

SLOWLY SCALING PER-RECORD DIFFERENTIAL PRIVACY

BRIAN FINLEY¹, ANTHONY M CARUSO¹, JUSTIN C DOTY², ASHWIN MACHANAVAJJHALA³,
MIKAELA R MEYER⁴, DAVID PUJOL³, WILLIAM SEXTON³, AND ZACHARY TERNER⁴

¹U.S. Census Bureau

e-mail address: brian.finley@census.gov, anthony.m.caruso@census.gov

²Formerly U.S. Census Bureau

³Formerly Tumult Labs

⁴The MITRE Corporation

e-mail address: mrmeyer@mitre.org

ABSTRACT. We develop formal privacy mechanisms for releasing statistics from data with many outlying values, such as income data. These mechanisms ensure that a per-record differential privacy guarantee degrades slowly in the protected records' influence on the statistics being released.

Formal privacy mechanisms generally add randomness, or "noise," to published statistics. If a noisy statistic's distribution changes little with the addition or deletion of a single record in the underlying dataset, an attacker looking at this statistic will find it plausible that any particular record was present or absent, preserving the records' privacy. More influential records – those whose addition or deletion would change the statistics' distribution more – typically suffer greater privacy loss. The per-record differential privacy framework quantifies these record-specific privacy guarantees, but existing mechanisms let these guarantees degrade rapidly (linearly or quadratically) with influence. While this may be acceptable in cases with some moderately influential records, it results in unacceptably high privacy losses when records' influence varies widely, as is common in economic data.

We develop mechanisms with privacy guarantees that instead degrade as slowly as logarithmically with influence. These mechanisms allow for the accurate, unbiased release of statistics, while providing meaningful protection for highly influential records. As an example, we consider the private release of sums of unbounded establishment data such as payroll, where our mechanisms extend meaningful privacy protection even to very large establishments. We evaluate these mechanisms empirically and demonstrate their utility on simulated employment data and the U.S. Department of Agriculture's Cattle Inventory Survey.

Key words and phrases: differential privacy, establishment data, Rényi divergence, long-tailed, exponential polylogarithmic, transformation mechanism, additive mechanism.

Works (articles, reports, speeches, software, etc.) created by U.S. Government employees are not subject to copyright in the United States, pursuant to 17 U.S.C. §105. International copyright, 2026, U.S. Department of Commerce, U.S. Government. Any opinions and conclusions expressed herein are those of the authors and do not reflect the views of the U.S. Census Bureau. Approved for Public Release by The MITRE Corporation; Distribution Unlimited. Public Release Case Number 24-2528.

1. INTRODUCTION

Many government organizations have an obligation to release statistics about the public without violating the public’s confidentiality. For example, although the U.S. Census Bureau releases data products for public use, Title 13 of the U.S. Code prohibits the Census Bureau from publishing information that could disclose the data provided by any specific individuals or businesses (83rd United States Congress, 1954). Title 26 similarly prohibits the IRS from disclosing tax data (99th United States Congress, 1986), and Title 7 protects the data and identities of data providers to the U.S. Department of Agriculture (National Agricultural Statistics Service, 2023). Ensuring that data releases from the government preserve the confidentiality of its respondents is critical to the overall mission of serving the public.

To achieve this mission, some statistical agencies have begun to use differential privacy (Dwork et al., 2006) to provide mathematically provable guarantees on the privacy of respondents. For example, the U.S. Census Bureau has adopted differential privacy for the release of some of its data products, including those from the 2020 Census (Haney et al., 2021; U.S. Census Bureau, 2023; Abowd et al., 2022). The IRS has also adopted differential privacy to facilitate the release of the Department of Education’s College Scorecard website (Berghel et al., 2022).

The goal of differential privacy is to ensure that the presence or absence of any single record cannot be detected with high certainty from published statistics or queries. To accomplish this, differentially private systems add noise to queries to obfuscate whether any particular record is present. When a query’s value can be dramatically altered by the presence or absence of a single record, a large amount of noise is required to hide such a record. Adding enough noise to hide very influential records can ruin the utility of the release.

In these situations, one may use *per-record zero-concentrated differential privacy*, or PRzCDP (Seeman et al., 2024). Instead of enforcing a uniform privacy guarantee for all possible records, PRzCDP lets less influential records receive stronger privacy guarantees. If only a small number of very influential records are expected to be present, then PRzCDP allows for an improvement in utility by adding noise at a scale that provides a desired level of protection for a vast majority of records, while maintaining quantifiable, albeit lower, privacy guarantees for influential outliers. This gives data holders a finer ability to control the tradeoff between privacy and utility.

This is particularly useful when large, influential records have a lower expectation of confidentiality or when the increased utility outweighs the increased privacy loss. One such example comes in the form of establishment data, where large firms often must produce public statements about confidential values, and as such, are less susceptible to the risks of additional disclosure. In these cases, utility can be greatly increased while reducing the overall privacy guarantee only for the few records which have reduced privacy risk.

Existing mechanisms for PRzCDP use a method known as unit splitting, which incurs privacy losses that scale with the square of a record’s magnitude. This can be helpful in cases where the dataset is moderately skewed. However, in cases with very large outliers (e.g., the largest record is several orders of magnitude larger than the other observations), these methods can easily result in unacceptable privacy losses.

To address this, we introduce a class of mechanisms we call *slowly scaling mechanisms*, in which the privacy guarantee degrades slowly in the records’ influence on the query. Rather

than scaling with the square of a record’s magnitude, these privacy losses scale sublinearly in a record’s influence, ensuring that even very large records incur a reasonable privacy loss.

The paper continues as follows. In Section 2, we provide some foundational definitions and theorems of differential privacy, discuss per-record differential privacy, and describe the unit splitting mechanism for per-record differential privacy (Seeman et al., 2024). Section 3 introduces slowly scaling mechanisms and Sections 4 and 5 respectively detail our two classes of slowly scaling mechanisms: the transformation mechanisms and the additive mechanisms. Section 6 empirically studies our mechanisms with applications to two datasets. The paper concludes with Section 7. Code used for the empirical applications in Section 6 is available in the supplemental materials in Finley et al. (2025).

We make several contributions in this paper:

- We demonstrate the inadequacy of unit splitting in heavily skewed data, where some records may be orders of magnitude larger than others.
- We introduce the idea of *slowly scaling mechanisms*, which incur privacy losses that scale sublinearly in the influence of a record.
- We develop two classes of slowly scaling mechanisms: the transformation and additive mechanisms.
- We empirically demonstrate the utility of these mechanisms when applied to real data.

2. PRELIMINARIES

2.1. Data Model. We consider the table schema $\mathcal{D}(a_1, a_2, \dots, a_d)$, where $\mathcal{A} = \{a_1, a_2, \dots, a_d\}$ denotes the set of attributes of \mathcal{D} . The domain of each attribute a_i , denoted by $\text{Dom}(a_i)$, need not be finite or bounded. The full domain of \mathcal{D} is $\text{Dom}(\mathcal{D}) = \text{Dom}(a_1) \times \text{Dom}(a_2) \times \dots \times \text{Dom}(a_d)$. A database D is an instance of relation \mathcal{D} . That is, D is a multi-set of tuples $r = (x_1, \dots, x_d)$ where $x_i \in \text{Dom}(a_i)$. The number of tuples in D is denoted as $|D| = n$. We sometimes refer to a single record in a database as a “unit.”

2.2. Zero-Concentrated Differential Privacy. Here, we discuss the properties of formal privacy protection. We turn to a framework of privacy standards called differential privacy (DP) (Dwork et al., 2006). Informally, a randomized mechanism M satisfies differential privacy if the output distribution of the mechanism does not change too much with the addition or removal of a single individual’s record in the input dataset. This ensures that an adversary observing the output of M cannot determine with certainty whether any particular individual’s data was present in the input of M .

If, instead, the distribution of M changes little with the *alteration* of a single individual’s record, then the resulting differential privacy notion ensures that an adversary cannot use M ’s output to determine the value (as opposed to the presence) of any particular record. In different applications, it may be important to protect either the presence or values of records, or both. While the main body of this paper develops mechanisms to protect the presence of records, Appendix E analyzes the use of these mechanisms to protect records’ values as well.

In this paper, the per-record differential privacy notion that we primarily work with is a generalization of a widely used DP variant called *zero-concentrated differential privacy* (zCDP) (Bun and Steinke, 2016). zCDP quantifies the difference between output distributions

induced by adding or removing a single individual's record in terms of the Rényi divergence, defined below.

Definition 2.1 (Rényi Divergence (Rényi, 1961)). Let P and Q be probability distributions on some probability space Ω and let them respectively have the density functions p and q . Let the random variables $X_P \sim P$ and $X_Q \sim Q$. For $\alpha \in (1, \infty)$, we define the Rényi divergence of order α between P and Q as

$$d_\alpha(X_P \parallel X_Q) \equiv d_\alpha(P \parallel Q) \equiv \frac{1}{\alpha - 1} \ln \int_{\Omega} p(x)^\alpha q(x)^{1-\alpha} dx.$$

We also define the Rényi divergence of order ∞ as

$$d_\infty(X_P \parallel X_Q) \equiv d_\infty(P \parallel Q) \equiv \ln(\text{ess sup}_P \frac{p}{q}).$$

We begin by formalizing the notion of neighboring databases, which are databases that differ only in the presence of a single record.

Definition 2.2 (Neighboring Databases). Where \ominus denotes the symmetric difference, two databases $D, D' \in \mathcal{D}$ are considered *neighboring databases* if $D \ominus D' = \{r\}$ for some record r .

From here, we can define the formal notion of privacy under zCDP.

Definition 2.3 (Zero-Concentrated Differential Privacy). The randomized mechanism M satisfies ρ -zCDP for $\rho \geq 0$ iff, for any two neighboring databases $D, D' \in \mathcal{D}$,

$$d_\alpha(M(D) \parallel M(D')) \leq \rho\alpha \text{ for all } \alpha \in (1, \infty).$$

A mechanism which satisfies zCDP ensures that no record contributes too much to the output of a mechanism by ensuring that the Rényi divergence between the output distributions of the mechanism with and without any arbitrary record scales linearly in the order term α . The parameter ρ is often called the privacy loss and governs how much the two distributions are allowed to differ from one another. Large values of ρ allow the distributions to vary more and correspond to weaker privacy protections. Smaller values of ρ restrict how much the distributions can vary and result in stronger privacy protections.

Complex zCDP mechanisms are usually constructed by composing smaller primitive mechanisms together. To do this, we employ the following properties of zCDP: sequential composition, parallel composition, and post-processing invariance. Sequential composition allows multiple mechanisms which satisfy zCDP to compose together and jointly satisfy zCDP. In these cases, the final privacy loss is the sum of the individual privacy losses of each mechanism.

Theorem 2.4 (zCDP Sequential Composition (Bun and Steinke, 2016)). *Let M_1, M_2 be randomized mechanisms which satisfy ρ_1 -zCDP and ρ_2 -zCDP respectively. Then the mechanism $M'(D) = (M_1(D), M_2(D))$ satisfies $(\rho_1 + \rho_2)$ -zCDP.*

Additionally, under parallel composition, if two zCDP mechanisms are run on disjoint subsets of the database, their combination satisfies zCDP. In these cases, the final privacy loss is the maximum privacy loss of the component mechanisms.

Theorem 2.5 (zCDP Parallel Composition (Bun and Steinke, 2016)). *Let M_1, M_2 be randomized mechanisms which respectively satisfy ρ_1 -zCDP and ρ_2 -zCDP. Let D_1, D_2 be two disjoint subsets of a database, D . The mechanism $M'(D) = (M_1(D_1), M_2(D_2))$ satisfies $\max(\rho_1, \rho_2)$ -zCDP.*

As with all formal privacy definitions, zCDP is invariant to post-processing. That is, if the output of a mechanism satisfies zCDP, any post-processing on that output (without accessing the input data) preserves the zCDP guarantee.

Theorem 2.6 (zCDP Post-processing (Bun and Steinke, 2016)). *Let M be a randomized mechanism which satisfies ρ -zCDP. Let $M'(D) = f(M(D))$ for some arbitrary function f . Then M' satisfies ρ -zCDP.*

Using the composition and post-processing theorems, complex zCDP mechanisms are often built from simple primitive mechanisms, such as the Gaussian mechanism (Bun and Steinke, 2016). To define this mechanism, we must first define the sensitivity of a query.

We call the maximum change to a function due to the removal or addition of a single record the *sensitivity* of the function. This can be thought of as the maximal influence that a single record can have on the function.

Definition 2.7 (ℓ_2 -Sensitivity). Denote the ℓ_2 -norm by $|\cdot|_2$. The sensitivity of the vector function f is

$$\Delta \equiv \sup_{D', D \in \mathcal{D} \text{ such that } |D' \ominus D|=1} |f(D) - f(D')|_2.$$

The Gaussian mechanism simply adds Gaussian noise to a query, achieving a zCDP guarantee proportional to the square of the query’s sensitivity.

Definition 2.8 (Gaussian Mechanism (Bun and Steinke, 2016)). Let q be a query with sensitivity Δ . Consider the mechanism M that, on input D , releases a sample from $\mathcal{N}(q(D), \sigma^2)$. Then M satisfies $\frac{\Delta^2}{2\sigma^2}$ -zCDP.

While the definition of zCDP protects the privacy of an individual, these protections extend to groups of arbitrary size. The group privacy property says that, like an individual record, groups of records receive similar privacy protection, though the strength of that protection decays quadratically in the size of the group.

Theorem 2.9 (zCDP Group Privacy (Bun and Steinke, 2016)). *Let M be a randomized mechanism which satisfies ρ -zCDP. Then M guarantees $(k^2\rho)$ -zCDP for groups of size k . That is, for every pair of databases $D, D' \in \mathcal{D}$ satisfying $|D \ominus D'| = k$, we have the following.*

$$d_\alpha(M(D) \| M(D')) \leq (k^2\rho) \cdot \alpha \text{ for all } \alpha \in (1, \infty).$$

2.3. Per-Record Zero-Concentrated Differential Privacy. Here, we describe *per-record zero-concentrated differential privacy* (PRzCDP) (Seeman et al., 2024), a relaxation of zCDP designed for skewed data tasks.

Under PRzCDP, the privacy loss is parameterized in the form of a function of a record’s confidential value. For example, if a record contains strictly positive numerical values, we can consider a privacy loss that grows proportionally to that record value. To capture this relationship, PRzCDP employs a *record-dependent policy function*, which returns the privacy loss of any arbitrary record.

Definition 2.10 (Record-dependent policy function). A record-dependent policy function $P : \mathcal{T} \rightarrow \mathbb{R}_{\geq 0}$ denotes a maximum allowable privacy loss associated with a particular record value $r \in \mathcal{T}$, where \mathcal{T} is the universe of possible records.

Since the policy function's value $P(r)$ is dependent on the private record r , the value $P(r)$ cannot be released for any particular record without violating that record's privacy. We can, however, release the policy function P itself, which can be used to compute the privacy loss of any arbitrary and well-formed record. Given a policy function, per-record zero-concentrated differential privacy is defined as follows.

Definition 2.11 (*P-Per-Record Zero-Concentrated DP (P-PRzCDP)*). The randomized mechanism M satisfies *P-per-record zero-concentrated differential privacy (P-PRzCDP)* iff, for any two neighboring databases $D, D' \in \mathcal{D}$ satisfying $D \ominus D' = \{r\}$,

$$d_\alpha(M(D) \| M(D')) \leq \alpha P(r) \text{ for all } \alpha \in (1, \infty).$$

In this definition, the privacy loss associated with each record scales according to the policy function, as opposed to having equal privacy loss for all records. Traditional zCDP can be stated as a special case of *P-PRzCDP*, where the policy function is a constant. In the remainder of the paper, we will refer to the value $P(r)$ for a particular record r as that record's privacy loss.

2.4. Properties of Per-Record Differential Privacy. PRzCDP satisfies many of the same properties as zCDP, which allows complex mechanisms to be built from smaller primitive mechanisms. Of these properties, PRzCDP satisfies post-processing invariance, basic adaptive sequential composition, parallel composition, and group privacy.

Lemma 2.12 (Closure under post-processing). *Let \mathcal{T}^* be a multiset of records in \mathcal{T} . Given $M : \mathcal{T}^* \rightarrow \mathcal{Y}$, a *P-PRzCDP* mechanism M , and any function f , it is the case that $f \circ M$ is also *P-PRzCDP*.*

PRzCDP satisfies *basic adaptive* sequential composition. That is, as long as the policy functions are chosen prior to running any of the mechanisms, any two adaptively chosen mechanisms in combination satisfy PRzCDP as follows.

Lemma 2.13 (Basic adaptive sequential composition for *P-PRzCDP*). *Let M_1 satisfy P_1 -PRzCDP, and let M_2 satisfy P_2 -PRzCDP. Then $M_3(D) = M_2(M_1(D), D)$ satisfies $(P_1(r) + P_2(r))$ -PRzCDP.*

Just as with zCDP, the privacy losses sum when the mechanisms are composed together in this way. In PRzCDP, since the privacy loss is in the form of the policy function, the equivalent is the sum of policy functions. PRzCDP also satisfies a form of parallel composition. When multiple mechanisms which satisfy *P-PRzCDP* are run on disjoint subsets of the database, the joint result also satisfies *P-PRzCDP*:

Lemma 2.14 (Parallel composition for *P-PRzCDP*). *Let T be a partition of size $J \in \mathbb{N} \cup \{\infty\}$ over the universe of possible records. That is,*

$$T \equiv \bigcup_{j=1}^J C_j, \quad C_i \cap C_j = \emptyset, i \neq j.$$

*Let \mathcal{D}_j be the space of all databases containing only records in C_j for $j \in [J]$. Let $\{M_j\}_{j=1}^J$ be mechanisms satisfying *P-PRzCDP* for databases $D_j \in \mathcal{D}_j$ for $j \in [J]$, respectively. Then:*

$$M(D) \equiv \{M_j(D \cap C_j) \mid j \in [J]\}$$

*satisfies *P-PRzCDP*. Note that the M_j s can depend on their respective C_j s, allowing for adaptivity.*

ID	Industry	Employees
1	Retail	5
2	Retail	5
3	Retail	10
4	Services	20
5	Hospitality	30
6	Technology	10,000

TABLE 1. An example database where each record represents a single establishment with its associated industry and number of employees. Applying unit splitting with a splitting threshold of 10 employees on Establishment 6 results in a privacy loss that is 1,000,000 times that of Establishment 1.

Note that Lemma 2.14 does not require the individual mechanisms $\{M_j\}_{j=1}^J$ to be the same for all $j \in [J]$, allowing for different analysis for each partition. PRzCDP satisfies the group privacy notion as well.

Lemma 2.15 (Simple group privacy for P -PRzCDP). *Consider a sequence of databases D_0, \dots, D_J where $D_0 = D$ and $D_j \ominus D_{j-1} = \{r_j\}$ for $j \in [J]$. Let M be a randomized mechanism satisfying P -PRzCDP. Then we have:*

$$d_\alpha(M(D_0) \| M(D_J)) \leq \alpha J \sum_{j=1}^J P(r_j). \quad (2.1)$$

While most of the results presented here will pertain to single queries, they can all be expanded to workloads of queries by using the composition theorems. For example, if a workload contains two queries, each query can be answered independently using a $P_1(r)$ and $P_2(r)$ PRzCDP mechanism, respectively. By Lemma 2.13 the combination of both mechanisms satisfies PRzCDP with a policy function $P(r) = P_1(r) + P_2(r)$. These lemmas apply to zCDP mechanisms as well, since they can be considered a special case of PRzCDP mechanisms with a constant policy function. For example, a low sensitivity query such as a count can be answered with a ρ -zCDP mechanism followed by a high sensitivity query such as a sum which is answered with a $P(r)$ -PRzCDP mechanism. The combination of the two satisfies $(\rho + P(r))$ -PRzCDP.

2.5. Unit Splitting. Unit splitting (Seeman et al., 2024) is a preprocessing step which can be used to ensure that traditional zCDP mechanisms satisfy PRzCDP. Informally, each record in the dataset is split into multiple records according to a splitting rule. Then, a traditional zCDP mechanism, such as the Gaussian mechanism, is used to compute the query on the split dataset. This results in a PRzCDP guarantee with a policy function of, $P(r) = \rho |A(r)|^2$ where $A(r)$ is the number of times record r was split and ρ is the privacy loss budget used for the Gaussian mechanism.

Example 2.16. Consider answering a sum over the Employees column on the sample data in Table 1 by unit splitting. Consider the splitting rule that splits each establishment into multiple rows containing at most 10 employees. In this case, the first three establishments are not split at all. Establishment 4 is split into two rows, Establishment 5 is split into three rows and Establishment 6 is split into 1,000 rows. Since the rows now have at maximum

10 employees, the sensitivity of the sum query is 10. If we answer the sum query using the Gaussian mechanism with $\sigma^2 = 50$, then the establishments that are not split incur a privacy loss of $P(r) = \rho = 1$. Establishment 4, which is split into 2 rows, incurs a privacy loss of $\rho|A(r)|^2 = 4$. Likewise, Establishment 5 incurs a privacy loss of 9, and Establishment 6 incurs a privacy loss of 1,000,000.

