Differentially Private Data-Driven Markov Chain Modeling

Anonymous authors
Paper under double-blind review

Abstract

Markov chains model a wide range of user behaviors. However, generating accurate Markov chain models requires substantial user data, and sharing these models without privacy protections may reveal sensitive information about the underlying user data. We introduce a method for protecting user data used to formulate a Markov chain model. First, we develop a method for privatizing database queries whose outputs are elements of the unit simplex, and we prove that this method is differentially private. We quantify its accuracy by bounding the expected KL divergence between private and non-private queries. We extend this method to privatize stochastic matrices whose rows are each a simplex-valued query of a database, which includes data-driven Markov chain models. To assess their accuracy, we analytically bound the change in the stationary distribution and the change in the convergence rate between a non-private Markov chain model and its private form. Simulations show that under a typical privacy implementation, our method yields less than 2% error in the stationary distribution, indicating that our approach to private modeling faithfully captures the behavior of the systems we study.

1 Introduction

Markov chain modeling is a powerful tool for understanding a variety of behaviors including domestic time use (Widén et al., 2009), traffic patterns (Gong et al., 2011), and internet browsing (Dong et al., 2022). Markov chains are often to model systems with random transitions among a finite set of states. In data-driven settings, transition probabilities are typically derived from databases of observed user behavior, as in Rendle et al. (2010) and Lundström et al. (2016). The development of these models often counts transitions between states in user data, and sharing the resulting models may leak information about the data used to generate the transition probabilities. Patterns such as home occupancy (Lundström et al., 2016) or individual shopping behavior (Rendle et al., 2010) may be inferred even from aggregated data. Thus, without some form of protection, these models present significant privacy risks for users. However, these models fundamentally require user data, which motivates our interest in protecting privacy while preserving model accuracy.

In this paper, we develop a framework for producing privacy-preserving Markov chain models that protects databases used to compute transition probabilities. We first develop a method for privatizing database queries whose outputs are stochastic vectors, i.e., vectors with non-negative entries that sum to 1. We then extend this method to stochastic matrices in which every row is equal to a query whose output is a stochastic vector.

We develop this framework using differential privacy. Differential privacy is a statistical notion of privacy designed to preserve the privacy of sensitive databases when queried (Dwork et al., 2006). We use differential privacy because (i) compositions of differentially private data remain differentially private, and (ii) it is immune to post-processing, in the sense that post-hoc computations on private data do not weaken differential privacy. These properties have led to the wide use of differential privacy for federated learning (Geyer et al., 2017; Agarwal et al., 2021; Chen et al., 2022; Noble et al., 2022), optimization (Hsu et al., 2014; Wang et al., 2016; Munoz et al., 2021; Benvenuti et al., 2024a), planning (Gohari et al., 2020; Chen et al., 2023a; Benvenuti et al., 2024b), and machine learning (Ponomareva et al., 2023; Jayaraman & Evans, 2019; Blanco-Justicia et al., 2022).

In the context of Markov chains, the compositional properties of differential privacy mechanisms allow us to privatize each row in a stochastic matrix to form a private Markov chain model, and the immunity of differentially private data to post-processing allows the private model to be shared and analyzed without harming the privacy of the underlying database. Additionally, the differential privacy literature has shown that highly accurate query outputs can sometimes be produced while enforcing strong privacy (Dwork et al., 2014). Thus, we use differential privacy to generate and share accurate Markov chain models while still protecting users with strong privacy.

To privatize stochastic vectors, we extend the Dirichlet mechanism first introduced in Gohari et al. (2021). Common privacy mechanisms such as the Gaussian and Laplace mechanisms (Dwork et al., 2014; McSherry, 2009) are ill-suited for stochastic vectors because they add noise with infinite support, which may cause private output vectors to no longer be stochastic. The resulting vectors could be projected onto the unit simplex to enforce that their entries are non-negative and sum to 1, but Gohari et al. (2021) showed such projections result in poor accuracy. We fundamentally differ from Gohari et al. (2021) because we privatize databases that are used to compute simplex-valued queries, whereas Gohari et al. (2021) privatized elements of the simplex themselves (without there being an underlying database that generated them). Our approach provides event-level privacy to data, in the sense that the presence of individual events is protected by differential privacy. This form of privacy is intended to be used for datasets in which individual users are not tracked over time, such as the taxi cab dataset in NYC Taxi Limousine Commision (2025), to which we apply the developed privacy mechanism in Section 6.

Of course, we expect the asymptotic and transient behaviors of a privatized Markov chain to differ from those of the comparable non-private Markov chain. To assess the impact of privacy on asymptotic behavior, we bound the change in the stationary distribution (i.e., the long-run probability distribution of the Markov chain's state) between a private and non-private Markov chain. To assess the impact on transient behavior, we bound the change in ergodicity coefficient (which is a proxy for the convergence rate to the stationary distribution) between a private and non-private Markov chain. These bounds quantify privacy's impact on Markov chain models, and they can be used to tune privacy parameters based on accuracy.

To summarize, our contributions are as follows:

- We develop a framework for privatizing queries whose outputs are stochastic vectors (Theorem 1).
- We bound the Kullback-Leibler divergence between a non-private stochastic vector and its privatized form (Theorem 2 and Corollary 1).
- Using Theorem 1, we develop a framework for privatizing the computation of Markov chain transition probabilities computed from user data (Theorem 3).
- We bound the changes in asymptotic and transient behaviors between the original, non-private Markov chain and its privatized form by bounding the differences in the stationary distribution and the ergodicity coefficient (Theorems 4 and 5).
- We empirically validate the impact of privacy on performance on two real datasets: a class grade distribution and transit data from New York City taxis. Results show less than 2% error in the stationary distribution with $(3.73, 3 \times 10^{-6})$-differential privacy (Section 6).

Related Work In this work, we build on the original Dirichlet mechanism from Gohari et al. (2021). More recently, Ponnoprat (2023) and Chitra et al. (2022) have developed alternate formulations of the Dirichlet Mechanism. However, these formulations introduce a bias in private outputs, which is undesirable for the Markov chain models we consider. Thus, we base our approach on Gohari et al. (2021) to yield unbiased private outputs.

Privacy for Markov processes has been extensively studied (Gohari et al., 2020; Chen et al., 2023b;a; Fallin et al., 2023; Benvenuti et al., 2024b), though only Gohari et al. (2020) and Fallin et al. (2023) privatize the transition probabilities, which we do in this work. However, those works consider the probabilities themselves as the sensitive data and apply input perturbation privacy to them. In contrast, we consider

transition probabilities that are computed from data, and we privatize the underlying data, not the values of the probabilities. Thus, Gohari et al. (2020) and Fallin et al. (2023) represent special cases of this work where the transition probabilities are the sensitive database.

Notation For $N \in \mathbb{N}$, we use $[N] := \{1, \ldots, N\}$. We use $|\cdot|$ to denote the cardinality of a set. We use $\text{diag}(a_1, \ldots, a_n)$ to denote an $n \times n$ diagonal matrix with diagonal entries $a_1, \ldots, a_n$. We use $\mathbb{I}\{\cdot\}$ to denote the indicator function and $\mathbb{1}_n$ to be a vector of all 1s of length n. For a random variable X, we use $\mathbb{E}[X]$ to denote its expectation. We use $\mathcal{K}(q||\ell) = \sum_{i=1}^{n} q_i \log\left(\frac{q_i}{\ell_i}\right)$ to denote the KL divergence between q and ℓ, where we use the convention $0 \log\left(\frac{0}{0}\right) = 0$. We use the following special functions: the Gamma function $\Gamma(x) = \int_0^\infty t^{x-1} e^{-t} dt$, the Beta function $\text{Beta}(x_1, x_2) = \int_0^1 t^{x_1-1}(1-t)^{x_2-1} dt = \frac{\Gamma(x_1)\Gamma(x_2)}{\Gamma(x_1+x_2)}$, the multi-variate Beta function $B(y) = \prod_{i=1}^{n} \Gamma(y_i) \left(\Gamma\left(\sum_{i=1}^{n} y_i\right)\right)^{-1}$, and the digamma function $\psi(x) = \frac{\Gamma'(x)}{\Gamma(x)}$.

2 Preliminaries and Problem Formulation

In this section, we provide background on Markov chains and differential privacy, define the structure of the databases and queries we consider, and formalize problem statements.

2.1 Unit Simplex and Stochastic Matrices

Definition 1 (Unit Simplex)**.** Let $n \in \mathbb{N}$. The unit simplex in $\mathbb{R}^n$ denoted by Δ_n is the set of all element-wise non-negative vectors of length n whose entries sum to 1, where

$$\Delta_n = \left\{ x \in \mathbb{R}^n \;\middle|\; \sum_{i=1}^{n} x_i = 1,\ x_j \geq 0 \text{ for all } j \in [n] \right\}. \blacktriangle$$

We also refer to elements of the unit simplex as “stochastic vectors". We also consider the bordered unit simplex.

Definition 2 (Bordered Unit Simplex)**.** Fix $n \geq 3$ and $\eta > 0$. Then the bordered unit simplex is

$$\Delta_n^{(\eta)} = \left\{ x \in \Delta_n \mid x_i \geq \eta \text{ for all } i \in [n] \right\}. \quad \blacktriangle$$

The bordered unit simplex is a subset of Δ_n whose components are a sufficient distance from 0.

A matrix whose rows are all stochastic vectors is called a “stochastic matrix". For $n \in \mathbb{N}$, we define $\mathbb{S}_n$ as the set of all $n \times n$ stochastic matrices.

2.2 Markov Chains

Consider a discrete time stochastic process $(Y_t)_{t \in \mathbb{N}}$ on the state space $\mathcal{Y}$. The process $(Y_t)_{t \in \mathbb{N}}$ is said to be a discrete time Markov chain if the sequence satisfies the Markov property, namely $\mathbb{P}(Y_{t+1} = y_{t+1} | Y_1 = y_1, \cdots, Y_t = y_t) = \mathbb{P}(Y_{t+1} | Y_t = y_t)$, which states that the conditional probability of transitioning from state y_t to y_{t+1} is independent of the sequence of states prior to t. See (Norris, 1998) for a detailed account of Markov chains. We study Markov chains that are (i) finite, (ii) irreducible, and (iii) homogeneous. These are Markov chains in which (i) the state space $\mathcal{Y}$ is finite, (ii) any state $y \in \mathcal{Y}$ can be reached from any other state $z \in \mathcal{Y}$ with positive probability, and (iii) the transition probabilities are independent of time. Given a finite-state Markov chain, we construct the transition probability matrix G such that $G_{ij} = \mathbb{P}(Y_{t+1} = y_j \mid Y_t = y_i)$. Row i of the matrix G, which we denote G_i, is the probability distribution of transitioning from state y_i to any other state in $\mathcal{Y}$. If $|\mathcal{Y}| = n$, then $G \in \mathbb{S}_n$.

2.3 Differential Privacy

We use differential privacy to provide privacy to a database $\mathcal{D}$, which we partition as $\mathcal{D} = \{D_1, \ldots, D_m\}$, where each set D_i contains $N_i \in \mathbb{N}$ entries and $D_i = \{d_1^i, \ldots, d_{N_i}^i\}$. The goal of differential privacy is to make

"similar" data appear "approximately indistinguishable" by using a randomized mapping or "privacy mechanism". Given two databases $\mathcal{D}$ and $\mathcal{D}'$, the notion of "similar" is defined by an adjacency relation (Dwork et al., 2014).

Definition 3 (Adjacency)**.** Consider two databases $\mathcal{D} = \{D_1, \ldots, D_m\}$ and $\mathcal{D}' = \{D'_1, \ldots, D'_m\}$, where D_i and D'_i each contain $N_i \in \mathbb{N}$ entries for all $i \in [m]$. Suppose $D_i = \{d^i_1, \ldots, d^i_{N_i}\}$ and similar for D'_i. We write $\text{Adj}(\mathcal{D}, \mathcal{D}') = 1$ if (a) there exist indices i, j such that $d^i_j \neq d'^i_j$ and (b) for all $(\ell, k) \neq (i, j)$, we have $d^\ell_k = d'^\ell_k$. That is, $\mathcal{D}$ and $\mathcal{D}'$ are adjacent if they differ in one entry. We write $\text{Adj}(\mathcal{D}, \mathcal{D}') = 0$ otherwise. ▲

A privacy mechanism $\mathcal{M}(\cdot)$ is a randomized mapping that enforces differential privacy.

Definition 4 (Differential Privacy; (Dwork et al., 2014))**.** Fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let $\delta \in [0, \frac{1}{2})$ and $\epsilon > 0$ be given. A mechanism $\mathcal{M} : \mathbb{R}^n \times \Omega \to \mathbb{R}^m$ is (ϵ, δ)-differentially private if for all $\mathcal{D}, \mathcal{D}'$ that are adjacent in the sense of Definition 3, we have $\mathbb{P}(\mathcal{M}(\mathcal{D}) \in T) \leq e^\epsilon \mathbb{P}(\mathcal{M}(\mathcal{D}') \in T) + \delta$ for all Borel measurable sets $T \subseteq \mathbb{R}^m$. ▲

The parameters ϵ and δ quantify the strength of privacy, and smaller ϵ and δ imply stronger privacy.

2.4 Databases and Queries

We consider databases $\mathcal{D}$ that are partitioned via $\mathcal{D} = \{D_1, \ldots, D_m\} \in \mathfrak{D}$, where $\mathfrak{D}$ is the set of all possible database realizations. We sometimes consider the degenerate case where $\mathcal{D} = \{D\}$. To differentiate between these cases, we use D without a subscript to mean $\mathcal{D} = D = \{d_1, \ldots, d_N\}$. A database entry d^i_j may come from one of n categories in a category set $\rho = \{\rho_1, \ldots, \rho_n\} \in \mathcal{P}$, where $\mathcal{P}$ is the collection of all category sets. We use $d^i_j \models \rho_k$ to indicate that entry d^i_j is from category ρ_k.

