Technical Report: Differential Privacy for Stochastic Matrices Using the Matrix Dirichlet Mechanism

Brandon Fallin*, Calvin Hawkins*, Bo Chen*, Parham Gohari†,
Alex Benvenuti*, Ufuk Topcu†, and Matthew Hale*

Abstract— **Stochastic matrices are commonly used to analyze Markov chains, but revealing them can leak sensitive information. Therefore, in this paper we introduce a technique to privatize stochastic matrices in a way that (i) conceals the probabilities they contain, and (ii) still allows for accurate analyses of Markov chains. Specifically, we use differential privacy, which is a statistical framework for protecting sensitive data. To implement it, we introduce the Matrix Dirichlet Mechanism, which is a probabilistic mapping that perturbs a stochastic matrix to provide privacy. We prove that this mechanism provides differential privacy, and we quantify the error induced in private stochastic matrices as a function of the strength of privacy being provided. We then bound the distance between the stationary distribution of the underlying, sensitive stochastic matrix and the stationary distribution of its privatized form. Numerical results show that, under typical conditions, privacy introduces error as low as 5.05% in the stationary distribution of a stochastic matrix.**

I. Introduction

Control applications have become increasingly reliant on user data, e.g., in smart power grids [1], smart transit applications [2], and networks of robots [3]. There are often privacy concerns associated with sharing user data because of what it can reveal. For example, smart appliance usage data [4], driving routines [5], and other sensitive data streams [6], [7] can be revealing about a user's past daily habits or locations, and allow inferences to be drawn about these behaviors in the future. Data is still needed in many applications, and it is therefore desirable to provide privacy to users while preserving the usefulness of their data.

In this paper, we provide privacy of this kind to Markov chain models of systems. Markov chains have been used to model smart power grids, users' online behavior, and devices on the Internet of Things (IoT) [8]–[10]. Markov chains model a system with random transitions between a finite number of states, and the probabilities of these transitions are represented by a stochastic matrix [11], i.e., a matrix with non-negative entries whose rows sum to one. The entries of a stochastic matrix are sensitive because they can reveal how often a user engages in a certain behavior or how likely they are to engage in it, e.g., by revealing the probability of home occupancy at a given time of day, browsing and shopping patterns, or trends in smart device usage [12], [13]. These privacy threats pose a significant risk to individuals, and they motivate us to privatize stochastic matrices.

We do so using differential privacy, which originates in the computer science literature, where it was originally used to protect sensitive data when databases are queried [14]. Differential privacy is appealing because (i) it is immune to post-processing, in the sense that arbitrary post-hoc computations on private data do not weaken its privacy protections, and (ii) it is robust to side information [15], in that gaining knowledge of some sensitive information does not weaken differential privacy by much [16]. It is also known that one can attain strong differential privacy protections while still providing accurate information [15]. In this paper, the queries that we consider are identity queries of stochastic matrices, i.e., releasing a stochastic matrix itself, and this is what we will privatize.

There is a growing body of work on differential privacy in decision systems, including in multi-agent control, convex optimization, filtering and estimation, and symbolic systems [17]–[22]. These works generally implement differential privacy for numerical data using the Laplace or Gaussian mechanisms, which add noise to sensitive data (or functions thereof) before it is shared. However, these mechanism are a poor fit for the privatization of stochastic matrices. Stochastic matrices have non-negative entries and row sums equal to one, but the outputs of the Laplace and Gaussian mechanisms will not preserve these properties. Therefore, projection onto the allowable set of data would be required, but this has been shown to destroy the accuracy of private data in similar contexts [23]. Thus, new developments are required.

Accordingly, to provide differential privacy to stochastic matrices, the contributions of this paper are as follows:

- We develop the Matrix Dirichlet Mechanism, the first differential privacy mechanism for stochastic matrices
- We bound the error induced by differential privacy between a privatized stochastic matrix and its non-private form in terms of the strength of privacy
- To quantify the utility of privatized stochastic matrices, we bound the distance between the stationary distribution of a stochastic matrix and the stationary distribution of its privatized form
- We show in simulation that, under typical conditions, the errors induced by privacy are as low as 5.05%

*Department of Mechanical and Aerospace Engineering at the University of Florida, Gainesville, FL USA. Emails: {brandonfallin,calvin.hawkins,bo.chen, abenvenuti,matthewhale} @ufl.edu. BF, CH, BC, AB, and MH were supported in part by NSF under CAREER Grant 1943275, by AFOSR under Grant FA9550-19-1-0169, by ONR under Grant N00014-21-1-2502, and by AFRL under Grant FA8651-23-F-A008.

†Department of Aerospace Engineering and Engineering Mechanics at the University of Texas at Austin, Austin, TX USA. Emails: {utopcu,pgohari} @utexas.edu. UT and PG were supported in part by NSF under CAREER Grant 1652113 and by ONR under Grant N00014-21-1-2502.

The rest of this paper is organized as follows. Section II provides background and problem statements. Section III implements differential privacy using the Matrix Dirichlet Mechanism and quantifies the error it induces. Section IV analyzes the trade-off between privacy and accuracy of the stationary distribution of a Markov chain. Section V provides simulations and Section VI concludes.

Notation: We use $\mathbb{R}$ and $\mathbb{N}$ to denote the real and natural numbers, respectively. The set $\mathbb{R}_+$ denotes the positive reals. We use $|S|$ to denote the cardinality of a finite set S. For $n \in \mathbb{N}$, let $[n] = \{1, \ldots, n\}$. We use $\mathbb{1}_n$ to denote the vector of all ones in $\mathbb{R}^n$.

II. Background and Problem Statements

This section provides background on Markov chains and differential privacy, and then it states the problems that are the focus of the remainder of the paper.

A. Unit Simplex and Stochastic Matrices

Each row of a stochastic matrix is an element of the unit simplex, which is the set of vectors with non-negative elements that sum to 1. Formally, let $n \in \mathbb{N}$. The unit simplex in $\mathbb{R}^n$ is denoted by Δ_n, where

$$\Delta_n = \left\{ x \in \mathbb{R}^n \mid \sum_{i=1}^{n} x_i = 1,\ x_i \geq 0 \text{ for all } i \in [n] \right\}.$$

Next, we define the bordered unit simplex. The bordered unit simplex is the set of vectors within the unit simplex whose components are a sufficient distance from 0 and 1.

Definition 1 (Bordered Unit Simplex)**.** Let Δ_n° denote the interior of Δ_n. Fix $\eta, \bar{\eta} > 0$ and let $W \subseteq [n-1]$ satisfy $|W| \geq 2$. Then the bordered unit simplex is defined as

$$\Delta_{n,W}^{(\eta,\bar{\eta})} = \Bigg\{ x \in \Delta_n^{\circ} \mid \sum_{i \in W} x_i \leq 1 - \bar{\eta}, \\ x_i \geq \eta \text{ for all } i \in W \Bigg\}. \quad \Diamond$$

We also establish some mathematical notation for special functions used throughout this work. $\mathbb{P}[\cdot]$ denotes the probability of an event. For a random variable, $\mathrm{E}[\cdot]$ denotes its expectation and $\mathrm{Var}[\cdot]$ denotes its variance. The notation $\|\cdot\|_1$ denotes the 1-norm of a vector or matrix. The space on which we use $\|\cdot\|_1$ will be clear from context. For $x, a, b \in \mathbb{R}_+$, we use the special functions $\Gamma(x) = \int_0^\infty z^{x-1} \exp(-z)\, dz$, $\psi(x) = \frac{\Gamma'(x)}{\Gamma(x)}$,

$$\text{beta}(a,b) = \int_0^1 t^{a-1} (1-t)^{b-1}\, dt = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)}, \quad (1)$$

$$I_z(a,b) = \frac{\int_0^z x^{a-1}(1-x)^{b-1}\, dx}{\text{beta}(a,b)}, \quad (2)$$

which are the gamma, digamma, beta, and regularized incomplete beta functions, respectively. Additionally, the gamma function has the property $\Gamma(k+1) = k\Gamma(k)$.

In this work, we implement privacy by generating random matrices whose rows are in the unit simplex. A building block of this technique is the Dirichlet distribution on the unit simplex. Formally, for a parameter $k \in \mathbb{R}_+$ and a vector $p \in \Delta_n^{\circ}$, the Dirichlet distribution with mean p and parameterized by k is denoted $\mathcal{M}_D^{(k)}$ and defined as

$$\mathbb{P}\left[\mathcal{M}_D^{(k)}(p) = x\right] = \frac{1}{B(kp)} \prod_{i=1}^{n} x_i^{kp_i - 1}, \quad (3)$$

where

$$B(kp) = \frac{\prod_{i=1}^{n} \Gamma(kp_i)}{\Gamma\left(k \sum_{i=1}^{n} p_i\right)}$$

is the multi-variate beta function.

A matrix whose rows all belong to the unit simplex is known as a stochastic matrix. Formally, fix $n \in \mathbb{N}$. We define $\mathcal{S}_n$ as the set of all $n \times n$ stochastic matrices, where

$$\mathcal{S}_n = \left\{ P \in \mathbb{R}^{n \times n} \mid P_{ij} \geq 0 \text{ for all } i, j \in [n],\ P\mathbb{1}_n = \mathbb{1}_n \right\}.$$

Next, we discuss the application of stochastic matrices to Markov chain models.

B. Markov Chains

We now review the necessary background on Markov chains. See [11] for a more detailed exposition. A Markov chain is a sequence of random variables $X_1, X_2, X_3, \ldots, X_k$ that possess the Markov property [11], namely

$$\mathbb{P}(X_{k+1} = x_{k+1} \mid X_1 = x_1,\ X_2 = x_2, \cdots, X_k = x_k) \\ = \mathbb{P}(X_{k+1} = x_{k+1} \mid X_k = x_k).$$

In words, the conditional probability of transitioning from state x_k to state x_{k+1} is independent of the sequence of states that came before the current state x_k. In this work, we consider finite, irreducible, homogeneous Markov chains. That is, Markov chains where (i) its random variables take values in a finite set, (ii) transition is possible from one state to any other state using only transitions of positive probability, and (iii) the transition probabilities are independent of time [11]. Since we consider finite Markov chains, we construct a transition probability matrix P where $P_{ij} = \mathbb{P}(X_{k+1} = x_j \mid X_k = x_i)$. This transition matrix P is stochastic by construction. Let P_i denote the i^{th} row of the transition probability matrix P.

