MIXING TIME ESTIMATION IN REVERSIBLE MARKOV CHAINS FROM A SINGLE SAMPLE PATH

DANIEL HSU, ARYEH KONTOROVICH, DAVID A. LEVIN, YUVAL PERES, AND CSABA SZEPESVÁRI

Abstract. The spectral gap $\gamma_\star$ of a finite, ergodic, and reversible Markov chain is an important parameter measuring the asymptotic rate of convergence. In applications, the transition matrix $\boldsymbol{P}$ may be unknown, yet one sample of the chain up to a fixed time n may be observed. We consider here the problem of estimating $\gamma_\star$ from this data. Let π be the stationary distribution of $\boldsymbol{P}$, and $\pi_\star = \min_x \pi(x)$. We show that if $n = \tilde{O}\big(\frac{1}{\gamma_\star \pi_\star}\big)$, then γ can be estimated to within multiplicative constants with high probability. When π is uniform on d states, this matches (up to logarithmic correction) a lower bound of $\tilde{\Omega}\big(\frac{d}{\gamma_\star}\big)$ steps required for precise estimation of $\gamma_\star$. Moreover, we provide the first procedure for computing a fully data-dependent interval, from a single finite-length trajectory of the chain, that traps the mixing time t_{mix} of the chain at a prescribed confidence level. The interval 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 $t_{\mathrm{relax}} = 1/\gamma_\star$, which is strongly related to the mixing time, and the width of the interval converges to zero roughly at a $1/\sqrt{n}$ rate, where n is the length of the sample path.

1. Introduction

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

$$\lim_{t\to\infty} \Pr_{\boldsymbol{q}}(X_t = i) = \lim_{t\to\infty} (\boldsymbol{q}\boldsymbol{P}^t)_i = \pi_i \quad \text{for each } i \in [d].$$

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:

$$t_{\mathrm{mix}} := \min\left\{ t \in \mathbb{N} : \sup_{\boldsymbol{q}} \max_{A\subset[d]} \left|\Pr_{\boldsymbol{q}}(X_t \in A) - \boldsymbol{\pi}(A)\right| \le 1/4 \right\}. \tag{1}$$

Here, $\boldsymbol{\pi}(A) = \sum_{i\in A} \pi_i$ is the probability assigned to set A by $\boldsymbol{\pi}$, and the supremum is over all possible initial distributions $\boldsymbol{q}$. The problem studied in this work is the construction of a non-trivial confidence interval $C_n = C_n(X_1, X_2, \ldots, X_n, \delta) \subset [0, \infty]$, based only on the observed sample path $(X_1, X_2, \ldots, X_n)$ 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 $\boldsymbol{\pi}(f) = \sum_i \pi_i f(i)$ for some function f of the states of a Markov chain. This is the setting of the celebrated Markov Chain Monte Carlo (MCMC) paradigm (J. S. Liu 2001), but the problem also arises in performance prediction involving time-correlated data, as is common in reinforcement learning (Sutton and Barto 1998). Observable, or *a posteriori* 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.

1.1. **Main results.** Consider a reversible ergodic Markov chain on d states with absolute spectral gap $\gamma_\star$ and stationary distribution minorized by $\pi_\star$. As is well-known (see, for example, Levin, Peres, and Wilmer (2009, Theorems 12.3 and 12.4)),

$$(t_{\text{relax}} - 1)\ln 2 \;\le\; t_{\text{mix}} \;\le\; t_{\text{relax}} \ln \frac{4}{\pi_\star} \tag{2}$$

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

(1) In Section 3.1, we show that in some problems $n = \Omega((d \log d)/\gamma_\star + 1/\pi_\star)$ observations are necessary for any procedure to guarantee constant multiplicative accuracy in estimating $\gamma_\star$ (Theorems 3.1 and 3.2). Essentially, in some problems *every* state may need to be visited about $\log(d)/\gamma_\star$ 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 $\hat{\gamma}_\star$ for $\gamma_\star$, based an *a single sample path*, and prove in Theorem 3.4 that $|\frac{\hat{\gamma}_\star}{\gamma_\star} - 1| < \varepsilon$ with high probability if the path is of length $\tilde{O}(1/(\pi_\star\gamma_\star\varepsilon^2))$. (The $\tilde{O}(\cdot)$ notation suppresses logarithmic factors.) We also provide and analyze a point estimator for $\pi_\star$. This establishes the feasibility of *estimating* the mixing time in this setting, and the dependence on $\pi_\star$ and $\gamma_\star$ in the path length matches our lower bound (up to logarithmic factors) in the case where $1/\pi_\star = \Omega(d)$. We note, however, that these results give only *a priori* confidence intervals that depend on the unknown quantities $\pi_\star$ and $\gamma_\star$. As such, the results do not lead to a universal (chain-independent) stopping rule for stopping the chain when the relative error is below the prescribed accuracy.
(3) In Section 4, we propose a procedure for *a posteriori* constructing confidence intervals for $\pi_\star$ and $\gamma_\star$ that depend only on the observed sample path and not on any unknown parameters. We prove that the intervals shrink at a $\tilde{O}(1/\sqrt{n})$ rate (Theorems 4.1 and 4.2). These confidence intervals trivially lead to a universal stopping rule to stop the chain when a prescribed relative error is achieved.

1.2. **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 $\boldsymbol{\pi}_n(f) := \frac{1}{n}\sum_{t=1}^n f(X_t)$ converges almost surely to $\boldsymbol{\pi}(f)$ (Meyn and Tweedie 1993), while the central limit theorem tells us that as $n \to \infty$, the distribution of the deviation $\sqrt{n}(\boldsymbol{\pi}_n(f) - \boldsymbol{\pi}(f))$ will be normal with mean zero and asymptotic variance $\lim_{n\to\infty} n \operatorname{Var}(\boldsymbol{\pi}_n(f))$ (Kipnis and Varadhan 1986).

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 (Kontoyiannis, Lastras-Montaño, and Meyn 2006; Flegal and Jones 2011; Gyori and Paulin 2014). To address this need, numerous works have developed Chernoff-type bounds on $\Pr(|\boldsymbol{\pi}_n(f) - \boldsymbol{\pi}(f)| > \epsilon)$, thus providing valuable tools for non-asymptotic probabilistic analysis (Gillman 1998; León and Perron 2004; Kontoyiannis, Lastras-Montaño, and Meyn 2006; Kontorovich and Weiss 2014; Paulin 2015). These probability bounds are larger than the 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 $\Theta(t_{\mathrm{mix}})$ time steps. Note that the bounds generally depend on distribution-specific properties of the Markov chain (e.g., $\boldsymbol{P}$, t_{mix}, $\gamma_\star$), 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., Garren and R. L. Smith 2000; Jones and Hobert 2001; Flegal and Jones 2011; Atchadé 2016; Gyori and Paulin 2014). 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. In particular, they also fail to provide universal stopping rules that allow the estimation of (for example) the mixing time with a fixed relative accuracy.

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 data (e.g., Yu 1994; Karandikar and Vidyasagar 2002; Gamarnik 2003; Mohri and Rostamizadeh 2008; Steinwart and Christmann 2009; Steinwart, Hush, and Scovel 2009), 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 (McDonald, Shalizi, and Schervish 2011). However, the convergence rates of the estimates from McDonald, Shalizi, and Schervish (2011), 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 by Batu, Fortnow, Rubinfeld, W. D. Smith, and White (2000), Batu, Fortnow, Rubinfeld, W. D. Smith, and White (2013), and Bhattacharya and Valiant (2015). Note that in this setting, Bhattacharya and Valiant (2015) asked if one could approximate t_{mix} (up to logarithmic factors) with a number of queries that is linear in both d and t_{mix}; our work answers the question affirmatively (up to logarithmic corrections) in the case when the stationary distribution is near uniform. Finally, 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 (Bhatnagar, Bogdanov, and Mossel 2011).