Assumption 1. Each database entry d^i_j belongs to exactly one category $\rho_k \in \rho$.

Assumption 1 is mild since the categories may represent physical events that cannot happen concurrently.

Example 1 (City Traffic)**.** Consider a database D that contains one record d_j for each time a driver turns from a street called First Street onto another street. Then $\mathfrak{D}$ is the set of all traffic patterns in the city. We use $\mathcal{P}$ to represent all streets in the city, and the category set ρ is the collection of streets that intersect First Street that drivers may turn onto. In this scenario, Assumption 1 says that a single driver must turn off somewhere and cannot turn off onto multiple different roads. △

For a database $\mathcal{D} = \{D\}$ with N entries, consider the counting query $C : \mathfrak{D} \times \mathcal{P} \to \mathbb{R}^n$ with category set $\rho = \{\rho_1, \ldots, \rho_n\}$, defined as

$$C(D, \rho) = \frac{1}{N} \left[\sum_{i=1}^N \mathbb{I}\{d_i \models \rho_1\}, \cdots, \sum_{i=1}^N \mathbb{I}\{d_i \models \rho_n\} \right]^T. \tag{1}$$

In words, $C(D, \rho)$ is a vector whose entries are the fraction of records in D that belong to category ρ_i for each $i \in [n]$.

Next, we consider the more general case where $\mathcal{D} = \{D_1, \ldots, D_m\}$, where each D_i contains N_i entries and we have the category set $\rho = \{\rho_1, \ldots, \rho_n\}$. We then consider the mapping $P : \mathfrak{D} \times \mathcal{P} \to \mathbb{R}^{m \times n}$, where each row of the output matrix is a count vector for a database D_i, namely

$$P(\mathcal{D}, \rho) = \begin{bmatrix} C(D_1, \rho) & \cdots & C(D_m, \rho) \end{bmatrix}^T. \tag{2}$$

When $m = n$, the matrix $P(\mathcal{D}, \rho)$ in (2) is a stochastic matrix, and it represents the transition probability matrix of a Markov chain. One interpretation of $P(\mathcal{D}, \rho)$ is as follows: D_i encodes all of the transitions of users from state i to some other state, and $\rho = \{\rho_1, \ldots, \rho_n\}$ is the set of states that they can transition to. Then $\sum_{j=1}^{N_i} \mathbb{1}\{d^i_j \models \rho_k\}$ is the count of how many users transitioned from state i to state k. Accordingly, each entry in $P(\mathcal{D}, \rho)$ is the empirical transition probability in a Markov chain.

2.5 Problem Statements

In this work, we solve the following problems:

Problem 1. *Develop a mechanism to privatize $\mathcal{D}$ when computing count vectors $C(\mathcal{D}, \rho)$ in the form of* (1).

Problem 2. *Bound the similarity between a privately generated count vector $\tilde{C}$ and the original, non-private count vector $C(\mathcal{D}, \rho)$.*

Problem 3. *Develop a privacy mechanism to privatize $\mathcal{D}$ when computing the transition matrix of a Markov chain, $P(\mathcal{D}, \rho)$.*

Problem 4. *Bound the difference in the asymptotic behavior (i.e., stationary distribution) and the transient behavior (i.e., convergence rate) between a Markov chain and its private form.*

3 Dirichlet Mechanism for Differential Privacy of Simplex-Valued Queries

In this section, we solve Problems 1 and 2. We consider the case where $\mathcal{D} = \{D\}$ and denote this as D. To solve Problem 1, we extend the Dirichlet mechanism, first introduced in Gohari et al. (2021).

Definition 5 (Dirichlet Mechanism; (Gohari et al., 2021))**.** A Dirichlet mechanism with parameter $k > 0$, which we denote $\mathcal{M}_{\text{Dir}}^{(k)}$, takes a vector $p \in \Delta_n^\circ$ as input and outputs $x \in \Delta_n$ according to the Dirichlet distribution centered on kp. That is,

$$\mathbb{P}\left(\mathcal{M}_{\text{Dir}}^{(k)}(p) = x\right) = \frac{1}{B(kp)} \prod_{i=1}^{n-1} x_i^{kp_i - 1} \left(1 - \sum_{i=1}^{n-1} x_i\right)^{kp_n - 1}. \blacktriangle$$

The Dirichlet mechanism takes a sensitive vector p and outputs a private vector $\tilde{p}$. The level of privacy is tuned using the parameter k. The Dirichlet mechanism in Definition 5 satisfies an alternate notion of privacy known as probabilistic differential privacy. To distinguish the two notions, we say "conventional differential privacy" to refer to Definition 4.

Definition 6 (Probabilistic Differential Privacy; (Machanavajjhala et al., 2008), (Meiser, 2018))**.** Let $\mathcal{M}$ be a randomized privacy mechanism, and let $\mathcal{S}$ be the set of possible outputs of the mechanism. The mechanism $\mathcal{M}$ satisfies (ϵ, δ)-probabilistic differential privacy if the set $\mathcal{S}$ can be partitioned into two disjoint sets Ω_1 and Ω_2 such that for all $\mathcal{D}$, we have $\mathbb{P}\left(\mathcal{M}(\mathcal{D}) \in \Omega_2\right) \leq \delta$, and for all $\mathcal{D}'$ adjacent to $\mathcal{D}$ and all $S \in \Omega_1$, we have $\log\left(\frac{\mathbb{P}(\mathcal{M}(\mathcal{D})=S)}{\mathbb{P}(\mathcal{M}(\mathcal{D}')=S)}\right) \leq \epsilon$. $\blacktriangle$

In this work, we partition the output space Δ_n into the sets $\Omega_1 = \{x \in \Delta_n \mid x_i \geq \gamma \text{ for all } i \in [n]\}$ and $\Omega_2 = \{x \in \Delta_n \mid x \notin \Omega_1\}$.

Assumption 2. When generating private outputs in Δ_n, γ satisfies $\gamma \leq \frac{1}{n-1}$.

If a mechanism provides (ϵ, δ)-probabilistic differential privacy, then it provides conventional (ϵ, δ)-differential privacy (Machanavajjhala et al., 2008). As mentioned in Meiser (2018), post-processing private data generated by mechanisms that satisfy (ϵ, δ)-probabilistic differential privacy only preserves conventional (ϵ, δ)-differential privacy. As a result, we use Definition 6 only as a means to design a mechanism that satisfies Definition 4.

To solve Problem 1, we input $C(D, \rho)$ into the Dirichlet mechanism to compute a privatized form $\tilde{C}$. We must then compute the privacy parameters ϵ and δ from Definition 6. The mechanism in Definition 5 does not take a database as input, and we must extend that mechanism to the case where the queries are functions of a database. This extension primarily involves ϵ. See the Appendix A for details on computing δ.

3.1 Computing ϵ

To compute ϵ, we first have two assumptions that ensure the boundedness of the distribution of $\tilde{C}$.

Assumption 3. In the bordered unit simplex $\Delta_n^{(\eta)}$ from Definition 2, we require $\eta \in (0, \frac{1}{4})$.

Algorithm 1 Privatizing counts of sensitive data

Inputs: D, ρ, N, γ, η, k
Outputs: Privacy-preserving count $\tilde{C}$, privacy parameters ϵ, δ

1: Compute the count $C(D, \rho)$ using (1)
2: Compute ϵ and δ via (3) and (4)
3: Generate $\tilde{C} = \mathcal{M}_{\text{Dir}}^{(k)}(C(D, \rho))$

Assumption 4. The Dirichlet parameter k from Definition 5 satisfies $k \geq \frac{3}{2\eta}$.

Assumption 3 is easily enforced by the user selecting an appropriate η, since $\min_{i\in[n]} p_i \geq \eta$ by the construction of $C(D, \rho)$. Assumption 4 is also enforced by the user choosing a suitable value for k. Assumption 4 highlights that not every count vector can be made arbitrarily private. Instead, there exists some strongest level of privacy (i.e., a smallest ϵ) for a given $C(D, \rho)$. We see in Section 6 that this strongest level of privacy is often well within the range of typically desired ϵ values. Below we use $\mathcal{C}(\rho)$ to denote the set of all possible outputs of $C(D, \rho)$ over all databases of size N whose entries satisfy Assumption 1. Next we provide (ϵ, δ)-differential privacy to $C(D, \rho)$. All proofs can be found in Appendix C.

Theorem 1 (Solution to Problem 1)**.** *Fix a database D with $N \in \mathbb{N}$ entries, a category set ρ such that $|\rho| = n \in \mathbb{N}$, and $\eta > 0$. Let $C(D, \rho) \in \Delta_n^{(\eta)}$ be the count vector as defined in* (1)*. Let Assumptions 1-4 hold. Then the Dirichlet mechanism with parameter $k > 0$ and input $C(D, \rho)$ is (ϵ, δ)-probabilistic differentially private, where*

$$\epsilon \geq \log\left(\frac{\text{Beta}(k\eta, k(1-2\eta))}{\text{Beta}\left(k\left(\eta + \frac{1}{N}\right), k(1-2\eta-\frac{1}{N})\right)}\right) + \frac{k}{N}\log\left(\frac{1-(n-1)\gamma}{\gamma}\right) \tag{3}$$

and

$$\delta = 1 - \min_{C(D,\rho)\in\mathcal{C}(\rho)} \mathbb{P}\left(\mathcal{M}_{Dir}^{(k)}(C(D, \rho)) \in \Omega_1\right). \tag{4}$$

Theorem 1 shows for the first time that the Dirichlet mechanism in Definition 5 provides (ϵ, δ)-probabilistic (and thus conventional) differential privacy for simplex-valued queries of databases. A typical differential privacy implementation begins by choosing a value of ϵ to implement. Theorem 1 and Assumption 4 together show that the smallest possible ϵ here is $\epsilon^* = \log\left(\frac{\text{Beta}(\frac{3}{2}, \frac{3}{2\eta}-3)}{\text{Beta}\left(\frac{3}{2}+\frac{3}{2\eta N}, \frac{3}{2\eta}-3-\frac{3}{2\eta N}\right)}\right) + \frac{3}{2\eta N}\log\left(\frac{1-(n-1)\gamma}{\gamma}\right)$. A user can select any $\epsilon \geq \epsilon^*$, and that choice can be implemented by tuning k. This workflow therefore largely mirrors the standard workflow for implementing differential privacy.

Algorithm 1 provides an overview of our method for privatizing D when computing $C(D, \rho)$.

To analyze the accuracy of Algorithm 1, we seek to compare $C(D, \rho)$ and $\tilde{C}$. Both $C(D, \rho)$ and $\tilde{C}$ can be interpreted as probability distributions because they are both in Δ_n. One way to measure the similarity of two probability distributions is the KL divergence, and we solve Problem 2 by computing the KL divergence between $C(D, \rho)$ and $\tilde{C}$.

Theorem 2 (Solution to Problem 2)**.** *Let the conditions of Theorem 1 hold and let $C = C(D, \rho)$. Then*

$$\mathbb{E}\left[\mathcal{K}(C||\tilde{C})\right] = \sum_{i=1}^{n} C_i(\log(C_i) + \psi(k) - \psi(kC_i)). \tag{5}$$

The result in (5) depends on the sensitive count vector $C(D, \rho)$, which means that it cannot be computed by someone who only has access to the private output $\tilde{C}$. To develop a bound that holds independent of the realization of the sensitive data, we upper bound (5) using the KL divergence-maximizing case, which is $C(D, \rho) = [\frac{1}{N}\mathbb{1}_{n-1}^T,\ 1 - \frac{n-1}{N}]^T$.

Corollary 1 (Alternative Solution to Problem 2)**.** *Let the conditions of Theorem 2 hold. Let $\zeta(x) = \log\left(\frac{x+1}{N}\right) - \psi\left(\frac{(x+1)k}{N}\right)$. Then $\mathbb{E}\left[\mathcal{K}(C(D, \rho)||\tilde{C})\right] \leq \frac{n-1}{N}\zeta(0) + \frac{N-n+1}{N}\zeta(N-n) + \psi(k)$.*

Corollary 1 provides users with a tool to tune the average error between a private count vector $\tilde{C}$ and the original, non-private count vector $C(D, \rho)$. Section 6.1 shows the tightness of the Corollary 1 relative to Theorem 2. Appendix D provides additional accuracy results for the privatized count vector $\tilde{C}$. Next we apply Algorithm 1 to a database to privatize stochastic matrices.

4 Differential Privacy for Markov Chain Modeling

In this section, we solve Problem 3. Each row of a Markov chain's transition matrix can be computed with a counting query on the unit simplex. We therefore use the mechanism developed in Section 3 to privatize each row of a transition matrix individually. We apply a property known as "parallel composition" (McSherry, 2009) to produce a privacy-preserving Markov chain model that protects the database $\mathcal{D}$ with (ϵ, δ)-differential privacy.

Parallel composition of ϵ-differentially private mechanisms was first proven in McSherry (2009). However, it has since been used in the literature to develop (ϵ, δ)-differential privacy mechanisms (Ponomareva et al., 2023), often without proof. Incomplete claims about parallel composition for (ϵ, δ)-differential privacy have also appeared in the literature (Chitra et al., 2022, Theorem 3).