At each time step, we compute the probability distribution of the states of a Markov chain. Let $\mu_k \in \Delta_n$ denote the probability distribution of states at time k. That is, $\mu_{k,i}$ is the probability that the Markov chain is in state i at time k. Let $\mu_0 \in \Delta_n$ denote the initial state distribution. Multiplying by the transition matrix P on the right updates the distribution by another time step, i.e., $\mu_k^T = \mu_{k-1}^T P$. In general, $\mu_k^T = \mu_0^T P^k$ for all $k \geq 1$.

A common way to analyze Markov chains is through their steady-state behavior. Specifically, when P is finite, irreducible, and homogeneous, there exists a limit π as $t \to \infty$ that must satisfy $\pi^T = \pi^T P$, where π is called the "stationary distribution" of the Markov chain.

C. Differential Privacy

We briefly review differential privacy here, and we refer the reader to [15] for a more complete exposition. Differential privacy protects sensitive data by randomizing it or functions of it. Differential privacy is enforced by a randomized mapping, or *mechanism*, that outputs statistically "similar" private values for "close" pieces of sensitive data. Formally, an adjacency relation allows us to quantify how "close" two pieces of data are. In the standard setup, two databases D and $D^{'}$ are adjacent if they differ in one entry [24]. Adjacency is a design choice we make that specifies what must be kept private. We state our adjacency relation for stochastic matrices in Section III-A.

The condition that private outputs be statistically "similar" is formalized by the definition of differential privacy itself. In this work, we utilize probabilistic differential privacy, which is defined as follows.

Definition 2 (Probabilistic Differential Privacy [25])**.** Let P and Q be two adjacent data sets, 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 we can partition the output space $\mathcal{S}$ into two disjoint sets, $\mathbf{\Omega_1}$ and $\mathbf{\Omega_2}$, such that for all P, we have $\mathbb{P}\left[\mathcal{M}(P) \in \mathbf{\Omega_2}\right] \leq \delta$, and, for all Q adjacent to P and all $S \in \mathbf{\Omega_1}$,

$$\log\left(\frac{\mathbb{P}\left[\mathcal{M}(P) = S\right]}{\mathbb{P}\left[\mathcal{M}(Q) = S\right]}\right) \leq \epsilon. \qquad \Diamond$$

The strength of differential privacy is quantified through two parameters: ϵ and δ. The value of ϵ controls the amount of information shared. In the literature, ϵ typically ranges from 0.01 to 10 [26]. The value of δ is the probability that more information is shared than should be allowed by ϵ, and this value typically ranges between 0 and 0.05 [27]. Smaller values of both ϵ and δ imply stronger privacy.

If a mechanism provides probabilistic (ϵ, δ)-differential privacy, then it provides conventional (ϵ, δ)-differential privacy [25]. For a privacy mechanism $\mathcal{M}$, conventional differential privacy states that for any measurable subset A of the range of $\mathcal{M}$ and all adjacent P and Q, we have the inequality $\mathbb{P}\left[\mathcal{M}\left(P\right) \in A\right] \leq e^{\epsilon}\mathbb{P}\left[\mathcal{M}\left(Q\right) \in A\right] + \delta$. It has been shown that the conventional differential privacy guarantees provided by probabilistic differential privacy are strictly stronger than what is implied by the privacy parameters used for probabilistic differential privacy [28, Section 4]. More specifically, it is shown that providing (ϵ, δ)-probabilistic differential privacy implies $(\epsilon, \delta^{'})$-conventional differential privacy, where $\delta^{'} < \delta$.

The level of privacy of a query (or several queries) of data can be calculated using methods of composition. While general sequences of private queries cause privacy to weaken [15, Theorem 3.14], this weakening does not occur if the queries are of disjoint subsets of the sensitive data [29]. Specifically, if the domain of the input to the mechanism is partitioned into disjoint sets and these disjoint sets are queried separately, then the ultimate privacy level is equal to the worst of the privacy guarantees of each query. This is formalized in the following lemma.

Lemma 1 (Parallel Composition [29])**.** Consider a database D partitioned into disjoint subsets $D_1, D_2, \dots, D_N$, and suppose that there are privacy mechanisms $\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_N$, where $\mathcal{M}_i$ is (ϵ_i, δ_i)-differentially private. Then the release of the queries $\mathcal{M}_1(D_1), \mathcal{M}_2(D_2), \dots, \mathcal{M}_N(D_N)$ provides D with $(\max_{i\in[N]} \epsilon_i, \max_{i\in[N]} \delta_i)$-differential privacy. □

We will use Lemma 1 to develop a differential privacy mechanism for stochastic matrices in Section III.

D. Problem Statements

We now formally state the problems that we solve.

Problem 1. *Develop a mechanism that provides (ϵ, δ)-differential privacy to a stochastic matrix P.*

Problem 2. *Bound the expected difference of the entries of the sensitive input P and a private output produced by the mechanism developed in solving Problem 1.*

Problem 3. *Apply the developed mechanism to transition matrices of Markov chains and develop a bound on the distance between the private and non-private stationary distributions to quantify the trade-off between the level of privacy and accuracy.*

We next solve Problem 1 and develop and analyze a privacy mechanism to generate private stochastic matrices.

III. Matrix Dirichlet Mechanism for Differential Privacy of Identity Queries

We begin this section by establishing a formal adjacency definition for stochastic matrices, as well as outlining the Matrix Dirichlet Mechanism in Section III-A. We then show the differential privacy guarantees provided by the Matrix Dirichlet Mechanism by computing δ in Section III-B and ϵ in Section III-C. This solves Problem 1. Lastly, the accuracy of private data is analyzed in Section III-D by bounding the expected difference between the entries of a sensitive stochastic matrix P and a private output produced by the Matrix Dirichlet Mechanism. This solves Problem 2.

A. Matrix Dirichlet Mechanism

For the use of the Matrix Dirichlet Mechanism, we require that row i of a sensitive stochastic matrix be in $\Delta_{n,W_i}^{(\eta,\bar{\eta})}$ for all $i \in [n]$; we emphasize that we let W in Definition 1 vary for each row, and we use W_i to denote the set of indices associated with row i. For convenience, we define a map

$$V : [n] \to \{W_1, \dots, W_n\}, \text{ where } V(i) = W_i. \qquad (4)$$

We now define the set of sensitive matrices we consider.

Definition 3 (Stochastic Matrices)**.** Fix $n \in \mathbb{N}$ and, for each $i \in [n]$, fix a collection of indices $W_i \subseteq [n-1]$. Fix $\eta, \bar{\eta} > 0$. Let $\mathcal{S}_{n,V}^{(\eta,\bar{\eta})}$ be the set of all stochastic matrices whose i^{th} row is in $\Delta^{(\eta,\bar{\eta})}_{n,W_i}$ from Definition 1 for all $i \in [n]$. Then

$$\mathcal{S}^{(\eta,\bar{\eta})}_{n,V} = \left\{ P \in \mathcal{S}_n \mid P_i \in \Delta^{(\eta,\bar{\eta})}_{n,V(i)} \text{ for all } i \in [n] \right\},$$

where V is from (4) and P_i is the i^{th} row of P. $\diamondsuit$

We impose the following assumption on η and $\bar{\eta}$ to ensure that ratios of probability distributions over stochastic matrices are bounded when showing that Matrix Dirichlet Mechanism provides differential privacy.

Assumption 1. For $\mathcal{S}^{(\eta,\bar{\eta})}_{n,V}$ in Definition 3, it holds that $\eta > 0$, $\bar{\eta} > 0$, and $\eta + \bar{\eta} < \frac{1}{2}$. It also holds that $W_i \subseteq [n-1]$ and $|V(i)| = |W_i| \geq 2$ for all $i \in [n]$. $\diamondsuit$

Next, we establish a formal adjacency relation for stochastic matrices.

Definition 4 (Adjacency)**.** Fix $n \in \mathbb{N}$, $\eta > 0$, and $\bar{\eta} > 0$.Let Assumption 1 hold. Let $P, Q \in \mathcal{S}^{(\eta,\bar{\eta})}_{n,V}$ be $n \times n$ stochastic matrices, let P_i be the i^{th} row of P, and let P_{ij} be the $i^{th}j^{th}$ entry of P; Q_i and Q_{ij} are defined analogously. For an adjacency parameter $b \in (0,1]$, P and Q are said to be b-adjacent if, for all $i \in [n]$, there exist $j, k \in W_i$ such that $P_{ij} \neq 0$, $P_{ik} \neq 0$, $P_{i\ell} = Q_{i\ell}$ for all $\ell \neq j, k,$ and $\|P_i - Q_i\|_1 \leq b$. $\diamondsuit$

In words, this means that P and Q are adjacent if for all $i \in [n]$, the i^{th} rows of P and Q differ in no more than two entries and the 1-norm of this difference is bounded by b. It was observed in Section II-C that the conventional definition of adjacency allows for databases to differ in one entry. Here, we must consider each row differing in two entries because rows must sum to 1 and thus it is not possible to change only a single entry. We now formalize the notion of differential privacy for stochastic matrices.

Definition 5 (Differential Privacy for Stochastic Matrices)**.** Fix $n \in \mathbb{N}$, $\eta > 0$, $\bar{\eta} > 0$, and $b \in (0,1]$. let Assumption 1 hold. Let P, $Q \in \mathcal{S}^{(\eta,\bar{\eta})}_{n,V}$ be two b-adjacent $n \times n$ stochastic matrices. Let V be defined as in (4). A mechanism $\mathcal{M} : \mathcal{S}^{(\eta,\bar{\eta})}_{n,V} \to \mathcal{S}_n$ is said to be (ϵ, δ)-differentially private, if, for for any measurable subset A of the range of $\mathcal{M}$ and all b-adjacent P, Q, we have

$$\mathbb{P}\left[\mathcal{M}\left(P\right) \in A\right] \leq e^{\epsilon}\mathbb{P}\left[\mathcal{M}\left(Q\right) \in A\right] + \delta. \qquad \diamondsuit$$

Formally, the query we privatize is the identity query, i.e., we privatize a single stochastic matrix in $\mathcal{S}_n$ to enable it to be shared. This approach is also called "input perturbation" in the literature. We now introduce our Matrix Dirichlet Mechanism that randomly maps elements of $\mathcal{S}^{(\eta,\bar{\eta})}_{n,V}$ to $\mathcal{S}_n$.