This paper is based on the conference paper of Hsu, Kontorovich, and Szepesvári (2015), combined with the results in the unpublished manuscript of Levin and Peres (2016).

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 > 1. Boldface symbols are used for vectors and matrices (e.g., $\boldsymbol{v}$, $\boldsymbol{M}$), and their entries are referenced by subindexing (e.g., v_i, $M_{i,j}$). For a vector $\boldsymbol{v}$, $\|\boldsymbol{v}\|$ denotes its Euclidean norm; for a matrix $\boldsymbol{M}$, $\|\boldsymbol{M}\|$ denotes its spectral norm. We use $\operatorname{Diag}(\boldsymbol{v})$ to denote the diagonal matrix whose (i,i)-th entry is v_i. The probability simplex is denoted by $\Delta^{d-1} = \{\boldsymbol{p} \in [0,1]^d : \sum_{i=1}^d p_i = 1\}$, and we regard vectors in Δ^{d-1} as row vectors.

2.2. **Setting.** Let $\boldsymbol{P} \in (\Delta^{d-1})^d \subset [0,1]^{d \times d}$ 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 $i \in [d]$ (Levin, Peres, and Wilmer 2009, Corollary 1.17). We also assume that $\boldsymbol{P}$ is *reversible* (with respect to $\boldsymbol{\pi}$):

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

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

Define the matrices

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

The (i,j)th entry of the matrix $M_{i,j}$ contains the *doublet probabilities* associated with $\boldsymbol{P}$: $M_{i,j} = \pi_i P_{i,j}$ is the probability of seeing state i followed by state j when the chain is started from its stationary distribution. The matrix $\boldsymbol{M}$ is symmetric on account of the reversibility of $\boldsymbol{P}$, and hence it follows that $\boldsymbol{L}$ is also symmetric. (We will strongly exploit the symmetry in our results.) Further, $\boldsymbol{L} = \operatorname{Diag}(\boldsymbol{\pi})^{1/2}\boldsymbol{P}\operatorname{Diag}(\boldsymbol{\pi})^{-1/2}$, hence $\boldsymbol{L}$ and $\boldsymbol{P}$ are similar and thus their eigenvalue systems are identical. Ergodicity and reversibility imply that the eigenvalues of $\boldsymbol{L}$ are contained in the interval $(-1, 1]$, and that 1 is an eigenvalue of $\boldsymbol{L}$ with multiplicity 1 (Levin, Peres, and Wilmer 2009, Lemmas 12.1 and 12.2). Denote and order the eigenvalues of $\boldsymbol{L}$ 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 $\boldsymbol{P}$. For each $t \in \mathbb{N}$, let $\boldsymbol{\pi}^{(t)} \in \Delta^{d-1}$ denote the marginal distribution of X_t, so

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

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 confidence intervals that $\pi_\star$ and $\gamma_\star$ with high probability; in particular, the construction of the intervals should be fully empirical and not depend on any unobservable quantities, including $\pi_\star$ and $\gamma_\star$ 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 $\pi_\star$ and $\gamma_\star$, 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 $\pi_\star$ and $\gamma_\star$ (as we will show in the proof of our first result). Therefore, valid confidence intervals for $\pi_\star$ and $\gamma_\star$ 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 $\gamma_\star$. By Eq. (2), the multiplicative accuracy lower bound for $\gamma_\star$ 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 $\Omega(1/\pi_\star)$ is necessary to achieve even a constant *additive* accuracy in estimating $\gamma_\star$.

Theorem 3.1. *Pick any $\bar{\pi} \in (0, 1/4)$. Consider any estimator $\hat{\gamma}_\star$ that takes as input a random sample path of length $n \le 1/(4\bar{\pi})$ 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 $\gamma_\star \ge 1/2$ and minimum stationary probability $\pi_\star \ge \bar{\pi}$ such that*

$$\Pr\left[|\hat{\gamma}_\star - \gamma_\star| \ge 1/8\right] \ge 3/8.$$

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

Theorem 3.2. *There is an absolute constant $c > 0$ such that the following holds. Pick any positive integer $d \ge 3$ and any $\bar{\gamma}_\star \in (0, 1/2)$. Consider any estimator $\hat{\gamma}_\star$ that takes as input a random sample path of length $n < cd\log(d)/\bar{\gamma}_\star$ 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 $\gamma_\star \in [\bar{\gamma}_\star, 2\bar{\gamma}_\star]$ and minimum stationary probability $\pi_\star \ge 1/(2d)$ such that*

$$\Pr\left[|\hat{\gamma}_\star - \gamma_\star| \ge \bar{\gamma}_\star/2\right] \ge 1/4.$$

The proofs of Theorems 3.1 and 3.2 are given in Section 5.

3.2. **A plug-in based point estimator and its accuracy.** Let us now consider the problem of estimating $\gamma_\star$. 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 $\widehat{\boldsymbol{M}} \in [0,1]^{d\times d}$ and random vector $\hat{\boldsymbol{\pi}} \in \Delta^{d-1}$ by

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

Furthermore, define

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

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

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

Let $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \cdots \geq \hat{\lambda}_d$ be the eigenvalues of $\operatorname{Sym}(\widehat{\boldsymbol{L}})$. Our estimator of the minimum stationary probability $\pi_\star$ is $\hat{\pi}_\star := \min_{i \in [d]} \hat{\pi}_i$, and our estimator of the spectral gap $\gamma_\star$ is $\hat{\gamma}_\star := 1 - \min\{1, \max\{\hat{\lambda}_2, |\hat{\lambda}_d|\}\} \in [0,1]$. The astute reader may notice that our estimator is ill-defined when $\hat{\boldsymbol{\pi}}$ is not positive valued. In this case, we can simply set $\hat{\gamma}_\star = 0$.

These estimators have the following accuracy guarantees:

Theorem 3.3. *There exists an absolute constant $C \geq 1$ such that the following holds. Let $(X_t)_{t=1}^n$ be an ergodic and reversible Markov chain with spectral gap $\gamma_\star$ and minimum stationary probability $\pi_\star > 0$. Let $\hat{\pi}_\star = \hat{\pi}_\star((X_t)_{t=1}^n)$ and $\hat{\gamma}_\star = \hat{\gamma}_\star((X_t)_{t=1}^n)$ be the estimators described above. For any $\delta \in (0,1)$, with probability at least $1-\delta$,*

$$|\hat{\pi}_\star - \pi_\star| \leq C \left(\sqrt{\frac{\pi_\star \log \frac{1}{\pi_\star \delta}}{\gamma_\star n}} + \frac{\log \frac{1}{\pi_\star \delta}}{\gamma_\star n} \right) \tag{4}$$

and

$$|\hat{\gamma}_\star - \gamma_\star| \leq C \sqrt{\frac{\log \frac{d}{\delta} \cdot \log \frac{n}{\pi_\star \delta}}{\pi_\star \gamma_\star n}}. \tag{5}$$

Theorem 3.3 implies that the sequence lengths sufficient to estimate $\pi_\star$ and $\gamma_\star$ to within constant multiplicative factors are, respectively,

$$\tilde{O}\left(\frac{1}{\pi_\star \gamma_\star}\right) \quad \text{and} \quad \tilde{O}\left(\frac{1}{\pi_\star \gamma_\star^3}\right).$$

The proof of Theorem 3.3 is based on analyzing the convergence of the sample averages $\widehat{\boldsymbol{M}}$ and $\hat{\boldsymbol{\pi}}$ to their expectation, and then using perturbation bounds for eigenvalues to derive a bound on the error of $\hat{\gamma}_\star$. 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 $\widehat{\boldsymbol{M}}$ 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 (Tropp 2015) to our non-iid setting by directly applying a blocking technique of Bernstein (1927) as described in the article of Yu (1994). Due to ergodicity, the convergence rate can be bounded without any dependence on the initial state distribution $\boldsymbol{\pi}^{(1)}$. The proof of Theorem 3.3 is given in Section 6.