This method has several benefits. First, it allows a practitioner to introduce an upper bound to the contribution of any record where there may not have been one previously. Additionally, it allows a practitioner to split larger records with large contributions into multiple smaller ones with smaller contributions. This results in less noise overall for the query, with the consequence of larger privacy losses for larger records. However, in cases with heavily skewed data like Example 2.16, some records can have extremely large privacy losses. To mitigate this effect, we introduce novel slowly scaling mechanisms which ensure that even large records incur reasonable privacy losses.

3. SLOWLY SCALING MECHANISMS

Unit splitting has been shown to provide high utility in cases where each record’s contribution may be unbounded but on average is small (Adeleye et al., 2023). As shown in Example 2.16, records that are split into a small number of additional rows still incur reasonable privacy losses, but records that are very large incur unacceptable privacy losses. This is a result of the policy function which scales in the square of the record size. Even a record which is split into only 10 rows incurs 100 times the privacy loss of a record which is not split. This is particularly problematic in cases where there are highly skewed records which may be several orders of magnitude larger than the median record.

To address the problem of rapidly scaling privacy loss, we introduce a class of PRzCDP mechanisms called *slowly scaling mechanisms*. These mechanisms satisfy PRzCDP with a policy function that is *sublinear*. In the following section, we will introduce two classes of slowly scaling mechanisms: transformation mechanisms and additive mechanisms.

The transformation mechanisms first pass the query result through a transformation function. They then add Gaussian noise to the result and invert the transformation function. These mechanisms result in policy functions which scale in the square of the transformation function applied to a quantification of a record’s influence. We demonstrate below that this can result in policy functions which scale as slowly as the log-squared rate.

The additive mechanisms act similarly to the Gaussian mechanism in that they simply draw a sample from a noise distribution and add it to the query. These mechanisms, however, use distributions which have thicker tails than the Gaussian distribution, resulting in policy functions that can scale in the records’ influence as slowly as the log rate.

4. TRANSFORMATION MECHANISMS

In this section, we describe the transformation mechanism, a mechanism which privatizes a univariate, nonnegative, real-valued query under the constraints of PRzCDP. As input, the mechanism takes a query value $q(D) \in [0, \infty)$; an offset parameter $a \in \mathbb{R}$; a noise scale parameter $\sigma \in (0, \infty)$; a concave, strictly increasing transformation function $f : [a, \infty) \rightarrow \mathcal{F} \subseteq \mathbb{R}$; and an estimator $g : \mathcal{F} \rightarrow \mathcal{G} \subseteq \mathbb{R}$ that “undoes” the transformation once noise has been added.

Haney et al. (2017) uses a similar mechanism that adds Laplace noise to transformed queries to provide protections under a different privacy concept which they develop specially for use with linked employer-employee data. Concurrent, independent work by Webb et al. (2025) uses mechanisms including special cases of the transformation mechanism to provide protections under yet another privacy concept. In this paper, we contribute novel estimators g and an analysis of the algorithm’s PRzCDP guarantees. See Appendix E for an analysis of our slowly scaling mechanisms using a privacy concept closely related to the one used by Webb et al. (2025). This privacy concept complements PRzCDP nicely, by protecting the values of records, rather than their presence in the data.

The transformation mechanism operates by mapping the private query value $q(D)$ to $f(q(D) + a)$ and adds noise sampled from a mean-zero Gaussian distribution with standard deviation σ . This quantity, called \tilde{v} in Algorithm 1, is essentially a privatized version of $f(q(D) + a)$, rather than of $q(D)$. The algorithm finishes, then, by returning $g(\tilde{v})$, an estimator of $q(D)$. Different estimators can be used, depending on what properties are desired and which transformation function is used. We provide mean-unbiased estimators in Section 4.3 and median-unbiased estimators in Appendix B.

Algorithm 1: Transformation-based mechanism

1: **procedure** TRANSFORMATIONPRIVATIZE(Private query answer $q(D)$, offset parameter a , scale parameter σ , transformation function $f : [a, \infty) \rightarrow \mathcal{F} \subseteq \mathbb{R}$, estimator $g : \mathcal{F} \rightarrow \mathcal{G} \subseteq \mathbb{R}$)
2: $v \leftarrow f(q(D) + a)$
3: $\tilde{v} \leftarrow N(v, \sigma^2)$
4: $\tilde{S} \leftarrow g(\tilde{v})$
5: **output** \tilde{S}
6: **end procedure**

To derive this algorithm’s PRzCDP guarantee, we first recall a bound on the Rényi divergence between two homoscedastic Gaussian distributions.

Lemma 4.1 (Lemma 2.4 of Bun and Steinke (2016)). *Let $\mu, \nu \in \mathbb{R}$, $\sigma \in \mathbb{R}_+$, and $\alpha \in [1, \infty)$. Then*

$$d_\alpha(N(\mu, \sigma^2) \| N(\nu, \sigma^2)) = \frac{\alpha(\mu - \nu)^2}{2\sigma^2}.$$

We also introduce the per-record sensitivity. In the same way that the sensitivity in Definition 2.7 captures the maximal influence that *any* record can have on a query, this captures the maximal influence that a *particular* record can have on a query.

Definition 4.2 (Per-Record Sensitivity). The per-record sensitivity of the query $q : \mathcal{D} \rightarrow \mathbb{R}$ for record r is

$$\Delta(r) \equiv \sup_{D, D' \text{ such that } D \oplus D' = r} |q(D) - q(D')|.$$

We now codify the basic assumptions on the inputs to the algorithm then state and prove the algorithm’s privacy guarantee.

Assumption 4.3 (Assumptions on Transformation Mechanisms). *Assume the query value $q(D) \in [0, \infty)$; the offset parameter $a \in \mathbb{R}$; the noise scale parameter $\sigma \in (0, \infty)$; the*

transformation function $f : [a, \infty) \rightarrow \mathcal{F} \subseteq \mathbb{R}$ is concave and strictly increasing; and the estimator $g : \mathcal{F} \rightarrow \mathcal{G} \subseteq \mathbb{R}$.

Theorem 4.4 (PRzCDP Guarantees for Transformation Mechanisms). *Under Assumption 4.3, Algorithm 1($q(D), a, \sigma, f, g$) satisfies P -PRzCDP for $P(r) = \frac{|f(\Delta(r)+a)-f(a)|^2}{2\sigma^2}$.*

Proof. By the definition of PRzCDP, for neighboring data sets D and D' where $D \ominus D' = \{r\}$, we would like to bound from above $d_\alpha(M(D) \parallel M(D'))$ and $d_\alpha(M(D') \parallel M(D))$ by $\alpha P(r)$, for all α , where $M(\cdot)$ is the notation for Algorithm 1 with transformation function f , estimator g , and privacy policy parameters σ and a . By the symmetry of the neighbor relationship, we can proceed WLOG by establishing a bound only on $d_\alpha(M(D) \parallel M(D'))$.

To start, we note that $M(D) = g(f(q(D) + a) + N(0, \sigma^2))$. By the data processing inequality (Rényi, 1961), we have

$$d_\alpha(M(D) \parallel M(D')) \leq d_\alpha(f(q(D) + a) + N(0, \sigma^2) \parallel f(q(D') + a) + N(0, \sigma^2)).$$

From Lemma 4.1, we get

$$d_\alpha(f(q(D) + a) + N(0, \sigma^2) \parallel f(q(D') + a) + N(0, \sigma^2)) = \frac{\alpha |f(q(D) + a) - f(q(D') + a)|^2}{2\sigma^2}.$$

We rewrite the right-hand side of the above equation as

$$\frac{\alpha |f(q(D) + a) - f(q(D') + a)|^2}{2\sigma^2} = \frac{\alpha |f(q(D) + a) - f(q(D) + a + [q(D') - q(D)])|^2}{2\sigma^2}.$$

Noting that f is strictly increasing, we can bound the right-hand side above using the per-record sensitivity, $\Delta(r)$ (see Definition 4.2):

$$\frac{\alpha |f(q(D) + a) - f(q(D) + a + [q(D') - q(D)])|^2}{2\sigma^2} \leq \frac{\alpha |f(q(D) + a) - f(q(D) + a + \Delta(r))|^2}{2\sigma^2}.$$

Since f is concave and increasing, the right-hand side above is bounded by $\alpha P(r)$, our policy function in the theorem statement:

$$\frac{\alpha |f(q(D) + a) - f(q(D) + a + \Delta(r))|^2}{2\sigma^2} \leq \frac{\alpha |f(a) - f(a + \Delta(r))|^2}{2\sigma^2}.$$

Therefore, Algorithm 1 satisfies P -PRzCDP. \square

We can see that the policy function of Theorem 4.4 grows in the per-record sensitivity at a rate determined by the transformation function, f . With the variety of transformation functions available, this can capture many possible policy functions. For example, we can achieve a quadratically scaling policy function, like that of the unit splitting mechanism, by choosing the identity function $f(x) = x$ as the transformation. This results in the policy function $P(r) = \frac{\Delta(r)^2}{2\sigma^2}$. To get a more slowly scaling policy function, the transformation itself must scale more slowly as well. In the following sections, we give two examples of such mechanisms, the k th root and log transformation mechanisms.

4.1. Kth Root Transformation Mechanism $f(x) = \sqrt[k]{x}$.

Corollary 4.5. *Under the conditions of Assumption 4.3, with transformation function $f(x) = \sqrt[k]{x}$, Algorithm 1($q(D), a, \sigma, f, g$) satisfies P -PRzCDP with*

$$P(r) = \frac{(\sqrt[k]{\Delta(r)} + a - \sqrt[k]{a})^2}{2\sigma^2}$$

for any estimator g and for all $a \geq 0$.

Proof. This follows from a straightforward application of Theorem 4.4, with $f(x) = \sqrt[k]{x}$. In this case, f is clearly concave and strictly increasing on the range $[a, \infty)$ for $a \geq 0$ and $k \geq 1$. Therefore, this mechanism satisfies P -PRzCDP with

$$P(r) = \frac{|f(\Delta(r) + a) - f(a)|^2}{2\sigma^2} = \frac{(\sqrt[k]{\Delta(r)} + a - \sqrt[k]{a})^2}{2\sigma^2}.$$

□

The k th root transformation mechanism is the result of using the transformation function $f(x) = \sqrt[k]{x}$ for any $k \geq 1$. This results in a policy function that scales in the square of the k th root. Notable special cases include the square root transformation mechanism and the fourth root transformation mechanism which result in linear and square root policy functions respectively. To see how this results in a slowly scaling policy function, we apply the fourth root transformation to the records from Example 2.16.

Example 4.6. Consider answering a sum over the Employees column on the sample data in Table 1 using the transformation mechanism. Let the transformation function $f(x) = \sqrt[4]{x}$ and the offset parameter $a = 0$. Likewise, assume Gaussian noise with scale $\sigma = 2$. The per-record sensitivity of a sum is simply the value of the record being summed over – in this case, the record’s employee count.

The first two establishments incur the same privacy loss of $\frac{\sqrt[4]{5^2}}{8} \approx 0.3$. Establishment 3 incurs a privacy loss of $\frac{\sqrt[4]{10^2}}{8} \approx 0.4$. Establishment 4 incurs a privacy loss of $\frac{\sqrt[4]{20^2}}{8} \approx 0.6$. Establishment 5 incurs a privacy loss of $\frac{\sqrt[4]{30^2}}{8} \approx 0.7$. Establishment 6 incurs a privacy loss of $\frac{\sqrt[4]{10000^2}}{8} = 12.5$.

Compared to the unit splitting results, the privacy losses incurred by each establishment are much closer together, and large establishments incur significantly smaller privacy losses.

4.2. Log Transformation Mechanism $f(x) = \ln(x)$.

In order to achieve a policy function which scales even more slowly, we can use the \ln function as the transformation.

Corollary 4.7. *Under the conditions of Assumption 4.3, with transformation function $f(x) = \ln(x)$, Algorithm 1($q(D), a, \sigma, f, g$) satisfies P -PRzCDP with*

$$P(r) = \frac{(\ln(\Delta(r) + a) - \ln(a))^2}{2\sigma^2}$$

for any estimator g and for all $a > 0$.

Proof. This follows from a straightforward application of Theorem 4.4, with $f(x) = \ln(x)$. In this case, f is clearly concave and strictly increasing on the range $[a, \infty)$ for $a > 0$. Therefore, this mechanism satisfies P -PRzCDP with

$$P(r) = \frac{|f(\Delta(r) + a) - f(a)|^2}{2\sigma^2} = \frac{(\ln(\Delta(r) + a) - \ln(a))^2}{2\sigma^2}.$$

□

Example 4.8. Consider answering a sum over the Employees column on the sample data in Table 1 using the transformation mechanism. Let the transformation function $f(x) = \ln(x)$ and the offset parameter $a = 1$. Likewise, assume Gaussian noise with scale $\sigma = 2$. The per-record sensitivity of a sum is simply the value of the record being summed over – in this case, the record’s employee count.

The first two establishments incur the same privacy loss of $\frac{\ln(6)^2}{8} \approx 0.4$. Establishment 3 incurs a privacy loss of $\frac{\ln(11)^2}{8} \approx 0.7$. Establishment 4 incurs a privacy loss of $\frac{\ln(21)^2}{8} \approx 1.2$. Establishment 5 incurs a privacy loss of $\frac{\ln(31)^2}{8} \approx 1.5$. Establishment 6 incurs a privacy loss of $\frac{\ln(10001)^2}{8} \approx 10.6$.

Like in the case of the fourth root transformation, the privacy losses for all records are closer together, and large records incur far less extreme privacy losses. In this case, smaller records incur a larger privacy loss under the log transformation when compared to the fourth root transformation. However, since the privacy loss scales more slowly than the fourth root transformation mechanism, larger records incur a smaller privacy loss.

4.3. Unbiased Estimators (g). In this section, we present mean-unbiased estimators of $q(D)$ for a rich set of transformation functions. Mean-unbiasedness may be essential in settings where data users are expected to aggregate a large number of queries. For example, if block-level populations are published via a mean-unbiased mechanism, users can sum these to obtain mean-unbiased estimators of state or national populations. We also develop median-unbiased estimators in Appendix B.

Per Algorithm 1, the noisy transformed query value, \tilde{v} , is distributed according to $\tilde{v} \sim N(v, \sigma^2)$, with $v \equiv f(q(D) + a)$. Our problem, then, is to find an unbiased estimator of $f^{-1}(v)$; simply subtracting a from this will give an unbiased estimator of $q(D)$. To do this, we use estimators from Washio et al. (1956) to obtain unbiased output from Algorithm 1 when the transformation function is of the forms $f(x) = \ln(x)$ and $f(x) = \sqrt[k]{x}$ for positive integers k .

Theorem 4.9 (Mean-Unbiased Estimators (Washio et al., 1956)). *Let Assumption 4.3 hold. Denote the output of Algorithm 1($q(D), a, \sigma, f, g$) by $M(D)$. Denote the probabilist’s k^{th} Hermite polynomial by $He_k(x) \equiv (-1)^k e^{\frac{x^2}{2}} \frac{\partial^k}{\partial x^k} e^{-\frac{x^2}{2}}$.*

- If $a \geq 0$, k is a positive integer, $f(x) = \sqrt[k]{x}$, and $g(x) = (-\sigma)^k He_k(-\frac{x}{\sigma}) - a$, then $\mathbb{E}[M(D)] = q(D)$ and $\mathbb{V}[M(D)] = \sum_{i=0}^{k-1} \binom{k}{i}^2 (k-i)! \sigma^{2(k-i)} (q(D) + a)^{\frac{2i}{k}}$.
- If $a > 0$, $f(x) = \ln(x)$, and $g(x) = e^{x - \frac{\sigma^2}{2}} - a$, then $\mathbb{E}[M(D)] = q(D)$ and $\mathbb{V}[M(D)] = (e^{\sigma^2} - 1)(q(D) + a)^2$.

It is worth noting that Theorem 1 of Washio et al. (1956) provides a formula that could be used to obtain mean-unbiased estimators for a very general class of transformation

Name	$f(x)$	$g(x)$	$\mathbb{V}[g(\tilde{v})]$
Identity	x	$x - a$	σ^2
Square Root	\sqrt{x}	$x^2 - \sigma^2 - a$	$2\sigma^4 + 4\sigma^2(q(D) + a)$
Fourth Root	$\sqrt[4]{x}$	$x^4 - 6x^2\sigma^2 + 3\sigma^4 - a$	$24\sigma^8 + 96\sigma^6(q(D) + a)^{\frac{1}{2}}$ $+ 72\sigma^4(q(D) + a) + 16\sigma^2(q(D) + a)^{\frac{3}{2}}$
k th Root	$\sqrt[k]{x}$	$(-\sigma)^k \text{He}_k(-\frac{x}{\sigma}) - a$	$\sum_{i=0}^{k-1} \binom{k}{i}^2 (k-i)! \sigma^{2(k-i)} (q(D) + a)^{\frac{2i}{k}}$
Log	$\ln(x)$	$e^{x - \frac{\sigma^2}{2}} - a$	$(e^{\sigma^2} - 1)(q(D) + a)^2$

TABLE 2. Mean-unbiased estimators and variances for selected transformation mechanisms

functions. This formula is unwieldy, however, involving an inverse Laplace transform. We restrict our analysis to the analytically convenient transformation functions treated already (see Table 2), and note that they induce an array of policy functions that should be wide enough for practical use.

An important feature of all the estimators discussed here is that, due to their nonlinearity, the shape and scale of the transformation mechanism’s distribution become dependent on the true query value, $q(D)$. In particular, the mechanisms’ variances tend to grow in $q(D)$. This can be seen in the variance formulae in Theorem 4.9. It is also important to see that the mechanism generally has thicker right tails when it uses a more slowly scaling transformation. Using the unbiased estimators, we can see this behavior in Table 2; as f becomes more slowly scaling, g raises the noisy transformed query to higher and higher powers and eventually applies an exponential function to it. This leads to an increasingly thick right tail, with the unbiased squared root mechanism having the tail thickness of a χ^2 distribution and the unbiased log mechanism having the tail thickness of a log-normal distribution. Independently of the variance, then, transformation mechanisms that scale more slowly will also be more likely to result in extreme positive outliers.

5. ADDITIVE MECHANISMS

In this section, we develop a class of additive mechanisms which achieve slowly scaling policy functions by drawing noise from relatively fat-tailed distributions. An additive mechanism simply adds data-independent random noise to the query. Algorithm 2 represents a generic additive mechanism.

Algorithm 2: Additive mechanism

- 1: **procedure** ADDITIVEPRIVATIZE(Private query answer $q(D)$, probability density function $f_Z : \mathbb{R} \rightarrow [0, \infty)$)
 - 2: $Z \sim f_Z$
 - 3: $\tilde{S} \leftarrow q(D) + Z$
 - 4: **output** \tilde{S}
 - 5: **end procedure**
-

As an example, the Gaussian mechanism from Definition 2.8 is an additive mechanism. As discussed in Section 4 (as a special case of the transformation mechanism), this has

the policy function $P(r) = \frac{\Delta(r)^2}{\sigma^2}$ which grows at the same rapid rate as the unit splitting algorithm’s policy function. The additive mechanisms developed in this section add noise, Z , with densities of the form $f_Z(z) \propto e^{f(|z|)}$, where f is a decreasing, convex function. These distributions are symmetric about 0, and the convexity of f gives these densities a kink at 0 as well as tails that are thicker than Gaussian. This tail thickness makes these mechanisms’ policy functions scale slowly. For intuition as to why, recall that privacy loss is determined by the worst-case divergence between the mechanism’s distributions at each of two neighboring databases. For an additive mechanism, the distributions realizing this worst-case divergence are the same up to a location shift equal to the per-record sensitivity. As the sensitivity grows, these distributions move apart, leading to greater divergence and greater privacy loss. When the distributions have fatter tails, however, their densities can still have substantial overlap even after large location shifts, leading to a lower divergence, and therefore lower privacy loss. This slows the growth of the privacy loss (i.e., the policy function) in the sensitivity, giving us slowly scaling behavior. Of course, as in the case of the transformation mechanism, fatter tails raise the probability of the mechanism returning extreme outlying values. This cost to utility is inherent to all the slowly scaling mechanisms in this paper.

The additive mechanisms offer an advantage in that they satisfy per-record differential privacy (PRDP), a stronger privacy concept than PRzCDP. This concept straightforwardly formalizes the suggestion by Seeman et al. (2024) (page 3140) of a “pure ϵ -DP” analogue of PRzCDP. In work published after the working version of this paper, Calmon et al. (2025) derive unbiased estimators for use with Laplace noise and introduce transformation mechanisms that satisfy PRDP by adding Laplace (rather than Gaussian) noise to transformed queries. See their Section 6.

We now define PRDP.

Definition 5.1 (*P-Per-Record Differential Privacy (P-PRDP)*). The randomized mechanism M satisfies P -per-record differential privacy (P -PRDP) iff, for any two neighboring databases $D, D' \in \mathcal{D}$ satisfying $D \ominus D' = \{r\}$,

$$d_\infty(M(D) \parallel M(D')) \leq P(r).$$

This privacy guarantee is stronger than PRzCDP in the sense that PRzCDP is implied by PRDP. To compare the additive mechanisms’ privacy guarantees with the PRzCDP guarantees of our transformation mechanisms, we derive the PRzCDP guarantees implied by the additive mechanisms’ PRDP guarantees.