Definition 6 (Matrix Dirichlet Mechanism)**.** The Matrix Dirichlet Mechanism with parameter $k \in \mathbb{R}_+$ is denoted by Dir_M, and it takes as input a stochastic matrix $P \in \mathcal{S}^{(\eta,\bar{\eta})}_{n,V}$, and outputs $\tilde{P} \in \mathcal{S}_n$ via $\tilde{P} \sim \text{Dir}_M(kP)$, where

$$\text{Dir}_M(kP) = \begin{pmatrix} \frac{1}{B(kP_1)} \prod_{j=1}^{n} X_{1j}^{kP_{1j}-1} \\ \frac{1}{B(kP_2)} \prod_{j=1}^{n} X_{2j}^{kP_{2j}-1} \\ \vdots \\ \frac{1}{B(kP_n)} \prod_{j=1}^{n} X_{nj}^{kP_{nj}-1} \end{pmatrix}. \qquad \diamondsuit$$

In words, the Matrix Dirichlet Mechanism takes as input a sensitive stochastic matrix P and outputs a private matrix $\tilde{P}$. Specifically, $\tilde{P}$ is generated row by row based on the rows of P. For a given $i \in [n]$, the process of generating $\tilde{P}_i$ is independent of $\tilde{P}_j$ for $j \neq i$. Definition 6 shows that the Matrix Dirichlet Mechanism outputs $\tilde{P}$ by applying (3) to each row of the matrix, i.e., to P_i for each $i \in [n]$. The parameter k in Definition 6 can be tuned to adjust the level of privacy provided. This will be used in Theorem 1 in Section III-C to quantify the strength of differential privacy. Given η and $\bar{\eta}$, we apply the following assumption to the privacy parameter k.

Assumption 2. The parameter k satisfies

$$k \geq \max\left\{\frac{1}{\eta}, \frac{1}{1-\eta-\bar{\eta}}\right\}. \qquad \diamondsuit$$

Sections III-B and III-C will show that the Matrix Dirichlet Mechanism from Definition 6 satisfies $(\epsilon,\delta)-$differential privacy in Definition 5. For analysis, we interpret the rows of the sensitive matrix P as a disjoint partition of the entire sensitive matrix. Then, the Matrix Dirichlet Mechanism can be viewed as privatizing the elements of this disjoint partition independently, which is a parallel composition as in Lemma 1. Thus, we analyze the privacy guarantees afforded to a single row, then we use parallel composition to make conclusions about the privacy guarantees for the entire sensitive matrix. Specifically, we show that the Matrix Dirichlet Mechanism provides conventional (ϵ,δ)-differential privacy for stochastic matrices by:

- Computing the (ϵ_i,δ_i)-probabilistic differential privacy guarantees from Definition 2 for P_i, for all $i \in [n]$.
- Once (ϵ_i,δ_i)-probabilistic differential privacy for row P_i is established, this implies that conventional differential privacy is satisfied for P_i, for all $i \in [n]$.
- Using parallel composition from Lemma 1, we conclude that the Matrix Dirichlet Mechanism provides conventional (ϵ,δ)-differential privacy for a stochastic matrix P, where $\epsilon = \max_{i\in[n]} \epsilon_i$ and $\delta = \max_{i\in[n]} \delta_i$.

B. Computing δ

In this section, we will compute δ_i for each row P_i. We begin by analyzing the rows P_i for $i \in [n]$, which form a disjoint partition of the input P. We choose W_i in Definition 1 such that it satisfies Assumption 1. We partition the output space of the Matrix Dirichlet Mechanism applied to a row P_i into two disjoint sets, $\mathbf{\Omega}_1^i$ and $\mathbf{\Omega}_2^i$. For all $i \in [n]$, fix $\gamma_i \in (0,1)$ and define the sets $\mathbf{\Omega}_1^i$ and $\mathbf{\Omega}_2^i$ as

$$\mathbf{\Omega}_1^i = \{x \in \Delta_n \mid x_j \geq \gamma_i \text{ for all } j \in W_i\}, \tag{5}$$

and $\mathbf{\Omega}_2^i = \Delta_n \backslash \mathbf{\Omega}_1^i$. We use the following assumption for γ_i.

Assumption 3. Fix $W_i \subseteq [n-1]$. Then $\gamma_i \leq \frac{1}{|W_i|}$. ◊

This assumption is easily satisfied since γ_i is a parameter that we choose. Since $\mathbf{\Omega}_1^i$ and $\mathbf{\Omega}_2^i$ are disjoint sets, we have

$$\mathbb{P}\left[\mathcal{M}_D^{(k)}(P_i) \in \mathbf{\Omega}_2^i\right] = 1 - \mathbb{P}\left[\mathcal{M}_D^{(k)}(P_i) \in \mathbf{\Omega}_1^i\right]. \tag{6}$$

Using Definition 2, we will compute the row-wise probabilities $\mathbb{P}[\mathcal{M}_D^{(k)}(P_i) \in \mathbf{\Omega}_2^i]$, which will give δ_i for each row P_i using probabilistic differential privacy. From (6), we first compute $\mathbb{P}[\mathcal{M}_D^{(k)}(P_i) \in \mathbf{\Omega}_1^i]$ for all $i \in [n]$.

Lemma 2. Fix $n \in \mathbb{N}$, $\eta > 0$ and $\bar{\eta} > 0$. For all $i \in [n]$, fix $W_i \subseteq [n-1]$, let $\mathcal{S}_{n,V}^{(\eta,\bar{\eta})}$ be defined as in Definition 3, and let Assumptions 1, 2, and 3 hold. For all $i \in [n]$, define

$$\mathcal{A}_{r_i} = \left\{ X_i \in \mathbb{R}^{r_i - 1} \mid \sum_{j \in [r_i - 1]} X_{ij} \leq 1, \right.$$
$$\left. X_{ij} \geq \gamma_i \text{ for all } j \in W_i \right\}$$

for all $r_i \geq |W_i| + 1$. Then, for the Matrix Dirichlet Mechanism Dir_M with parameter $k \in \mathbb{R}_+$, we have that

$$\mathbb{P}[\mathcal{M}_D^{(k)}(P_i) \in \mathbf{\Omega}_1^i] = \frac{1}{B\left(k\tilde{P}_{W_i}\right)} \int_{\mathcal{A}_{|W_i|+1}} f\left(X_{ij}\right) \prod_{j \in W_i} dX_{ij},$$

where $f\left(X_{ij}\right)$ equals

$$\prod_{j \in W_i} X_{ij}^{kP_{ij}-1} \left(1 - \sum_{j \in W_i} X_{ij}\right)^{k(1 - \sum\limits_{j \in W_i} P_{ij}) - 1}$$

and $\tilde{P}_{W_i} \in \Delta_{|W_i|+1}$ is equal to P_i after removing entries with indices outside W_i and an entry equal to $1 - \sum_{j \in W_i} P_{ij}$ is appended as its final entry.

Proof. This follows from applying [23, Lemma 1] to each row of P. □

Lemma 2 shows that instead of an $(n-1)$-fold integral of the Dirichlet PDF in (3), the computation of $\mathbb{P}[\mathcal{M}_D^{(k)}(P_i) \in \mathbf{\Omega}_1^i]$ can be reduced to a $|W_i|$-fold integral. However, this result still depends on P_i. This cannot yet be used to show that differential privacy is provided, because doing so requires that privacy is guaranteed for any adjacent input data, not just a specific input. Below, Lemma 3 shows that $\mathbb{P}[\mathcal{M}_D^{(k)}(P_i) \in \mathbf{\Omega}_1^i]$ is a log-concave function of P_i over $\Delta_{n,W_i}^{(\eta,\bar{\eta})}$. Using this, we will compute a value of δ_i that holds for all P_i of interest.

Lemma 3. Fix $n \in \mathbb{N}$, $\eta > 0$ and $\bar{\eta} > 0$. For all $i \in [n]$, fix $W_i \subseteq [n-1]$. Let Assumptions 1, 2, and 3 hold. Let $\mathcal{M}_D^{(k)}$ be the Dirichlet distribution with parameter k from (3). Then $\mathbb{P}[\mathcal{M}_D^{(k)}(P_i) \in \mathbf{\Omega}_1^i]$ is a log-concave function of P_i over the domain $\Delta_{n,W_i}^{(\eta,\bar{\eta})}$.

Proof. The result follows from applying [23, Lemma 2] to each row of P. □

From (6), an upper bound for $\mathbb{P}[\mathcal{M}_D^{(k)}(P_i) \in \mathbf{\Omega}_2^i]$ can be found by minimizing $\mathbb{P}[\mathcal{M}_D^{(k)}(P_i) \in \mathbf{\Omega}_1^i]$, where

$$\mathbb{P}\left[\mathcal{M}_D^{(k)}(P_i) \in \mathbf{\Omega}_2^i\right] \leq 1 - \min_{P_i \in \Delta_{n,W_i}^{(\eta,\bar{\eta})}} \mathbb{P}\left[\mathcal{M}_D^{(k)}(P_i) \in \mathbf{\Omega}_1^i\right] =: \delta_i. \tag{7}$$

From Lemma 3, $\mathbb{P}[\mathcal{M}_D^{(k)}(P_i) \in \mathbf{\Omega}_1^i]$ is a log-concave function and can be easily minimized numerically over the domain $\Delta_{n,W_i}^{(\eta,\bar{\eta})}$. Therefore, δ_i for each row P_i can be readily obtained. This value of δ_i is for probabilistic differential privacy from Definition 2. This implies that the same δ_i can be used in conventional differential privacy from Definition 5. We know that the rows $P_1, \ldots, P_n$ give a disjoint partition of P. Using Lemma 1 we see that, for privacy of the entire stochastic matrix as given in Definition 5, we have $\delta = \max_{i \in [n]} \delta_i$. Next, we will compute ϵ.