3.3. **Improving the plug-in estimator.** We can bootstrap the plug-in estimator in Eq. (5) to show that in fact, to obtain any prescribed multiplicative accuracy, $\tilde{O}(1/(\pi_\star \gamma_\star))$ steps suffice to estimate $\gamma_\star$. The idea is to apply the estimator $\hat{\gamma}_\star$ from Eq. (5) to the *a-skipped chain* $(X_{as})_{s=1}^{n/a}$ for some $a \geq 1$. This chain has spectral gap $\gamma_\star(a) = 1 - (1-\gamma_\star)^a$. Thus, letting $\hat{\gamma}_\star(a)$ be the plug-in estimator for $\gamma_\star(a)$ based on the a-skipped chain, a natural estimator of $\gamma_\star$ is $1 - (1 - \hat{\gamma}_\star(a))^{1/a}$.

Why may this improve on the original plug-in estimator from Section 3.2? Observe that $\gamma_\star(a) = \Omega(\gamma_\star a)$ for $a \leq 1/\gamma_\star$, so the additive accuracy bound from

Eq. (5) for the plug-in estimator on $(X_{as})_{s=1}^{n/a}$ is roughly the same for all $a \leq 1/\gamma_\star$. However, when $\gamma_\star(a)$ is bounded away from 0 and 1, a small additive error in estimating $\gamma_\star(a)$ with $\hat{\gamma}_\star(a)$ translates to a small multiplicative error in estimating $\gamma_\star$ using $1-(1-\hat{\gamma}_\star(a))^{1/a}$. So it suffices to use the skipped chain estimator with some $a = O(1/\gamma_\star)$. Since $\gamma_\star$ is not known (of course), we use a doubling trick to find a suitable value of a.

The estimator is defined as follow. For simplicity, assume n is a power of two. Initially, set $k := 0$. Let $a := 2^k$ and $\hat{\gamma}_\star(a) := \hat{\gamma}_\star((X_{as})_{s=1}^{n/a})$. If $\hat{\gamma}_\star(a) > 0.31$ or $a = n$, then set $A := a$ and return $\tilde{\gamma}_\star := 1-(1-\hat{\gamma}_\star(A))^{1/A}$. Otherwise, increment k by one and repeat.

Theorem 3.4. *There exists a polynomial function $\mathcal{L}$ of the logarithms of $\gamma_\star^{-1}$, $\pi_\star^{-1}$, δ^{-1}, and d such that the following holds. Let $(X_t)_{t=1}^n$ be an ergodic and reversible Markov chain with spectral gap $\gamma_\star$ and minimum stationary probability $\pi_\star > 0$. Let $\tilde{\gamma}_\star = \tilde{\gamma}_\star((X_t)_{t=1}^n)$ be the estimator defined above. For any $\varepsilon, \delta \in (0,1)$, if $n \geq \mathcal{L}/(\pi_\star\gamma_\star\varepsilon^2)$, then with probability at least $1-\delta$,*

$$\left|\frac{\tilde{\gamma}_\star}{\gamma_\star} - 1\right| \leq \varepsilon.$$

The definition of $\mathcal{L}$ is in Eq. (34). The proof of Theorem 3.4 is given in Section 7. The result shows that to estimate *both* $\pi_\star$ and $\gamma_\star$ to within constant multiplicative factors, a single sequence of length $\tilde{O}(1/(\pi_\star\gamma_\star))$ suffices.

4. A posteriori confidence intervals

In this section, we describe and analyze a procedure for constructing confidence intervals for the stationary probabilities and the spectral gap $\gamma_\star$.

4.1. Procedure.

We first note that the point estimators from Theorem 3.3 and Theorem 3.4 fall short of being directly suitable for obtaining a fully empirical, a posteriori confidence interval for $\gamma_\star$ and $\pi_\star$. This is because the deviation terms themselves depend inversely both on $\gamma_\star$ and $\pi_\star$, and hence can never rule out 0 (or an arbitrarily small positive value) as a possibility for $\gamma_\star$ or $\pi_\star$.[1] In effect, the fact that the Markov chain could be slow mixing and the long-term frequency of some states could be small makes it difficult to be confident in the estimates provided by $\hat{\gamma}_\star$ and $\hat{\pi}_\star$.

The main idea behind our procedure, given as Algorithm 1, is to use the Markov property to eliminate the dependence of the confidence intervals on the unknown quantities (including $\pi_\star$ and $\gamma_\star$). Specifically, we estimate the transition probabilities from the sample path using simple state visit counts: 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.

With confidence intervals for the entries of $\boldsymbol{P}$ in hand, it is possible to form a confidence interval for $\gamma_\star$ based on the eigenvalues of an estimated transition probability matrix by appealing to the Ostrowski-Elsner theorem (cf. Theorem 1.4 on Page 170 of Stewart and Sun (1990).) However, directly using this perturbation

[1] Using Theorem 3.3, it is possible to trap $\gamma_\star$ in the union of *two* empirical confidence intervals—one around $\hat{\gamma}_\star$ and the other around zero, both of which shrink in width as the sequence length increases.

Algorithm 1 Confidence intervals

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

1: Compute state visit counts and smoothed transition probability estimates:
$$N_i := |\{t \in [n-1] : X_t = i\}|, \quad i \in [d];$$
$$N_{i,j} := |\{t \in [n-1] : (X_t, X_{t+1}) = (i,j)\}|, \quad (i,j) \in [d]^2;$$
$$\widehat{P}_{i,j} := \frac{N_{i,j} + 1/d}{N_i + 1}, \quad (i,j) \in [d]^2.$$

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

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

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

5: Spectral gap estimate:
$$\hat{\gamma}_\star := 1 - \max\{\hat{\lambda}_2, |\hat{\lambda}_d|\}.$$

6: Bounds for $|\widehat{P}_{i,j} - P_{i,j}|$ for $(i,j) \in [d]^2$: $c := 1.1$, $\tau_{n,\delta} := \inf\{t \geq 0 : 2d^2(1 + \lceil \log_c \frac{2n}{t} \rceil_+) e^{-t} \leq \delta\}$, and
$$\widehat{B}_{i,j} := \left(\sqrt{\frac{c\tau_{n,\delta}}{2N_i}} + \sqrt{\frac{c\tau_{n,\delta}}{2N_i} + \sqrt{\frac{2c\widehat{P}_{i,j}(1-\widehat{P}_{i,j})\tau_{n,\delta}}{N_i}} + \frac{\frac{4}{3}\tau_{n,\delta} + |\widehat{P}_{i,j} - \frac{1}{d}|}{N_i}} \right)^2.$$

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

8: Bounds for $\max_{i \in [d]} |\hat{\pi}_i - \pi_i|$ and $\max \bigcup_{i \in [d]} \{ |\sqrt{\pi_i / \hat{\pi}_i} - 1|, |\sqrt{\hat{\pi}_i / \pi_i} - 1| \}$:
$$\hat{b} := \hat{\kappa} \max\left\{ \widehat{B}_{i,j} : (i,j) \in [d]^2 \right\}, \qquad \hat{\rho} := \frac{1}{2} \max \bigcup_{i \in [d]} \left\{ \frac{\hat{b}}{\hat{\pi}_i}, \frac{\hat{b}}{[\hat{\pi}_i - \hat{b}]_+} \right\}.$$

9: Bounds for $|\hat{\gamma}_\star - \gamma_\star|$:
$$\hat{w} := 2\hat{\rho} + \hat{\rho}^2 + (1 + 2\hat{\rho} + \hat{\rho}^2) \left(\sum_{(i,j) \in [d]^2} \frac{\hat{\pi}_i}{\hat{\pi}_j} \widehat{B}_{i,j}^2 \right)^{1/2}.$$

---

result leads to very wide intervals, shrinking only at a rate of $O(n^{-1/(2d)})$. We avoid this slow rate by constructing confidence intervals for the symmetric matrix $\boldsymbol{L}$, so that we can use a stronger perturbation result (namely Weyl's inequality, as in the proof of Theorem 3.3) available for symmetric matrices.

