



Balancing Privacy and Utility in Correlated Data: A Study of Bayesian Differential Privacy

Martin Lange*

Karlsruhe Institute of Technology, KASTEL SRL
lange@martin-lange.eu

Javier Parra-Arnau

Universitat Politècnica de Catalunya
javier.parra@upc.edu

Patricia Guerra-Balboa*

Karlsruhe Institute of Technology, KASTEL SRL
patricia.balboa@kit.edu

Thorsten Strufe

Karlsruhe Institute of Technology, KASTEL SRL
thorsten.strufe@kit.edu

ABSTRACT

Privacy risks in differentially private (DP) systems increase significantly when data is correlated, as standard DP metrics often underestimate the resulting privacy leakage, leaving sensitive information vulnerable. Given the ubiquity of dependencies in real-world databases, this oversight poses a critical challenge for privacy protections. Bayesian differential privacy (BDP) extends DP to account for these correlations, yet current BDP mechanisms indicate a notable utility loss, limiting its adoption.

In this work, we address whether BDP can be realistically implemented in common data structures without sacrificing utility—a key factor for its applicability. By analyzing arbitrary and structured correlation models, including Gaussian multivariate distributions and Markov chains, we derive practical utility guarantees for BDP. Our contributions include theoretical links between DP and BDP and a novel methodology to adapt DP mechanisms to meet the requirements of BDP. Through evaluations on real-world databases, we demonstrate that our novel theorems enable the design of BDP mechanisms that maintain competitive utility, paving the way for practical privacy-preserving data practices in correlated settings.

PVLDB Reference Format:

Martin Lange, Patricia Guerra-Balboa, Javier Parra-Arnau, and Thorsten Strufe. Balancing Privacy and Utility in Correlated Data: A Study of Bayesian Differential Privacy. PVLDB, 18(11): 4090 – 4103, 2025.
doi:10.14778/3749646.3749679

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/lange-martin/privacy-utility-bdp>.

1 INTRODUCTION

Differential privacy (DP) [18] has become the leading framework for preserving privacy in data analysis, providing formal guarantees that protect individuals' sensitive information. However, its protection guarantees are limited to statistically independent data

records, i.e., DP mechanisms can leak private information when the underlying data is correlated. The limitations of DP for protecting correlated data have been theoretically exposed [26, 34, 38, 49] and empirically confirmed with attacks on real databases [25]. This is a significant issue, as correlations among data records are common in real-world databases, such as those induced by friendships in social networks [33] or genetic similarities among family members [1].

As a response to the limitations of DP in the presence of correlation, several instantiations of the Pufferfish framework—a general methodology to define privacy notions—have been proposed to specifically address this challenge [12, 23, 28, 32, 34]. Among them, *Bayesian Differential Privacy* (BDP) [53] stands out for its simplicity and generality: it provides a strict strengthening of DP, supports arbitrary correlation structures, and preserves the composability properties of DP—capabilities that are not generally achievable within the Pufferfish framework. BDP also underlies extensions such as prior DP [32] and correlated DP for location data [12].

While DP assumes that the adversary knows all records except the target, BDP considers arbitrary priors, including those where unknown records are correlated. It ensures bounded changes in output distributions even when the target record is part of a correlated subset. When data are independent, BDP and DP coincide. Under correlation, however, BDP quantifies worst-case leakage by integrating the mechanism's output with the data distribution via Bayes' rule, capturing adversarial advantages that DP overlooks. Hence, BDP mitigates correlation-driven reconstruction attacks that breach DP guarantees, as empirically shown in [8].

Although BDP provides a robust framework for assessing privacy leakage under data dependencies, its practical applicability remains uncertain. The few mechanisms that satisfy this notion [8, 53] are limited to specific correlation models, such as Gaussian Markov random fields—a subclass of multivariate Gaussian distributions forming a Markov random field where missing edges correspond to zeros in the inverse covariance matrix [44]—and binary-state Markov chains with a symmetric transition matrix. Given the scarcity of mechanisms and their applicability restrictions, it remains unclear whether BDP can serve as a usable privacy notion. Moreover, the only solution for Gaussian Markov fields reported highly conservative utility, since noise addition scales linearly with the number of records in the database and their only mitigation is to weaken BDP privacy by incorporating assumptions about the adversary [53].

In summary, DP privacy leakage estimation does not provide sufficient protection under data dependencies, and there is a need for improved utility with the robust BDP framework. Motivated by

*These authors contributed equally.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 18, No. 11 ISSN 2150-8097.
doi:10.14778/3749646.3749679

this issue, this paper examines BDP’s utility from both theoretical and practical perspectives, analyzing its limitations and proposing new strategies to reduce utility loss while maintaining BDP privacy guarantees. Particularly, we present theoretical bounds on the accuracy of BDP mechanisms and derive specific utility guarantees when certain correlation models are assumed. To formally analyze utility, we use the standard utility metric for DP mechanisms, (α, β) -accuracy [18, 51], due to its mathematical formalism and broad applicability. For the experimental results, we focus on two specific, albeit common, tasks: counting and sum queries [18].

Prior impossibility results [27, 28] show that strong utility under BDP without distributional assumptions is fundamentally limited. We extend this insight by proving that, without any assumption on the data correlation model, no BDP mechanism can simultaneously guarantee meaningful (α, β) -accuracy and valid privacy. Thus, the rest of our work examines whether targeting specific correlation models can improve utility.

Particularly, we analyze the impact of limiting the amount of correlated records, and we investigate the applicability of BDP to both discrete and continuous correlation models. For the discrete case, we analyze data following a Markov chain and, for continuous data, we analyze multivariate Gaussian correlation. We focus on these two particular correlation models following previous work in BDP [32, 53] and due to their relevance in many real-world applications such as medical [6], location [20], or activity data [16].

For each correlation model studied, we prove novel theorems that bound the BDP leakage of a DP mechanism. Notably, our BDP leakage bound for Gaussian multivariate models is tighter than that provided in [53], and our correlation model is broader. These privacy bounds provide a systematic way to build BDP mechanisms by adjusting the parameters of existing DP mechanisms. Using this approach, we propose novel BDP mechanisms based on Laplace noise. Furthermore, we calculate the accuracy of our BDP mechanisms showing the improved accuracy compared to scenarios where protection is required against any correlation.

Finally, we provide insight into how our theoretical results apply in practice to real-world data containing Gaussian and Markov correlations. This allows us to confirm that our results enhance the utility of BDP mechanisms in actual applications.

In summary, this work makes the following main contributions:

- We prove a bound on the BDP leakage of a DP mechanism with a fixed number of arbitrarily correlated records, showing it is tight. We call this the *general bound*.
- We derive a tighter BDP leakage bound for DP mechanisms under multivariate Gaussian correlations, improving on the general bound and prior work. This provides a systematic method for constructing more accurate BDP mechanisms tailored to Gaussian dependencies.
- We derive a BDP leakage bound for DP mechanisms under Markovian correlations, improving the general bound when transition probabilities are similar. This enables the design of mechanisms that are more accurate than prior approaches in Markov settings.

The paper is organized as follows: In Sections 2 and 3, we review relevant prior work and provide the necessary preliminaries. We

then present our analysis of arbitrary correlation limiting the number of correlated records in Section 4. In Section 5, we analyze the impact of Gaussian correlation on BDP and provide our improved bound in Theorem 5.8. In Section 6, we present analogous results for the Markov scenario. Finally, we discuss our empirical study in Section 7, demonstrating the practical relevance of our theoretical results, and conclude with a brief summary in Section 8.

We provide detailed proofs in the long version of this paper (arXiv:2506.21308) together with the code used for our experiments accessible in <https://github.com/linge-martin/privacy-utility-bdp>.

2 RELATED WORK

The challenge of designing privacy mechanisms that remain robust under arbitrary correlations has been a central concern in the development of privacy frameworks. Foundational work by Kifer and Machanavajjhala [26] introduced free-lunch Privacy, the first formalism to consider the impact of correlations on privacy guarantees. Their no-free-lunch theorem shows that, under arbitrary data distributions, achievable utility is fundamentally constrained. However, they express utility in terms of discriminants—an abstraction that is neither intuitively interpretable nor translatable into practical utility metrics. Kifer and Machanavajjhala [28] further raise this concern defining the general Pufferfish framework for privacy notions proving that any Pufferfish notion protecting against arbitrary correlations will face the same free-lunch utility challenge.

The existing strategy for obtaining Pufferfish privacy [46] mechanisms requires noise calibration based on the Wasserstein distance. It does not, however, provide a closed-form solution, but requires computing the Wasserstein distance between the conditional output distributions corresponding to all pairs of sensitive values. This is computationally intractable [41, 46] in the general case. While a closed-form mechanism is derived for specific Markov chain models, it relies on a weakened instantiation of Pufferfish that assumes limited adversarial background knowledge and, therefore, cannot be meaningfully compared to BDP.

The only concrete evidence of the potential applicability of pure BDP in practice has been provided in the context of Gaussian and Markov correlation models. In their foundational work, Yang et al. proposed adapting the Laplace mechanism to defend against correlated leakage in Gaussian Markov Random Fields. They also established preliminary theoretical connections between DP and BDP in this setting. Despite these important contributions, the proposed mechanisms face several limitations: (1) the approach is restricted to Gaussian Markov models, which greatly limits its practical scope. (2) Even within this narrow domain, privacy guarantees degrade linearly with the number of correlated records, resulting in excessive noise that renders the mechanism impractical. Although the authors suggest mitigating this by limiting the adversary’s knowledge, such a compromise weakens the privacy model and undermines the core guarantees of BDP. (3) The proposed mechanisms remain purely theoretical and have not been evaluated in real-world scenarios, leaving their practical effectiveness uncertain.

A more recent effort by Chakrabarti et al. [8] proposes an adaptation of the randomized response to BDP on binary Markov chains. However, this mechanism is extremely constrained: it only applies to lazy, binary, stationary Markov chains and does not provide any

Notation	Description
\mathcal{X}	Domain of a single record $x \in \mathcal{X}$.
$\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{Y}$	Randomized mechanism with input from domain \mathcal{X}^n and output in codomain \mathcal{Y} .
$\mathbf{X} = (X_1, \dots, X_n)$	Random vector representing the input of \mathcal{M} .
Y	Random variable representing output of \mathcal{M} .
$[n]$	Set $\{1, \dots, n\}$ for $n \in \mathbb{N}$.
$\mathbf{X}_K = (X_{i_1}, \dots, X_{i_k})$	Random vector formed by a subset $K = \{i_1, \dots, i_k\} \subseteq [n]$ of the random variables X_1, \dots, X_n .
$\mathbf{x}_K = (x_{i_1}, \dots, x_{i_k})$	Database with k records belonging to \mathcal{X}^k .

Table 1: Notation summary

general bounds relating DP and BDP leakage. Moreover, the only closed-form expression for mechanism parameters holds under the restrictive assumption of a symmetric transition matrix, limiting its usability even further.

In response to these limitations, several relaxed privacy notions have been proposed to strike a better balance between privacy and utility. Mutual Information Privacy (MI DP) [13] and its extension to Pufferfish [41], for example, can be viewed as a relaxation of Pufferfish, offering a framework where traditional mechanisms such as Laplace and Gaussian can be calibrated to account for correlation. These methods yield promising theoretical utility guarantees. However, MI guarantees are weaker; in particular, MI characterizes average-case privacy leakage rather than worst-case guarantees, and therefore cannot substitute the BDP framework when worst-case guarantees are desired.

In conclusion, while previous work highlights the limitations of DP protection and the need for BDP as a privacy standard, the challenge of providing utility with BDP protection remains unsolved, and the relationship between DP and BDP is not fully understood.

3 BACKGROUND

In this section, we present the fundamental definitions and notation (summarized in Table 1) necessary to understand this work.