C. Computing ϵ

As above, we begin by analyzing a row P_i for some $i \in [n]$. Fix $\eta, \bar{\eta} \in (0,1]$ satisfying Assumption 1, $b \in (0,1]$, and $W_i \subseteq [n-1]$ for all $i \in [n]$. For a given $k \in \mathbb{R}_+$, we must bound the following term to compute ϵ_i:

$$\log\left(\frac{\mathbb{P}\left[\mathcal{M}_D^{(k)}(P_i) = X_i\right]}{\mathbb{P}\left[\mathcal{M}_D^{(k)}(Q_i) = X_i\right]}\right),$$

where $X_i \in \mathbf{\Omega}_1^i$ and P_i and Q_i are b-adjacent in the sense of Definition 4.

The following theorem establishes the differential privacy guarantees of the Matrix Dirichlet distribution and hence solves Problem 1.

Theorem 1. Fix $n \in \mathbb{N}$, $\eta > 0$, $\bar{\eta} > 0$, $b \in (0,1]$, and $W_i \subseteq [n-1]$ for all $i \in [n]$. Let Assumptions 1, 2, and 3 hold. Let the adjacency relation in Definition 4 hold. Then, the Matrix Dirichlet Mechanism with parameter $k \in \mathbb{R}_+$, defined in Definition 6 and denoted as $\mathrm{Dir}_M(kP)$, is (ϵ, δ)-differentially private, where

$$\epsilon = \log\left(\frac{\mathrm{beta}\left(k\eta, k\left(1 - \bar{\eta} - \eta\right)\right)}{\mathrm{beta}\left(k\left(\eta + \frac{b}{2}\right), k\left(1 - \bar{\eta} - \eta - \frac{b}{2}\right)\right)}\right)$$
$$+ \max_{i \in [n]} \frac{kb}{2} \log\left(\frac{1 - \left(|W_i| - 1\right)\gamma_i}{\gamma_i}\right),$$

and $\delta = 1 - \max\limits_{i \in [n]} \min\limits_{P_i \in \Delta_{n,W_i}^{(\eta,\bar{\eta})}} \mathbb{P}\left[\mathcal{M}_D^{(k)}(P_i) \in \mathbf{\Omega}_1^i\right]$.

Proof. We begin by evaluating ϵ_i for a row P_i using probabilistic differential privacy. Fix an arbitrary $i \in [n]$. Let $\ell, m \in W_i$ be the indices from Definition 4 in which P_i and

Q_i can differ. Using Definition 6 and (3), we find that

$$\begin{aligned}
&\log\left(\frac{\mathbb{P}\left[\mathcal{M}_D^{(k)}(P_i)=X_i\right]}{\mathbb{P}\left[\mathcal{M}_D^{(k)}(Q_i)=X_i\right]}\right)\\
&=\log\left(\frac{B\left(kQ_i\right)\prod_{j=1}^{n}X_{ij}^{kP_{ij}-1}}{B\left(kP_i\right)\prod_{j=1}^{n}X_{ij}^{kQ_{ij}-1}}\right)\\
&=\log\left(\frac{\Gamma\left(kQ_{i\ell}\right)\Gamma\left(kQ_{im}\right)X_{i\ell}^{kP_{i\ell}-1}X_{i\ell}^{kP_{im}-1}}{\Gamma\left(kP_{i\ell}\right)\Gamma\left(kP_{im}\right)X_{i\ell}^{kQ_{i\ell}-1}X_{i\ell}^{kQ_{im}-1}}\right)\\
&=\log\left(\frac{\Gamma\left(kQ_{i\ell}\right)\Gamma\left(kQ_{im}\right)}{\Gamma\left(kP_{i\ell}\right)\Gamma\left(kP_{im}\right)}\cdot\left(\frac{X_{i\ell}}{X_{im}}\right)^{k(P_{i\ell}-Q_{i\ell})}\right). \quad (8)
\end{aligned}$$

We find (8) using adjacency, where P_i and Q_i are rows of b-adjacent stochastic matrices. Therefore $P_{i\ell}+P_{im}=Q_{i\ell}+Q_{im}$. For the sake of brevity, we let $G=\frac{\Gamma(kQ_{i\ell})\Gamma(kQ_{im})}{\Gamma(kP_{i\ell})\Gamma(kP_{im})}$. To calculate ϵ, we must upper bound (8). To that end, let

$$\begin{aligned}
\nu=\max_{P_i,Q_i,X_i\in\mathbb{R}^n}&\log\left(G\cdot\left(\frac{X_{i\ell}}{X_{im}}\right)^{k(P_{i\ell}-Q_{i\ell})}\right) \quad (9)\\
\text{subject to }&|P_{i\ell}-Q_{i\ell}|\leq\frac{b}{2}\\
&P_{i\ell}+P_{im}=Q_{i\ell}+Q_{im}\\
&P_{i\ell}+P_{im}\leq 1-\bar{\eta}\\
&P_{(i\ell,im)}\in[\eta,1-\bar{\eta}-\eta]^2\\
&Q_{(i\ell,im)}\in[\eta,1-\bar{\eta}-\eta]^2\\
&X_{(i\ell,im)}\in[\gamma_i,1-(|W_i|-1)\gamma_i]^2.
\end{aligned}$$

Let $\mathcal{C}$ denote the set of feasible points of the optimization problem. The first two constraints enforce adjacency, while the others ensure that $P_i,Q_i\in\Delta_{n,W_i}^{(\eta,\bar{\eta})}$ and $X_i\in\mathbf{\Omega}_1^i$. By sub-additivity of the maximum, we find

$$\nu\leq\max_{P_i,Q_i,X_i\in\mathcal{C}}\log\left(G\right)+\max_{P_i,Q_i,X_i\in\mathcal{C}}\left(\frac{X_{i\ell}}{X_{im}}\right)^{k(P_{i\ell}-Q_{i\ell})}.$$

Let

$$\begin{aligned}
\nu_1&=\max_{P_i,Q_i,X_i\in\mathcal{C}}\left(\frac{X_{i\ell}}{X_{im}}\right)^{k(P_{i\ell}-Q_{i\ell})} &(10)\\
\nu_2&=\max_{P_i,Q_i,X_i\in\mathcal{C}}\log(G). &(11)
\end{aligned}$$

To bound ν_1, we use the fact that $|P_{i\ell}-Q_{i\ell}|\leq\frac{b}{2}$ from adjacency in Definition 4, as well as the inequalities $X_{i\ell}\leq(1-(|W_i|-1))\gamma_i$ and $X_{im}\geq\gamma_i$, which follow from the definition of $\mathbf{\Omega}_1$ in (5). Applying these bounds to (10) gives

$$\begin{aligned}
\nu_1&\leq\max_{P_i,Q_i,X_i\in\mathcal{C}}|k\left(P_{i\ell}-Q_{i\ell}\right)|\cdot\left|\log\left(\frac{X_{i\ell}}{X_{im}}\right)\right|\\
&\leq\frac{kb}{2}\log\left(\frac{1-\left(|W_i|-1\right)\gamma_i}{\gamma_i}\right). \quad (12)
\end{aligned}$$

We now bound ν_2 in (11). Let $c_i=P_{i\ell}+P_{im}=Q_{i\ell}+Q_{im}$, and replace Q_{im} and P_{im} with $c_i-Q_{i\ell}$ and $c_i-P_{i\ell}$, respectively. This encodes the constraints imposed by $\mathcal{C}$ in (11), allowing us to optimize over $\mathbb{R}$. Let

$$\begin{aligned}
\nu_2=\max_{P_{i\ell},Q_{i\ell},c_i\in\mathbb{R}}&\log\left(G\right)\\
\text{subject to }&|P_{i\ell}-Q_{i\ell}|\leq\frac{b}{2}\\
&c_i\in[2\eta,1-\bar{\eta}]\\
&P_{i\ell}\in[\eta,1-\bar{\eta}-\eta]\\
&Q_{i\ell}\in[\eta,1-\bar{\eta}-\eta].
\end{aligned}$$

To compute ν_2, [23] shows that the Karush-Kuhn-Tucker (KKT) conditions of optimality are not satisfied in the interior of the set of feasible points, and we only need to consider the extreme $(P_{i\ell},Q_{i\ell})$'s in the set

$$\begin{aligned}
\Bigg\{&\left(\eta+\frac{b}{2},\eta\right),\left(1-\bar{\eta}-\eta,1-\bar{\eta}-\eta\right),\\
&\left(\eta,\eta+\frac{b}{2}\right),\left(1-\bar{\eta}-\eta,1-\bar{\eta}-\eta-\frac{b}{2}\right)\Bigg\}. \quad (13)
\end{aligned}$$

These points are the vertices of the feasible region $\mathcal{C}$. Note that since $\text{beta}(a,b)=\text{beta}(b,a)$, the points in the first row give equal positive objectives, and the points in the second row give equal negative objectives. Therefore, we can choose the first point in (13) to find

$$\nu_2=\log\left(\frac{\text{beta}\left(k\eta,k\left(1-\bar{\eta}-\eta\right)\right)}{\text{beta}\left(k\left(\eta+\frac{b}{2}\right),k\left(1-\bar{\eta}-\eta-\frac{b}{2}\right)\right)}\right). \quad (14)$$

From (9)-(11), we see that $\epsilon_i\leq\nu_1+\nu_2$. Substituting ν_1 from (12) and ν_2 from (14) gives

$$\begin{aligned}
\epsilon_i\leq\log&\left(\frac{\text{beta}\left(k\eta,k\left(1-\bar{\eta}-\eta\right)\right)}{\text{beta}\left(k\left(\eta+\frac{b}{2}\right),k\left(1-\bar{\eta}-\eta-\frac{b}{2}\right)\right)}\right)\\
&+\frac{kb}{2}\log\left(\frac{1-\left(|W_i|-1\right)\gamma_i}{\gamma_i}\right).
\end{aligned}$$

We use δ_i from (7), which is

$$\delta_i=1-\min_{P_i\in\Delta_{n,W_i}^{(\eta,\bar{\eta})}}\mathbb{P}\left[\mathcal{M}_D^{(k)}(P_i)\in\mathbf{\Omega}_1^i\right].$$

We have shown that the Matrix Dirichlet Mechanism provides (ϵ_i,δ_i)-probabilistic differential privacy for P_i for all $i\in[n]$. This implies that conventional differential privacy is provided for every P_i as well. When we apply Dir_M to P, we see that, for all $i\in[n]$, the output row $\tilde{P}_i$ depends only on the input row P_i. Thus, each row of the private output matrix is equivalent to a query on a disjoint subset of the sensitive input matrix. Therefore, by Lemma 1, the Matrix Dirichlet Mechanism provides (ϵ,δ)-differential privacy to P with $\epsilon=\max_{i\in[n]}\epsilon_i$ and $\delta=\max_{i\in[n]}\delta_i$. □

We have established the (ϵ,δ)-differential privacy guarantees provided by the Matrix Dirichlet Mechanism. Generally, we wish to implement privacy in a way that preserves the usefulness of private data. To that end, we next quantify the trade-off between privacy and accuracy for our mechanism.