We therefore state that parallel composition holds for (ϵ, δ)-probabilistic differentially private mechanisms, and we present what is, to the best of our knowledge, the first proof of this result in the Supplementary Material, Section C.

Lemma 1 (Parallel Composition for Probabilistic Differential Privacy)**.** Consider a database $\mathcal{D}$ partitioned into disjoint subsets $D_1, \ldots, D_m$, and suppose that there are privacy mechanisms $\mathcal{M}_1, \ldots, \mathcal{M}_m$, where $\mathcal{M}_i$ satisfies (ϵ_i, δ_i)-probabilistic differential privacy with respect to the adjacency relation in Definition 3. Then the release of $\mathcal{M}(\mathcal{D}) = (\mathcal{M}_1(D_1), \ldots, \mathcal{M}_m(D_m))$ provides $\mathcal{D}$ with conventional $(\max_{i\in[m]} \epsilon_i, \max_{i\in[m]} \delta_i)$-differential privacy with respect to the adjacency relation in Defintion 3. ♦

Lemma 1 relies on the fact that only a single entry is different between adjacent databases $\mathcal{D}$ and $\mathcal{D}'$. Given a collection of mechanisms $\{\mathcal{M}_i\}_{i\in[m]}$ where $\mathcal{M}_i$ is (ϵ_i, δ_i)-differentially private for all $i \in [m]$, their parallel composition generates outputs that are differentially private with privacy parameters no worse than those of the individual mechanisms, namely $\max_{i\in[m]} \epsilon_i$ and $\max_{i\in[m]} \delta_i$. Next, we show an example of how this parallel composition property fails when multiple entries may change between adjacent databases, which is the case in Gohari et al. (2021).

Example 2 (Failure of Parallel Composition)**.** Let the databases $\mathcal{D}$ and $\mathcal{D}'$ be collections of bits that differ in two entries, namely $\mathcal{D} = \{D_1, D_2, D_3\} = \{1, 1, 0\}$ and $\mathcal{D}' = \{D'_1, D'_2, D'_3\} = \{0, 0, 0\}$. Let $\mathcal{M}(\mathcal{D}) = (\mathcal{M}_1(D_1), \mathcal{M}_2(D_2), \mathcal{M}_3(D_3))$, where $\mathcal{M}_i$ is conventional $(\epsilon_i, 0)$-differentially private for all i. One such mechanism is

$$\mathcal{M}_i(D_i) = \begin{cases} D_i & \text{with probability } \frac{e^{\epsilon_i}}{1+e^{\epsilon_i}} \\ 1 - D_i & \text{with probability } \frac{1}{1+e^{\epsilon_i}}. \end{cases}$$

Let $T_i = \{D_i\}$ for all $i \in \{1, 2, 3\}$ and $T = T_1 \times T_2 \times T_3$. Then, we have

$$\mathbb{P}\left(\mathcal{M}(\mathcal{D}) \in T\right) = \frac{e^{\epsilon_1+\epsilon_2+\epsilon_3}}{(1+e^{\epsilon_1})(1+e^{\epsilon_2})(1+e^{\epsilon_3})} \tag{6}$$

$$\mathbb{P}\left(\mathcal{M}(\mathcal{D}') \in T\right) = \frac{e^{\epsilon_3}}{(1+e^{\epsilon_1})(1+e^{\epsilon_2})(1+e^{\epsilon_3})}. \tag{7}$$

For parallel composition, we seek to show that

$$\mathbb{P}\left(\mathcal{M}(\mathcal{D}) \in T\right) \leq \max_{i\in\{1,2,3\}} e^{\epsilon_i} \mathbb{P}\left(\mathcal{M}(\mathcal{D}') \in T\right). \tag{8}$$

Substituting (6) and (7) into (8) then yields the requirement that

$$e^{\epsilon_1+\epsilon_2+\epsilon_3} \leq \max_{i\in\{1,2,3\}} e^{\epsilon_i+\epsilon_3},$$

which is never satisfied for non-zero ϵ_1 and ϵ_2. △

Algorithm 2 Private Markov chain modeling

Inputs: $\mathcal{D}$, ρ, N_i, γ_i, η_i, k_i for $i \in [n]$
Outputs: Privacy-preserving Markov chain model $\tilde{P}$, privacy parameters ϵ, δ

1: **for all** $i \in [n]$ **do**
2: Compute $\tilde{C}^i, \epsilon_i, \delta_i$ with Algorithm 1
3: **end for**
4: Compute $\epsilon = \max_{i\in[n]} \epsilon_i$ and $\delta = \max_{i\in[n]} \delta_i$
5: Generate $\tilde{P}$ by setting row i equal to $\tilde{C}^i$ for $i \in [n]$

Example 2 shows that if adjacent databases can differ in more than one entry, then the parallel composition property may fail to hold. As a result, a naïve implementation of Definition 5 from Gohari et al. (2021) to rows of an $n \times n$ Markov chain transition matrix may not yield conventional $(\max_{i\in[n]} \epsilon_i, \max_{i\in[n]} \delta_i)$-differential privacy. However, Lemma 1 does guarantee the parallel composition property for the adjacency relation in Definition 3. Therefore, we use Definition 3 (and the accompanying result in Lemma 1) to preserve the privacy of $\mathcal{D}$ by privatizing the rows of $P(\mathcal{D}, \rho)$ individually to generate a private transition matrix $\tilde{P}$. That is, we compute

$$\tilde{P} = \left[\mathcal{M}_{\text{Dir}}^{(k_1)}(C(D_1, \rho)),\ \cdots,\ \mathcal{M}_{\text{Dir}}^{(k_n)}(C(D_n, \rho))\right]^T. \tag{9}$$

We now state the privacy guarantee of computing $\tilde{P}$.

Theorem 3 (Solution to Problem 3)**.** *Fix a database $\mathcal{D} = \{D_1, D_2, \ldots, D_n\}$, a category set ρ such that $|\rho| = n \in \mathbb{N}$, and Dirichlet parameters $k_i > 0$ for $i \in [n]$. Let $P(\mathcal{D}, \rho) \in \mathbb{S}_n$ be defined as in* (2)*. Let Assumptions 1-4 hold. Then computing $\tilde{P} \in \mathbb{S}_n$ using* (9) *keeps $\mathcal{D}$ $(\max_{i\in[n]} \epsilon_i, \max_{i\in[n]} \delta_i)$-differentially private, where ϵ_i is from* (3) *and δ_i is from* (4)*.*

Algorithm 2 presents our unified framework for computing a privacy-preserving Markov chain model.

5 Stationary Distribution and Convergence Rate

In this section we solve Problem 4. At each time t, there exists a probability distribution over the state space of the Markov chain, which we denote $\mu_t \in \Delta_n$. Let $\mu_0 \in \Delta_n$ denote the initial distribution. The distribution μ_t can be computed by right-multiplying by the transition matrix G to find $\mu_t^T = \mu_{t-1}^T G$.

We analyze the steady-state behavior of a Markov chain by analyzing the limiting behavior of μ_t. When a Markov chain is finite, irreducible, and homogeneous, $\lim_{t\to\infty} \mu_t$ exists and is called the "stationary distribution", which we denote by π. The stationary distribution π satisfies $\pi^T = \pi^T G$. A Markov chain in which μ_t converges to π is called "ergodic". We ensure $P(\mathcal{D}, \rho)$ is ergodic via Assumptions 5 and 6 below. First we establish a key definition.

To analyze the transient behavior of a Markov chain, it is common to analyze the convergence rate to π. Many metrics exist for doing so, and we consider the "∞-norm ergodicity coefficient" of ergodic Markov chains. See Ipsen & Selee (2011) for a detailed account of ergodicity coefficients.

Definition 7 (∞-norm Ergodicity Coefficient; (Ipsen & Selee, 2011))**.** Let $G \in \mathbb{S}_n$ be the transition matrix of an ergodic Markov chain with stationary distribution π. Then the ∞-norm ergodicity coefficient of the Markov chain is defined as $\tau_\infty(G) = \max_{\|z\|_\infty = 1, z^T \mathbb{1} = 0} \left\|G^T z\right\|_\infty$. ▲

Let the eigenvalue moduli of $G \in \mathbb{S}_n$ be $1 = |\lambda_1| \geq \ldots \geq |\lambda_n|$. The eigenvalue modulus $|\lambda_2|$ governs the rate of convergence of G to its stationary distribution if G is ergodic. Definition 7 provides a bound on $|\lambda_2|$ when it is difficult to compute exactly.

To solve Problem 4, we compute the total variation distance in the stationary distribution between a private Markov chain $\tilde{P}$ and its non-private form $P(\mathcal{D}, \rho)$, as well as the absolute difference in their ∞-norm ergodicity coefficients. For a Markov chain to converge to its stationary distribution, we require it to be (i) finite, (ii) irreducible, and (iii) homogeneous. Condition (i) holds trivially since the state space has the same cardinality

as ρ, which is finite. Condition (iii) holds since the database is not a function of time. Let the state space of the Markov chain model be $\mathcal{Y}$, and note that each category $\rho_i \in \rho$ corresponds to a state $y_i \in \mathcal{Y}$. We apply the following assumptions on $\mathcal{Y}$ and $\mathcal{D}$ to ensure condition (ii) holds:

Assumption 5. Each state $y_i \in \mathcal{Y}$ is reachable from any other state $y_j \in \mathcal{Y}$ in finite time.

Assumption 6. In each D_i, every feasible transition has occurred at least once.

Returning to the traffic scenario in Example 1, Assumption 5 enforces that all roads whose data are in $\mathcal{D}$ can be reached in finite time from any other road with data in $\mathcal{D}$. Assumption 6 enforces that every possible turn on the roads with data in $\mathcal{D}$ has been taken at least once.

The next two theorems both rely on a bound on $\|P(\mathcal{D}, \rho) - \tilde{P}\|_1$, which results in them having some terms in common. See the Supplementary Material, Section B for the derivation of a bound on $\|P(\mathcal{D}, \rho) - \tilde{P}\|_1$. Given a database and a category set that satisfy Assumptions 5 and 6, we first bound the change in the stationary distribution, thus solving the first part of Problem 4.

Theorem 4 (Solution to Problem 4, Part 1)**.** *Let $\tilde{\pi}$ denote the stationary distribution of $\tilde{P}$. Let Assumptions 1-6 hold. Then $P(D, \rho)$ converges to a stationary distribution π. Let $Z = (I - P(\mathcal{D}, \rho) - \mathbb{1}_n \pi^T)^{-1}$ and let $\zeta_i(x) = \log\left(\frac{x+1}{N_i}\right) - \psi\left(\frac{(x+1)k_i}{N_i}\right)$ for all $i \in [n]$. Then we have $\mathbb{E}\left[\|\pi - \tilde{\pi}\|_{TV}\right] \leq \frac{1}{2}\|Z\|_1 \sqrt{2L}$, where $L = \sum_{i=1}^{n} \pi(i)\left(\frac{n-1}{N_i}\zeta_i(0) + \frac{N_i - n + 1}{N_i}\zeta_i(N_i - n) + \psi(k_i)\right)$.*

Theorem 4 provides a tool for users to understand the quantitative change in the stationary distribution as both the number of individuals in the database and the strength of privacy change. Since the digamma function $\psi(x)$ grows as $\mathcal{O}(\log(x))$, Theorem 4 shows that the error in the stationary distribution is $\mathcal{O}(\log(k_i^{-1}))$, which implies that accuracy can be substantially improved with small changes in privacy. Similarly, we develop a utility guarantee on the transient behavior (and solve the remaining part of Problem 4) by bounding the change in the ergodicity coefficient.

Theorem 5 (Solution to Problem 4, Part 2)**.** *Let the conditions of Theorem 4 hold. Let $\tau_\infty(P(\mathcal{D}, \rho))$ be the ergodicity coefficient of $P(\mathcal{D}, \rho)$ in the sense of Definition 7 and let $\tau_\infty(\tilde{P})$ be the ergodicity coefficient of $\tilde{P}$. Then $\mathbb{E}\left[|\tau_\infty(P(\mathcal{D}, \rho)) - \tau_\infty(\tilde{P})|\right] \leq \sqrt{2L}$.*

Theorem 5, similar to Theorem 4, also provides users with a tool to understand how privacy affects the privatized Markov chain transitions and also behaves as $\mathcal{O}(\log(k_i^{-1}))$. Combined, Theorems 4 and 5 provide insight into how the rate of convergence is affected by privacy and into the final error in the stationary distribution.

6 Numerical Simulations

In this section we apply Algorithm 1 to a database of grades for an undergraduate course and Algorithm 2 to a database of taxi trips in New York City.

6.1 Class Grade Distribution

Many universities publish the distribution of grades for classes at the end of every semester, along with the instructor's name, to help students assess the relative difficulty of a course with a particular instructor. However, given side information about a student's performance on individual assignments, observers such as classmates may be able to make inferences about that student's final grade from this published distribution, which indicates a need for privacy.