Lemma 5.2 (*PRDP Implies PRzCDP*). *If a mechanism, M , is P -PRDP, it is also P -PRzCDP and P' -PRzCDP, where $P'(r) = \tanh(P(r)/2)P(r)$.*

Proof. By Theorem 5 of Steinke (2024), for any fixed neighboring databases and for all $\alpha \in (1, \infty)$ the following holds.

$$d_\infty(M(D) \parallel M(D')) \leq \epsilon \implies d_\alpha(M(D) \parallel M(D')) \leq \alpha \frac{e^\epsilon - 1}{e^\epsilon + 1} \epsilon = \alpha \tanh(\epsilon/2) \epsilon \leq \alpha \epsilon$$

Since this applies to any arbitrary neighboring databases the following is immediate. For any neighboring databases D, D' where $D \ominus D' = \{r\}$ the following holds:

$$d_\infty(M(D) \parallel M(D')) \leq P(r) \implies d_\alpha(M(D) \parallel M(D')) \leq \alpha \tanh(P(r)/2) P(r) \leq \alpha P(r)$$

This yields P' -PRzCDP where $P'(r) = \tanh(P(r)/2)P(r)$ or, as a looser bound, $P'(r) = P(r)$. \square

To derive the privacy guarantees of our additive mechanisms, we use intermediate results from the following two theorems. The first of these is a technical result establishing that, to determine the PRDP guarantee of an additive mechanism, we need to find bounds on the absolute log probability ratio of the noise distribution centered at 0 and the noise distribution centered at a value no greater than the per-record sensitivity.

Theorem 5.3. *An additive mechanism $M(D) \equiv q(D) + Z$, with a continuous random variable Z having density f_Z , satisfies P -PRDP if and only if, for all neighboring databases $D, D' \in \mathcal{D}$ satisfying $D \ominus D' = \{r\}$,*

$$\sup_z |\ln(f_Z(z)) - \ln(f_Z(z + q(D) - q(D')))| \leq P(r).$$

Proof. First, we note that the ∞ -Rényi divergence between the distributions of the mechanism evaluated at the two databases D and D' is

$$d_\infty(M(D) \| M(D')) = \ln \left(\sup_y \frac{f_Z(y - q(D))}{f_Z(y - q(D'))} \right) \quad (5.1)$$

$$= \sup_y \ln(f_Z(y - q(D))) - \ln(f_Z(y - q(D'))). \quad (5.2)$$

By the symmetry of the neighbor relation used in defining P -PRDP (that is, the fact that $D \ominus D' = D' \ominus D$), M is P -PRDP if and only if, for all D, D' such that $D \ominus D' = r$,

$$\max [d_\infty(M(D) \| M(D')), d_\infty(M(D') \| M(D))] \leq P(r). \quad (5.3)$$

Substituting Equation 5.2 into the left-hand-side of Inequality (5.3), we can simplify and substitute in $z = y - q(D)$ to obtain Theorem 5.3 as follows:

$$\begin{aligned} & \max [d_\infty(M(D) \| M(D')), d_\infty(M(D') \| M(D))] \\ &= \max \left[\sup_y \ln(f_Z(y - q(D))) - \ln(f_Z(y - q(D'))), \right. \\ & \quad \left. \sup_y \ln(f_Z(y - q(D'))) - \ln(f_Z(y - q(D))) \right]. \quad (5.4) \end{aligned}$$

Noting that the sup and max commute, and factoring -1 out of the second expression, we rearrange the above to get

$$\begin{aligned} &= \sup_y \max [\ln(f_Z(y - q(D))) - \ln(f_Z(y - q(D'))), \\ & \quad - [\ln(f_Z(y - q(D))) - \ln(f_Z(y - q(D')))]] \quad (5.5) \end{aligned}$$

$$= \sup_y |\ln(f_Z(y - q(D))) - \ln(f_Z(y - q(D')))| \quad (5.6)$$

$$= \sup_z |\ln(f_Z(z)) - \ln(f_Z(z + q(D) - q(D')))|. \quad (5.7)$$

□

The second result, in Theorem 5.4 below, is the cornerstone of our analysis. It shows that the supremal absolute log probability ratio from Theorem 5.3 is easy to find when the noise distribution's density is symmetric and log-convex on either side of the mode. This is the class of distributions we consider for use in our slowly scaling additive mechanisms.

Theorem 5.4 (Bounds on Log-Probability Ratio for Exponential Absolute Convex Densities). *Consider the probability density $f_Z(z) \propto e^{f(|z|)}$ for $z \in \mathbb{R}$, where $f : [0, \infty) \rightarrow \mathbb{R}$ is (weakly) decreasing and (weakly) convex. The log probability ratio of this distribution and this distribution location-shifted by l is tightly bounded as follows:*

$$\sup_z |\ln(f_Z(z)) - \ln(f_Z(z-l))| = f(0) - f(|l|).$$

Proof. Denote the supremand in Theorem 5.4 by $g(z;l) \equiv |\ln(f_Z(z)) - \ln(f_Z(z-l))|$. First, note that

$$g(z;l) = |f(|z|) - f(|z-l|)|.$$

Now, observe that $g(z;l) = g(-z;-l)$. This implies that $\sup_z g(z;l) = \sup_z g(z;|l|)$ because $\sup_z g(z;l) = \sup_z g(-z;-l) = \sup_z g(z;-l)$.

Similarly, observe that $g(z;l)$ is symmetric about $z = .5l$. That is, $g(z;l) = g(.5l - (z - .5l); l) = g(l - z; l)$. We can therefore limit our search for the maximizing value of z to $z \geq .5l$.

Combining these two observations, our problem is now to solve

$$\sup_{z \geq .5|l|} |f(|z|) - f(|z - |l||)|.$$

To do this, we will show that $g(z;|l|) \leq g(l;|l|)$ separately for the two cases where $z \in [.5|l|, |l|]$ and where $z \in [|l|, \infty]$. This will conclude our proof.

To simplify notation in analyzing the two cases, we will assume WLOG that $l \geq 0$, so that $|l| = l$.

First case: $z \in [.5l, l]$. For $z \in [.5l, l]$, we have

$$g(z;l) = |f(|z|) - f(|z-l|)| = |f(z) - f(l-z)|.$$

Further, $z \in [.5l, l]$ and $l \geq 0$ implies that $0 \leq l-z \leq .5l \leq z$. Because f is a decreasing function, $f(z) \leq f(l-z)$, and we have

$$g(z;l) = |f(z) - f(l-z)| = f(l-z) - f(z).$$

Because f is decreasing, both $f(l-z)$ and $-f(z)$ are increasing in z . For $z \in [.5l, l]$, then, $g(z;l)$ is maximized at $z = l$, where

$$g(l;l) = f(0) - f(l). \tag{5.8}$$

Second case: $z \in [l, \infty]$. In the case where $z \in [l, \infty)$, we have, by reasoning like that in the first case,

$$g(z; l) = |f(|z|) - f(|z - l|)| = |f(z) - f(z - l)| = f(z - l) - f(z). \quad (5.9)$$

We want to show that $g(z; l) \leq g(l; l)$ for all $z \in [l, \infty)$. Plugging equations 5.8 and 5.9 into this inequality, this means we want to demonstrate that

$$f(z - l) - f(z) \leq f(0) - f(l) \text{ or, equivalently,} \quad (5.10)$$

$$f(l) + f(z - l) \leq f(z) + f(0). \quad (5.11)$$

We now invoke the convexity of f (and recall that $z \geq l$) to derive the following two inequalities.

$$f(l) = f\left(\frac{l}{z}z + \left(1 - \frac{l}{z}\right)0\right) \leq \frac{l}{z}f(z) + \left(1 - \frac{l}{z}\right)f(0). \quad (5.12)$$

$$f(z - l) = f\left(\frac{z - l}{z}z + \left(1 - \frac{z - l}{z}\right)0\right) \leq \frac{z - l}{z}f(z) + \left(1 - \frac{z - l}{z}\right)f(0). \quad (5.13)$$

Summing inequalities 5.12 and 5.13 gives us

$$\begin{aligned} f(l) + f(z - l) &\leq \frac{l}{z}f(z) + \left(1 - \frac{l}{z}\right)f(0) + \frac{z - l}{z}f(z) + \left(1 - \frac{z - l}{z}\right)f(0) \\ &= \frac{l}{z}f(z) + \frac{z - l}{z}f(0) + \frac{z - l}{z}f(z) + \left(1 - \frac{z - l}{z}\right)f(0) \\ &= f(z) + f(0). \end{aligned} \quad (5.14)$$

This is (5.11), proving the second case. □

The following is our main result, giving the PRDP and PRzCDP guarantees for the additive mechanisms we consider.

Theorem 5.5 (Privacy of Symmetric, Absolute-Log-Convex Densities). *Let q be a query with per-record sensitivity $\Delta(r)$, as defined in Definition 4.2. Let f_Z be a noise distribution that has density proportional to $e^{f(|z|)}$, where $f : [0, \infty) \rightarrow \mathbb{R}$ is (weakly) decreasing and (weakly) convex.*

For all records r , Algorithm $\mathcal{Q}(q(D), f_Z)$, denoted below by M , satisfies P -PRDP with $P(r) = f(0) - f(\Delta(r))$. This policy function is tight, in the sense that M is not P' -PRDP for an alternative policy function P' such that there exists an r satisfying $P'(r) < P(r)$. M also satisfies P'' -PRzCDP with $P''(r) = \tanh(P(r)/2)P(r)$.

Proof. Applying Theorem 5.3 and then Theorem 5.4, we see that this algorithm is P -PRDP if and only if, for all r ,

$$P(r) \geq \sup_{D, D' \text{ such that } D \ominus D' = r} \sup_z |\ln(f_Z(z)) - \ln(f_Z(z + q(D) - q(D')))| \quad (5.15)$$

$$= \sup_{D, D' \text{ such that } D \ominus D' = r} f(0) - f(|q(D) - q(D')|). \quad (5.16)$$

Because f is a decreasing function, $f(0) - f(|q(D) - q(D')|)$ is increasing in $|q(D) - q(D')|$ and has its supremum at $|q(D) - q(D')| = \Delta(r)$. Therefore, it satisfies P -PRDP with

$P(r) = f(0) - f(\Delta(r))$. Furthermore, by Lemma 5.2, it also satisfies P' -PRzCDP with $P'(r) = \tanh(P(r)/2)P(r)$. \square

The appeal of this class of mechanisms lies in the ability to flexibly choose $P(r)$ by choosing f . Intuitively, the more convex f is, the fatter the noise distribution's tails are, and the more slowly privacy loss grows. The next sections provide examples of slowly scaling additive mechanisms for different choices of f . In the first example, we show that a slowly scaling mechanism can be achieved by adding noise from a generalized Gaussian distribution, while our second example uses a distribution that, to our knowledge, is novel.

5.1. Generalized Gaussian Mechanism $\left(f(|z|) = -\left(\frac{|z|}{\sigma}\right)^p\right)$.

Corollary 5.6. *Under the conditions of Theorem 5.5, with density function $f_Z(z) = \frac{p}{2\sigma\Gamma(\frac{1}{p})}e^{-\left(\frac{|z|}{\sigma}\right)^p}$, where $0 < \sigma$ and $0 \leq p \leq 1$, Algorithm $\mathfrak{A}(q(D), f_Z)$ satisfies P -PRDP and P' -PRzCDP where $P(r) = \left(\frac{\Delta(r)}{\sigma}\right)^p$ and $P'(r) = \tanh(P(r)/2)P(r)$.*

Proof. This follows from a straightforward application of Theorem 5.5, with $f(|z|) = -\left(\frac{|z|}{\sigma}\right)^p$. In this case, $f(|z|)$ is clearly decreasing in $|z| \in [0, \infty)$, and, for $0 \leq p \leq 1$, $f(|z|)$ is convex. For this range of p , then, this mechanism satisfies P -PRDP and P' -PRzCDP where $P(r) = f(0) - f\left(\frac{\Delta(r)}{\sigma}\right) = \left(\frac{\Delta(r)}{\sigma}\right)^p$ and $P'(r) = \tanh(P(r)/2)P(r)$. \square

The mechanism in Corollary 5.6 has a generalized Gaussian distribution with variance $\sigma^2 \frac{\Gamma(\frac{3}{p})}{\Gamma(\frac{1}{p})}$, per Nadarajah (2005). The generalized Gaussian mechanism has been used in the context of differential privacy before (Liu, 2019; Ganesh and Zhao, 2021), but our use of it to achieve a slowly scaling policy function for per-record variants of DP is novel.

5.2. Exponential Polylogarithmic Distribution $\left(f(|z|) = -d \ln\left(\frac{|z|}{\sigma} + a\right)^p\right)$.

Corollary 5.7. *Under the conditions of Theorem 5.5, with density function $f_Z(z) \propto e^{-d \ln\left(\frac{|z|}{\sigma} + a\right)^p}$, where $\sigma > 0$, $a \geq e^{p-1}$, and either $(p > 1$ and $d > 0)$ or $(p = 1$ and $d > 1)$, Algorithm $\mathfrak{A}(q(D), f_Z)$ satisfies P -PRDP and P' -PRzCDP with $P(r) = d \left[\ln\left(\frac{\Delta(r)}{\sigma} + a\right)^p - \ln(a)^p \right]$ and $P'(r) = \tanh(P(r)/2)P(r)$.*

Proof. This result follows from Theorem 5.5 once we show that $f(|z|) = -d \ln\left(\frac{|z|}{\sigma} + a\right)^p$ is decreasing and convex. First, the function is decreasing in $|z|$ due to $-\ln\left(\frac{|z|}{\sigma} + a\right)$ being non-positive and decreasing whenever $1 \leq a$. This holds because $p \geq 1$ and $e^{p-1} \leq a$. Second, the function's second derivative,

$$-\frac{dp}{\left(\frac{|z|}{\sigma} + a\right)^2} \left[(p-1) \ln\left(\frac{|z|}{\sigma} + a\right)^{p-2} - \ln\left(\frac{|z|}{\sigma} + a\right)^{p-1} \right],$$

is non-negative when $p-1 \leq \ln\left(\frac{|z|}{\sigma} + a\right)$. This condition holds when $e^{p-1} - a \leq \frac{|z|}{\sigma}$, which follows from the assumption that $e^{p-1} \leq a$ in Corollary 5.7.

□

We call the noise distribution used in Corollary 5.7 the *exponential polylogarithmic distribution*. This distribution is novel, to the best of our knowledge. It has a tractable density function, CDF, quantile function, and variance when $p \in \{1, 2\}$, as seen in Table 3. See Appendix F for the general case, where these quantities involve obscure special functions.

	$p = 1$	$p = 2$
Density	$\frac{d-1}{2\sigma} a^{d-1} \left(\frac{ z }{\sigma} + a \right)^{-d} \{d > 1\}$	$\frac{\sqrt{d} e^{-d \ln(\frac{ z }{\sigma} + a)^2}}{2e^{\frac{1}{4d}} \sigma \sqrt{\pi} [1 - \kappa(a, \frac{1}{2})]}$
CDF	$\frac{1}{2} + \frac{\text{sign}(z)}{2} \times (1 - (\frac{ z }{\sigma a} + 1)^{1-d}) \{d > 1\}$	$\frac{1}{2} + \frac{\text{sign}(z)}{2} \left[\frac{\kappa(\frac{ z }{\sigma} + a, \frac{1}{2}) - \kappa(a, \frac{1}{2})}{1 - \kappa(a, \frac{1}{2})} \right]$
Quantile	$\sigma a \text{sign}(q - 0.5) \left([1 - 2 q - 0.5]^{\frac{1}{1-d}} - 1 \right) \{d > 1\}$	$\sigma \text{sign}(q - 0.5) \left(\exp \left[(2d)^{-\frac{1}{2}} \times \Phi^{-1} \left(2 q - 0.5 \left[1 - \kappa \left(a, \frac{1}{2} \right) \right] + \kappa \left(a, \frac{1}{2} \right) \right) + (2d)^{-1} \right] - a \right)$
Variance	$\sigma^2 a^2 (d - 1) \left[\frac{1}{d-3} - \frac{2}{d-2} + \frac{1}{d-1} \right] \{d > 3\}$	$\sigma^2 \frac{e^{\frac{2}{d}} [1 - \kappa(a, \frac{3}{2})] - 2ae^{\frac{3}{4d}} [1 - \kappa(a, 1)]}{[1 - \kappa(a, \frac{1}{2})]} + a^2 \sigma^2$

TABLE 3. Special cases of the density, CDF, quantile, and variance of the exponential polylogarithmic distribution for $p = 1$ and $p = 2$ where $q \in [0, 1]$, $\kappa(x, y) = \Phi \left(\frac{\ln(x) - yd^{-1}}{(2d)^{-1/2}} \right)$, and $\Phi(\cdot)$ is the CDF of a standard normal distribution. The notation $\{d > .\}$ in the second column indicates conditions required for each cell's quantity to exist or be finite.

Given that this is an atypical family of distributions, it is helpful to visualize its density alongside the more commonly used normal distribution. Figure 1 plots PDFs for random values sampled from equal variance and equal privacy loss distributions. Tail weights for noise from equal privacy loss mechanisms are given in Table 4.




FIGURE 1. Probability densities for two exponential polylogarithmic distributions and a normal distribution. Distributions in plot (a) have unit variance. Additive mechanisms using distributions in plot (b) have PRzCDP privacy loss of 1 when $\Delta(r) = 1$ (that is, mechanism i satisfies P_i -PRzCDP such that $P_i(1) = 1$). Note the thicker tails on the exponential polylogarithmic distributions.

Mechanism	$P[z > 1]$	$P[z > 2]$	$P[z > 3]$	$P[z > 4]$
Exp polylog, $\sigma \approx 0.708$, $a = 3$, $d = 4$, $p = 1$	0.314	0.137	0.071	0.042
Exp polylog, $\sigma \approx 1.877$, $a = 3$, $d = 4$, $p = 2$	0.221	0.051	0.013	0.003
Gaussian, $\sigma \approx 0.707$	0.157	0.005	0.000	0.000

TABLE 4. Tail weights for noise from additive mechanisms applied to sum queries with parameters set such that PRzCDP privacy loss is 1 when $\Delta(r) = 1$.

6. EMPIRICAL RESULTS

In this section, we empirically evaluate the mechanisms presented above. In the following subsections, we briefly specify the datasets and queries we use, the mechanisms we chose, and the metrics we calculate before discussing the results.

6.1. Datasets. We conduct these experiments using groupby sums on two different datasets. The first, provided to us by the U.S. Census Bureau, is a simulated version of the data used to compute the Census Bureau’s County Business Patterns (CBP) data product. Each record in the data corresponds to a single establishment (i.e., a single location of a single business).

The groupby sums for the CBP data are grouped by the attributes NAICS3 (the three-digit North American Industry Classification System Code) and COUNTY. These sums

Transformation Mechanism with Transformation f

Name	$f(x)$	Variance	PRzCDP Policy Function
Identity	x	σ^2	$\frac{\Delta(r)^2}{2\sigma^2}$
Square Root	\sqrt{x}	$2\sigma^4 + 4\sigma^2(q(D) + a)$	$\frac{(\sqrt{\Delta(r)+a}-\sqrt{a})^2}{2\sigma^2}$
Fourth Root	$\sqrt[4]{x}$	$24\sigma^8 + 96\sigma^6(q(D) + a)^{\frac{1}{2}} + 72\sigma^4(q(D) + a) + 16\sigma^2(q(D) + a)^{\frac{3}{2}}$	$\frac{(\sqrt[4]{\Delta(r)+a}-\sqrt[4]{a})^2}{2\sigma^2}$
k th Root	$\sqrt[k]{x}$	$\sum_{i=0}^{k-1} \binom{k}{i}^2 (k-i)! \sigma^{2(k-i)} (q(D) + a)^{\frac{2i}{k}}$	$\frac{(\sqrt[k]{\Delta(r)+a}-\sqrt[k]{a})^2}{2\sigma^2}$
Log	$\ln(x)$	$(e^{\sigma^2} - 1)(q(D) + a)^2$	$\frac{(\ln(\Delta(r)+a)-\ln(a))^2}{2\sigma^2}$

Additive Mechanism with Noise Density $\propto e^{f(|z|)}$

Name	$f(z)$	Variance	PRDP Policy Function
Generalized Gaussian	$-\left(\frac{ z }{\sigma}\right)^p$	$\sigma^2 \frac{\Gamma(\frac{3}{p})}{\Gamma(\frac{1}{p})}$	$\left(\frac{\Delta(r)}{\sigma}\right)^p$
Exponential Polylog ($p = 2$)	$-d \ln\left(\frac{ z }{\sigma} + a\right)^2$	$\sigma^2 \frac{e^{\frac{2}{d}[1-\kappa(a, \frac{3}{2})]} - 2ae^{\frac{3}{4d}[1-\kappa(a, 1)]}}{[1-\kappa(a, \frac{1}{2})]} + a^2\sigma^2$	$d(\ln(\frac{\Delta(r)}{\sigma} + a)^2 - \ln(a)^2)$

TABLE 5. Variance and policy functions for mean-unbiased transformation mechanisms and additive mechanisms. Here, $\kappa(x, y) = \Phi\left(\frac{\ln(x)-yd^{-1}}{(2d)^{-1/2}}\right)$, and $\Phi(\cdot)$ is the CDF of a standard normal distribution. The PRDP policy functions of the additive mechanisms are also valid PRzCDP policy functions (see Lemma 5.2), though tighter PRzCDP policy functions are given in Corollaries 5.6 and 5.7.

are computed for three attributes: EMP (number of employees), PAYQTR1 (first quarter payroll, in thousands of dollars), and PAYANN (annual payroll, in thousands of dollars). For this experiment, we take a 5% sample of all NAICS3×COUNTY combinations (groups) in the dataset and use all records in the sampled groups. This results in a sample of 1,095 groups, totaling 323,402 records.

