
Mixing Time Estimation in Reversible Markov Chains from a Single Sample Path

Daniel Hsu
Columbia University
djhsu@cs.columbia.edu

Aryeh Kontorovich
Ben-Gurion University
karyeh@cs.bgu.ac.il

Csaba Szepesvári
University of Alberta
szepesva@cs.ualberta.ca

Abstract

This article provides the first procedure for computing a fully data-dependent interval that traps the mixing time t_{mix} of a finite reversible ergodic Markov chain at a prescribed confidence level. The interval is computed from a single finite-length sample path from the Markov chain, and does not require the knowledge of any parameters of the chain. This stands in contrast to previous approaches, which either only provide point estimates, or require a reset mechanism, or additional prior knowledge. The interval is constructed around the relaxation time t_{relax} , which is strongly related to the mixing time, and the width of the interval converges to zero roughly at a $\sqrt{n}$ rate, where n is the length of the sample path. Upper and lower bounds are given on the number of samples required to achieve constant-factor multiplicative accuracy. The lower bounds indicate that, unless further restrictions are placed on the chain, no procedure can achieve this accuracy level before seeing each state at least $\Omega(t_{\text{relax}})$ times on the average. Finally, future directions of research are identified.

1 Introduction

This work tackles the challenge of constructing fully empirical bounds on the mixing time of Markov chains based on a single sample path. Let $(X_t)_{t=1,2,\dots}$ be an irreducible, aperiodic time-homogeneous Markov chain on a finite state space $[d] := \{1, 2, \dots, d\}$ with transition matrix P . Under this assumption, the chain converges to its unique stationary distribution $\pi = (\pi_i)_{i=1}^d$ regardless of the initial state distribution q :

$$\lim_{t \rightarrow \infty} \Pr_q(X_t = i) = \lim_{t \rightarrow \infty} (qP^t)_i = \pi_i \quad \text{for each } i \in [d].$$

The *mixing time* t_{mix} of the Markov chain is the number of time steps required for the chain to be within a fixed threshold of its stationary distribution:

$$t_{\text{mix}} := \min \left\{ t \in \mathbb{N} : \sup_q \max_{A \subset [d]} |\Pr_q(X_t \in A) - \pi(A)| \leq 1/4 \right\}. \quad (1)$$

Here, $\pi(A) = \sum_{i \in A} \pi_i$ is the probability assigned to set A by π , and the supremum is over all possible initial distributions q . The problem studied in this work is the construction of a non-trivial confidence interval $C_n = C_n(X_1, X_2, \dots, X_n, \delta) \subset [0, \infty]$, based only on the observed sample path $(X_1, X_2, \dots, X_n)$ and $\delta \in (0, 1)$, that succeeds with probability $1 - \delta$ in trapping the value of the mixing time t_{mix} .

This problem is motivated by the numerous scientific applications and machine learning tasks in which the quantity of interest is the mean $\pi(f) = \sum_i \pi_i f(i)$ for some function f of the states of a Markov chain. This is the setting of the celebrated Markov Chain Monte Carlo (MCMC) paradigm [1], but the problem also arises in performance prediction involving time-correlated data, as is common in reinforcement learning [2]. Observable bounds on mixing times are useful in the

design and diagnostics of these methods; they yield effective approaches to assessing the estimation quality, even when *a priori* knowledge of the mixing time or correlation structure is unavailable.

Main results. We develop the first procedure for constructing non-trivial and fully empirical confidence intervals for Markov mixing time. Consider a reversible ergodic Markov chain on d states with absolute spectral gap γ_* and stationary distribution minorized by π_* . As is well-known [3, Theorems 12.3 and 12.4],

$$(t_{\text{relax}} - 1) \ln 2 \leq t_{\text{mix}} \leq t_{\text{relax}} \ln \frac{4}{\pi_*} \quad (2)$$

where $t_{\text{relax}} := 1/\gamma_*$ is the *relaxation time*. Hence, it suffices to estimate γ_* and π_* . Our main results are summarized as follows.

1. In Section 3.1, we show that in some problems $n = \Omega((d \log d)/\gamma_* + 1/\pi_*)$ observations are necessary for any procedure to guarantee constant multiplicative accuracy in estimating γ_* (Theorems 1 and 2). Essentially, in some problems *every* state may need to be visited about $\log(d)/\gamma_*$ times, on average, before an accurate estimate of the mixing time can be provided, regardless of the actual estimation procedure used.
2. In Section 3.2, we give a point-estimator for γ_* , and prove in Theorem 3 that it achieves multiplicative accuracy from a *single sample path* of length $\tilde{O}(1/(\pi_* \gamma_*^3))$.¹ We also provide a point-estimator for π_* that requires a sample path of length $\tilde{O}(1/(\pi_* \gamma_*))$. This establishes the feasibility of *estimating* the mixing time in this setting. However, the valid confidence intervals suggested by Theorem 3 depend on the unknown quantities π_* and γ_* . We also discuss the importance of reversibility, and some possible extensions to non-reversible chains.
3. In Section 4, the construction of valid *fully empirical confidence intervals* for π_* and γ_* are considered. First, the difficulty of the task is explained, i.e., why the standard approach of turning the finite time confidence intervals of Theorem 3 into a fully empirical one fails. Combining several results from perturbation theory in a novel fashion we propose a new procedure and prove that it avoids slow convergence (Theorem 4). We also explain how to combine the empirical confidence intervals from Algorithm 1 with the non-empirical bounds from Theorem 3 to produce valid empirical confidence intervals. We prove in Theorem 5 that the width of these new intervals converge to zero asymptotically at least as fast as those from either Theorem 3 and Theorem 4.

Related work. There is a vast statistical literature on estimation in Markov chains. For instance, it is known that under the assumptions on $(X_t)_t$ from above, the law of large numbers guarantees that the sample mean $\pi_n(f) := \frac{1}{n} \sum_{t=1}^n f(X_t)$ converges almost surely to $\pi(f)$ [4], while the central limit theorem tells us that as $n \rightarrow \infty$, the distribution of the deviation $\sqrt{n}(\pi_n(f) - \pi(f))$ will be normal with mean zero and asymptotic variance $\lim_{n \rightarrow \infty} n \text{Var}(\pi_n(f))$ [5].

Although these asymptotic results help us understand the limiting behavior of the sample mean over a Markov chain, they say little about the finite-time non-asymptotic behavior, which is often needed for the prudent evaluation of a method or even its algorithmic design [6–13]. To address this need, numerous works have developed Chernoff-type bounds on $\Pr(|\pi_n(f) - \pi(f)| > \epsilon)$, thus providing valuable tools for non-asymptotic probabilistic analysis [6, 14–16]. These probability bounds are larger than corresponding bounds for independent and identically distributed (iid) data due to the temporal dependence; intuitively, for the Markov chain to yield a fresh draw $X_{t'}$ that behaves as if it was independent of X_t , one must wait $\Theta(t_{\text{mix}})$ time steps. Note that the bounds generally depend on distribution-specific properties of the Markov chain (e.g., $\mathbf{P}$, t_{mix} , γ_*), which are often unknown *a priori* in practice. Consequently, much effort has been put towards estimating these unknown quantities, especially in the context of MCMC diagnostics, in order to provide data-dependent assessments of estimation accuracy [e.g., 11, 12, 17–19]. However, these approaches generally only provide asymptotic guarantees, and hence fall short of our goal of empirical bounds that are valid with any finite-length sample path.

Learning with dependent data is another main motivation to our work. Many results from statistical learning and empirical process theory have been extended to sufficiently fast mixing, dependent

¹The $\tilde{O}(\cdot)$ notation suppresses logarithmic factors.

data [e.g., 20–26], providing learnability assurances (e.g., generalization error bounds). These results are often given in terms of mixing coefficients, which can be consistently estimated in some cases [27]. However, the convergence rates of the estimates from [27], which are needed to derive confidence bounds, are given in terms of unknown mixing coefficients. When the data comes from a Markov chain, these mixing coefficients can often be bounded in terms of mixing times, and hence our main results provide a way to make them fully empirical, at least in the limited setting we study.

It is possible to eliminate many of the difficulties presented above when allowed more flexible access to the Markov chain. For example, given a sampling oracle that generates independent transitions from any given state (akin to a “reset” device), the mixing time becomes an efficiently testable property in the sense studied in [28, 29]. On the other hand, when one only has a circuit-based description of the transition probabilities of a Markov chain over an exponentially-large state space, there are complexity-theoretic barriers for many MCMC diagnostic problems [30].

2 Preliminaries

2.1 Notations

We denote the set of positive integers by $\mathbb{N}$, and the set of the first d positive integers $\{1, 2, \dots, d\}$ by $[d]$. The non-negative part of a real number x is $[x]_+ := \max\{0, x\}$, and $\lceil x \rceil_+ := \max\{0, \lceil x \rceil\}$. We use $\ln(\cdot)$ for natural logarithm, and $\log(\cdot)$ for logarithm with an arbitrary constant base. Bold-face symbols are used for vectors and matrices (e.g., $\mathbf{v}$, $\mathbf{M}$), and their entries are referenced by subindexing (e.g., v_i , $M_{i,j}$). For a vector $\mathbf{v}$, $\|\mathbf{v}\|$ denotes its Euclidean norm; for a matrix $\mathbf{M}$, $\|\mathbf{M}\|$ denotes its spectral norm. We use $\text{Diag}(\mathbf{v})$ to denote the diagonal matrix whose (i, i) -th entry is v_i . The probability simplex is denoted by $\Delta^{d-1} = \{\mathbf{p} \in [0, 1]^d : \sum_{i=1}^d p_i = 1\}$, and we regard vectors in Δ^{d-1} as row vectors.

2.2 Setting

Let $\mathbf{P} \in (\Delta^{d-1})^d \subset [0, 1]^{d \times d}$ be a $d \times d$ row-stochastic matrix for an ergodic (i.e., irreducible and aperiodic) Markov chain. This implies there is a unique stationary distribution $\boldsymbol{\pi} \in \Delta^{d-1}$ with $\pi_i > 0$ for all $i \in [d]$ [3, Corollary 1.17]. We also assume that $\mathbf{P}$ is *reversible* (with respect to $\boldsymbol{\pi}$):

$$\pi_i P_{i,j} = \pi_j P_{j,i}, \quad i, j \in [d]. \quad (3)$$

The minimum stationary probability is denoted by $\pi_\star := \min_{i \in [d]} \pi_i$.

Define the matrices

$$\mathbf{M} := \text{Diag}(\boldsymbol{\pi})\mathbf{P} \quad \text{and} \quad \mathbf{L} := \text{Diag}(\boldsymbol{\pi})^{-1/2}\mathbf{M}\text{Diag}(\boldsymbol{\pi})^{-1/2}.$$

The (i, j) th entry of the matrix $M_{i,j}$ contains the *doublet probabilities* associated with $\mathbf{P}$: $M_{i,j} = \pi_i P_{i,j}$ is the probability of seeing state i followed by state j when the chain is started from its stationary distribution. The matrix $\mathbf{M}$ is symmetric on account of the reversibility of $\mathbf{P}$, and hence it follows that $\mathbf{L}$ is also symmetric. (We will strongly exploit the symmetry in our results.) Further, $\mathbf{L} = \text{Diag}(\boldsymbol{\pi})^{1/2}\mathbf{P}\text{Diag}(\boldsymbol{\pi})^{-1/2}$, hence $\mathbf{L}$ and $\mathbf{P}$ are similar and thus their eigenvalue systems are identical. Ergodicity and reversibility imply that the eigenvalues of $\mathbf{L}$ are contained in the interval $(-1, 1]$, and that 1 is an eigenvalue of $\mathbf{L}$ with multiplicity 1 [3, Lemmas 12.1 and 12.2]. Denote and order the eigenvalues of $\mathbf{L}$ as

$$1 = \lambda_1 > \lambda_2 \geq \dots \geq \lambda_d > -1.$$

Let $\lambda_\star := \max\{\lambda_2, |\lambda_d|\}$, and define the (absolute) spectral gap to be $\gamma_\star := 1 - \lambda_\star$, which is strictly positive on account of ergodicity.

Let $(X_t)_{t \in \mathbb{N}}$ be a Markov chain whose transition probabilities are governed by $\mathbf{P}$. For each $t \in \mathbb{N}$, let $\boldsymbol{\pi}^{(t)} \in \Delta^{d-1}$ denote the marginal distribution of X_t , so

$$\boldsymbol{\pi}^{(t+1)} = \boldsymbol{\pi}^{(t)}\mathbf{P}, \quad t \in \mathbb{N}.$$

Note that the initial distribution $\boldsymbol{\pi}^{(1)}$ is arbitrary, and need not be the stationary distribution $\boldsymbol{\pi}$.

The goal is to estimate $\pi_\star$ and $\gamma_\star$ from the length n sample path $(X_t)_{t \in [n]}$, and also to construct fully empirical confidence intervals that $\pi_\star$ and $\gamma_\star$ with high probability; in particular, the construction

of the intervals should not depend on any unobservable quantities, including π_* and γ_* themselves. As mentioned in the introduction, it is well-known that the *mixing time* of the Markov chain t_{mix} (defined in Eq. 1) is bounded in terms of π_* and γ_* , as shown in Eq. (2). Moreover, convergence rates for empirical processes on Markov chain sequences are also often given in terms of mixing coefficients that can ultimately be bounded in terms of π_* and γ_* (as we will show in the proof of our first result). Therefore, valid confidence intervals for π_* and γ_* can be used to make these rates fully observable.

3 Point estimation

In this section, we present lower and upper bounds on achievable rates for estimating the spectral gap as a function of the length of the sample path n .

3.1 Lower bounds

The purpose of this section is to show lower bounds on the number of observations necessary to achieve a fixed multiplicative (or even just additive) accuracy in estimating the spectral gap γ_* . By Eq. (2), the multiplicative accuracy lower bound for γ_* gives the same lower bound for estimating the mixing time. Our first result holds even for two state Markov chains and shows that a sequence length of $\Omega(1/\pi_*)$ is necessary to achieve even a constant *additive* accuracy in estimating γ_* .

Theorem 1. *Pick any $\bar{\pi} \in (0, 1/4)$. Consider any estimator $\hat{\gamma}_*$ that takes as input a random sample path of length $n \leq 1/(4\bar{\pi})$ from a Markov chain starting from any desired initial state distribution. There exists a two-state ergodic and reversible Markov chain distribution with spectral gap $\gamma_* \geq 1/2$ and minimum stationary probability $\pi_* \geq \bar{\pi}$ such that*

$$\Pr[|\hat{\gamma}_* - \gamma_*| \geq 1/8] \geq 3/8.$$

Next, considering d state chains, we show that a sequence of length $\Omega(d \log(d)/\gamma_*)$ is required to estimate γ_* up to a constant multiplicative accuracy. Essentially, the sequence may have to visit all d states at least $\log(d)/\gamma_*$ times each, on average. This holds *even* if π_* is within a factor of two of the *largest* possible value of $1/d$ that it can take, i.e., when π is nearly uniform.