3.1 Differential Privacy and Metric Privacy

The bounded formulation of DP [18] assumes that the database consists of a finite number n of rows, $D = (x_1, \dots, x_n) \in \mathcal{X}^n$, drawn from the joint distribution of the random vector $\mathbf{X} = (X_1, \dots, X_n)$, where each row represents data associated with an individual, sampled from a universe of records \mathcal{X} . We use $[n] := \{1, \dots, n\}$ to denote the set of indices. For a subset $K = \{i_1, \dots, i_k\} \subseteq [n]$, we define the subvector $\mathbf{X}_K \in \mathcal{X}^k$ as $\mathbf{X}_K := (X_{i_1}, \dots, X_{i_k})$. In particular, \mathbf{X}_{-i} denotes \mathbf{X}_K with $K = [n] \setminus \{i\}$. The attacker is assumed to know all records except for a target index $i \in [n]$, for which all possible values x_i and x'_i must be indistinguishable. Formally,

Definition 3.1 (Differential Privacy [18]). A randomized mechanism $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{Y}$ is called ϵ -differentially private, if for all measurable sets $S \subseteq \mathcal{Y}$ any target index $i \in [n]$, any target values $x_i, x'_i \in \mathcal{X}$, and any remaining values $\mathbf{x} \in \mathcal{X}^{n-1}$, we have

$$\Pr[Y \in S \mid \mathbf{X}_{-i} = \mathbf{x}, X_i = x_i] \leq e^\epsilon \Pr[Y \in S \mid \mathbf{X}_{-i} = \mathbf{x}, X_i = x'_i].$$

The output of \mathcal{M} is represented by the random variable Y , which depends on the input data. The DP leakage ϵ governs the privacy-utility trade-off: a smaller ϵ means that the output distributions for neighboring inputs are “closer together”, resulting in higher privacy with an opposing effect on utility (see Proposition 3.5).

We focus on a bounded DP due to its broad applicability and its close relation to BDP. However, other neighboring definitions, i.e., specifications of which information can change while ensuring that the output probabilities remain similar up to e^ϵ , exist [15]. For instance, in streaming data applications, it is common to use *event-level DP* [18]: While each stream belongs to an individual, two streams are neighbors if they differ in one single time step value. We will see an example of the application of this neighborhood in Section 7. The change of neighborhood allows to encode protection against different privacy threats [9, 15]. To obtain a general framework suitable to model a large variety of privacy problems, Chatzikokolakis et al. [9] introduce *metric privacy* as a generalization of DP that encapsulates the neighborhood notion and privacy leakage ϵ into a single parameter d , which determines the level of indistinguishability between databases:

Definition 3.2 (Metric Privacy [9]). Given $d : \mathcal{X}^{2n} \rightarrow \mathbb{R}$ a pseudo-metric, a randomized mechanism $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{Y}$ is called d -private if for all databases $D, D' \in \mathcal{X}^n$ and all measurable sets $S \subseteq \mathcal{Y}$ we have

$$\Pr_{\mathcal{M}}[Y \in S \mid \mathbf{X} = D] \leq e^{d(D, D')} \Pr_{\mathcal{M}}[Y \in S \mid \mathbf{X} = D'].$$

This definition makes it challenging for an adversary to distinguish between databases D and D' that are “close” according to the metric d . However, if the two databases are significantly different, the output distributions can differ more, making it easier for the adversary to distinguish them. Note that d -privacy is equivalent to DP when considering the Hamming distance scaled by ϵ .

One of the earliest and most common methods proven to satisfy ϵ -DP is the Laplace mechanism [18]:

Definition 3.3 (Laplace Mechanism [18]). Let $f : \mathcal{X}^n \rightarrow \mathbb{R}^k$ be a function and its *sensitivity* defined as

$$\Delta f := \sup_{d_H(D, D')=1} \|f(D) - f(D')\|_1.$$

Given that sensitivity $\Delta f < \infty$ and $\epsilon > 0$, the Laplace mechanism is defined for all $D \in \mathcal{X}^n$ as $\mathcal{M}_{\epsilon, f}(D) = f(D) + (Z_1, \dots, Z_k)$ where Z_i are i.i.d. random variables that follow the Laplace distribution centered at 0 and with scale $\frac{\Delta f}{\epsilon}$.

While $\mathcal{M}_{\epsilon, f}$ provides ϵ -DP, adding noise to the output of a function f undoubtedly has an impact on utility. A well-established metric for quantifying the utility of a private mechanism is the (α, β) -accuracy [5, 34]. It provides a measure of how well the mechanism approximates a true statistic or function while considering the inherent randomness introduced by the mechanism:

Definition 3.4 ((α, β) -Accuracy [5]). A randomized mechanism \mathcal{M} is (α, β) -accurate with respect to function f if for all databases $D \in \mathcal{X}^n$ we have

$$\Pr[|\mathcal{M}(D) - f(D)| \geq \alpha] \leq \beta.$$

A randomized mechanism \mathcal{M} is (α, β) -accurate if an error of magnitude α has a probability of at most β . Thus, the smaller α and/or β , the better the accuracy of mechanism \mathcal{M} . Here, α quantifies the error tolerance, and β the failure probability. More precisely, it refers to the utility guarantee that with probability at least $1 - \beta$, the mechanism's output is within an interval of radius α centered on the true value. For example, the Laplace mechanism verifies:

PROPOSITION 3.5 ([18]). *Let $\mathcal{M}_{\epsilon, f}$ be the Laplace mechanism. Let $\beta \in (0, 1]$ be a probability. Then $\mathcal{M}_{\epsilon, f}$ is (α, β) -accurate with respect to f with $\alpha = \ln(\beta^{-1}) \frac{\Delta f}{\epsilon}$.*

This accuracy result for the Laplace mechanism is tight [18].

3.2 Bayesian Differential Privacy

BDP [53] is an instantiation of the general Pufferfish framework that extends DP privacy guarantees to settings with correlated data. It assumes the adversary is uncertain between two possible records x_i, x'_i , analogously to DP. However, it eliminates the notion of neighboring databases in order to consider different possible adversaries with different background knowledge. Formally, the adversary (K, i) is targeting the record at position i and already knows the values of the sub vector \mathbf{x}_K on the database. Then, for each adversary, Bayesian leakage is defined as follows:

Definition 3.6 (Adversary-specific BDPL [53]). Given $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{Y}$ a randomized mechanism, \mathbf{X} the input random vector following the distribution π , the targeted record index $i \in [n]$, and the known record indices $K \subseteq [n] \setminus \{i\}$, the *adversary-specific Bayesian differential privacy leakage* is

$$\text{BDPL}_{(K, i)} = \sup_{x_i, x'_i, \mathbf{x}_K, S} \ln \frac{\Pr[Y \in S \mid \mathbf{X}_K = \mathbf{x}_K, X_i = x_i]}{\Pr[Y \in S \mid \mathbf{X}_K = \mathbf{x}_K, X_i = x'_i]},$$

where the supremum is taken over all the possible target values $x_i, x'_i \in \mathcal{X}$, all the possible known vector values $\mathbf{x}_K \in \mathcal{X}^K$ and all the measurable sets $S \subseteq \mathcal{Y}$.

When computing the adversary-specific BDPL, the correlation between the unknown and known records modifies the final leakage since given the unknown remaining indices U , we have

$$\Pr[Y \in S \mid \mathbf{x}_K, x_i] = \sum_{\mathbf{x}_U \in \mathcal{X}^U} \Pr[Y \in S \mid \mathbf{x}_K, x_i, \mathbf{x}_U] \Pr[\mathbf{x}_U \mid \mathbf{x}_K, x_i],$$

where $u = |U| = n - k - 1$. The sum must be substituted by an integral in the continuous case.

While the adversary-specific BDPL only accounts for a particular case, we aim to protect against any possible adversary. Therefore, to compute the worst-case leakage we take the supremum:

Definition 3.7 (Bayesian DP [53]). A mechanism \mathcal{M} satisfies ϵ -Bayesian differentially privacy if

$$\text{BDPL}(\mathcal{M}) = \sup_{K, i} \text{BDPL}_{(K, i)}(\mathcal{M}) \leq \epsilon,$$

where the supremum is taken over all the possible set of indices $i \in [n]$ and $K \subseteq [n] \setminus \{i\}$. $\text{BDPL}(\mathcal{M})$ is called *Bayesian differential privacy leakage*.

The BDPL has a similar role to the privacy leakage ϵ in DP: It measures the extent of a possible privacy violation by comparing the difference in the output probabilities of mechanism \mathcal{M} . A lower

BDPL corresponds to higher privacy because any adversary will be less likely to differentiate between any two target values $x_i, x'_i \in \mathcal{X}$. Particularly, if X_i, X_j are mutually independent for all $i \neq j \in [n]$ then ϵ -DP and ϵ -BDP are equivalent [53].

While we have results on the accuracy loss associated with using DP mechanisms [48], the impact of BDP protection on utility remains unclear. The following sections aim to address this question by analyzing various correlation scenarios.

4 LIMITED NUMBER OF CORRELATED VARIABLES

To protect against potential correlations without making distributional assumptions—which are often unclear or hard to estimate [47]—a mechanism must satisfy BDP with respect to all possible correlation distributions π , a condition we call protection under *arbitrary correlation*. However, Kifer and Machanavajjhala [28] showed that under this assumption, any Pufferfish notion—including BDP—collapses to free-lunch privacy [27, 53]. This corresponds to a metric privacy model where all dataset pairs are at distance ϵ , forcing all query outputs $f(D)$ and $f(D')$ to be ϵ -indistinguishable [17]—intuitively implying a complete loss of utility. To our knowledge, we are the first to formalize this limitation using the standard (α, β) -accuracy metric, offering a concrete, interpretable, and widely used measure of utility loss that enables clearer reasoning and meaningful comparison across mechanisms.

PROPOSITION 4.1. *Let $\mathcal{M} : \mathcal{X}^n \rightarrow \mathbb{R}$ be an ϵ -BDP mechanism protecting against arbitrary correlation. Let $0 \leq \beta < \frac{1}{e^\epsilon + 1}$ be a real number and let $f : \mathcal{X}^n \rightarrow \mathbb{R}$ be a deterministic function. If \mathcal{M} is (α, β) -accurate w.r.t. f , then*

$$\alpha > \frac{1}{2} \max_{D, D'} |f(D) - f(D')|.$$

PROOF. We proceed by *reductio ad absurdum*. We assume that \mathcal{M} fulfills an (α, β) -accuracy respect to f with $\alpha \leq \frac{1}{2} |f(D) - f(D')|$ and $\beta < \frac{1}{e^\epsilon + 1}$ and derive a contradiction for D' :

$$\begin{aligned} \Pr[|f(D') - \mathcal{M}(D')| \geq \alpha] &= \Pr[\mathcal{M}(D') \in \mathbb{R} \setminus (f(D') - \alpha, f(D') + \alpha)] \\ &\geq \Pr[\mathcal{M}(D') \in (f(D) - \alpha, f(D) + \alpha)] \\ &\stackrel{[28]}{\geq} e^{-\epsilon} \Pr[\mathcal{M}(D) \in (f(D) - \alpha, f(D) + \alpha)] \\ &= e^{-\epsilon} (1 - \Pr[\mathcal{M}(D) \in \mathbb{R} \setminus (f(D) - \alpha, f(D) + \alpha)]) \\ &= e^{-\epsilon} (1 - \Pr[|f(D) - \mathcal{M}(D)| \geq \alpha]) \stackrel{(*)}{\geq} e^{-\epsilon} (1 - \beta) = \frac{1}{e^\epsilon + 1} > \beta, \end{aligned}$$

where $(*)$ follows from the (α, β) -accuracy assumption. \square

Specifically, the result from Proposition 4.1 indicates that for theoretically relevant privacy levels $\epsilon \in (0, 4)$ [30], the only confidence interval where we can reliably estimate the actual value of our query function f , with standard confidence levels (e.g., between 90% and 99%), includes almost all possible query values. For instance, consider a free-lunch algorithm used to compute $f(D)$, where f counts the number of infections in a database of n individuals. If the algorithm outputs $\frac{n}{2}$, it suggests that half of the population is infected. However, with a 90% confidence interval, we cannot tell whether there is no infection at all, or whether the entire population is infected.