Using publicly available (Georgia Tech, 2025) grade distributions $\mathcal{D}$ and class sizes N, we compute privatized forms of the grade distribution from a Calculus 1 course with $N = 98$ students. The categories are $\rho = \{\texttt{A}, \texttt{B}, \texttt{C}, \texttt{D}, \texttt{F}\}$, and we take $\eta = 0.073$ and $\gamma = 0.0004$. We simulate $2,000$ privatized forms for each $k \in \{20.6, 21.6, \ldots, 110.6\}$. The smallest value of k gives the strongest allowable privacy according to Assumption 4. Figure 1 shows the average KL divergence between the private and non-private distributions with the strength of privacy ϵ corresponding to each k. We see that as privacy weakens (i.e., as ϵ grows), the bound in Corollary 1 approaches the true average KL divergence. Additionally, even at the smallest ϵ value,

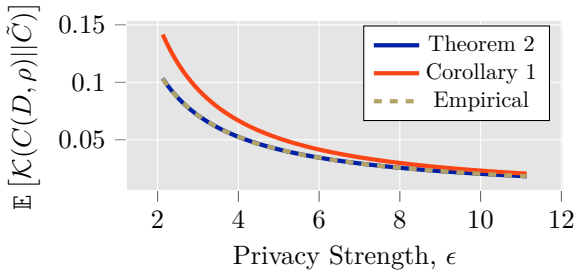

Figure 1: KL divergence between a privatized grade distribution and its true grade distribution. The upper bound in Corollary 1 becomes increasingly tight as privacy weakens, though it is close to the true value for all ϵ. The distributions of $\tilde{C}$ and $C(D, \rho)$ remain similar even under strong privacy, highlighting that Algorithm 1 produces outputs with strong privacy that exhibit high accuracy.

$\epsilon = 2.255$, the private distributions remain accurate, with a KL divergence of 0.103. At this smallest ϵ value, we find that $\delta = 0.0026$, and we can provide $(2.255, 0.0026)$-differential privacy to this grade distribution. We may also achieve smaller δ values at larger ϵ values, e.g., at $\epsilon = 3.31$, we obtain $\delta = 1.3 \times 10^{-4}$. Thus, Algorithm 1 produces private distributions with high similarity to the original, non-private distribution, which highlights the applicability to this scenario.

6.2 New York City Taxis

We apply Algorithm 2 to a Markov chain model of New York City Yellow Taxi trips. Each state in the Markov chain represents a pickup/drop-off zone, and the dataset in NYC Taxi Limousine Commision (2025) contains the pickup and drop-off locations of taxi trips in January 2025. We empirically compute the transition probabilities by counting the number of trips from a given pickup location to a corresponding drop-off location. Here, ρ_i is drop-off location i, with $|\rho| = 40$, and D_j contains all departures from pickup location j. We have $d_k^j \models \rho_i$ if taxi trip k begins in pickup location j and drops off in drop-off location i. See Appendix E for discussion on how the taxi data was pre-processed for this example.

To simulate $1,000$ privatizations of the Markov chain, we find $k_{min,i}$ via Assumption 4 for every $i \in [n]$, where n is the number of pickup/drop-off locations. Then we compute a set of k_i values by incrementing k_i by $\frac{1}{10}k_{min,i}$ up to $3k_{min,i}$ for all $i \in [n]$. Additionally, we set $\gamma_i = 10^{-8}$ for all $i \in [n]$. We then find the corresponding strength of privacy ϵ at each increment according to Algorithm 2 for the entire transition matrix. We find the strongest possible privacy to be $(2.68, 0.0214)$-differential privacy. Figure 2 illustrates the stationary distribution in the borough of Manhattan, where there were $N = 2,933,898$ taxi trips, under different values of ϵ. As privacy weakens (i.e., as ϵ grows), the stationary distribution approaches that of the original, non-private Markov chain. This is demonstrated numerically in Figure 3, which shows the expected TV distance between the stationary distributions with varying privacy strength. We find that even at the strongest possible privacy, $\epsilon = 3.73$ and $\delta = 3 \times 10^{-6}$, the change induced by privacy on average is only $\mathbb{E}\left[\|\pi - \tilde{\pi}\|_{TV}\right] = 0.017$. We may attain even smaller δ values by increasing the value of ϵ, e.g., we find $\delta = 5 \times 10^{-9}$ at $\epsilon = 5.19$. Thus, we find that Algorithm 2 induces only small differences in the stationary distribution while provably protecting the data that is used to generate the underlying Markov chain model.

7 Conclusion

We have presented a framework privatizing databases that are used to compute both stochastic vectors and stochastic matrices. We proved that this approach is differentially private and provided utility guarantees that allow users to trade-off privacy, convergence rates, and accuracy of stationary distributions. Results show with $(3.73, 3 \times 10^{-6})$-differential privacy, our approach induces less than 2% error in the stationary distribution on a large dataset ($N = 2,933,898$). Future work will address user-level privacy for data-driven Markov chain modeling and privacy for data-driven Markov decision processes.

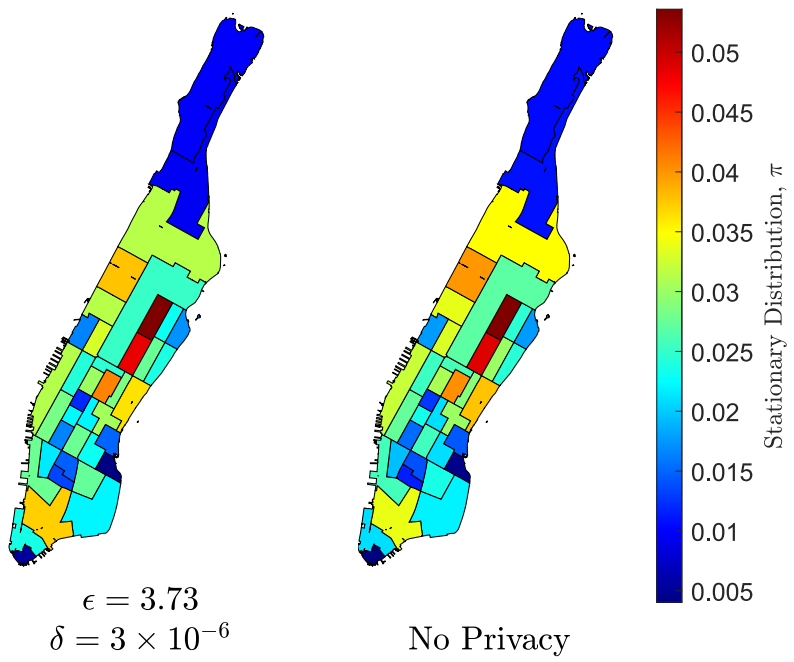

Figure 2: Stationary distribution of the Markov chain formed by New York City taxi drop-offs and pickups. At the strongest privacy level of $\epsilon = 3.73$, the stationary distribution of the privatized Markov chain model remains close to that of the non-private model.

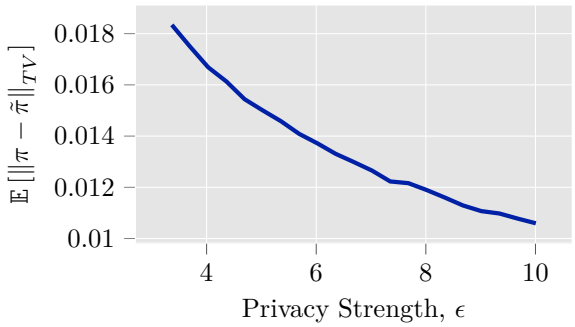

Figure 3: Change in the stationary distribution with varying privacy strength. Even with $\epsilon = 3.73$, we find the average TV distance change in the stationary distribution is minimal.

References

Naman Agarwal, Peter Kairouz, and Ziyu Liu. The Skellam mechanism for differentially private federated learning. *Advances in Neural Information Processing Systems*, 34:5052–5064, 2021.

Alexander Benvenuti, Brendan Bialy, Miriam Dennis, and Matthew Hale. Guaranteed feasibility in differentially private linearly constrained convex optimization. *IEEE Control Systems Letters*, 8:2745–2750, 2024a.

Alexander Benvenuti, Calvin Hawkins, Brandon Fallin, Bo Chen, Brendan Bialy, Miriam Dennis, and Matthew Hale. Differentially private reward functions for Markov decision processes. In *2024 IEEE Conference on Control Technology and Applications (CCTA)*, pp. 631–636. IEEE, 2024b.

Alberto Blanco-Justicia, David Sánchez, Josep Domingo-Ferrer, and Krishnamurty Muralidhar. A critical review on the use (and misuse) of differential privacy in machine learning. *ACM Computing Surveys*, 55 (8):1–16, 2022.

Bo Chen, Calvin Hawkins, Mustafa O Karabag, Cyrus Neary, Matthew Hale, and Ufuk Topcu. Differential privacy in cooperative multiagent planning. In *Uncertainty in Artificial Intelligence*, pp. 347–357. PMLR, 2023a.

Bo Chen, Kevin Leahy, Austin Jones, and Matthew Hale. Differential privacy for symbolic systems with application to Markov chains. *Automatica*, 152:110908, 2023b.

Wei-Ning Chen, Christopher A Choquette Choo, Peter Kairouz, and Ananda Theertha Suresh. The fundamental price of secure aggregation in differentially private federated learning. In *International Conference on Machine Learning*, pp. 3056–3089. PMLR, 2022.

Tarun Chitra, Guillermo Angeris, and Alex Evans. Differential privacy in constant function market makers. In *International Conference on Financial Cryptography and Data Security*, pp. 149–178. Springer, 2022.

Jie Dong, Gui Li, Wenkai Ma, and Jianshun Liu. Personalized recommendation system based on social tags in the era of internet of things. *Journal of Intelligent Systems*, 31(1):681–689, 2022.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pp. 265–284. Springer, 2006.

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

Brandon Fallin, Calvin Hawkins, Bo Chen, Parham Gohari, Alexander Benvenuti, Ufuk Topcu, and Matthew Hale. Differential privacy for stochastic matrices using the matrix Dirichlet mechanism. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pp. 5067–5072. IEEE, 2023.

Georgia Tech. Course critique. `https://lite.gatech.edu/lite_script/dashboards/grade_distribution.html`, 2025.

Robin C Geyer, Tassilo Klein, and Moin Nabi. Differentially private federated learning: A client level perspective. *arXiv preprint arXiv:1712.07557*, 2017.

Parham Gohari, Matthew Hale, and Ufuk Topcu. Privacy-preserving policy synthesis in Markov decision processes. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 6266–6271. IEEE, 2020.

Parham Gohari, Bo Wu, Calvin Hawkins, Matthew Hale, and Ufuk Topcu. Differential privacy on the unit simplex via the Dirichlet mechanism. *IEEE Transactions on Information Forensics and Security*, 16: 2326–2340, 2021.

Qiuming Gong, Shawn Midlam-Mohler, Vincenzo Marano, and Giorgio Rizzoni. An iterative Markov chain approach for generating vehicle driving cycles. *SAE International Journal of Engines*, 4(1):1035–1045, 2011.

Justin Hsu, Marco Gaboardi, Andreas Haeberlen, Sanjeev Khanna, Arjun Narayan, Benjamin C Pierce, and Aaron Roth. Differential privacy: An economic method for choosing epsilon. In *2014 IEEE 27th Computer Security Foundations Symposium*, pp. 398–410. IEEE, 2014.

Ilse CF Ipsen and Teresa M Selee. Ergodicity coefficients defined by vector norms. *SIAM Journal on Matrix Analysis and Applications*, 32(1):153–200, 2011.

Bargav Jayaraman and David Evans. Evaluating differentially private machine learning in practice. In *28th USENIX security symposium (USENIX security 19)*, pp. 1895–1912, 2019.

Samuel Kotz, Narayanaswamy Balakrishnan, and Norman L Johnson. *Continuous multivariate distributions, Volume 1: Models and applications*, volume 1. John Wiley & Sons, 2019.

Jens Lundström, Eric Järpe, and Antanas Verikas. Detecting and exploring deviating behaviour of smart home residents. *Expert Systems with Applications*, 55:429–440, 2016.

Ashwin Machanavajjhala, Daniel Kifer, John Abowd, Johannes Gehrke, and Lars Vilhuber. Privacy: Theory meets practice on the map. In *2008 IEEE 24th International Conference on Data Engineering*, pp. 277–286. IEEE, 2008.

Frank D McSherry. Privacy integrated queries: an extensible platform for privacy-preserving data analysis. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of Data*, pp. 19–30, 2009.

Sebastian Meiser. Approximate and probabilistic differential privacy definitions. *Cryptology ePrint Archive*, 2018.

Andres Munoz, Umar Syed, Sergei Vassilvtiskii, and Ellen Vitercik. Private optimization without constraint violations. In *International Conference on Artificial Intelligence and Statistics*, pp. 2557–2565. PMLR, 2021.

Maxence Noble, Aurélien Bellet, and Aymeric Dieuleveut. Differentially private federated learning on heterogeneous data. In *International Conference on Artificial Intelligence and Statistics*, pp. 10110–10145. PMLR, 2022.

James R Norris. *Markov chains*. Cambridge University Press, 1998.

NYC Taxi Limousine Commision. January 2025 Yellow taxi trip record data. `https://www.nyc.gov/site/tlc/about/tlc-trip-record-data.page`, 2025.

Donlapark Ponnoprat. Dirichlet mechanism for differentially private KL divergence minimization. *Transactions on Machine Learning Research*, 2023.

Natalia Ponomareva et al. How to DP-fy ML: A practical guide to machine learning with differential privacy. *Journal of Artificial Intelligence Research*, 77:1113–1201, 2023.

Feng Qi, Senlin Guo, and Bai-Ni Guo. Complete monotonicity of some functions involving polygamma functions. *Journal of Computational and Applied Mathematics*, 233(9):2149–2160, 2010.

Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. Factorizing personalized Markov chains for next-basket recommendation. In *Proceedings of the 19th International Conference on World Wide Web*, pp. 811–820, 2010.

E Seneta. Sensitivity to perturbation of the stationary distribution: Some refinements. *Linear Algebra and its Applications*, 108:121–126, 1988.

Youjia Wang and Michael CH Choi. Information divergences of Markov chains and their applications. *arXiv preprint arXiv:2312.04863*, 2023.