To form an estimate of $\boldsymbol{L}$ based on an estimate of the transition probabilities, one possibility is to estimate $\boldsymbol{\pi}$ using state visit counts as was done in Section 3, and appeal to the relation $\boldsymbol{L} = \operatorname{Diag}(\boldsymbol{\pi})^{1/2} \boldsymbol{P} \operatorname{Diag}(\boldsymbol{\pi})^{-1/2}$ to form a plug-in estimate of $\boldsymbol{L}$. However, it is not clear how to construct a confidence interval for the entries of $\boldsymbol{\pi}$ because the accuracy of this estimator depends on the unknown mixing time.

We adopt a different strategy for estimating $\boldsymbol{\pi}$. We form the matrix $\widehat{\boldsymbol{P}}$ using smoothed frequency estimates of $\boldsymbol{P}$ (Step 1), then compute the *group inverse* $\widehat{\boldsymbol{A}}^{\#}$ of $\widehat{\boldsymbol{A}} = \boldsymbol{I} - \widehat{\boldsymbol{P}}$ (Step 2), followed by finding the unique stationary distribution $\hat{\boldsymbol{\pi}}$ of $\widehat{\boldsymbol{P}}$ (Step 3), this way decoupling the bound on the accuracy of $\hat{\boldsymbol{\pi}}$ from the mixing time. The group inverse $\widehat{\boldsymbol{A}}^{\#}$ of $\widehat{\boldsymbol{A}}$ is uniquely defined; and if $\widehat{\boldsymbol{P}}$ defines an ergodic chain (which is the case here due to the use of the smoothed estimates), $\widehat{\boldsymbol{A}}^{\#}$ can be computed at the cost of inverting an $(d-1)\times(d-1)$ matrix (Meyer Jr. 1975, Theorem 5.2).[2] Further, given $\widehat{\boldsymbol{A}}^{\#}$, the unique stationary distribution $\hat{\boldsymbol{\pi}}$ of $\widehat{\boldsymbol{P}}$ can be read out from the last row of $\widehat{\boldsymbol{A}}^{\#}$ (Meyer Jr. 1975, Theorem 5.3). The group inverse is also used to determine the relative sensitivity of $\hat{\boldsymbol{\pi}}$ to $\widehat{\boldsymbol{P}}$, which is quantified by

$$\hat{\kappa} := \frac{1}{2} \max \left\{ \widehat{A}^{\#}_{j,j} - \min \left\{ \widehat{A}^{\#}_{i,j} : i \in [d] \right\} : j \in [d] \right\} . \tag{6}$$

We can regard $\hat{\kappa}$ as a plug-in estimator for κ, which is defined by substituting the group inverse $\boldsymbol{A}^{\#}$ of $\boldsymbol{A}$ in for $\widehat{\boldsymbol{A}}^{\#}$ in Eq. (6).

We can now follow the strategy based on estimating $\boldsymbol{L}$ alluded to above. Using $\hat{\boldsymbol{\pi}}$ and $\widehat{\boldsymbol{P}}$, we construct the plug-in estimate $\widehat{\boldsymbol{L}}$ of $\boldsymbol{L}$, and use the eigenvalues of its symmetrization to form the estimate $\hat{\gamma}_\star$ of the spectral gap (Steps 4 and 5). In the remaining steps, we use matrix perturbation analyses to relate $\hat{\boldsymbol{\pi}}$ and $\boldsymbol{\pi}$, viewing $\boldsymbol{P}$ as the perturbation of $\widehat{\boldsymbol{P}}$; and also to relate $\hat{\gamma}_\star$ and $\gamma_\star$, viewing $\boldsymbol{L}$ as a perturbation of $\mathrm{Sym}(\widehat{\boldsymbol{L}})$. Both analyses give error bounds entirely in terms of observable quantities (e.g., $\hat{\kappa}$), tracing back to empirical error bounds for the estimate of $\boldsymbol{P}$.

The most computationally expensive step in Algorithm 1 is the computation of the group inverse $\widehat{\boldsymbol{A}}^{\#}$, which, as noted earlier, 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 $O(d^2)$.

4.2. **Main result.** We now state our main theorems. Below, the big-O notation should be interpreted as follows. For a random sequence $(Y_n)_{n\geq 1}$ and a (non-random) positive sequence $(\varepsilon_{\theta,n})_{n\geq 1}$ parameterized by θ, we say "$Y_n = O(\varepsilon_{\theta,n})$ holds almost surely as $n \to \infty$" if there is some universal constant $C > 0$ such that for all θ, $\limsup_{n\to\infty} Y_n/\varepsilon_{\theta,n} \leq C$ holds almost surely.

Theorem 4.1. *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 $\gamma_\star > 0$ denote the spectral gap, $\boldsymbol{\pi}$ the unique stationary distribution, and $\pi_\star > 0$ the minimum stationary probability. Then, on an event of probability at least $1-\delta$,*

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

Moreover,

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

almost surely as $n \to \infty$.

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

The proof of Theorem 4.1 is given in Section 8. As mentioned above, the obstacle encountered in Theorem 3.3 is avoided by exploiting the Markov property. We establish fully observable upper and lower bounds on the entries of $\boldsymbol{P}$ that converge at a $\sqrt{(\log\log n)/n}$ rate using standard martingale tail inequalities; this justifies the validity of the bounds from Step 6. Properties of the group inverse (Meyer Jr. 1975; Cho and Meyer 2001) and eigenvalue perturbation theory (Stewart and Sun 1990) are used to validate the empirical bounds on π_i and $\gamma_\star$ developed in the remaining steps of the algorithm.

The first part of Theorem 4.1 provides valid empirical confidence intervals for each π_i and for $\gamma_\star$, which are simultaneously valid at confidence level δ. The second part of Theorem 4.1 shows that the width of the intervals decrease as the sequence length increases. The rate at which the widths shrink is given in terms of $\boldsymbol{P}$, $\boldsymbol{\pi}$, κ, and n. We show in Section 8.5 (Lemma 8.8) that

$$\kappa \leq \frac{1}{\gamma_\star}\min\{d, 8+\log(4/\pi_\star)\},$$

and hence

$$\hat{b} = O\left(\max_{(i,j)\in[d]^2} \frac{\min\{d,\log(1/\pi_\star)\}}{\gamma_\star}\sqrt{\frac{P_{i,j}\log\log n}{\pi_i n}}\right),$$

$$\hat{w} = O\left(\frac{\min\{d,\log(1/\pi_\star)\}}{\pi_\star\gamma_\star}\sqrt{\frac{\log\log n}{\pi_\star n}}\right).$$

It is easy to combine Theorems 3.3 and 4.1 to yield intervals whose widths shrink at least as fast as both the non-empirical intervals from Theorem 3.3 and the empirical intervals from Theorem 4.1. Specifically, determine lower bounds on $\pi_\star$ and $\gamma_\star$ using Algorithm 1, $\pi_\star \geq \min_{i\in[d]}[\hat{\pi}_i - \hat{b}]_+$, $\gamma_\star \geq [\hat{\gamma}_\star - \hat{w}]_+$; then plug-in these lower bounds for $\pi_\star$ and $\gamma_\star$ in the deviation bounds in Eq. (5) from Theorem 3.3. This yields a new interval centered around the estimate of $\gamma_\star$ from Theorem 3.3 and the new interval no longer depends on unknown quantities. The interval is a valid $1-2\delta$ probability confidence interval for $\gamma_\star$, and for sufficiently large n, the width shrinks at the rate given in Eq. (5). We can similarly construct an empirical confidence interval for $\pi_\star$ using Eq. (4), which is valid on the same $1-2\delta$ probability event.[3] 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 Section 9.