While designing accurate BDP mechanisms is infeasible under arbitrary correlation—potentially involving all records—it is often reasonable in practice to assume that only subsets of records are correlated. For instance, in the context of genomic data, an individual’s genome is strongly correlated with that of their relatives, but not with the entire population [1]. Hence, we assume that only m of n records in the database are correlated with each other, formally:

Definition 4.2. We say the random vector $\mathbf{X} = (X_1, \dots, X_n)$ has at most $m \leq n$ correlated random variables if there exist disjoint sets of indices C_1, \dots, C_r that make up $[n] = \bigcup_{l=1}^r C_l$ so that each set C_l has maximum cardinality $m \geq |C_l|$ for any $l \in [r]$, and for any $l \in [r]$, the random variables $\{X_j \mid j \in C_l\}$ are independent of the remaining random variables $\{X_j \mid j \in [n] \setminus C_l\}$.

This definition considers multiple groups of up to m correlated records as long as they do not “overlap”, i.e., the records in one group are independent of the records in the other groups. Otherwise, we do not make any further assumptions about the distribution of the data. This allows us to find acceptable utility guarantees in Corollary 4.5 as long as m is sufficiently small.

4.1 Relationship between DP and BDP

We begin by introducing and proving a general bound on the BDP leakage of an ε -DP mechanism. Specifically, we show that if an ε -DP mechanism operates on data drawn from a distribution involving at most m correlated random variables, then it satisfies $m\varepsilon$ -BDP. The practice of scaling the DP leakage by the number of correlated records to estimate worst-case leakage under correlation has been used in prior work [10, 34], but to our knowledge, this approach had not been formally shown to satisfy the BDP definition. We further prove that this bound is tight.

THEOREM 4.3 (THE GENERAL BOUND). *Let $\mathbf{X} = (X_1, \dots, X_n)$ be a random vector with at most $m \leq n$ correlated random variables that follows a distribution π . Then, any ε -DP mechanism with input data drawn from distribution π is $m\varepsilon$ -BDP.*

PROOF SKETCH. Consider any adversary (K, i) with $i \in [n]$, $K \subseteq [n] \setminus \{i\}$ and $k = |K|$. Since $\{C_j\}_{j \in [r]}$ is a partition of $[n]$, there exists an $l \in [r]$ so that we have target index $i \in C_l$. Thus, C_l contains the index i and all indices of random variables potentially correlated with X_i . Let set $\tilde{C} := C_l \setminus K$ be the indices of random variables correlated with X_i that are not already included in K . Then, we first show that the adversary-specific BDPL can be upper bounded as:

$$\text{BDPL}_{(K,i)} \leq \sup_{S, \mathbf{x}_K, \mathbf{x}_{\tilde{C}}, \mathbf{x}'_{\tilde{C}}} \ln \frac{\Pr[Y \in S \mid \mathbf{x}_K, \mathbf{x}_{\tilde{C}}]}{\Pr[Y \in S \mid \mathbf{x}_K, \mathbf{x}'_{\tilde{C}}]}.$$

Let the set $U = [n] \setminus (K \cup \tilde{C})$, with $u = |U|$, include all remaining indices. Since by hypotheses $|\tilde{C}| \leq |C_l| \leq m$, for any known values $\mathbf{x}_K \in \mathcal{X}^k$, the correlated values $\mathbf{x}_{\tilde{C}} \in \mathcal{X}^{|\tilde{C}|}$ and $\mathbf{x}'_{\tilde{C}} \in \mathcal{X}^{|\tilde{C}|}$ we have

$$\begin{aligned} \Pr_M[Y \in S \mid \mathbf{x}_K, \mathbf{x}_{\tilde{C}}] &= \int_{\mathcal{X}^u} \Pr_M[Y \in S \mid \mathbf{x}_K, \mathbf{x}_{\tilde{C}}, \mathbf{x}_U] p_{\mathbf{X}_U}(\mathbf{x}_U \mid \mathbf{x}_K, \mathbf{x}_{\tilde{C}}) d\mathbf{x}_U \\ &\leq \int_{\mathcal{X}^u} e^{m\varepsilon} \Pr_M[Y \in S \mid \mathbf{x}_K, \mathbf{x}'_{\tilde{C}}, \mathbf{x}_U] p_{\mathbf{X}_U}(\mathbf{x}_U \mid \mathbf{x}_K, \mathbf{x}_{\tilde{C}}) d\mathbf{x}_U \\ &= e^{m\varepsilon} \int_{\mathcal{X}^u} \Pr_M[Y \in S \mid \mathbf{x}_K, \mathbf{x}'_{\tilde{C}}, \mathbf{x}_U] p_{\mathbf{X}_U}(\mathbf{x}_U \mid \mathbf{x}_K) d\mathbf{x}_U \end{aligned}$$

	$X_1 = 0$	$X_1 = 1$	Total
$X_2 = 0$	$\frac{1}{r^2}$	$\frac{r-1}{r^2}$	$\frac{1}{r}$
$X_2 = 1$	$\frac{1}{r^3}$	$\frac{r^3-r^2-1}{r^3}$	$\frac{r-1}{r}$
Total	$\frac{1+r}{r^3}$	$\frac{r^3-r-1}{r^3}$	1

Table 2: Probability distribution of Example 4.4

$$\begin{aligned} &= e^{m\varepsilon} \int_{\mathcal{X}^u} \Pr[Y \in S \mid \mathbf{x}_K, \mathbf{x}'_{\tilde{C}}, \mathbf{x}_U] p_{\mathbf{X}_U}(\mathbf{x}_U \mid \mathbf{x}_K, \mathbf{x}'_{\tilde{C}}) d\mathbf{x}_U \\ &= e^{m\varepsilon} \Pr[Y \in S \mid \mathbf{x}_K, \mathbf{x}'_{\tilde{C}}]. \end{aligned}$$

Combining both inequalities we obtain the result. \square

This bound may seem overly pessimistic, seemingly assuming perfect positive correlation—the records are fully dependent, changing in lockstep: when one variable changes, the other changes in the same direction by exactly the same amount. This corresponds to the extreme case of linear dependence, where the Pearson correlation coefficient is $\rho = 1$, an edge case among all possible (including non-linear) correlation models. However, as we show in the following example, the bound remains tight even when this extreme case is excluded. Specifically, we provide a counterexample in which the bound holds even when ρ is arbitrarily small—i.e., the variables do not deterministically determine one another. This confirms both the tightness of our result and that the bound cannot be improved, even in the absence of perfect correlation.

Example 4.4. Let $r \in \mathbb{N}$. Table 2 shows a valid probability distribution π for $\mathbf{X} = (X_1, X_2)$ where the Pearson correlation coefficient satisfies:

$$\rho_{X_1, X_2} = \sqrt{\frac{r^4 - 2r^3 - r^2 + 2r + 1}{r^5 - 2r^3 - r^2 + r + 1}} \xrightarrow{r \rightarrow \infty} 0$$

Moreover, if \mathcal{M} is ε -DP, then there are two neighboring databases $D, D' \in \{0, 1\}^2$ for which the privacy loss reaches ε ; as is the case, for instance, with the Generalized Randomized Response mechanism [50]. Without loss of generality we assume they differ in the first coordinate, otherwise by inverting Table 2 we get the same result. Then, computing the BDPL we obtain for all $S \subseteq \{0, 1\}^2$:

$$\varepsilon^{\text{BDPL}} \geq \frac{\Pr[Y \in S \mid X_1=0]}{\Pr[Y \in S \mid X_1=1]} = \frac{e^{2\varepsilon} \frac{r}{r+1} + e^\varepsilon \frac{1}{r+1}}{e^\varepsilon \frac{r^2-r}{r^3-r-1} + \frac{r^3-r^2-1}{r^3-r-1}},$$

for all $r \in \mathbb{N}$, therefore, taking the limit when r tends to infinity we have $\text{BDPL} \geq 2\varepsilon$. According to the general bound $\text{BDPL} \leq 2\varepsilon$ hence we have $\text{BDPL} = 2\varepsilon$. Therefore, taking arbitrary large r , we have ρ arbitrary close to zero—making impossible perfect correlation— and BDPL arbitrary close to 2ε .

Example 4.4 proves that, without additional hypotheses, the general bound from Theorem 4.3 is tight, even if we limit the Pearson correlation coefficient ρ to be arbitrarily small.

4.2 Accuracy

Theorem 4.3 enables to use $(\frac{\varepsilon}{m})$ -DP mechanisms as ε -BDP mechanisms. However, reducing ε in a DP mechanism often has a negative impact on utility. In particular, we investigate the impact on the accuracy of the Laplace mechanism. As a consequence of our result Theorem 4.3 and Proposition 3.5 we obtain the following result:

COROLLARY 4.5. Let $\mathcal{M}_{\epsilon, f}$ be the Laplace ϵ -DP mechanism that approximates the query $f: \mathcal{X}^n \rightarrow \mathbb{R}$ with input described by the random vector $\mathbf{X} = (X_1, \dots, X_n)$ with at most $m \leq n$ correlated random variables that follows distribution π . Then, if $\mathcal{M}_{\epsilon, f}$ is (α, β) -accurate w.r.t. f , there exists an ϵ -BDP mechanism \mathcal{B} whose input is drawn from π and that is $(m\alpha, \beta)$ -accurate w.r.t. f .

This result shows that the error α of the Laplace mechanism increases proportionally with the number of correlated records when moving from ϵ -DP to ϵ -BDP, and while making no assumption about the distribution of the records. This may be acceptable when the number of correlated records m is small. For example, if $m = 2$, the error α doubles when transitioning from DP to BDP. If the DP mechanism's error is small, this increase may be acceptable. However, utility sharply decreases as m grows.

Since we have shown that our bound on BDPL is tight under the assumption of arbitrary correlation, the utility bound cannot be improved, even when the Pearson correlation coefficient is close to zero. This motivates the next two sections, where we investigate whether additional assumptions on the correlation model can lead to tighter bounds, enabling reduced noise and improved utility while still protecting against correlation attacks.

5 MULTIVARIATE GAUSSIAN CORRELATION

A wide variety of phenomena are effectively modeled using a Gaussian distribution [42]. For example, physiological measures such as height and weight are correlated among family members, and the joint distribution of height and weight in a large population is well fit by a bivariate Gaussian distribution [6]. Consequently, we explore the applicability of BDP to multivariate Gaussian data.

When we are dealing with a database of n records, and each record is drawn from a Gaussian distribution, we can model the joint distribution of all records as a multivariate Gaussian distribution. This model also captures linear correlation between records [45].

Definition 5.1 (Multivariate Gaussian Distribution [45]). Let $\mathbf{X} = (X_1, \dots, X_n)$ be a random vector, let vector $\mu \in \mathbb{R}^n$ be real and let matrix $\Sigma \in \mathbb{R}^{n \times n}$ be symmetric and positive definite. We say \mathbf{X} follows the *multivariate Gaussian distribution with mean μ and covariance Σ* if the probability density of \mathbf{X} for any point $\mathbf{x} \in \mathbb{R}^n$ is

$$p_{\mathbf{X}}(\mathbf{x}) = \frac{1}{\sqrt{(2\pi)^n |\Sigma|}} \exp\left(-\frac{1}{2}(\mathbf{x} - \mu)^T \Sigma^{-1}(\mathbf{x} - \mu)\right),$$

where $|\Sigma|$ is the determinant of Σ . We write $\mathbf{X} \sim \mathcal{N}(\mu, \Sigma)$.