D. Accuracy

Through providing (ϵ, δ)-differential privacy, the Matrix Dirichlet Mechanism randomizes the entries of the matrix P. Here, we solve Problem 2 and provide an upper bound on how much privacy perturbs the individual entries of P.

Theorem 2. Fix $n \in \mathbb{N}$, $\eta > 0$, $\bar{\eta} > 0$, $b \in (0,1]$, and $W_i \subseteq [n-1]$ for all $i \in [n]$. Let Assumptions 1, 2, and 3 hold. Let the adjacency relation in Definition 4 hold. Fix a sensitive stochastic matrix $P \in \mathcal{S}_{n,V}^{(\eta,\bar{\eta})}$ and a parameter $k \in \mathbb{R}_+$. Let $\tilde{P} \sim \text{Dir}_M(kP)$, where the Matrix Dirichlet Mechanism Dir_M is given in Definition 6. Then

$$\text{E}\left[|P_{ij} - \tilde{P}_{ij}|\right] \leq \frac{\Gamma(k)\, 2^{1-k}}{\Gamma(k/2)^2 k}, \tag{15}$$

and

$$\text{E}\left[|P_{ij} - \tilde{P}_{ij}|^2\right] \leq \frac{k}{4\,(k^2+k)}. \tag{16}$$

Proof. The expectation of error over the randomness of the Matrix Dirichlet Mechanism takes the form

$$\begin{aligned}\text{E}\left[|P_{ij} - \tilde{P}_{ij}|\right] &= \int_0^1 |P_{ij} - x| \frac{x^{kP_{ij}-1}(1-x)^{k(1-P_{ij})-1}}{\text{beta}(kP_{ij}, k(1-P_{ij}))} dx \\ &= \int_0^{P_{ij}} (P_{ij} - x) \frac{x^{kP_{ij}-1}(1-x)^{k(1-P_{ij})-1}}{\text{beta}(kP_{ij}, k(1-P_{ij}))} dx \\ &- \int_{P_{ij}}^1 (P_{ij} - x) \frac{x^{kP_{ij}-1}(1-x)^{k(1-P_{ij})-1}}{\text{beta}(kP_{ij}, k(1-P_{ij}))} dx.\end{aligned}$$

Expanding, we see that

$$\begin{aligned}\text{E}\left[|P_{ij} - \tilde{P}_{ij}|\right] = P_{ij} &\int_0^{P_{ij}} \frac{x^{kP_{ij}-1}(1-x)^{k(1-P_{ij})-1}}{\text{beta}(kP_{ij}, k(1-P_{ij}))} dx \\ &- \int_0^{P_{ij}} \frac{x^{kP_{ij}}(1-x)^{k(1-P_{ij})-1}}{\text{beta}(kP_{ij}, k(1-P_{ij}))} dx \\ &- P_{ij} \int_{P_{ij}}^1 \frac{x^{kP_{ij}-1}(1-x)^{k(1-P_{ij})-1}}{\text{beta}(kP_{ij}, k(1-P_{ij}))} dx \\ &+ \int_{P_{ij}}^1 \frac{x^{kP_{ij}}(1-x)^{k(1-P_{ij})-1}}{\text{beta}(kP_{ij}, k(1-P_{ij}))} dx. \end{aligned} \tag{17}$$

From the properties of the regularized incomplete beta function in (2), we have

$$\int_0^z x^{a-1}(1-x)^{b-1} dx = \text{beta}(a,b) I_z(a,b),$$

and

$$\int_z^1 x^{a-1}(1-x)^{b-1} dx = \text{beta}(a,b)(1 - I_z(a,b)).$$

Applying these properties to the integrals in (17) gives

$$\begin{aligned}\text{E}\left[|P_{ij} - \tilde{P}_{ij}|\right] &= P_{ij} I_{P_{ij}}(kP_{ij}, k(1-P_{ij})) \\ &- \frac{\text{beta}(kP_{ij}+1, k(1-P_{ij}))}{\text{beta}(kP_{ij}, k(1-P_{ij}))} I_{P_{ij}}(kP_{ij}+1, k(1-P_{ij})) \\ &- P_{ij}(1 - I_{P_{ij}}(kP_{ij}, k(1-P_{ij}))) \\ &+ \frac{\text{beta}(kP_{ij}+1, k(1-P_{ij}))}{\text{beta}(kP_{ij}, k(1-P_{ij}))} \\ &\cdot (1 - I_{P_{ij}}(kP_{ij}+1, k(1-P_{ij}))).\end{aligned}$$

Factoring and collecting like terms we find

$$\begin{aligned}\text{E}\left[|P_{ij} - \tilde{P}_{ij}|\right] &= P_{ij}\left(2I_{P_{ij}}(kP_{ij}, k(1-P_{ij})) - 1\right) \\ &+ \frac{\text{beta}(kP_{ij}+1, k(1-P_{ij}))}{\text{beta}(kP_{ij}, k(1-P_{ij}))}(1 - 2I_{P_{ij}}(kP_{ij}+1, k(1-P_{ij})).\end{aligned}$$

Using the gamma function representation of the beta function from (1) we have

$$\begin{aligned}\text{E}\left[|P_{ij} - \tilde{P}_{ij}|\right] &= P_{ij}\left(2I_{P_{ij}}(kP_{ij}, k(1-P_{ij})) - 1\right) \\ &+ \frac{\Gamma(kP_{ij}+1)\Gamma(k)}{\Gamma(kP_{ij})\Gamma(k+1)}\left(1 - 2I_{P_{ij}}(kP_{ij}+1, k(1-P_{ij}))\right).\end{aligned} \tag{18}$$

We know that $\Gamma(k+1) = k\Gamma(k)$, and therefore

$$\frac{\Gamma(kp+1)\Gamma(k)}{\Gamma(k+1)\Gamma(kp)} = \frac{kp\Gamma(kp)\Gamma(k)}{k\Gamma(k)\Gamma(kp)} = p.$$

Substituting this result into (18) gives

$$\begin{aligned}\text{E}\left[|P_{ij} - \tilde{P}_{ij}|\right] &= 2P_{ij}(I_{P_{ij}}(kP_{ij}, k(1-P_{ij})) \\ &- I_{P_{ij}}(kP_{ij}+1, k(1-P_{ij})).\end{aligned} \tag{19}$$

From [30], the regularized incomplete beta function satisfies

$$I_z(a,b) = I_z(a+1,b) + \frac{z^a(1-z)^b}{a \cdot \text{beta}(a,b)}.$$

Applying this to (19) gives

$$\text{E}\left[|P_{ij} - \tilde{P}_{ij}|\right] = 2\frac{P_{ij}^{kP_{ij}}(1-P_{ij})^{k(1-P_{ij})}}{k \cdot \text{beta}(kP_{ij}, k(1-P_{ij}))}. \tag{20}$$

The maximum value of the right-hand side of (20) occurs at $P_{ij} = 0.5$. Therefore

$$\text{E}\left[|P_{ij} - \tilde{P}_{ij}|\right] \leq \frac{2 \cdot (0.5)^k}{k \cdot \text{beta}(kP_{ij}, k(1-P_{ij}))}.$$

Using the gamma function representation of the beta function from (1) gives

$$\text{E}\left[|P_{ij} - \tilde{P}_{ij}|\right] \leq \frac{\Gamma(k)\, 2^{1-k}}{\Gamma(k/2)^2 k}.$$

This completes the proof for (15).

Next, we will prove (16). By definition,

$$\begin{aligned}
&\mathrm{E}\left[|P_{ij}-\tilde{P}_{ij}|^2\right] \\
&= \int_0^1 (x-P_{ij})^2 \frac{x^{kP_{ij}-1}(1-x)^{k(1-P_{ij})-1}}{\text{beta}\left(kP_{ij},k\left(1-P_{ij}\right)\right)}dx \\
&= \int_0^1 \frac{x^{kP_{ij}+1}(1-x)^{k(1-P_{ij})-1}}{\text{beta}\left(kP_{ij},k\left(1-P_{ij}\right)\right)}dx \\
&- 2P_{ij}\int_0^1 \frac{x^{kP_{ij}}(1-x)^{k(1-P_{ij})-1}}{\text{beta}\left(kP_{ij},k\left(1-P_{ij}\right)\right)}dx \\
&+ P_{ij^2}\int_0^1 \frac{x^{kP_{ij}-1}(1-x)^{k(1-P_{ij})-1}}{\text{beta}\left(kP_{ij},k\left(1-P_{ij}\right)\right)}dx.
\end{aligned}$$

Using the gamma function representation of the beta function from (1) gives

$$\begin{aligned}
\mathrm{E}\left[|P_{ij}-\tilde{P}_{ij}|^2\right] &= \frac{\Gamma\left(kP_{ij}+2\right)\Gamma\left(k\right)}{\Gamma\left(k+2\right)\Gamma\left(kP_{ij}\right)} \\
&- 2P_{ij}\left[\frac{\Gamma\left(kP_{ij}+1\right)\Gamma\left(k\right)}{\Gamma\left(k+1\right)\Gamma\left(kP_{ij}\right)}\right] + P_{ij}^2.
\end{aligned}$$