The second dataset is from the U.S. Department of Agriculture’s (USDA’s) Cattle Inventory Survey (CIS), which is a product of the National Agricultural Statistical Service (National Agricultural Statistics Service, 2023). This dataset contains county-level survey records of total cattle inventory. In contrast to the simulated CBP dataset, records on individual establishments (farms) are not available. Therefore, we consider totals at the agricultural district level, which is nested within state, as our privacy unit. Our groupby sums are computed using State as the key, which groups 35,556 values into 43 groups (states). We perform this experiment on the attribute Cattle Value (total cattle inventory).

Table 6 provides summary statistics for each of the variables we consider in these experiments, where Cattle Value corresponds to the attribute from the USDA Cattle Inventory Survey and the remaining three attributes correspond to the simulated CBP subsample. Figures 2 and 3 also display the distributions of the variables we sum over as well as the grouped sums for each dataset. No noise has been added to the data or sums in these figures.

	Min	1st Quartile	Median	Mean	3rd Quartile	Max	N
Cattle Value	100	3,100	10,200	22,480	26,000	1,341,000	5,926
EMP	0	1	2	17	9	19,589	323,402
PAYANN	0	30	106	919	388	1,553,359	323,402
PAYQTR1	0	7	25	239	93	402,324	323,402

TABLE 6. Summary statistics for the columns that are summed over in each of the experiments for both the cattle and sampled simulated CBP data set.

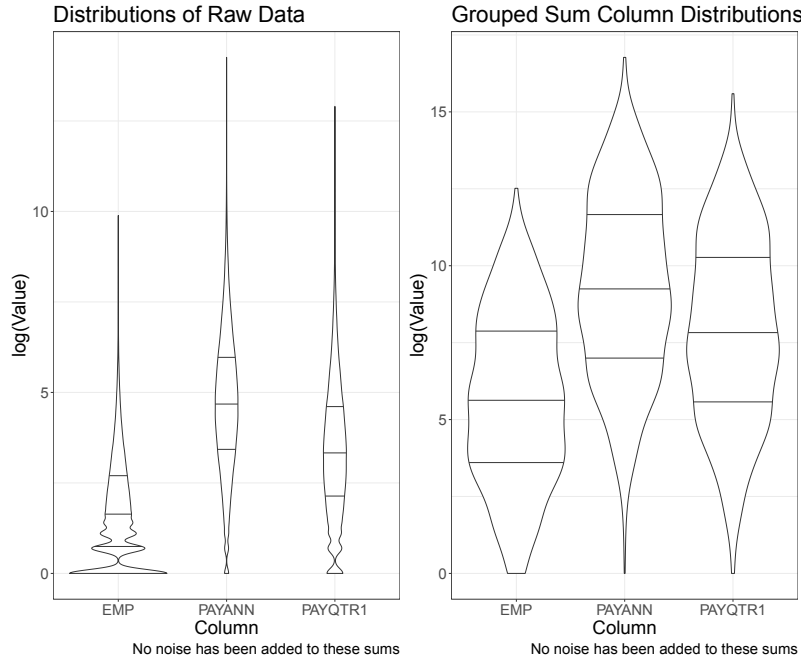


FIGURE 2. This figure shows violin plots for the three variables we consider from the sampled simulated CBP data. The image on the left shows the raw data and the image on the right shows the data grouped and summed by $\text{NAICS3} \times \text{COUNTY}$. There are 1,095 groups in total. The y -axis is on the log scale since all three variables contain large outliers. The black horizontal lines in each of the violins represent the 25%, 50%, and 75% quantiles, respectively.

6.2. Mechanisms. We apply our mechanisms - developed for univariate queries - to the groupby sums by applying them independently to each variable’s sum for each group.

We select transformation and additive mechanisms whose PRzCDP policy functions scale at each of the $O(\Delta(r)^2)$, $O(\sqrt{\Delta(r)})$, and $O(\ln(\Delta(r))^2)$ rates. The transformation mechanisms use the identity, fourth root, and log transformations, and the additive mechanisms use the Gaussian, generalized Gaussian, and exponential polylog noise distributions. See Table 5 for a summary of these mechanisms’ policy and variance functions.

For each attribute summed over in our groupby sums, we set the parameters of the identity transformation mechanism and the equivalent Gaussian additive mechanism to have a privacy loss of 1 for the median record. This entails setting these mechanisms’ standard

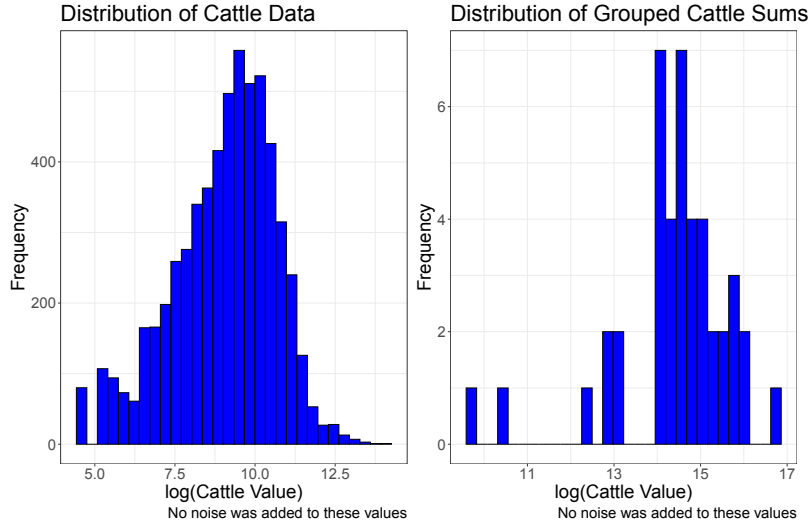


FIGURE 3. This histogram displays the distribution of the number of cattle in the USDA Cattle Inventory Survey dataset. The image on the left shows the raw data and the image on the right shows the cattle data grouped and summed by state. There are 43 groups in total. The x -axis is on the log scale since the dataset contains many large outliers.

Mechanism	Parameter Values (Cattle)	Parameter Values (EMP)	Parameter Values (PAYQTR1)	Parameter Values (PAYANN)
Gaussian	$\sigma = 7212$	$\sigma = 1.41$	$\sigma = 17.7$	$\sigma = 75$
Generalized Gaussian	$\sigma = 658,$ $p = 0.5$	$\sigma = 0.129,$ $p = 0.5$	$\sigma = 1.61,$ $p = 0.5$	$\sigma = 6.84,$ $p = 0.5$
Exponential Polylog	$\sigma = 1,$ $p = 2,$ $d = 0.113,$ $a = e$	$\sigma = 1,$ $p = 2,$ $d = 0.595,$ $a = e$	$\sigma = 1,$ $p = 2,$ $d = 0.337,$ $a = e$	$\sigma = 1,$ $p = 2,$ $d = 0.231,$ $a = e$
Identity Transformation	$\sigma = 7212$	$\sigma = 1.41$	$\sigma = 17.7$	$\sigma = 75$
Fourth Root	$\sigma = 1.67,$ $a = 0$	$\sigma = 0.198,$ $a = 0$	$\sigma = 0.372,$ $a = 0$	$\sigma = 0.534,$ $a = 0$
Log	$\sigma = 0.637,$ $a = 1$	$\sigma = 0.448,$ $a = 1$	$\sigma = 0.616,$ $a = 1$	$\sigma = 0.632,$ $a = 1$

TABLE 7. This table shows the mechanism parameters that were used for each attribute in the experiments.

deviations to $\sqrt{.5}$ times the median of each attribute’s raw data and provides a reasonably strong privacy guarantee to at least half of the records. The parameters for the other mechanisms are set to provide comparable utility; the slowly scaling additive mechanisms have the same standard deviation as the Gaussian additive mechanism, and the slowly

scaling transformation mechanisms have that same standard deviation when run on a query equal to the median value of each attribute.

For all mechanisms except the exponential polylog mechanism, all parameters except σ are set to sensible defaults or to realize the desired slowly scaling behavior. We then solved for σ to obtain the standard deviations described above. For the exponential polylog mechanism, we solved for d to obtain the standard deviations, while the other parameters are set to sensible defaults or to obtain the desired scaling. See Table 7 for the mechanism parameters used for each attribute.

Since these parameter settings are based on the medians of the raw, confidential data, they are not feasible in a traditional DP setting, which does not permit access to the raw data. We use these settings here because they can suggest how well the mechanisms might perform with parameters set similarly, based either on private estimates or reasonably accurate prior knowledge of the medians.

6.3. Metrics. We compute three metrics to evaluate the performance of the different mechanisms. In Figures 4 and 5, we compute every record’s PRzCDP privacy loss from each groupby sum and plot empirical CDFs of these values. For any point (x, y) in a CDF plot, the corresponding value on the y -axis represents the proportion of records with privacy losses that are less than or equal to the x -value. A plot with a CDF that reaches 1 at a given value of x means that no records have a privacy loss greater than that x value. Thus, CDFs which peak quickly indicate that a large proportion of records have small privacy losses.

To assess the mechanisms’ utility, we compute the empirical absolute relative error (ARE), which, for a given mechanism output $M(D)$ and non-noisy query value $q(D)$, is defined as

$$ARE(M(D), q(D)) = \frac{|M(D) - q(D)|}{|q(D)|}.$$

Figures 6 and 7 show boxplots of the ARE for different quintiles of the true data to show how the distribution of the error changes as the size of the record grows.

Finally, in Figure 8, we plot the mechanisms’ standard deviations as functions of the true query value, $q(D)$.

6.4. Discussion. We present our results below, discussing the mechanisms’ privacy guarantees in Section 6.4.1 and then their utility in Sections 6.4.2 and 6.4.3.

6.4.1. Empirical Privacy Loss CDFs. Figure 4 displays the empirical CDFs of the privacy losses of the additive and transformation mechanisms when applied to the sampled simulated CBP data. To build the CDFs, we computed the privacy loss for each record and then determined, over a range of privacy loss thresholds, the proportion of privacy losses which were less than or equal to each threshold value. This provides a finer view on the shape of the CDF than would be attainable using just the data.

Recall that a CDF indicates stronger privacy guarantees when it has a larger y value for any given x value. The left-hand panel of Figure 4 shows that, in the simulated CBP data, both the generalized Gaussian and the exponential polylog mechanisms dominate the Gaussian mechanism for all values of privacy loss except for small values near 0. The CDFs for the exponential polylog mechanism peak (reach 1) earliest, indicating that, by using this

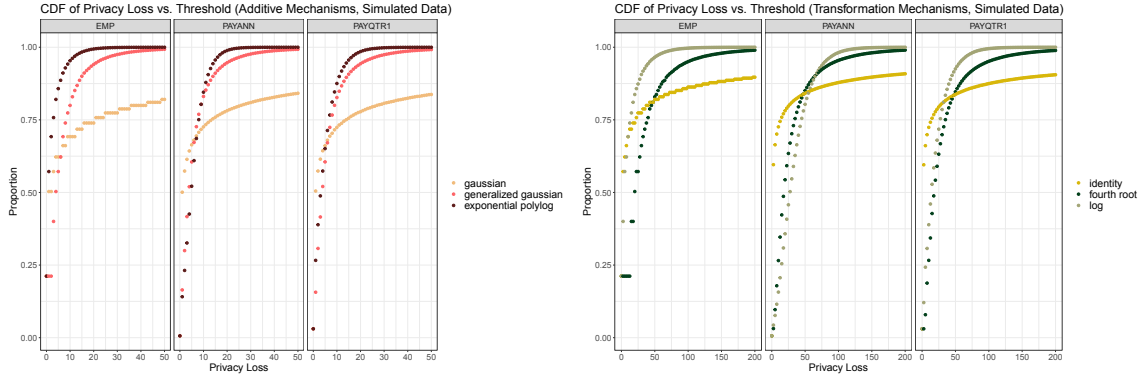


FIGURE 4. This figure shows empirical CDFs of the privacy loss for each mechanism for the sampled simulated data. For a given value on the x -axis, the y -axis shows the proportion of records which have at most that level of privacy loss. These CDFs were built by getting the privacy loss for each record and then computing, over a range of privacy loss thresholds, the proportion of privacy losses which were less than or equal to each threshold value. For each attribute, the mechanisms’ parameters were calculated such that their standard deviations are equal to $\sqrt{0.5} * \text{median}(\text{attribute value})$ on a query equal to $\text{median}(\text{attribute value})$.

mechanism for these datasets, we can ensure that the maximum privacy loss is lower than the maximum privacy loss for other mechanisms.

In the right-hand panel of Figure 4, we see that the transformation mechanisms exhibit similar patterns, but the privacy losses of the slowly scaling transformation mechanisms are generally larger than their counterpart additive mechanisms – note the larger scale of the x -axis in this panel. As a result, the identity mechanism dominates for privacy losses which are lower than about 50 for the PAYANN variable and about 30 for PAYQTR1. After these intersection points, the CDFs of the slowly scaling mechanisms again dominate.

Figure 5 presents the empirical CDFs of the same mechanisms using the cattle data. For the additive mechanisms, the Gaussian mechanism has the highest CDF until about $x = 6$ before being overtaken by the generalized Gaussian and later the exponential polylog. This follows a similar trend as we observed on the simulated CBP data.

The transformation mechanisms show similar results compared to the simulated CBP data. Here, the identity mechanism CDF is highest until about $x = 50$, when it is overtaken by the fourth root mechanism. The largest privacy losses for the fourth root mechanism and the log mechanism appear to be smaller than those for the identity mechanism, since the identity mechanism CDF does not reach 1 on this graph.

The CDFs for the cattle data differ from the CDFs for the simulated CBP data in that the less slowly scaling mechanisms (the generalized Gaussian and fourth root mechanisms) dominate the more slowly scaling mechanisms (the exponential polylog and log mechanisms) over a relatively wide range of privacy losses. Even when the more slowly scaling mechanisms eventually catch up, they dominate the less slowly scaling mechanisms’ CDFs by a very thin margin, the latter already being very close to 1 at this point. On the whole, the privacy guarantees of the less slowly scaling mechanisms look substantially better for the vast majority of records.

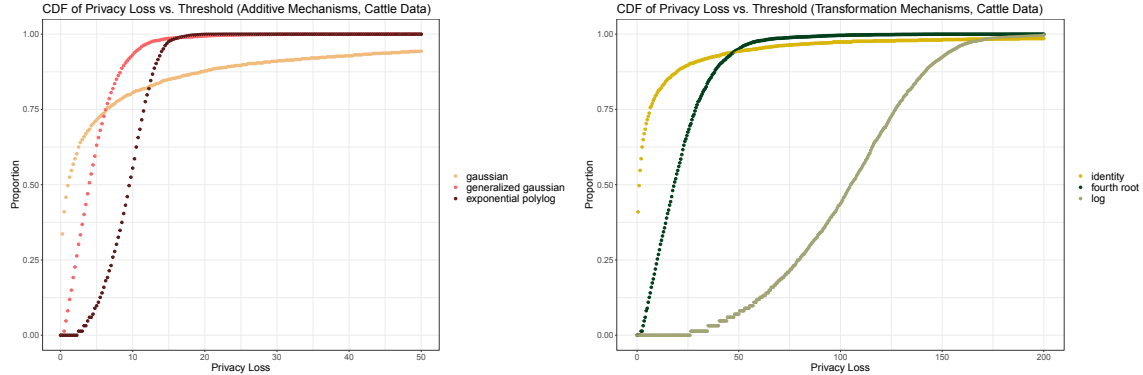


FIGURE 5. This figure shows an empirical CDF of the privacy loss for each mechanism for the cattle data. For every value on the x -axis, the y -axis shows the proportion of records which have at most that privacy loss. These CDFs were built by getting the privacy loss for each record and then computing, over a range of privacy loss thresholds, the proportion of privacy losses which were less than or equal to each threshold value. The mechanisms’ parameters were calculated such that their standard deviations are equal to $\sqrt{0.5} * \text{median}(\text{cattle value})$ on a query equal to $\text{median}(\text{cattle value})$.

6.4.2. *Absolute Relative Error.* Figure 6 displays boxplots of the empirical ARE on the \log_{10} scale for the additive and transformation mechanisms. Separate boxplots are shown for each quintile of the true sum values, as indicated on the x -axis. In these plots, the blue line represents 100% ARE and the orange line represents 10% ARE. One may consider these as fitness-for-use targets, since 100% error may be intolerable but 10% error may be tolerable.

For the additive mechanisms, whose variances are constant with respect to the true query value, we can see that the ARE decreases markedly as the quintile increases for all mechanisms and all attributes. In these plots, the additive mechanisms all have comparable error and exhibit similar patterns as the quintile increases. The generalized Gaussian and exponential polylog mechanisms typically have lower median ARE than the Gaussian mechanism, but with a wider spread.

Because the identity transformation mechanism is the same as the Gaussian additive mechanism, its ARE boxplots follow the same pattern, declining rapidly with the quintile on the x -axis. The slowly scaling transformation mechanisms, however, have variances that increase with the true query value, offsetting this effect. The log mechanism’s standard deviation asymptotically grows linearly in the query value (see Figure 6.4.3 for an illustration), leading the expected relative error to bottom out at a nonzero value. In Figure 6, this leads the log mechanism’s empirical ARE distributions to have 25th percentiles that never fall below 10%. In this application, this would likely rule out the log mechanism as a viable choice. The fourth root mechanism fares only somewhat better, having 75th percentiles below 10% only for the largest quintile.

Figure 7 presents the AREs for the cattle data. For the additive mechanisms, we still observe that ARE decreases as the quintile increases as we did with the simulated CBP dataset. Additionally, we note that the exponential polylog generally has the lowest error, followed closely by the generalized Gaussian and then the Gaussian mechanisms.

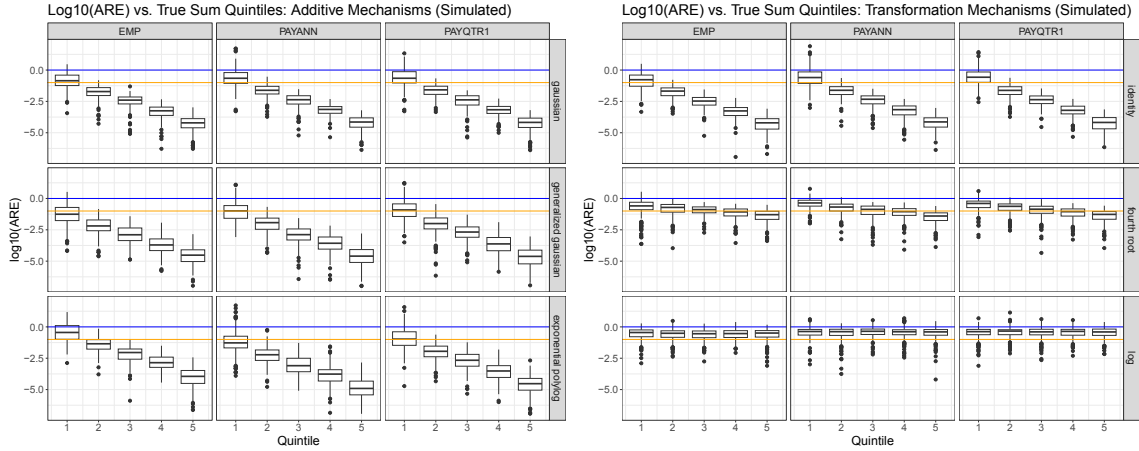


FIGURE 6. This figure shows the distribution of absolute relative error (ARE) on the \log_{10} scale for each mechanism using each of our three variables in the CBP dataset. The quintile represents where the true sum (without noise added) was in the first 20%, the second 20%, etc. The blue line represents where $\text{ARE} = 100\%$ and the orange line represents where $\text{ARE} = 10\%$. The horizontal lines in the boxplot represent the 25th, 50th (the bold middle line), and 75th percentiles of error within that quintile. For each attribute, the mechanisms’ parameters were calculated such that their standard deviations are equal to $\sqrt{0.5} * \text{median}(\text{attribute value})$ on a query equal to $\text{median}(\text{attribute value})$.

We also see that the ARE is generally lower for this dataset compared to the simulated dataset; each quintile’s range of ARE values lies almost entirely below the orange 10% line. This is a result of the generally larger query values in the cattle dataset, which follows from its groupby sum separating the data into fewer groups. This demonstrates that the fitness for use of the additive mechanisms is heavily dependent on the query values, much like other traditional differentially private mechanisms.