We establish a relationship between DP and BDP for data drawn from a multivariate Gaussian distribution, based on the maximum Pearson correlation coefficient, which is calculated directly from the covariance matrix [45]. This provides a new, tighter upper bound for the BDPL that improves upon the specific Gaussian bound given in [53] and upon the general bound $n\epsilon$ for any correlation model.

However, our bound applies only to a specific class of mechanisms: those that satisfy both DP and metric privacy under the ℓ_1 metric. We show in Section 5.2 that the clipped Laplace mechanism meets these criteria and develop a practical application in Section 7.1. To establish this result, we first connect metric privacy with an analogous form of BDP, termed Bayesian metric privacy, which we define below.

5.1 Relationship between Metric Privacy and Bayesian Metric Privacy

Unbounded continuous data domains, such as \mathbb{R}^n , usually produce challenges on DP application due to infinite sensitivities [2]. In the context of BDP, Yang, Sato, and Nakagawa [53] defined a relaxation to work in those domains: If the data domain is equivalent to the real numbers (i.e., $\mathcal{X}^n = \mathbb{R}^n$), they defined a modified leakage, $\text{BDPL}(\mathcal{M}; M)$, where they only take into account the leakage between points with a distance smaller than $M \in \mathbb{R}$, i.e.,

$$\sup_{|x_i - x'_i| \leq M, \mathbf{x}_K, S} \ln \frac{\Pr[Y \in S \mid \mathbf{X}_K = \mathbf{x}_K, X_i = x_i]}{\Pr[Y \in S \mid \mathbf{X}_K = \mathbf{x}_K, X_i = x'_i]}.$$

Applying this BDP relaxation leaves indistinguishability between records at distances greater than M entirely uncontrolled. While this may increase applicability, it reduces privacy and limits insights into the impact of correlation.

However, metric privacy provides a solution to quantify privacy leakage as the distance $d(D, D')$ for each pair of databases D, D' when the maximum privacy leakage cannot be bounded [9]. Therefore, we define Bayesian metric privacy as equivalent to metric privacy where the indistinguishability between two records x, x' depends on the distance $d(x, x')$ between them. Note that the change from databases to records is necessary because BDP does not apply to neighboring databases, but to target records, as we describe in Section 3.2. In this way, we can work with \mathbb{R}^n as the data domain without losing information about the privacy leakage.

Definition 5.2 (Target Dependent Leakage). Given a randomized mechanism $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{Y}$, \mathbf{X} the input random vector following the distribution π , the targeted record index $i \in [n]$, and the known record indices $K \subseteq [n] \setminus \{i\}$, the *adversary-specific target dependent BDPL* of \mathcal{M} w.r.t. adversary (K, i) for any target values $x, x' \in \mathcal{X}$ is

$$\text{BDPL}_{(K, i)}(x, x') = \sup_{\mathbf{x}_K, s} \ln \frac{p_Y(s \mid \mathbf{X}_K = \mathbf{x}_K, X_i = x)}{p_Y(s \mid \mathbf{X}_K = \mathbf{x}_K, X_i = x')}.$$

Given that we understand the leakage for each pair of data records we can simply define Bayesian metric privacy analogously to the original metric privacy notion:

Definition 5.3 (Bayesian Metric Privacy). Let d be a (pseudo)metric on \mathcal{X}^2 . A mechanism \mathcal{M} is *Bayesian d -private* if for all $x, x' \in \mathcal{X}$,

$$\text{BDPL}(x, x') = \sup_{i, K} \text{BDPL}_{(K, i)}(x, x') \leq d(x, x'),$$

where the supremum is taken over all the possible set of indices $i \in [n]$ and $K \subseteq [n] \setminus \{i\}$.

The only difference between BDP and Bayesian d -privacy is that Bayesian d -privacy does not take the supremum over x, x' . Moreover, both notions are equivalent when the distance metric is defined as $d(x, x') = \epsilon$ for $x \neq x'$ and $d(x, x') = 0$ otherwise.

Now we can prove the relation between a d -private and a Bayesian d -private mechanism when the data distribution is a multivariate Gaussian. Particularly, we focus on the ℓ_1 distance due to its direct application to the Gaussian case. We formalize the conditions needed to obtain our bound:

Definition 5.4. For $\rho \in [0, 1]$ and $n \in \mathbb{N}$, we call the matrix $\Sigma_\rho \in \mathbb{R}^{n \times n}$ a *limited covariance matrix* if

- the matrix Σ_ρ is symmetric and positive definite,
- the diagonal of Σ_ρ is constant, i.e., there is a variance $\sigma^2 > 0$ so that $\Sigma_{\rho,ii} = \sigma^2$ for all $i \in [n]$ and,
- any pairwise correlation is limited by ρ . I.e., for all $i \neq j$ we have $|\Sigma_{\rho,ij}| \leq \rho\sigma^2$.

The first condition is required to be a valid covariance matrix for a Gaussian distribution (see Definition 4.2). The second condition ensures that no records have a deviating variance, i.e., all records are drawn from the same one-dimensional distribution. The final condition imposes that the maximum Pearson correlation coefficient between any two random variables X_i and X_j is bounded by ρ . If we limit ρ to be small enough, we get the following bound:

THEOREM 5.5. *Let \mathcal{M} with data domain \mathbb{R}^n be an $(\varepsilon\ell_1)$ -private mechanism where $\varepsilon > 0$ with input data drawn from a multivariate Gaussian distribution $\mathcal{N}(\mu, \Sigma_\rho)$ with mean $\mu \in \mathbb{R}^n$, limited covariance matrix $\Sigma_\rho \in \mathbb{R}^{n \times n}$ (Def. 5.4) and a maximum of $m \leq n$ correlated variables such that $\rho(m-2) < 1$ is the maximum Pearson coefficient. Then, for any $x, x' \in \mathbb{R}$ we have*

$$\text{BDPL}(x, x') \leq \left(\frac{m^2}{4(\frac{1}{\rho} - m + 2)} + 1 \right) |x' - x|.$$

PROOF SKETCH. If from the set of unknown records V , none is correlated with the target, $U = \emptyset$, then the BDPL is the metric privacy leakage and the inequality trivially holds. Otherwise, for any adversary H , without loss of generality we reorder the subset of known records correlated with the target as $K = \{m-k, \dots, m-1\}$ and $i = m$, denote $T = K \cup \{m\}$ and we show that if the principal submatrix Σ_T is invertible, then

$$\text{BDPL}_{(H,m)}(x_m, x'_m) \leq \varepsilon |x_m - x'_m| \left(\|\Sigma_{U;T} \Sigma_T^{-1} \mathbf{e}_{k+1}\|_1 + 1 \right) \quad (1)$$

where $\mathbf{e}_{k+1} \equiv (0, \dots, 0, 1)^\top \in \mathbb{R}^{k+1}$ and the notation of the Gaussian distribution $\mathcal{N}(\mu, \Sigma)$ is reordered as

$$\Sigma_\rho = \begin{pmatrix} \Sigma_U & \Sigma_{U;T} & \mathbf{0} \\ \Sigma_{U;T}^\top & \Sigma_T & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \Sigma_S \end{pmatrix}.$$

If $m = 2$, given that $U \neq \emptyset$, then $k = 0$ and $\Sigma_{U;T} \Sigma_T^{-1} = \frac{\rho\sigma_1\sigma_2}{\sigma_2^2} \leq \rho$. Using Eq. 1 we obtain: $\text{BDPL}_{(H,i)}(x_i, x'_i) \leq (\rho + 1)\varepsilon |x'_i - x_i|$.

If $m > 2$, we prove that Σ_T is invertible applying the Gershgorin circle Theorem [21], to prove that the eigenvalues of Σ_T obey the following inequality

$$\lambda \stackrel{(**)}{\geq} (1 - k\rho)\sigma^2 > (1 - (m-2)\frac{1}{m-2})\sigma^2 = 0. \quad (2)$$

This bound is positive since the number of known correlated records k must be $m-2$ or smaller because the target record is correlated with at most m others, $U \neq \emptyset$, and $\rho(m-1) < 1$. Hence we can apply Eq. 1 obtaining:

$$\begin{aligned} \text{BDPL}_{(H,i)}(x_i, x'_i) &\leq (\|\Sigma_{U;T} \Sigma_T^{-1} \mathbf{e}_{k+1}\|_1 + 1) |x_i - x'_i| \varepsilon \\ &\stackrel{(*)}{\leq} \left(\sum_{j=1}^u \left| \sum_{l=1}^{k+1} \frac{\rho\sigma^2}{(1-k\rho)\sigma^2} \right| + 1 \right) |x_i - x'_i| \varepsilon \\ &= \left(\frac{u(k+1)\rho}{1-k\rho} + 1 \right) |x_i - x'_i| \varepsilon \leq \left(\frac{m^2}{4(\frac{1}{\rho} - m + 2)} + 1 \right) |x_i - x'_i| \varepsilon. \end{aligned}$$

Where $(*)$ holds since the entries of $\Sigma_{U;T}$ are bounded by $\rho\sigma^2$ and the one of Σ_T^{-1} by $\frac{1}{\lambda_-} \leq \frac{1}{(1-k\rho)\sigma^2}$ as derived in Eq. 2. \square

Theorem 5.5 provides a concrete formula for the increase in privacy leakage due to linear correlation relative to independent data. Higher Pearson coefficients lead to greater leakage. Additionally, we can extend this result to derive a relation between DP and BDP.

5.2 Relationship between DP and BDP

Observe that any d -private mechanism is an ε -DP mechanism with $\varepsilon = \sup_{D \sim D'} d(D, D')$. Moreover, any Bayesian d -private mechanism is an ε -BDP mechanism with $\varepsilon = \sup_{x, x'} d(x, x')$. By leveraging these relationships between privacy notions we can establish a connection between DP and BDP. However, since this supremum may be unbounded, it can lead to undesirable privacy guarantees. To manage this relationship effectively we apply clipping techniques, resulting in Theorem 5.8, which enables the construction of BDP mechanisms from DP mechanisms. Formally, clipping is defined as:

Definition 5.6. For any interval $I = [a, b] \subset \mathbb{R}$, we define the *clipping function* $c_I : \mathbb{R}^n \rightarrow \mathbb{R}^n$, which, for all $D \in \mathbb{R}^n$ and all $i \in [n]$, outputs

$$c_I(D)_i = \max(a, \min(b, D_i)).$$

Let $\mathcal{M} : \mathbb{R}^n \rightarrow \mathbb{R}$ be a mechanism. We define its *clipped version* $\mathcal{M}_I : \mathbb{R}^n \rightarrow \mathbb{R}$ as $\mathcal{M}_I = \mathcal{M} \circ c_I$.

Due to the data domain reduction, we can bound the DP leakage of $\varepsilon\ell_1$ -private mechanisms.

LEMMA 5.7. *If $\mathcal{M} : \mathbb{R}^n \rightarrow \mathbb{R}$ is $\varepsilon\ell_1$ -private, then its clipped version \mathcal{M}_I is $\varepsilon\ell_1$ -private and $(M\varepsilon)$ -DP with $M = |b - a|$.*

With Lemma 5.7 and Theorem 5.5, we can directly show that this class of DP mechanisms has a limited BDPL.