Theorem 4.2. *The following holds under the same conditions as Theorem 4.1. For any $\delta \in (0,1)$, the confidence intervals $\widehat{U}$ and $\widehat{V}$ described above for $\pi_\star$ and $\gamma_\star$, respectively, satisfy $\pi_\star \in \widehat{U}$ and $\gamma_\star \in \widehat{V}$ with probability at least $1-2\delta$. Furthermore, $|\widehat{U}| = O\left(\sqrt{\frac{\pi_\star \log\frac{d}{\pi_\star\delta}}{\gamma_\star n}}\right)$ and $|\widehat{V}| = O\left(\min\left\{\sqrt{\frac{\log\frac{d}{\delta}\cdot\log(n)}{\pi_\star\gamma_\star n}}, \hat{w}\right\}\right)$ almost surely as $n\to\infty$, where $\hat{w}$ is the width from Algorithm 1.*

Finally, note that a stopping rule that stops when $\gamma_\star$ and $\pi_\star$ are estimated with a given relative error ϵ can be obtained as follows. At time n:

[3]For the $\pi_\star$ interval, we only plug-in lower bounds on $\pi_\star$ and $\gamma_\star$ only where these quantities appear as $1/\pi_\star$ and $1/\gamma_\star$ in Eq. (4). It is then possible to "solve" for observable bounds on $\pi_\star$. See Section 9 for details.

if $n = 2^k$ for an integer k **then**
  Run Algorithm 1 (or the improved variant from Theorem 4.2) with inputs $(X_1, X_2, \dots, X_n)$ and $\delta/(k(k+1))$ to obtain intervals for $\pi_\star$ and $\gamma_\star$.
  Stop if, for each interval, the interval width divided by the lower bound on estimated quantity falls below ϵ.
end if

It is easy to see then that with probability $1-\delta$, the algorithm only stops when the relative accuracy of its estimate is at least ϵ. Combined with the lower bounds, we conjecture that the expected stopping time of the resulting procedure is optimal up to log factors.

5. Proofs of Theorems 3.1 and 3.2

In this section, we prove Theorem 3.1 and Theorem 3.2.

5.1. Proof of Theorem 3.1. Fix $\bar{\pi} \in (0, 1/4)$. Consider two Markov chains given by the following stochastic matrices:

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

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

In order to guarantee $|\hat{\gamma}_\star - \gamma_\star| < 1/8 < |1 - (1/2 + \bar{\pi})|/2$, 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 $\bar{\pi} < 1/4$). □

5.2. Proof of Theorem 3.2. We consider d-state Markov chains of the following form:

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

for some $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_d \in (0,1)$. Such a chain is ergodic and reversible, and its unique stationary distribution $\boldsymbol{\pi}$ 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:

- $\boldsymbol{P}^{(0)}$: $\varepsilon_1 = \dots = \varepsilon_d = \varepsilon$. For this Markov chain, $\lambda_2 = \lambda_d = \lambda_\star = 1 - \frac{d}{d-1}\varepsilon$.
- $\boldsymbol{P}^{(i)}$ 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 $\lambda_\star = 1 - \frac{d/2}{d-1}\varepsilon$.

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

In order to guarantee $|\hat{\gamma}_\star - \gamma_\star| < \bar{\gamma}/2$, it must be possible to distinguish $\boldsymbol{P}^{(0)}$ from each $\boldsymbol{P}^{(i)}$, $i \in [d]$. But $\boldsymbol{P}^{(0)}$ is identical to $\boldsymbol{P}^{(i)}$ 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 $\boldsymbol{P}^{(0)}$ from each $\boldsymbol{P}^{(i)}$, $i \in [d]$. 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 $T_1, T_2, \ldots, T_{d-1}$ are independent, and are geometrically distributed, $T_i \sim \mathrm{Geom}(\varepsilon - (i-1)\varepsilon/(d-1))$. We have that

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

Therefore

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

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

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

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. ☐

6. Proof of Theorem 3.3

In this section, we prove Theorem 3.3.

6.1. Accuracy of $\hat{\pi}_\star$.

We start by proving the deviation bound on $\pi_\star - \hat{\pi}_\star$, from which we may easily deduce Eq. (4) in Theorem 3.3.

Lemma 6.1. *Pick any $\delta \in (0,1)$, and let*

$$\varepsilon_n := \frac{\ln\left(\frac{d}{\delta}\sqrt{\frac{2}{\pi_\star}}\right)}{\gamma_\star n}. \tag{7}$$

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

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

$$|\hat{\pi}_\star - \pi_\star| \leq 4\sqrt{\pi_\star\varepsilon_n} + 47\varepsilon_n. \tag{9}$$

Proof. We use the following Bernstein-type inequality for Markov chains of Paulin (2015, Theorem 3.3): letting $\mathbb{P}^{\boldsymbol{\pi}}$ denote the probability with respect to the stationary chain (where the marginal distribution of each X_t is $\boldsymbol{\pi}$), we have for every $\epsilon > 0$,

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

To handle possibly non-stationary chains, as is our case, we combine the above inequality with Paulin (2015, Proposition 3.10), to obtain for any $\epsilon > 0$,

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

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. (8) hold with probability at least $1-\delta$.

Now assume this $1-\delta$ probability event holds; it remains to prove that Eq. (9) also holds in this event. Without loss of generality, we assume that $\pi_\star = \pi_1 \leq \pi_2 \leq \cdots \leq \pi_d$. Let $j \in [d]$ be such that $\hat{\pi}_\star = \hat{\pi}_j$. By Eq. (8), 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., Bousquet, Boucheron, and Lugosi 2004) 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 $a, b \geq 0$. □

6.2. **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 $\mathrm{Sym}(\widehat{\boldsymbol{L}})$ in approximating $\boldsymbol{L}$ via Weyl's inequality:

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

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

$$\|\mathrm{Sym}(\widehat{\boldsymbol{L}}) - \boldsymbol{L}\| \leq \|\widehat{\boldsymbol{L}} - \boldsymbol{L}\|.$$

Therefore, we can deduce Eq. (5) in Theorem 3.3 from the following lemma.

Lemma 6.2. *There exists an absolute constant $C > 0$ such that the following holds. For any $\delta \in (0,1)$, if*

$$n \geq C\left(\frac{\log\frac{1}{\pi_\star\delta}}{\pi_\star\gamma_\star} + \frac{\log n}{\gamma_\star}\right), \tag{10}$$

then with probability at least $1-\delta$, the bounds from Lemma 6.1 hold, and

$$\|\widehat{\boldsymbol{L}} - \boldsymbol{L}\| \leq C\left(\sqrt{\varepsilon} + \varepsilon + \varepsilon^2\right),$$

where

$$\varepsilon := \frac{\left(\log\frac{d}{\delta}\right)\left(\log\frac{n}{\pi_\star\delta}\right)}{\pi_\star\gamma_\star n}.$$

We briefly describe how to obtain the bound on $|\hat{\gamma}_\star - \gamma_\star|$ that appears in Eq. (5), which is of the form $C'\sqrt{\varepsilon}$. Observe that if $\varepsilon > 1/C'$, then, owing to $C' \geq 1$, the bound on $|\hat{\gamma}_\star - \gamma_\star|$ is trivial. So we may assume that $\varepsilon \leq 1/C'$, which implies $n/\log n \geq C'(\log(d/\delta))/(\pi_\star\gamma_\star)$ (and thus $n \geq 2$), and also $n \geq C'(\log(d/\delta))(\log(1/(\pi_\star\delta)))/(\pi_\star\gamma_\star)$. These inequalities imply that n satisfies the condition in Eq. (10), so by Lemma 6.2, we have $|\hat{\gamma}_\star - \gamma_\star| \leq \|\widehat{\boldsymbol{L}} - \boldsymbol{L}\| \leq C(\sqrt{\varepsilon} + \varepsilon + \varepsilon^2) \leq C'\sqrt{\varepsilon}$.

