighboring datasets* means that x and x' differ by exactly one individual's record, i.e. for some $j = 1, \ldots, n$, $x_j \neq x'_j$ but for all $i \neq j$, $x_i = x'_i$. Write $\mathrm{d}(x, x') = 1$ if x and x' are neighbors. This concept of neighbor is employed in the definition of ϵ-bounded differential privacy [1], and is distinct from the original formulation which defined neighbors as a pair that differ from each other by the addition or deletion of a single record. There are many ways to design an ϵ-differentially private algorithm M, among which the most widely known and implemented are the Laplace and the Double Geometric algorithms [5, 8], both are additive ϵ-differentially private mechanisms.

Definition 1.1 (Laplace mechanism). Let $s : \mathcal{X} \to \mathbb{R}^d$ be a deterministic query function. The Laplace mechanism is given by

$$M(x) := s(x) + (U_1, \ldots, U_d), \tag{1.2}$$

where U_i's are independent zero-mean Laplacian random variables each with dispersion parameter $\epsilon^{-1}\nabla(s)$, and

$$\nabla(s) = \sup_{(x,x') \in \mathcal{X} \times \mathcal{X}} \left\{ \left\| s(x) - s\left(x'\right) \right\| : \mathrm{d}\left(x, x'\right) = 1 \right\}$$

is the global sensitivity of s. When the database consists of binary records in an unrestricted domain, and s is the counting or histogram query, $\nabla(s) = 1$.

Definition 1.2 (Double Geometric mechanism). A random query mechanism M is called the Double Geometric mechanism if it has the same functional form as (1.2), with U_i random variables defined on the integers with probability mass function

$$p_i(u \mid \epsilon) = \frac{1-a}{1+a} \cdot a^{|u|}, \tag{1.3}$$

where the parameter $a = a(\epsilon, s) = \exp(-\epsilon/\nabla(s))$.

Note that the definition of either the Laplace or the Double Geometric mechanism presents not just one, but a collection of mechanisms that can be written as $\{M_\epsilon\}$, indexed by $\epsilon > 0$ the *privacy loss budget* allocated to the mechanism in question. When regarded as a sequence of statistical procedures, ϵ serves as an indicator of the statistical quality of the output (with larger ϵ for higher quality) that can be used to offer interpretation and to guide its own choice. This point will be immediately useful in Section 1.4.

1.4 Statistical intelligibility

An important reason that differential privacy is embraced by the statistics community is that it defines privacy in the language of probability, induced by the mechanism that injects randomness in the data release. An impactful consequence is that the distributional specification of the mechanism can be made fully transparent without sabotaging the promised protection. This opens the door for systematic analysis of the statistical property of mechanisms, which is in turn crucial to the accurate interpretation of statistical inference from privatized data releases [10]. The clarity both in definition and in implementation makes up the statistical intelligibility of differential privacy as a data processing procedure.

We discuss the statistical interpretation of the privacy mechanism, which is what served as inspiration for the conditioning approach to construct the invariant-respecting mechanism in the first place. The degree of protection exerted by a privacy mechanism on the confidential database is seen as a calculated limit on the statistical knowledge it is able to supply, as a function of the privacy loss budget allotted to the mechanism. Compare quantities

$$\pi\left(x_i = \omega\right) \quad \text{and} \quad \pi\left(x_i = \omega \mid M_\epsilon\left(x\right) \in B\right), \tag{1.4}$$

which are the analyst's prior probability about the value of the ith entry of the dataset, versus her posterior probability if an ϵ-differentially private query released a report B. Thus, the statistical meaning of M_ϵ can be explained as follows.

Theorem 1.3. *Let $x = (\cdots, x_i, \cdots) \in \mathcal{X}$ be the database, and $\{M_\epsilon : \mathcal{X} \to \mathbb{R}^d, \epsilon > 0\}$ a class of ϵ-differentially private procedures operating on x. Then for every $B \in \mathscr{B}\left(\mathbb{R}^d\right)$, $\epsilon > 0$ and every prior probability π the analyst harbors about x_i,*

$$\pi\left(x_i = \omega \mid M_\epsilon\left(x\right) \in B\right) \in \left[\exp\left(-\epsilon\right)\pi\left(x_i = \omega\right),\ \exp\left(\epsilon\right)\pi\left(x_i = \omega\right)\right]. \tag{1.5}$$

Proof. The posterior probability $\pi\left(x_i = \omega \mid M_\epsilon\left(x\right) \in B\right)$ can be written as

$$\frac{P\left(M_\epsilon\left(x\right) \in B \mid x_i = \omega\right)\pi\left(x_i = \omega\right)}{\sum_{\omega'} P\left(M_\epsilon\left(x\right) \in B \mid x_i = \omega'\right)\pi\left(x_i = \omega'\right)}.$$

The result then follows immediately from the fact that M_ϵ is ϵ-private, which means that for any $B \in \mathscr{B}\left(\mathbb{R}^d\right)$,

$$\exp\left(-\epsilon\right) \le \frac{P\left(M_\epsilon\left(x\right) \in B \mid x_i = \omega'\right)}{P\left(M_\epsilon\left(x\right) \in B \mid x_i = \omega\right)} \le \exp\left(\epsilon\right).$$

□

Theorem 1.3 says that, any release generated by an ϵ-differentially private procedure sharpens the analyst's knowledge about x_i by at most a factor of $\exp(\epsilon)$. This interpretation provides a direct link between the differential privacy promise and the actual posterior risk of disclosure due to the release of the random query M_ϵ.

Recall that the definition of the Laplace and the Double Geometric mechanisms are both well-defined for any $\epsilon > 0$. However, in the limiting case of $\epsilon \to 0$, i.e. the privacy loss budget becomes increasingly restrictive, both algorithms amount to adding noise with increasingly large variance to the confidential query. At $\epsilon = 0$, neither mechanism remains well-defined since the distributions of the noise component become improper due to the infinite cardinality of their respective domains. Nevertheless, the definition of ϵ-differential privacy allows for the expression with $\epsilon = 0$. A mechanism is 0-differentially private if one cannot gain any discriminatory knowledge from its release about the underlying database whatsoever. In other words, the analyst's knowledge about the individual state of x_i must remain the same as her prior. This notion can be explained consistently with Theorem 1.3, by observing that

$$\lim_{\epsilon \to 0} \pi\left(x_i = \omega \mid M_\epsilon\left(x\right) \in B\right) = \pi\left(x_i = \omega\right), \tag{1.6}$$

where the limit is implied by (1.5). This inspires the following deliberate construction of M_0 as a 0-differentially private procedure.

Definition 1.4 (0-differentially private procedure). For $\{M_\epsilon\}$ a class of ϵ-differentially private procedures well-defined for $\epsilon > 0$ but not $\epsilon = 0$, define M_0 as the 0-differentially private procedure such that for every prior probability π,

$$\pi\left(x_i = \omega \mid M_0\left(x\right) \in B\right) = \pi\left(x_i = \omega\right),\ \forall B \in \mathscr{B}\left(\mathbb{R}^d\right). \tag{1.7}$$

Definition 1.4 grants conceptual continuity to M_0, a perfectly meaningful object in the privacy sense but lacking statistical intelligibility from the mechanistic point of view. For practical purposes, M_0 should be taken to mean the *suppression* procedure, which supplies *vacuous* knowledge to the analyst about the state of affairs of the database. Theoretically, the meaning of M_0 as a probabilistic mechanism cannot be supplied by ordinary probabilities, because any probability specification represents a set of specific knowledge about relative frequencies of any pair of states. However, its meaning can be quantified precisely in the more general framework of *imprecise probability*, as Section 5 will discuss.

2 CONSTRUCTING CONGENIAL PRIVACY

2.1 Privacy with invariants

While privacy protection is called for, the data curator may be simultaneously mandated to disclose certain aspects of the data as they are exactly observed, without subjecting them to any privacy protection. This collection of information is referred to as *invariant information*, or *invariants*. In practice, invariants are often defined according to a set of exact statistics calculated based on the confidential database [2]. For example, suppose s is the histogram query which tabulates the population residing in each county of the state of New Jersey from the 2020 U.S. Census. When producing a differentially private version of the histogram, the Census Bureau is constitutionally mandated to report the total population of each state as enumerated. This means that the privatized histogram $M(x^*)$ must possess the same total population size as $s(x^*)$, where x^* the confidential Census microdata; or in notation, $\|M\left(x^*\right)\| = \|s\left(x^*\right)\|$.

Suppose that the Double Geometric mechanism is to be applied to the histogram query s. Due to the random nature of the perturbations, a single realization of the mechanism will with high probability produce $\|M\left(x^*\right)\| \neq \|s\left(x^*\right)\|$. Furthermore $M\left(x^*\right)$ has a positive probability of consisting negative cell counts, which is logically impossible of the confidential query $s\left(x^*\right)$. The challenge of privacy preservation under mandated disclosure is thus to find an alternative mechanism, say $\tilde{M}$, such that every realization of

$\tilde{M}$ meets all the requirement of mandated information disclosure, while preserving the promise of differential privacy.