Theorem 2. *There is an absolute constant $c > 0$ such that the following holds. Pick any positive integer $d \geq 3$ and any $\bar{\gamma} \in (0, 1/2)$. Consider any estimator $\hat{\gamma}_*$ that takes as input a random sample path of length $n < cd \log(d)/\bar{\gamma}$ from a d -state reversible Markov chain starting from any desired initial state distribution. There is an ergodic and reversible Markov chain distribution with spectral gap $\gamma_* \in [\bar{\gamma}, 2\bar{\gamma}]$ and minimum stationary probability $\pi_* \geq 1/(2d)$ such that*

$$\Pr[|\hat{\gamma}_* - \gamma_*| \geq \bar{\gamma}/2] \geq 1/4.$$

The proofs of Theorems 1 and 2 are given in Appendix A.²

3.2 A plug-in based point estimator and its accuracy

Let us now consider the problem of estimating γ_* . For this, we construct a natural plug-in estimator. Along the way, we also provide an estimator for the minimum stationary probability, allowing one to use the bounds from Eq. (2) to trap the mixing time.

Define the random matrix $\widehat{\mathbf{M}} \in [0, 1]^{d \times d}$ and random vector $\hat{\pi} \in \Delta^{d-1}$ by

$$\begin{aligned} \widehat{M}_{i,j} &:= \frac{|\{t \in [n-1] : (X_t, X_{t+1}) = (i, j)\}|}{n-1}, \quad i, j \in [d], \\ \hat{\pi}_i &:= \frac{|\{t \in [n] : X_t = i\}|}{n}, \quad i \in [d]. \end{aligned}$$

Furthermore, define

$$\text{Sym}(\widehat{\mathbf{L}}) := \frac{1}{2}(\widehat{\mathbf{L}} + \widehat{\mathbf{L}}^\top)$$

²A full version of this paper, with appendices, is available on arXiv [31].

to be the symmetrized version of the (possibly non-symmetric) matrix

$$\widehat{L} := \text{Diag}(\widehat{\pi})^{-1/2} \widehat{M} \text{Diag}(\widehat{\pi})^{-1/2}.$$

Let $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \dots \geq \hat{\lambda}_d$ be the eigenvalues of $\text{Sym}(\widehat{L})$. Our estimator of the minimum stationary probability π_* is $\hat{\pi}_* := \min_{i \in [d]} \hat{\pi}_i$, and our estimator of the spectral gap γ_* is $\hat{\gamma}_* := 1 - \max\{\hat{\lambda}_2, |\hat{\lambda}_d|\}$.

These estimators have the following accuracy guarantees:

Theorem 3. *There exists an absolute constant $C > 0$ such that the following holds. Assume the estimators $\hat{\pi}_*$ and $\hat{\gamma}_*$ described above are formed from a sample path of length n from an ergodic and reversible Markov chain. Let $\gamma_* > 0$ denote the spectral gap and $\pi_* > 0$ the minimum stationary probability. For any $\delta \in (0, 1)$, with probability at least $1 - \delta$,*

$$|\hat{\pi}_* - \pi_*| \leq C \left(\sqrt{\frac{\pi_* \log \frac{d}{\pi_* \delta}}{\gamma_* n}} + \frac{\log \frac{d}{\pi_* \delta}}{\gamma_* n} \right) \quad (4)$$

and

$$|\hat{\gamma}_* - \gamma_*| \leq C \left(\sqrt{\frac{\log \frac{d}{\delta} \cdot \log \frac{n}{\pi_* \delta}}{\pi_* \gamma_* n}} + \frac{\log \frac{1}{\gamma_*}}{\gamma_* n} \right). \quad (5)$$

Theorem 3 implies that the sequence lengths required to estimate π_* and γ_* to within constant multiplicative factors are, respectively, $\tilde{O}\left(\frac{1}{\pi_* \gamma_*}\right)$ and $\tilde{O}\left(\frac{1}{\pi_* \gamma_*^3}\right)$. By Eq. (2), the second of these is also a bound on the required sequence length to estimate t_{mix} .

The proof of Theorem 3 is based on analyzing the convergence of the sample averages $\widehat{M}$ and $\widehat{\pi}$ to their expectation, and then using perturbation bounds for eigenvalues to derive a bound on the error of $\hat{\gamma}_*$. However, since these averages are formed using a *single sample path* from a (possibly) non-stationary Markov chain, we cannot use standard large deviation bounds; moreover applying Chernoff-type bounds for Markov chains to each entry of $\widehat{M}$ would result in a significantly worse sequence length requirement, roughly a factor of d larger. Instead, we adapt probability tail bounds for sums of independent random matrices [32] to our non-iid setting by directly applying a blocking technique of [33] as described in the article of [20]. Due to ergodicity, the convergence rate can be bounded without any dependence on the initial state distribution $\pi^{(1)}$. The proof of Theorem 3 is given in Appendix B.

Note that because the eigenvalues of L are the same as that of the transition probability matrix P , we could have instead opted to estimate P , say, using simple frequency estimates obtained from the sample path, and then computing the second largest eigenvalue of this empirical estimate $\widehat{P}$. In fact, this approach is a way to extend to non-reversible chains, as we would no longer rely on the symmetry of M or L . The difficulty with this approach is that P lacks the structure required by certain strong eigenvalue perturbation results. One could instead invoke the Ostrowski-Elsner theorem [cf. Theorem 1.4 on Page 170 of 34], which bounds the *matching distance* between the eigenvalues of a matrix A and its perturbation $A + E$ by $O(\|E\|^{1/d})$. Since $\|\widehat{P} - P\|$ is expected to be of size $O(n^{-1/2})$, this approach will give a confidence interval for γ_* whose width shrinks at a rate of $O(n^{-1/(2d)})$ —an exponential slow-down compared to the rate from Theorem 3. As demonstrated through an example from [34], the dependence on the d -th root of the norm of the perturbation cannot be avoided in general. Our approach based on estimating a symmetric matrix affords us the use of perturbation results that exploit more structure.

Returning to the question of obtaining a fully empirical confidence interval for γ_* and π_* , we notice that, unfortunately, Theorem 3 falls short of being directly suitable for this, at least without further assumptions. This is because the deviation terms themselves depend inversely both on γ_* and π_* , and hence can never rule out 0 (or an arbitrarily small positive value) as a possibility for γ_* or π_* .³ In effect, the fact that the Markov chain could be slow mixing and the long-term frequency of some

³Using Theorem 3, it is possible to trap γ_* in the union of *two* empirical confidence intervals—one around $\hat{\gamma}_*$ and the other around zero, both of which shrink in width as the sequence length increases.

Algorithm 1 Empirical confidence intervals

Input: Sample path $(X_1, X_2, \dots, X_n)$, confidence parameter $\delta \in (0, 1)$.

1: Compute state visit counts and smoothed transition probability estimates:

$$N_i := |\{t \in [n-1] : X_t = i\}|, \quad i \in [d];$$

$$N_{i,j} := |\{t \in [n-1] : (X_t, X_{t+1}) = (i, j)\}|, \quad \hat{P}_{i,j} := \frac{N_{i,j} + 1/d}{N_i + 1}, \quad (i, j) \in [d]^2.$$

2: Let $\hat{\mathbf{A}}^\#$ be the group inverse of $\hat{\mathbf{A}} := \mathbf{I} - \hat{\mathbf{P}}$.

3: Let $\hat{\pi} \in \Delta^{d-1}$ be the unique stationary distribution for $\hat{\mathbf{P}}$.

4: Compute eigenvalues $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \dots \geq \hat{\lambda}_d$ of $\text{Sym}(\hat{\mathbf{L}})$, where $\hat{\mathbf{L}} := \text{Diag}(\hat{\pi})^{1/2} \hat{\mathbf{P}} \text{Diag}(\hat{\pi})^{-1/2}$.

5: Spectral gap estimate:

$$\hat{\gamma}_* := 1 - \max\{\hat{\lambda}_2, |\hat{\lambda}_d|\}.$$

6: Empirical bounds for $|\hat{P}_{i,j} - P_{i,j}|$ for $(i, j) \in [d]^2$: $c := 1.1$, $\tau_{n,\delta} := \inf\{t \geq 0 : 2d^2(1 + \lceil \log_c \frac{2n}{t} \rceil_+) e^{-t} \leq \delta\}$,

$$\text{and } \hat{B}_{i,j} := \left(\sqrt{\frac{c\tau_{n,\delta}}{2N_i}} + \sqrt{\frac{c\tau_{n,\delta}}{2N_i}} + \sqrt{\frac{2c\hat{P}_{i,j}(1 - \hat{P}_{i,j})\tau_{n,\delta}}{N_i}} + \frac{(4/3)\tau_{n,\delta} + |\hat{P}_{i,j} - 1/d|}{N_i} \right)^2.$$

7: Relative sensitivity of π :

$$\hat{\kappa} := \frac{1}{2} \max \left\{ \hat{A}_{j,j}^\# - \min \left\{ \hat{A}_{i,j}^\# : i \in [d] \right\} : j \in [d] \right\}.$$

8: Empirical bounds for $\max_{i \in [d]} |\hat{\pi}_i - \pi_i|$ and $\max \bigcup_{i \in [d]} \{|\sqrt{\pi_i/\hat{\pi}_i} - 1|, |\sqrt{\hat{\pi}_i/\pi_i} - 1|\}$:

$$\hat{b} := \hat{\kappa} \max \left\{ \hat{B}_{i,j} : (i, j) \in [d]^2 \right\}, \quad \hat{\rho} := \frac{1}{2} \max \bigcup_{i \in [d]} \left\{ \frac{\hat{b}}{\hat{\pi}_i}, \frac{\hat{b}}{[\hat{\pi}_i - \hat{b}]_+} \right\}.$$

9: Empirical bounds for $|\hat{\gamma}_* - \gamma_*|$:

$$\hat{w} := 2\hat{\rho} + \hat{\rho}^2 + (1 + 2\hat{\rho} + \hat{\rho}^2) \left(\sum_{(i,j) \in [d]^2} \frac{\hat{\pi}_i}{\hat{\pi}_j} \hat{B}_{i,j}^2 \right)^{1/2}.$$

states could be small makes it difficult to be confident in the estimates provided by $\hat{\gamma}_*$ and $\hat{\pi}_*$. This suggests that in order to obtain fully empirical confidence intervals, we need an estimator that is not subject to such effects—we pursue this in Section 4. Theorem 3 thus primarily serves as a point of comparison for what is achievable in terms of estimation accuracy when one does not need to provide empirical confidence bounds.

4 Fully empirical confidence intervals

In this section, we address the shortcoming of Theorem 3 and give fully empirical confidence intervals for the stationary probabilities and the spectral gap γ_* . The main idea is to use the Markov property to eliminate the dependence of the confidence intervals on the unknown quantities (including π_* and γ_*). Specifically, we estimate the transition probabilities from the sample path using simple frequency estimates: as a consequence of the Markov property, for each state, the frequency estimates converge at a rate that depends only on the number of visits to the state, and in particular the rate (given the visit count of the state) is independent of the mixing time of the chain.

As discussed in Section 3, it is possible to form a confidence interval for γ_* based on the eigenvalues of an estimated transition probability matrix by appealing to the Ostrowski-Elsner theorem. However, as explained earlier, this would lead to a slow $O(n^{-1/(2d)})$ rate. We avoid this slow rate by using an estimate of the symmetric matrix $\mathbf{L}$, so that we can use a stronger perturbation result (namely Weyl's inequality, as in the proof of Theorem 3) available for symmetric matrices.

To form an estimate of $\mathbf{L}$ based on an estimate of the transition probabilities, one possibility is to estimate π using a frequency-based estimate for π as was done in Section 3, and appeal to the relation $\mathbf{L} = \text{Diag}(\pi)^{1/2} \mathbf{P} \text{Diag}(\pi)^{-1/2}$ to form a plug-in estimate. However, as noted in Section 3.2, confidence intervals for the entries of π formed this way may depend on the mixing time. Indeed, such an estimate of π does not exploit the Markov property.

We adopt a different strategy for estimating π , which leads to our construction of empirical confidence intervals, detailed in Algorithm 1. We form the matrix $\hat{\mathbf{P}}$ using smoothed frequency estimates of $\mathbf{P}$ (Step 1), then compute the so-called group inverse $\hat{\mathbf{A}}^\#$ of $\hat{\mathbf{A}} = \mathbf{I} - \hat{\mathbf{P}}$ (Step 2), followed by finding the unique stationary distribution $\hat{\pi}$ of $\hat{\mathbf{P}}$ (Step 3), this way decoupling the bound on the accuracy of $\hat{\pi}$ from the mixing time. The group inverse $\hat{\mathbf{A}}^\#$ of $\hat{\mathbf{A}}$ is uniquely defined; and if $\hat{\mathbf{P}}$ defines an ergodic chain (which is the case here due to the use of the smoothed estimates), $\hat{\mathbf{A}}^\#$ can be computed at the cost of inverting an $(d-1) \times (d-1)$ matrix [35, Theorem 5.2].⁴ Further, once given $\hat{\mathbf{A}}^\#$, the unique stationary distribution $\hat{\pi}$ of $\hat{\mathbf{P}}$ can be read out from the last row of $\hat{\mathbf{A}}^\#$ [35, Theorem 5.3]. The group inverse is also used to compute the sensitivity of π . Based on $\hat{\pi}$ and $\hat{\mathbf{P}}$, we construct the plug-in estimate $\hat{\mathbf{L}}$ of $\mathbf{L}$, and use the eigenvalues of its symmetrization to form the estimate $\hat{\gamma}_*$ of the spectral gap (Steps 4 and 5). In the remaining steps, we use perturbation analyses to relate $\hat{\pi}$ and π , viewing $\mathbf{P}$ as the perturbation of $\hat{\mathbf{P}}$; and also to relate $\hat{\gamma}_*$ and γ_* , viewing $\mathbf{L}$ as a perturbation of $\text{Sym}(\hat{\mathbf{L}})$. Both analyses give error bounds entirely in terms of observable quantities (e.g., $\hat{\kappa}$), tracing back to empirical error bounds for the smoothed frequency estimates of $\mathbf{P}$.

The most computationally expensive step in Algorithm 1 is the computation of the group inverse $\hat{\mathbf{A}}^\#$, which, as noted reduces to matrix inversion. Thus, with a standard implementation of matrix inversion, the algorithm's time complexity is $O(n + d^3)$, while its space complexity is $O(d^2)$.

To state our main theorem concerning Algorithm 1, we first define κ to be analogous to $\hat{\kappa}$ from Step 7, with $\hat{\mathbf{A}}^\#$ replaced by the group inverse $\mathbf{A}^\#$ of $\mathbf{A} := \mathbf{I} - \mathbf{P}$. The result is as follows.

Theorem 4. *Suppose Algorithm 1 is given as input a sample path of length n from an ergodic and reversible Markov chain and confidence parameter $\delta \in (0, 1)$. Let $\gamma_* > 0$ denote the spectral gap, π the unique stationary distribution, and $\pi_* > 0$ the minimum stationary probability. Then, on an event of probability at least $1 - \delta$,*

$$\pi_i \in [\hat{\pi}_i - \hat{b}, \hat{\pi}_i + \hat{b}] \quad \text{for all } i \in [d], \quad \text{and} \quad \gamma_* \in [\hat{\gamma}_* - \hat{w}, \hat{\gamma}_* + \hat{w}].$$

Moreover, $\hat{b}$ and $\hat{w}$ almost surely satisfy (as $n \rightarrow \infty$)

$$\hat{b} = O\left(\max_{(i,j) \in [d]^2} \kappa \sqrt{\frac{P_{i,j} \log \log n}{\pi_i n}}\right), \quad \hat{w} = O\left(\frac{\kappa}{\pi_*} \sqrt{\frac{\log \log n}{\pi_* n}} + \sqrt{\frac{d \log \log n}{\pi_* n}}\right).^5$$

The proof of Theorem 4 is given in Appendix C. As mentioned above, the obstacle encountered in Theorem 3 is avoided by exploiting the Markov property. We establish fully observable upper and lower bounds on the entries of $\mathbf{P}$ that converge at a $\sqrt{n/\log \log n}$ rate using standard martingale tail inequalities; this justifies the validity of the bounds from Step 6. Properties of the group inverse [35, 36] and eigenvalue perturbation theory [34] are used to validate the empirical bounds on π_i and γ_* developed in the remaining steps of the algorithm.