THEOREM 5.8 (THE GAUSSIAN BOUND). *Let \mathcal{M}_I with data domain \mathbb{R}^n be the clipped version of an $\varepsilon\ell_1$ -private mechanism \mathcal{M} where $\varepsilon > 0$ and input data drawn from a multivariate Gaussian distribution $\mathcal{N}(\mu, \Sigma_\rho)$ with mean $\mu \in \mathbb{R}^n$, maximum number of correlated records $m \leq n$ and limited covariance matrix $\Sigma_\rho \in \mathbb{R}^{n \times n}$ (Def. 5.4) such that $\rho(m-2) < 1$ is the maximum correlation coefficient. Then, the clipped mechanism \mathcal{M}_I is*

$$\left(\frac{m^2}{4(\frac{1}{\rho} - m + 2)} + 1 \right) M\varepsilon\text{-BDP}.$$

where M is the diameter of the interval I .

The proof follows directly from Theorem 5.5 taking the supremum over all data records, subject to the clipping constraint.

Theorem 5.8 allows us to systematically build a BDP mechanism by recalibrating the noise of a DP mechanism when $\rho(m-2) < 1$. For instance, given the clipped Laplace mechanism \mathcal{M}_I that adds noise to a data point $x \in \mathbb{R}$ following $\text{Lap}(\frac{M}{\tau})$, where

$$\tau = \varepsilon \frac{4(\frac{1}{\rho} - m + 2)}{m^2 + 4(\frac{1}{\rho} - m + 2)}, \quad (3)$$

we obtain an ε -BDP mechanism. Moreover,

$$\frac{m^2}{4(\frac{1}{\rho} - m + 2)} + 1 \leq m \text{ if and only if } \rho \leq \frac{m-1}{\frac{5}{4}m^2 - 3m + 2}. \quad (4)$$

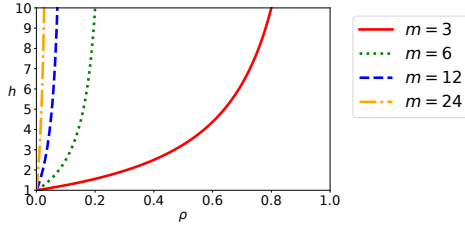


Figure 1: Relative accuracy of an ε -BDP mechanism to an ε -DP mechanism for a Multivariate Gaussian distribution.

Hence, the Gaussian bound improves on the general bound if ρ is on the order of $\rho \approx \frac{1}{m}$. The higher the number of correlated records m , the better the relative improvement of the Gaussian-specific bound compared to the general bound for small correlation.

Importantly, Yang et al. [53] establish a bound for Gaussian Markov random fields. They establish that a clipped $\mathcal{M}_{\varepsilon,f}$ satisfies $(nM\varepsilon)$ -BDP, which coincides with the general bound when all records are correlated. Theorem 5.8 applies to this particular case since a Gaussian Markov random field is an example of Gaussian Multivariate distribution. Moreover, our bound improves over theirs in the same cases it improves over the general bound.

5.3 Accuracy

When the Pearson correlation is bounded as specified in Equation (4), it is guaranteed that a larger ε' than $\frac{\varepsilon}{m}$ is sufficient to guarantee ε -BDP via an ε' -DP mechanism. Since a larger privacy budget generally correlates with improved utility, we can therefore anticipate enhanced utility results. In particular, we express the accuracy improvement of the Laplace mechanism calibrated to protect data drawn from a multivariate Gaussian distribution. As a consequence of our Theorem 5.8 and Proposition 3.5 from [18] we obtain the following result:

COROLLARY 5.9. *Let $\mathcal{M}_{\varepsilon,f_1}$ be the clipped Laplace ε -DP mechanism that approximates the query f_1 as defined in 5.6 with input data drawn from a multivariate Gaussian distribution $\mathcal{N}(\mu, \Sigma_\rho)$ with mean $\mu \in \mathbb{R}^n$ and limited covariance $\Sigma_\rho \in \mathbb{R}^{n \times n}$ with a maximum number of correlated variables $m \leq n$ such that $\rho(m-2) < 1$. Then, if the Laplace mechanism $\mathcal{M}_{\varepsilon,f_1}$ is (α, β) -accurate w.r.t. f_1 , there exists an ε -BDP mechanism \mathcal{B} whose input is drawn from π and that is $(h\alpha, \beta)$ -accurate w.r.t. f_1 with*

$$h = \frac{m^2}{4(\frac{1}{\rho} - m + 2)} + 1.$$

The statement of Corollary 5.9 is visualized in Figure 1. This figure shows that in order to provide similar utility to DP, ρ must be small. The larger the number of correlated records m , the smaller ρ has to be to provide similar utility. The results in this section enable the protection of weakly correlated data drawn from a multivariate Gaussian distribution. Furthermore, a comparison of the accuracy achieved by our method versus the state-of-the-art bound from [53] and the general BDP bound is presented in Figure 4, demonstrating a consistent improvement enabled by our approach.

6 MARKOV CHAIN CORRELATION MODEL

In streaming processes or time series data, states at successive time steps are often correlated, meaning that the state at a given time step depends on the state at the previous one. For example, a user's location at time step t is correlated with their location at $t-1$. This dependency pattern is commonly modeled using Markov chains [4].

Consequently, in this section we investigate the impact of correlations following a Markov model on the privacy leakage and utility of BDP mechanisms. Particularly, we prove Theorem 6.2, a new bound on the BDPL of any ε -DP mechanism when data is correlated corresponding to a Markov chain. Additionally, we use our results to elaborate on the utility gain compared to protecting against arbitrary correlation.

For the remainder of this work, we adopt the definition of a Markov chain from [4], which specifically refers to finite, time-homogeneous Markov chains, i.e., those with finite state spaces and time-invariant transition probabilities. Formally,

Definition 6.1 (Markov Chain [4]). Let \mathcal{S} be a finite set of possible states of size $s \in \mathbb{N}$ and let $\mathbf{X} = (X_1, \dots, X_n)$ be a random vector. We say \mathbf{X} is a *Markov chain* with transition matrix $P \in \mathbb{R}^{s \times s}$ and initial distribution $w \in \mathbb{R}^s$ if all of the following holds.

- (1) For all states $x, y \in \mathcal{S}$ and all indices $i \in [n-1]$ we have $\Pr[X_{i+1} = x | X_i = y] = P_{y,x}$.
- (2) For all states $x \in \mathcal{S}$ we have $\Pr[X_1 = x] = w_x$.
- (3) The Markov property: For all indices $i \in [n-1]$ and for all states $x_1, \dots, x_i, x_{i+1} \in \mathcal{S}$ we have

$$\begin{aligned} \Pr[X_{i+1} = x_{i+1} | X_1 = x_1, \dots, X_i = x_i] \\ = \Pr[X_{i+1} = x_{i+1} | X_i = x_i]. \end{aligned}$$

6.1 Relationship between DP and BDP

In this subsection, we show that it is possible to obtain a bound on the BDPL of any DP mechanism based on the maximum ratio between the largest and smallest transition probabilities in the Markov chain. The intuition is that if all transition probabilities are similar, changing the random variable X_i from state x_i to state x'_i will have minimal impact on the subsequent time steps of the Markov chain. However, if the transition probabilities differ significantly, this change could have a large effect over many time steps. Formally, we bound the BDPL of a DP mechanism on data that follows a Markov chain as follows:

THEOREM 6.2 (THE MARKOV CHAIN BOUND). *Let $s \in \mathbb{N}$ be the number of states. Let $\mathcal{M} : \mathcal{S}^n \rightarrow \mathcal{Y}$ be an ε -DP mechanism. Let the databases follow a Markov chain with transition matrix $P \in \mathbb{R}^{s \times s}$ and initial distribution $w \in \mathbb{R}^s$ with the following properties:*

- (H1) *For all $x, y \in \mathcal{S}$ we have $P_{x,y} > 0$ and,*
- (H2) *$wP = w$.*

Then, \mathcal{M} is an $(\varepsilon + 4 \ln \gamma)$ -BDP mechanism where

$$\gamma := \frac{\max_{x,y \in \mathcal{S}} P_{xy}}{\min_{x,y \in \mathcal{S}} P_{xy}}.$$

PROOF SKETCH. If there are no unknown indices, $\text{BDPL}_{(K,i)}$ is the DP leakage [53] and the inequality is trivially satisfied. Otherwise, combining Bayes' rule, Markov property, (H1) and (H2) we

prove that:

$$\begin{aligned} \frac{\Pr[\mathbf{x}_U | \mathbf{x}_K, x_i]}{\Pr[\mathbf{x}_U | \mathbf{x}_K, x'_i]} &= \frac{\Pr[x_i | \mathbf{x}_K, \mathbf{x}_U] \Pr[x'_i | \mathbf{x}_K]}{\Pr[x'_i | \mathbf{x}_K, \mathbf{x}_U] \Pr[x_i | \mathbf{x}_K]} \\ &= \frac{\Pr[x_i | \mathbf{x}_{-i}] \Pr[x'_i | \mathbf{x}_K]}{\Pr[x'_i | \mathbf{x}_{-i}] \Pr[x_i | \mathbf{x}_K]} \leq \gamma^2 \gamma^2 = \gamma^4. \end{aligned}$$

Note that if $K = \emptyset$ the previous expression gets simplified to

$$\frac{\Pr[\mathbf{x}_U | x_i]}{\Pr[\mathbf{x}_U | x'_i]} = \frac{\Pr[x_i | \mathbf{x}_U] \Pr[x'_i]}{\Pr[x'_i | \mathbf{x}_U] \Pr[x_i]} \leq \gamma^3 \leq \gamma^4.$$

Therefore,

$$\begin{aligned} \Pr_M[Y \in S | \mathbf{x}_K, x_i] &= \sum_{\mathbf{x}_U \in \mathcal{S}^u} \Pr_M[Y \in S | \mathbf{x}_K, x_i, \mathbf{x}_U] \Pr_\pi[\mathbf{x}_U | \mathbf{x}_K, x_i] \\ &= \sum_{\mathbf{x}_U \in \mathcal{S}^u} \Pr[Y \in S | \mathbf{x}_K, x_i, \mathbf{x}_U] \Pr[\mathbf{x}_U | \mathbf{x}_K, x'_i] \frac{\Pr[\mathbf{x}_U | \mathbf{x}_K, x_i]}{\Pr[\mathbf{x}_U | \mathbf{x}_K, x'_i]} \\ &\leq \gamma^4 e^\epsilon \Pr[Y \in S | \mathbf{x}_K, x'_i]. \quad \square \end{aligned}$$

(H1) states that all entries in the transition matrix are strictly positive, while (H2) requires that the initial distribution is a *stationary distribution*, meaning the distribution over states w_t (without considering the previous one) remains constant at each time—a common modeling assumption in various data mining tasks such as weather forecasting [52] or electricity consumption [3]. Notably, condition (H1) implies that the chain is both irreducible and aperiodic, which in turn guarantees the existence of a unique stationary distribution w [31] satisfying (H2). Moreover, for any initial distribution w' , the distribution at time t converges geometrically fast to w as t increases [31]. Hence, even when the initial distribution is not stationary, it asymptotically approaches the stationary distribution, satisfying (H2) after discarding sufficient initial events.

While prior work provides a mechanism for BDP protection of lazy binary Markov chains with a symmetric transition matrix [8], we present the first direct and general relationship between DP and BDP leakage for arbitrary Markov chains, including non-binary ones. When comparing this novel bound with the general one, for any $\epsilon > 0$ and maximum transition probability ratio $\gamma \geq 1$, we have

$$\epsilon + 4 \ln \gamma < n\epsilon \text{ if and only if } \gamma < \exp\left(\frac{n-1}{4}\epsilon\right). \quad (5)$$

Therefore, the Markov chain bound outperforms the general bound in most cases. For instance, with an ϵ -DP mechanism where $\epsilon = 0.5$ and a database size of $n = 80$, it remains tighter even when the largest transition probability is 10,000 times the smallest. For the same $\epsilon = 0.5$, the Markov bound only becomes looser than the general one when the number of correlated records is small, e.g., $n = 20$, and the transition probability ratio γ is as high as 100, which still represents a significant disparity.