Let $\mathcal{X}^* \subset \mathcal{X}$ be the set of x's that obey the given invariants. In turn, $\mathcal{X}^*$ defines the set of values that the query must satisfy as

$$\mathcal{S}^* = \left\{ s(x) \in \mathbb{R}^d : x \in \mathcal{X}^* \right\}. \tag{2.1}$$

Note that implicitly, $\mathcal{S}^* = \mathcal{S}^*(x^*)$ is a set-valued function of the confidential dataset x^*, because the invariant knowledge we intend to impose on the private release is implied by x^*.

A random mechanism $\tilde{M}$ satisfies the mandated disclosure if

$$\tilde{M}(x) \in \mathcal{S}^*, \quad \forall x \in \mathcal{X}^*. \tag{2.2}$$

That is, whenever applied to a database conformal to the mandated disclosure, with mathematical certainty $\tilde{M}$ is also conformal to the mandated disclosure. The size and complexity of the restricted $\mathcal{S}^*$ (and $\mathcal{X}^*$) relative to their original spaces are crucial to the overall extent to which privacy of the residual information in the database can be protected, a point we will revisit in Section 5.1.

There may be many ways to construct a random mechanism $\tilde{M}$, but all are not equally desirable. We argue that $\tilde{M}$ should be constructed in a principled manner, and a constructive way to achieve that is to use conditional distributions of unconstrained privacy mechanisms. The resulting mechanism can be easily tuned to retain its differential privacy promise, while ensuring its congenial integration into the data processing pipeline, preserving the statistical intelligibility of its releases.

2.2 Imposing invariants via conditioning

Let M be a valid ϵ-differentially private mechanism, which generates outputs that typically do not obey the invariant requirement $M(x) \in \mathcal{S}^*$, even if $x \in \mathcal{X}^*$. A natural idea to force the requirement onto M is via conditioning. Define a modified privatization mechanism M^*, such that the probability distribution it induces is the same as the conditional distribution of M subject to the constraint that $M(x) \in \mathcal{S}^*$. For what's next, we'll use the notation $Z \overset{L}{=} W$ to mean Z and W are identically distributed, and $M(x) \mid M(x) \in \mathcal{S}^*$ denotes a well-specified conditional distribution $P(M(x) \mid M(x) \in \mathcal{S}^*)$. Also assume for now $P(M(x) \in \mathcal{S}^*) > 0$. We have the following theorem.

Theorem 2.1. *Let $x^* \in \mathcal{X}$ be the confidential database, and $\mathcal{X}^* \subset \mathcal{X}$ the invariant subset to which x^* conforms. The deterministic function $s : \mathcal{X} \to \mathbb{R}^d$ is a query, and the implied $\mathcal{S}^* \in \mathscr{B}\left(\mathbb{R}^d\right)$ is defined by* (2.1)*. Let M be an ϵ-differentially private mechanism based on s, and M^* be a constrained mechanism such that*

$$M^*(x) \overset{L}{=} M(x) \mid M(x) \in \mathcal{S}^*. \tag{2.3}$$

Then for all invariant-conforming pairs of datasets that are k-neighbors, i.e. $(x, x') \in \mathcal{X}^* \times \mathcal{X}^*$ such that $\mathsf{d}(x, x') = k$, there exists a real-valued $\gamma \in [-1, 1]$ such that for all $B \in \mathscr{B}\left(\mathbb{R}^d\right)$,

$$P\left(M^*(x) \in B\right) \leq \exp\left((1+\gamma)\, k\epsilon\right) P\left(M^*\left(x'\right) \in B\right). \tag{2.4}$$

Proof. For a pair of k-neighboring and $\mathcal{S}^*$-conforming datasets (x, x') and any $B \in \mathscr{B}\left(\mathbb{R}^d\right)$,

$$\begin{aligned}\frac{P(M^*(x) \in B)}{P(M^*(x') \in B)} &= \frac{P(M(x) \in B \mid M(x) \in \mathcal{S}^*)}{P(M(x') \in B \mid M(x') \in \mathcal{S}^*)} \\ &= \frac{P(M(x) \in B \cap \mathcal{S}^*)}{P(M(x') \in B \cap \mathcal{S}^*)} \cdot \frac{P(M(x') \in \mathcal{S}^*)}{P(M(x) \in \mathcal{S}^*)}.\end{aligned}$$

Clearly each of the last two ratios above is bounded above by $\exp(k\epsilon)$ and below by $\exp(-k\epsilon)$ because M is ϵ-differentially private. Consequently, if we let

$$\gamma^* = \frac{1}{k\epsilon} \log\left[\max_{\substack{(x,x') \in \mathcal{X}^* \times \mathcal{X}^* \\ \mathsf{d}(x,x')=k}} \frac{P(M(x') \in \mathcal{S}^*)}{P(M(x) \in \mathcal{S}^*)} \right], \tag{2.5}$$

then $\gamma^* \in [0, 1]$ and it is a known constant because $\mathcal{S}^*$ is public. Then, (2.4) holds for any $\gamma \in [\gamma^*, 1]$, and certainly for $\gamma = 1$. □

Our proof above might create an impression that only $\gamma \geq 0$ is permissible, but Sections 4 and 5.1 will supply two examples both with $\gamma < 0$. Negative γ may sound paradoxical, for it seems to suggest that better privacy protection can be achieved by disclosing some information. However, we must be mindful that differential privacy is not about protecting privacy in absolute terms. Rather, it is about controlling the *additional* disclosure risk from releasing the privatized data to the users (or hackers), relative to what they know before the release. The presence or absence of mandated invariants would *ex ante* constitute two different bodies of knowledge, hence any additional privacy protection would carry different interpretations too. We also emphasize that Theorem 2.1 generalizes to cases where $P(M(x) \in \mathcal{S}^*) = 0$, such as when it is a linear subspace of $\mathbb{R}^d$ and the privacy mechanism is continuous. The proof is a bit more involved in order to properly define $P(M(x) \mid M(x) \in \mathcal{S}^*)$, which we will discuss in future work. However, this complication is not a concern for discrete privatization mechanisms, such as within the Census TopDown algorithm [1].

Theorem 2.1 holds broadly for arbitrary kinds of ϵ-differentially private mechanisms, as well as any deterministic invariant information about either the database or the query function that can be expressed in a set form. It lends itself to the same kind of posterior interpretation enjoyed by unconstrained differentially private mechanisms. Specifically, if M^*_ϵ is the constrained differentially private procedure constructed based on the unconstrained procedure M_ϵ, according to the specifications of Theorem 2.1, then for all $x \in \mathcal{X}^*$ such that $\exists x' \in \mathcal{X}^*$ so that $\mathsf{d}(x, x') = 1$, and $\forall B \in \mathscr{B}(\mathcal{S}^*)$, the analyst's posterior probability $\pi\left(x_i = \omega \mid x \in \mathcal{X}^*, M^*_\epsilon(x) \in B\right)$ is bounded in between

$$\Big[\exp\left(-(1+\gamma)\,\epsilon\right) \pi\left(x_i = \omega \mid x \in \mathcal{X}^*\right), \\ \exp\left((1+\gamma)\,\epsilon\right) \pi\left(x_i = \omega \mid x \in \mathcal{X}^*\right) \Big].$$

This an interval that bears structural resemblance to (1.5), thanks to the conditional nature of M^* which allows for statistical information from privacy mechanisms, constrained or otherwise, to be interpreted in the same (hence congenial) way.

The definition of ϵ-differential privacy has the property that, if a mechanism M is ϵ_1-differentially private, then for all $\epsilon_2 \geq \epsilon_1$, M is also ϵ_2-differentially private. If the invariant information does not *substantially disrupt* the neighboring structure of the sample space of the database, a notion we will make precise later in Section 5, what Theorem 2.1 says is that enforcing the invariant $\mathcal{X}^*$ onto the unconstrained mechanism M via conditioning costs $(1+\gamma)$ times – and *at most twice* since γ can always be set to 1 – the privacy loss budget allotted to M. When a more cost-effective value of γ is hard to determine, a simplest way to ensure privacy loss budget ϵ for M^* is to use a budget of $\epsilon_0 = \epsilon/2$ for the unconstrained mechanism M to begin with; see Section 3.