The first part of Theorem 4 provides valid empirical confidence intervals for each π_i and for γ_* , which are simultaneously valid at confidence level δ . The second part of Theorem 4 shows that the

⁴ The group inverse of a square matrix $\mathbf{A}$, a special case of the *Drazin inverse*, is the unique matrix $\mathbf{A}^\#$ satisfying $\mathbf{A}\mathbf{A}^\#\mathbf{A} = \mathbf{A}$, $\mathbf{A}^\#\mathbf{A}\mathbf{A}^\# = \mathbf{A}^\#$ and $\mathbf{A}^\#\mathbf{A} = \mathbf{A}\mathbf{A}^\#$.

⁵ In Theorems 4 and 5, our use of big- O notation is as follows. For a random sequence $(Y_n)_n$ and a (non-random) positive sequence $(\varepsilon_{\theta,n})_n$ parameterized by θ , we say “ $Y_n = O(\varepsilon_{\theta,n})$ holds almost surely as $n \rightarrow \infty$ ” if there is some universal constant $C > 0$ such that for all θ , $\limsup_{n \rightarrow \infty} Y_n/\varepsilon_{\theta,n} \leq C$ holds almost surely.

width of the intervals decrease as the sequence length increases. We show in Appendix C.5 that $\kappa \leq d/\gamma_*$, and hence $\hat{b} = O\left(\max_{(i,j) \in [d]^2} \frac{d}{\gamma_*} \sqrt{\frac{P_{i,j} \log \log n}{\pi_i n}}\right)$, $\hat{w} = O\left(\frac{d}{\pi_* \gamma_*} \sqrt{\frac{\log \log n}{\pi_* n}}\right)$.

It is easy to combine Theorems 3 and 4 to yield intervals whose widths shrink at least as fast as both the non-empirical intervals from Theorem 3 and the empirical intervals from Theorem 4. Specifically, determine lower bounds on π_* and γ_* using Algorithm 1, $\pi_* \geq \min_{i \in [d]} [\hat{\pi}_i - \hat{b}]_+$, $\gamma_* \geq [\hat{\gamma}_* - \hat{w}]_+$; then plug-in these lower bounds for π_* and γ_* in the deviation bounds in Eq. (5) from Theorem 3. This yields a new interval centered around the estimate of γ_* from Theorem 3, and it no longer depends on unknown quantities. The interval is a valid $1 - 2\delta$ probability confidence interval for γ_* , and for sufficiently large n , the width shrinks at the rate given in Eq. (5). We can similarly construct an empirical confidence interval for π_* using Eq. (4), which is valid on the same $1 - 2\delta$ probability event.⁶ Finally, we can take the intersection of these new intervals with the corresponding intervals from Algorithm 1. This is summarized in the following theorem, which we prove in Appendix D.

Theorem 5. *The following holds under the same conditions as Theorem 4. For any $\delta \in (0, 1)$, the confidence intervals $\hat{U}$ and $\hat{V}$ described above for π_* and γ_* , respectively, satisfy $\pi_* \in \hat{U}$ and $\gamma_* \in \hat{V}$ with probability at least $1 - 2\delta$. Furthermore, the widths of these intervals almost surely satisfy (as $n \rightarrow \infty$) $|\hat{U}| = O\left(\sqrt{\frac{\pi_* \log \frac{d}{\pi_* \delta}}{\gamma_* n}}\right)$, $|\hat{V}| = O\left(\min\left\{\sqrt{\frac{\log \frac{d}{\delta} \cdot \log(n)}{\pi_* \gamma_* n}}, \hat{w}\right\}\right)$, where $\hat{w}$ is the width from Algorithm 1.*

5 Discussion

The construction used in Theorem 5 applies more generally: Given a confidence interval of the form $I_n = I_n(\gamma_*, \pi_*, \delta)$ for some confidence level δ and a fully empirical confidence set $E_n(\delta)$ for (γ_*, π_*) for the same level, $I'_n = E_n(\delta) \cap \bigcup_{(\gamma, \pi) \in E_n(\delta)} I_n(\gamma, \pi, \delta)$ is a valid fully empirical 2δ -level confidence interval whose asymptotic width matches that of I_n up to lower order terms under reasonable assumptions on E_n and I_n . In particular, this suggests that future work should focus on closing the gap between the lower and upper bounds on the accuracy of point-estimation. Another interesting direction is to reduce the computation cost: The current cubic cost in the number of states can be too high even when the number of states is only moderately large.

Perhaps more important, however, is to extend our results to large state space Markov chains: In most practical applications the state space is continuous or is exponentially large in some natural parameters. As follows from our lower bounds, without further assumptions, the problem of fully data dependent estimation of the mixing time is intractable for information theoretical reasons. Interesting directions for future work thus must consider Markov chains with specific structure. Parametric classes of Markov chains, including but not limited to Markov chains with factored transition kernels with a few factors, are a promising candidate for such future investigations. The results presented here are a first step in the ambitious research agenda outlined above, and we hope that they will serve as a point of departure for further insights in the area of fully empirical estimation of Markov chain parameters based on a single sample path.

References

- [1] J. S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer Series in Statistics. Springer-Verlag, 2001.
- [2] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning)*. A Bradford Book, 1998.
- [3] D. Levin, Y. Peres, and E. Wilmer. *Markov Chains and Mixing Times*. AMS, 2008.
- [4] S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Springer, 1993.
- [5] C. Kipnis and S. R. S. Varadhan. Central limit theorem for additive functionals of reversible markov processes and applications to simple exclusions. *Comm. Math. Phys.*, 104(1):1–19, 1986.

⁶For the π_* interval, we only plug-in lower bounds on π_* and γ_* only where these quantities appear as $1/\pi_*$ and $1/\gamma_*$ in Eq. (4). It is then possible to “solve” for observable bounds on π_* . See Appendix D for details.

- [6] I. Kontoyiannis, L. A. Lastras-Montaño, and S. P. Meyn. Exponential bounds and stopping rules for MCMC and general Markov chains. In *VALUETOOLS*, page 45, 2006.
- [7] M.-F. Balcan, A. Beygelzimer, and J. Langford. Agnostic active learning. In *ICML*, pages 65–72, 2006.
- [8] V. Mnih, Cs. Szepesvári, and J.-Y. Audibert. Empirical Bernstein stopping. In *ICML*, pages 672–679, 2008.
- [9] A. Maurer and M. Pontil. Empirical Bernstein bounds and sample-variance penalization. In *COLT*, 2009.
- [10] L. Li, M. L. Littman, T. J. Walsh, and A. L. Strehl. Knows what it knows: a framework for self-aware learning. *Machine Learning*, 82(3):399–443, 2011.
- [11] J. M. Flegal and G. L. Jones. Implementing MCMC: estimating with confidence. In *Handbook of Markov chain Monte Carlo*, pages 175–197. Chapman & Hall/CRC, 2011.
- [12] B. M. Gyori and D. Paulin. Non-asymptotic confidence intervals for MCMC in practice. arXiv:1212.2016, 2014.
- [13] A. Swaminathan and T. Joachims. Counterfactual risk minimization: Learning from logged bandit feedback. In *ICML*, 2015.
- [14] D. Gillman. A Chernoff bound for random walks on expander graphs. *SIAM Journal on Computing*, 27(4):1203–1220, 1998.
- [15] C. A. León and F. Perron. Optimal Hoeffding bounds for discrete reversible Markov chains. *Annals of Applied Probability*, pages 958–970, 2004.
- [16] D. Paulin. Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electronic Journal of Probability*, 20:1–32, 2015.
- [17] S. T. Garren and R. L. Smith. Estimating the second largest eigenvalue of a Markov transition matrix. *Bernoulli*, 6:215–242, 2000.
- [18] G. L. Jones and J. P. Hobert. Honest exploration of intractable probability distributions via markov chain monte carlo. *Statist. Sci.*, 16(4):312–334, 11 2001.
- [19] Y. Atchadé. Markov Chain Monte Carlo confidence intervals. *Bernoulli*, 2015. (to appear).
- [20] B. Yu. Rates of convergence for empirical processes of stationary mixing sequences. *The Annals of Probability*, 22(1):94–116, January 1994.
- [21] R. L. Karandikar and M. Vidyasagar. Rates of uniform convergence of empirical means with mixing processes. *Statistics and Probability Letters*, 58(3):297–307, 2002.
- [22] D. Gamarnik. Extension of the PAC framework to finite and countable Markov chains. *IEEE Transactions on Information Theory*, 49(1):338–345, 2003.
- [23] M. Mohri and A. Rostamizadeh. Stability bounds for non-iid processes. In *NIPS*, 2008.
- [24] M. Mohri and A. Rostamizadeh. Rademacher complexity bounds for non-i.i.d. processes. In *NIPS*, 2009.
- [25] I. Steinwart and A. Christmann. Fast learning from non-i.i.d. observations. In *NIPS*, 2009.
- [26] I. Steinwart, D. Hush, and C. Scovel. Learning from dependent observations. *Journal of Multivariate Analysis*, 100(1):175–194, 2009.
- [27] D. McDonald, C. Shalizi, and M. Schervish. Estimating beta-mixing coefficients. In *AISTATS*, pages 516–524, 2011.
- [28] T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White. Testing that distributions are close. In *FOCS*, pages 259–269. IEEE, 2000.
- [29] T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White. Testing closeness of discrete distributions. *Journal of the ACM (JACM)*, 60(1):4:2–4:25, 2013.
- [30] N. Bhatnagar, A. Bogdanov, and E. Mossel. The computational complexity of estimating MCMC convergence time. In *RANDOM*, pages 424–435. Springer, 2011.
- [31] D. Hsu, A. Kontorovich, and C. Szepesvári. Mixing time estimation in reversible Markov chains from a single sample path. *CoRR*, abs/1506.02903, 2015.
- [32] J. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 2015.
- [33] S. Bernstein. Sur l’extension du theoreme limite du calcul des probabilites aux sommes de quantites dependantes. *Mathematische Annalen*, 97:1–59, 1927.
- [34] G. W. Stewart and J. Sun. *Matrix perturbation theory*. Academic Press, Boston, 1990.
- [35] C. D. Meyer Jr. The role of the group generalized inverse in the theory of finite Markov chains. *SIAM Review*, 17(3):443–464, 1975.
- [36] G. Cho and C. Meyer. Comparison of perturbation bounds for the stationary distribution of a Markov chain. *Linear Algebra and its Applications*, 335:137–150, 2001.

A Proofs of the lower bounds

Theorem 6 (Theorem 1 restated). *Pick any $\bar{\pi} \in (0, 1/4)$. Consider any estimator $\hat{\gamma}_*$ that takes as input a random sample path of length $n \leq 1/(4\bar{\pi})$ from a Markov chain starting from any desired initial state distribution. There exists a two-state ergodic and reversible Markov chain distribution with spectral gap $\gamma_* \geq 1/2$ and minimum stationary probability $\pi_* \geq \bar{\pi}$ such that*

$$\Pr[|\hat{\gamma}_* - \gamma_*| \geq 1/8] \geq 3/8.$$

Proof. Fix $\bar{\pi} \in (0, 1/4)$. Consider two Markov chains given by the following stochastic matrices:

$$\mathbf{P}^{(1)} := \begin{bmatrix} 1 - \bar{\pi} & \bar{\pi} \\ 1 - \bar{\pi} & \bar{\pi} \end{bmatrix}, \quad \mathbf{P}^{(2)} := \begin{bmatrix} 1 - \bar{\pi} & \bar{\pi} \\ 1/2 & 1/2 \end{bmatrix}.$$

Each Markov chain is ergodic and reversible; their stationary distributions are, respectively, $\pi^{(1)} = (1 - \bar{\pi}, \bar{\pi})$ and $\pi^{(2)} = (1/(1 + 2\bar{\pi}), 2\bar{\pi}/(1 + 2\bar{\pi}))$. We have $\pi_* \geq \bar{\pi}$ in both cases. For the first Markov chain, $\lambda_* = 0$, and hence the spectral gap is 1; for the second Markov chain, $\lambda_* = 1/2 - \bar{\pi}$, so the spectral gap is $1/2 + \bar{\pi}$.

In order to guarantee $|\hat{\gamma}_* - \gamma_*| < 1/8 < |1 - (1/2 + \bar{\pi})|/2$, it must be possible to distinguish the two Markov chains. Assume that the initial state distribution has mass at least $1/2$ on state 1. (If this is not the case, we swap the roles of states 1 and 2 in the constructions above.) With probability at least half, the initial state is 1; and both chains have the same transition probabilities from state 1. The chains are indistinguishable unless the sample path eventually reaches state 2. But with probability at least $3/4$, a sample path of length $n < 1/(4\bar{\pi})$ starting from state 1 always remains in the same state (this follows from properties of the geometric distribution and the assumption $\bar{\pi} < 1/4$). $\square$

Theorem 7 (Theorem 2 restated). *There is an absolute constant $c > 0$ such that the following holds. Pick any positive integer $d \geq 3$ and any $\bar{\gamma} \in (0, 1/2)$. Consider any estimator $\hat{\gamma}_*$ that takes as input a random sample path of length $n < cd \log(d)/\bar{\gamma}$ from a d -state reversible Markov chain starting from any desired initial state distribution. There is an ergodic and reversible Markov chain distribution with spectral gap $\gamma_* \in [\bar{\gamma}, 2\bar{\gamma}]$ and minimum stationary probability $\pi_* \geq 1/(2d)$ such that*

$$\Pr[|\hat{\gamma}_* - \gamma_*| \geq \bar{\gamma}/2] \geq 1/4.$$

Proof. We consider d -state Markov chains of the following form:

$$P_{i,j} = \begin{cases} 1 - \varepsilon_i & \text{if } i = j; \\ \frac{\varepsilon_i}{d-1} & \text{if } i \neq j \end{cases}$$

for some $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_d \in (0, 1)$. Such a chain is ergodic and reversible, and its unique stationary distribution π satisfies

$$\pi_i = \frac{1/\varepsilon_i}{\sum_{j=1}^d 1/\varepsilon_j}.$$

We fix $\varepsilon := \frac{d-1}{d/2} \bar{\gamma}$ and set $\varepsilon' := \frac{d/2-1}{d-1} \varepsilon < \varepsilon$. Consider the following $d + 1$ different Markov chains of the type described above:

- $\mathbf{P}^{(0)}$: $\varepsilon_1 = \dots = \varepsilon_d = \varepsilon$. For this Markov chain, $\lambda_2 = \lambda_d = \lambda_* = 1 - \frac{d}{d-1} \varepsilon$.
- $\mathbf{P}^{(i)}$ for $i \in [d]$: $\varepsilon_j = \varepsilon$ for $j \neq i$, and $\varepsilon_i = \varepsilon'$. For these Markov chains, $\lambda_2 = 1 - \varepsilon' - \frac{1}{d-1} \varepsilon = 1 - \frac{d/2}{d-1} \varepsilon$, and $\lambda_d = 1 - \frac{d}{d-1} \varepsilon$. So $\lambda_* = 1 - \frac{d/2}{d-1} \varepsilon$.