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

When $\hat{\boldsymbol{\pi}}$ is positive valued, the error $\widehat{\boldsymbol{L}} - \boldsymbol{L}$ may be written as

$$\widehat{\boldsymbol{L}} - \boldsymbol{L} = \boldsymbol{\mathcal{E}}_{\boldsymbol{M}} + \boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}}\boldsymbol{L} + \boldsymbol{L}\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}} + \boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}}\boldsymbol{L}\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}} + \boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}}\boldsymbol{\mathcal{E}}_{\boldsymbol{M}} + \boldsymbol{\mathcal{E}}_{\boldsymbol{M}}\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}} + \boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}}\boldsymbol{\mathcal{E}}_{\boldsymbol{M}}\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}}\,,$$

where

$$\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}} := \operatorname{Diag}(\hat{\boldsymbol{\pi}})^{-1/2}\operatorname{Diag}(\boldsymbol{\pi})^{1/2} - \boldsymbol{I} \quad \text{and}$$
$$\boldsymbol{\mathcal{E}}_{\boldsymbol{M}} := \operatorname{Diag}(\boldsymbol{\pi})^{-1/2}\left(\widehat{\boldsymbol{M}} - \boldsymbol{M}\right)\operatorname{Diag}(\boldsymbol{\pi})^{-1/2}\,.$$

Therefore

$$\|\widehat{\boldsymbol{L}} - \boldsymbol{L}\| \leq \|\boldsymbol{\mathcal{E}}_{\boldsymbol{M}}\| + (\|\boldsymbol{\mathcal{E}}_{\boldsymbol{M}}\| + \|\boldsymbol{L}\|)\left(2\|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}}\| + \|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}}\|^2\right).$$

If $\|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}}\| \leq 1$ also holds, then, thanks to $\|\boldsymbol{L}\| \leq 1$,

$$\|\widehat{\boldsymbol{L}} - \boldsymbol{L}\| \leq \|\boldsymbol{\mathcal{E}}_{\boldsymbol{M}}\| + \|\boldsymbol{\mathcal{E}}_{\boldsymbol{M}}\|^2 + 3\|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}}\|. \tag{11}$$

6.3. A bound on $\|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}}\|$.

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

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

Assume that

$$n \geq \frac{108\ln\left(\frac{d}{\delta}\sqrt{\frac{2}{\pi_\star}}\right)}{\pi_\star\gamma_\star}, \tag{12}$$

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. (7). Therefore, on the $1-\delta$ probability event from Lemma 6.1, 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 $i \in [d]$. In particular, it also holds that $\hat{\boldsymbol{\pi}}$ is positive valued. Further, 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. (12), then on this $1-\delta$ probability event from Lemma 6.1, $\hat{\boldsymbol{\pi}}$ is positive valued and

$$\|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}}\| \leq \max_{i\in[d]}\left|\frac{\hat{\pi}_i}{\pi_i} - 1\right| \leq \max_{i\in[d]}\frac{\sqrt{8\pi_i(1-\pi_i)\varepsilon_n} + 20\varepsilon_n}{\pi_i} \tag{13}$$
$$\leq \sqrt{\frac{8\varepsilon_n}{\pi_\star}} + \frac{20\varepsilon_n}{\pi_\star} = \sqrt{\frac{8\ln\left(\frac{d}{\delta}\sqrt{\frac{2}{\pi_\star}}\right)}{\pi_\star\gamma_\star n}} + \frac{20\ln\left(\frac{d}{\delta}\sqrt{\frac{2}{\pi_\star}}\right)}{\pi_\star\gamma_\star n} \leq \min\{C'(\sqrt{\varepsilon}+\varepsilon), 1\}$$

for some suitable constant $C' > 0$, where ε as defined in Lemma 6.2.

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

6.4.1. *Bounding $\|\mathbb{E}(\boldsymbol{\mathcal{E}}_{\boldsymbol{M}})\|$: The price of a non-stationary start.* Let $\boldsymbol{\pi}^{(t)}$ be the distribution of states at time step t. We will make use of the following proposition, which can be derived by following Montenegro and Tetali (2006, Proposition 1.12):

Proposition 6.3. *For $t \geq 1$, let $\boldsymbol{\Upsilon}^{(t)}$ be the vector with $\Upsilon_i^{(t)} = \frac{\pi_i^{(t)}}{\pi_i}$ and let $\|\cdot\|_{2,\boldsymbol{\pi}}$ denote the $\boldsymbol{\pi}$-weighted 2-norm*

$$\|\boldsymbol{v}\|_{2,\boldsymbol{\pi}} := \left(\sum_{i=1}^{d} \pi_i v_i^2\right)^{1/2} . \tag{14}$$

Then,

$$\|\boldsymbol{\Upsilon}^{(t)} - \mathbf{1}\|_{2,\boldsymbol{\pi}} \leq \frac{(1-\gamma_\star)^{t-1}}{\sqrt{\pi_\star}} . \tag{15}$$

An immediate corollary of this result is that

$$\left\|\operatorname{Diag}(\boldsymbol{\pi}^{(t)})\operatorname{Diag}(\boldsymbol{\pi})^{-1} - \boldsymbol{I}\right\| \leq \frac{(1-\gamma_\star)^{t-1}}{\pi_\star} . \tag{16}$$

Now note that

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

and thus

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

Combining this, $\|\boldsymbol{L}\| \leq 1$ and Eq. (16), we get

$$\|\mathbb{E}(\boldsymbol{\mathcal{E}}_{\boldsymbol{M}})\| \leq \frac{1}{(n-1)\pi_\star}\sum_{t=1}^{n-1}(1-\gamma_\star)^{t-1} \leq \frac{1}{(n-1)\gamma_\star\pi_\star}. \tag{17}$$

6.4.2. *Bounding $\|\boldsymbol{\mathcal{E}}_{\boldsymbol{M}} - \mathbb{E}(\boldsymbol{\mathcal{E}}_{\boldsymbol{M}})\|$: Application of a matrix tail inequality.* In this section we analyze the deviations of $\boldsymbol{\mathcal{E}}_{\boldsymbol{M}} - \mathbb{E}(\boldsymbol{\mathcal{E}}_{\boldsymbol{M}})$. By the definition of $\boldsymbol{\mathcal{E}}_{\boldsymbol{M}}$,

$$\|\boldsymbol{\mathcal{E}}_{\boldsymbol{M}} - \mathbb{E}(\boldsymbol{\mathcal{E}}_{\boldsymbol{M}})\| = \|\operatorname{Diag}(\boldsymbol{\pi})^{-1/2}\left(\widehat{\boldsymbol{M}} - \mathbb{E}\widehat{\boldsymbol{M}}\right)\operatorname{Diag}(\boldsymbol{\pi})^{-1/2}\| . \tag{18}$$

The matrix $\widehat{\boldsymbol{M}} - \mathbb{E}\left(\widehat{\boldsymbol{M}}\right)$ is defined as a sum of dependent centered random matrices. We will use the blocking technique of Bernstein (1927) to relate the likely deviations of this matrix to that of a sum of independent centered random matrices. The deviations of these will then bounded with the help of a Bernstein-type matrix tail inequality due to Tropp (2015).