2.3 An Monte Carlo Implementation

Let p denote the probability distribution, either a mass function or a density function, induced by the unconstrained differentially private algorithm M which depends on the confidential query s^*. Further denote p^* to be the corresponding conditional distribution of p constrained on the invariant set $\mathcal{S}^*$. The constrained privacy mechanism requires samples from p^*. A simplest, though often inefficient or event impractical, method is rejection sampling. Since

$$p^*(s) = \begin{cases} c^{-1}p(s) & \text{if } s \in \mathcal{S}^* \\ 0 & \text{otherwise,} \end{cases}$$

where $c = \int_{\mathcal{S}^*} p(s)\,ds$, one can opt for a proposal density q with support $\mathcal{S}^*$ such that $\sup_{s\in\mathcal{S}^*}[p(s)/q(s)] \leq R$, and accept a sample $s \sim q$ with probability $p(s)/Rq(s)$. This encompasses the option to set $q = p$, the unconstrained privacy kernel itself, and keep sampling until the sample falls into $\mathcal{S}^*$. This strategy is clearly inefficient in general, and impossible when $c = 0$.

Efficient algorithm tailor-made to specifications of $\mathcal{S}^*$ are possible. Here, we present in Algorithm 1 a generic approach based on Metropolized Independent Sampling [MIS; 18], for the most common case in which the mandated invariants are expressed in terms of a consistent system of linear equalities and inequalities

$$\mathcal{S}^* = \left\{ s \in \mathbb{R}^d : As = a, Bs \geq b \right\}.$$

Here A and B are $d_A \times d$ and $d_B \times d$ matrices with ranks $d_A < d$ and $d_B < d$ respectively, and a and b vectors with length d_A and d_B respectively. The algorithm requires a proposal index set $\mathcal{I}$, a subset of $\{1,\ldots,d\}$ of size $d - d_A$ such that $\text{rank}\left(A_{[\mathcal{I}^c]}\right) = d_A$, where $A_{[\mathcal{J}]}$ is the submatrix of A consisting of all columns whose indices belong to $\mathcal{J}$. For each A, the choice of $\mathcal{I}$ may not be unique, and can have a potential impact on the efficiency of the algorithm.

This algorithm is applicable to both discrete and continuous data and privatization schemes, and it does not require the normalizing constant for p^*. In Section 3 below, we use it to construct a mechanism for differentially private demographic contingency tables subject to both linear equality and inequality invariant constraints.

Algorithm 1 Metropolized Independent Differentially Private Sampler with Invariants

Input: unconstrained privacy mechanism p,
confidential query s^*, invariant parameters (A, a, B, b),
proposal distribution q, proposal index set $\mathcal{I}$,
initial value $s^{(0)} \in \mathcal{S}^*$, integer nsim;

Iterate: for $t = 0, 1, \ldots, \text{nsim} - 1$, at $t+1$:

step 1, propose $\tilde{s}$:
1-1. sample $\tilde{s}_{\mathcal{I}} \sim q$;
1-2. solve for $\tilde{s}_{\mathcal{I}^c}$ in $A_{[\mathcal{I}]}\tilde{s}_{\mathcal{I}} + A_{[\mathcal{I}^c]}\tilde{s}_{\mathcal{I}^c} = a$;
1-3. write $\tilde{s} = (\tilde{s}_{\mathcal{I}}, \tilde{s}_{\mathcal{I}^c})$;

step 2, compute $\alpha(s^{(t)}, \tilde{s}) = \min\left\{1, \frac{p(\tilde{s})\mathbf{1}(B\tilde{s}\geq b)q\left(s^{(t)}_{\mathcal{I}}\right)}{p\left(s^{(t)}\right)q(\tilde{s}_{\mathcal{I}})}\right\}$;

step 3, set $s^{(t+1)} = \tilde{s}$ with probability $\alpha(s^{(t)}, \tilde{s})$,
otherwise set $s^{(t+1)} = s^{(t)}$.

Output: a set of draws $\{s^{(t)}\}$, $t = 1, \ldots, \text{nsim}$.

3 CONTINGENCY TABLE WITH INVARIANTS

The table we consider is of dimension 2×23, with rows representing sex (male/female), and columns representing age bucketed roughly every four years, with finer buckets around key age ranges such as 18, 21 and 60. This data structure corresponds to the 2010 Census Table #P12 and 2020 Census Table #P7 [21], one of the most referenced type of contingency table releases by the Census Bureau at various geographic levels. For the purpose of computational illustration, this example will use simulation to construct synthetic datasets that represent the confidential Census demographic data.

Let s be a vector of length 46, denoting the row-vectorized contingency table. The constraints to be imposed on the differentially private table include

(1) total population;
(2) proportion of female population;
(3) total voting age population (18+ age); and
(4) nonnegative table entries.

Items (1) to (3) constitute equality constraints, and item (4) inequality constraints. The unconstrained privacy mechanism that serves as the basis of our construction is the Double Geometric mechanism, with distribution function

$$p(s) = \prod_{i=1}^{46} p_i\left(s_i - s_i^* \mid \epsilon\right),$$

where p_i is as defined in (1.3), and the privacy loss budget set to $\epsilon = 0.5$ per cell. The proposal distribution is set to be of the same family as does the unconstrained privatization algorithm, but it is given a distinct dispersion parameter $\tilde{\epsilon}$ to tune for best algorithmic performance, in this case the acceptance probability of the algorithm. The proposal distribution function for $\tilde{s}_{\mathcal{I}}$, the $(d-d_A)$-length subvector of the kth proposal $\tilde{s}$, is

$$q(s) = \prod_{i\in\mathcal{I}} p_i\left(s_i - s_i^* \mid \tilde{\epsilon}\right),$$

where $\mathcal{I}$ is chosen to be $\{2,\ldots,22,24,\ldots 45\}$. The remainder coordinates of the kth proposal, $\tilde{s}_{\mathcal{I}^c}$, is solved according to the equality constraint $A\tilde{s} = a$.

Table 1 displays a simulated confidential table (top) and a constrained differentially private table (bottom) based on the confidential table, where the draw is produced by Algorithm 1. In this case, setting $\tilde{\epsilon} = 0.6$ yields the best acceptance probability, with acceptance probability at 1.68%, as shown in Figure 2 in Appendix C

	< 5	6-10	11-15	16-17	18-19	20	21	22-24	25-29	30-34	35-39	40-44	45-49	50-54
Female	8	6	3	6	4	4	4	8	5	7	7	6	1	5
Male	3	4	5	8	6	4	5	5	5	6	10	7	3	2
					55-59	60-61	62-64	65-66	67-69	70-74	75-79	80-84	85+	Total
Female					4	4	9	6	2	8	8	8	7	**130**
Male					5	11	6	4	7	4	5	3	8	**126**
Voting				**43**									**213**	**256**

	< 5	6-10	11-15	16-17	18-19	20	21	22-24	25-29	30-34	35-39	40-44	45-49	50-54
Female	6	6	4	3	2	10	2	5	6	7	5	6	0	5
Male	9	4	5	6	3	4	5	5	3	4	7	8	5	3
					55-59	60-61	62-64	65-66	67-69	70-74	75-79	80-84	85+	Total
Female					4	5	10	8	3	8	8	12	5	**130**
Male					2	23	8	2	6	3	0	3	8	**126**
Voting				**43**									**213**	**256**

Table 1: A confidential sex × age contingency table (top) and a corresponding constrained differentially private (bottom) release, subject to total population, proportion female population and voting age population constraints (bold).

alongside traceplots for the second and the first cells of the table, with the former index belonging to $\mathcal{I}$ and the latter not.

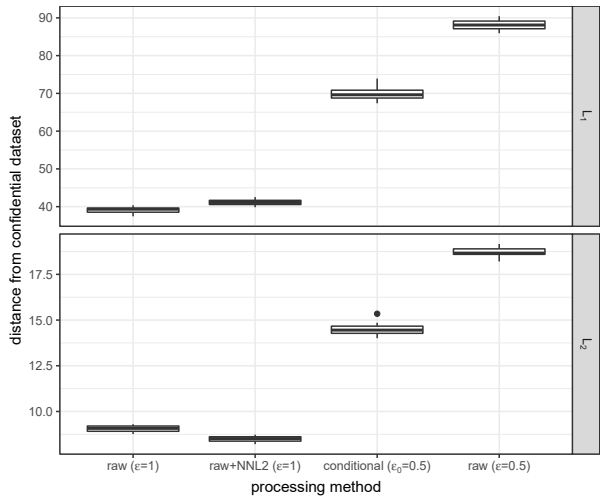

Figure 1: Average L_1 (top) and L_2 (bottom) distances of a simulated confidential dataset from its privatized releases using four different processing methods.