The spectral gap in each chain satisfies $\gamma_* \in [\bar{\gamma}, 2\bar{\gamma}]$; in $\mathbf{P}^{(i)}$ for $i \in [d]$, it is half of what it is in $\mathbf{P}^{(0)}$. Also $\pi_i \geq 1/(2d)$ for each $i \in [d]$.

In order to guarantee $|\hat{\gamma}_* - \gamma_*| < \bar{\gamma}/2$, it must be possible to distinguish $\mathbf{P}^{(0)}$ from each $\mathbf{P}^{(i)}$, $i \in [d]$. But $\mathbf{P}^{(0)}$ is identical to $\mathbf{P}^{(i)}$ except for the transition probabilities from state i . Therefore, regardless

of the initial state, the sample path must visit all states in order to distinguish $\mathbf{P}^{(0)}$ from each $\mathbf{P}^{(i)}$, $i \in [d]$. For any of the $d + 1$ Markov chains above, the earliest time in which a sample path visits all d states stochastically dominates a generalized coupon collection time $T = 1 + \sum_{i=1}^{d-1} T_i$, where T_i is the number of steps required to see the $(i + 1)$ -th distinct state in the sample path beyond the first i . The random variables $T_1, T_2, \dots, T_{d-1}$ are independent, and are geometrically distributed, $T_i \sim \text{Geom}(\varepsilon - (i - 1)\varepsilon/(d - 1))$. We have that

$$\mathbb{E}[T_i] = \frac{d - 1}{\varepsilon(d - i)}, \quad \text{var}(T_i) = \frac{1 - \varepsilon \frac{d-i}{d-1}}{\left(\varepsilon \frac{d-i}{d-1}\right)^2}.$$

Therefore

$$\mathbb{E}[T] = 1 + \frac{d - 1}{\varepsilon} H_{d-1}, \quad \text{var}(T) \leq \left(\frac{d - 1}{\varepsilon}\right)^2 \frac{\pi^2}{6}$$

where $H_{d-1} = 1 + 1/2 + 1/3 + \dots + 1/(d - 1)$. By the Paley-Zygmund inequality,

$$\Pr\left(T > \frac{1}{3}\mathbb{E}[T]\right) \geq \frac{1}{1 + \frac{\text{var}(T)}{(1-1/3)^2\mathbb{E}[T]^2}} \geq \frac{1}{1 + \frac{\left(\frac{d-1}{\varepsilon}\right)^2 \frac{\pi^2}{6}}{(4/9)\left(\frac{d-1}{\varepsilon} H_2\right)^2}} \geq \frac{1}{4}.$$

Since $n < cd \log(d)/\bar{\gamma} \leq (1/3)(1 + (d - 1)H_{d-1}/(2\bar{\gamma})) = \mathbb{E}[T]/3$ (for an appropriate absolute constant c), with probability at least $1/4$, the sample path does not visit all d states. $\square$

We claim in Section 1 that a sample path of length $\Omega((d \log d)/\gamma_* + 1/\pi_*)$ is required to guarantee constant multiplicative accuracy in estimating γ_* . This follows by combining Theorems 1 and 2 in a standard, straightforward way. Specifically, if the length of the sample path n is smaller than $(n_1 + n_2)/2$ —where n_1 is the lower bound from Theorem 1, and n_2 is the lower bound from Theorem 2—then n is smaller than $\max\{n_1, n_2\}$. So at least one of Theorem 1 and Theorem 2 implies the existence of an ergodic and reversible Markov chain distribution with (say) $1/4 \leq \gamma_* \leq 1/2$ and $\pi_* \geq 1/(2d)$ such that

$$\Pr[|\hat{\gamma}_* - \gamma_*| \geq \gamma_*/4] \geq 1/4.$$

B Proof of Theorem 3

In this section, we prove Theorem 3.

B.1 Accuracy of $\hat{\pi}_*$

We start by proving the deviation bound on $\pi_* - \hat{\pi}_*$, from which we may easily deduce Eq. (4) in Theorem 3.

Lemma 1. *Pick any $\delta \in (0, 1)$, and let*

$$\varepsilon_n := \frac{\ln\left(\frac{d}{\delta} \sqrt{\frac{2}{\pi_*}}\right)}{\gamma_* n}. \quad (6)$$

With probability at least $1 - \delta$, the following inequalities hold simultaneously:

$$|\hat{\pi}_i - \pi_i| \leq \sqrt{8\pi_i(1 - \pi_i)\varepsilon_n} + 20\varepsilon_n \quad \text{for all } i \in [d]; \quad (7)$$

$$|\hat{\pi}_* - \pi_*| \leq 4\sqrt{\pi_*\varepsilon_n} + 47\varepsilon_n. \quad (8)$$

Proof. We use the following Bernstein-type inequality for Markov chains from [1, Theorem 3.8]: letting $\mathbb{P}^\pi$ denote the probability with respect to the stationary chain (where the marginal distribution of each X_t is π), we have for every $\epsilon > 0$,

$$\mathbb{P}^\pi(|\hat{\pi}_i - \pi_i| > \epsilon) \leq 2 \exp\left(-\frac{n\gamma_*\epsilon^2}{4\pi_i(1 - \pi_i) + 10\epsilon}\right), \quad i \in [d].$$

To handle possibly non-stationary chains, as is our case, we combine the above inequality with [1, Proposition 3.14], to obtain for any $\epsilon > 0$,

$$\mathbb{P}(|\hat{\pi}_i - \pi_i| > \epsilon) \leq \sqrt{\frac{1}{\pi_\star} \mathbb{P}^\pi(|\hat{\pi}_i - \pi_i| > \epsilon)} \leq \sqrt{\frac{2}{\pi_\star} \exp\left(-\frac{n\gamma_\star \epsilon^2}{8\pi_i(1 - \pi_i) + 20\epsilon}\right)}.$$

Using this tail inequality with $\epsilon := \sqrt{8\pi_i(1 - \pi_i)\varepsilon_n} + 20\varepsilon_n$ and a union bound over all $i \in [d]$ implies that the inequalities in Eq. (7) hold with probability at least $1 - \delta$.

Now assume this $1 - \delta$ probability event holds; it remains to prove that Eq. (8) also holds in this event. Without loss of generality, we assume that $\pi_\star = \pi_1 \leq \pi_2 \leq \dots \leq \pi_d$. Let $j \in [d]$ be such that $\hat{\pi}_\star = \hat{\pi}_j$. By Eq. (7), we have $|\pi_i - \hat{\pi}_i| \leq \sqrt{8\pi_i\varepsilon_n} + 20\varepsilon_n$ for each $i \in \{1, j\}$. Since $\hat{\pi}_\star \leq \hat{\pi}_1$,

$$\hat{\pi}_\star - \pi_\star \leq \hat{\pi}_1 - \pi_1 \leq \sqrt{8\pi_\star\varepsilon_n} + 20\varepsilon_n \leq \pi_\star + 22\varepsilon_n$$

where the last inequality follows by the AM/GM inequality. Furthermore, using the fact that $a \leq b\sqrt{a} + c \Rightarrow a \leq b^2 + b\sqrt{c} + c$ for nonnegative numbers $a, b, c \geq 0$ [see, e.g., 2] with the inequality $\pi_j \leq \sqrt{8\varepsilon_n}\sqrt{\pi_j} + (\hat{\pi}_j + 20\varepsilon_n)$ gives

$$\pi_j \leq \hat{\pi}_j + \sqrt{8(\hat{\pi}_j + 20\varepsilon_n)\varepsilon_n} + 28\varepsilon_n.$$

Therefore

$$\pi_\star - \hat{\pi}_\star \leq \pi_j - \hat{\pi}_j \leq \sqrt{8(\hat{\pi}_\star + 20\varepsilon_n)\varepsilon_n} + 28\varepsilon_n \leq \sqrt{8(2\pi_\star + 42\varepsilon_n)\varepsilon_n} + 28\varepsilon_n \leq 4\sqrt{\pi_\star\varepsilon_n} + 47\varepsilon_n$$

where the second-to-last inequality follows from the above bound on $\hat{\pi}_\star - \pi_\star$, and the last inequality uses $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for nonnegative $a, b \geq 0$. $\square$

B.2 Accuracy of $\hat{\gamma}_\star$

Let us now turn to proving Eq. (5), i.e., the bound on the error of the spectral gap estimate $\hat{\gamma}_\star$. The accuracy of $\hat{\gamma}_\star$ is based on the accuracy of $\text{Sym}(\hat{\mathbf{L}})$ in approximating $\mathbf{L}$ via Weyl's inequality:

$$|\hat{\lambda}_i - \lambda_i| \leq \|\text{Sym}(\hat{\mathbf{L}}) - \mathbf{L}\| \quad \text{for all } i \in [d].$$

Moreover, the triangle inequality implies that symmetrizing $\hat{\mathbf{L}}$ can only help:

$$\|\text{Sym}(\hat{\mathbf{L}}) - \mathbf{L}\| \leq \|\hat{\mathbf{L}} - \mathbf{L}\|.$$

Therefore, we can deduce Eq. (5) in Theorem 3 from the following lemma.

Lemma 2. *There exists an absolute constant $C > 0$ such that the following holds. For any $\delta \in (0, 1)$, with probability at least $1 - \delta$, the bounds from Lemma 1 hold, and*

$$\|\hat{\mathbf{L}} - \mathbf{L}\| \leq C \sqrt{\frac{\log\left(\frac{d}{\delta}\right) \log\left(\frac{n}{\pi_\star\delta}\right)}{\pi_\star\gamma_\star n}} + C \frac{\log\left(\frac{1}{\gamma_\star}\right)}{\gamma_\star n}$$

The remainder of this section is devoted to proving this lemma.

The error $\hat{\mathbf{L}} - \mathbf{L}$ may be written as

$$\hat{\mathbf{L}} - \mathbf{L} = \mathcal{E}_M + \mathcal{E}_\pi \mathbf{L} + \mathbf{L} \mathcal{E}_\pi + \mathcal{E}_\pi \mathbf{L} \mathcal{E}_\pi + \mathcal{E}_\pi \mathcal{E}_M + \mathcal{E}_M \mathcal{E}_\pi + \mathcal{E}_\pi \mathcal{E}_M \mathcal{E}_\pi,$$

where

$$\begin{aligned} \mathcal{E}_\pi &:= \text{Diag}(\hat{\boldsymbol{\pi}})^{-1/2} \text{Diag}(\boldsymbol{\pi})^{1/2} - \mathbf{I}, \\ \mathcal{E}_M &:= \text{Diag}(\boldsymbol{\pi})^{-1/2} (\hat{\mathbf{M}} - \mathbf{M}) \text{Diag}(\boldsymbol{\pi})^{-1/2}. \end{aligned}$$

Therefore

$$\|\hat{\mathbf{L}} - \mathbf{L}\| \leq \|\mathcal{E}_M\| + (\|\mathcal{E}_M\| + \|\mathbf{L}\|) (2\|\mathcal{E}_\pi\| + \|\mathcal{E}_\pi\|^2).$$

Since $\|\mathbf{L}\| \leq 1$ and $\|\hat{\mathbf{L}}\| \leq 1$ [3, Lemma 12.1], we also have $\|\hat{\mathbf{L}} - \mathbf{L}\| \leq 2$. Therefore,

$$\|\hat{\mathbf{L}} - \mathbf{L}\| \leq \min \left\{ \|\mathcal{E}_M\| + (\|\mathcal{E}_M\| + \|\mathbf{L}\|) (2\|\mathcal{E}_\pi\| + \|\mathcal{E}_\pi\|^2), 2 \right\} \leq 3(\|\mathcal{E}_M\| + \|\mathcal{E}_\pi\|). \quad (9)$$

B.3 A bound on $\|\mathcal{E}_\pi\|$

Since $\mathcal{E}_\pi$ is diagonal,

$$\|\mathcal{E}_\pi\| = \max_{i \in [d]} \left| \sqrt{\frac{\pi_i}{\hat{\pi}_i}} - 1 \right|.$$

Assume that

$$n \geq \frac{108 \ln \left(\frac{d}{\delta} \sqrt{\frac{2}{\pi_*}} \right)}{\pi_* \gamma_*}, \quad (10)$$

in which case

$$\sqrt{8\pi_i(1-\pi_i)\varepsilon_n} + 20\varepsilon_n \leq \frac{\pi_i}{2}$$

where ε_n is as defined in Eq. (6). Therefore, in the $1 - \delta$ probability event from Lemma 1, we have $|\pi_i - \hat{\pi}_i| \leq \pi_i/2$ for each $i \in [d]$, and moreover, $2/3 \leq \pi_i/\hat{\pi}_i \leq 2$ for each $i \in [d]$. For this range of $\pi_i/\hat{\pi}_i$, we have

$$\left| \sqrt{\frac{\pi_i}{\hat{\pi}_i}} - 1 \right| \leq \left| \frac{\hat{\pi}_i}{\pi_i} - 1 \right|.$$

We conclude that if n satisfies Eq. (10), then in this $1 - \delta$ probability event from Lemma 1,

$$\begin{aligned} \|\mathcal{E}_\pi\| &\leq \max_{i \in [d]} \left| \frac{\hat{\pi}_i}{\pi_i} - 1 \right| \leq \max_{i \in [d]} \frac{\sqrt{8\pi_i(1-\pi_i)\varepsilon_n} + 20\varepsilon_n}{\pi_i} \\ &\leq \sqrt{\frac{8\varepsilon_n}{\pi_*}} + \frac{20\varepsilon_n}{\pi_*} = \sqrt{\frac{8 \ln \left(\frac{d}{\delta} \sqrt{\frac{2}{\pi_*}} \right)}{\pi_* \gamma_* n}} + \frac{20 \ln \left(\frac{d}{\delta} \sqrt{\frac{2}{\pi_*}} \right)}{\pi_* \gamma_* n}. \end{aligned} \quad (11)$$

B.4 Accuracy of doublet frequency estimates (bounding $\|\mathcal{E}_M\|$)

In this section we prove a bound on $\|\mathcal{E}_M\|$. For this, we decompose $\mathcal{E}_M = \text{Diag}(\pi)^{-1/2}(\widehat{M} - M) \text{Diag}(\pi)^{-1/2}$ into $\mathbb{E}(\mathcal{E}_M)$ and $\mathcal{E}_M - \mathbb{E}(\mathcal{E}_M)$, the first measuring the effect of a non-stationary start of the chain, while the second measuring the variation due to randomness.

B.4.1 Bounding $\|\mathbb{E}(\mathcal{E}_M)\|$: The price of a non-stationary start.