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

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

for $s = 1, 2, \ldots$. 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

$$\begin{aligned} \widehat{\boldsymbol{M}} &= \frac{1}{n-1} \sum_{t=1}^{n-1} \boldsymbol{e}_{X_t} \boldsymbol{e}_{X_{t+1}}^\top \\ &= \frac{a'}{n-1} \cdot \underbrace{\frac{1}{a'} \sum_{t \in F} \boldsymbol{e}_{X_t} \boldsymbol{e}_{X_{t+1}}^\top}_{\widehat{\boldsymbol{M}}_F} + \frac{n_H}{n-1} \cdot \underbrace{\frac{1}{\mu_H} \sum_{s=1}^{\mu_H} \left(\frac{1}{a} \sum_{t \in H_s} \boldsymbol{e}_{X_t} \boldsymbol{e}_{X_{t+1}}^\top \right)}_{\widehat{\boldsymbol{M}}_H} \\ &\quad + \frac{n_T}{n-1} \cdot \underbrace{\frac{1}{\mu_T} \sum_{s=1}^{\mu_T} \left(\frac{1}{a} \sum_{t \in T_s} \boldsymbol{e}_{X_t} \boldsymbol{e}_{X_{t+1}}^\top \right)}_{\widehat{\boldsymbol{M}}_T}. \end{aligned} \tag{19}$$

Here, $\boldsymbol{e}_i$ is the i-th coordinate basis vector, so $\boldsymbol{e}_i \boldsymbol{e}_j^\top \in \{0,1\}^{d \times d}$ 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:

$$\begin{aligned} &\frac{a'}{n-1} \left\| \operatorname{Diag}(\boldsymbol{\pi})^{-1/2} \left(\widehat{\boldsymbol{M}}_F - \mathbb{E}(\widehat{\boldsymbol{M}}_F) \right) \operatorname{Diag}(\boldsymbol{\pi})^{-1/2} \right\| \\ &\qquad \le \frac{1}{n-1} \sum_{t \in F} \left\{ \left\| \frac{\boldsymbol{e}_{X_t} \boldsymbol{e}_{X_{t+1}}^\top}{\sqrt{\pi_{X_t} \pi_{X_{t+1}}}} \right\| + \left\| \mathbb{E}\left(\frac{\boldsymbol{e}_{X_t} \boldsymbol{e}_{X_{t+1}}^\top}{\sqrt{\pi_{X_t} \pi_{X_{t+1}}}} \right) \right\| \right\} \le \frac{2a'}{\pi_\star (n-1)}. \end{aligned} \tag{20}$$

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

Let

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

so

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

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

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

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

$$a \geq a_\delta := \frac{1}{\gamma_\star} \ln \frac{2(n-2)}{\delta \pi_\star} \tag{21}$$

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

Define

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

Observe that

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

so

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

Define

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

We apply a matrix tail inequality to the average of *independent* copies of the $\boldsymbol{Z}_s$'s. More precisely, we will apply the tail inequality to independent copies $\widetilde{\boldsymbol{Z}}_s$, $s \in [\mu_H]$ of the random variables $\boldsymbol{Z}_s$ and then relate the average of $\widetilde{\boldsymbol{Z}}_s$ to that of $\boldsymbol{Z}_s$. The following probability inequality is from Tropp (2015, Theorem 6.1.1.).

Theorem 6.4 (Matrix Bernstein inequality)**.** *Let $\boldsymbol{Q}_1, \boldsymbol{Q}_2, \dots, \boldsymbol{Q}_m$ be a sequence of independent, random $d_1 \times d_2$ matrices. Assume that $\mathbb{E}(\boldsymbol{Q}_i) = \boldsymbol{0}$ and $\|\boldsymbol{Q}_i\| \leq R$ for each $1 \leq i \leq m$. Let $\boldsymbol{S} = \sum_{i=1}^m \boldsymbol{Q}_i$ and let*

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

Then, for all $t \geq 0$,

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

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

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

To apply Theorem 6.4, it suffices to bound the spectral norms of $\boldsymbol{Z}_s$ (almost surely), $\mathbb{E}(\boldsymbol{Z}_s \boldsymbol{Z}_s^\top)$, and $\mathbb{E}(\boldsymbol{Z}_s^\top \boldsymbol{Z}_s)$.

Range bound. By the triangle inequality,

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

For the first term, we have

$$\|\operatorname{Diag}(\boldsymbol{\pi})^{-1/2} \boldsymbol{Y}_s \operatorname{Diag}(\boldsymbol{\pi})^{-1/2}\| \leq \frac{1}{\pi_\star}. \tag{22}$$

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

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

Then, using Eq. (16),

$$\|\operatorname{Diag}(\boldsymbol{\pi}^{(H_s)})\operatorname{Diag}(\boldsymbol{\pi})^{-1} - \boldsymbol{I}\| \leq \frac{(1-\gamma_\star)^{a'+2(s-1)a}}{\pi_\star} \leq \frac{(1-\gamma_\star)^{a}}{\pi_\star} \leq 1\,, \tag{23}$$

where the last inequality follows from the assumption that the block length a satisfies Eq. (21). Combining this with $\|\operatorname{Diag}(\boldsymbol{\pi})^{-1/2}\boldsymbol{M}\operatorname{Diag}(\boldsymbol{\pi})^{-1/2}\| = \|\boldsymbol{L}\| \leq 1$, it follows that

$$\|\operatorname{Diag}(\boldsymbol{\pi})^{-1/2}\mathbb{E}(\boldsymbol{Y}_s)\operatorname{Diag}(\boldsymbol{\pi})^{-1/2}\| \leq 2 \tag{24}$$

by the triangle inequality. Therefore, together with Eq. (22), we obtain the range bound

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

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

$$\begin{aligned}&\mathbb{E}(\boldsymbol{Z}_s\boldsymbol{Z}_s^\top)\\ &= \frac{1}{a^2}\sum_{t\in H_s}\mathbb{E}\left(\operatorname{Diag}(\boldsymbol{\pi})^{-1/2}\boldsymbol{e}_{X_t}\boldsymbol{e}_{X_{t+1}}^\top\operatorname{Diag}(\boldsymbol{\pi})^{-1}\boldsymbol{e}_{X_{t+1}}\boldsymbol{e}_{X_t}^\top\operatorname{Diag}(\boldsymbol{\pi})^{-1/2}\right) && (25)\\ &\quad+\frac{1}{a^2}\sum_{\substack{t\neq t'\\ t,t'\in H_s}}\mathbb{E}\left(\operatorname{Diag}(\boldsymbol{\pi})^{-1/2}\boldsymbol{e}_{X_t}\boldsymbol{e}_{X_{t+1}}^\top\operatorname{Diag}(\boldsymbol{\pi})^{-1}\boldsymbol{e}_{X_{t'+1}}\boldsymbol{e}_{X_{t'}}^\top\operatorname{Diag}(\boldsymbol{\pi})^{-1/2}\right) && (26)\\ &\quad-\operatorname{Diag}(\boldsymbol{\pi})^{-1/2}\mathbb{E}(\boldsymbol{Y}_s)\operatorname{Diag}(\boldsymbol{\pi})^{-1}\mathbb{E}(\boldsymbol{Y}_s^\top)\operatorname{Diag}(\boldsymbol{\pi})^{-1/2}. && (27)\end{aligned}$$

The first sum, Eq. (25), easily simplifies to the diagonal matrix

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

For the second sum, Eq. (26), a symmetric matrix, consider

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

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

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

which simplifies to

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

The expectation is the same as that for the first term, Eq. (25).

Finally, the spectral norm of the third term, Eq. (27), is bounded using Eq. (24):

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

Therefore, by the triangle inequality, the bound $\pi_i^{(H)} / \pi_i \leq 2$ from Eq. (23), and simplifications,

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

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

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

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

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

again using the bound $\pi_i^{(H)} / \pi_i \leq 2$ from Eq. (23).