We simplify using the fact $\Gamma\left(k+1\right)=k\Gamma\left(k\right)$, which gives

$$\begin{aligned}
\mathrm{E}\left[|P_{ij}-\tilde{P}_{ij}|^2\right] &= \frac{(kP_{ij}+1)(kP_{ij})(\Gamma(kP_{ij}))\Gamma(k)}{(k+1)(k)\Gamma(kP_{ij})\Gamma(k+1)} \\
&- 2P_{ij}\left[\frac{kP_{ij}\Gamma(kP_{ij})\Gamma(k)}{k\Gamma(kP_{ij})\Gamma(k)}\right].
\end{aligned}$$

Combining like terms, we find

$$\mathrm{E}\left[|P_{ij}-\tilde{P}_{ij}|^2\right] = \frac{kP_{ij}-P_{ij}^2k}{k^2+k}. \tag{21}$$

The maximum value of (21) occurs at $P_{ij}=0.5$. Therefore,

$$\mathrm{E}\left[|P_{ij}-\tilde{P}_{ij}|^2\right] \leq \frac{k}{4\left(k^2+k\right)}.$$

□

We now consider a numerical example to illustrate Theorem 2. The sensitive matrix is a 10×10 stochastic matrix in which each entry is 0.1. Using an adjacency parameter $b=0.05$, a value of $\gamma_i=0.001$, and W_i so that $|W_i|=4$ for all $i\in[n]$, we apply privacy in the range $\epsilon\in[2.5,22.5]$ and empirically quantify the accuracy from (15) as a function of ϵ. Figure 1 was generated by considering $k\in[10,100]$. For a given value of k, the values of ϵ and δ are calculated and $10,000$ private responses are simulated. To compute the empirical average of the entry-wise error, we first average the error over the entries of each private response, then we average over the $10,000$ responses for the given level of privacy. Figure 1 shows that the upper bound calculated in (15) and the simulated error both monotonically decrease as ϵ increases. Bounding the perturbation of the individual entries of P due to privacy is essential for accurate analysis and implementation of privacy, and Figure 1 shows that Theorem 2 provides an accurate bound on this error.

Solving Problem 2 provides a means of extending the analysis of privatized stochastic matrices to Markov chains, which are widely used models in systems and control theory.

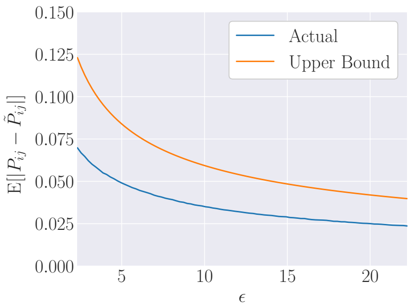

Fig. 1. Average absolute difference between entries of a stochastic matrix P and the private matrix $\tilde{P}$ over $10,000$ applications of Dir_M from Definition 6 as a function of privacy strength. The simulated value is compared to the upper bound presented in (15). For all values of ϵ in this plot, we have $0<\delta\leq 0.027$.

The stationary distribution of a Markov chain is one of its fundamental properties and often used in the analysis of Markov chains. Therefore, we will next bound the perturbation of the stationary distribution as a function of the strength of privacy.

IV. Stationary Distribution Perturbation Bound

In this section, we solve Problem 3 by quantifying how the implementation of privacy using the Matrix Dirichlet Mechanism alters the stationary distribution of a finite, irreducible, homogeneous Markov chain (see Section II-B for a discussion of these terms). The change between the private and non-private stationary distribution of a transition probability matrix can be bounded using perturbation theory. We first define the matrix A as $A=I-P$, where I is the identity matrix and P is the transition probability matrix.

Using A, we can solve for the fundamental matrix of the Markov chain, Z, which is defined as $Z=(A-\mathbb{1}_n\pi^T)^{-1}$, where π is the stationary distribution of the Markov chain. The fundamental matrix Z always exists for a finite, irreducible, homogeneous Markov chain. Suppose the transition probability matrix P is perturbed to $\tilde{P}$ using the Matrix Dirichlet Mechanism. The stationary distribution of the perturbed matrix is denoted as $\tilde{\pi}$. We use the following bound for the distance between the private and non-private stationary distribution.

Lemma 4 (Stationary Distribution Perturbation Bound [31])**.** The norm-wise perturbation bound of the stationary distribution of a finite, homogeneous, irreducible Markov chain is of the form:

$$\|\pi-\tilde{\pi}\|_1 \leq \|Z\|_1\,\|P-\tilde{P}\|_1. \tag{22}$$

□

The privatized transition probability matrix, $\tilde{P}$, is generated by randomizing the original transition probability matrix using the Matrix Dirichlet Mechanism from Definition 6. We therefore now bound $\mathrm{E}\left[\|\pi-\tilde{\pi}\|_1\right]$.

Theorem 3. Fix $n\in\mathbb{N}$, $\eta>0$, $\bar{\eta}>0$, $b\in(0,1]$, and $W_i\subseteq[n-1]$ for all $i\in[n]$. Let Assumptions 1, 2, and 3 hold. Fix $n\in\mathbb{N}$ and consider a stochastic matrix $P\in\mathcal{S}_{n,V}^{(\eta,\bar{\eta})}$ that corresponds to a finite, homogeneous, irreducible Markov chain. Let π denote its stationary distribution. Let $\tilde{P} \sim \mathrm{Dir}_M(kP)$ be its privatized form with stationary distribution $\tilde{\pi}$. Then the expected distance between the private and non-private stationary distribution can be upper bounded as

$$\mathrm{E}\left[\|\pi-\tilde{\pi}\|_1\right] \leq \|Z\|_1 \cdot \left\{ \frac{n\Gamma\left(k\right)2^{1-k}}{\Gamma\left(k/2\right)^2 k} + \left(\frac{n-1}{n}\right)^{\frac{1}{2}} \right. \cdot \left\{ \frac{n^2 k}{4(k^2+k)} - \max_i \left[4\frac{\eta^{2k[1-\bar{\eta}-(|W_i|-1)\eta]}\cdot[\bar{\eta}+(|W_i|-1)\eta]^{2k(1-\eta)}}{k^2\cdot \mathrm{beta}(k\eta, k[\bar{\eta}+(|W_i|-1)\eta])^2} \right] \right\}^{\frac{1}{2}} \left. \vphantom{\frac{1}{1}} \right\},$$

where $Z=(A-\mathbb{1}_n\pi^T)^{-1}$ and $|W_i|$ is the number of indices in W_i for a given row P_i.

Proof. The matrix 1-norm is the maximum absolute column sum of its argument. The 1-norm of the perturbation induced by the Matrix Dirichlet Mechanism can be expressed as $\|P-\tilde{P}\|_1 = \max_{j\in[n]}\sum_{i=1}^n |P_{ij}-\tilde{P}_{ij}|$. The perturbation of P by the Matrix Dirichlet Mechanism is a random variable. As a result, we begin this proof by taking the expected value of both sides of (22) in Lemma 4 to find

$$\mathrm{E}\left[\|\pi-\tilde{\pi}\|_1\right] \leq \|Z\|_1\, \mathrm{E}\left[\max_{j\in[n]}\sum_{i=1}^n |P_{ij}-\tilde{P}_{ij}|\right]. \tag{23}$$

Now we bound the expectation of the maximum of a collection of random variables. Let $X_1, X_2, \ldots, X_N$ $(2 \leq N < \infty)$ be an arbitrary sequence of real-valued random variables with finite mean and variance. The bound in [32, Theorem 2.1] gives

$$\mathrm{E}\left[\max_{j\in[N]} X_j\right] \leq \max_{j\in[N]} E[X_j] + \left(\frac{N-1}{N}\sum_{j=1}^N \mathrm{Var}[X_j]\right)^{\frac{1}{2}} \tag{24}$$

Now let $X_j = \sum_{i=1}^n |P_{ij}-\tilde{P}_{ij}|$. Then applying (24) to (23) with $N=n$ gives

$$\mathrm{E}\left[\max_{j\in[n]}\sum_{i=1}^n |P_{ij}-\tilde{P}_{ij}|\right] \leq \max_{j\in[n]} \mathrm{E}\left[\sum_{i=1}^n |P_{ij}-\tilde{P}_{ij}|\right] + \left(\frac{n-1}{n}\right)^{\frac{1}{2}} \cdot \left\{\sum_{j=1}^n \mathrm{Var}\left[\sum_{i=1}^n |P_{ij}-\tilde{P}_{ij}|\right]\right\}^{\frac{1}{2}}. \tag{25}$$

Applying the upper bound from (15) to the first term in (25), we can further simplify the above equation as

$$\mathrm{E}\left[\max_{j\in[n]}\sum_{i=1}^n |P_{ij}-\tilde{P}_{ij}|\right] \leq \frac{n\Gamma\left(k\right)2^{1-k}}{\Gamma\left(k/2\right)^2 k} + \left(\frac{n-1}{n}\right)^{\frac{1}{2}} \cdot \left\{\sum_{j=1}^n \mathrm{Var}\left[\sum_{i=1}^n |P_{ij}-\tilde{P}_{ij}|\right]\right\}^{\frac{1}{2}}. \tag{26}$$

We now focus on the variance term in brackets in (26). The sum inside the variance is over the rows of P, which are randomized in a mutually independent way. Thus, the expression inside the variance is the sum of independent random variables and is equivalent to the sum of the variances of the individual terms. Therefore,

$$\sum_{j=1}^n \mathrm{Var}\left[\sum_{i=1}^n |P_{ij}-\tilde{P}_{ij}|\right] = \sum_{i=1}^n\sum_{j=1}^n \mathrm{Var}[|P_{ij}-\tilde{P}_{ij}|]. \tag{27}$$