Let $\pi^{(t)}$ be the distribution of states at time step t . We will make use of the following proposition, which can be derived by following [4, Proposition 1.12]:

Proposition 1. For $t \geq 1$, let $\Upsilon^{(t)}$ be the vector with $\Upsilon_i^{(t)} = \frac{\pi_i^{(t)}}{\pi_i}$ and let $\|\cdot\|_{2,\pi}$ denote the π -weighted 2-norm

$$\|v\|_{2,\pi} := \left(\sum_{i=1}^d \pi_i v_i^2 \right)^{1/2}. \quad (12)$$

Then,

$$\|\Upsilon^{(t)} - \mathbf{1}\|_{2,\pi} \leq \frac{(1 - \gamma_*)^{t-1}}{\sqrt{\pi_*}}. \quad (13)$$

An immediate corollary of this result is that

$$\left\| \text{Diag}(\pi^{(t)}) \text{Diag}(\pi)^{-1} - \mathbf{I} \right\| \leq \frac{(1 - \gamma_*)^{t-1}}{\pi_*}. \quad (14)$$

Now note that

$$\mathbb{E}(\widehat{M}) = \frac{1}{n-1} \sum_{t=1}^{n-1} \text{Diag}(\pi^{(t)}) P$$

and thus

$$\begin{aligned}
\mathbb{E}(\mathcal{E}_M) &= \text{Diag}(\boldsymbol{\pi})^{-1/2} \left(\mathbb{E}(\widehat{\mathbf{M}}) - \mathbf{M} \right) \text{Diag}(\boldsymbol{\pi})^{-1/2} \\
&= \frac{1}{n-1} \sum_{t=1}^{n-1} \text{Diag}(\boldsymbol{\pi})^{-1/2} (\text{Diag}(\boldsymbol{\pi}^{(t)}) - \text{Diag}(\boldsymbol{\pi})) \mathbf{P} \text{Diag}(\boldsymbol{\pi})^{-1/2} \\
&= \frac{1}{n-1} \sum_{t=1}^{n-1} \text{Diag}(\boldsymbol{\pi})^{-1/2} (\text{Diag}(\boldsymbol{\pi}^{(t)}) \text{Diag}(\boldsymbol{\pi})^{-1} - \mathbf{I}) \mathbf{M} \text{Diag}(\boldsymbol{\pi})^{-1/2} \\
&= \frac{1}{n-1} \sum_{t=1}^{n-1} (\text{Diag}(\boldsymbol{\pi}^{(t)}) \text{Diag}(\boldsymbol{\pi})^{-1} - \mathbf{I}) \mathbf{L}.
\end{aligned}$$

Combining this, $\|\mathbf{L}\| \leq 1$ and Eq. (14), we get

$$\|\mathbb{E}(\mathcal{E}_M)\| \leq \frac{1}{(n-1)\pi_\star} \sum_{t=1}^{n-1} (1 - \gamma_\star)^{t-1} \leq \frac{1}{(n-1)\gamma_\star\pi_\star}. \quad (15)$$

B.4.2 Bounding $\|\mathcal{E}_M - \mathbb{E}(\mathcal{E}_M)\|$: Application of a matrix tail inequality

In this section we analyze the deviations of $\mathcal{E}_M - \mathbb{E}(\mathcal{E}_M)$. By the definition of $\mathcal{E}_M$,

$$\|\mathcal{E}_M - \mathbb{E}(\mathcal{E}_M)\| = \|\text{Diag}(\boldsymbol{\pi})^{-1/2} (\widehat{\mathbf{M}} - \mathbb{E}\widehat{\mathbf{M}}) \text{Diag}(\boldsymbol{\pi})^{-1/2}\|. \quad (16)$$

The matrix $\widehat{\mathbf{M}} - \mathbb{E}(\widehat{\mathbf{M}})$ is defined as a sum of dependent centered random matrices. We will use the blocking technique of [5] to relate the likely deviations of this matrix to that of a sum of independent centered random matrices. The deviations of these will then be bounded with the help of a Bernstein-type matrix tail inequality due to [6].

We divide $[n-1]$ into contiguous blocks of time steps; each has size $a \leq n/3$ except possibly the first block, which has size between a and $2a-1$. Formally, let $a' := a + ((n-1) \bmod a) \leq 2a-1$, and define

$$\begin{aligned}
F &:= [a'], \\
H_s &:= \{t \in [n-1] : a' + 2(s-1)a + 1 \leq t \leq a' + (2s-1)a\}, \\
T_s &:= \{t \in [n-1] : a' + (2s-1)a + 1 \leq t \leq a' + 2sa\},
\end{aligned}$$

for $s = 1, 2, \dots$. Let μ_H (resp., μ_T) be the number of non-empty H_s (resp., T_s) blocks. Let $n_H := a\mu_H$ (resp., $n_T := a\mu_T$) be the number of time steps in $\cup_s H_s$ (resp., $\cup_s T_s$). We have

$$\begin{aligned}
\widehat{\mathbf{M}} &= \frac{1}{n-1} \sum_{t=1}^{n-1} \mathbf{e}_{X_t} \mathbf{e}_{X_{t+1}}^\top \\
&= \frac{a'}{n-1} \cdot \underbrace{\frac{1}{a'} \sum_{t \in F} \mathbf{e}_{X_t} \mathbf{e}_{X_{t+1}}^\top}_{\widehat{\mathbf{M}}_F} + \frac{n_H}{n-1} \cdot \underbrace{\frac{1}{\mu_H} \sum_{s=1}^{\mu_H} \left(\frac{1}{a} \sum_{t \in H_s} \mathbf{e}_{X_t} \mathbf{e}_{X_{t+1}}^\top \right)}_{\widehat{\mathbf{M}}_H} \\
&\quad + \frac{n_T}{n-1} \cdot \underbrace{\frac{1}{\mu_T} \sum_{s=1}^{\mu_T} \left(\frac{1}{a} \sum_{t \in T_s} \mathbf{e}_{X_t} \mathbf{e}_{X_{t+1}}^\top \right)}_{\widehat{\mathbf{M}}_T}. \quad (17)
\end{aligned}$$

Here, $\mathbf{e}_i$ is the i -th coordinate basis vector, so $\mathbf{e}_i \mathbf{e}_j^\top \in \{0, 1\}^{d \times d}$ is a $d \times d$ matrix of all zeros except for a 1 in the (i, j) -th position.

The contribution of the first block is easily bounded using the triangle inequality:

$$\begin{aligned}
&\frac{a'}{n-1} \left\| \text{Diag}(\boldsymbol{\pi})^{-1/2} (\widehat{\mathbf{M}}_F - \mathbb{E}(\widehat{\mathbf{M}}_F)) \text{Diag}(\boldsymbol{\pi})^{-1/2} \right\| \\
&\leq \frac{1}{n-1} \sum_{t \in F} \left\{ \left\| \frac{\mathbf{e}_{X_t} \mathbf{e}_{X_{t+1}}^\top}{\sqrt{\pi_{X_t} \pi_{X_{t+1}}}} \right\| + \left\| \mathbb{E} \left(\frac{\mathbf{e}_{X_t} \mathbf{e}_{X_{t+1}}^\top}{\sqrt{\pi_{X_t} \pi_{X_{t+1}}}} \right) \right\| \right\} \leq \frac{2a'}{\pi_\star(n-1)}. \quad (18)
\end{aligned}$$

It remains to bound the contributions of the H_s blocks and the T_s blocks. We just focus on the H_s blocks, since the analysis is identical for the T_s blocks.

Let

$$\mathbf{Y}_s := \frac{1}{a} \sum_{t \in H_s} \mathbf{e}_{X_t} \mathbf{e}_{X_{t+1}}^\top, \quad s \in [\mu_H],$$

so

$$\widehat{\mathbf{M}}_H = \frac{1}{\mu_H} \sum_{s=1}^{\mu_H} \mathbf{Y}_s,$$

an average of the random matrices $\mathbf{Y}_s$. For each $s \in [\mu_H]$, the random matrix $\mathbf{Y}_s$ is a function of

$$(X_t : a' + 2(s-1)a + 1 \leq t \leq a' + (2s-1)a + 1)$$

(note the $+1$ in the upper limit of t), so $\mathbf{Y}_{s+1}$ is a time steps ahead of $\mathbf{Y}_s$. When a is sufficiently large, we will be able to effectively treat the random matrices $\mathbf{Y}_s$ as if they were independent. In the sequel, we shall always assume that the block length a satisfies

$$a \geq a_\delta := \frac{1}{\gamma_\star} \ln \frac{2(n-2)}{\delta \pi_\star} \quad (19)$$

for $\delta \in (0, 1)$.

Define

$$\boldsymbol{\pi}^{(H_s)} := \frac{1}{a} \sum_{t \in H_s} \boldsymbol{\pi}^{(t)}, \quad \boldsymbol{\pi}^{(H)} := \frac{1}{\mu_H} \sum_{s=1}^{\mu_H} \boldsymbol{\pi}^{(H_s)}.$$

Observe that

$$\mathbb{E}(\mathbf{Y}_s) = \text{Diag}(\boldsymbol{\pi}^{(H_s)}) \mathbf{P}$$

so

$$\mathbb{E} \left(\frac{1}{\mu_H} \sum_{s=1}^{\mu_H} \mathbf{Y}_s \right) = \text{Diag}(\boldsymbol{\pi}^{(H)}) \mathbf{P}.$$

Define

$$\mathbf{Z}_s := \text{Diag}(\boldsymbol{\pi})^{-1/2} (\mathbf{Y}_s - \mathbb{E}(\mathbf{Y}_s)) \text{Diag}(\boldsymbol{\pi})^{-1/2}.$$

We apply a matrix tail inequality to the average of *independent* copies of the $\mathbf{Z}_s$'s. More precisely, we will apply the tail inequality to independent copies $\tilde{\mathbf{Z}}_s$, $s \in [\mu_H]$ of the random variables $\mathbf{Z}_s$ and then relate the average of $\tilde{\mathbf{Z}}_s$ to that of $\mathbf{Z}_s$. The following probability inequality is from [6, Theorem 6.1.1.].

Theorem 8 (Matrix Bernstein inequality). *Let $\mathbf{Q}_1, \mathbf{Q}_2, \dots, \mathbf{Q}_m$ be a sequence of independent, random $d_1 \times d_2$ matrices. Assume that $\mathbb{E}(\mathbf{Q}_i) = \mathbf{0}$ and $\|\mathbf{Q}_i\| \leq R$ for each $1 \leq i \leq m$. Let $\mathbf{S} = \sum_{i=1}^m \mathbf{Q}_i$ and let*

$$v = \max \left\{ \|\mathbb{E} \sum_i \mathbf{Q}_i \mathbf{Q}_i^\top\|, \|\mathbb{E} \sum_i \mathbf{Q}_i^\top \mathbf{Q}_i\| \right\}.$$

Then, for all $t \geq 0$,

$$\mathbb{P}(\|\mathbf{S}\| \geq t) \leq 2(d_1 + d_2) \exp \left(-\frac{t^2/2}{v + Rt/3} \right).$$

In other words, for any $\delta \in (0, 1)$,

$$\mathbb{P} \left(\|\mathbf{S}\| > \sqrt{2v \ln \frac{2(d_1 + d_2)}{\delta}} + \frac{2R}{3} \ln \frac{2(d_1 + d_2)}{\delta} \right) \leq \delta.$$

To apply Theorem 8, it suffices to bound the spectral norms of $\mathbf{Z}_s$ (almost surely), $\mathbb{E}(\mathbf{Z}_s \mathbf{Z}_s^\top)$, and $\mathbb{E}(\mathbf{Z}_s^\top \mathbf{Z}_s)$.

Range bound. By the triangle inequality,

$$\|\mathbf{Z}_s\| \leq \|\text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbf{Y}_s \text{Diag}(\boldsymbol{\pi})^{-1/2}\| + \|\text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbb{E}(\mathbf{Y}_s) \text{Diag}(\boldsymbol{\pi})^{-1/2}\|.$$

For the first term, we have

$$\|\text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbf{Y}_s \text{Diag}(\boldsymbol{\pi})^{-1/2}\| \leq \frac{1}{\pi_\star}. \quad (20)$$

For the second term, we use the fact $\|\mathbf{L}\| \leq 1$ to bound

$$\begin{aligned} \|\text{Diag}(\boldsymbol{\pi})^{-1/2} (\mathbb{E}(\mathbf{Y}_s) - \mathbf{M}) \text{Diag}(\boldsymbol{\pi})^{-1/2}\| &= \|(\text{Diag}(\boldsymbol{\pi}^{(H_s)}) \text{Diag}(\boldsymbol{\pi})^{-1} - \mathbf{I}) \mathbf{L}\| \\ &\leq \|\text{Diag}(\boldsymbol{\pi}^{(H_s)}) \text{Diag}(\boldsymbol{\pi})^{-1} - \mathbf{I}\|. \end{aligned}$$

Then, using Eq. (14),

$$\|\text{Diag}(\boldsymbol{\pi}^{(H_s)}) \text{Diag}(\boldsymbol{\pi})^{-1} - \mathbf{I}\| \leq \frac{(1 - \gamma_\star)^{a' + 2(s-1)a}}{\pi_\star} \leq \frac{(1 - \gamma_\star)^a}{\pi_\star} \leq 1, \quad (21)$$

where the last inequality follows from the assumption that the block length a satisfies Eq. (19). Combining this with $\|\text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbf{M} \text{Diag}(\boldsymbol{\pi})^{-1/2}\| = \|\mathbf{L}\| \leq 1$, it follows that

$$\|\text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbb{E}(\mathbf{Y}_s) \text{Diag}(\boldsymbol{\pi})^{-1/2}\| \leq 2 \quad (22)$$

by the triangle inequality. Therefore, together with Eq. (20), we obtain the range bound

$$\|\mathbf{Z}_s\| \leq \frac{1}{\pi_\star} + 2.$$

Variance bound. We now determine bounds on the spectral norms of $\mathbb{E}(\mathbf{Z}_s \mathbf{Z}_s^\top)$ and $\mathbb{E}(\mathbf{Z}_s^\top \mathbf{Z}_s)$. Observe that

$$\begin{aligned} \mathbb{E}(\mathbf{Z}_s \mathbf{Z}_s^\top) &= \frac{1}{a^2} \sum_{t \in H_s} \mathbb{E} \left(\text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbf{e}_{X_t} \mathbf{e}_{X_{t+1}}^\top \text{Diag}(\boldsymbol{\pi})^{-1} \mathbf{e}_{X_{t+1}} \mathbf{e}_{X_t}^\top \text{Diag}(\boldsymbol{\pi})^{-1/2} \right) \end{aligned} \quad (23)$$

$$+ \frac{1}{a^2} \sum_{\substack{t \neq t' \\ t, t' \in H_s}} \mathbb{E} \left(\text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbf{e}_{X_t} \mathbf{e}_{X_{t+1}}^\top \text{Diag}(\boldsymbol{\pi})^{-1} \mathbf{e}_{X_{t'+1}} \mathbf{e}_{X_{t'}}^\top \text{Diag}(\boldsymbol{\pi})^{-1/2} \right) \quad (24)$$

$$- \text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbb{E}(\mathbf{Y}_s) \text{Diag}(\boldsymbol{\pi})^{-1} \mathbb{E}(\mathbf{Y}_s^\top) \text{Diag}(\boldsymbol{\pi})^{-1/2}. \quad (25)$$