We compare the outputs of our congenial mechanism with the unconstrained privacy mechanism, with and without post-processing via nonnegative L_2 minimization onto the invariant set $\mathcal{S}^*$ defined by (1)-(4). A total of 20 confidential contingency tables are simulated, each with cells following

$$s_i \overset{i.i.d.}{\sim} \textit{Negative Binomial}\,(100, .05)\,,$$

which has mean $E(s_i) = 5.26$ and variance $Var(s_i) = 5.54$. For each confidential table, 100 privatized releases were created using each of the following methods:

a) the unconstrained (raw) Double Geometric mechanism with privacy loss budget $\epsilon = 1$ per cell;
b) the above mechanism followed by nonnegative L_2 (NNL2) minimization onto the subspace defined by the four constraints, i.e. $\mathcal{S}^*$;
c) Our proposed $\mathcal{S}^*$-conditional algorithm per Algorithm 1, constructed based on an unconstrained Double Geometric mechanism with $\epsilon_0 = 0.5$; and
d) the same unconstrained Double Geometric mechanism as in (a), but with $\epsilon = 0.5$ per cell.

The L_1 distance and L_2 distance between a confidential table s and $\tilde{s}$, a privatization of s, are given respectively by

$$L_1\,(s, \tilde{s}) = \sum_i |s_i - \tilde{s}_i| \quad \text{and} \quad L_2\,(s, \tilde{s}) = \sqrt{\sum_i (s_i - \tilde{s}_i)^2}. \quad (3.1)$$

For each of the four types of privatization and processing mechanisms, we compute averages of both distances over 100 realizations of $\tilde{s}$. Figure 1 displays the box plots of these average distances over the 20 simulated copies of private table s. We observe that the conditional mechanism, constructed from an unconstrained mechanism with $\epsilon_0 = 0.5$, exhibits a degree of variability between the unconstrained mechanisms with privacy loss budget $\epsilon = 0.5$ and $\epsilon = 1$. On the other hand, the nonnegative L_2 projection of the unconstrained mechanism with $\epsilon = 1$ achieves a level of accuracy mostly on par with it. Note that this observation should not be taken as a suggestion of relative accuracy between the nonnegative L_2 minimization and the constrained mechanism, because the effective privacy guarantee that either mechanism enjoys is undetermined, an issue we will discuss further at the end of Section 4. That said, by Theorem 2.1 that the conditional mechanism inflates the privacy loss budget of a unconstrained algorithm by $(1 + \gamma)$, the empirical

observation suggests that the effective γ may be somewhere in between 0 and 1. In the following sections, we will see examples where $\gamma = 0$ or even $\gamma < 0$.

4 CURATOR'S POST-PROCESSING MAY NOT BE INNOCENT PROCESSING

A common practice to ensure unconstrained differentially private releases respect the mandated disclosure requirements is through optimization-based post-processing, which takes the general form

$$f\left(M; \mathcal{S}^*\right) = \text{argmin}_{s \in \mathcal{S}^*} \Delta\left(M, s\right). \tag{4.1}$$

That is, f is the element in $\mathcal{S}^*$ that is the closest to M according to some discrepancy measure Δ, typically a distance, such as the two given in (3.1). In the case of the Census Bureau's TopDown algorithm, f is a composite post-processing procedure consisting of first a nonnegative L_2 minimization followed by an L_1 minimization onto the integer solutions; see [1].

In the literature of differential privacy, there is a widely referenced theorem which establishes that differentially privatized releases are "immune" to post-processing [6]. The theorem states that if M is an ϵ-differentially private mechanism and g is an arbitrary function, then $g \circ M$ is still ϵ-differentially private. Indeed for any g-measurable set B,

$$P\left(g(M\left(x\right)) \in B\right) = P(M\left(x\right) \in g^{-1}(B)). \tag{4.2}$$

Thus for every $x \in \mathcal{X}$, the maximal increased risk of disclosure from releasing $g(M(x))$ cannot exceed that from releasing $M(x)$ (but the reversed is guaranteed only when g is one-to-one). Intuitively, further blurring an already blurred picture can only make it harder, not easier, to see what is in the original picture.

This intuition, however, is based on the assumption that the further blurring process does not use any knowledge about the original picture. We need to make clear here that imposing invariants on differentially private releases via optimization-based post-processing, in the sense of the operation discussed here, does *not* in general fall under the jurisdiction of the post-processing theorem. This is because f, the function used to impose invariants on the unconstrained output $M\left(x\right)$, is implicitly dependent on the confidential dataset x^*, with the dependence induced via $\mathcal{S}^*$, or equivalently $\mathcal{X}^*$ to which x^* belongs. Since $\mathcal{S}^*$ supplies information about the confidential database, whereas the unconstrained mechanism M is by design not preferential towards $\mathcal{S}^*$, any further processing of M that makes nontrivial use of $\mathcal{S}^*$ risks violating the privacy guarantee that M deserves.

The post-processing theorem guarantees that no loss of privacy will be induced to the privatized query via any functional transformation that may be carried out by an ordinary analyst or data user. However, imposing invariants is the kind of post-processing that only the data curator – one who has access to the confidential data – is capable of performing. In the extreme scenario (see Example 5.2) that the invariant forces the privatized disclosure to be precisely equal to the confidential query, for the data curator to achieve this algorithmically is as simple as taking the privatized query M and projecting it to the single point in $\mathbb{R}^d$ defined by the confidential value $s\left(x^*\right)$. But this is impossible for a data user who do not know what $s\left(x^*\right)$ is.

One may wonder the following question. While the invariant $\mathcal{S}^*$ has a dependence on the confidential x^*, itself is nevertheless public information. Doesn't that make $f(\cdot; \mathcal{S}^*)$ a fully specified function, just like g in the post-processing theorem? The answer is no in general, and the distinction here is a subtle one. When talking about the value of the invariants, it suffices to regard $\mathcal{S}^*$ merely as an announced description of the confidential data. However as alluded to previously, $\mathcal{S}^*$ is a set-valued map from the database space to subsets of the query space, i.e. $\mathcal{S}^* : \mathcal{X} \to \mathscr{B}(\mathbb{R}^d)$. Almost always is the case in practice that the functional form of map of $\mathcal{S}^*$ is *a priori* determined, but its value – namely $\mathcal{S}^*\left(x^*\right)$ – can be calculated only after the confidential data is observed. Indeed, the actual specification of $\mathcal{S}^*$ would almost certainly change if x^* were observed differently. This means for an f-measurable set B and a database $x \in \mathcal{X}$, the equivalent events in the f and the M spaces are now

$$f\left(M\left(x\right); \mathcal{S}^*\left(x\right)\right) \in B \;\Leftrightarrow\; M\left(x\right) \in f^{-1}\left(B; \mathcal{S}^*(x)\right), \tag{4.3}$$

noting that the inverse function $f^{-1}(\cdot; \mathcal{S}^*(x))$ now depends on x.

To see the complication caused by this dependence, write $\tilde{B}(x) = f^{-1}(B; \mathcal{S}^*(x))$ and $f(x) = f(M\left(x\right); \mathcal{S}^*(x))$. We then have

$$\frac{P(f(x) \in B)}{P(f(x') \in B)} = \frac{P(M(x) \in \tilde{B}(x))}{P(M(x') \in \tilde{B}(x'))}. \tag{4.4}$$

Although both $\tilde{B}(x)$ and $B(x')$ are measurable sets, they are not necessarily the same when $x' \neq x$. Hence we cannot use (1.1) to conclude that the right hand side of (4.4) is bounded above by $\exp\left(\epsilon\right)$. This does not necessarily imply that the post-processing f as defined in (4.1) is not ϵ-differentially private. Indeed, we prove in Appendix A that both L_2 and (a class of) L_1 post-processing in Example 4.1 below is ϵ-differentially private. But it does imply that in general, the statistical and privacy properties of f are not straightforwardly inherited from that of M, and hence they need to be established on a case by case basis.

Another major drawback of using optimization-based methods to impose invariants is that the statistical intelligibility of differential privacy is obscured. The post-processing function f is often procedurally defined, hence a complex and confidential data-dependent map from the unconstrained query space to the constrained query space, with almost impenetrable statistical properties, and certainly so for any given database. In contrast, using conditioning to realize differential privacy with mandated disclosure, despite often computationally demanding by construction, preserves the statistical intelligibility of the privacy mechanism. The constrained privacy mechanism is distributionally – as opposed to procedurally – constructed, preserving the possibility of transparent downstream analysis. It furthermore delivers privacy guarantee in the same format as does differential privacy without constraints, offering a congenial statistical interpretation that resembles the original.

Below we use an example to compare congenial privacy with two approaches of post-processing for a same query function. The example is simple enough for analytical derivations of the distributions of post-processing mechanisms to be possible. As we will see, the three approaches to impose invariant constraints yield distinct theoretical behaviors.