Moreover, Theorem 6.2 enables the systematic design of BDP mechanisms by adjusting the noise of an existing DP mechanism. Noise calibration depends only on the maximum ratio between the Markov transition probabilities of the model, γ , and the adjusted mechanism must be calibrated to $\epsilon' = \epsilon - 4 \ln(\gamma)$. Note that the best achievable BDPL using Theorem 6.2 is $\epsilon = 4 \ln(\gamma)$, since $\epsilon' \geq 0$. Consequently, the minimum achievable ϵ is fundamentally constrained by the structure of the underlying Markov model—specifically by the maximum transition ratio γ . We illustrate how the transition matrix changes the minimum ϵ in theoretical settings in Figure 2, and in real-world data in Section 7.

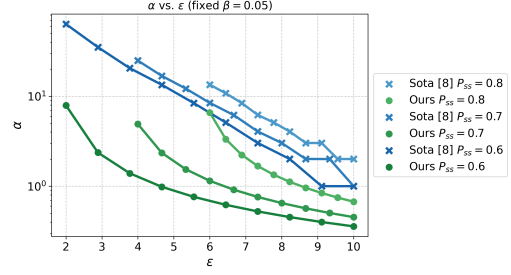


Figure 2: Accuracy of our mechanism vs. the one proposed in [8] for $n = 700$ and various self-transition probabilities P_{ss} .

6.2 Accuracy

The Markov chain bound enables us to derive improved utility guarantees for the Laplace mechanism when γ is sufficiently small.

COROLLARY 6.3. *Let $\mathcal{M}_{\epsilon, f}$ be the ϵ -Laplace mechanism that approximates the query $f : \mathcal{S}^n \rightarrow \mathbb{R}$ and inputs a database drawn from a Markov chain satisfying (H1) and (H2). If $\mathcal{M}_{\epsilon, f}$ is (α, β) -accurate w.r.t. f and $\epsilon \geq 4 \ln(\gamma)$ then, there exists an ϵ -BDP mechanism \mathcal{B} that is $(h\alpha, \beta)$ -accurate w.r.t. f with*

$$h = \frac{\epsilon}{\epsilon - 4 \ln(\gamma)}.$$

The statement of Corollary 6.3 is visualized in Figure 3. This figure shows that in order to provide similar utility guarantees to DP, either the BDPL bound ϵ has to be larger than 5, or the ratio γ between different transition probabilities must be smaller than 3.

The only previous BDP mechanism for Markov chains [8], focuses on lazy binary Markov models with a symmetric transition matrix, i.e., the probability of staying in the same state P_{ss} —self-transition probability—is constant for $s \in \{0, 1\}$, and $P_{ss} > 0.5$. In this regime, $\gamma = \frac{P_{ss}}{1-P_{ss}}$. Restricting ourselves to this setting, we find that our mechanism achieves superior (α, β) -accuracy, as shown in Figure 2. The detailed formal analysis can be found in the long version of this paper. It is important to note that while their mechanism supports arbitrary BDPL, ours applies only for $\epsilon \geq 4 \ln(\gamma)$. However, our approach generalizes to arbitrary Markov chains, whereas theirs is limited to lazy, symmetric binary models. In the intersection of both applicability domains, our use of Laplace-based recalibration yields improved utility.

In conclusion, the Markov-specific bound improves upon the general bound under certain conditions and enables improved utility (Figure 2) compared to prior work [8]. Its advantage is most notable when the number of correlated records is large, as it remains independent of dataset size—unlike the general bound, which grows linearly. However, this comes at the cost of a minimum privacy level determined by the data distribution, a limitation absent in the general bound and [8].

7 UTILITY EXPERIMENTS

Theoretical bounds on privacy and utility do not always translate directly to practical implementations. For instance, while it may be theoretically feasible to achieve a given (α, β) -accuracy, designing or implementing a mechanism that attains this in practice can

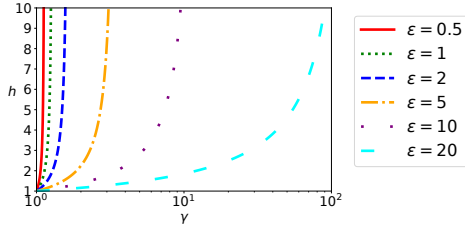


Figure 3: Relative accuracy h of an ε -BDP to an ε -DP mechanism for Markov chain data respect to γ .

be challenging. In this section, we use our theoretical results to construct a BDP mechanism and empirically evaluate its utility on real-world databases that follow either multivariate Gaussian correlations or Markov chains. Our objective is to demonstrate that the utility gains predicted under specific correlation structures, rather than arbitrary ones, are indeed achievable in practice as well as measure the improvement over previous approaches.

We calibrated the Laplace mechanism using Theorem 5.8 and Theorem 6.2 to derive BDP mechanisms. We then ran these BDP mechanisms on the selected databases and compared the utility results with those of BDP mechanisms designed to protect against arbitrary correlation, in order to assess the improvements offered by the correlation-specific approach. Moreover, we also plot, when applicable, the accuracy results of the state-of-the-art solutions for Gaussian BDP [53]. Unfortunately, none of the evaluated datasets meet the strict assumptions needed to apply the only prior mechanism for Markov models [8]. Finally, we plot the utility of the classical DP Laplace mechanism as a baseline, representing the best-case utility achievable ignoring correlation.

7.1 Databases

We use four real-world databases, two for each correlation model. Additionally, we use a synthetic dataset to test scalability for Gaussian correlations. The selection criteria are public availability, quality of the databases, and the fulfillment of the theoretical assumptions, namely, following the correlation model and fulfilling the extra hypotheses of the corresponding theorem in each case, regarding the Pearson correlation coefficient and the transition matrix.

7.1.1 Multivariate Gaussian: We use two datasets that align well with our modeling framework: the Galton Height Data [19], a historical dataset originally compiled to study the correlation between parents’ and children’s heights, and the FamilyIQ dataset [22], which includes IQ scores of gifted children and their parents.

The Galton Height Data—considered a classical example of linear correlation modeling, where regression and correlation are interpreted within the framework of a multivariate Gaussian distribution [35]—is especially well known in statistical analysis for introducing the very concept of regression [6]. In contrast, several studies provide evidence that IQ scores in the general population are standardized to follow a multivariate Gaussian distribution, where non-zero correlations are observed only among close relatives [43]. These properties make both datasets well-suited for evaluating the practical transferability of our Gaussian-based bounds. Additionally,

Database	n	m	Parameters	Sensitivity
Galton	897	3	$\rho = 0.275$	$\Delta q = 254 \text{ cm}$
FamilyIQ	868	2	$\rho = 0.4483$	$\Delta q = 120$
SyntheticIQ	20000	2	$\rho = 0.45$	$\Delta q = 120$
Activity	17568	n	$\gamma = 7.54$	$\Delta q = 1$
Electricity	731	n	70 kWh, $\gamma = 3.29$	$\Delta q = 1$
			80 kWh, $\gamma = 4.49$	
			90 kWh, $\gamma = 8.43$	

Table 3: Data description. m is the max number of correlated records and n the total amount.

we generate the SyntheticIQ dataset to test the scalability of our approach. Following the findings among several populations summarized in [43], we generate data following a Gaussian distribution with $\mu = 100$, $\sigma^2 = 15$ and $\rho = 0.45$ for parent-child.

To ensure bounded sensitivities, all records are clipped to the range of 1cm to 254cm (0 to 100 inches) for Galton, and from 40 to 160 for IQ datasets as summarized in Table 3.

All explored datasets fulfill the conditions of our Theorem 5.8: Galton Pearson correlation coefficient of $\rho = 0.275$, satisfies the condition $\rho = 0.275 < 1 = \frac{1}{m-2}$, hence our bounded-correlation assumptions hold. For $m = 2$, the condition trivially holds for all ρ values, so in particular for FamilyIQ and SyntheticIQ.

7.1.2 Markov Model: We study two use cases—human activity and electricity consumption—well-suited for Markov modeling. Human activity representations such as “inactive” versus “active” are modeled by Markov chains [24]. Similarly, electricity usage patterns, particularly transitions between high and low consumption periods, have been effectively modeled using Markov processes [3, 14, 39]. We select a representative database for each domain to evaluate our framework. For human activity, we use Activity Data [37], which contains the time series of step counts recorded every 5 minutes from a personal activity monitoring device worn by a single individual during October and November 2012. This allows us to extract the “active” state if any steps are recorded and the “inactive” when the user does not move. Besides, to assess the data size impact, we split Activity data into 61 unique subdatabases, each corresponding to the activity states of a single day and report the results in the long version of this paper. For electricity usage, we use the Electricity Dataset [36], which captures a single residence electricity usage in Canada from 2012 to 2014. We classify each hour as low or high consumption depending on whether the usage falls below or exceeds a fixed threshold of 80 kWh—the central value of the range. Additionally, we study different threshold values, 70 and 90 kWh, to assess their impact on utility. In all cases, we evaluate event-level local privacy guarantees, assuming no trusted curator and focusing on user-side privacy protection [18]. The technical details of the three datasets are summarized in Table 3.

In order to fulfill the conditions of Theorem 6.2 we require the transition probabilities of the Markov chain to be positive. We calculate them empirically and receive the following transition

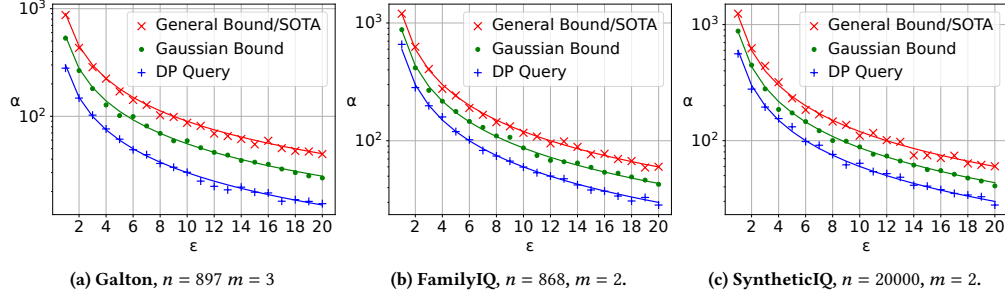


Figure 4: Gaussian data results. Lines show theoretical error at $\beta = 5\%$ and markers indicate empirical 95% upper bounds.

matrices for Activity and Electricity 70, 80, 90 kWh in this order:

$$\begin{pmatrix} 0.882 & 0.117 \\ 0.305 & 0.695 \end{pmatrix}, \begin{pmatrix} 0.445 & 0.555 \\ 0.149 & 0.850 \end{pmatrix}, \begin{pmatrix} 0.818 & 0.182 \\ 0.371 & 0.629 \end{pmatrix}, \begin{pmatrix} 0.894 & 0.106 \\ 0.478 & 0.522 \end{pmatrix},$$

representing $P_{00}, P_{01}, P_{10}, P_{11}$ with $s = 0$ inactive/low consumption and $s = 1$ active/high consumption. Our theorem also requires w to be a stationary distribution. While w can not be empirically computed since we only have one initial state, both Markov chains are irreducible, since both states are reachable from each other, aperiodic, since $P_{ss} \neq 0$ for both $s \in \{0, 1\}$, and $P_{st} > 0$ hence there exists a stationary initial distribution [11]. Therefore, we conclude that the databases fulfill the conditions for testing our results.