For a random variable X we have $\mathrm{Var}[X]=\mathrm{E}[X^2]-\mathrm{E}[X]^2$. Applying this fact to (27) gives

$$\sum_{i=1}^n\sum_{j=1}^n \mathrm{Var}[|P_{ij}-\tilde{P}_{ij}|] = \tag{28}$$
$$\sum_{i=1}^n\left\{\sum_{j=1}^n \mathrm{E}\left[|P_{ij}-\tilde{P}_{ij}|^2\right] - \sum_{j=1}^n \mathrm{E}\left[|P_{ij}-\tilde{P}_{ij}|\right]^2\right\}.$$

To upper bound the first term in (28), we apply (16) from Theorem 2, which gives

$$\sum_{i=1}^n\sum_{j=1}^n \mathrm{Var}[|P_{ij}-\tilde{P}_{ij}|] \leq \sum_{i=1}^n \frac{nk}{4(k^2+k)} - \sum_{i=1}^n\sum_{j=1}^n \mathrm{E}\left[|P_{ij}-\tilde{P}_{ij}|\right]^2. \tag{29}$$

Next, we substitute (20) for the second term, and upper bound the sum of the first term in (29) to find

$$\sum_{i=1}^n\sum_{j=1}^n \mathrm{Var}[|P_{ij}-\tilde{P}_{ij}|] \leq \frac{n^2k}{4(k^2+k)} - \sum_{i=1}^n\sum_{j=1}^n 4\frac{P_{ij}^{2kP_{ij}}(1-P_{ij})^{2k(1-P_{ij})}}{k^2\cdot \mathrm{beta}(kP_{ij}, k(1-P_{ij}))^2}. \tag{30}$$

We now focus on lower bounding the second term in (30). By definition, for any $j\in W_i$ we have $P_{ij}\geq\eta$. All other $P_{ij}\geq 0$. Therefore, we solely examine P_{ij} for $j\in W_i$ to find a lower bound, where

$$\sum_{j=1}^n 4\frac{P_{ij}^{2kP_{ij}}(1-P_{ij})^{2k(1-P_{ij})}}{k^2\cdot \mathrm{beta}(kP_{ij}, k(1-P_{ij}))^2} \geq \sum_{j\in W_i} 4\frac{P_{ij}^{2kP_{ij}}(1-P_{ij})^{2k(1-P_{ij})}}{k^2\cdot \mathrm{beta}(kP_{ij}, k(1-P_{ij}))^2}. \tag{31}$$

From Assumption 1, we know that $\eta < 1$. Definition 1 implies that for P_{ij} with $j\in W_i$ we have

$$\eta \leq P_{ij} \leq 1-\bar{\eta}-(|W_i|-1)\eta, \tag{32}$$

and

$$\bar{\eta}+(|W_i|-1)\eta \leq 1-P_{ij} \leq 1-\eta. \tag{33}$$

Next, we find a lower bound for the numerator and an upper bound for the denominator of (31). Making the numerator as small as possible by substituting (32) and (33), we find

$$P_{ij}^{2kP_{ij}}(1-P_{ij})^{2k(1-P_{ij})} \geq \eta^{2k[1-\bar{\eta}-(|W_i|-1)\eta]} \cdot [\bar{\eta}+(|W_i|-1)\eta]^{2k(1-\eta)}. \tag{34}$$

Next, we lower bound the denominator of (31). To do so, we show that $\text{beta}(a, b)$ is monotonically decreasing in a and b, where $a, b > 0$. We begin by finding the partial derivatives of $\text{beta}(a, b)$ with respect to a and b, namely

$$\frac{\partial}{\partial a}\text{beta}(a, b) = \text{beta}(a, b)(\psi(a) - \psi(a + b)) \tag{35}$$

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

where $\psi(x)$ is the digamma function. Note that $\text{beta}(a, b)$ is strictly positive for $a, b > 0$. From [33], $\psi(x)$ is strictly increasing for $x > 0$, which implies $\psi(a) - \psi(a + b) < 0$ and $\psi(b) - \psi(a + b) < 0$. Thus, (35) and (36) are strictly negative and $\text{beta}(a, b)$ is monotonically decreasing in a and b. We apply this property to the denominator of (31), where

$$k^2 \cdot \text{beta}(kP_{ij}, k(1 - P_{ij}))^2 \leq k^2 \cdot \text{beta}(k\eta, k[\bar{\eta} + (|W_i| - 1)\eta])^2. \tag{37}$$

Substituting (34) and (37) into (31), we find

$$\sum_{j\in W_i} 4\frac{P_{ij}^{2kP_{ij}}(1 - P_{ij})^{2k(1-P_{ij})}}{k^2 \cdot \text{beta}(kP_{ij}, k(1 - P_{ij}))^2} \geq \sum_{j\in W_i} 4\frac{\eta^{2k[1-\bar{\eta}-(|W_i|-1)\eta]} \cdot [\bar{\eta} + (|W_i| - 1)\eta]^{2k(1-\eta)}}{k^2 \cdot \text{beta}(k\eta, k[\bar{\eta} + (|W_i| - 1)\eta])^2}. \tag{38}$$

Next, we substitute (38) into (30), which gives

$$\sum_{i=1}^{n}\sum_{j=1}^{n}\text{Var}[|P_{ij} - \tilde{P}_{ij}|] \leq \frac{n^2k}{4(k^2 + k)} - \sum_{i=1}^{n}\sum_{j\in W_i} 4\frac{\eta^{2k[1-\bar{\eta}-(|W_i|-1)\eta]} \cdot [\bar{\eta} + (|W_i| - 1)\eta]^{2k(1-\eta)}}{k^2 \cdot \text{beta}(k\eta, k[\bar{\eta} + (|W_i| - 1)\eta])^2}. \tag{39}$$

To upper bound (39) we lower bound the sum in the second term. A sum of non-negative terms is lower bounded by any one of the summands. Taking the single largest value gives the tightest lower bound of this form, and this results in

$$\sum_{i=1}^{n}\sum_{j=1}^{n}\text{Var}[|P_{ij} - \tilde{P}_{ij}|] \leq \frac{n^2k}{4(k^2 + k)} - \max_i\left[4\frac{\eta^{2k[1-\bar{\eta}-(|W_i|-1)\eta]} \cdot [\bar{\eta} + (|W_i| - 1)\eta]^{2k(1-\eta)}}{k^2 \cdot \text{beta}(k\eta, k[\bar{\eta} + (|W_i| - 1)\eta])^2}\right]. \tag{40}$$

Next, we substitute the result from (40) into (26) to find

$$\text{E}\left[\max_{j\in[n]}\sum_{i=1}^{n}|P_{ij} - \tilde{P}_{ij}|\right] \leq \frac{n\Gamma(k)\,2^{1-k}}{\Gamma(k/2)^2\,k} + \left(\frac{n-1}{n}\right)^{\frac{1}{2}} \cdot \left\{\frac{n^2k}{4(k^2 + k)} - \max_i\left[4\frac{\eta^{2k[1-\bar{\eta}-(|W_i|-1)\eta]} \cdot [\bar{\eta} + (|W_i| - 1)\eta]^{2k(1-\eta)}}{k^2 \cdot \text{beta}(k\eta, k[\bar{\eta} + (|W_i| - 1)\eta])^2}\right]\right\}^{\frac{1}{2}}. \tag{41}$$

To conclude the proof, we substitute (41) into (23) to find the expression of interest. □

The result presented in Theorem 3 provides us with a quantitative measure of how much the computation of the stationary distribution is expected to change as a function of the strength of privacy. The accuracy of the stationary distribution is crucial in Markov chain models, especially when predicting the future behavior of a Markov chain, and we illustrate this on a practical example in the next section.

V. Simulation Results

This section presents simulation results. We consider a Markov chain model of a traffic system generated from the Annual Average Daily Traffic (AADT) of some of the major streets in Gainesville, Florida from 2021 [34]. A trajectory produced by this Markov chain represents a user's route through Gainesville, and such routes are sensitive. For example, they may reveal a user's place of work, their home, or other private activities conducted in their daily routine. We implement differential privacy for this system using the Matrix Dirichlet Mechanism.

The AADT data used to compute the transition probabilities was provided by Florida Traffic Online, a mapping application that shows historical traffic count site locations and data. AADT numbers present the average daily traffic volume on segments of roads over the course of one year. The transition probabilities for the Markov chain model were found using frequency analysis, where the transition probability from one state to another feasible state is equal to the number of times a driver transitions from the first state to the second state divided by the total number of times a driver transitions away from the first state.

The Markov chain model we construct contains 32 states. In this example, the support of the Markov chain model is assumed to be public knowledge. For example, any outside observer can conclude that there is no probability of transitioning from Old Archer Rd to SW 2nd Ave simply by observing the layout of the streets. To this end, the entries in P equal to 0 are not perturbed using the Matrix Dirichlet Mechanism. Analysis in this example requires the input to be in the interior of the unit simplex. To accommodate this, we define a mechanism on $\Delta_{|S_i|}$, where $|S_i|$ is the number of non-zero entries in row i.

As ϵ shrinks and privacy strengthens, the entries of P are perturbed more by privacy, which will affect predictions of routes taken by users. The first transition probability matrix P was derived directly from the AADT traffic statistics and is treated as the sensitive data. The second transition probability matrix $\tilde{P}_1$ was generated by privatizing P using Dir_M with a parameter $k = 9.87$ corresponding to $\epsilon = 1.16$ and $\delta = 0.011$. The final transition probability matrix $\tilde{P}_2$ was generated by privatizing P using Dir_M with a parameter $k = 98.7$ corresponding to $\epsilon = 11.12$ and $\delta = 0.019$. For both applications of Dir_M, an adjacency parameter $b = 0.025$ was selected. We fix $\eta = 0.10$ and $\bar{\eta} = 0.051$. The set W_i was selected as the indices of the $n - 1$ largest non-zero

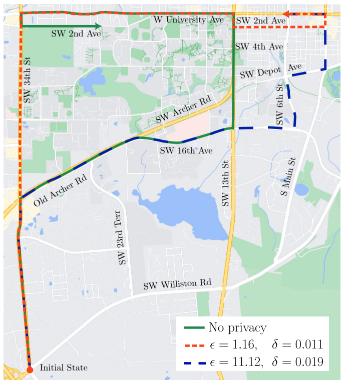

Fig. 2. A random walk over 10 states of the Markov chain model of Gainesville, Florida conducted for varying levels of privacy. As ϵ increases, the sampled path becomes more similar to the random walk without privacy.