Example 4.1 (A two-bin histogram with constrained total). Suppose the confidential database x is a binary vector, and the query of interest tabulates the number of 0 and 1 entries in x, i.e.

$$s(x) = (s_1(x), s_2(x)) = \left(\sum_i \mathbf{1}(x_i = 0), \sum_i \mathbf{1}(x_i = 1) \right).$$

Employ the Laplace mechanism as the unconstrained privatization mechanism to protect the two-bin histogram, i.e.

$$M(x) = (m_1 = s_1 + u_1, m_2 = s_2 + u_2),\ u_i \overset{i.i.d.}{\sim} Lap(2\epsilon^{-1}),$$

expending in total ϵ privacy loss budget. The induced probability density of M is

$$p(m_1, m_2) = \left(\frac{\epsilon}{4}\right)^2 \exp\left\{-\frac{\epsilon}{2}(|m_1 - s_1| + |m_2 - s_2|)\right\}.$$

Suppose the invariant to be imposed is that the total of the privatized histogram shall be the same as that the the confidential query itself. That is, for any given x, the associated invariant set is

$$\mathcal{S}^*(x) = \left\{(a_1, a_2) \in \mathbb{R}^2 : a_1 + a_2 = s_1(x) + s_2(x)\right\}. \tag{4.5}$$

In the calculations below, a certain database x is fixed, and we write the invariant total $n = \|x\|$, the length of x.

Congenial privacy. Under the constraint of histogram total, our congenial $M^*(x)$ is obtained from the conditional distribution

$$(s_1 + u_1, s_2 + u_2) \mid u_1 + u_2 = 0.$$

It turns out that the probability density of M^* is

$$P(m_1 = m, m_2 = n - m) = \frac{\epsilon}{2} \exp\{-\epsilon |m - s_1|\} \tag{4.6}$$

and 0 otherwise; see Appendix A. That is, our congenial mechanism M^* is simply to draw a u from $Lap(\epsilon^{-1})$, and then release $(m_1 = s_1 + u, m_2 = s_2 - u)$. Clearly, the privacy property of M^* is the same as its first component, call it M_1^*, which protects s_1 by releasing m_1 because setting $m_2 = n - m_1$ is a deterministic step with no implication on privacy when n is known. But M_1^* is simply the Laplace mechanism with ϵ budget. Consequently, our congenial mechanism maintains the same ϵ guarantee as the original unconstrained mechanism, even though the meaning of protection is different, as we emphasized in Section 2.2.

Post-processing with L_2 minimization. Here we minimize the L_2 distance between $M(x)$ and the post-processed histogram release, denoted f_{L_2}, subject to its sum being n. The solution is

$$f_{L_2}\left(M(x), \mathcal{S}^*\right) = \text{argmin}_{s \in \mathcal{S}^*} \|M(x) - s\|_2 = (\bar{x} + \tilde{u}, \bar{x} - \tilde{u}),$$

where $\tilde{u}$ is the average of two independent Laplace random variables with scale $2\epsilon^{-1}$. As can be easily seen (for example from its characteristic function), $\tilde{u}$ is not a Laplace random variable, but in fact follows distribution

$$\frac{1}{2} Lap(\epsilon^{-1}) + \frac{1}{2} SgnGamma(2, \epsilon^{-1}), \tag{4.7}$$

that is a 50-50 mixture of a Laplace distribution with scale ϵ^{-1} and a signed Gamma distribution (i.e. a regular Gamma distribution multiplied by a fair random sign) with shape $k = 2$ and scale ϵ^{-1}; see Appendix A for derivation. It is worth noting that since a signed Gamma distribution of shape $k = 2$ can be written as the sum of two independent Laplace distributions of the same scale, $\tilde{u}$ is more variable than a single independent Laplace random variable of the same scale. Hence for any x, the privatized release using f_{L_2} will be more variable than that of the congenial privatization M^*. Intuitively, this suggests that f_{L_2} should not do worse than M^* in terms of privacy protection. Indeed as we will show in Appendix A, f_{L_2} also achieves the same ϵ-differentially private guarantee.

	$E(m_1, m_2)$	$Var(m_1)$
M^*	$(s_1(x), s_2(x))^\top$	$\frac{2}{\epsilon^2}$
f_{L_2}	$(\bar{s}(x), \bar{s}(x))^\top$	$\frac{4}{\epsilon^2}$
f_{L_1}	$(s_1(x), s_2(x))^\top$	$\left[\frac{4}{\epsilon^2}, \frac{8}{\epsilon^2}\right]$

Table 2: Differentially private two-bin histogram with invariant total: expectation and first-component variance of the conditional (M^*) and post-processed (f_{L_2} and f_{L_1}) histograms.

Post-processing with L_1 minimization. If we change the L_2 norm to L_1 norm in the above, the privatization mechanism is no longer unique. There will be infinitely many solutions in the form of

$$f_{L_1}\left(M(x), \mathcal{S}^*\right) = \text{argmin}_{s \in \mathcal{S}^*} \|M(x) - s\|_1 := (\tilde{s}, n - \tilde{s}),$$

where $\tilde{s}$ only needs to satisfy

$$\begin{aligned} \tilde{s} &\in [\min\{s_1 + u_1, n - (s_2 + u_2)\}, \max\{s_1 + u_1, n - (s_2 + u_2)\}] \\ &\overset{L}{=} [s_1 + \min(u_1, u_2), s_1 + \max(u_1, u_2)], \end{aligned}$$

where $\min(u_1, u_2)$ and $\max(u_1, u_2)$ are the minimum and the maximum of two i.i.d. Laplace random variables. In particular, choosing any convex combination of u_1 and u_2 as the additive noise term to the first entry constitutes a solution, i.e., $\tilde{s}_1 = s_1 + \beta u_1 + (1-\beta)u_2$ for some $\beta \in [0, 1]$, and then set $\tilde{s}_2 = n - \tilde{s}_1$. For the rest of the article, L_1 post-processing refers to this convex combination strategy.

For ease of reference, Table 2 collects a comparison of the constrained differentially private histogram M^* and two the post-processing approaches, f_{L_2} and f_{L_1}, in terms of the expectation and variance of the resulting release for a given database $x \in \mathcal{X}$ and confidential query s. All expectations are taken with respect to the relevant mechanism.

We can see that our congenial mechanism has the smallest variance. Because the congenial mechanism and f_{L_2} both carry the same ϵ-privacy guarantee which *cannot* be further improved, we can comfortably declare that f_{L_2} is *inadmissible* because it is dominated by the congenial mechanism, providing less utility (in terms of statistical precision) without the benefit of increased privacy protection. However, we cannot say that the congenial mechanism dominates f_{L_1} even though it still leads to smaller variance. This is because, as we will prove in Appendix A, the attained level of privacy guarantee of f_{L_1} is $\epsilon/(2\max\{\beta, 1-\beta\})$, which is never worse than ϵ. Hence the increased variance under f_{L_1} might be acceptable as a price for gaining more privacy protection. In general, comparing the utility of two privatization mechanisms with the same nominal but different actual privacy loss budget is as thorny an issue as comparing the statistical power of two testing procedures with the same nominal, but different actual, Type I error rates.

5 DISCUSSION

5.1 Finding better γ

While Theorem 2.1 always holds with $\gamma = 1$, it likely sets a loose bound on the ratio between $P(M^*(x) \in B)$ and $P(M^*(x') \in B)$, hence declaring an overly "conservative" nominal level of privacy loss induced by M^*. Depending on how the invariant $\mathcal{S}^*$ interacts with the distributional property of the unconstrained mechanism M in a specific instance, γ can be shown to take smaller values, adding more "bang of the buck" to the privacy loss budget, so to speak. Three examples are given below.

Example 5.1 (trivial invariants). Consider the trivial case where the set of invariants does not actually impose any restriction, i.e., $\mathcal{X}^* = \mathcal{X}$. It is then necessarily true that $\mathcal{S}^* = \mathcal{S}$, and the "constrained" differentially private mechanism is identical in distribution to the unconstrained one: $M^* \overset{L}{=} M$. In this case, $\gamma = 0$ and M^* is ϵ-differentially private.

Example 5.2 (rounding and secrecy). Let x be a binary vector of length n indicating a group of individuals' possession of a certain feature (yes 1/no 0), and the query of interest is $s(x) = \lceil \sum x_i / 10 \rceil$, or the number of groups of size 10 that can be formed by people who possesses the feature. A Double Geometric mechanism $M(x) = s(x) + U$ is used to protect the query, with a privacy loss budget of ϵ (under the global sensitivity of $\nabla(s) = 1$).