The first sum, Eq. (23), easily simplifies to the diagonal matrix

$$\begin{aligned} &\frac{1}{a^2} \sum_{t \in H_s} \sum_{i=1}^d \sum_{j=1}^d \Pr(X_t = i, X_{t+1} = j) \cdot \frac{1}{\pi_i \pi_j} \mathbf{e}_i \mathbf{e}_j^\top \mathbf{e}_j \mathbf{e}_i^\top \\ &= \frac{1}{a^2} \sum_{t \in H_s} \sum_{i=1}^d \sum_{j=1}^d \pi_i^{(t)} P_{i,j} \cdot \frac{1}{\pi_i \pi_j} \mathbf{e}_i \mathbf{e}_i^\top = \frac{1}{a} \sum_{i=1}^d \frac{\pi_i^{(H_s)}}{\pi_i} \left(\sum_{j=1}^d \frac{P_{i,j}}{\pi_j} \right) \mathbf{e}_i \mathbf{e}_i^\top. \end{aligned}$$

For the second sum, Eq. (24), a symmetric matrix, consider

$$\mathbf{u}^\top \left(\frac{1}{a^2} \sum_{\substack{t \neq t' \\ t, t' \in H_s}} \mathbb{E} \left(\text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbf{e}_{X_t} \mathbf{e}_{X_{t+1}}^\top \text{Diag}(\boldsymbol{\pi})^{-1} \mathbf{e}_{X_{t'+1}} \mathbf{e}_{X_{t'}}^\top \text{Diag}(\boldsymbol{\pi})^{-1/2} \right) \right) \mathbf{u}$$

for an arbitrary unit vector $\mathbf{u}$. By Cauchy-Schwarz and AM/GM, this is bounded from above by

$$\begin{aligned} &\frac{1}{2a^2} \sum_{\substack{t \neq t' \\ t, t' \in H_s}} \left[\mathbb{E} \left(\mathbf{u}^\top \text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbf{e}_{X_t} \mathbf{e}_{X_{t+1}}^\top \text{Diag}(\boldsymbol{\pi})^{-1} \mathbf{e}_{X_{t+1}} \mathbf{e}_{X_t}^\top \text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbf{u} \right) \right. \\ &\quad \left. + \mathbb{E} \left(\mathbf{u}^\top \text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbf{e}_{X_{t'}} \mathbf{e}_{X_{t'+1}}^\top \text{Diag}(\boldsymbol{\pi})^{-1} \mathbf{e}_{X_{t'+1}} \mathbf{e}_{X_{t'}}^\top \text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbf{u} \right) \right], \end{aligned}$$

which simplifies to

$$\frac{a-1}{a^2} \mathbf{u}^\top \mathbb{E} \left(\sum_{t \in H_s} \text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbf{e}_{X_t} \mathbf{e}_{X_{t+1}}^\top \text{Diag}(\boldsymbol{\pi})^{-1} \mathbf{e}_{X_{t+1}} \mathbf{e}_{X_t}^\top \text{Diag}(\boldsymbol{\pi})^{-1/2} \right) \mathbf{u}.$$

The expectation is the same as that for the first term, Eq. (23).

Finally, the spectral norm of the third term, Eq. (25), is bounded using Eq. (22):

$$\| \text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbb{E}(\mathbf{Y}_s) \text{Diag}(\boldsymbol{\pi})^{-1/2} \|^2 \leq 4.$$

Therefore, by the triangle inequality, the bound $\pi_i^{(H)}/\pi_i \leq 2$ from Eq. (21), and simplifications,

$$\| \mathbb{E}(\mathbf{Z}_s \mathbf{Z}_s^\top) \| \leq \max_{i \in [d]} \left(\sum_{j=1}^d \frac{P_{i,j}}{\pi_j} \right) \frac{\pi_i^{(H)}}{\pi_i} + 4 \leq 2 \max_{i \in [d]} \left(\sum_{j=1}^d \frac{P_{i,j}}{\pi_j} \right) + 4.$$

We can bound $\mathbb{E}(\mathbf{Z}_s^\top \mathbf{Z}_s)$ in a similar way; the only difference is that the reversibility needs to be used at one place to simplify an expectation:

$$\begin{aligned} & \frac{1}{a^2} \sum_{t \in H_s} \mathbb{E} \left(\text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbf{e}_{X_{t+1}} \mathbf{e}_{X_t}^\top \text{Diag}(\boldsymbol{\pi})^{-1} \mathbf{e}_{X_t} \mathbf{e}_{X_{t+1}}^\top \text{Diag}(\boldsymbol{\pi})^{-1/2} \right) \\ &= \frac{1}{a^2} \sum_{t \in H_s} \sum_{i=1}^d \sum_{j=1}^d \Pr(X_t = i, X_{t+1} = j) \cdot \frac{1}{\pi_i \pi_j} \mathbf{e}_j \mathbf{e}_j^\top \\ &= \frac{1}{a^2} \sum_{t \in H_s} \sum_{i=1}^d \sum_{j=1}^d \pi_i^{(t)} P_{i,j} \cdot \frac{1}{\pi_i \pi_j} \mathbf{e}_j \mathbf{e}_j^\top \\ &= \frac{1}{a^2} \sum_{t \in H_s} \sum_{j=1}^d \left(\sum_{i=1}^d \frac{\pi_i^{(t)}}{\pi_i} \cdot \frac{P_{j,i}}{\pi_i} \right) \mathbf{e}_j \mathbf{e}_j^\top \end{aligned}$$

where the last step uses Eq. (3). As before, we get

$$\| \mathbb{E}(\mathbf{Z}_s^\top \mathbf{Z}_s) \| \leq \max_{i \in [d]} \left(\sum_{j=1}^d \frac{P_{i,j}}{\pi_j} \cdot \frac{\pi_j^{(H)}}{\pi_j} \right) + 4 \leq 2 \max_{i \in [d]} \left(\sum_{j=1}^d \frac{P_{i,j}}{\pi_j} \right) + 4$$

again using the bound $\pi_i^{(H)}/\pi_i \leq 2$ from Eq. (21).

Independent copies bound. Let $\tilde{\mathbf{Z}}_s$ for $s \in [\mu_H]$ be independent copies of $\mathbf{Z}_s$ for $s \in [\mu_H]$. Applying Theorem 8 to the average of these random matrices, we have

$$\mathbb{P} \left(\left\| \frac{1}{\mu_H} \sum_{s=1}^{\mu_H} \tilde{\mathbf{Z}}_s \right\| > \sqrt{\frac{4(d_{\mathbf{P}} + 2) \ln \frac{4d}{\delta}}{\mu_H}} + \frac{2 \left(\frac{1}{\pi_\star} + 2 \right) \ln \frac{4d}{\delta}}{3\mu_H} \right) \leq \delta \quad (26)$$

where

$$d_{\mathbf{P}} := \max_{i \in [d]} \sum_{j=1}^d \frac{P_{i,j}}{\pi_j} \leq \frac{1}{\pi_\star}.$$

The actual bound. To bound the probability that $\| \sum_{s=1}^{\mu_H} \mathbf{Z}_s / \mu_H \|$ is large, we appeal to the following result, which is a consequence of [7, Corollary 2.7]. For each $s \in [\mu_H]$, let $X^{(H_s)} := (X_t : a' + 2(s-1)a + 1 \leq t \leq a' + (2s-1)a + 1)$, which are the random variables determining $\mathbf{Z}_s$. Let $\mathbb{P}$ denote the joint distribution of $(X^{(H_s)} : s \in [\mu_H])$; let $\mathbb{P}_s$ be its marginal over $X^{(H_s)}$, and let $\mathbb{P}_{1:s+1}$ be its marginal over $(X^{(H_1)}, X^{(H_2)}, \dots, X^{(H_{s+1})})$. Let $\tilde{\mathbb{P}}$ be the product distribution

formed from the marginals $\mathbb{P}_1, \mathbb{P}_2, \dots, \mathbb{P}_{\mu_H}$, so $\tilde{\mathbb{P}}$ governs the joint distribution of $(\tilde{\mathbf{Z}}_s : s \in [\mu_H])$. The result from [7, Corollary 2.7] implies for any event E ,

$$|\mathbb{P}(E) - \tilde{\mathbb{P}}(E)| \leq (\mu_H - 1)\beta(\mathbb{P})$$

where

$$\beta(\mathbb{P}) := \max_{1 \leq s \leq \mu_H - 1} \mathbb{E} \left(\left\| \mathbb{P}_{1:s+1}(\cdot | X^{(H_1)}, X^{(H_2)}, \dots, X^{(H_s)}) - \mathbb{P}_{s+1} \right\|_{\text{tv}} \right).$$

Here, $\|\cdot\|_{\text{tv}}$ denotes the total variation norm. The number $\beta(\mathbb{P})$ can be recognized to be the β -mixing coefficient of the stochastic process $\{X^{(H_s)}\}_{s \in [\mu_H]}$. This result implies that the bound from Eq. (26) for $\|\sum_{s=1}^{\mu_H} \tilde{\mathbf{Z}}_s / \mu_H\|$ also holds for $\|\sum_{s=1}^{\mu_H} \mathbf{Z}_s / \mu_H\|$, except the probability bound increases from δ to $\delta + (\mu_H - 1)\beta(\mathbb{P})$:

$$\mathbb{P} \left(\left\| \frac{1}{\mu_H} \sum_{s=1}^{\mu_H} \mathbf{Z}_s \right\| > \sqrt{\frac{4(d\mathbf{P} + 2) \ln \frac{4d}{\delta}}{\mu_H}} + \frac{2 \left(\frac{1}{\pi_*} + 2 \right) \ln \frac{4d}{\delta}}{3\mu_H} \right) \leq \delta + (\mu_H - 1)\beta(\mathbb{P}). \quad (27)$$

By the triangle inequality,

$$\beta(\mathbb{P}) \leq \max_{1 \leq s \leq \mu_H - 1} \mathbb{E} \left(\left\| \mathbb{P}_{1:s+1}(\cdot | X^{(H_1)}, X^{(H_2)}, \dots, X^{(H_s)}) - \mathbb{P}^\pi \right\|_{\text{tv}} + \left\| \mathbb{P}_{s+1} - \mathbb{P}^\pi \right\|_{\text{tv}} \right)$$

where $\mathbb{P}^\pi$ is the marginal distribution of $X^{(H_1)}$ under the stationary chain. Using the Markov property and integrating out X_t for $t > \min H_{s+1} = a' + 2sa + 1$,

$$\left\| \mathbb{P}_{1:s+1}(\cdot | X^{(H_1)}, X^{(H_2)}, \dots, X^{(H_s)}) - \mathbb{P}^\pi \right\|_{\text{tv}} = \left\| \mathcal{L}(X_{a'+2sa+1} | X_{a'+(2s-1)a+1}) - \pi \right\|_{\text{tv}}$$

where $\mathcal{L}(Y|Z)$ denotes the conditional distribution of Y given Z . We bound this distance using standard arguments for bounding the mixing time in terms of the *relaxation time* $1/\gamma_*$ [see, e.g., the proof of Theorem 12.3 of 3]: for any $i \in [d]$,

$$\left\| \mathcal{L}(X_{a'+2sa+1} | X_{a'+(2s-1)a+1} = i) - \pi \right\|_{\text{tv}} = \left\| \mathcal{L}(X_{a+1} | X_1 = i) - \pi \right\|_{\text{tv}} \leq \frac{\exp(-a\gamma_*)}{\pi_*}.$$

The distance $\|\mathbb{P}_{s+1} - \mathbb{P}^\pi\|_{\text{tv}}$ can be bounded similarly:

$$\begin{aligned} \|\mathbb{P}_{s+1} - \mathbb{P}^\pi\|_{\text{tv}} &= \|\mathcal{L}(X_{a'+2sa+1}) - \pi\|_{\text{tv}} \\ &= \left\| \sum_{i=1}^d \mathbb{P}(X_1 = i) \mathcal{L}(X_{a'+2sa+1} | X_1 = i) - \pi \right\|_{\text{tv}} \\ &\leq \sum_{i=1}^d \mathbb{P}(X_1 = i) \|\mathcal{L}(X_{a'+2sa+1} | X_1 = i) - \pi\|_{\text{tv}} \\ &\leq \frac{\exp(-(a' + 2sa)\gamma_*)}{\pi_*} \leq \frac{\exp(-a\gamma_*)}{\pi_*}. \end{aligned}$$

We conclude

$$(\mu_H - 1)\beta(\mathbb{P}) \leq (\mu_H - 1) \frac{2 \exp(-a\gamma_*)}{\pi_*} \leq \frac{2(n-2) \exp(-a\gamma_*)}{\pi_*} \leq \delta$$

where the last step follows from the block length assumption Eq. (19).

We return to the decomposition from Eq. (17). We apply Eq. (27) to both the H_s blocks and the T_s blocks, and combine with Eq. (18) to obtain the following probabilistic bound. Pick any $\delta \in (0, 1)$, let the block length be

$$a := \lceil a_\delta \rceil = \left\lceil \frac{1}{\gamma_*} \ln \frac{2(n-2)}{\pi_* \delta} \right\rceil,$$

so

$$\min\{\mu_H, \mu_T\} = \left\lfloor \frac{n-1-a'}{2a} \right\rfloor \geq \frac{n-1}{2 \left(1 + \frac{1}{\gamma_*} \ln \frac{2(n-2)}{\pi_* \delta} \right)} - 2 =: \mu.$$

If

$$n \geq 7 + \frac{6}{\gamma_\star} \ln \frac{2(n-2)}{\pi_\star \delta} \geq 3a, \quad (28)$$

then with probability at least $1 - 4\delta$,

$$\begin{aligned} & \left\| \text{Diag}(\boldsymbol{\pi})^{-1/2} \left(\widehat{\mathbf{M}} - \mathbb{E}[\widehat{\mathbf{M}}] \right) \text{Diag}(\boldsymbol{\pi})^{-1/2} \right\| \\ & \leq \frac{4 \left\lceil \frac{1}{\gamma_\star} \ln \frac{2(n-2)}{\pi_\star \delta} \right\rceil}{\pi_\star(n-1)} + \sqrt{\frac{4(d_{\mathbf{P}} + 2) \ln \frac{4d}{\delta}}{\mu}} + \frac{2 \left(\frac{1}{\pi_\star} + 2 \right) \ln \frac{4d}{\delta}}{3\mu}. \end{aligned}$$

B.4.3 The bound on $\|\mathcal{E}_M\|$

Combining the probabilistic bound from above with the bound on the bias from Eq. (15), we obtain the following. Assuming the condition on n from Eq. (28), with probability at least $1 - 4\delta$,

$$\|\mathcal{E}_M\| \leq \frac{1}{(n-1)\gamma_\star\pi_\star} + \frac{4 \left\lceil \frac{1}{\gamma_\star} \ln \frac{2(n-2)}{\pi_\star \delta} \right\rceil}{\pi_\star(n-1)} + \sqrt{\frac{4(d_{\mathbf{P}} + 2) \ln \frac{4d}{\delta}}{\mu}} + \frac{2 \left(\frac{1}{\pi_\star} + 2 \right) \ln \frac{4d}{\delta}}{3\mu}. \quad (29)$$

B.5 Overall error bound

Assume the sequence length n satisfies Eq. (10) and Eq. (28). Consider the $1 - 5\delta$ probability event in which Eqs. (7), (8) and (29) hold. In this event, Eq. (11) also holds, and hence by Eq. (9),

$$\begin{aligned} \|\widehat{\mathbf{L}} - \mathbf{L}\| & \leq 3 \left(\sqrt{\frac{8 \ln \left(\frac{d}{\delta} \sqrt{\frac{2}{\pi_\star}} \right)}{\pi_\star \gamma_\star n}} + \frac{20 \ln \left(\frac{d}{\delta} \sqrt{\frac{2}{\pi_\star}} \right)}{\pi_\star \gamma_\star n} \right) \\ & + 3 \left(\frac{1}{(n-1)\gamma_\star\pi_\star} + \frac{4 \left\lceil \frac{1}{\gamma_\star} \ln \frac{2(n-2)}{\pi_\star \delta} \right\rceil}{\pi_\star(n-1)} + \sqrt{\frac{4(d_{\mathbf{P}} + 2) \ln \frac{4d}{\delta}}{\mu}} + \frac{2 \left(\frac{1}{\pi_\star} + 2 \right) \ln \frac{4d}{\delta}}{3\mu} \right) \\ & = O \left(\sqrt{\frac{\log \left(\frac{d}{\delta} \right) \log \left(\frac{n}{\pi_\star \delta} \right)}{\pi_\star \gamma_\star n}} \right). \end{aligned}$$

To finish the proof of Lemma 2, we replace δ with $\delta/5$, and now observe that the bound on $\|\widehat{\mathbf{L}} - \mathbf{L}\|$ is trivial if Eq. (10) is violated (recalling that $\|\widehat{\mathbf{L}} - \mathbf{L}\| \leq 2$ always holds). We tack on an additional term $\log(1/(\pi_\star \gamma_\star \delta))/(\gamma_\star n)$ to also ensure a trivial bound if Eq. (28) is violated. $\square$

C Proof of Theorem 4

In this section, we derive Algorithm 1 and prove Theorem 4.