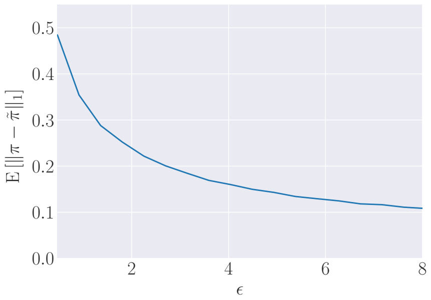

Fig. 3. Using an adjacency parameter of $b = 0.025$ and $\gamma_i = 0.001$ for all $i \in [n]$, the expected value of the error between the private and non-private stationary distribution was found over $10,000$ applications of the Matrix Dirichlet Mechanism at each value of $\epsilon \in [0.5, 9]$, with $0 < \delta \leq 0.011$. As ϵ increases, there is a clear decreasing trend in this error, which shows that weaker privacy provides greater accuracy for the stationary distribution, and vice versa.

entries in each row of P, and we set $\gamma_i = 0.001$ for all $i \in [n]$.

A random walk for 10 steps was then performed for each of the 3 transition probability matrices and the results are illustrated in Figure 2. Figure 2 shows that as ϵ increases and the strength of privacy decreases, the random walk will tend to have similar state transitions to those in the original transition probability matrix P. This is intuitive since as ϵ increases, the strength of privacy protections decreases, and P is perturbed by a smaller amount.

We now quantify the accuracy of the perturbed transition probability matrices by evaluating their corresponding stationary distributions. Figure 3 shows the relationship between the strength of privacy and the 1-norm of the difference between the private and non-private stationary distributions. Figure 3 was generated by considering $k \in [10, 200]$. For a given value of k, the values of ϵ and δ were calculated and $10,000$ private matrices were generated. The stationary distribution $\tilde{\pi}$ was calculated for each one. To compute the empirical average of the 1-norm difference, the average stationary distribution over the $10,000$ private responses was calculated and subtracted from the original stationary distribution, π.

Figure 3 shows that the simulated error monotonically decreases as privacy increases. Figure 3 shows a maximum expected error of 0.485 corresponding to $\epsilon = 0.5$ and a minimum error of 0.101 corresponding to $\epsilon = 8$. With respect to the maximum possible error (which is 2), the private stationary distribution error ranges from 5.05% to 24.25% under typical privacy conditions. This error is spread over the 32 entries of the stationary distribution implying that the error associated with a single state is quite small, even for strong privacy guarantees.

VI. Conclusion

This paper introduced the Matrix Dirichlet Mechanism as a means of providing differential privacy to stochastic matrices. We proved that this mechanism satisfies conventional differential privacy guarantees, and quantified the error induced in private stochastic matrices as a function of the strength of privacy provided. We then applied the Matrix Dirichlet Mechanism to Markov chains and calculated the distance between the stationary distribution of the sensitive stochastic matrix and the stationary distribution of its privatized form. Future work includes an extension of the Matrix Dirichlet Mechanism to provide differential privacy to doubly stochastic matrices.

References

[1] F. Fioretto, T. W. Mak, and P. Van Hentenryck, “Differential privacy for power grid obfuscation,” *IEEE Transactions on Smart Grid*, vol. 11, no. 2, pp. 1356–1366, 2019.

[2] Y. Li, D. Yang, and X. Hu, “A differential privacy-based privacy-preserving data publishing algorithm for transit smart card data,” *Transportation Research Part C: Emerging Technologies*, vol. 115, p. 102634, 2020.

[3] A. Prorok and V. Kumar, “A macroscopic model for differential privacy in dynamic robotic networks,” *arXiv preprint arXiv:1703.04797*, 2017.

[4] M. R. Alam, M. B. I. Reaz, and M. M. Ali, “Speed: An inhabitant activity prediction algorithm for smart homes,” *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, vol. 42, no. 4, pp. 985–990, 2011.

[5] Q. Gong, S. Midlam-Mohler, V. Marano, and G. Rizzoni, “An iterative markov chain approach for generating vehicle driving cycles,” *SAE International Journal of Engines*, vol. 4, no. 1, pp. 1035–1045, 2011.

[6] A. Asahara, K. Maruyama, A. Sato, and K. Seto, “Pedestrian-movement prediction based on mixed markov-chain model,” in *Proceedings of the 19th ACM SIGSPATIAL international conference on advances in geographic information systems*, 2011, pp. 25–33.

[7] M. Hale, P. Barooah, K. Parker, and K. Yazdani, “Differentially private smart metering: Implementation, analytics, and billing,” in *Proceedings of the 1st ACM International Workshop on Urban Building Energy Sensing, Controls, Big Data Analysis, and Visualization*. Association for Computing Machinery, 2019, p. 33–42.

[8] J. Widén, A. M. Nilsson, and E. Wäckelgård, “A combined markov-chain and bottom-up approach to modelling of domestic lighting demand,” *Energy and Buildings*, vol. 41, no. 10, pp. 1001–1012, 2009.

[9] S. Vermeer and D. Trilling, “Toward a better understanding of news user journeys: A markov chain approach,” *Journalism Studies*, vol. 21, no. 7, pp. 879–894, 2020.

[10] J. Dong, G. Li, W. Ma, and J. Liu, “Personalized recommendation system based on social tags in the era of internet of things,” *Journal of Intelligent Systems*, vol. 31, no. 1, pp. 681–689, 2022.

[11] D. A. Levin and Y. Peres, *Markov chains and mixing times*. American Mathematical Soc., 2017, vol. 107.

[12] J. Lundström, E. Järpe, and A. Verikas, “Detecting and exploring deviating behaviour of smart home residents,” *Expert Systems with Applications*, vol. 55, pp. 429–440, 2016.

[13] S. Rendle, C. Freudenthaler, and L. Schmidt-Thieme, “Factorizing personalized markov chains for next-basket recommendation,” in *Proceedings of the 19th international conference on World wide web*, 2010, pp. 811–820.

[14] C. Dwork, F. McSherry, K. Nissim, and A. Smith, “Calibrating noise to sensitivity in private data analysis,” in *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*. Springer, 2006, pp. 265–284.

[15] C. Dwork, A. Roth *et al.*, “The algorithmic foundations of differential privacy,” *Foundations and Trends in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

[16] S. P. Kasiviswanathan and A. Smith, “On the’semantics’ of differential privacy: A bayesian formulation,” *Journal of Privacy and Confidentiality*, vol. 6, no. 1, 2014.

[17] C. Hawkins and M. Hale, “Differentially private formation control,” in *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 2020, pp. 6260–6265.

[18] G. Yuan, Y. Yang, Z. Zhang, and Z. Hao, “Convex optimization for linear query processing under approximate differential privacy,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 2005–2014.

[19] J. Le Ny, “Differentially private kalman filtering,” *Differential Privacy for Dynamic Data*, pp. 55–75, 2020.

[20] B. Chen, K. Leahy, A. Jones, and M. Hale, “Differential privacy for symbolic systems with application to markov chains,” *Automatica*, vol. 152, p. 110908, 2023.

[21] Y. Wang, M. Hale, M. Egerstedt, and G. E. Dullerud, “Differentially private objective functions in distributed cloud-based optimization,” in *2016 IEEE 55th Conference on Decision and Control (CDC)*, 2016, pp. 3688–3694.

[22] P. Gohari, M. Hale, and U. Topcu, “Privacy-preserving policy synthesis in markov decision processes,” in *2020 59th IEEE Conference on Decision and Control (CDC)*, 2020, pp. 6266–6271.

[23] P. Gohari, B. Wu, C. Hawkins, M. Hale, and U. Topcu, “Differential privacy on the unit simplex via the dirichlet mechanism,” *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 2326–2340, 2021.

[24] C. Dwork and J. Lei, “Differential privacy and robust statistics,” in *Proceedings of the forty-first annual ACM symposium on Theory of computing*, 2009, pp. 371–380.

[25] A. Machanavajjhala, D. Kifer, J. Abowd, J. Gehrke, and L. Vilhuber, “Privacy: Theory meets practice on the map,” in *2008 IEEE 24th international conference on data engineering*. IEEE, 2008, pp. 277–286.

[26] J. Hsu, M. Gaboardi, A. Haeberlen, S. Khanna, A. Narayan, B. C. Pierce, and A. Roth, “Differential privacy: An economic method for choosing epsilon,” in *2014 IEEE 27th Computer Security Foundations Symposium*. IEEE, 2014, pp. 398–410.

[27] C. Hawkins, B. Chen, K. Yazdani, and M. Hale, “Node and edge differential privacy for graph laplacian spectra: Mechanisms and scaling laws,” *arXiv preprint arXiv:2211.15366*, 2022.

[28] M. Gotz, A. Machanavajjhala, G. Wang, X. Xiao, and J. Gehrke, “Publishing search logs—a comparative study of privacy guarantees,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 24, no. 3, pp. 520–532, 2011.

[29] N. Ponomareva, H. Hazimeh, A. Kurakin, Z. Xu, C. Denison, H. B. McMahan, S. Vassilvitskii, S. Chien, and A. Thakurta, “How to dp-fy ml: A practical guide to machine learning with differential privacy,” *arXiv preprint arXiv:2303.00654*, 2023.

[30] “8.17 incomplete beta functions.” [Online]. Available: https://dlmf.nist.gov/8.17

[31] E. Seneta, “Sensitivity to perturbation of the stationary distribution: some refinements,” *Linear Algebra and its Applications*, vol. 108, pp. 121–126, 1988.

[32] T. Aven, “Upper (lower) bounds on the mean of the maximum (minimum) of a number of random variables,” *Journal of applied probability*, vol. 22, no. 3, pp. 723–728, 1985.

[33] H. Alzer and G. Jameson, “A harmonic mean inequality for the digamma function and related results,” *Rendiconti del Seminario Matematico della Università di Padova*, vol. 137, pp. 203–209, 2017.

[34] “Florida traffic online.” [Online]. Available: https://tdaappsprod.dot.state.fl.us/fto/