Suppose the following invariant set is mandated for disclosure:

$$\mathcal{X}^* = \left\{ (x_1, \ldots, x_n) \in \{0,1\}^n : \sum x_i \in [41, 50] \right\},$$

or equivalently, $\mathcal{S}^* = \{5\}$ is the singleton set that contains nothing but the true value $s(x^*) = 5$. In this case, the implied constrained privacy mechanism M^* is equivalent to a degenerate distribution: $P(M^*(x) = 5) = 1$ for all $x \in \mathcal{X}^*$. Furthermore, for all neighboring datasets $(x, x') \in \mathcal{X}^* \times \mathcal{X}^*$, and any B a measurable subset of $\mathbb{N}$,

$$P\left(M^*(x) \in B\right) = \exp(0) P\left(M^*\left(x'\right) \in B\right) = \begin{cases} 1 & \text{if } 5 \in B \\ 0 & \text{otherwise.} \end{cases}$$

Therefore in this particular instance, M^* is in fact 0-differentially private, corresponding to $\gamma = -1$ in Theorem 2.1. This means that for those databases conformal to the invariant $\mathcal{X}^*$, M^* supplies no discriminatory information among them whatsoever. Indeed, if the value of the supposedly private query is already public knowledge, no mechanism can further increase its disclosure risk, therefore achieving complete differential privacy.

Our third example is a less trivial example of $\gamma < 0$, which is actually provided by the congenial mechanism in Example 4.1. There, although the guaranteed privacy loss budget is still ϵ, in applying Theorem 2.1, k must be set to 2 or greater, because under the constraint of fixed sum, the nearest neighbors among binary vectors (x, x') must have $\mathsf{d}(x, x') = 2$. Hence the ϵ privacy bound implies $k(1+\gamma) = 2(1+\gamma) = 1$, yielding $\gamma = -0.5$.

This example also reminds us that a major cause of information leakage due to invariants is the structural erosion to the underlying data space, such as making $\mathsf{d}(x, x') = 1$ (as measured on the original space $\mathcal{X}$) impossible. In reality, the unconstrained data space $\mathcal{X}$ is typically regular, and contains $\mathcal{X}^*$ as a proper subset. We should expect to find many $x \in \mathcal{X}^*$, and many (if not many more) $x' \in \mathcal{X} \backslash \mathcal{X}^*$ such that x and x' are neighbors, near or far. Knowing that the confidential dataset must belong to $\mathcal{X}^*$ categorically rules out the possibility that all the x''s can be the confidential dataset, weakening the differential privacy promise by eliminating the neighbors. If the invariant is sufficiently restrictive such that $\mathcal{X}^*$ becomes topologically small relative to $\mathcal{X}$, it may be the case that for some $x \in \mathcal{X}^*$, all of its original immediate neighboring datasets are not in $\mathcal{X}^*$:

$$\left\{ (x, x') \in \mathcal{X}^* \times \mathcal{X}^* : \mathsf{d}(x, x') = 1 \right\} = \emptyset,$$

in which case we say that the neighboring structure of the original data space of the database is *substantially disrupted*, as seen in Example 4.1. If the disruption is so substantial that neighbors of any distance cease to exist, we say that the neighborhood structure is completely *destroyed*:

$$\left\{ (x, x') \in \mathcal{X}^* \times \mathcal{X}^* : \mathsf{d}(x, x') \geq 1 \right\} = \emptyset.$$

Then, even for the constrained mechanism M^*, the ϵ-differential privacy promise becomes vacuously true, since no possible neighboring pairs remain for which the concept of privacy is applicable. However, Example 5.2 demonstrates that vacuous privacy promise can occur without the neighborhood structures completely destroyed.

In general, it is conceptually difficult to parse out the share of responsibility on privacy attributable to the data curator under any scenario of mandated disclosure. If certain information is made public, then any information that it logically implies cannot be expected to be protected, either. The best that we can expect any privacy mechanism to deliver, then, is protection over information that truly remains. Notions that serve the equivalent purposes as $\mathcal{X}^*$ and $\mathcal{S}^*$ have been proposed in the literature for expositions of new notions of differential privacy, including *blowfish* and *pufferfish* privacy [14, 16], where the privacy guarantee is re-conceptualized on the restricted space *modulo* any structural erosion to the original sample space due to external or auxiliary information available to an attacker. When interpreting the promise of Theorem 2.1, we shall pay due diligence to the case in which immediate neighbors no longer exists, and talk about the ϵ-differential privacy guarantee only for those k-neighbors that actually do.

5.2 Other interpretations of privacy

The literature has seen other lines of work that offer interpretations of differential privacy in statistical terms. Notably, the *posterior-to-posterior semantics* of differential privacy [2, 4, 15] explains the effect of privacy also in the vocabulary of Bayesian posteriors. The posterior-to-posterior semantics establishes differential privacy as a bound for the ratio of posterior probabilities assigned to an individual confidential data entry, when the private mechanism is applied to neighboring datasets that differ in only one entry. The said ratio is between the two quantities

$$\pi\left(x_i^* = \omega \mid M_\epsilon(x) \in B\right) \text{ and } \pi\left(x_i^* = \omega \mid M_\epsilon\left(x'\right) \in B\right), \quad (5.1)$$

where x and x' are neighboring datasets. What varies between the two posterior quantities is the confidential dataset on which the private query is applied. The datasets x and x' are neighboring datasets, one of which presumably (but not necessarily) contains the true value of the ith confidential data entry x_i^*, and the other contains a fabricated value of it.

The comparison in (5.1) raises the question of what it means by the conditional probability of x_i^* given a private query constructed from a database that does *not* contain this true value, as this conditional probability hinges on external knowledge about how a fabricated database may inform the actual confidential database. Our *prior-to-posterior semantics* formulated in Theorem 1.3 takes a practical point of view and avoids such conceptual complication. We compare the disclosure risk before and after an *actual* release, reflecting the core idea behind differential privacy.

5.3 Full privacy or vacuous knowledge

As alluded to in Section 1.4, the notion of vacuous knowledge cannot be appropriately captured by ordinary probabilities. The defect reflects a fundamental inability of the language of probability in expressing a true lack of knowledge, a central struggle in the Bayesian literature that motivated endeavors in search for the so-called "objective priors" [9]. Neither the uniform distribution nor any other reference distributions are truly information-free, as they all invariably invoke some principle of indifference in relation to a specific criterion (such as the Lebesgue measure, the counting measure, or the likelihood function) which is subject to debate.

To supply the rigorous definition needed to define probabilistically $M_0(x)$ in Section 1.4, we invoke the concept of lower probability functions, and as a special case belief functions [see e.g. 11], both generalized versions of a probability function which amounts to a set of probability distributions on a given space. The statistical information contained in M_0 is represented by the vacuous lower probability function, denoted $\underline{P}$, which takes the value $\underline{P}(B) = 1$ only when $B = \mathbb{R}^d$, and 0 everywhere else. Equivalently stated in terms of its conjugate upper probability function $\overline{P}(B) = 1 - \underline{P}(B^c)$,

$$\overline{P}(B) = \begin{cases} 1 & \text{if } B \in \mathscr{B}\left(\mathbb{R}^d\right) \setminus \{\emptyset\}, \\ 0 & \text{if } B = \emptyset. \end{cases} \tag{5.2}$$

That is, the statistical information contained in M_0 can be (but is not known to be) concordant with *any* Borel-measurable probability function, thus the probability of any B is as low as 0 and as high as 1, as long as B is neither the full set nor the empty set.

Generally, the conditioning operation involving lower probability functions is not trivial, and it is not unique due to the existence of several applicable rules. But if the lower probability functions being conditioned on is vacuous, there is consensus among different rules as to what posterior distribution should result, namely precisely as stated in (1.7). See [13] for an extended exposition of conditioning rules involving lower probability and belief functions.

5.4 Future directions

This work points to several future directions of pursuit. On the computational front, how do we construct efficient algorithms to realize congenial privacy, by drawing possibly high-dimensional releases subject to complex constraints? When we use non-perfect Markov chain Monte Carlo to accomplish this task, how do we ensure the declared privacy guarantee is not destroyed because a chain cannot run indefinitely? On the privacy front, for every constrained mechanism constructed through conditioning, how to find the best γ value that tracks as closely as possible the effective privacy loss budget, which in turn enables fair performance comparisons among invariant-respecting algorithms? Furthermore, how to achieve an orthogonal decomposition of the public, invariant information from the free, residual information that remains in the confidential microdata, in a logical sense without having to resort to the probabilistic vocabulary of statistical independence?

ACKNOWLEDGMENTS

The authors thank Jeremy Seeman, Salil Vadhan, Guanyang Wang and three anonymous reviewers for helpful suggestions. Ruobin Gong gratefully acknowledges research support by the National Science Foundation (NSF) DMS-1916002. Xiao-Li Meng also thanks NSF for partial financial support while completing this article.

REFERENCES

[1] John Abowd, Robert Ashmead, Garfinkel Simson, Daniel Kifer, Philip Leclerc, Ashwin Machanavajjhala, and William Sexton. 2019. *Census TopDown: Differentially Private Data, Incremental Schemas, and Consistency with Public Knowledge.* Technical Report. US Census Bureau.

[2] Robert Ashmead, Daniel Kifer, Philip Leclerc, Ashwin Machanavajjhala, and William Sexton. 2019. *Effective Privacy After Adjusting for Invariants with Applications to the 2020 Census.* Technical Report. US Census Bureau.