Yu Wang, Matthew Hale, Magnus Egerstedt, and Geir E Dullerud. Differentially private objective functions in distributed cloud-based optimization. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 3688–3694. IEEE, 2016.

William C Waterhouse. Do symmetric problems have symmetric solutions? *The American Mathematical Monthly*, 90(6):378–387, 1983.

Joakim Widén, Annica M Nilsson, and Ewa Wäckelgård. A combined Markov-chain and bottom-up approach to modelling of domestic lighting demand. *Energy and Buildings*, 41(10):1001–1012, 2009.

A Computing δ

In this section, we detail our approach for computing δ when $\mathcal{D} = \{D\}$. To enforce (ϵ, δ)-probabilistic differential privacy in the sense of Definition 6, we partition the output space of the Dirichlet mechanism into two sets in via

$$\Omega_1 = \{x \in \Delta_n \mid x_i \geq \gamma \text{ for all } i \in [n]\}$$

and

$$\Omega_2 = \{x \in \Delta_n \mid x \notin \Omega_1\},$$

as in Section 3. We state the following lemma on the computation of δ.

Lemma 2 (Computation of δ; (Gohari et al., 2021))**.** Fix a database D, a category set ρ such that $|\rho| = n \in \mathbb{N}$, a Dirichlet parameter $k > 0$, and a bordered simplex parameter $\eta \in (0, 1)$. Let $C(D, \rho) \in \Delta_n^{(\eta)}$ be the count vector on the database D with category set ρ. Let Assumptions 1 and 2 hold. Define

$$\mathcal{A}_n = \left\{ x \in \mathbb{R}^{n-1} \mid \sum_{i \in [n-1]} x_i \leq 1, x_i \geq \gamma \text{ for all } i \in [n-1] \right\}.$$

Then for a Dirichlet mechanism with parameter $k > 0$, we have

$$\mathbb{P}\left(\mathcal{M}_{\text{Dir}}^{(k)}(C(D,\rho)) \in \Omega_1\right) = \frac{\int_{\mathcal{A}_n} \prod_{i \in [n-1]} x_i^{kC(D,\rho)_i - 1} \left(1 - \sum_{i \in [n-1]} x_i\right)^{k(1 - \sum_{i \in [n-1]} C(D,\rho)_i) - 1} \prod_{i \in [n-1]} dx_i}{B(kC(D,\rho))}.$$

♦

Since $C(D, \rho)$ appears in this computation of δ, this form of δ cannot be used to guarantee differential privacy since the value of δ must hold for all adjacent pairs of input data. However, we may minimize $\mathbb{P}\left(\mathcal{M}_{\text{Dir}}^{(k)}(C(D,\rho)) \in \Omega_1\right)$ over all $C(D, \rho) \in \Delta_n^{(\eta)}$ to find an upper bound on δ which holds for all adjacent input data. That is, we find

$$\mathbb{P}\left(\mathcal{M}_{\text{Dir}}^{(k)}(C(D,\rho)) \in \Omega_2\right) = 1 - \mathbb{P}\left(\mathcal{M}_{\text{Dir}}^{(k)}(C(D,\rho)) \in \Omega_1\right) \leq 1 - \min_{C(D,\rho) \in \mathcal{C}(\rho)} \mathbb{P}\left(\mathcal{M}_{\text{Dir}}^{(k)}(C(D,\rho)) \in \Omega_1\right) = \delta.$$

Next, we state a result showing that computation of $\min_{C(D,\rho) \in \mathcal{C}(\rho)} \mathbb{P}\left(\mathcal{M}_{\text{Dir}}^{(k)}(C(D,\rho)) \in \Omega_1\right)$ can be performed efficiently.

Lemma 3 (δ log-concavity; (Gohari et al., 2021))**.** Let all of the conditions from Lemma 2 hold. Then $\mathbb{P}\left(\mathcal{M}_{\text{Dir}}^{(k)}(C(D,\rho)) \in \Omega_1\right)$ is a log-concave function of $C(D, \rho)$ over the domain $\Delta_n^{(\eta)}$. ♦

Lemma 3 implies that the minimizer lies on the extreme points of the feasible region. As noted in (Gohari et al., 2021), Lemma 3 then implies that the time complexity of computing δ is $\mathcal{O}((n-1)^2)$. Thus, it remains tractable to compute δ in this fashion. However, as $n - 1$ increases, the objective function in (4) becomes increasingly computationally intensive to compute because it requires the evaluation of an increasing number of nested integrals. Many publicly available numerical integration libraries do not support calculations with an arbitrarily large number of nested integrals. For example, `MATLAB` only supports up to 3 integrals (though publicly shared frameworks exist supporting up to 6). As a result, we take a different approach to computing δ when $n - 1$ is large.

To address these issues, we introduce a Monte-Carlo approach to compute δ. Lemma 3 shows that we need only evaluate the objective function at the extreme points of the feasible region. Let $\{q_1, \ldots, q_{n-1}\}$ denote the set of extreme points of the feasible region. Using a Monte-Carlo approach, we approximate the probability in (4) at each of the extreme points q_i by drawing samples of $\mathcal{M}_{\text{Dir}}^{(k)}(q_i)$ for all $i \in [n-1]$ and evaluating if each sample is in the set Ω_1. The fraction of samples of $\mathcal{M}_{\text{Dir}}^{(k)}(q_i)$ in Ω_1 is then an approximation for $\mathbb{P}(\mathcal{M}_{\text{Dir}}^{(k)}(q_i) \in \Omega_1)$. We substitute the smallest approximated probability into (4) to approximate δ. We

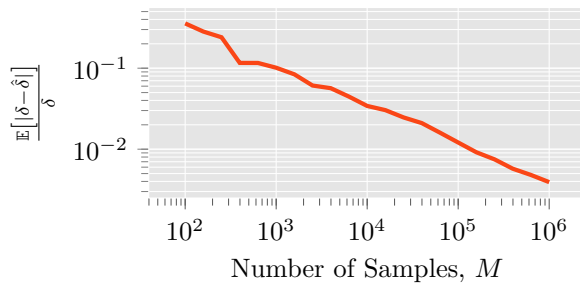

Figure 4: Average relative error between δ and $\hat{\delta}$.

denote this approximated value of δ by $\hat{\delta}$. The accuracy of this approach depends on the number of samples M used in the Monte Carlo simulation. Figure 4 shows how the relative error between δ and $\hat{\delta}$ decreases with the number of samples. We see that with 10^6 samples, $\hat{\delta}$ has less than 1% error, highlighting the practicality of this approach for estimating δ.

B Additional Technical Lemmas

The omitted proofs in Section C make use of the lemmas and definitions in this section. In addition to the special functions presented in Section 1, we make use of the following in the subsequent definitions, lemmas, and proofs:

$$\psi^{(1)}(x) = \psi'(x) = \sum_{n=0}^{\infty} \frac{1}{(x+n)^2}$$
$$\text{Beta}(x;\alpha,\beta) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{\text{Beta}(\alpha,\beta)} \text{ with } \alpha > 0, \beta > 0,$$

where $\psi^{(1)}$ denotes the trigamma function, and $\text{Beta}(x;\alpha_1,\beta)$ denotes the probability density function of the Beta distribution.

First, we define the KL divergence for Markov chains.

Definition 8 (Markov chain KL divergence)**.** Let $L, Q \in \mathbb{S}_n$ be two stochastic matrices with a common state space $\mathcal{X}$. The Kullback-Leibler (KL) divergence from L to Q with respect to a probability distribution π on $\mathcal{X}$ is

$$\mathcal{K}(Q||L) = \sum_{i=1}^{n} \pi(i) \sum_{j=1}^{n} Q_{ij} \log\left(\frac{Q_{ij}}{L_{ij}}\right),$$

where we use the convention $0 \log\left(\frac{0}{0}\right) = 0$. ▲

Lemma 4 (Markov chain divergence bound)**.** Fix Dirichlet parameters $k_i > 0$ for $i \in [n]$. Fix a database $\mathcal{D} = \{D_1, D_2, \ldots, D_n\}$ where each D_i has $N_i \in \mathbb{N}$ entries, and category set ρ such that $|\rho| = n \in \mathbb{N}$. Let $P(\mathcal{D}, \rho)$ be defined as in (2). Let Assumptions 1-6 hold so that $P(\mathcal{D}, \rho)$ is a transition matrix of a finite, irreducible, aperiodic, time-homogeneous Markov chain. Let $\tilde{P}$ denote the privatized form of $P(\mathcal{D}, \rho)$ and let π denote the stationary distribution of $P(\mathcal{D}, \rho)$. Let $\mathcal{K}(\cdot||\cdot)$ be defined as in Definition 8. Then

$$\mathbb{E}\left[\mathcal{K}(P(D,\rho)||\tilde{P})\right] = \sum_{i=1}^{n} \pi(i) \sum_{j=1}^{n} P(\mathcal{D},\rho)_{ij}\big(\log(P(\mathcal{D},\rho)_{ij}) + \psi(k_i) - \psi(k_i P(\mathcal{D},\rho)_{ij})\big).$$

Additionally, we have that

$$\begin{aligned}&\mathbb{E}\left[\mathcal{K}(P(D,\rho)||\tilde{P})\right]\\ &\leq \sum_{i=1}^{n} \pi(i) \left(\frac{n-1}{N_i}\left(\log\left(\frac{1}{N_i}\right) - \psi\left(\frac{k_i}{N_i}\right)\right) + \frac{N_i - n + 1}{N_i}\left(\log\left(\frac{N_i - n + 1}{N_i}\right) - \psi\left(\frac{(N_i - n + 1)k_i}{N_i}\right)\right) + \psi(k_i)\right).\end{aligned} \tag{10}$$

♦

Proof. From Definition 8, we have

$$\mathcal{K}(P(\mathcal{D},\rho)||\tilde{P}) = \sum_{i=1}^{n}\pi(i)\sum_{j=1}^{n}P(\mathcal{D},\rho)_{ij}\log\left(\frac{P(\mathcal{D},\rho)_{ij}}{\tilde{P}_{ij}}\right).$$

Taking the expectation and exploiting linearity yields

$$\mathbb{E}\left[\mathcal{K}(P(\mathcal{D},\rho)||\tilde{P})\right] = \sum_{i=1}^{n}\pi(i)\mathbb{E}\left[\sum_{j=1}^{n}P(\mathcal{D},\rho)_{ij}\log\left(\frac{P(\mathcal{D},\rho)_{ij}}{\tilde{P}_{ij}}\right)\right]. \tag{11}$$

The expression

$$\mathbb{E}\left[\sum_{j=1}^{n}P(\mathcal{D},\rho)_{ij}\log\left(\frac{P(\mathcal{D},\rho)_{ij}}{\tilde{P}_{ij}}\right)\right]$$

is the expected KL divergence from $P(\mathcal{D},\rho)_i$, the ith row of $P(\mathcal{D},\rho)$, to $\tilde{P}_i$, i.e., $\mathbb{E}\left[\mathcal{K}(P(\mathcal{D},\rho)_i||\tilde{P}_i)\right]$. Substituting $\mathbb{E}\left[\mathcal{K}(P(\mathcal{D},\rho)_i||\tilde{P}_i)\right]$ into (11) gives

$$\mathbb{E}\left[\mathcal{K}(P(\mathcal{D},\rho)||\tilde{P})\right] = \sum_{i=1}^{n}\pi(i)\mathbb{E}\left[\mathcal{K}(P(\mathcal{D},\rho)_i||\tilde{P}_i)\right].$$

We note that $\tilde{P}_i$ is a Dirichlet-distributed random vector, and thus substituting (5) from Theorem 2 gives

$$\mathbb{E}\left[\mathcal{K}(P(\mathcal{D},\rho)||\tilde{P})\right] = \sum_{i=1}^{n}\pi(i)\sum_{j=1}^{n}P(\mathcal{D},\rho)_{ij}\big(\log(P(\mathcal{D},\rho)_{ij}) + \psi(k_i) - \psi(k_iP(\mathcal{D},\rho)_{ij})\big),$$

as desired.