7.2 Target Queries and Utility metrics

We focus our utility study on two concrete although commonly used queries: sum and counting queries. Formally, given a database $D = (x_1, x_2, \dots, x_n)$, where each x_i represents a numerical value, a sum query is defined as: $q_S(D) = \sum_{i=1}^n x_i$. In the case of the Gaussian data, each x_i corresponds to an individual's height or IQ. If each record is binary, i.e., $x_i \in \{0, 1\}$, as is the case for the activity and electricity datasets, $q_S(D)$ is called a counting query since it outputs the count of states with the attribute 1.

Our theoretical results are expressed in terms of (α, β) -accuracy. To evaluate empirical utility, we use the upper bound of a $(1 - \beta)$ confidence interval for the absolute query error, which serves as a practical counterpart. Specifically, we report the upper limit of a 95% confidence interval (i.e., $\beta = 0.05$), a standard choice in practice [29]. A smaller upper bound indicates higher utility. When this bound is close to the theoretical error α , it demonstrates a strong alignment between empirical and theoretical results, highlighting their practical applicability. To facilitate comparison with our theoretical results, we plot the theoretical error tolerance α for each mechanism, derived from Proposition 3.5 for the baseline DP mechanism and Corollary 4.5, Corollary 5.9, and Corollary 6.3 for the general bound, the Gaussian bound and the Markov chain bound respectively. Additionally, to give an idea of the impact on utility in practice, we report the mean absolute percentage error (MAPE) in the long version of this paper to estimate the expected relative error for a single execution.

7.3 Mechanism and Experiment Design

In order to provide BDP mechanisms that approximate the target queries presented in Section 7.2, we use the Laplace mechanism with the noise calibrated through Theorem 4.3 for the DP baseline, Theorem 5.8 for Gaussian data and Theorem 6.2 for Markov data. In this section, we refer to the DP privacy leakage by τ , to avoid confusion with the actual maximum BDPL denoted by ε .

7.3.1 Gaussian Data. As explained in Section 7.1, we assume that the dataset is drawn from a multivariate Gaussian distribution with maximum number of correlated variables m respectively. Both the general bound and state of the art [53] indicate that for the Laplace mechanism $\mathcal{M}_{\tau, f}$, we have $\varepsilon = m\tau$, i.e., $\varepsilon = 3\tau$ for Galton and $\varepsilon = 2\tau$ for IQ datasets. Alternatively, according to the Pearson coefficients described in Table 3, Theorem 5.8 tells us that $\mathcal{M}_{\tau, f}$ is ε -BDP, with $\varepsilon \approx 1.853\tau, 1.45\tau$ for Galton and IQ datasets respectively. Consequently, we fix BDPL values $\varepsilon \in (0, 20]$ and compute the corresponding τ using Eq. 3 for the Gaussian-specific correlation approach and $\tau = \frac{\varepsilon}{3}$ for the general correlation and state of the art. For $\varepsilon \in (0, 5)$, we ensure strong theoretical privacy guarantees, while also considering the higher range $\varepsilon \in [5, 20]$, which has shown empirical resilience to certain privacy attacks [7, 40].

7.3.2 Markov Data. As discussed in Section 7.1 we assume that the data follows a Markov chain. According to the γ values summarized in Table 3, Theorem 6.2 tells us that the Laplace mechanism $\mathcal{M}_{\tau, f}$ applied to a counting query f is ε -BDP, with

$$\varepsilon_A = \tau + 8.05, \varepsilon_{E,70} \approx \tau + 4.7, \varepsilon_{E,80} \approx \tau + 6.03, \varepsilon_{E,90} \approx \tau + 8.54, \quad (6)$$

In comparison, with the general bound we have $\varepsilon = n\tau$ for mechanism $\mathcal{M}_{\tau, f}$. Similar to Gaussian data, we apply the Laplace mechanism to compute the sum query of each subgroup with BDPL values $\varepsilon \in (0, 20]$ and compute the corresponding τ using Eq. 6 for the Markov-specific mechanism and taking $\tau = \frac{\varepsilon}{n}$ for the general correlation approach. However, none of the datasets provide a symmetric transition matrix, which means that the proposal in [8] is not applicable, making an empirical comparison impossible.

Note that while ε -BDP can be provided for all values using the general bound and state of the art [8], Eq.6 only allows for $\varepsilon \geq 8.05, 6.9, 4.7$ and 8.45 for Activity and Electricity data respectively, since τ must be positive (see Section 6).

In all experiments, we calculate empirical confidence intervals executing the mechanism for each dataset 1000 times.

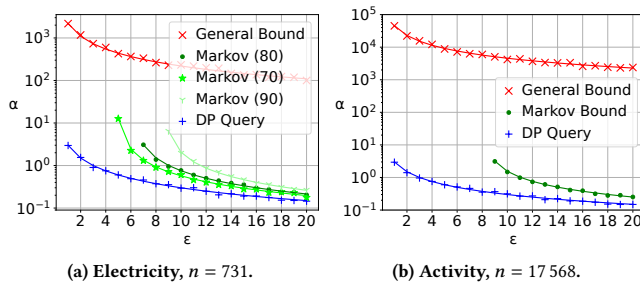


Figure 5: Markov data results. Lines show theoretical error at $\beta = 5\%$ and markers indicate empirical 95% upper bounds.

7.4 Results and Discussion

Figure 4 presents the results for the Gaussian models, including our Gaussian-specific bound, the state-of-the-art bound from [53] (which coincides with the general bound), and the DP Laplace mechanism for sum queries. We plot the DP mechanism as the baseline for the best possible utility; however, it is important to note that DP does not offer meaningful protection in this experiment, given correlation. Among the correlation-protecting mechanisms, those that use the Gaussian bound consistently outperform the s-o-t-a mechanism [53] for all ϵ in all datasets. Note that we plot all results on a logarithmic scale. This makes it harder to visually see the substantial reduction of error achieved by our mechanisms—particularly for small values of ϵ . For instance, for $\epsilon = 1$ the error is reduced by more than 400 units for both IQ datasets and 200 inches for the Galton. Note that the Galton height data uses imperial units (inches), thus the errors are also interpreted in inches.

The results for Markov chains are shown in Figure 5. Again, we use the DP mechanism as the baseline for the best possible utility, not as a comparable protective mechanism. For BDP mechanisms, we observe that the different Markov models tested lead to varying minimum achievable BDPL levels, as determined by our Markov-based bound: Electricity 70 kWh yields the most favorable case with a minimum $\epsilon = 4.9$, while 90 kWh imposes the weakest bound with a minimum $\epsilon = 8.45$. In contrast, the general bound supports all $\epsilon > 0$. In all cases where the Markov chain bound is applicable, mechanisms using it significantly outperform those relying on the general bound. While the error of mechanisms based on the general bound increases sharply, the error of both the Markov chain-based mechanism and the standard DP mechanism remains stable. In particular, in the Activity dataset the general bound results in a 10^5 times larger error than that of our proposed Markov chain bound. This is because the general bound scales with the size of the database n , while the Markov bound is independent of n , highlighting the huge benefit of using our novel bound for large datasets.

The results demonstrate that BDP mechanisms calibrated with our newly proven Gaussian and Markov chain bounds outperform prior BDP mechanisms and mechanisms calibrated with the general bound in terms of utility on real-world data. Moreover, the empirical errors from our experiments closely align with our theoretical utility results, validating the practical applicability of our theorems. We

extend this study with the analysis of the relative error in the long version of this paper obtaining similar results.

We acknowledge certain limitations when extrapolating our results. The validity of our experimental findings is constrained by the specific databases used. While the Galton height data serves as a well-known example of record correlation, it reflects only one of many possible correlation patterns. Similarly, most practical applications of a Markov chain would involve more than two states, introducing complexity beyond the binary-state model used in our study. Nevertheless, our results provide valuable insight into the practical applicability of our theorems and indicate their potential for real-world scenarios. Furthermore, these experiments demonstrate that achieving meaningful utility while protecting against correlation is feasible in practice.

8 CONCLUSION

In this paper, we explored the utility of BDP mechanisms for correlated data. We addressed prior limitations by analyzing broader correlation models and providing a detailed study of privacy-utility trade-offs, supported by theoretical results and empirical evidence. Specifically, we established new connections between DP and BDP mechanisms and demonstrated how they can be leveraged for privacy protection under correlation.

We proved that any ϵ -DP mechanism satisfies $m\epsilon$ -BDP, where m is the size of the correlated group, and showed this bound is tight. We then improved upon it by considering multivariate Gaussian and Markov models, deriving novel bounds on BDP leakage that provide stronger utility guarantees than the s-o-t-a approaches under the same privacy constraints. The advantage of our correlation-specific bounds is particularly evident under Markov-modeled correlations. While mechanisms based on the general bound exhibit high sensitivity to the number of correlated records, our Markov-based bound remains robust and stable regardless of the dataset size.

While it remains a futile attempt to apply BDP without assuming a specific correlation model, both our theoretical and experimental results demonstrate that it is possible to achieve better utility without weakening the adversary model in practical scenarios: (a) when the number of correlated records is small, (b) when the data follows a weakly correlated Gaussian model, or (c) when the data is a time series following a Markov chain with sufficiently similar transition probabilities.

Overall, our Theorems 4.3, 5.8 and 6.2 advance the theoretical and practical understanding of BDP, enabling the reuse of DP mechanisms in correlated settings. This opens future directions for deriving correlation-specific bounds, allowing the design of more accurate BDP mechanisms that protect against real-world, correlation-based attacks.

ACKNOWLEDGMENTS

This work was funded by the Topic Engineering Secure Systems of the Helmholtz Association (HGF) and supported by KASTEL Security Research Labs, Karlsruhe, and Germany’s Excellence Strategy (EXC 2050/1 ‘CeTI’; ID 390696704). J.P.-A. is a ‘‘Ram3n y Cajal’’ fellow (RYC2021-034256-I) funded by NextGenerationEU/PRTR and MCIN/AEI/10.13039/501100011033.