[3] Alexander W Blocker and Xiao-Li Meng. 2013. The potential and perils of preprocessing: Building new foundations. *Bernoulli* 19, 4 (2013), 1176–1211.

[4] Irit Dinur and Kobbi Nissim. 2003. Revealing information while preserving privacy. In *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems.* 202–210.

[5] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference.* Springer, 265–284.

[6] Cynthia Dwork and Aaron Roth. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science* 9, 3–4 (2014), 211–407.

[7] Robert E Fay. 1992. *When Are Inferences from Multiple Imputation Valid?* US Census Bureau.

[8] Ferdinando Fioretto and Pascal Van Hentenryck. 2019. Differential privacy of hierarchical census data: An optimization approach. In *International Conference on Principles and Practice of Constraint Programming.* Springer, 639–655.

[9] Malay Ghosh. 2011. Objective priors: An introduction for frequentists. *Statist. Sci.* 26, 2 (2011), 187–202.

[10] Ruobin Gong. 2019. Exact inference with approximate computation for differentially private data via perturbations. *arXiv preprint arXiv:1909.12237* (2019).

[11] Ruobin Gong. 2019. Simultaneous Inference under the Vacuous Orientation Assumption. *Proceedings of Machine Learning Research* 103 (2019), 225–234.

[12] Ruobin Gong. 2020. Transparent Privacy is Principled Privacy. *arXiv preprint arXiv:2006.08522* (2020).

[13] Ruobin Gong and Xiao-Li Meng. 2020. Judicious judgment meets unsettling updating: dilation, sure loss, and Simpson's paradox. *Statist. Sci.* (2020).

[14] Xi He, Ashwin Machanavajjhala, and Bolin Ding. 2014. Blowfish privacy: Tuning privacy-utility trade-offs using policies. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data.* 1447–1458.

[15] Shiva P Kasiviswanathan and Adam Smith. 2014. On the'semantics' of differential privacy: A bayesian formulation. *Journal of Privacy and Confidentiality* 6, 1 (2014).

[16] Daniel Kifer and Ashwin Machanavajjhala. 2012. A rigorous and customizable framework for privacy. In *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGAI symposium on Principles of Database Systems.* 77–88.

[17] PS Kott. 1995. A paradox of multiple imputation. In *Proceedings of the Section on Survey Research Methods.* American Statistical Association, 384–389.

[18] Jun S Liu. 1996. Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Statistics and computing* 6, 2 (1996), 113–119.

[19] Xiao-Li Meng. 1994. Multiple-imputation inferences with uncongenial sources of input. *Statist. Sci.* (1994), 538–558.

[20] Donald B Rubin. 2004. *Multiple imputation for nonresponse in surveys.* John Wiley & Sons.

[21] US Census Bureau. 2020. 2010 Demonstration Data Products. https://www.census.gov/programs-surveys/decennial-census/2020-census/planning-management/2020-census-data-products/2010-demonstration-data-products.html [Accessed: 04-08-2020].

[22] Xianchao Xie and Xiao-Li Meng. 2017. Dissecting Multiple Imputation from a Multi-Phase Inference Perspective: What Happens When God's, Imputer's and Analyst's Models Are Uncongenial? *Statistica Sinica* (2017), 1485–1545.

A PRIVACY GUARANTEES FOR CONGENIAL, L_1 AND L_2 METHODS

We first derive the conditional distribution of two i.i.d. Laplace random variables, given that their sum is zero. Let $u_1, u_2 \overset{i.i.d.}{\sim} Lap\left(2\epsilon^{-1}\right)$ and denote $v = u_1$, $w = u_1 + u_2$. Since (v, w) is linear in (u_1, u_2), their joint probability density function is given by.

$$\begin{aligned} p(v,w) &\propto p(u_1(v,w), u_2(v,w)) \\ &\propto \exp\left(-0.5\epsilon|v| - 0.5\epsilon|w-v|\right), \end{aligned}$$

This implies that

$$\begin{aligned} p(v \mid w=0) &\propto p(v, w=0) \\ &\propto \exp\left(-\epsilon|v|\right) \sim Lap(\epsilon^{-1}), \end{aligned}$$

which leads to (4.6).

We then derive the density for $\tilde{u} = \beta u_1 + (1-\beta)u_2$, where $\beta \in (0,1)$ (the case of $\beta = 0$ or 1 is trivial). This covers the L_1 projection case, where any $\beta \in [0,1]$ is acceptable, and the L_2 projection case, where $\beta = 1/2$. Since $u_1 = \beta^{-1}[\tilde{u} - (1-\beta)u_2]$, the Jacobian from $(\tilde{u}, u_2)$ to (u_1, u_2) is β^{-1}. Consequently,

$$p_\epsilon(\tilde{u}, u_2) = \frac{\epsilon^2}{16\beta} \exp\left\{-\frac{\epsilon}{2\beta}\left[|\tilde{u} - (1-\beta)u_2| + \beta|u_2|\right]\right\}.$$

To derive $p_\epsilon(\tilde{u})$, we assume without loss of generality $\tilde{u} \geq 0$. Consider $p_\epsilon(\tilde{u}) = \int p_\epsilon(\tilde{u}, u_2)\, du_2$ on three regions:

$$\begin{aligned} I_1(\beta) &= \frac{\epsilon^2}{16\beta}\int_{-\infty}^{0} \exp\left\{-\frac{\epsilon}{2\beta}(\tilde{u} - u_2)\right\} du_2 = \frac{\epsilon}{8}\exp\left\{-\frac{\epsilon}{2\beta}\tilde{u}\right\}; \\ I_2(\beta) &= \frac{\epsilon^2}{16\beta}\int_{0}^{\frac{\tilde{u}}{1-\beta}} \exp\left\{-\frac{\epsilon}{2\beta}\left[\tilde{u} + (2\beta-1)u_2\right]\right\} du_2 \\ &= \frac{\epsilon}{8(2\beta-1)}\left[\exp\left\{-\frac{\epsilon}{2\beta}\tilde{u}\right\} - \exp\left\{-\frac{\epsilon}{2(1-\beta)}\tilde{u}\right\}\right]; \\ I_3(\beta) &= \frac{\epsilon^2}{16\beta}\int_{\frac{\tilde{u}}{1-\beta}}^{\infty} \exp\left\{-\frac{\epsilon}{2\beta}\left[u_2 - \tilde{u}\right]\right\} du_2 = \frac{\epsilon}{8}\exp\left\{-\frac{\epsilon}{2(1-\beta)}\tilde{u}\right\}. \end{aligned}$$

Summing up these terms and noting the symmetry of p_ϵ, we obtain

$$\begin{aligned} p_\epsilon(\tilde{u}) &= \frac{\epsilon}{4(2\beta-1)}\left[\beta\exp\left\{-\frac{\epsilon}{2\beta}|\tilde{u}|\right\} - (1-\beta)\exp\left\{-\frac{\epsilon}{2(1-\beta)}|\tilde{u}|\right\}\right] \\ &= \frac{\beta^2}{(2\beta-1)}Lap\left(2\beta\epsilon^{-1}\right) - \frac{(1-\beta)^2}{(2\beta-1)}Lap\left(2(1-\beta)\epsilon^{-1}\right). \quad \text{(A.1)} \end{aligned}$$

Remark I. The expression (A.1) is fascinating. It shows that the density of a convex combination of i.i.d. Laplace random variables is a "mixture" but non-convex combination of two Laplace densities with different scale parameters, because

$$\frac{\beta^2}{(2\beta-1)} + \left[-\frac{(1-\beta)^2}{(2\beta-1)}\right] = 1.$$

That is, although the two weights add up to one, they always take the opposite sign when $\beta \neq 1/2$.

Remark II. When $\beta = 1/2$, the expression (A.1) is of 0/0 appearance, but is well-defined once taking the limit $\beta \to 1/2$ and using the L'Hospital's rule, yielding

$$p_\epsilon(\tilde{u}) = \frac{\epsilon}{4}\left[1 + \epsilon|\tilde{u}|\right]\exp\left\{-\epsilon|\tilde{u}|\right\}. \quad \text{(A.2)}$$

This can be written as

$$\begin{aligned} p(\tilde{u}) &= \frac{1}{2}\cdot\frac{\epsilon}{2}\exp\left\{-\epsilon|\tilde{u}|\right\} + \frac{1}{2}\cdot\frac{\epsilon^2}{2}|\tilde{u}|\exp\left\{-\epsilon|\tilde{u}|\right\} \\ &\sim \frac{1}{2}Lap\left(\epsilon^{-1}\right) + \frac{1}{2}SgnGamma\left(2, \epsilon^{-1}\right), \end{aligned}$$

suggesting that it is more variable that $Lap(\epsilon^{-1})$. We now prove that (A.2), moreover the entire family of distributions given by (A.1) as indexed by $\beta \in (0,1)$, is ϵ-differentially private. In fact, they attain a level of privacy protection more stringent than ϵ whenever $\beta \neq 1/2$. Our proof relies on the following general result, which can be useful for verifying differential privacy guarantees in other situations.