Alternatively, for a bound that does not depend on the sensitive data $\mathcal{D}$, we apply the upper bound on $C(D,\rho)_i\mathbb{E}\left[\log\left(\frac{C(\mathcal{D},\rho)_i}{\tilde{C}_i}\right)\right]$ from (29) in the proof of Corollary 1 to obtain

$$\begin{aligned}&\mathbb{E}\left[\mathcal{K}(P(\mathcal{D},\rho)_i||\tilde{P}_i)\right]\\ &\quad\leq \frac{n-1}{N_i}\left(\log\left(\frac{1}{N_i}\right) - \psi\left(\frac{k_i}{N_i}\right)\right) + \frac{N_i-n+1}{N_i}\left(\log\left(\frac{N_i-n+1}{N_i}\right) - \psi\left(\frac{(N_i-n+1)k_i}{N_i}\right)\right) + \psi(k_i),\end{aligned}$$

which we substitute into (11) to obtain

$$\begin{aligned}&\mathbb{E}\left[\mathcal{K}(P(\mathcal{D},\rho)||\tilde{P})\right]\\ &\leq \sum_{i=1}^{n}\pi(i)\left(\frac{n-1}{N_i}\left(\log\left(\frac{1}{N_i}\right) - \psi\left(\frac{k_i}{N_i}\right)\right) + \frac{N_i-n+1}{N_i}\left(\log\left(\frac{N_i-n+1}{N_i}\right) - \psi\left(\frac{(N_i-n+1)k_i}{N_i}\right)\right) + \psi(k_i)\right),\end{aligned}$$

which completes the proof. □

Lemma 5 (Stationary Distribution Perturbation Bound; (Seneta, 1988))**.** Let P be the transition matrix of a finite, homogeneous, irreducible Markov chain with stationary distribution π, and let $\tilde{P}$ be a perturbed form of P. Let $\tilde{\pi}$ be the stationary distribution of $\tilde{P}$. Then,

$$\|\pi - \tilde{\pi}\|_1 \leq \|Z\|_1 \left\|P - \tilde{P}\right\|_1,$$

where $Z = (I - P - \mathbb{1}_n\pi^T)^{-1}$. ♦

Lemma 6 (1-norm Accuracy Bound)**.** Let the conditions of Lemma 4 hold. Then

$$\mathbb{E}\left[\left\|P-\tilde{P}\right\|_1\right] \leq \sqrt{2}\cdot\left(\sum_{i=1}^{n}\pi(i)\left(\frac{n-1}{N_i}\left(\log\left(\frac{1}{N_i}\right)-\psi\left(\frac{k_i}{N_i}\right)\right)+\right.\right.$$
$$\left.\left.\frac{N_i-n+1}{N_i}\left(\log\left(\frac{N_i-n+1}{N_i}\right)-\psi\left(\frac{(N_i-n+1)k_i}{N_i}\right)\right)+\psi(k_i)\right)\right)^{\frac{1}{2}}. \quad (12)$$

♦

Proof. From Pinsker's inequality (which also holds for matrix-valued f-divergences (Wang & Choi, 2023)) it follows that

$$\left\|P-\tilde{P}\right\|_1 \leq \sqrt{2}\sqrt{\mathcal{K}(P(\mathcal{D},\rho)||\tilde{P})}.$$

Taking the expectation on both sides gives

$$\mathbb{E}\left[\left\|P-\tilde{P}\right\|_1\right] \leq \sqrt{2}\mathbb{E}\left[\sqrt{\mathcal{K}(P(\mathcal{D},\rho)||\tilde{P})}\right] \leq \sqrt{2}\sqrt{\mathbb{E}\left[\mathcal{K}(P(\mathcal{D},\rho)||\tilde{P})\right]},$$

where the second inequality follows from Jensen's inequality and the fact that the square root function is concave. Substituting (10) from Lemma 4 into the above bound gives

$$\mathbb{E}\left[\left\|P-\tilde{P}\right\|_1\right] \leq \sqrt{2}\cdot$$
$$\left(\sum_{i=1}^{n}\pi(i)\left(\frac{n-1}{N_i}\left(\log\left(\frac{1}{N_i}\right)-\psi\left(\frac{k_i}{N_i}\right)\right)+\frac{N_i-n+1}{N_i}\left(\log\left(\frac{N_i-n+1}{N_i}\right)-\psi\left(\frac{(N_i-n+1)k_i}{N_i}\right)\right)+\psi(k_i)\right)\right)^{\frac{1}{2}},$$

as desired. □

Lemma 7 (Ratio of Beta Functions; Gohari et al. (2021))**.** Let $C(D,\rho)$ and $C(D',\rho)$ be two adjacent count vectors in the sense of Definition 3, and let i and j be the indices where they differ. Then, for a constant $k>0$, we have that

$$\frac{\mathrm{Beta}(kC(D',\rho)_i, kC(D',\rho)_j)}{\mathrm{Beta}(kC(D,\rho)_i, kC(D,\rho)_j)} \leq \frac{\mathrm{Beta}(kC(D',\rho)_i, k(1-\eta-C(D',\rho)_i)}{\mathrm{Beta}(kC(D,\rho)_i, k(1-\eta-C(D,\rho)_i)}.$$

♦

C Omitted Proofs

C.1 Proof of Theorem 1

We seek to bound

$$\log\left(\frac{\mathbb{P}\left(\mathcal{M}_{\mathrm{Dir}}^{(k)}(C(D,\rho))=S\right)}{\mathbb{P}\left(\mathcal{M}_{\mathrm{Dir}}^{(k)}(C(D',\rho))=S\right)}\right). \quad (13)$$

The adjacency relation in Definition 3 implies that adjacent databases D and D' have the same number of entries, one of which differs between them. Thus, there exists $d\in D$ and $d'\in D'$ such that $d\neq d'$ and $D\setminus\{d\}=D'\setminus\{d'\}$. Under Assumption 1, $d\models\rho_i$ for exactly one i, and $d'\models\rho_j$ for exactly one j. If $i=j$, then $C(D,\rho)=C(D',\rho)$ and the term in (13) is 0. Suppose $i\neq j$. Then $C(D,\rho)\neq C(D',\rho)$ and in particular they differ in two entries. Without loss of generality, consider $i,j\in[n]$ as the indices of the entries in which $C(D,\rho)$ and $C(D',\rho)$ differ. From the definition of the Dirichlet mechanism in Definition 5,

we have

$$\log\left(\frac{\mathbb{P}\left(\mathcal{M}_{\mathrm{Dir}}^{(k)}(C(D,\rho))=S\right)}{\mathbb{P}\left(\mathcal{M}_{\mathrm{Dir}}^{(k)}(C(D',\rho))=S\right)}\right)=\log\left(\frac{B(kC(D',\rho))}{B(kC(D,\rho))}\prod_{i=1}^{n}\frac{x_i^{kC(D,\rho)_i-1}}{x_i^{kC(D',\rho)_i-1}}\right)$$
$$=\log\left(\frac{\Gamma\left(kC(D',\rho)_i\right)\Gamma\left(kC(D',\rho)_j\right)x_i^{kC(D,\rho)_i-1}x_j^{kC(D,\rho)_j-1}}{\Gamma\left(kC(D,\rho)_i\right)\Gamma\left(kC(D,\rho)_j\right)x_i^{kC(D',\rho)_i-1}x_j^{kC(D',\rho)_j-1}}\right)$$
$$=\log\left(\frac{\Gamma\left(kC(D',\rho)_i\right)\Gamma\left(kC(D',\rho)_j\right)}{\Gamma\left(kC(D,\rho)_i\right)\Gamma\left(kC(D,\rho)_j\right)}\left(\frac{x_i}{x_j}\right)^{k(C(D,\rho)_i-C(D',\rho)_i)}\right),$$

where the final equality follows from the fact that $C(D,\rho)_i+C(D,\rho)_j=C(D',\rho)_i+C(D',\rho)_j$. To bound the term in (13), we formulate the optimization problem

$$\begin{aligned}\nu=\max_{C(D,\rho),C(D',\rho),x}\quad & \log\left(\frac{\Gamma\left(kC(D',\rho)_i\right)\Gamma\left(kC(D',\rho)_j\right)}{\Gamma\left(kC(D,\rho)_i\right)\Gamma\left(kC(D,\rho)_j\right)}\left(\frac{x_i}{x_j}\right)^{k(C(D,\rho)_i-C(D',\rho)_i)}\right)\\ \text{subject to}\quad & |C(D,\rho)_i-C(D',\rho)_i|\le\frac{1}{N}\\ & C(D,\rho)_i+C(D,\rho)_j=C(D',\rho)_i+C(D',\rho)_j\\ & C(D,\rho)_i,C(D,\rho)_j\in[\eta,1-2\eta]\\ & C(D',\rho)_i,C(D',\rho)_j\in[\eta,1-2\eta]\\ & x_i,x_j\in[\gamma,1-(n-1)\gamma],\end{aligned}\tag{14}$$

where the first two constraints follow from adjacency and the remaining constraints enforce that $C(D,\rho),C(D',\rho)\in\Delta_n^{(\eta)}$ and $x\in\Omega_1$, where

$$\Omega_1=\{x\in\Delta_n\mid x_i\ge\gamma\text{ for all }i\in[n]\}.$$

Assumptions 2 and 3 ensure that all intervals in the constraints are non-empty. Let Ξ denote the feasible region of (14). From the sub-additivity of the maximum, we have that

$$\begin{aligned}\nu\le\max_{C(D,\rho),C(D',\rho)_i,x\in\Xi}\log\left(\left(\frac{x_i}{x_j}\right)^{k(C(D,\rho)_i-C(D',\rho)_i)}\right)+&\\ \max_{C(D,\rho),C(D',\rho),x\in\Xi}\log\left(\frac{\Gamma\left(kC(D',\rho)_i\right)\Gamma\left(kC(D',\rho)_j\right)}{\Gamma\left(kC(D,\rho)_i\right)\Gamma\left(kC(D,\rho)_j\right)}\right)=\nu_1+\nu_2.&\end{aligned}\tag{15}$$

Starting with ν_1, we have

$$\begin{aligned}\nu_1&=\max_{C(D,\rho),C(D',\rho),x\in\Xi}\log\left(\left(\frac{x_i}{x_j}\right)^{k(C(D,\rho)_i-C(D',\rho)_i)}\right)\\ &\le\max_{C(D,\rho),C(D',\rho),x\in\Xi}|k(C(D,\rho)_i-C(D',\rho)_i)|\left|\log\left(\frac{x_i}{x_j}\right)\right|\\ &=\frac{k}{N}\log\left(\frac{1-(n-1)\gamma}{\gamma}\right).\end{aligned}\tag{16}$$

Next we bound ν_2 from (15). Let $p=C(D,\rho)_i+C(D,\rho)_j=C(D',\rho)_i+C(D',\rho)_j$. Then we have

$$\begin{aligned}\nu_2=\max_{C(D,\rho),C(D',\rho)}\quad & \log\left(\frac{\Gamma\left(kC(D',\rho)_i\right)\Gamma\left(k(p-C(D',\rho)_i)\right)}{\Gamma\left(kC(D,\rho)_i\right)\Gamma\left(k(p-C(D,\rho)_i)\right)}\right)\\ \text{subject to}\quad & |C(D,\rho)_i-C(D',\rho)_i|\le\frac{1}{N}\\ & p\in[2\eta,1-\eta]\\ & C(D,\rho)_i,C(D,\rho)_j\in[\eta,1-2\eta]\\ & C(D',\rho)_i,C(D',\rho)_j\in[\eta,1-2\eta].\end{aligned}\tag{17}$$

Assume without loss of generality that $C(D',\rho)_i \leq C(D,\rho)_i$ (since adjacency is a symmetric relation, the following also holds when $C(D,\rho)_i \leq C(D',\rho)_i$). Then from Lemma 7, we can upper bound the objective in (17) by replacing p with its upper bound. Then we have

$$\begin{aligned} \nu_2 \leq \max_{C(D,\rho),C(D',\rho)} \quad & \log\left(\frac{\text{Beta}(kC(D',\rho)_i, k(1-\eta-C(D',\rho)_i))}{\text{Beta}(kC(D,\rho)_i, k(1-\eta-C(D,\rho)_i))}\right) \\ \text{subject to} \quad & |C(D,\rho)_i - C(D',\rho)_i| \leq \frac{1}{N} \\ & C(D,\rho)_i, C(D,\rho)_j \in [\eta, 1-2\eta] \\ & C(D',\rho)_i, C(D',\rho)_j \in [\eta, 1-2\eta]. \end{aligned} \tag{18}$$

We then have that the Karush-Kuhn-Tucker (KKT) conditions for optimality are not satisfied in the interior of the feasible region for a problem of the form of (18). That same argument applies here, and therefore only the extreme values of $C(D,\rho)_i$ and $C(D',\rho)_i$ need to be considered to find the optimum, thus giving the set of possible solutions

$$\left\{\left(\eta, \eta+\frac{1}{N}\right), \left(1-2\eta, 1-2\eta-\frac{1}{N}\right), \left(\eta+\frac{1}{N}, \eta\right), \left(1-2\eta+\frac{1}{N}, 1-2\eta\right)\right\}.$$

Substituting each solution into (18), we find that $\left(1-2\eta, 1-2\eta-\frac{1}{N}\right)$ and $\left(1-2\eta+\frac{1}{N}, 1-2\eta\right)$ both yield negative objective values, while $\left(\eta, \eta+\frac{1}{N}\right)$ and $\left(\eta+\frac{1}{N}, \eta\right)$ yield identical positive objective values. We select the optimal pair $\left(C(D,\rho)_i^*, C(D',\rho)_i^*\right) = \left(\eta+\frac{1}{N}, \eta\right)$ to obtain

$$\nu_2 \leq \log\left(\frac{\text{Beta}(k\eta, k(1-2\eta))}{\text{Beta}\left(k\left(\eta+\frac{1}{N}\right), k(1-2\eta-\frac{1}{N})\right)}\right). \tag{19}$$

Substituting (16) and (19) into (15) gives

$$\nu \leq \log\left(\frac{\text{Beta}(k\eta, k(1-2\eta))}{\text{Beta}\left(k\left(\eta+\frac{1}{N}\right), k(1-2\eta-\frac{1}{N})\right)}\right) + \frac{k}{N}\log\left(\frac{1-(n-1)\gamma}{\gamma}\right).$$

□

C.2 Proof of Theorem 2

From the definition of the KL divergence, we have

$$\mathcal{K}(C(D,\rho)||\tilde{C}) = \sum_{i=1}^{n} C(D,\rho)_i \log\left(\frac{C(D,\rho)_i}{\tilde{C}_i}\right).$$

Taking the expectation and using linearity yields

$$\mathbb{E}\left[\mathcal{K}(C(D,\rho)||\tilde{C})\right] = \sum_{i=1}^{n} C(D,\rho)_i \mathbb{E}\left[\log\left(\frac{C(D,\rho)_i}{\tilde{C}_i}\right)\right] = \sum_{i=1}^{n} C(D,\rho)_i \left(\log\left(C(D,\rho)_i\right) - \mathbb{E}\left[\log(\tilde{C}_i)\right]\right). \tag{20}$$