REFERENCES

- [1] Nour Almadhoun, Erman Ayday, and Özgür Ulusoy. 2020. Differential privacy under dependent tuples—the case of genomic privacy. *Bioinformatics* 36, 6 (2020), 1696–1703. <https://doi.org/10.1093/bioinformatics/btz837>
- [2] Miguel E. Andrés, Nicolás E. Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. 2013. Geo-indistinguishability: differential privacy for location-based systems. In *ACM SIGSAC Conference on Computer & Communications Security (CCS)*. ACM, New York, NY, USA, 901–914. <https://doi.org/10.1145/2508859.2516735>
- [3] Omid Ardakanian, Srinivasan Keshav, and Catherine Rosenberg. 2011. Markovian models for home electricity consumption. In *Proceedings of the 2nd ACM SIGCOMM workshop on Green networking*. ACM, New York, USA, 31–36. <https://doi.org/10.1145/2018536.2018544>
- [4] Ehrhard Behrends. 2000. *Introduction to Markov Chains*. Vieweg+Teubner Verlag, Wiesbaden. <https://doi.org/10.1007/978-3-322-90157-6>
- [5] Avrim Blum, Katrina Ligett, and Aaron Roth. 2013. A learning theory approach to noninteractive database privacy. *J. ACM* 60, 2 (2013), 1–25. <https://doi.org/10.1145/2450142.2450148>
- [6] Jennifer Brainard and David E. Burmaster. 1992. Bivariate Distributions for Height and Weight of Men and Women in the United States. *Risk Analysis* 12, 2 (1992), 267–275. <https://doi.org/10.1111/j.1539-6924.1992.tb00674.x>
- [7] Nicholas Carlini, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tramèr. 2022. Membership Inference Attacks From First Principles. In *IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, 1897–1914. <https://doi.org/10.1109/sp46214.2022.9833649>
- [8] Darshan Chakrabarti, Jie Gao, Aditya Saraf, Grant Schoenebeck, and Fang-Yi Yu. 2022. Optimal Local Bayesian Differential Privacy over Markov Chains. *arXiv:2206.11402 [cs.CR]* <https://arxiv.org/abs/2206.11402>
- [9] Konstantinos Chatzikokolakis, Miguel E Andrés, Nicolás Emilio Bordenabe, and Catuscia Palamidessi. 2013. Broadening the scope of differential privacy using metrics. In *Proceedings on Privacy Enhancing Technologies Symposium*. Springer,loomington, Indiana, United States., 82–102. https://doi.org/10.1007/978-3-642-39077-7_5
- [10] Rui Chen, Benjamin CM Fung, Philip S Yu, and Bipin C Desai. 2014. Correlated network data publication via differential privacy. *The VLDB Journal* 23, 4 (2014), 653–676. <https://doi.org/10.1007/s00778-013-0344-8>
- [11] Wai-Ki Ching and Michael K Ng. 2006. *Markov chains: models, algorithms and applications*. Springer, Boston, MA. https://doi.org/10.1007/0-387-29337-X_7
- [12] Kah Meng Chong and Amizah Malip. 2024. May the privacy be with us: Correlated differential privacy in location data for ITS. *Computer Networks* 241 (2024), 110214. <https://doi.org/10.1016/j.comnet.2024.110214>
- [13] Paul Cuff and Lanqing Yu. 2016. Differential Privacy as a Mutual Information Constraint. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS '16)*. ACM, New York, NY, USA, 43–54. <https://doi.org/10.1145/2976749.2978308>
- [14] Hadis Dalkani, Musa Mojarad, and Hassan Arfaeina. 2021. Modelling electricity consumption forecasting using the markov process and hybrid features selection. *International Journal of Intelligent Systems and Applications* 10, 5 (2021), 14. <https://doi.org/10.5815/ijisa.2021.05.02>
- [15] Damien Desfontaines and Balázs Pejó. 2020. SoK: Differential privacies. *Proceedings on Privacy Enhancing Technologies* 2020 (2020), 288–313. <https://doi.org/10.2478/popets-2020-0028>
- [16] Thi V. Duong, Hung H. Bui, Dinh Q. Phung, and Svetha Venkatesh. 2005. Activity Recognition and Abnormality Detection with the Switching Hidden Semi-Markov Model. In *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) (CVPR '05)*. IEEE Computer Society, USA, 838–845. <https://doi.org/10.1109/CVPR.2005.61>
- [17] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). Springer, Berlin, Heidelberg, 265–284. https://doi.org/10.1007/11681878_14
- [18] Cynthia Dwork and Aaron Roth. 2014. *The Algorithmic Foundations of Differential Privacy*. Now Publishers, Inc., Hanover, MA, USA. <https://doi.org/10.1561/04000000042>
- [19] Francis Galton. 2017. Galton height data. <https://doi.org/10.7910/DVN/T0HSJ1>
- [20] Sébastien Gams, Marc-Olivier Killijian, and Miguel Núñez del Prado Cortez. 2012. Next place prediction using mobility Markov chains. In *Proceedings of the First Workshop on Measurement, Privacy, and Mobility (Bern, Switzerland) (MPM '12)*. ACM, New York, NY, USA, Article 3, 6 pages. <https://doi.org/10.1145/2181196.2181199>
- [21] Semyon Aranovich Gershgorin. 1931. Über die Abgrenzung der Eigenwerte einer Matrix. *Izvestija Rossijskoj akademii nauk. Serija matematičeskaja* 1, 6 (1931), 749–754.
- [22] Franklin A Graybill and Hariharan K Iyer. 1994. Retrieved May 13, 2025 from <https://www.kaggle.com/datasets/jacopoferretti/child-vs-mother-iq/data?select=gifted.csv>
- [23] Xi He, Ashwin Machanavajjhala, and Bolin Ding. 2014. Blowfish privacy: tuning privacy-utility trade-offs using policies. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data (Snowbird, Utah, USA) (SIGMOD '14)*. ACM, New York, NY, USA, 1447–1458. <https://doi.org/10.1145/2588555.2588581>
- [24] Qi Huang, Dwayne Cohen, Sandra Komarzynski, Xiao-Mei Li, Pasquale Innominato, Francis Lévi, and Bärbel Finkenstädt. 2018. Hidden Markov models for monitoring circadian rhythmicity in telemetric activity data. *Journal of The Royal Society Interface* 15, 139 (2018), 20170885.
- [25] Thomas Humphries, Simon Oya, Lindsey Tulloch, Matthew Rafuse, Ian Goldberg, Urs Hengartner, and Florian Kerschbaum. 2023. Investigating Membership Inference Attacks under Data Dependencies. In *IEEE Computer Security Foundations Symposium (CSF)*. IEEE, Dubrovnik, Croatia, 473–488. <https://doi.org/10.1109/csf57540.2023.00013>
- [26] Daniel Kifer and Ashwin Machanavajjhala. 2011. No free lunch in data privacy. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data (Athens, Greece) (SIGMOD '11)*. ACM, New York, NY, USA, 193–204. <https://doi.org/10.1145/1989323.1989345>
- [27] Daniel Kifer and Ashwin Machanavajjhala. 2011. No free lunch in data privacy. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data (SIGMOD '11)*. ACM, New York, USA, 193–204. <https://doi.org/10.1145/1989323.1989345>
- [28] Daniel Kifer and Ashwin Machanavajjhala. 2014. Pufferfish: A framework for mathematical privacy definitions. *ACM Trans. Database Syst.* 39, 1 (2014), 3:1–3:36. <https://doi.org/10.1145/2514689>
- [29] Dong Kyu Lee. 2016. Alternatives to P value: confidence interval and effect size. *Korean Journal of Anesthesiology* 69, 6 (2016), 555–562. <https://doi.org/10.4097/kjae.2016.69.6.555>
- [30] Jaewoo Lee and Chris Clifton. 2011. How Much Is Enough? Choosing Epsilon for Differential Privacy. In *Information Security*. Springer, Berlin, Heidelberg, 325–340. https://doi.org/10.1007/978-3-642-24861-0_22
- [31] David A Levin and Yuval Peres. 2017. *Markov chains and mixing times*. Vol. 107. American Mathematical Soc., USA. <https://doi.org/10.1090/mbk/107>
- [32] Yanan Li, Xuebin Ren, Shusen Yang, and Xinyu Yang. 2019. Impact of prior knowledge and data correlation on privacy leakage: A unified analysis. *IEEE Transactions on Information Forensics and Security* 14, 9 (2019), 2342–2357.
- [33] David Liben-Nowell and Jon Kleinberg. 2003. The link prediction problem for social networks. In *Proceedings of the Twelfth International Conference on Information and Knowledge Management (CIKM '03)*. ACM, New York, NY, USA, 556–559. <https://doi.org/10.1145/956863.956972>
- [34] Changchang Liu, Supriyo Chakraborty, and Prateek Mittal. 2016. Dependence Makes You Vulnerable: Differential Privacy Under Dependent Tuples. In *23rd Annual Network and Distributed System Security Symposium, NDSS 2016*, Vol. 16. The Internet Society, San Diego, USA, 21–24. <https://doi.org/10.14722/ndss.2016.23279>
- [35] Zhong Cheng Luo, Kerstin Albertsson-Wikland, and Johan Karlberg. 1998. Target Height as Predicted by Parental Heights in a Population-Based Study. *Pediatric Research* 44(4) (1998), 563–571. <https://doi.org/10.1203/00006450-199810000-00016>
- [36] Stephen Makonin, Bradley Ellert, Ivan V Bajić, and Fred Popowich. 2016. Electricity, water, and natural gas consumption of a residential house in Canada from 2012 to 2014. *Scientific data* 3, 1 (2016), 1–12. <https://doi.org/10.1038/sdata.2016.37>
- [37] Shambavi Malik. 2020. Activity Data. <https://www.kaggle.com/datasets/shambhvimalik/activity-data/data>. Accessed: 2024-06-17.
- [38] Àlex Miranda-Pascual, Patricia Guerra-Balboa, Javier Parra-Arnau, Jordi Forné, and Thorsten Strufe. 2023. SoK: differentially private publication of trajectory data. *Proceedings on Privacy Enhancing Technologies* 2023 (2023), 496–516. <https://doi.org/10.56553/popets-2023-0065>
- [39] Joakim Munkhammar, Dennis van der Meer, and Joakim Widén. 2021. Very short term load forecasting of residential electricity consumption using the Markov-chain mixture distribution (MCM) model. *Applied Energy* 282 (2021), 116180. <https://doi.org/10.1016/j.apenergy.2020.116180>
- [40] Joseph Near and David Darais. 2022. Differential Privacy: Future Work & Open Challenges. <https://www.nist.gov/blogs/cybersecurity-insights/differential-privacy-future-work-open-challenges>. Accessed: 2024-06-11.
- [41] Theshani Nuradha and Ziv Goldfeld. 2023. Pufferfish Privacy: An Information-Theoretic Study. *IEEE Transactions on Information Theory* 69, 11 (2023), 7336–7356. <https://doi.org/10.1109/TIT.2023.3296288>
- [42] Victor M. Panaretos. 2016. *Statistics for Mathematicians*. Springer International Publishing, Switzerland. <https://doi.org/10.1007/978-3-319-28341-8>
- [43] Robert Plomin, J. C. DeFries, Valerie S. Knopik, and Jenae M. Neiderhiser. 2013. *Behavioral genetics: a primer* (sixth edition ed.). Worth Publishers, New York.
- [44] Havard Rue and Leonhard Held. 2005. *Gaussian Markov random fields: theory and applications*. Chapman and Hall/CRC, New York. <https://doi.org/10.1201/9780203492024>
- [45] Jun Shao. 2003. *Mathematical Statistics*. Springer, New York. <https://doi.org/10.1007/b97553>

- [46] Shuang Song, Yizhen Wang, and Kamalika Chaudhuri. 2017. Pufferfish Privacy Mechanisms for Correlated Data. In *Proceedings of the 2017 ACM International Conference on Management of Data SIGMOD*. ACM, New York, USA, 1291–1306. <https://doi.org/10.1145/3035918.3064025>
- [47] Mikael Sunnåker and Joerg Stelling. 2015. Model extension and model selection. In *Uncertainty in Biology: A Computational Modeling Approach*, Vol. 17. Springer, Cham, Switzerland, 213–241.
- [48] Salil Vadhan. 2017. *The complexity of differential privacy*. Springer, Cham, Switzerland. 347–450 pages. https://doi.org/10.1007/978-3-319-57048-8_7
- [49] Hao Wang, Zhengquan Xu, Shan Jia, Ying Xia, and Xu Zhang. 2021. Why current differential privacy schemes are inapplicable for correlated data publishing? *World Wide Web* 24 (2021), 1–23. <https://doi.org/10.1007/s11280-020-00825-8>
- [50] Tianhao Wang, Jeremiah Blocki, Ninghui Li, and Somesh Jha. 2017. Locally Differentially Private Protocols for Frequency Estimation. In *26th USENIX Security Symposium*. USENIX Association, Vancouver, BC, 729–745. <https://doi.org/10.5555/3241189.3241247>
- [51] Larry Wasserman and Shuheng Zhou. 2010. A statistical framework for differential privacy. *J. Amer. Statist. Assoc.* 105, 489 (2010), 375–389.
- [52] Daniel S Wilks. 2011. *Statistical methods in the atmospheric sciences*. Vol. 100. Academic Press, Oxford, UK. <https://doi.org/10.1016/C2017-0-03921-6>
- [53] Bin Yang, Issei Sato, and Hiroshi Nakagawa. 2015. Bayesian Differential Privacy on Correlated Data. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (Melbourne, Victoria, Australia) (SIGMOD '15)*. ACM, New York, NY, USA, 747–762. <https://doi.org/10.1145/2723372.2747643>