The slowly scaling transformation mechanisms again suffer from their variances growing with the query value. In every quintile, the fourth root and log mechanisms have median AREs above or near the 10% line. The failure of the fourth root mechanism to improve much with the quintile appears to be due to the distribution of the cattle value queries being more closely clustered than those in the simulated CBP dataset. The result is a reduced difference in relative error between queries.

6.4.3. Standard Deviation. In Figure 8, we plot the standard deviations for our selected mechanisms as $q(D)$ grows, holding the privacy loss constant. We compute the variances using the formulas in Table 5 but plot the standard deviations to ease interpretation. In the left-hand panel in Figure 8, we use the parameters used for the EMP attribute, and in the right-hand panel, we use the cattle value parameters (see Table 7). In these plots, the horizontal lines labeled “Constant Variance Mechanisms” plot the variances of the additive mechanisms and the identity transformation mechanism, since their variances are equal and do not depend on $q(D)$.

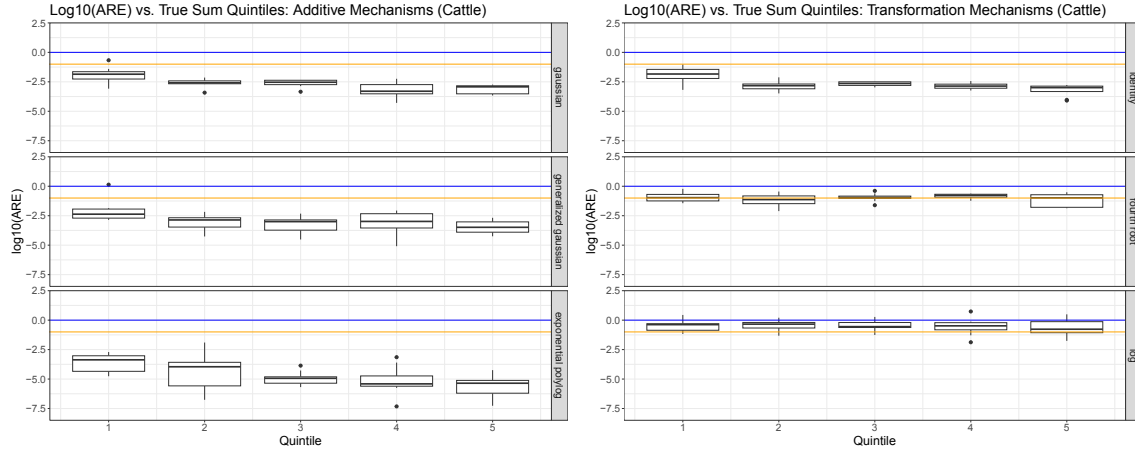


FIGURE 7. This figure shows the distribution of absolute relative error (ARE) on the \log_{10} scale for each mechanism using the cattle value variable in the cattle dataset. The quintile represents where the true sum (without noise added) was in the first 20%, the second 20%, etc. The blue line represents where $\text{ARE} = 100\%$, and the orange line represents where $\text{ARE} = 10\%$. The horizontal lines in the boxplot represent the 25th, 50th (the bold middle line), and 75th percentiles of error within that quintile. The mechanisms' parameters were calculated such that their standard deviations are equal to $\sqrt{0.5} * \text{median}(\text{cattle value})$ on a query equal to $\text{median}(\text{cattle value})$.

Because the mechanisms are calibrated to have the same variance for a query equal to the median record, and because the slowly scaling transformation mechanisms have variances that grow in the query value, the slowly scaling transformation mechanisms at these settings will have lower variances than the additive mechanisms for queries smaller than the median record. In Figure 8, we see that this leaves a fairly large space of query values (all values below 10,200) for the cattle data where the transformation mechanisms have lower variances. For the EMP data, by contrast, this space only ranges from 0 to 2 and is invisibly small in the graph. It is clear from the graphs that the transformation mechanisms can have far worse utility than the additive mechanisms when the query value is large.

While the results of our experiments make the additive mechanisms look decidedly better for the datasets we use, the slowly scaling transformation mechanisms could well have lower variances and better ARE distributions than the additive mechanisms if one were to apply these mechanisms to very sparse data – data with lots of queries summing over a single record or even no records. With this improvement in utility, the apparently worse privacy guarantees of the slowly scaling transformation mechanisms in Figures 4 and 5 could also be remedied by increasing the mechanisms' variances. In Appendix A, we discuss some conditions under which the transformation mechanism has a lower variance than the additive mechanism for a comparable privacy guarantee.

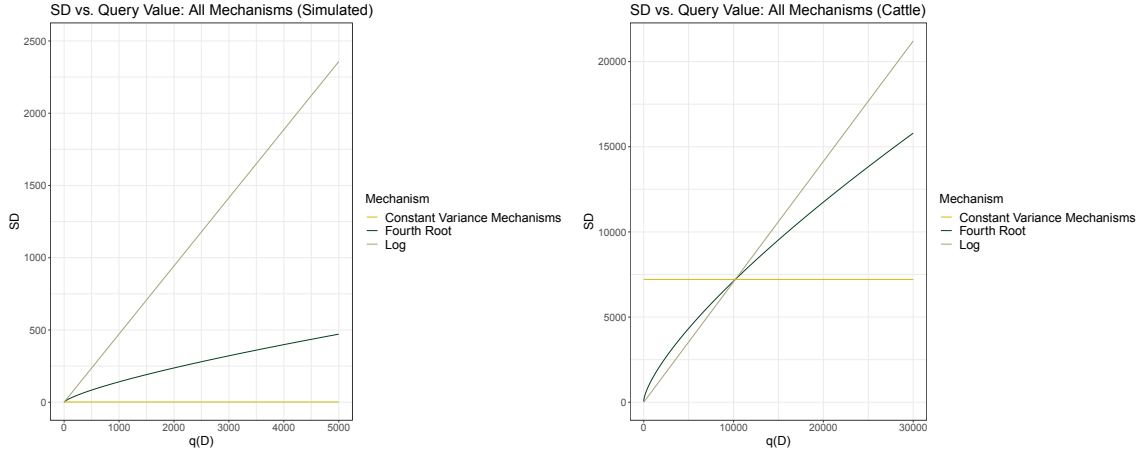


FIGURE 8. These plots display how the standard deviation for each mechanism changes as a function of the query value $q(D)$. See Table 5 for the mechanisms’ variance functions. See Table 7 for the mechanisms’ parameters. On the left, we use the parameters for the EMP attribute, and on the right, we use the parameters for the cattle value attribute. Every additive mechanism has constant standard deviation regardless of query value, as does the identity transformation mechanism. These mechanisms’ standard deviations are all reflected in the yellow, horizontal “constant variance mechanisms” line.

7. CONCLUSION

In this paper, we have developed two classes of slowly scaling PRzCDP mechanisms – mechanisms whose policy functions scale sub-linearly in a record-specific measure of influence. We first demonstrated that unit splitting – the previous method for satisfying PRzCDP – results in unacceptably high privacy losses when individual records can be very large. This was because the policy function incurred by unit splitting scales in the square of the record’s magnitude.

The first class of slowly scaling mechanisms, the transformation mechanisms, add Gaussian noise to a concave transformation of the query. This results in a policy function that scales in the square of the chosen transformation function. For example, the square root transformation results in a linear policy function and the log transformation results in a policy function that scales in the square of the logarithm.

The second class of mechanisms, the additive mechanisms, draw noise from specially chosen distributions with fat tails. These mechanisms can achieve a wide range of policy functions, including linear and polylogarithmic policy functions like those achieved by the transformation mechanisms.

Last, we conducted experiments using our mechanisms with two different datasets: a sample of a simulated version of the County Business Patterns dataset and a dataset of cattle inventory from the U.S. Department of Agriculture. Both of these experiments showed that the empirical CDFs of our slowly scaling mechanisms generally peak earlier than a standard (i.e., Gaussian) mechanism, indicating that the maximum privacy loss (for the largest records) incurred by these mechanisms is lower.

Our results generally indicate that slowly scaling mechanisms can sharply reduce the privacy losses incurred by very influential records. With this check on extreme privacy losses, it becomes tenable to publish queries without clipping influential records. This, in turn, can reduce or eliminate bias. Future work could include the derivation of additional slowly scaling mechanisms; the derivation of tighter PRzCDP policy functions for the additive mechanisms, since our current results are based on the mechanisms' PRDP policy functions and may be loose; and the development of frequentist hypothesis testing semantics (Kifer et al., 2022) for PRzCDP and PRDP.

ACKNOWLEDGMENT

We would like to acknowledge the following individuals for their helpful discussion, feedback, and contributions toward this paper: Paul Bartholomew, Christina Bassis, Noah Fine, Monica Mendoza, Theresa Nguyen, Sayi Sathyavanan, and Rich Stevenson from The MITRE Corporation; Jordan Awan from Purdue University; Skye Berghel, Bayard Carlson, Casey Meehan, and Ruchit Shrestha from Tumult Labs; Margaret Beckom, William C. Davie Jr., and Rolando Rodriguez from the U.S. Census Bureau; and Philip Leclerc. We thank Ian Schmutte for suggesting a mechanism that directly inspired the transformation mechanisms developed here.

REFERENCES

- 83rd United States Congress. Title 13, 1954.
- 99th United States Congress. Title 26 / Internal Revenue Code of 1986, 1986.
- J. M. Abowd, R. Ashmead, R. Cumings-Menon, S. Garfinkel, M. Heineck, C. Heiss, R. Johns, D. Kifer, P. Leclerc, A. Machanavajjhala, B. Moran, W. Sexton, M. Spence, and P. Zhuravlev. The 2020 Census Disclosure Avoidance System TopDown Algorithm. *Harvard Data Science Review*, (Special Issue 2), jun 24 2022. <https://hdsr.mitpress.mit.edu/pub/7evz361i>.
- T. Adeleye, S. Berghel, D. Desfontaines, M. Hay, I. Johnson, C. Lemoisson, A. Machanavajjhala, T. Magerlein, G. Modena, D. Pujol, D. Simmons-Marengo, and H. Trieman. Publishing Wikipedia usage data with strong privacy guarantees, 2023. <https://arxiv.org/abs/2308.16298>.
- S. Berghel, P. Bohannon, D. Desfontaines, C. Estes, S. Haney, L. Hartman, M. Hay, A. Machanavajjhala, T. Magerlein, G. Miklau, A. Pai, W. Sexton, and R. Shrestha. Tumult Analytics: a robust, easy-to-use, scalable, and expressive framework for differential privacy, 2022. <https://arxiv.org/abs/2212.04133>.
- M. Bun and T. Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pages 635–658. Springer, 2016. https://doi.org/10.1007/978-3-662-53641-4_24.
- F. Calmon, E. Du, C. Dwork, B. Finley, and G. Franguridi. Debiasing functions of private statistics in postprocessing. volume 329, pages 17:1–17:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. <https://doi.org/10.4230/LIPICS.FORC.2025.17>.
- K. Chatzikokolakis, M. E. Andrés, N. E. Bordenabe, and C. Palamidessi. Broadening the scope of differential privacy using metrics. In E. D. Cristofaro and M. K. Wright, editors, *Privacy Enhancing Technologies - 13th International Symposium, PETS 2013, Bloomington, IN, USA, July 10-12, 2013. Proceedings*, volume 7981 of *Lecture Notes in Computer*

- Science*, pages 82–102. Springer, 2013. https://doi.org/10.1007/978-3-642-39077-7_5.
- J. Dong, A. Roth, and W. J. Su. Gaussian differential privacy. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(1):3–37, 02 2022. ISSN 1369-7412. <https://doi.org/10.1111/rssb.12454>.
- C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pages 265–284. Springer, 2006. https://doi.org/10.1007/11681878_14.
- B. Finley, A. Caruso, J. Doty, A. Machanavajjhala, M. Meyer, D. Pujol, W. Sexton, and Z. Terner. Supplemental materials for "slowly scaling per-record differential privacy", Mar. 2025. <https://doi.org/10.5281/zenodo.15085814>.
- A. Ganesh and J. Zhao. Privately answering counting queries with generalized Gaussian mechanisms. In K. Ligett and S. Gupta, editors, *2nd Symposium on Foundations of Responsible Computing (FORC 2021)*, volume 192 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 1:1–1:18, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-187-0. <https://doi.org/10.4230/LIPIcs.FORC.2021.1>.
- M. Griffin. *gnorm: Generalized Normal/Exponential Power Distribution*, 2018a. R package version 1.0.0.
- M. Griffin. *Working with the Exponential Power Distribution Using gnorm*, 2018b. <https://cran.r-project.org/web/packages/gnorm/vignettes/gnormUse.html>.
- S. Haney, A. Machanavajjhala, J. M. Abowd, M. Graham, M. Kutzbach, and L. Vilhuber. Utility cost of formal privacy for releasing national employer-employee statistics. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pages 1339–1354, 2017. <https://doi.org/10.1145/3035918.3035940>.
- S. Haney, W. Sexton, A. Machanavajjhala, M. Hay, and G. Miklau. Differentially private algorithms for 2020 Census Detailed DHC Race & Ethnicity. Technical report, arXiv, July 2021. <http://arxiv.org/abs/2107.10659>.
- S. Haney, D. Desfontaines, L. Hartman, R. Shrestha, and M. Hay. Precision-based attacks and interval refining: How to break, then fix, differential privacy on finite computers, 2022. <https://arxiv.org/abs/2207.13793>.
- D. Kifer and A. Machanavajjhala. No free lunch in data privacy. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data*, SIGMOD '11, page 193–204, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450306614. <https://doi.org/10.1145/1989323.1989345>.
- D. Kifer, J. M. Abowd, R. Ashmead, R. Cumings-Menon, P. Leclerc, A. Machanavajjhala, W. Sexton, and P. Zhuravlev. Bayesian and frequentist semantics for common variations of differential privacy: Applications to the 2020 Census. *arXiv preprint arXiv:2209.03310*, 2022. <https://arxiv.org/abs/2209.03310>.
- A. A. Kilbas. *H-transforms: Theory and Applications*. CRC Press, 2004. <https://doi.org/10.1201/9780203487372>.
- F. Koufogiannis, S. Han, and G. J. Pappas. Optimality of the Laplace mechanism in differential privacy, 2015. <https://arxiv.org/abs/1504.00065>.
- F. Liu. Generalized Gaussian mechanism for differential privacy. *IEEE Trans. on Knowl. and Data Eng.*, 31(4):747–756, apr 2019. ISSN 1041-4347. <https://doi.org/10.1109/TKDE.2018.2845388>.

- S. Nadarajah. A generalized Normal distribution. *Journal of Applied Statistics*, 32(7): 685–694, 2005. <https://doi.org/10.1080/02664760500079464>.
- National Agricultural Statistics Service. Census of Agriculture, 2023. <https://www.nass.usda.gov/AgCensus/>.
- A. Rényi. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, volume 4, pages 547–562. University of California Press, 1961. https://digitalassets.lib.berkeley.edu/math/ucb/text/math_s4_v1_article-27.pdf.
- J. Seeman, W. Sexton, D. Pujol, and A. Machanavajhala. Privately answering queries on skewed data via per-record differential privacy. *Proc. VLDB Endow.*, 17(11):3138–3150, 2024. <https://doi.org/10.14778/3681954.3681989>.
- H. Srivastava, R. Saxena, and R. Parmar. Some families of the incomplete H -functions and the incomplete \bar{H} -functions and associated integral transforms and operators of fractional calculus with applications. *Russian Journal of Mathematical Physics*, 25:116–138, 2018. <https://doi.org/10.1134/S1061920818010119>.
- T. Steinke. Tight RDP and zCDP bounds from pure DP. DifferentialPrivacy.org, 05 2024. <https://differentialprivacy.org/pdp-to-zcdp/>.
- Team OpenDP. The opendp white paper, 2020. https://salil.seas.harvard.edu/sites/g/files/omnuum4266/files/salil/files/opendp_white_paper_11may2020.pdf.
- U.S. Census Bureau. Why the Census Bureau chose differential privacy, 2023. <https://www.census.gov/library/publications/2023/decennial/c2020br-03.html>.
- Y. Washio, H. Morimoto, and N. Ikeda. Unbiased estimation based on sufficient statistics. *Bulletin of Mathematical Statistics*, 6:69–93, 1956. <https://doi.org/10.5109/12969>.
- K. Webb, P. Protivash, J. Durrell, D. Toth, A. Slavković, and D. Kifer. Statistics-friendly confidentiality protection for establishment data, with applications to the QCEW, 2025.

APPENDIX A. ANALYTICAL COMPARISON OF ADDITIVE AND TRANSFORMATION MECHANISMS

When choosing a mechanism for a PRzCDP release, there are multiple tradeoffs. The choice of policy function not only decides how the privacy loss scales across records, but also the amount of error required to achieve specific privacy losses. Meanwhile, the class of mechanism (either transformation or additive) determines whether error will scale as a function of query value. In this section, we discuss these tradeoffs and how specific choices of policy function and mechanism class influence the utility of a data release. A summary of the variances and policy functions of selected mechanisms can be found in Table 5.

A.1. Policy function. The choice of policy function, for a given mechanism class, informs how privacy loss will be distributed across all the records. While more slowly scaling policy functions ensure that larger records incur less privacy loss, they may also result in smaller records incurring larger privacy losses. We visualize this in Figure 9. For this reason, a data curator must use their knowledge of the data distribution (or private exploratory data analysis) in order to select a policy function that results in an appropriate distribution of privacy loss for the given context.

Example A.1. Consider a sum being computed by different transformation mechanisms, whose parameters are set such that they have a standard deviation of 1,000 for a sum of size

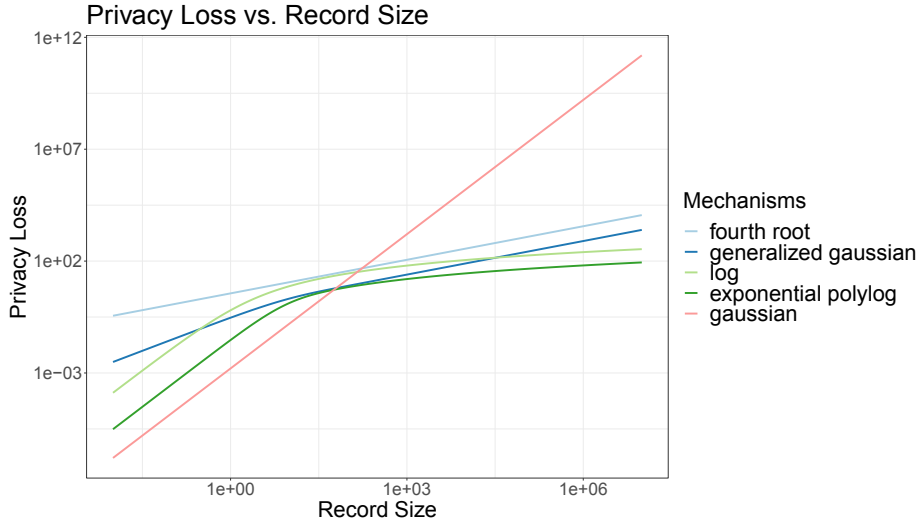


FIGURE 9. This figure shows the PRzCDP policy functions for several mechanisms over a large domain. The parameters are set for PAYQTR1 as indicated in Table 7. Both the x and y axes are on the \log_{10} scale.

10,000. That is, the square root transformation has parameters $a = 0$ and $\sigma \approx 5$, resulting in the PRzCDP policy function $\frac{\Delta(r)}{50}$. The fourth root transformation has parameters $a = 0$ and $\sigma \approx 0.2$, resulting in the policy function $\frac{\sqrt{\Delta(r)}}{0.08}$. Finally, the log transformation has parameters $a = 1$ and $\sigma \approx 0.1$, resulting in the policy function $\frac{\log(\Delta(r)+1)}{0.02}$. The square root transformation function incurs the lowest privacy losses for records whose value being summed over lies in $[0, 390625]$. Then the fourth root transformation incurs the lowest privacy loss for records those value being summed over lies in $[390625, 463584]$. Any records with values larger than 463584 incur the lowest privacy loss when the log transformation mechanism is used.

This demonstrates that small record values are best protected by faster scaling policy functions, such as that obtained with the square root transformation. The fourth root transformation is preferable for medium-sized records, and the log transformation results in the lowest privacy loss for large records. As such, faster scaling policy functions may be preferable when the data's tails are thinner (i.e., when outliers are few and relatively small). We conduct a similar analysis on the additive mechanisms below.

Example A.2. Consider a sum being computed by the class of additive mechanisms, whose parameters are set such that they have a standard deviation of 1,000. That is, the generalized Gaussian with $p = 1$ has parameter $\sigma \approx 707$, resulting in the PRDP policy function $\frac{\Delta(r)}{707}$. The generalized Gaussian with $p = 1/2$ has parameter $\sigma \approx 91$, resulting in the PRDP policy function $\sqrt{\frac{\Delta(r)}{91}}$. Finally, the exponential polylog has parameters $p = 2, a = e, \sigma = 1$, and $d \approx 0.145$, resulting in the policy function $0.145(\ln(\Delta(r) + e)^2 - 1)$.