THEOREM A.1. *Suppose $f(x)$ is a positive real-valued function on a normed vector space $\mathcal{X}$, with its norm denoted by $|x|$. Suppose $f(x)$ has the following properties:*

- **(i)** *$f(x)$ is monotone decreasing in $|x|$;*
- **(ii)** *$g_\alpha(x) = f(x)e^{\alpha|x|}$ is monotone increasing in $|x|$, where α is a positive constant.*

Then for any $a \in \mathcal{X}$ and $b \in \mathcal{X}$, we have

$$\sup_{x\in\mathcal{X}} \frac{f(x-a)}{f(x-b)} \leq e^{\alpha|a-b|}. \quad \text{(A.3)}$$

PROOF. For any $x \in \mathcal{X}$, if $|x-a| > |x-b|$, then $f(x-a) \leq f(x-b)$ by (i) and hence (A.3) holds trivially. If $|x-a| \leq |x-b|$, then $g_\alpha(x-a) \leq g_\alpha(x-b)$ by (ii), and hence

$$\begin{aligned} \frac{f(x-a)}{f(x-b)} &= \frac{g_\alpha(|x-a|)}{g_\alpha(|x-b|)}e^{\alpha(|x-b|-|x-a|)} \\ &\leq e^{\alpha(|x-b|-|x-a|)} \leq e^{\alpha|a-b|}. \end{aligned}$$

□

To apply this result, we first note that for $p_\epsilon(x)$ of (A.1) with $x > 0$, its derivative is given by

$$\frac{dp_\epsilon(x)}{dx} = \frac{\epsilon^2}{8(2\beta-1)}\left[\exp\left\{-\frac{\epsilon}{2(1-\beta)}x\right\} - \exp\left\{-\frac{\epsilon}{2\beta}x\right\}\right] < 0,$$

for any $\beta \neq 1/2$. For $\beta = 1/2$, we can directly verify from (A.2) that

$$\frac{dp_\epsilon(x)}{dx} = -\frac{\epsilon^3}{4}x\exp\left\{-\epsilon x\right\} < 0,$$

Hence, condition (i) holds for (A.1) for all $\beta \in (0,1)$.

To establish condition (ii), the choice α is the key since it directly governs the degree of privacy guarantee. From the expression (A.3), we want the smallest α such that condition (ii) holds, which gives us the tightest bound hence better privacy guarantee. A good strategy here is to start with $\alpha = c\epsilon$ and let the mathematics tell us how to minimize over c. We start with the simplest case with $\beta = 1/2$. From

the expression (A.2), the smallest c that can make g_α monotone increasing is clearly 1. The resulting g_α also has the property that

$$\lim_{|x|\to\infty} \frac{g_\alpha(|x-a|)}{g_\alpha(|x-b|)} = 1. \tag{A.4}$$

This implies that the bound $e^{\epsilon|a-b|}$ can be approached arbitrarily closely by letting $|x| \to \infty$, which means that the privacy loss budget ϵ cannot be reduced. For our current application, this means the post-processing by L_2 projection is also differentially private at level ϵ, but not more stringent than that.

When $\beta \neq 1/2$, we assume without loss of generality $\beta > 1/2$. Then for $g_\alpha(x) = e^{c\epsilon|x|}p_\epsilon(x)$, it is easy to verify that for any $x > 0$,

$$\frac{dg_\alpha(x)}{dx} = \frac{\epsilon^2}{8(2\beta-1)}\left[w_c \exp\left\{\frac{w_c}{2\beta}\epsilon x\right\} + [w_c + c']\exp\left\{-\frac{w_c + c'}{2(1-\beta)}\epsilon x\right\}\right],$$

where $w_c = 2c\beta - 1$ and $c' = 2(1-c)$. Our job is to seek the smallest c such that this derivative is non-negative regardless of the value of x. Clearly the positivity holds when we set $w_c = 0$, that is $c = (2\beta)^{-1} < 1$, and hence $c' > 0$. To show that this is the smallest possible c, we see that when setting $\alpha = \epsilon/(2\beta)$, we have

$$\frac{g_\alpha(|x-a|)}{g_\alpha(|x-b|)} = \frac{\beta - (1-\beta)\exp\{-\tau|x-a|\}}{\beta - (1-\beta)\exp\{-\tau|x-b|\}},$$

where $\tau = (2\beta-1)/(2\beta(1-\beta)) > 0$. Clearly as $|x| \to \infty$, the ratio above goes to 1 regardless of the value a and b as long as they are fixed. Consequently, the same implication from (A.4) follows, that the bound $e^{\alpha|a-b|}$ can be approached arbitrarily closely with $\alpha = \epsilon/(2\beta)$, hence it cannot be further improved. That is, for post-processing via L_1 projection, the actual differential privacy protection achieved is $\epsilon/(2\beta)$ when $\beta > 1/2$ (and $\epsilon/(2(1-\beta))$ when $\beta < 1/2$). This makes intuitive sense. For example, when $\beta = 1$ the injected noise is drawn from a single $Lap(2\epsilon^{-1})$ distribution, corresponding to a privacy loss budget of $\epsilon/2$.

In summary, for any $\beta \in [0,1]$, the attained privacy loss budget for $M(x) = s(x) + \tilde{u}$ is $\epsilon/(2\max\{\beta, 1-\beta\})$.

B SAMPLING SCHEME FOR THE DOUBLE GEOMETRIC DISTRIBUTION

The Double Geometric mechanism, as introduced in Definition 1.2, utilizes additive noise whose cumulative mass function is given by

$$F(u) = P(U \le u) = \begin{cases} \frac{a^{-u}}{1+a} & u \le 0, \\ 1 - \frac{a^{u+1}}{1+a} & u > 0, \end{cases}$$

with quantile function

$$F^{-1}(v) = \begin{cases} \left\lceil \frac{-\log v - \log(1+a)}{\log(a)} \right\rceil & v \le \frac{1}{1+a}, \\ \left\lfloor \frac{\log(1-v) + \log(1+a)}{\log(a)} \right\rfloor & v > \frac{1}{1+a}. \end{cases}$$

Hence, one way to sample a Double geometric random variable is via inverse probability sampling. That is,

$$U_i \sim Unif(0,1), \qquad F^{-1}(U_i) \sim p_i(\cdot \mid \epsilon),$$

where p_i is given by (1.3). This method is implemented for all numerical examples illustrated in this paper.

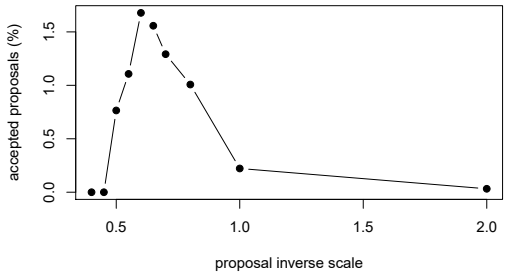

Figure 2: Algorithm 1 acceptance rate as a function of the proposal parameter $\tilde{\epsilon}$.

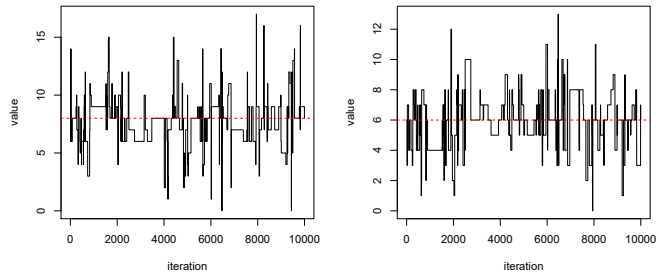

Figure 3: Traceplots of $10,000$ draws from Algorithm 1 of the second (left) and the first (right) cell of the constrained differentially private contingency table.

C PERFORMANCE DIAGNOSTICS OF THE MIS ALGORITHM

The acceptance rate of Algorithm 1 rate is shown in Figure 2 as a function of the proposal inverse scale parameter $\tilde{\epsilon}$. The acceptance rate is the highest in this example when $\tilde{\epsilon}$ is set to 0.6, just slightly larger than the privacy loss budget of the unconstrained privacy mechanism ($\epsilon = 0.5$ per cell). The acceptance rate achieved is about 1.68%.

Figure 3 shows traceplots of $10,000$ draws from Algorithm 1 of respectively the second (in proposal index set $\mathcal{I}$) and the first (not in proposal index set $\mathcal{I}$) cells of the constrained differentially private contingency table, when $\tilde{\epsilon} = 0.6$.