C.1 Estimators for $\boldsymbol{\pi}$ and $\gamma_\star$

The algorithm forms the estimator $\widehat{\mathbf{P}}$ of $\mathbf{P}$ using Laplace smoothing:

$$\widehat{P}_{i,j} := \frac{N_{i,j} + \alpha}{N_i + d\alpha}$$

where

$$N_{i,j} := |\{t \in [n-1] : (X_t, X_{t+1}) = (i, j)\}|, \quad N_i := |\{t \in [n-1] : X_t = i\}|$$

and $\alpha > 0$ is a positive constant, which we set beforehand as $\alpha := 1/d$ for simplicity.

As a result of the smoothing, all entries of $\widehat{\mathbf{P}}$ are positive, and hence $\widehat{\mathbf{P}}$ is a transition probability matrix for an ergodic Markov chain. We let $\widehat{\pi}$ be the unique stationary distribution for $\widehat{\mathbf{P}}$. Using $\widehat{\pi}$, we form an estimator $\text{Sym}(\widehat{\mathbf{L}})$ of $\mathbf{L}$ using:

$$\text{Sym}(\widehat{\mathbf{L}}) := \frac{1}{2}(\widehat{\mathbf{L}} + \widehat{\mathbf{L}}^\top), \quad \widehat{\mathbf{L}} := \text{Diag}(\widehat{\pi})^{1/2} \widehat{\mathbf{P}} \text{Diag}(\widehat{\pi})^{-1/2}.$$

Let $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \dots \geq \hat{\lambda}_d$ be the eigenvalues of $\text{Sym}(\widehat{\mathbf{L}})$ (and in fact, we have $1 = \hat{\lambda}_1 > \hat{\lambda}_2$ and $\hat{\lambda}_d > -1$). The algorithm estimates the spectral gap γ_* using

$$\hat{\gamma}_* := 1 - \max\{\hat{\lambda}_2, |\hat{\lambda}_d|\}.$$

C.2 Empirical bounds for P

We make use of a simple corollary of Freedman's inequality for martingales [8, Theorem 1.6].

Theorem 9 (Freedman's inequality). *Let $(Y_t)_{t \in \mathbb{N}}$ be a bounded martingale difference sequence with respect to the filtration $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \mathcal{F}_2 \subset \dots$; assume for some $b > 0$, $|Y_t| \leq b$ almost surely for all $t \in \mathbb{N}$. Let $V_k := \sum_{t=1}^k \mathbb{E}(Y_t^2 | \mathcal{F}_{t-1})$ and $S_k := \sum_{t=1}^k Y_t$ for $k \in \mathbb{N}$. For all $s, v > 0$,*

$$\Pr[\exists k \in \mathbb{N} \text{ s.t. } S_k > s \wedge V_k \leq v] \leq \left(\frac{v/b^2}{s/b + v/b^2} \right)^{s/b + v/b^2} e^{s/b} = \exp\left(-\frac{v}{b^2} \cdot h\left(\frac{bs}{v}\right)\right),$$

where $h(u) := (1+u)\ln(1+u) - u$.

Observe that in Theorem 9, for any $x > 0$, if $s := \sqrt{2vx} + bx/3$ and $z := b^2x/v$, then the probability bound on the right-hand side becomes

$$\exp\left(-x \cdot \frac{h(\sqrt{2z} + z/3)}{z}\right) \leq e^{-x}$$

since $h(\sqrt{2z} + z/3)/z \geq 1$ for all $z > 0$ (see, e.g., [9, proof of Lemma 5]).

Corollary 1. *Under the same setting as Theorem 9, for any $n \geq 1$, $x > 0$, and $c > 1$,*

$$\Pr[\exists k \in [n] \text{ s.t. } S_k > \sqrt{2cV_kx} + 4bx/3] \leq (1 + \lceil \log_c(2n/x) \rceil_+) e^{-x}.$$

Proof. Define $v_i := c^i b^2 x / 2$ for $i = 0, 1, 2, \dots, \lceil \log_c(2n/x) \rceil_+$, and let $v_{-1} := -\infty$. Then, since $V_k \in [0, b^2 n]$ for all $k \in [n]$,

$$\begin{aligned} & \Pr[\exists k \in [n] \text{ s.t. } S_k > \sqrt{2 \max\{v_0, cV_k\}x} + bx/3] \\ &= \sum_{i=0}^{\lceil \log_c(2n/x) \rceil_+} \Pr[\exists k \in [n] \text{ s.t. } S_k > \sqrt{2 \max\{v_0, cV_k\}x} + bx/3 \wedge v_{i-1} < V_k \leq v_i] \\ &\leq \sum_{i=0}^{\lceil \log_c(2n/x) \rceil_+} \Pr[\exists k \in [n] \text{ s.t. } S_k > \sqrt{2 \max\{v_0, cv_{i-1}\}x} + bx/3 \wedge v_{i-1} < V_k \leq v_i] \\ &\leq \sum_{i=0}^{\lceil \log_c(2n/x) \rceil_+} \Pr[\exists k \in [n] \text{ s.t. } S_k > \sqrt{2v_i x} + bx/3 \wedge V_k \leq v_i] \\ &\leq (1 + \lceil \log_c(2n/x) \rceil_+) e^{-x}, \end{aligned}$$

where the final inequality uses Theorem 9. The conclusion now follows because

$$\sqrt{2cV_kx} + 4bx/3 \geq \sqrt{2 \max\{v_0, cV_k\}x} + bx/3$$

for all $k \in [n]$. □

Lemma 3. *The following holds for any constant $c > 1$ with probability at least $1 - \delta$: for all $(i, j) \in [d]^2$,*

$$|\hat{P}_{i,j} - P_{i,j}| \leq \sqrt{\left(\frac{N_i}{N_i + d\alpha}\right) \frac{2cP_{i,j}(1 - P_{i,j})\tau_{n,\delta}}{N_i + d\alpha}} + \frac{(4/3)\tau_{n,\delta}}{N_i + d\alpha} + \frac{|\alpha - d\alpha P_{i,j}|}{N_i + d\alpha}, \quad (30)$$

where

$$\tau_{n,\delta} := \inf \{t \geq 0 : 2d^2 (1 + \lceil \log_c(2n/t) \rceil_+) e^{-t} \leq \delta\} = O\left(\log\left(\frac{d \log(n)}{\delta}\right)\right). \quad (31)$$

Proof. Let $\mathcal{F}_t$ be the σ -field generated by $X_1, X_2, \dots, X_t$. Fix a pair $(i, j) \in [d]^2$. Let $Y_1 := 0$, and for $t \geq 2$,

$$Y_t := \mathbb{1}\{X_{t-1} = i\} (\mathbb{1}\{X_t = j\} - P_{i,j}),$$

so that

$$\sum_{t=1}^n Y_t = N_{i,j} - N_i P_{i,j}.$$

The Markov property implies that the stochastic process $(Y_t)_{t \in [n]}$ is an $(\mathcal{F}_t)$ -adapted martingale difference sequence: Y_t is $\mathcal{F}_t$ -measurable and $\mathbb{E}(Y_t | \mathcal{F}_{t-1}) = 0$, for each t . Moreover, for all $t \in [n]$,

$$Y_t \in [-P_{i,j}, 1 - P_{i,j}],$$

and for $t \geq 2$,

$$\mathbb{E}(Y_t^2 | \mathcal{F}_{t-1}) = \mathbb{1}\{X_{t-1} = i\} P_{i,j}(1 - P_{i,j}).$$

Therefore, by Corollary 1 and union bounds, we have

$$|N_{i,j} - N_i P_{i,j}| \leq \sqrt{2cN_i P_{i,j}(1 - P_{i,j})\tau_{n,\delta}} + \frac{4\tau_{n,\delta}}{3}$$

for all $(i, j) \in [d]^2$. □

Equation (30) can be viewed as constraints on the possible value that $P_{i,j}$ may have (with high probability). Since $P_{i,j}$ is the only unobserved quantity in the bound from Eq. (30), we can numerically maximize $|\hat{P}_{i,j} - P_{i,j}|$ subject to the constraint in Eq. (30) (viewing $P_{i,j}$ as the optimization variable). Let $B_{i,j}^*$ be this maximum value, so we have

$$P_{i,j} \in [\hat{P}_{i,j} - B_{i,j}^*, \hat{P}_{i,j} + B_{i,j}^*]$$

in the same event where Eq. (30) holds.

In the algorithm, we give a simple alternative to computing $B_{i,j}^*$ that avoids numerical optimization, derived in the spirit of empirical Bernstein bounds [9]. Specifically, with $c := 1.1$ (an arbitrary choice), we compute

$$\hat{B}_{i,j} := \left(\sqrt{\frac{c\tau_{n,\delta}}{2N_i}} + \sqrt{\frac{c\tau_{n,\delta}}{2N_i} + \sqrt{\frac{2c\hat{P}_{i,j}(1 - \hat{P}_{i,j})\tau_{n,\delta}}{N_i}} + \frac{(4/3)\tau_{n,\delta} + |\alpha - d\alpha\hat{P}_{i,j}|}{N_i}} \right)^2 \quad (32)$$

for each $(i, j) \in [d]^2$, where $\tau_{n,\delta}$ is defined in Eq. (31). We show in Lemma 4 that

$$P_{i,j} \in [\hat{P}_{i,j} - \hat{B}_{i,j}, \hat{P}_{i,j} + \hat{B}_{i,j}]$$

again, in the same event where Eq. (30) holds. The observable bound in Eq. (32) is not too far from the unobservable bound in Eq. (30).

Lemma 4. *In the same $1 - \delta$ event as from Lemma 3, we have $P_{i,j} \in [\hat{P}_{i,j} - \hat{B}_{i,j}, \hat{P}_{i,j} + \hat{B}_{i,j}]$ for all $(i, j) \in [d]^2$, where $\hat{B}_{i,j}$ is defined in Eq. (32).*

Proof. Recall that in the $1 - \delta$ probability event from Lemma 3, we have for all $(i, j) \in [d]^2$,

$$|\hat{P}_{i,j} - P_{i,j}| = \left| \frac{N_{i,j} - N_i P_{i,j}}{N_i + d\alpha} + \frac{\alpha - d\alpha P_{i,j}}{N_i + d\alpha} \right| \leq \sqrt{\frac{2cN_i P_{i,j}(1 - P_{i,j})\tau_{n,\delta}}{(N_i + d\alpha)^2}} + \frac{(4/3)\tau_{n,\delta}}{N_i + d\alpha} + \frac{|\alpha - d\alpha P_{i,j}|}{N_i + d\alpha}.$$

Applying the triangle inequality to the right-hand side, we obtain

$$\begin{aligned} |\hat{P}_{i,j} - P_{i,j}| &\leq \sqrt{\frac{2cN_i(\hat{P}_{i,j}(1 - \hat{P}_{i,j}) + |\hat{P}_{i,j} - P_{i,j}|)\tau_{n,\delta}}{(N_i + d\alpha)^2}} + \frac{(4/3)\tau_{n,\delta}}{N_i + d\alpha} \\ &\quad + \frac{|\alpha - d\alpha \hat{P}_{i,j}| + d\alpha|\hat{P}_{i,j} - P_{i,j}|}{N_i + d\alpha}. \end{aligned}$$

Since $\sqrt{A+B} \leq \sqrt{A} + \sqrt{B}$ for non-negative A, B , we loosen the above inequality and rearrange it to obtain

$$\begin{aligned} \left(1 - \frac{d\alpha}{N_i + d\alpha}\right) |\hat{P}_{i,j} - P_{i,j}| &\leq \sqrt{|\hat{P}_{i,j} - P_{i,j}|} \cdot \sqrt{\frac{2cN_i\tau_{n,\delta}}{(N_i + d\alpha)^2}} \\ &\quad + \sqrt{\frac{2cN_i\hat{P}_{i,j}(1 - \hat{P}_{i,j})\tau_{n,\delta}}{(N_i + d\alpha)^2}} + \frac{(4/3)\tau_{n,\delta} + |\alpha - d\alpha \hat{P}_{i,j}|}{N_i + d\alpha}. \end{aligned}$$

Whenever $N_i > 0$, we can solve a quadratic inequality to conclude $|\hat{P}_{i,j} - P_{i,j}| \leq \hat{B}_{i,j}$. $\square$

C.3 Empirical bounds for π

Recall that $\hat{\pi}$ is obtained as the unique stationary distribution for $\hat{P}$. Let $\hat{A} := I - \hat{P}$, and let $\hat{A}^\#$ be the *group inverse* of $\hat{A}$ —i.e., the unique square matrix satisfying the following equalities:

$$\hat{A}\hat{A}^\#\hat{A} = \hat{A}, \quad \hat{A}^\#\hat{A}\hat{A}^\# = \hat{A}^\#, \quad \hat{A}^\#\hat{A} = \hat{A}\hat{A}^\#.$$

The matrix $\hat{A}^\#$, which is well defined no matter what transition probability matrix $\hat{P}$ we start with [10], is a central quantity that captures many properties of the ergodic Markov chain with transition matrix $\hat{P}$ [10]. We denote the (i, j) -th entry of $\hat{A}^\#$ by $\hat{A}_{i,j}^\#$. Define

$$\hat{\kappa} := \frac{1}{2} \max \left\{ \hat{A}_{j,j}^\# - \min \left\{ \hat{A}_{i,j}^\# : i \in [d] \right\} : j \in [d] \right\}.$$

Analogously define

$$\begin{aligned} A &:= I - P, \\ A^\# &:= \text{group inverse of } A, \\ \kappa &:= \frac{1}{2} \max \left\{ A_{j,j}^\# - \min \left\{ A_{i,j}^\# : i \in [d] \right\} : j \in [d] \right\}. \end{aligned}$$

We now use the following perturbation bound from [11, Section 3.3] (derived from [12, 13]).