The generalized Gaussian with $p = 1$ incurs the lowest privacy losses for records whose values lie in approximately $[0, 5492]$. Then the generalized Gaussian with $p = 1/2$ incurs the

lowest privacy loss for records with values in approximately $[5492, 17387]$. Any records with values larger than 17387 incur the lowest privacy loss when the exponential polylog is used.

Much like the transformation mechanism, the faster scaling policy functions incur less privacy loss for the smaller records, despite their larger privacy loss for larger records. This means that regardless of mechanism class, the data curator must decide which policy function results in an appropriate distribution of privacy loss across records. Depending on the context, that may mean choosing a faster scaling policy function that results in low privacy loss for the majority of the records, while few records incur high privacy loss. In other contexts, that may mean choosing a slowly scaling policy function which may incur larger privacy loss for smaller records but ensure all records incur reasonable privacy loss.

A.2. Mechanism Class. The transformation and additive mechanisms come with different advantages and disadvantages. A principal strength of the transformation mechanisms are their ease of implementation and development. Software implementations of privacy mechanisms often require that special floating-point-safe sampling routines be used to draw any random noise; naively implemented pseudorandom number generators that suffice for other purposes often draw from an approximate distribution in a way that can catastrophically degrade any privacy guarantees (Haney et al., 2022). The transformation mechanism can be implemented as long as Gaussian noise can be safely drawn, which existing software platforms already support (see, for example, Berghel et al. (2022) and Team OpenDP (2020)). This extends to transformation functions beyond those described in this work. A new transformation mechanism can easily be developed for any suitable concave, strictly increasing function, allowing for precise targeting of specific policy functions.

On the other hand, the additive mechanisms require the development of specific distributions for each policy function. For each of these distributions, a new floating-point safe sampling method must be developed in order to be used safely. This may limit the range of policy functions one can practically attain with the additive mechanism. For the additive mechanisms we implement in this paper, their quantile functions depend on standard functions, which makes implementation of samplers for these distributions straightforward with the use of an arbitrary precision library. While this is true for the exponential polylogarithmic distribution when $p \in \{1, 2\}$, for any other p the distribution function becomes intractable and quantile functions are not easily derived.

The additive noise mechanisms have the advantage in terms of bias. Being symmetrically distributed about $q(D)$, these are unbiased for the mean, median, and mode. Transformation mechanisms exist that are either mean- or median-unbiased, but typically are not both (see Appendix B for median-unbiased transformation mechanisms).

A notable drawback of the transformation mechanisms is that the mechanisms' variances typically grow with the query value, $q(D)$. By contrast, the variances of the additive mechanisms are constant. This reduces the transparency of the transformation mechanisms, as data users cannot be told exactly what variance the added noise has. Users can still estimate these variances, however. For example, a user could trim the mechanism's output to lie within the query's domain and plug the result into the variance formulas in Table 2 in place of $q(D)$.

This dependence of the transformation mechanisms' variance on $q(D)$ also immediately implies that the additive mechanism can always offer a lower variance for a given privacy guarantee if $q(D)$ is large enough. In the case of the transformation mechanism with the log




FIGURE 10. (a) Standard deviation and (b) relative standard deviation of noise added from mean-unbiased transformation mechanisms for sum queries, with transformation functions specified in the legend. Relative standard deviation is defined as the standard deviation divided by $q(D)$. These mechanisms are tuned to achieve PRzCDP privacy loss of 1 when $\Delta(r) = 100$.

transform, the mechanism’s standard deviation grows at the $O(q(D))$ rate, and the ratio of $q(D)$ to this standard deviation limits to a positive constant, as shown in Figure 10.

In this sense, then, the noise added by this mechanism is non-negligible, even for very large $q(D)$. For the additive mechanism, by contrast, the amount of noise can become arbitrarily small, relative to an arbitrarily large $q(D)$.

The transformation mechanism, however, may offer lower variance for small query values. To illustrate this, we take a mechanism of each type that asymptotically achieves the same policy function and compare their variances in Examples A.3 and A.4, as well as Figure 11.

Example A.3. Consider the transformation mechanism with transformation $f(x) = \sqrt{x}$, $a = 0$, and $\sigma = \sqrt{.5\sigma_c}$. Likewise, consider the additive mechanism using the generalized Gaussian distribution with $p = 1$ and $\sigma = \sigma_c$ (i.e., the Laplace mechanism). Per Theorem 4.4, the transformation mechanism has the policy function $P(r) = \frac{\Delta(r)}{\sigma_c}$. By Corollary 5.6, the additive mechanism has the policy function $P'(r) = \tanh(\frac{\Delta(r)}{2\sigma_c}) \frac{\Delta(r)}{\sigma_c}$. The difference between these policy functions approaches zero rapidly as $\Delta(r)$ grows, so we say that these mechanisms have the same asymptotic policy functions.

The transformation mechanism has variance $2\sqrt{.5\sigma_c^4} + 4\sqrt{.5\sigma_c^2}q(D) = .5\sigma_c^2 + 2\sigma_c q(D)$ and the additive mechanism has variance $\sigma_c^2 \frac{\Gamma(3)}{\Gamma(1)} = 2\sigma_c^2$. The transformation mechanism therefore has larger variance than the additive mechanism when $2\sigma_c^2 \leq .5\sigma_c^2 + 2q(D)\sigma_c$, or $.75\sigma_c \leq q(D)$. This tradeoff is visualized in Figure 11. Note that the region where the transformation mechanism dominates the additive mechanism is larger when the privacy guarantee is stronger (i.e., when σ_c is larger).

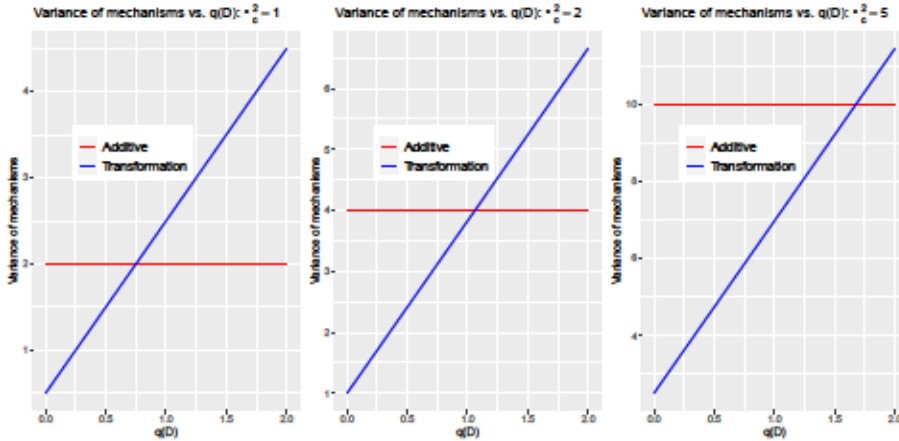


FIGURE 11. Variances of the generalized Gaussian mechanism with $p = 1$, $\sigma = \sigma_c$ and the square root transformation mechanism with $a = 0$, $\sigma = \sqrt{.5\sigma_c}$. These mechanisms have the same asymptotic policy functions (see Example A.3). For all values of σ_c , the transformation mechanism has lower variance for small query values, while the additive mechanism has lower variance for large query values.

Likewise, we compare the log transformation mechanism with the additive mechanism using the exponential polylogarithmic distribution when $p = 2$.

Example A.4. To compare the log transformation mechanism and the additive mechanism with the exponential polylogarithmic distribution we set parameters such that the mechanisms' policy functions nowhere exceed $\ln(\Delta(r) + a)^2/2\sigma^2$. In order to achieve this, we set $d = 1/2\sigma^2$ for the exponential polylog distribution and compare to the log transformation mechanism with noise parameter σ . The offset parameter $a \geq e$ is the same for both mechanisms. Given this setting, the log-transformation mechanism has variance $(e^{2d} - 1)(q(D) + a)^2$.

Let σ_g^2 denote the variance of the exponential polylogarithmic distribution from Table 3. Then the transformation mechanism will have larger variance when $\sqrt{\frac{\sigma_g^2}{e^{2d}-1}} - a \leq q(D)$.

With similar reasoning as above, the additive mechanism will dominate when the true query value is large. This means that the choice of mechanism class should be decided based on prior knowledge of the data release. If the data curator knows that query values are likely to be small, the transformation mechanisms may incur less error. However, if query values are likely to be large, then additive mechanisms will result in lower overall error.

APPENDIX B. TRANSFORMATION MECHANISMS WITH MEDIAN UNBIASED ESTIMATORS

In Section 4.3, we developed mean-unbiased estimators for use in transformation mechanisms. Another class of estimators sets $g(x) = f^{-1}(x) - a$ for x where f^{-1} exists - that is, for $x \in [f(a), \infty)$, provided that $\lim_{x \rightarrow \infty} f(x) = \infty$. The intuitive appeal is clear; if no noise were added to $f(q(D))$ in Line 3 of Algorithm 1, then this would return $q(D)$ exactly. However, while $f(q(D) + a)$ lies in $[0, \infty)$, the noisy transformed query value \tilde{v} may take

any real value, so care must be taken to define $g(x)$ for $x \in (-\infty, f(a))$, as well. Different estimators in this class take different values in this region.

One option that may be intuitively appealing is to set g using some "canonical" inverse of f that is defined over all of \mathbb{R} . For example, if $f(x) = \sqrt[k]{x}$ for a positive integer k , one could use $g(x) = x^k - a$. These estimators do not necessarily have any unbiasedness properties, but other estimators in this class are median-unbiased and can be readily obtained for any f ; when g is defined to always be non-positive where f^{-1} does not exist, the estimator is median-unbiased. Among these median-unbiased estimators, the risk-minimizing estimator simply sets $g(x) = 0$ in this region.

Theorem B.1. *Let Assumption 4.3 hold. Suppose further that $\lim_{t \rightarrow \infty} f(t) = \infty$, $g(x) = f^{-1}(x) - a$ for $x \in [f(a), \infty)$, and $g(x) \leq 0$ for $x < f(a)$. Denote Algorithm 1($q(D)$, a , σ , f , g) by $M(D)$ and the median of a random variable, X , by $\mathbb{M}(X)$. Let $L(M(D), q(D))$ be any loss function that is weakly decreasing in $M(D)$ for $M(D) < q(D)$.*

- $\mathbb{M}[M(D)] = q(D)$.
- If $g(x) = 0$ for $x < f(a)$, then the risk, $\mathbb{E}[L(M(D), q(D))]$, is minimized among estimators subject to the above assumptions.

Proof. Recall that the noisy transformed query value created in Line 3 of Algorithm 1 is denoted by \tilde{v} and the algorithm's output is $M(D) = g(\tilde{v})$. Note that f^{-1} is monotonically increasing. For the first item in Theorem B.1, we have the following:

$$\begin{aligned}
P[q(D) < g(\tilde{v})] &= P[f(q(D) + a) \leq \tilde{v}]P[q(D) < f^{-1}(\tilde{v}) - a | f(q(D) + a) \leq \tilde{v}] \\
&\quad + P[f(a) \leq \tilde{v} < f(q(D) + a)]P[q(D) < f^{-1}(\tilde{v}) - a | f(a) \leq \tilde{v} < f(q(D) + a)] \\
&\quad + P[\tilde{v} < f(q(D) + a)]P[q(D) < g(\tilde{v}) | \tilde{v} < f(a)] \\
&= .5P[f(q(D) + a) < \tilde{v} | f(q(D) + a) \leq \tilde{v}] \\
&\quad + P[f(a) \leq \tilde{v} < f(q(D) + a)]P[f(q(D) + a) < \tilde{v} | f(a) \leq \tilde{v} < f(q(D) + a)] \\
&\quad + P[\tilde{v} < f(q(D) + a)]P[q(D) < g(\tilde{v}) | \tilde{v} < f(a)] \\
&= .5
\end{aligned}$$

To see the last equality, recall our assumption that, when $\tilde{v} < f(a)$, we have $g(\tilde{v}) \leq 0 \leq q(D)$, so $P[q(D) < g(\tilde{v}) | \tilde{v} < f(a)] = 0$. Likewise, it is clear that $P[f(q(D) + a) < \tilde{v} | f(a) \leq \tilde{v} < f(q(D) + a)] = 0$ and $P[f(q(D) + a) < \tilde{v} | f(q(D) + a) \leq \tilde{v}] = 1$.

It follows that $q(D)$ is the median of $g(\tilde{v}) = M(D)$.

For the second item in Theorem B.1, recall that, by assumption, $g(\tilde{v}) \leq 0 \leq q(D)$ for $\tilde{v} < f(a)$, so $L(M(D), q(D)) = L(g(\tilde{v}), q(D))$ is weakly decreasing in \tilde{v} over the same range.

The risk of the mechanism can therefore be bounded below by

$$\begin{aligned}
\mathbb{E}[L(M(D), q(D))] &= P[\tilde{v} < f(a)] \mathbb{E}[L(g(\tilde{v}), q(D)) | \tilde{v} < f(a)] \\
&\quad + P[f(a) \leq \tilde{v}] \mathbb{E}[L(f^{-1}(\tilde{v}) - a, q(D)) | f(a) \leq \tilde{v}] \\
&\geq P[\tilde{v} < f(a)] \mathbb{E}[L(0, q(D)) | \tilde{v} < f(a)] \\
&\quad + P[f(a) \leq \tilde{v}] \mathbb{E}[L(f^{-1}(\tilde{v}) - a, q(D)) | f(a) \leq \tilde{v}].
\end{aligned}$$

This lower bound on the risk is clearly achieved by

$$g(x) = \begin{cases} 0 & \text{if } x < f(a) \\ f^{-1}(x) - a & \text{if } f(a) \leq x. \end{cases}$$

This completes the proof. \square

APPENDIX C. TRANSFORMATION MECHANISM ERROR BOUNDS

Here, we present quantitative measures of error incurred by the transformation mechanisms. These are presented in the form of prediction intervals for the noisy count, given the true count. That is, given the true count, the noisy count will be within a range $[a, b]$ with some probability p .

Theorem C.1 (Transformation Mechanism Error Bounds). *Let*

$$\Phi^{-1}(p, \mu, \sigma) = \mu + \sigma\sqrt{2} \operatorname{erf}^{-1}(2p - 1)$$

be the inverse CDF of the normal distribution with mean μ and standard deviation σ . Let X be the range given by $[\Phi^{-1}((1-p)/2, f(q(D) + a), \sigma), \Phi^{-1}((1+p)/2, f(q(D) + a), \sigma)]$. Under Assumption 4.3, the output of Algorithm 1($q(D), a, \sigma, f, g$), \tilde{S} , lies within $[\min_{x \in X} g(x), \max_{x \in X} g(x)]$ with probability p .

Proof. Recall that the noisy transformed query value created in Line 3 of Algorithm 1 is denoted by \tilde{v} and the algorithm's output is $\tilde{S} = g(\tilde{v})$. Note that \tilde{v} is distributed according to $N(f(q(D) + a), \sigma^2)$ and, as such, has the prediction interval $[\Phi^{-1}((1-p)/2, f(q(D) + a), \sigma), \Phi^{-1}((1+p)/2, f(q(D) + a), \sigma)]$ with probability p . Since $\tilde{S} = g(\tilde{v})$, the prediction interval for \tilde{S} has bounds given simply by the minimum and maximum of the estimator $g(x)$ within the prediction interval for \tilde{v} . \square

These bounds give the prediction interval for the output of the transformation mechanism for any valid transformation and estimator. Note that the estimators we discuss here are typically asymmetrically distributed, resulting in asymmetric prediction intervals. This can be seen below, where we derive the 95% prediction intervals for the square root transformation and log transformation.

Corollary C.2 (Square Root Transformation 95% Error Bound Mean Unbiased Estimator). *Let $a > 0$, $f(x) = \sqrt{x}$, $g(x) = x^2 - \sigma^2 - a$, $p = .95$, and $q(D) \in [0, \infty)$. Let X be the range given by $[\Phi^{-1}((1-p)/2, f(q(D) + a), \sigma), \Phi^{-1}((1+p)/2, f(q(D) + a), \sigma)]$.*

The 95% prediction interval for the transformation mechanism is $[-\sigma^2 - a, q(D) + 2\sqrt{q(D) + a(1.96\sigma) + (1.96^2 - 1)\sigma^2}]$ when $0 \in X$ and $[q(D) - 2\sqrt{q(D) + a(1.96\sigma) + (1.96^2 - 1)\sigma^2}, q(D) + 2\sqrt{q(D) + a(1.96\sigma) + (1.96^2 - 1)\sigma^2}]$ when $0 \notin X$.

Proof. Note that, with $p = .95$, we have $\Phi^{-1}((1-p)/2, \sqrt{q(D) + a}, \sigma) \approx \sqrt{q(D) + a} - 1.96\sigma$ and $\Phi^{-1}((1+p)/2, \sqrt{q(D) + a}, \sigma) \approx \sqrt{q(D) + a} + 1.96\sigma$. Note that the global minimum of $g(x)$ is at $x = 0$. Similarly, since $q(D) \geq 0$, the maximum of $g(x)$ is always at the largest value in X . Likewise, since $q(D) \geq 0$, if X does not contain the global minimum at $x = 0$, then the minimum of $g(x)$ over X is realized at the minimum value in X . Therefore, when $0 \in X$, the lower bound is $g(0) = -\sigma^2 - a$ and the upper bound is as follows.

$$g(\sqrt{q(D) + a} + 1.96\sigma) =$$

$$\begin{aligned}
& (\sqrt{q(D) + a} + 1.96\sigma)^2 - \sigma^2 - a = \\
& q(D) + a + 2\sqrt{q(D) + a}(1.96\sigma) + (1.96\sigma)^2 - \sigma^2 - a = \\
& q(D) + 2\sqrt{q(D) + a}(1.96\sigma) + (1.96\sigma)^2 - \sigma^2 = \\
& q(D) + 2\sqrt{q(D) + a}(1.96\sigma) + (1.96^2 - 1)\sigma^2.
\end{aligned}$$

When $0 \notin X$, the upper bound remains the same, but the lower bound is the lower boundary of X , as follows.

$$\begin{aligned}
& g(\sqrt{q(D) + a} - 1.96\sigma) = \\
& (\sqrt{q(D) + a} - 1.96\sigma)^2 - \sigma^2 - a = \\
& q(D) + a - 2\sqrt{q(D) + a}(1.96\sigma) + (1.96\sigma)^2 - \sigma^2 - a = \\
& q(D) - 2\sqrt{q(D) + a}(1.96\sigma) + (1.96\sigma)^2 - \sigma^2 = \\
& q(D) - 2\sqrt{q(D) + a}(1.96\sigma) + (1.96^2 - 1)\sigma^2.
\end{aligned}$$

□

For the parameters $a, \sigma = 1$ and query value $q(D) = 1000$, these bounds result in a range of approximately $[879, 1127]$ which is 127 in the positive direction and 121 in the negative direction. Due to the asymmetry of the estimator, these bounds are larger in the positive direction than in the negative direction. This can be seen for the log transformation, as well, in the following.

Corollary C.3 (Log Transformation 95% Error Bound Mean Unbiased Estimator). *Let $a > 0$, $f(x) = \ln(x)$, $g(x) = e^{x - \frac{\sigma^2}{2}} - a$, $p = .95$, and $q(D) \in [0, \infty)$. Let X be the range given by $[\Phi^{-1}((1-p)/2, f(q(D)+a), \sigma), \Phi^{-1}((1+p)/2, f(q(D)+a), \sigma)]$. The 95% prediction interval for the transformation mechanism is $[(q(D) + a)e^{-1.96\sigma - \frac{\sigma^2}{2}} - a, (q(D) + a)e^{1.96\sigma - \frac{\sigma^2}{2}} - a]$.*

Proof. Note that $\Phi^{-1}((1-p)/2, \ln(q(D) + a), \sigma) \approx \ln(q(D) + a) - 1.96\sigma$ and $\Phi^{-1}((1+p)/2, \ln(q(D) + a), \sigma) \approx \ln(q(D) + a) + 1.96\sigma$. Since $g(x)$ is monotonic in x , the minimum and maximum are respectively at the lower and upper boundaries of X . Start with the lower bound.

$$\begin{aligned}
& g(\ln(q(D) + a) - 1.96\sigma) = \\
& e^{\ln(q(D)+a) - 1.96\sigma - \frac{\sigma^2}{2}} - a = \\
& (q(D) + a)e^{-1.96\sigma - \frac{\sigma^2}{2}} - a.
\end{aligned}$$

Similarly, for the upper bound.

$$\begin{aligned}
& g(\ln(q(D) + a) + 1.96\sigma) = \\
& e^{\ln(q(D)+a) + 1.96\sigma - \frac{\sigma^2}{2}} - a = \\
& (q(D) + a)e^{1.96\sigma - \frac{\sigma^2}{2}} - a.
\end{aligned}$$

□

For the parameters $a, \sigma = 1$ and query value $q(D) = 1000$, these bounds result in a range of approximately $[85, 4309]$, which is 3309 in the positive direction and 915 in the negative direction. This again results in a much higher range in the positive direction than in the negative direction. This is a result of the asymmetry of the estimator, which stretches positively noised values to even larger values while negatively noised values are compressed together.