The marginal distribution for a single element of a Dirichlet random vector is the Beta distribution (Kotz et al., 2019). The expectation of the log of a $\text{Beta}(x;\alpha,\beta)$ distributed random variable X is

$$\mathbb{E}\left[\log(X)\right] = \psi(\alpha) - \psi(\alpha+\beta).$$

Here, $\tilde{C}_i \sim \text{Beta}(x; kC(D,\rho)_i, k-kC(D,\rho)_i)$. Then we have

$$\mathbb{E}\left[\log\left(\tilde{C}_i\right)\right] = \psi\big(kC(D,\rho)_i\big) - \psi\big(k-kC(D,\rho)_i+kC(D,\rho)_i\big) = \psi(kC(D,\rho)_i) - \psi(k). \tag{21}$$

Substituting (21) into (20) gives us

$$\mathbb{E}\left[\mathcal{K}(C(D,\rho)||\tilde{C})\right] = \sum_{i=1}^{n} C(D,\rho)_i\big(\log(C(D,\rho)_i) + \psi(k) - \psi(kC(D,\rho)_i)\big)$$

as desired. □

C.3 Proof of Corollary 1

To determine an upper bound that is independent of the sensitive data, we consider the optimization problem

$$\begin{aligned} f^* = \max_p \quad & \mathbb{E}\left[\mathcal{K}(p||\tilde{C})\right] \\ \text{subject to} \quad & p_i \in \left[\frac{1}{N}, \frac{N-n+1}{N}\right] \text{ for all } i \in [n] \\ & \sum_{i=1}^{n} p_i = 1, \end{aligned} \tag{22}$$

where the first constraint enforces that p satisfies Assumption 1 (i.e., it can be the output of $C(\cdot, \rho)$ for some some $\mathcal{D} \in \mathfrak{D}$) and the second constraint enforces that p is still a probability mass function. It then follows that

$$\mathbb{E}\left[\mathcal{K}(C(D,\rho)||\tilde{C})\right] \leq f^* \tag{23}$$

for all $D \in \mathfrak{D}$, and we are left to solve the problem in (22) to obtain an upper bound. Substituting (5) into the cost function in (22) gives us

$$\begin{aligned} f^* = \max_p \quad & \sum_{i=1}^{n} p_i(\log(p_i) - \psi(kp_i) + \psi(k)) \\ \text{subject to} \quad & p_i \in \left[\frac{1}{N}, \frac{N-n+1}{N}\right] \text{ for all } i \in [n] \\ & \sum_{i=1}^{n} p_i = 1. \end{aligned}$$

Next, we split the sum in the objective function into two terms:

$$\begin{aligned} f^* = \max_{p,q_2} \quad & \sum_{i=1}^{n-1} \left[p_i(\log(p_i) - \psi(kp_i) + \psi(k))\right] + q_2(\log(q_2) - \psi(kq_2) + \psi(k)) \\ \text{subject to} \quad & p_i, q_2 \in \left[\frac{1}{N}, \frac{N-n+1}{N}\right] \text{ for all } i \in [n-1] \\ & \sum_{i=1}^{n-1} p_i + q_2 = 1. \end{aligned}$$

From the sub-additivity of the maximum, we then have

$$\begin{aligned} f^* \leq \max_p \quad & \sum_{i=1}^{n-1} p_i(\log(p_i) - \psi(kp_i) + \psi(k)) + \max_{q_2} q_2(\log(q_2) - \psi(kq_2) + \psi(k)) \\ \text{subject to} \quad & p_i, q_2 \in \left[\frac{1}{N}, \frac{N-n+1}{N}\right] \text{ for all } i \in [n-1] \\ & q_2 = 1 - \sum_{i=1}^{n-1} p_i. \end{aligned}$$

For the maximization over p we have

$$\begin{aligned} f_1^* = \max_p \quad & \sum_{i=1}^{n-1} p_i(\log(p_i) - \psi(kp_i) + \psi(k)) \\ \text{subject to} \quad & p_i \in \left[\frac{1}{N}, \frac{N-n+1}{N}\right] \text{ for all } i \in [n-1]. \end{aligned} \tag{24}$$

Note that

$$\sum_{i=1}^{n-1} p_i(\log(p_i) - \psi(kp_i) + \psi(k))$$

is a symmetric function of the p_i's in the sense that an arbitrary permutation of the indices of $\{1, 2, \ldots, n-1\}$ does not change its value. From (Waterhouse, 1983), it follows that every p_i in (24) must be equal at the optimum. Then the optimization problem in (24) becomes

$$\begin{aligned} f_1^* = f_1(q_1^*) = (n-1) \max_{q_1} \quad & q_1(\log(q_1) - \psi(kq_1) + \psi(k)) \\ \text{subject to} \quad & q_1 \in \left[\frac{1}{N}, \frac{N-n+1}{N}\right], \end{aligned}$$

and we are left to determine what the value of q_1 must be. To do so, we will show that f_1 is monotonically decreasing by analyzing the derivative of f_1:

$$f_1'(q_1) = \frac{d}{dq_1}(q_1 \log(q_1) - q_1\psi(kq_1) + q_1\psi(k)) = \log(q_1) + 1 - \psi(kq_1) - kq_1\psi^{(1)}(kq_1) + \psi(k), \tag{25}$$

where $\psi^{(1)}(\cdot)$ is the trigamma function, and we neglect the constant $(n-1)$ since it is a positive constant and has no impact on the sign of $f'(q_1)$. Now define $x = kq_1$. Noting that $\log(q_1) = \log(x) - \log(k)$, we substitute x into (25) to obtain

$$\log(q_1) + 1 - \psi(kq_1) - kq_1\psi^{(1)}(kq_1) + \psi(k) = (\log(x) - \psi(x) - x\psi^{(1)}(x) + 1) - (\log(k) - \psi(k)).$$

We define $H(x) = \log(x) - \psi(x) - x\psi^{(1)}(x) + 1$, which we substitute above to obtain

$$f_1'(q_1) = H(kq_1) - (\log(k) - \psi(k)). \tag{26}$$

We will now show that $H(kq_1)$ is monotonically increasing in q_1. Since $\psi^{(1)}(x)$ is strictly negative and $x > 0$, we may rewrite $H(x)$ as

$$H(x) = \log(x) + |x\psi^{(1)}(x)| - \psi(x) + 1.$$

We lower bound $x = kq_1$ using Assumption 4. We know that $\eta \leq \min_{i\in[n]} q_i$, and thus $\eta^{-1} \geq \frac{1}{\min_{i\in[n]} q_i}$. Since $k \geq \frac{3}{2\eta}$ from Assumption 4, substituting the bound $\eta^{-1} \geq \frac{1}{\min_{i\in[n]} q_i}$ yields that $k \geq \frac{3}{2}\frac{1}{\min_{i\in[n]} q_i}$. Additionally, since $q_1 \geq \min_{i\in[n]} q_i$, we have that $x = kq_1 \geq \frac{3}{2}\frac{1}{\min_{i\in[n]} q_i}\min_{i\in[n]} q_i = \frac{3}{2}$, implying $\psi(x) > 0$. Thus, we may rewrite $H(x)$ again as

$$H(x) = \log(x) + |x\psi^{(1)}(x)| - |\psi(x)| + 1.$$

From (Qi et al., 2010, Theorem 2), the function $-|x\psi^{(1)}(x)| + |\psi(x)|$ is monotonically decreasing. Then $|x\psi^{(1)}(x)| - |\psi(x)|$ is monotonically increasing. Since $\log(x)$ is also monotonically increasing and the sum of monotonically increasing functions is monotonically increasing, $H(x)$ is monotonically increasing. Since $q_1 < 1$, it follows that $kq_1 < k$, and since $H(x)$ is monotonically increasing, $H(kq_1) < H(k)$. Substituting this bound into (26) yields

$$f_1'(q_1) < H(k) - (\log(k) - \psi(k)) = \log(k) - \psi(k) - k\psi^{(1)}(k) + 1 - \log(k) + \psi(k) = 1 - k\psi^{(1)}(k).$$

We are left to show that $1 - k\psi^{(1)}(k) < 0$ or equivalently $k\psi^{(1)}(k) > 1$. Recall the integral representation of the trigamma function:

$$\psi^{(1)}(k) = \int_0^\infty \frac{te^{-xt}}{1-e^{-t}}dt. \tag{27}$$

We may bound the integrand using the fact that $1 - e^{-t} < t$ for $t \in (0, \infty)$, which implies that

$$\frac{1}{1-e^{-t}} > \frac{1}{t}. \tag{28}$$

Substituting (28) into (27) yields

$$k\psi^{(1)}(k) = k\int_0^\infty \frac{te^{-kt}}{1-e^{-t}}dt > \int_0^\infty e^{-kt}dt = k\left(\frac{1}{k}\right) = 1.$$

Thus, $f_1'(q_1) < 0$ and $f_1(q_1)$ is monotonically decreasing in q_1. Therefore, q_1^* is equal to the lower bound of the domain of q_1, which is $\frac{1}{N}$. Additionally, since $q_2 = 1 - \sum_{i=1}^{n-1} p_i = 1 - (n-1)q_1$, we have $q_2^* = 1 - \frac{n-1}{N}$. This gives $p^* = [\frac{1}{N}\mathbb{1}_{n-1}{}^T,\ 1 - \frac{n-1}{N}]^T$. Thus,

$$f^* \leq \frac{n-1}{N}\left(\log\left(\frac{1}{N}\right) - \psi\left(\frac{k}{N}\right)\right) + \frac{N-n+1}{N}\left(\log\left(\frac{N-n+1}{N}\right) - \psi\left(\frac{(N-n+1)k}{N}\right)\right) + \psi(k). \quad (29)$$

Substituting (29) into (23) gives

$$\mathbb{E}\left[\mathcal{K}(C(D,\rho)||\tilde{C})\right] \leq \frac{n-1}{N}\left(\log\left(\frac{1}{N}\right) - \psi\left(\frac{k}{N}\right)\right) + \frac{N-n+1}{N}\left(\log\left(\frac{N-n+1}{N}\right) - \psi\left(\frac{(N-n+1)k}{N}\right)\right) + \psi(k).$$

Defining $\zeta(x) = \log\left(\frac{x+1}{N}\right) - \psi\left(\frac{(x+1)k}{N}\right)$, substituting in $\zeta(0)$ and $\zeta(N-n)$ completes the proof. □

C.4 Proof of Lemma 1

From Meiser (2018), we have that (ϵ, δ)-probabilistic differential privacy implies conventional (ϵ, δ)-differential privacy. Thus, it is sufficient to show that parallel composition holds for conventional (ϵ, δ)-differential privacy. Consider a database $\mathcal{D} = \{D_1, \ldots, D_i, \ldots, D_m\}$, where the collection $\{D_i\}_{i\in[m]}$ is a partition of $\mathcal{D}$, and a mechanism $\mathcal{M}(\mathcal{D}) = (\mathcal{M}_1(D_1), \ldots, \mathcal{M}_m(D_m))$, where $\mathcal{M}_i$ is conventional (ϵ_i, δ_i)-differentially private. For any Borel measurable set $T = T_1 \times \cdots \times T_m$, we have that

$$\mathbb{P}\left(\mathcal{M}(\mathcal{D}) \in T\right) = \prod_{k\in[m]} \mathbb{P}\left(\mathcal{M}_k(D_k) \in T_k\right).$$

Let $\mathcal{D}' = \{D_1', \ldots, D_i', \ldots, D_m'\}$ be an adjacent database to $\mathcal{D}$, and assume without loss of generality that $\mathcal{D}, \mathcal{D}'$ differ in entry i. Then, because $D_k = D_k'$ for all $k \in [m] \setminus \{i\}$, we have that

$$\prod_{k\in[m]} \mathbb{P}\left(\mathcal{M}_k(D_k) \in T_k\right) = \mathbb{P}\left(\mathcal{M}_i(D_i) \in T_i\right) \prod_{k\in[m]\setminus i} \mathbb{P}\left(\mathcal{M}_k(D_k') \in T_k\right). \quad (30)$$

Since the mechanism $\mathcal{M}_i$ is conventional (ϵ_i, δ_i)-differentially private, we bound (30) as

$$\begin{aligned}\mathbb{P}\left(\mathcal{M}_i(D_i) \in T_i\right) \prod_{k\in[m]\setminus i} \mathbb{P}\left(\mathcal{M}_k(D_k') \in T_k\right) &\leq \left(e^{\epsilon_i}\mathbb{P}\left(\mathcal{M}_i(D_i') \in T_i\right) + \delta_i\right) \prod_{k\in[m]\setminus i} \mathbb{P}\left(\mathcal{M}_k(D_k') \in T_k\right) \\ &= e^{\epsilon_i}\mathbb{P}\left(\mathcal{M}(\mathcal{D}') \in T\right) + \delta_i \prod_{k\in[m]\setminus i} \mathbb{P}\left(\mathcal{M}_k(D_k') \in T_k\right). \quad (31)\end{aligned}$$

Then, because $\mathbb{P}\left(\mathcal{M}_k(D_k') \in T_k\right) \leq 1$ for all $k \in [m] \setminus i$, we can bound the righthand side of (31) by

$$e^{\epsilon_i}\mathbb{P}\left(\mathcal{M}(\mathcal{D}') \in T\right) + \delta_i \prod_{k\in[m]\setminus i} \mathbb{P}\left(\mathcal{M}_k(D_k') \in T_k\right) \leq e^{\epsilon_i}\mathbb{P}\left(\mathcal{M}(\mathcal{D}') \in T\right) + \delta_i$$

and thus

$$\mathbb{P}\left(\mathcal{M}(\mathcal{D}) \in T\right) \leq e^{\epsilon_i}\mathbb{P}\left(\mathcal{M}(\mathcal{D}') \in T\right) + \delta_i. \quad (32)$$

The inequality in (32) must hold for all $i \in [m]$, and thus we have that

$$\mathbb{P}\left(\mathcal{M}(\mathcal{D}) \in T\right) \leq \max_{i\in[m]} e^{\epsilon_i}\mathbb{P}\left(\mathcal{M}(\mathcal{D}') \in T\right) + \max_{i\in[m]} \delta_i.$$

Thus, releasing $\mathcal{M}(\mathcal{D})$ provides $\mathcal{D}$ with conventional $(\max_{i\in[m]} \epsilon_i, \max_{i\in[m]} \delta_i)$-differential privacy.