Lemma 5 ([12, 13]). *If $|\hat{P}_{i,j} - P_{i,j}| \leq \hat{B}_{i,j}$ for each $(i, j) \in [d]^2$, then*

$$\begin{aligned} \max \{ |\hat{\pi}_i - \pi_i| : i \in [d] \} &\leq \min \{ \kappa, \hat{\kappa} \} \max \{ \hat{B}_{i,j} : (i, j) \in [d]^2 \} \\ &\leq \hat{\kappa} \max \{ \hat{B}_{i,j} : (i, j) \in [d]^2 \}. \end{aligned}$$

This establishes the validity of the confidence intervals for the π_i in the same event from Lemma 3.

We now establish the validity of the bounds for the ratio quantities $\sqrt{\hat{\pi}_i/\pi_i}$ and $\sqrt{\pi_i/\hat{\pi}_i}$.

Lemma 6. *If $\max \{ |\hat{\pi}_i - \pi_i| : i \in [d] \} \leq \hat{b}$, then*

$$\max \bigcup_{i \in [d]} \{ |\sqrt{\pi_i/\hat{\pi}_i} - 1|, |\sqrt{\hat{\pi}_i/\pi_i} - 1| \} \leq \frac{1}{2} \max \bigcup_{i \in [d]} \left\{ \frac{\hat{b}}{\hat{\pi}_i}, \frac{\hat{b}}{[\hat{\pi}_i - \hat{b}]_+} \right\}.$$

Proof. By Lemma 5, we have for each $i \in [d]$,

$$\frac{|\hat{\pi}_i - \pi_i|}{\hat{\pi}_i} \leq \frac{\hat{b}}{\hat{\pi}_i}, \quad \frac{|\hat{\pi}_i - \pi_i|}{\pi_i} \leq \frac{\hat{b}}{\pi_i} \leq \frac{\hat{b}}{[\hat{\pi}_i - \hat{b}]_+}.$$

Therefore, using the fact that for any $x > 0$,

$$\max \left\{ |\sqrt{x} - 1|, |\sqrt{1/x} - 1| \right\} \leq \frac{1}{2} \max \{ |x - 1|, |1/x - 1| \}$$

we have for every $i \in [d]$,

$$\begin{aligned} \max \left\{ |\sqrt{\pi_i/\hat{\pi}_i} - 1|, |\sqrt{\hat{\pi}_i/\pi_i} - 1| \right\} &\leq \frac{1}{2} \max \{ |\pi_i/\hat{\pi}_i - 1|, |\hat{\pi}_i/\pi_i - 1| \} \\ &\leq \frac{1}{2} \max \left\{ \frac{\hat{b}}{\hat{\pi}_i}, \frac{\hat{b}}{[\hat{\pi}_i - \hat{b}]_+} \right\}. \end{aligned} \quad \square$$

C.4 Empirical bounds for L

By Weyl's inequality and the triangle inequality,

$$\max_{i \in [d]} |\lambda_i - \hat{\lambda}_i| \leq \|L - \text{Sym}(\hat{L})\| \leq \|L - \hat{L}\|.$$

It is easy to show that $|\hat{\gamma}_* - \gamma_*|$ is bounded by the same quantity. Therefore, it remains to establish an empirical bound on $\|L - \hat{L}\|$.

Lemma 7. *If $|\hat{P}_{i,j} - P_{i,j}| \leq \hat{B}_{i,j}$ for each $(i, j) \in [d]^2$ and $\max\{|\hat{\pi}_i - \pi_i| : i \in [d]\} \leq \hat{b}$, then*

$$\|\hat{L} - L\| \leq 2\hat{\rho} + \hat{\rho}^2 + (1 + 2\hat{\rho} + \hat{\rho}^2) \left(\sum_{(i,j) \in [d]^2} \frac{\hat{\pi}_i}{\hat{\pi}_j} \hat{B}_{i,j}^2 \right)^{1/2},$$

where

$$\hat{\rho} := \frac{1}{2} \max_{i \in [d]} \bigcup \left\{ \frac{\hat{b}}{\hat{\pi}_i}, \frac{\hat{b}}{[\hat{\pi}_i - \hat{b}]_+} \right\}.$$

Proof. We use the following decomposition of $L - \hat{L}$:

$$L - \hat{L} = \mathcal{E}_P + \mathcal{E}_{\pi,1} \hat{L} + \hat{L} \mathcal{E}_{\pi,2} + \mathcal{E}_{\pi,1} \mathcal{E}_P + \mathcal{E}_P \mathcal{E}_{\pi,2} + \mathcal{E}_{\pi,1} \hat{L} \mathcal{E}_{\pi,2} + \mathcal{E}_{\pi,1} \mathcal{E}_P \mathcal{E}_{\pi,2}$$

where

$$\begin{aligned} \mathcal{E}_P &:= \text{Diag}(\hat{\pi})^{1/2} (P - \hat{P}) \text{Diag}(\hat{\pi})^{-1/2}, \\ \mathcal{E}_{\pi,1} &:= \text{Diag}(\pi)^{1/2} \text{Diag}(\hat{\pi})^{-1/2} - I, \\ \mathcal{E}_{\pi,2} &:= \text{Diag}(\hat{\pi})^{1/2} \text{Diag}(\pi)^{-1/2} - I. \end{aligned}$$

Therefore

$$\begin{aligned} \|L - \hat{L}\| &\leq \|\mathcal{E}_{\pi,1}\| + \|\mathcal{E}_{\pi,2}\| + \|\mathcal{E}_{\pi,1}\| \|\mathcal{E}_{\pi,2}\| \\ &\quad + (1 + \|\mathcal{E}_{\pi,1}\| + \|\mathcal{E}_{\pi,2}\| + \|\mathcal{E}_{\pi,1}\| \|\mathcal{E}_{\pi,2}\|) \|\mathcal{E}_P\|. \end{aligned}$$

Observe that for each $(i, j) \in [d]^2$, the (i, j) -th entry of $\mathcal{E}_P$ is bounded in absolute value by

$$|(\mathcal{E}_P)_{i,j}| = \hat{\pi}_i^{1/2} \hat{\pi}_j^{-1/2} |P_{i,j} - \hat{P}_{i,j}| \leq \hat{\pi}_i^{1/2} \hat{\pi}_j^{-1/2} \hat{B}_{i,j}.$$

Since the spectral norm of $\mathcal{E}_P$ is bounded above by its Frobenius norm,

$$\|\mathcal{E}_P\| \leq \left(\sum_{(i,j) \in [d]^2} (\mathcal{E}_P)_{i,j}^2 \right)^{1/2} \leq \left(\sum_{(i,j) \in [d]^2} \frac{\pi_i}{\pi_j} \hat{B}_{i,j}^2 \right)^{1/2}.$$

Finally, the spectral norms of $\mathcal{E}_{\pi,1}$ and $\mathcal{E}_{\pi,2}$ satisfy

$$\max \{ \|\mathcal{E}_{\pi,1}\|, \|\mathcal{E}_{\pi,2}\| \} = \max_{i \in [d]} \bigcup \{ |\sqrt{\pi_i/\hat{\pi}_i} - 1|, |\sqrt{\hat{\pi}_i/\pi_i} - 1| \},$$

which can be bounded using Lemma 6. □

This establishes the validity of the confidence interval for γ_* in the same event from Lemma 3.

C.5 Asymptotic widths of intervals

Let us now turn to the asymptotic behavior of the interval widths (regarding $\hat{b}$, $\hat{\rho}$, and $\hat{w}$ all as functions of n).

A simple calculation gives that, almost surely, as $n \rightarrow \infty$,

$$\begin{aligned}\sqrt{\frac{n}{\log \log n}} \hat{b} &= O\left(\max_{i,j} \kappa \sqrt{\frac{P_{i,j}}{\pi_i}}\right), \\ \sqrt{\frac{n}{\log \log n}} \hat{\rho} &= O\left(\frac{\kappa}{\pi_\star^{3/2}}\right).\end{aligned}$$

Here, we use the fact that $\hat{\kappa} \rightarrow \kappa$ as $n \rightarrow \infty$ since $\hat{\mathbf{A}}^\# \rightarrow \mathbf{A}^\#$ as $\hat{\mathbf{P}} \rightarrow \mathbf{P}$ [14, 15].

Further, since

$$\sqrt{\frac{n}{\log \log n}} \left(\sum_{i,j} \frac{\hat{\pi}_i}{\hat{\pi}_j} \hat{B}_{i,j}^2 \right)^{1/2} = O\left(\left(\sum_{i,j} \frac{\pi_i}{\pi_j} \cdot \frac{P_{i,j}(1-P_{i,j})}{\pi_i} \right)^{1/2} \right) = O\left(\sqrt{\frac{d}{\pi_\star}} \right),$$

we thus have

$$\sqrt{\frac{n}{\log \log n}} \hat{w} = O\left(\frac{\kappa}{\pi_\star^{3/2}} + \sqrt{\frac{d}{\pi_\star}} \right).$$

This completes the proof of Theorem 4. $\square$

The following claim provides a bound on κ in terms of the number of states and the spectral gap.

Claim 1. $\kappa \leq d/\gamma_\star$.

Proof. It is established by Cho and Meyer [11] that

$$\kappa \leq \max_{i,j} |\mathbf{A}_{i,j}^\#| \leq \sup_{\|\mathbf{v}\|_1=1, \langle \mathbf{v}, \mathbf{1} \rangle=0} \|\mathbf{v}^\top \mathbf{A}^\#\|_1$$

(our κ is the κ_4 quantity from [11]), and Seneta [16] establishes

$$\sup_{\|\mathbf{v}\|_1=1, \langle \mathbf{v}, \mathbf{1} \rangle=0} \|\mathbf{v}^\top \mathbf{A}^\#\|_1 \leq \frac{d}{\gamma_\star}. \quad \square$$

D Proof of Theorem 5

Let $\hat{\pi}_{\star, \text{lb}}$ and $\hat{\gamma}_{\star, \text{lb}}$ be the lower bounds on $\pi_\star$ and $\gamma_\star$, respectively, computed from Algorithm 1. Let $\hat{\pi}_\star$ and $\hat{\gamma}_\star$ be the estimates of $\pi_\star$ and $\gamma_\star$ computed using the estimators from Theorem 3. By a union bound, we have by Theorem 3 and Theorem 4 that with probability at least $1 - 2\delta$,

$$|\hat{\pi}_\star - \pi_\star| \leq C \left(\sqrt{\frac{\pi_\star \log \frac{d}{\hat{\pi}_{\star, \text{lb}} \delta}}{\hat{\gamma}_{\star, \text{lb}} n}} + \frac{\log \frac{d}{\hat{\pi}_{\star, \text{lb}} \delta}}{\hat{\gamma}_{\star, \text{lb}} n} \right) \quad (33)$$

and

$$|\hat{\gamma}_\star - \gamma_\star| \leq C \left(\sqrt{\frac{\log \frac{d}{\delta} \cdot \log \frac{n}{\hat{\pi}_{\star, \text{lb}} \delta}}{\hat{\pi}_{\star, \text{lb}} \hat{\gamma}_{\star, \text{lb}} n}} + \frac{\log \frac{d}{\delta} \cdot \log \frac{n}{\hat{\pi}_{\star, \text{lb}} \delta}}{\hat{\pi}_{\star, \text{lb}} \hat{\gamma}_{\star, \text{lb}} n} + \frac{\log \frac{1}{\hat{\gamma}_{\star, \text{lb}} \delta}}{\hat{\gamma}_{\star, \text{lb}} n} \right). \quad (34)$$

The bound on $|\hat{\gamma}_\star - \gamma_\star|$ in Eq. (34)—call it $\hat{w}'$ —is fully observable and hence yields a confidence interval for $\gamma_\star$. The bound on $|\hat{\pi}_\star - \pi_\star|$ in Eq. (33) depends on $\pi_\star$, but from it one can derive

$$|\hat{\pi}_\star - \pi_\star| \leq C' \left(\sqrt{\frac{\hat{\pi}_\star \log \frac{d}{\hat{\pi}_{\star, \text{lb}} \delta}}{\hat{\gamma}_{\star, \text{lb}} n}} + \frac{\log \frac{d}{\hat{\pi}_{\star, \text{lb}} \delta}}{\hat{\gamma}_{\star, \text{lb}} n} \right)$$

using the approach from the proof of Lemma 4. Here, $C' > 0$ is an absolute constant that depends only on C . This bound—call it $\hat{b}'$ —is now also fully observable. We have established that in the $1 - 2\delta$ probability event from above,

$$\pi_\star \in \widehat{U} := [\hat{\pi}_\star - \hat{b}', \hat{\pi}_\star + \hat{b}'], \quad \gamma_\star \in \widehat{V} := [\hat{\gamma}_\star - \hat{w}', \hat{\gamma}_\star + \hat{w}'].$$

It is easy to see that almost surely (as $n \rightarrow \infty$),

$$\sqrt{\frac{n}{\log n}} \hat{w}' = O\left(\sqrt{\frac{\log(d/\delta)}{\pi_\star \gamma_\star}}\right)$$

and

$$\sqrt{n} \hat{b}' = O\left(\sqrt{\frac{\pi_\star \log \frac{d}{\pi_\star \delta}}{\gamma_\star}}\right).$$

This completes the proof of Theorem 5. □

References

- [1] D. Paulin. Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electronic Journal of Probability*, 20:1–32, 2015.
- [2] O. Bousquet, S. Boucheron, and G. Lugosi. Introduction to statistical learning theory. *Lecture Notes in Artificial Intelligence*, 3176:169–207, 2004.
- [3] D. Levin, Y. Peres, and E. Wilmer. *Markov Chains and Mixing Times*. AMS, 2008.
- [4] R. Montenegro and P. Tetali. *Mathematical Aspects of Mixing Times in Markov Chains*. Now Publishers, 2006.
- [5] S. Bernstein. Sur l’extension du theoreme limite du calcul des probabilites aux sommes de quantites dependantes. *Mathematische Annalen*, 97:1–59, 1927.
- [6] J. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 2015.
- [7] B. Yu. Rates of convergence for empirical processes of stationary mixing sequences. *The Annals of Probability*, 22(1):94–116, January 1994.
- [8] D. Freedman. On tail probabilities for martingales. *The Annals of Probability*, 3(1):100–118, 1975.
- [9] J.-Y. Audibert, R. Munos, and C. Szepesvári. Exploration-exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19):1876–1902, 2009.
- [10] C. D. Meyer Jr. The role of the group generalized inverse in the theory of finite Markov chains. *SIAM Review*, 17(3):443–464, 1975.
- [11] G. Cho and C. Meyer. Comparison of perturbation bounds for the stationary distribution of a Markov chain. *Linear Algebra and its Applications*, 335:137–150, 2001.
- [12] M. Haviv and L. Van der Heyden. Perturbation bounds for the stationary probabilities of a finite Markov chain. *Advances in Applied Probability*, 16:804–818, 1984.
- [13] S. Kirkland, M. Neumann, and B. Shader. Applications of Paz’s inequality to perturbation bounds for Markov chains. *Linear Algebra and its Applications*, 268:183–196, 1998.
- [14] X. Li and Y. Wei. An improvement on the perturbation of the group inverse and oblique projection. *Linear Algebra and its Applications*, 338:53–66, 2001.
- [15] J. Benítez and X. Liu. On the continuity of the group inverse. *Operators and Matrices*, 6(4):859–868, 2012.
- [16] E. Seneta. Sensitivity of finite Markov chains under perturbation. *Statistics & Probability Letters*, 17:163–168, 1993.