Among the mechanisms described here, the log transformation mechanism is unique in having error bounds that are linear in $q(D)$, for any choice of parameters or the confidence value, p . This is due to a natural property of exponential functions, where $\exp(a + b) = \exp(a)\exp(b)$. This property transforms the additive differences between the transformed query and its bounds into multiplicative differences when the estimator is applied.

The error bounds for all slowly scaling transformation mechanisms are dependent on the true query values and cannot be released directly in a privacy-preserving manner. Much like the policy function, the functional form of the error bounds can be released instead. This would allow a data curator to find the prediction interval for any known query value without directly leaking the true query values.

APPENDIX D. ADDITIVE MECHANISM ERROR BOUNDS

Here, we present quantitative measures of error incurred by the additive mechanisms. These are presented in the form of prediction intervals for the noisy count given the true count. That is, given the true count, the noisy count will be within a range $[a, b]$ with probability p . Since the additive mechanisms are sampled directly from known distributions, it is easy to find these bounds using the inverse CDF function as follows.

Theorem D.1 (Additive Mechanism Error Bounds). *Let f_Z be a noise distribution that has density proportional to $e^{f(|z|)}$, where $f : [0, \infty) \rightarrow \mathbb{R}$ is (weakly) decreasing and (weakly) convex.*

Let $\text{CDF}^{-1}(p, \mu)$ be the inverse CDF of f_Z when centered at μ . The output of Algorithm 2 ($q(D), f_Z$) lies within $[\text{CDF}^{-1}((1-p)/2, q(D)), \text{CDF}^{-1}((1+p)/2, q(D))]$ with probability at least p .

Proof. This follows directly from the definition of the inverse CDF function of a distribution and the fact that $((1+p)/2 - (1-p)/2) = p$. \square

Unlike the transformation mechanism, these bounds are symmetric about the true query answer. Further, the error bounds for an additive mechanism (that is, the interval obtained by subtracting the true query answer from every point in the prediction interval from Theorem D.1) are independent of the true query answer and can be published directly. We use Theorem D.1 to derive the bounds for two examples of additive mechanisms below.

Corollary D.2 (Generalized Gaussian $p = 1/2$). *Let $f_Z(z) = \frac{1}{4\sigma} e^{-\left(\frac{|z|}{\sigma}\right)^{1/2}}$. The 95% prediction interval for Algorithm 2($q(D), f_Z$) is $[q(D) - 22.50425\sigma, q(D) + 22.50425\sigma]$.*

Proof. First, note that if the random variable X has CDF $F(x)$, then the location scale family $\mu + \delta X$ has CDF $F\left(\frac{x-\mu}{\delta}\right)$ and therefore inverse CDF (with argument denoted by y) $\mu + \sigma F^{-1}(y)$. The documentation (Griffin, 2018b) for the R package gnorm (Griffin, 2018a)

gives the following formula for the inverse CDF of the generalized Gaussian distribution with $\sigma = 1$:

$$\text{sign}(y - .5)g(2|y - .5|; \frac{1}{p}, 1)^{\frac{1}{p}}, \quad (\text{D.1})$$

where $g(2|y - .5|; \frac{1}{p}, 1)$ is the inverse CDF of the gamma distribution with rate 1 and shape $\frac{1}{p}$, evaluated at $2|y - .5|$. It follows that the inverse CDF of the generalized Gaussian distribution centered at μ and with scale parameter σ is

$$\mu + \sigma \text{sign}(y - .5)g(2|y - .5|; \frac{1}{p}, 1)^{\frac{1}{p}}. \quad (\text{D.2})$$

Evaluating this inverse CDF at $\mu = q(D)$, $p = .5$, $y \in (.025, .975)$ and applying Theorem D.1 gives us the result. \square

The same can be done for special cases of the exponential polylogarithmic distribution.

Corollary D.3 (Exponential Polylogarithmic Distribution $p = 2$). *Let $f_Z(z) \propto e^{-d \ln(\frac{|z|}{\sigma} + a)^2}$, with $\sigma > 0$, $a \geq e$, and $d > 0$. Let $\Phi^{-1}(\cdot)$ denote the standard normal quantile function. The 95% prediction interval for Algorithm 2($q(D)$, f_Z) is as follows:*

$$\left[q(D) - \sigma \exp\left([2d]^{-1/2} \Phi^{-1}\left([1 - \Phi\left(\frac{\ln(a) - (2d)^{-1}}{(2d)^{-1/2}}\right)][0.95]\right) + \Phi\left(\frac{\ln(a) - (2d)^{-1}}{(2d)^{-1/2}}\right)\right) + (2d)^{-1} + \sigma a, \right. \\ \left. q(D) + \sigma \exp\left([2d]^{-1/2} \Phi^{-1}\left([1 - \Phi\left(\frac{\ln(a) - (2d)^{-1}}{(2d)^{-1/2}}\right)][0.95]\right) + \Phi\left(\frac{\ln(a) - (2d)^{-1}}{(2d)^{-1/2}}\right)\right) + (2d)^{-1} - \sigma a \right]$$

Proof. The inverse CDF for the exponential poly-logarithmic distribution with $p = 2$ centered at 0 is provided in Table 3 and can be expanded into the following expression.

$$\begin{cases} -\sigma \exp\left([2d]^{-1/2} \Phi^{-1}\left[t\left([1 - \Phi\left(\frac{\ln(a) - (2d)^{-1}}{(2d)^{-1/2}}\right)][1 - 2x]\right) + \Phi\left(\frac{\ln(a) - (2d)^{-1}}{(2d)^{-1/2}}\right)\right) + (2d)^{-1} + \sigma a & x < \frac{1}{2} \\ \sigma \exp\left([2d]^{-1/2} \Phi^{-1}\left([1 - \Phi\left(\frac{\ln(a) - (2d)^{-1}}{(2d)^{-1/2}}\right)][2x - 1]\right) + \Phi\left(\frac{\ln(a) - (2d)^{-1}}{(2d)^{-1/2}}\right)\right) + (2d)^{-1} - \sigma a & x > \frac{1}{2} \\ 0 & x = \frac{1}{2} \end{cases} \quad (\text{D.3})$$

We then plug in the values for $x = 0.025$ for the lower bound and $x = 0.975$ for the upper bound, and shift the range to be centered at $q(D)$, to obtain the desired bounds. \square

For parameters $\sigma = 1$, $d = 1$, $a = e$ this results in a bound of approximately $[q(D) - 5.418, q(D) + 5.418]$

APPENDIX E. PROTECTING VALUES – BOUNDED NEIGHBOR PRIVACY GUARANTEES

The privacy concepts treated thus far - those in Definitions 2.3, 2.11, and 5.1 - are designed to keep an adversary from distinguishing between pairs of neighboring databases which differ in the *presence* of a single record. These privacy concepts keep each record's presence private from an adversary who may already know that record's value (e.g., an adversary who knows a certain business's revenue, but does not know whether that business is included in a particular dataset). In this section, we develop privacy concepts that protect the privacy of records' *values* from an adversary who may already know of their presence (e.g., an adversary who knows that a certain business is in a dataset, but does not know its revenue).

With these privacy concepts, we derive privacy guarantees that express how well our slowly scaling mechanisms inhibit adversaries from distinguishing between a database containing the record r and an otherwise identical neighboring database that instead contains the record r' .

Denote attribute c of record r by $r.c \in \mathbb{R}$. We focus particularly on the case where $r'.c$ equals either $r.c + \delta$ or $\delta r.c$ for some scalar δ , and where our slowly scaling mechanisms are used to protect a sum of attribute c over all records in the database. Using these neighboring databases, we can quantify the difficulty an adversary has in distinguishing between a record's true value r and nearby values. When the privacy guarantee is high for a wide range of values of δ , a data curator may be satisfied that adversaries will be unable to accurately estimate the record r .

Using language from Kifer and Machanavajjhala (2011), we call databases that differ only in the value of one record "bounded" neighbors.

Definition E.1 (Bounded Neighboring Databases). Two databases D and D' in \mathcal{D} are considered *bounded neighboring databases* with the *differing pair* (r, r') if, for some database $D_0 \in \mathcal{D}$ and records r, r' , we have either $D = D_0 \cup \{r\}$ and $D' = D_0 \cup \{r'\}$, or $D = D_0 \cup \{r'\}$ and $D' = D_0 \cup \{r\}$.

Using this terminology, the neighbors defined in Definition 2.2 and used elsewhere in this paper are called "unbounded" neighbors.

We start by defining privacy concepts that bound the privacy loss between bounded neighboring databases with a given differing pair, (r, r') . We give these bounded analogues of both PRzCDP (Definition 2.11) and PRDP (Definition 5.1) below.

Definition E.2 (Bounded Per-Record Differential Privacy). The randomized mechanism M satisfies the following privacy concepts with the policy function P iff the following conditions hold for all bounded neighboring databases D, D' with differing pair (r, r') .

- Bounded P -Per-Record Zero-Concentrated Differential Privacy (Bounded P -PRzCDP):

$$d_\alpha (M(D) \| M(D')) \leq \alpha P(r, r') \text{ for all } \alpha \in (1, \infty).$$

- Bounded P -Per-Record Differential Privacy (Bounded P -PRDP):

$$d_\infty (M(D) \| M(D')) \leq P(r, r').$$

When the policy function $P(r, r')$ is a metric, the definition of bounded PRDP in Definition E.2 is the same as that of another privacy concept variously called "metric DP", "Lipschitz DP", and " d -DP" (see, for example, Chatzikokolakis et al. (2013) and Koufogiannis et al. (2015)). A closely related privacy concept is Gaussian Establishment DP, introduced in independent, concurrent work by Webb et al. (2025) for use with the same sort of establishment data that motivated our work. Gaussian Establishment DP considers bounded neighbors and bounds a notion of privacy loss based on the Gaussian DP concept introduced in Dong et al. (2022). Webb et al. (2025) also propose Gaussian Establishment DP mechanisms which include special cases of the transformation mechanisms described in Algorithm 1.

Like in the unbounded per-record DP variants, where a policy function bounds a notion of privacy loss as a function of the record r , the bounded per-record privacy concepts have a policy function that bounds the privacy loss as a function of the differing pair r, r' . Just as we needed the per-record sensitivity (Definition 4.2) to derive the (unbounded) PRzCDP and PRDP guarantees above, we need an analogous sensitivity notion that captures the

maximal sensitivity of the query q to switching between bounded neighbors with a particular differing pair.

Definition E.3 (Bounded Per-Record Sensitivity). The bounded per-record sensitivity of the query $q : \mathcal{D} \rightarrow \mathbb{R}$ to the differing pair (r, r') is

$$\Delta(r, r') \equiv \sup_{\substack{D, D' | D, D' \text{ bounded neighbors} \\ \text{with differing pair } (r, r')}} |q(D) - q(D')|.$$

With these concepts in hand, we now derive the bounded PRzCDP and bounded PRDP guarantees of our slowly scaling mechanisms, starting with the transformation mechanisms.

Theorem E.4 (Bounded PRzCDP Guarantees for Transformation Mechanisms). *Let Assumption 4.3 hold and denote by $\Delta_f(r, r')$ the bounded sensitivity of the query $f(q(D) + a)$, defined in Definition E.3. The output of Algorithm 1($q(D), a, \sigma, f, g$) satisfies bounded P -PRzCDP for $P(r, r') = \frac{\Delta_f(r, r')^2}{2\sigma^2}$.*

Proof. Let $M(D)$ denote Algorithm 1($q(D), a, \sigma, f, g$). Note that $M(D) = g(f(q(D) + a) + N(0, \sigma^2))$. It follows from Definition E.2 that M satisfies bounded PRzCDP for any policy function that is lower bounded by

$$\sup_{\substack{\alpha \in (1, \infty), D, D' | D, D' \text{ bounded neighbors} \\ \text{with differing pair } (r, r')}} \frac{d_\alpha(M(D) \| M(D'))}{\alpha}$$

To derive $P(r, r')$, we bound this supremum using the data processing inequality (Rényi, 1961), Lemma 4.1, and Definition E.3, as follows.

$$\begin{aligned} & \sup_{\substack{\alpha \in (1, \infty), D, D' | D, D' \text{ bounded neighbors} \\ \text{with differing pair } (r, r')}} \frac{d_\alpha(M(D) \| M(D'))}{\alpha} \\ & \leq \sup_{\substack{\alpha \in (1, \infty), D, D' | D, D' \text{ bounded neighbors} \\ \text{with differing pair } (r, r')}} \frac{d_\alpha(f(q(D) + a) + N(0, \sigma^2) \| f(q(D') + a) + N(0, \sigma^2))}{\alpha} \\ & = \sup_{\substack{\alpha \in (1, \infty), D, D' | D, D' \text{ bounded neighbors} \\ \text{with differing pair } (r, r')}} \frac{|f(q(D) + a) - f(q(D') + a)|^2}{2\sigma^2} \\ & = \left(\sup_{\substack{D, D' | D, D' \text{ bounded neighbors} \\ \text{with differing pair } (r, r')}} |f(q(D) + a) - f(q(D') + a)| \right)^2 / 2\sigma^2 \\ & = \frac{\Delta_f(r, r')^2}{2\sigma^2} \\ & = P(r, r'). \end{aligned}$$

□

We now similarly derive the bounded PRDP and bounded PRzCDP guarantees for the additive mechanisms. This requires that we first state extensions of Lemma 5.2 and Theorem 5.3 to our bounded DP concepts. The proofs of these extensions closely follow the proofs of Lemma 5.2 and Theorem 5.3, and so are omitted here.

Lemma E.5 (Bounded PRDP Implies Bounded PRzCDP). *If a mechanism, M , satisfies bounded P -PRDP, it also satisfies bounded P -PRzCDP and bounded P' -PRzCDP, where $P'(r) = \tanh(P(r)/2)P(r)$.*

Theorem E.6. *An additive mechanism $M(D) \equiv q(D) + Z$, with a continuous random variable Z having density f_Z , satisfies bounded P -PRDP if and only if, for all bounded neighboring databases D, D' with differing pair (r, r')*

$$\sup_z |\ln(f_Z(z)) - \ln(f_Z(z + q(D) - q(D')))| \leq P(r, r').$$

The following theorem is an analogue of Theorem 5.4 for our bounded DP concepts.

Theorem E.7 (Bounded Privacy of Symmetric, Absolute-Log-Convex Densities). *Let q be a query with bounded sensitivity $\Delta(r, r')$, as defined in Definition E.3. Let f_Z be a noise distribution that has density proportional to $e^{f(|z|)}$, where $f : [0, \infty) \rightarrow \mathbb{R}$ is (weakly) decreasing and (weakly) convex.*

For all pairs of records r, r' , Algorithm 2($q(D), f_Z$), denoted below by M , satisfies bounded P -PRDP with $P(r, r') = f(0) - f(\Delta(r, r'))$. This policy function is tight, in the sense that M is not bounded P' -PRDP for an alternative policy function P' if there exists a pair of records r, r' such that $P'(r, r') < P(r, r')$. M also satisfies bounded P'' -PRzCDP with $P''(r, r') = \tanh(P(r, r')/2)P(r, r')$.

Proof. Applying Theorem E.6 and then Theorem 5.4, we see that M is bounded P -PRDP if and only if, for all r, r' ,

$$P(r, r') \geq \sup_{\substack{z, D, D' | D, D' \text{ bounded neighbors} \\ \text{with differing pair } (r, r')}} |\ln(f_Z(z)) - \ln(f_Z(z + q(D) - q(D')))| \quad (\text{E.1})$$

$$= \sup_{\substack{D, D' | D, D' \text{ bounded neighbors} \\ \text{with differing pair } (r, r')}} f(0) - f(|q(D) - q(D')|). \quad (\text{E.2})$$

Because f is a decreasing function, $f(0) - f(|q(D) - q(D')|)$ is increasing in $|q(D) - q(D')|$ and has its supremum at $|q(D) - q(D')| = \Delta(r, r')$. This yields the result. Furthermore, by Lemma E.5, M also satisfies bounded P'' -PRzCDP with $P''(r, r') = \tanh(P(r, r')/2)P(r, r')$. \square

The bounded privacy guarantees allow us to determine the greatest possible privacy loss between bounded neighbors where one neighbor contains a real record r and another contains a hypothetical alternative record r' that is within some neighborhood of r .

E.1. Global Sums. In this section, we derive the mechanisms' bounded privacy guarantees for global sum queries. Denoting the scalar attribute c of record r by $r.c \in \mathbb{R}$, a global sum query q of attribute c is simply the sum of $r.c$ over every record in the database; $q(D) = \sum_{r \in D} r.c$. We obtain the bounded PRzCDP and bounded PRDP policy functions and further highlight these policy functions' values for differing pairs r, r' satisfying $r'.c = r.c + \delta$

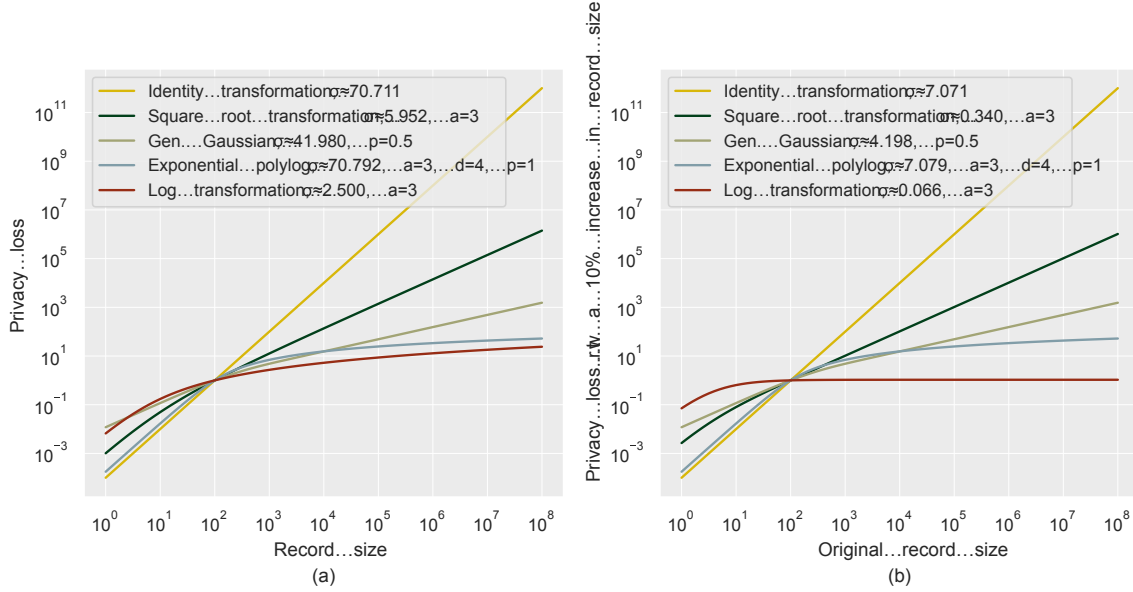


FIGURE 12. Unbounded and bounded PRzCDP policy functions under several mechanisms applied to the global sum query $q(D) = \sum_{r \in D} r.c$. The x -axes show values of $r.c$. Panel (a) shows unbounded PRzCDP policy functions. Panel (b) shows bounded PRzCDP policy functions with respect to a 10% increase in record value (i.e., with a differing pair such that $r'.c = 1.1r.c$). All mechanisms are tuned to have policy functions that equal 1 for records with $r.c = 100$.

or $r'.c = \delta r.c$ for some scalar δ . These expressions let us easily determine the privacy loss against an adversary trying to distinguish between a true value $r.c$ and other, nearby values.

We start by stating the bounded sensitivity of the global sum query. The proof is trivial.

Lemma E.8 (Bounded Sensitivity of Sum Query). *The bounded sensitivity of the global sum query $q(D) = \sum_{r \in D} r.c$ is $|r.c - r'.c|$.*

The above suffices to obtain the privacy guarantees for the additive mechanisms, but, for the transformation mechanisms, Theorem E.4 requires that we obtain the sensitivities of the transformation function f composed with the global sum query. These sensitivities are given in the following theorem.

Lemma E.9 (Bounded Sensitivity of Transformed Sum Query). *Let $a \geq 0$ and $f : [a, \infty) \rightarrow \mathcal{F} \subseteq \mathbb{R}$ be an increasing, concave function. Let $r.c \geq 0$ for all records r . For the global sum query $q(D) = \sum_{r \in D} r.c$, the bounded sensitivity of $f(q(D) + a)$ is $\Delta_f(r, r') = |f(r.c + a) - f(r'.c + a)|$.*

Proof. The bounded sensitivity of $f(q(D) + a)$ is defined as

$$\Delta_f(r, r') \equiv \sup_{D, D' | D, D' \text{ bounded neighbors with differing pair } (r, r')} \left| f\left(a + \sum_{s \in D} s.c\right) - f\left(a + \sum_{s \in D'} s.c\right) \right| \quad (\text{E.3})$$

$$= \sup_{D_0 \in \mathcal{D}; r, r'} \left| f\left(a + r.c + \sum_{s \in D_0} s.c\right) - f\left(a + r'.c + \sum_{s \in D_0} s.c\right) \right|. \quad (\text{E.4})$$

Because f is concave and increasing, the final supremum is decreasing in $\sum_{s \in D_0} s.c.$. $\sum_{s \in D_0} s.c.$ has a minimum of 0, so the supremum is

$$\Delta_f(r, r') = |f(r.c + a) - f(r'.c + a)|.$$

□

The following theorem lays out the bounded PRzCDP guarantees for the transformation mechanisms applied to a global sum query. It also provides easily interpretable bounds on the policy function when the differing pairs satisfy $r'.c = r.c + \delta$ and $r'.c = \delta r.c$.