□

C.5 Proof of Theorem 3

From Theorem 1, $\mathcal{M}_{\text{Dir}}^{(k_i)}(C(D_i,\rho))$ is (ϵ_i,δ_i)-differentially private for all $i \in [n]$. Since the D_i's partition the database $\mathcal{D}$, $\tilde{P}$ is the output of the parallel composition of n mechanisms, where the ith is (ϵ_i,δ_i)-probabilistic differentially private. Then from Lemma 1 the computation of $\tilde{P}$ is $(\max_{i\in[n]}\epsilon_i, \max_{i\in[n]}\delta_i)$-differentially private. □

C.6 Proof of Theorem 4

From the definition of the total variation distance,

$$\mathbb{E}\left[\|\pi-\tilde{\pi}\|_{TV}\right] = \frac{1}{2}\mathbb{E}\left[\|\pi-\tilde{\pi}\|_1\right].$$

From Lemma 5,

$$\mathbb{E}\left[\|\pi-\tilde{\pi}\|_1\right] \leq \|Z\|_1 \,\mathbb{E}\left[\left\|P-\tilde{P}\right\|_1\right]. \tag{33}$$

Substituting (12) from Lemma 6 into (33) gives

$$\mathbb{E}\left[\|\pi-\tilde{\pi}\|_{TV}\right] = \frac{1}{2}\mathbb{E}\left[\|\pi-\tilde{\pi}\|_1\right] \leq \frac{\sqrt{2}}{2}\|Z\|_1 \cdot$$
$$\left(\sum_{i=1}^{n}\pi(i)\left(\frac{n-1}{N_i}\left(\log\left(\frac{1}{N_i}\right)-\psi\left(\frac{k}{N_i}\right)\right)+\frac{N_i-n+1}{N_i}\left(\log\left(\frac{N_i-n+1}{N_i}\right)-\psi\left(\frac{(N_i-n+1)k_i}{N_i}\right)\right)+\psi(k_i)\right)\right)^{\frac{1}{2}}.$$

Defining $\zeta_i(x) = \log\left(\frac{x+1}{N_i}\right) - \psi\left(\frac{(x+1)k_i}{N_i}\right)$ and substituting in $\zeta_i(0)$ and $\zeta_i(N_i-n)$ above completes the proof. □

C.7 Proof of Theorem 5

From (Ipsen & Selee, 2011, Theorem 4.1), for given stochastic matrices S, S_1, and S_2, the ∞-norm ergodicity coefficient satisfies $\tau_\infty(S) \leq \|S\|_1$ and $|\tau_\infty(S_1)-\tau_\infty(S_2)| \leq \tau_\infty(S_1-S_2)$. Combining these properties for $P(\mathcal{D},\rho)$ and $\tilde{P}$ gives

$$|\tau_\infty(P(\mathcal{D},\rho))-\tau_\infty(\tilde{P})| \leq \left\|P(\mathcal{D},\rho)-\tilde{P}\right\|_1.$$

Taking the expectation yields

$$\mathbb{E}\left[|\tau_\infty(P(\mathcal{D},\rho))-\tau_\infty(\tilde{P})|\right] \leq \mathbb{E}\left[\left\|P(\mathcal{D},\rho)-\tilde{P}\right\|_1\right].$$

Substituting (12) from Lemma 6 then gives

$$\mathbb{E}\left[|\tau_\infty(P(\mathcal{D},\rho))-\tau_\infty(\tilde{P})|\right] \leq \sqrt{2}\cdot$$
$$\left(\sum_{i=1}^{n}\pi(i)\left(\frac{n-1}{N_i}\left(\log\left(\frac{1}{N_i}\right)-\psi\left(\frac{k_i}{N_i}\right)\right)+\frac{N_i-n+1}{N_i}\left(\log\left(\frac{N_i-n+1}{N_i}\right)-\psi\left(\frac{(N_i-n+1)k_i}{N_i}\right)\right)+\psi(k_i)\right)\right)^{\frac{1}{2}}.$$

Substituting $\zeta_i(0)$ and $\zeta_i(N_i-n)$ above completes the proof. □

D Additional Accuracy Results

Lemma 8 (Accuracy bounds; (Fallin et al., 2023))**.** Fix a database D with $N \in \mathbb{N}$ entries, a category set ρ such that $|\rho| = n \in \mathbb{N}$, and $\eta \in (0,1)$. Let $C(D,\rho) \in \Delta_n^{(\eta)}$ be defined as in (1), and let $\tilde{C}$ be its privatized form. Let Assumptions 1-4 hold. Then

$$\mathbb{E}\left[|C(D,\rho)_i-\tilde{C}_i|\right] = 2\frac{C(D,\rho)_i^{kC(D,\rho)_i}(1-C(D,\rho)_i)^{k(1-C(D,\rho)_i)}}{k\,\text{Beta}(kC(D,\rho)_i, k(1-C(D,\rho)_i))}, \quad \mathbb{E}\left[|C(D,\rho)_i-\tilde{C}_i|\right] \leq \frac{\Gamma(k)\,2^{1-k}}{\Gamma\left(\frac{k}{2}\right)^2 k}$$

and

$$\mathbb{E}\left[|C(D,\rho)_i - \tilde{C}_i|^2\right] = \frac{C(D,\rho)_i - C(D,\rho)_i^2}{k+1}, \quad \mathbb{E}\left[|C(D,\rho)_i - \tilde{C}_i|^2\right] \leq \frac{1}{4(k+1)}.$$

♦

Lemma 9 (Accuracy Lower Bound)**.** Let the conditions of Lemma 8 hold. Then

$$\mathbb{E}\left[|C(D,\rho)_i - \tilde{C}_i|\right] \geq 2\frac{\left(\frac{1}{N}\right)^{k(1-\frac{n-1}{N})}\left(\frac{n-1}{N}\right)^{k(1-\frac{1}{N})}}{k\,\mathrm{Beta}\left(\frac{k}{N}, \frac{k(n-1)}{N}\right)}.$$

♦

Proof. From Lemma 8,

$$\mathbb{E}\left[|C(D,\rho)_i - \tilde{C}_i|\right] = 2\frac{C(D,\rho)_i^{kC(D,\rho)_i}(1-C(D,\rho)_i)^{k(1-C(D,\rho)_i)}}{k\,\mathrm{Beta}(kC(D,\rho)_i, k(1-C(D,\rho)_i))}. \tag{34}$$

To lower bound (34), we seek a lower bound on the numerator and an upper bound on the denominator. We begin with a lower bound on the numerator. For all $i \in [n]$,

$$\frac{1}{N} \leq C(D,\rho)_i \leq 1 - \frac{n-1}{N}$$

and

$$1 - \frac{1}{N} \geq 1 - C(D,\rho)_i \geq \frac{n-1}{N}.$$

For the numerator,

$$C(D,\rho)_i^{kC(D,\rho)_i}(1-C(D,\rho)_i)^{k(1-C(D,\rho)_i)} \geq \left(\frac{1}{N}\right)^{k(1-\frac{n-1}{N})}\left(\frac{n-1}{N}\right)^{k(1-\frac{1}{N})}. \tag{35}$$

Next, we upper bound the denominator. We proceed by showing that the denominator is monotonically decreasing in $C(D,\rho)_i$. Consider the following partial derivatives of the Beta function:

$$\frac{\partial}{\partial a}\mathrm{Beta}(a,b) = \mathrm{Beta}(a,b)(\psi(a) - \psi(a+b)) \tag{36}$$

$$\frac{\partial}{\partial b}\mathrm{Beta}(a,b) = \mathrm{Beta}(a,b)(\psi(b) - \psi(a+b)). \tag{37}$$

For all $a, b > 0$, $\mathrm{Beta}(a,b) > 0$. Additionally, $\psi(x)$ is strictly increasing for $x > 0$. Then $\psi(a) - \psi(a+b) < 0$ and $\psi(b) - \psi(a+b) < 0$. These inequalities imply that (36) and (37) are negative for all $a, b > 0$ and $\mathrm{Beta}(a,b)$ is therefore monotonically decreasing in both a and b. We then upper bound the denominator of (34) by

$$k\,\mathrm{Beta}(kC(D,\rho)_i, k(1-C(D,\rho)_i)) \leq k\,\mathrm{Beta}\left(\frac{k}{N}, \frac{k(n-1)}{N}\right). \tag{38}$$

We then substitute (35) and (38) into (34) to obtain

$$\mathbb{E}\left[|C(D,\rho)_i - \tilde{C}_i|\right] \geq 2\frac{\left(\frac{1}{N}\right)^{k(1-\frac{n-1}{N})}\left(\frac{n-1}{N}\right)^{k(1-\frac{1}{N})}}{k\,\mathrm{Beta}\left(\frac{k}{N}, \frac{k(n-1)}{N}\right)},$$

which completes the proof. □

Table 1: The strongest possible privacy at $\gamma = 10^{-8}$ for varying numbers of taxi zones, where there are originally 63 zones. As more zones are combined, users attain stronger privacy guarantees while still maintaining model accuracy.

Number of Taxi Zones	Smallest Possible ϵ
63	45.39
60	45.39
55	19.82
50	4.690
45	4.690
40	3.730
35	3.078
30	3.078

E Additional Simulation Details

E.1 Class Grade Distribution

We first provide an example of two class grade databases which satisfy the adjacency relation in Definition 3. Consider two databases $\mathcal{A}, \mathcal{B}$ which each contain a list of letter grades $\rho \triangleq \{\texttt{A}, \texttt{B}, \texttt{C}, \texttt{D}, \texttt{F}\}$ and a corresponding number of students in the course who achieved each letter grade. In this example, each database has a single partition given by A_1, B_1, respectively. Each database is given as

$$\mathcal{A} = \{A_1 : [(\texttt{A}, 10), (\texttt{B}, 20), (\texttt{C}, 22), (\texttt{D}, 15), (\texttt{F}, 4)]\},$$

and

$$\mathcal{B} = \{B_1 : [(\texttt{A}, 10), (\texttt{B}, 19), (\texttt{C}, 22), (\texttt{D}, 15), (\texttt{F}, 4)]\}.$$

We have that $\mathcal{A}, \mathcal{B}$ are adjacent since they differ only in the number of students which earned a grade of $\texttt{B}$.

E.2 New York City Taxis

We first provide an example of a simplified NYC Taxi Limousine Commission database representing trips between the stops at Alphabet City and Battery Park. Define $\texttt{AC} \triangleq$ Alphabet City and $\texttt{BP} \triangleq$ Battery Park. Consider the database $\mathcal{C}$ with partitions C_1, C_2 and the database $\mathcal{D}$ with partitions D_1, D_2 such that

$$\mathcal{C} = \{C_1 : [((\texttt{AC}, \texttt{AC}), 100), ((\texttt{AC}, \texttt{BP}), 133)],\ C_2 : [((\texttt{BP}, \texttt{AC}), 567), ((\texttt{BP}, \texttt{BP}), 217)]\}$$

and

$$\mathcal{D} = \{D_1 : [((\texttt{AC}, \texttt{AC}), 100), ((\texttt{AC}, \texttt{BP}), 132)],\ D_2 : [((\texttt{BP}, \texttt{AC}), 567), ((\texttt{BP}, \texttt{BP}), 217)]\}.$$

The entry $((\texttt{AC}, \texttt{BP}), 133)$ in database $\mathcal{C}$ indicates that 133 trips were taken originating from Alphabet City and terminating at Battery Park. We have that $\mathcal{C}, \mathcal{D}$ are adjacent because they differ only in the number of trips taken from Alphabet City to Battery Park.

Next, we detail the pre-processing steps performed on the NYC taxi trips database in Section 6. First, we considered only trips between taxi zones on the island of Manhattan since this is where the majority of taxi trips take place. Additionally, while some smaller islands are considered taxi zones in Manhattan, we neglect to include these since there is no clear interpretation of a taxi trip between islands where there is no bridge, such as from Manhattan to the Statue of Liberty.

Some taxi zones are significantly larger than others, and due to the dependency of ϵ on the number of participants in each independent subset the database N_i, they achieve much stronger privacy than smaller taxi zones. As a result, smaller taxi zones have an out-sized impact on weakening the privacy guarantee of the

entire Markov chain model. To that end, we combine low data taxi zones with adjacent taxi zones to achieve a more equitable distribution of participants. Table 1 shows the strongest possible privacy (i.e., smallest ϵ) for a varying number of taxi zones with $\gamma = 10^{-8}$. As the number of zones decreases, it is possible to attain stronger privacy, highlighting a fundamental tradeoff: as a database is partitioned into smaller subsets with more categories, it becomes more difficult to guarantee strong privacy for the entire dataset. Thus, database curators interested in strong privacy must consider this tradeoff when selecting categories and partitioning a database.