Corollary E.10 (Bounded PRzCDP Guarantee for Transformation Mechanisms on Global Sums). *Let Assumption 4.3 hold and additionally assume that $a \geq 0$ and $r.c \geq 0$ for all records r . With the global sum query $q(D) = \sum_{r \in D} r.c$, the output of Algorithm 1($q(D), a, \sigma, f, g$) satisfies bounded P-PRzCDP for $P(r, r') = \frac{|f(r.c+a) - f(r'.c+a)|^2}{2\sigma^2}$. Easily interpretable bounds on special cases of this policy function are given below.*

- (1) Let $r'.c = r.c + \delta \geq 0$.
 - (a) $P(r, r') = \frac{|f(r.c+a) - f(r.c+\delta+a)|^2}{2\sigma^2} \leq \frac{|f(a) - f(|\delta|+a)|^2}{2\sigma^2}$.
 - (b) If $f(x) = x^k$ for some $0 < k \leq 1$, then $P(r, r') \leq \frac{|\delta|^{2k}}{2\sigma^2}$.
 - (c) If $a > 0$ and $f(x) = \ln(x)$, then $P(r, r') \leq \frac{|\ln(a) - \ln(|\delta|+a)|^2}{2\sigma^2}$.
- (2) Let $r'.c = \delta r.c$, with $\delta \geq 0$.
 - (a) $P(r, r') = \frac{|f(r.c+a) - f(\delta r.c+a)|^2}{2\sigma^2}$.
 - (b) If $f(x) = x^k$ for some $0 < k \leq 1$, then $P(r, r') \leq \frac{|(1-\delta^k)r^k|^2}{2\sigma^2}$.
 - (c) If $r.c > 0$, $a > 0$, $\delta > 0$, and $f(x) = \ln(x)$, then $P(r, r') \leq \frac{|\ln(\delta)|^2}{2\sigma^2}$.

Proof. The policy function $P(r, r') = \frac{|f(r.c+a) - f(r'.c+a)|^2}{2\sigma^2}$ follows immediately from Lemma E.9 and Theorem E.4.

Recall that, if f is increasing and concave and if $x_0 < x_1$, then $f(x_1 + b) - f(x_0 + b)$ is nonnegative and decreasing in b .

If $\delta \geq 0$, then the inequality in Result 1a follows from $r.c \geq 0$ and f being concave and increasing. If $\delta < 0$, then we use these facts and the fact that $r.c + \delta \geq 0$ to obtain $\frac{|f(r.c+a) - f(r.c+\delta+a)|^2}{2\sigma^2} = \frac{|f(r.c+\delta-\delta+a) - f(r.c+\delta+a)|^2}{2\sigma^2} \leq \frac{|f(-\delta+a) - f(a)|^2}{2\sigma^2} = \frac{|f(|\delta|+a) - f(a)|^2}{2\sigma^2}$.

Result 1b follows from Result 1a, $a \geq 0$, and f being concave, increasing, and defined at 0.

Result 1c follows from Result 1a.

Result 2b follows from $P(r, r') = \frac{|(r.c+a)^k - (\delta r.c+a)^k|^2}{2\sigma^2} \leq \frac{|r.c^k - \delta^k r.c^k|^2}{2\sigma^2} = \frac{|(1-\delta^k)r^k|^2}{2\sigma^2}$.

Result 2c follows from $P(r, r') = \frac{|\ln(r.c+a) - \ln(\delta r.c+a)|^2}{2\sigma^2} \leq \frac{|\ln(r.c) - \ln(\delta r.c)|^2}{2\sigma^2} = \frac{|\ln(\delta)|^2}{2\sigma^2}$. In both these derivations, the inequalities use the fact that f is concave and increasing. □

The following theorem gives corresponding privacy guarantees for the additive mechanisms and follows immediately from Lemma E.8 and Theorem E.7.

Corollary E.11 (Bounded Privacy Guarantees for Additive Mechanisms on Global Sums). *Let f_Z be a noise distribution that has density proportional to $e^{f(|z|)}$, where $f : [0, \infty) \rightarrow \mathbb{R}$ is (weakly) decreasing and (weakly) convex. With the global sum query $q(D) = \sum_{r \in D} r.c$, the output of Algorithm 2($q(D), f_Z$) satisfies bounded P-PRDP for $P(r, r') = f(0) - f(|r.c - r'.c|)$.*

Transformation Mechanism with Transformation f – Bounded PRzCDP Policy Functions

Name	f	Additive: $r'.c = r.c + \delta$	Multiplicative: $r'.c = \delta r.c$
Generic	$f(q(D) + a)$	$\frac{ f(r.c+a) - f(r.c+\delta+a) ^2}{2\sigma^2}$	$\frac{ f(r.c+a) - f(\delta r.c+a) ^2}{2\sigma^2}$
Identity	$q(D) + a$	$\frac{ \delta ^2}{2\sigma^2}$	$\frac{(1-\delta)^2 r^2}{2\sigma^2}$
Square Root*	$\sqrt{q(D) + a}$	$\frac{ \delta }{2\sigma^2}$	$\frac{(1-\sqrt{\delta})^2 r}{2\sigma^2}$
Fourth Root*	$\sqrt[4]{q(D) + a}$	$\frac{\sqrt{ \delta }}{2\sigma^2}$	$\frac{(1-\sqrt[4]{\delta})^2 \sqrt{r}}{2\sigma^2}$
Log*	$\ln(q(D) + a)$	$\frac{(\ln(\delta +a) - \ln(a))^2}{2\sigma^2}$	$\frac{\ln(\delta)^2}{2\sigma^2}$

Additive Mechanism with Noise Density $\propto e^{f(|z|)}$ – Bounded PRDP Policy Functions

Name	f	Additive: $r'.c = r.c + \delta$	Multiplicative: $r'.c = \delta r.c$
Generic	$f(z)$	$f(0) - f(\delta)$	$f(0) - f((1-\delta)r.c)$
Gen Gaussian	$-\left(\frac{ z }{\sigma}\right)^p$	$\left(\frac{ \delta }{\sigma}\right)^p$	$\left(\frac{ (1-\delta)r.c }{\sigma}\right)^p$
Exp Polylog	$-d \ln\left(\frac{ z }{\sigma} + a\right)^p$	$d[\ln\left(\frac{ \delta }{\sigma} + a\right)^p - \ln(a)^p]$	$d[\ln\left(\frac{ (1-\delta)r.c }{\sigma} + a\right)^p - \ln(a)^p]$

TABLE 8. Bounded PRzCDP and bounded PRDP policy functions for a global sum query privatized by our transformation and additive mechanisms. Policy functions are given for differing pairs satisfying $r'.c = r.c + \delta$ and $r'.c = \delta r.c$. The policy functions shown for transformation mechanisms with stars in the Name column are loose, but tractable, bounds on privacy loss. The generic policy function in the first row gives tighter privacy loss bounds when applied to these mechanisms. See Corollaries E.10 and E.11 for details.

This algorithm also satisfies bounded P' -PRzCDP for $P'(r, r') = \tanh(P(r, r')/2)P(r, r')$. Expressions for special cases of this policy function are given below.

- (1) Let $r'.c = r.c + \delta$, with $\delta \in \mathbb{R}$. $P(r, r') = f(0) - f(|\delta|)$.
 - (a) (Generalized Gaussian) If $f(|z|) = -\left(\frac{|z|}{\sigma}\right)^p$ for $0 < \sigma$ and $0 < p \leq 1$, then $P(r, r') = \left(\frac{|\delta|}{\sigma}\right)^p$.
 - (b) (Exp polylog) If $f(|z|) = -d \ln\left(\frac{|z|}{\sigma} + a\right)^p$, for $\sigma > 0$, $a \geq e^{p-1}$, and either $(p > 1$ and $d > 0)$ or $(p = 1$ and $d > 1)$, then $P(r, r') = d[\ln\left(\frac{|\delta|}{\sigma} + a\right)^p - \ln(a)^p]$.
- (2) Let $r'.c = \delta r.c$ with $\delta \in \mathbb{R}$. $P(r, r') = f(0) - f(|(1-\delta)r.c|)$.
 - (a) (Generalized Gaussian) If $f(|z|) = -\left(\frac{|z|}{\sigma}\right)^p$ for $0 < \sigma$ and $0 < p \leq 1$, then $P(r, r') = \left(\frac{|(1-\delta)r.c|}{\sigma}\right)^p$.
 - (b) (Exp polylog) If $f(|z|) = -d \ln\left(\frac{|z|}{\sigma} + a\right)^p$, for $\sigma > 0$, $a \geq e^{p-1}$, and either $(p > 1$ and $d > 0)$ or $(p = 1$ and $d > 1)$, then $P(r, r') = d[\ln\left(\frac{|(1-\delta)r.c|}{\sigma} + a\right)^p - \ln(a)^p]$.

See Table 8 for a summary of the results in Corollaries E.10 and E.11. From these corollaries, we see that every mechanism's policy function has an upper bound that depends on δ , but not r when $r'.c = r.c + \delta$. This means that, with appropriately set parameters, all of these mechanisms can guarantee any desired level of privacy against an attacker who is trying to determine the value of $r.c$ up to an additive error δ .

However, these protections may be inadequate when $r.c$ is very large. Guaranteeing that an attacker cannot learn a business' true annual revenue without a \$1 million margin of

error may be adequate protection for, say, a local florist, but this may still be very disclosive for a large business with revenue in the billions. In that case, the large business may care more about protecting its revenue up to a multiplier δ , and so we look at differing pairs with $r'.c = \delta r.c$.

For these differing pairs, only the log transformation mechanism has an upper bound that does not depend on $r.c$. All the other mechanisms have protections that weaken as $r.c$ grows. For this reason, the log transformation mechanism may be acceptable even when its unbounded PRzCDP guarantees can be matched with greater utility by an additive mechanism, as in our examples in Section 6.

Despite that advantage, the other mechanisms' bounded policy functions can be quite slowly scaling in $r.c$. For the transformation mechanisms with $f(x) = x^k$, the policy function is bounded above by $\frac{|(1-\delta^k)r.c|^2}{2\sigma^2}$, so the privacy loss can be made to scale more slowly by picking a smaller value of k . Likewise, the generalized Gaussian mechanism has the policy function $(\frac{|(1-\delta)r.c|}{\sigma})^p$, which scales more slowly with smaller values of p . The exponential polylog mechanism has the policy function $d[\ln(\frac{|(1-\delta)r.c|}{\sigma} + a)^p - \ln(a)^p]$, which scales at a polylogarithmic rate. To see this slowly scaling behavior illustrated, see Figure 12 for graphs of unbounded and bounded PRzCDP policy functions for a variety of mechanisms. Even though these mechanisms do not impose a hard limit on the privacy losses between these types of bounded neighbors, the slow growth rates of the privacy loss may be acceptable for the same reasons as the unbounded slowly scaling privacy guarantees.

APPENDIX F. PROPERTIES OF THE EXPONENTIAL POLYLOGARITHMIC DISTRIBUTION

F.1. Density Function. For the general case, with $p \geq 1$, the density of the exponential polylogarithmic distribution is given by

$$f_Z(z) = \frac{pd^{\frac{1}{p}}}{2\sigma\Psi^{(\Gamma)}\left(p, d\ln(a)^p, d^{-\frac{1}{p}}\right)} e^{-d\ln\left(\frac{|z|}{\sigma} + a\right)^p}, \quad (\text{F.1})$$

where

$$\Psi^{(\Gamma)}(p, d\ln(x)^p, c) = {}_1\Psi_1^{(\Gamma)}\left[\begin{matrix} \left(\frac{1}{p}, \frac{1}{p}, d\ln(x)^p\right) \\ (1, 0) \end{matrix} \middle| c\right] = \sum_{n=0}^{\infty} \Gamma\left(\frac{1}{p} + \frac{n}{p}, d\ln(x)^p\right) \frac{c^n}{n!} \quad (\text{F.2})$$

is the upper-incomplete Fox-Wright Psi function, introduced by Srivastava et al. (2018), evaluated at $c = d^{-\frac{1}{p}}$ and $x = a$. This converges when $p > 1$ and $0 < d < \infty$ or when $p = 1$ and $d > 1$. Here, $\Gamma\left(\frac{1}{p} + \frac{n}{p}, d\ln(x)^p\right)$ is the upper incomplete gamma function with parameter $\frac{n+1}{p}$ with lower limit $d\ln(x)^p$.

Proof. First, we find the constant of integration C such that $C \int_{-\infty}^{\infty} f_Z(z) dz = 1$. Then, $C \int_{-\infty}^{\infty} f_Z(z) dz = 2C \int_0^{\infty} e^{-d\ln(\frac{z}{\sigma} + a)^p} dz$ by symmetry of $f_Z(z)$ around zero. Then, letting $u = \ln(\frac{z}{\sigma} + a)$ we have $2C\sigma \int_{\ln(a)}^{\infty} e^{-du^p + u} dz = 1$. Substituting the Taylor series for e^u and changing the order of summation and integration, we have

$$2C\sigma \sum_{n=0}^{\infty} \frac{1}{n!} \int_{\ln(a)}^{\infty} u^n e^{-du^p} = 1.$$

Substituting $t = du^p$ and simplifying

$$\frac{2\sigma C}{pd^{\frac{1}{p}}} \sum_{n=0}^{\infty} \Gamma\left(\frac{1}{p} + \frac{n}{p}, d \ln(a)^p\right) \frac{\left(d^{-\frac{1}{p}}\right)^n}{n!} = 1.$$

The series in the above equation is the upper-incomplete Fox-Wright function. Since the incomplete Fox-Wright function is bounded above by the more familiar complete Fox-Wright function, conditions for convergence are summarized by Kilbas (2004). The series converges when $p > 1$ and $d < \infty$ or when $p = 1$ and $d > 1$. Therefore, $C = \frac{pd^{\frac{1}{p}}}{2\sigma\Psi^{(\Gamma)}\left(p, d \ln(a)^p, d^{-\frac{1}{p}}\right)}$. \square

F.2. Higher Order Moments. Due to the symmetry of $f_Z(z)$ around zero, the odd moments are zero. The even moments are given by

$$E[z^{2k}] = \frac{\sigma^{2k}}{\Psi^{(\Gamma)}\left(p, d \ln(x)^p, d^{-\frac{1}{p}}\right)} \sum_{j=0}^{2k} \binom{2k}{j} (-a)^{2k-j} \Psi^{(\Gamma)}\left(p, d \ln(x)^p, (j+1)d^{-\frac{1}{p}}\right), \quad k = 1, 2, 3, \dots, \quad (\text{F.3})$$

where $\Psi^{(\Gamma)}\left(p, d \ln(x)^p, (j+1)d^{-\frac{1}{p}}\right)$ is the upper incomplete Fox-Wright function (F.2) evaluated at $c = (j+1)d^{-\frac{1}{p}}$ and $x = a$. The even moments are finite when $p > 1$ and $0 < d < \infty$ or when $p = 1$ and $d > 2k + 1$.

Proof. The even moments for $f_Z(z)$ are given by

$$\begin{aligned} E[z^{2k}] &= \int_{-\infty}^{\infty} z^{2k} f_Z(z) dz \\ &= \frac{pd^{\frac{1}{p}}}{\sigma\Psi^{(\Gamma)}\left(p, d \ln(a)^p, d^{-\frac{1}{p}}\right)} \int_0^{\infty} z^{2k} e^{-d \ln\left(\frac{z}{\sigma} + a\right)^p} dz \\ &= \frac{pd^{\frac{1}{p}}}{\Psi^{(\Gamma)}} \int_{\ln(a)}^{\infty} e^u (e^u - a)^{2k} e^{-du^p} du \\ &= \frac{\sigma^{2k} pd^{\frac{1}{p}}}{\Psi^{(\Gamma)}} \int_{\ln(a)}^{\infty} e^u \sum_{j=0}^{2k} \binom{2k}{j} e^{ju} (-a)^{2k-j} e^{-du^p} du \\ &= \frac{\sigma^{2k} pd^{\frac{1}{p}}}{\Psi^{(\Gamma)}} \sum_{j=0}^{2k} \binom{2k}{j} (-a)^{2k-j} \int_{\ln(a)}^{\infty} e^{u(j+1)} e^{-du^p} du \\ &= \frac{\sigma^{2k}}{\Psi^{(\Gamma)}} \sum_{j=0}^{2k} \binom{2k}{j} (-a)^{2k-j} \sum_{n=0}^{\infty} \Gamma\left(\frac{1}{p} + \frac{n}{p}, d \ln(a)^p\right) \frac{\left((j+1)d^{-\frac{1}{p}}\right)^n}{n!} \\ &= \frac{\sigma^{2k}}{\Psi^{(\Gamma)}} \sum_{j=0}^{2k} \binom{2k}{j} (-a)^{2k-j} \Psi^{(\Gamma)}\left(p, d \ln(a)^p, (j+1)d^{-\frac{1}{p}}\right). \end{aligned}$$

The second line uses the fact that the integrand is even, and the third line uses the substitution $u = \ln\left(\frac{z}{\sigma} + a\right)$. The fourth line uses the binomial theorem, and the fifth through sixth line

interchange summation and integration. The final line uses the definition of the upper-incomplete Fox-Wright function. \square

F.3. Cumulative Distribution Function. The CDF is given by

$$F_Z(z) = \frac{1}{2} + \frac{\text{sign}(z)}{2\Psi^{(\Gamma)}(p, d\ln(a)^p, d^{-\frac{1}{p}})} \times (\Psi^{(\gamma)}(p, d\ln(\frac{|z|}{\sigma} + a)^p, d^{-\frac{1}{p}}) - \Psi^{(\gamma)}(p, d\ln(a)^p, d^{-\frac{1}{p}})), \quad (\text{F.4})$$

where

$$\Psi^{(\gamma)}(p, d\ln(x)^p, c) = {}_1\Psi_1^{(\gamma)} \left[\begin{matrix} \left(\frac{1}{p}, \frac{1}{p}, d\ln(x)^p \right) \\ (1, 0) \end{matrix} \middle| c \right] = \sum_{n=0}^{\infty} \gamma\left(\frac{1}{p} + \frac{n}{p}, d\ln(x)^p\right) \frac{c^n}{n!}, \quad (\text{F.5})$$

is the lower-incomplete Fox-Wright function as a function of x , evaluated at $c = d^{-\frac{1}{p}}$, where $\gamma\left(\frac{1}{p} + \frac{n}{p}, d\ln(x)^p\right)$ is the lower incomplete gamma function with parameter $\frac{n+1}{p}$ whose upper limit is $d\ln(x)^p$. This function converges with the same constraints as the upper-incomplete Fox-Wright function.

Proof. By definition, $F_Z(t) = \int_{-\infty}^t f_Z(z) dz$. Depending on the sign of t we have the following cases

$$F_Z(t) = \begin{cases} \int_{-\infty}^t f_Z(z) dz & t \leq 0 \\ \int_{-\infty}^0 f_Z(z) dz + \int_0^t f_Z(z) dz & t > 0. \end{cases}$$

Or, equivalently

$$F_Z(t) = \begin{cases} \frac{1}{2} - \int_0^{|t|} f_Z(z) dz & t \leq 0 \\ \frac{1}{2} + \int_0^{|t|} f_Z(z) dz & t > 0. \end{cases}$$

So we may write $F_Z(t) = \frac{1}{2} + \text{sign}(t) \int_0^{|t|} f_Z(z) dz$, where

$$\begin{aligned} & \int_0^{|t|} f_Z(z) dz \\ &= \frac{pd^{\frac{1}{p}}}{2\sigma\Psi^{(\Gamma)}(p, d\ln(a)^p, d^{-\frac{1}{p}})} \int_0^{|t|} e^{-d\ln(\frac{|z|}{\sigma} + a)^p} dz \\ &= \frac{pd^{\frac{1}{p}}}{2\Psi^{(\Gamma)}} \int_{\ln(a)}^{\ln(\frac{|t|}{\sigma} + a)} e^u e^{-du^p} du \\ &= \frac{pd^{\frac{1}{p}}}{2\Psi^{(\Gamma)}} \times \left(\int_0^{\ln(\frac{|t|}{\sigma} + a)} e^u e^{-du^p} du - \int_0^{\ln(a)} e^u e^{-du^p} du \right) \\ &= \frac{1}{2\Psi^{(\Gamma)}} \times \left(\Psi^{(\gamma)}\left(p, d\ln\left(\frac{|z|}{\sigma} + a\right)^p, d^{-\frac{1}{p}}\right) - \Psi^{(\gamma)}\left(p, d\ln(a)^p, d^{-\frac{1}{p}}\right) \right), \end{aligned}$$

where the second line follows from substitution and the fourth line follows from the definition of the lower-incomplete Fox-Wright function. \square