Independent copies bound. Let $\widetilde{\boldsymbol{Z}}_s$ for $s \in [\mu_H]$ be independent copies of $\boldsymbol{Z}_s$ for $s \in [\mu_H]$. Applying Theorem 6.4 to the average of these random matrices, we have

$$\mathbb{P}\left(\left\| \frac{1}{\mu_H} \sum_{s=1}^{\mu_H} \widetilde{\boldsymbol{Z}}_s \right\| > \sqrt{\frac{4\,(d_{\boldsymbol{P}} + 2) \ln \frac{4d}{\delta}}{\mu_H}} + \frac{2\left(\frac{1}{\pi_\star} + 2 \right) \ln \frac{4d}{\delta}}{3\mu_H} \right) \leq \delta \tag{28}$$

where

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

The actual bound. To bound the probability that $\|\sum_{s=1}^{\mu_H} \boldsymbol{Z}_s/\mu_H\|$ is large, we appeal to the following result (a consequence of Yu 1994, 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 $\boldsymbol{Z}_s$. 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 $(X^{(H_1)}, X^{(H_2)}, \ldots, X^{(H_{s+1})})$. Let $\widetilde{\mathbb{P}}$ be the product distribution formed from the marginals $\mathbb{P}_1, \mathbb{P}_2, \ldots, \mathbb{P}_{\mu_H}$, so $\widetilde{\mathbb{P}}$ governs the joint distribution of $(\widetilde{\boldsymbol{Z}}_s : s \in [\mu_H])$. The result from Yu (1994, Corollary 2.7) implies for any event E,

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

where

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

Here, $\|\cdot\|_{\mathrm{tv}}$ 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. (28) for $\|\sum_{s=1}^{\mu_H} \widetilde{\boldsymbol{Z}}_s/\mu_H\|$ also holds for $\|\sum_{s=1}^{\mu_H} \boldsymbol{Z}_s/\mu_H\|$, except the probability bound increases from δ to $\delta + (\mu_H - 1)\beta(\mathbb{P})$:

(29)
$$\mathbb{P}\left(\left\|\frac{1}{\mu_H}\sum_{s=1}^{\mu_H} \boldsymbol{Z}_s\right\| > \sqrt{\frac{4\,(d_{\boldsymbol{P}}+2)\ln\frac{4d}{\delta}}{\mu_H}} + \frac{2\left(\frac{1}{\pi_\star}+2\right)\ln\frac{4d}{\delta}}{3\mu_H}\right) \leq \delta + (\mu_H - 1)\beta(\mathbb{P}).$$

By the triangle inequality,

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

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

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

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* $1/\gamma_\star$ (see, e.g., the proof of Theorem 12.3 of Levin, Peres, and Wilmer 2009): for any $i \in [d]$,

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

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

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

We conclude

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

where the last step follows from the block length assumption Eq. (21).

We return to the decomposition from Eq. (19). We apply Eq. (29) to both the H_s blocks and the T_s blocks, and combine with Eq. (20) to obtain the following probabilistic bound. Pick any $\delta \in (0, 1)$, let the block length be

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

so

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

If

$$n \geq 7 + \frac{6}{\gamma_\star} \ln \frac{2(n-2)}{\pi_\star \delta} \geq 3a, \tag{30}$$

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

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

6.4.3. *The bound on* $\|\boldsymbol{\mathcal{E}}_{\boldsymbol{M}}\|$. Combining the probabilistic bound from above with the bound on the bias from Eq. (17), we obtain the following. Assuming the condition on n from Eq. (30), with probability at least $1 - 4\delta$,

$$\begin{aligned}
\|\boldsymbol{\mathcal{E}}_{\boldsymbol{M}}\| &\leq \frac{1}{(n-1)\gamma_\star\pi_\star} + \frac{4\left\lceil \frac{1}{\gamma_\star} \ln \frac{2(n-2)}{\pi_\star\delta} \right\rceil}{\pi_\star(n-1)} \\
&\quad + \sqrt{\frac{4\,(d_{\boldsymbol{P}} + 2)\ln\frac{4d}{\delta}}{\mu}} + \frac{2\left(\frac{1}{\pi_\star} + 2\right)\ln\frac{4d}{\delta}}{3\mu} \leq C'\left(\sqrt{\varepsilon} + \varepsilon\right),
\end{aligned} \tag{31}$$

for some suitable constant $C' > 0$, where ε as defined in Lemma 6.2.

6.5. **Overall error bound.** Observe that the assumption on the sequence length in Eq. (10) implies the conditions in Eq. (12) and Eq. (30) for a suitable choice of $C > 0$. With this assumption, there is a $1-5\delta$ probability event in which Eqs. (8), (9) and (31) hold; in particular, we have the bound on $\|\boldsymbol{\mathcal{E}}_{\boldsymbol{M}}\|$ from Eq. (31). In this event, the bound on $\|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}}\|$ in Eq. (13) also holds, and the claimed bound on $\|\widehat{\boldsymbol{L}}-\boldsymbol{L}\|$ follows from combining the bound in Eq. (11) with the bounds on $\|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}}\|$ and $\|\boldsymbol{\mathcal{E}}_{\boldsymbol{M}}\|$:

$$\begin{aligned}\|\widehat{\boldsymbol{L}}-\boldsymbol{L}\| &\le \|\boldsymbol{\mathcal{E}}_{\boldsymbol{M}}\| + \|\boldsymbol{\mathcal{E}}_{\boldsymbol{M}}\|^2 + 3\|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi}}\| \\ &\le 4C'\left(\sqrt{\varepsilon}+\varepsilon\right) + C'^2\left(\sqrt{\varepsilon}+\varepsilon\right)^2 \le C\left(\sqrt{\varepsilon}+\varepsilon+\varepsilon^2\right),\end{aligned}$$

where ε is defined in the statement of Lemma 6.2. The proof of Lemma 6.2 now follows by replacing δ with $\delta/5$. □

7. Proof of Theorem 3.4

In this section, we prove Theorem 3.4.

We call $\hat{\gamma}_\star$ of Theorem 3.3 the *initial estimator*. Let C be the constant from Theorem 3.3, and define

$$n_1 = n_1(\varepsilon;\delta,\gamma_\star) := \frac{3C^2}{\varepsilon^2\pi_\star\gamma_\star}\cdot\left(\log\frac{d}{\delta}\right)\cdot\left(\log\frac{3C^2}{\varepsilon^2\pi_\star^2\gamma_\star\delta}\right)$$

and

$$M(n;\delta,\gamma_\star) := C\sqrt{\frac{\log\frac{d}{\delta}\cdot\log\frac{n}{\pi_\star\delta}}{\pi_\star\gamma_\star n}},$$

which is the right-hand side of Eq. (5). Observe that

$$M(n_1;\delta,\gamma_\star) \le \varepsilon\sqrt{\frac{\log\frac{3C^2}{\varepsilon^2\pi_\star^2\gamma_\star\delta} + \log\log\frac{d}{\delta} + \log\log\frac{3C^2}{\varepsilon^2\pi_\star^2\gamma_\star\delta}}{3\log\frac{3C^2}{\varepsilon^2\pi_\star^2\gamma_\star^2\delta}}} \le \varepsilon.$$

(Each term in the numerator under the radical is at most a third of the denominator. We have used that $\pi_\star \le 1/d$ in comparing the second term in the numerator to the denominator.)

For $a > 0$, the spectral gap of the chain with transition matrix $\boldsymbol{P}^a$ is denoted by $\gamma_\star(a)$, and the initial estimator of $\gamma_\star(a)$, based on n/a steps of $\boldsymbol{P}^a$, is denoted by $\hat{\gamma}_\star(a)$. Note that

$$\gamma_\star(a) = 1-(1-\gamma_\star)^a\,.$$

Define $K_{\gamma_\star} := \lfloor\log_2(1/\gamma_\star)\rfloor$ and, for any $\delta\in(0,1)$, $\delta_{\gamma_\star} = \delta_{\gamma_\star}(\delta) := \delta/(K_{\gamma_\star}+1)$.

Proposition 7.1. *Fix $\varepsilon \in (0, 0.01)$ and $\delta \in (0,1)$. Let A be the random variable defined in the estimator of Theorem 3.4 (which depends on $(X_t)_{t=1}^n$). If $n > n_1(\varepsilon/\sqrt{2};\delta_{\gamma_\star},\gamma_\star)$, then there is an event $G(\varepsilon)$ having probability at least $1-\delta$, such that on $G(\varepsilon)$,*

$$\begin{aligned} 0.30 < \gamma_\star(A) < 0.54 &\quad \text{if } \gamma_\star < 1/2\,, \\ A = 1 &\quad \text{if } \gamma_\star \ge 1/2\,. \end{aligned}$$

Moreover, on $G(\varepsilon)$, the initial estimator $\hat{\gamma}_\star(A)$ applied to the chain $(X_{As})_{s=1}^{n/A}$ satisfies

$$|\hat{\gamma}_\star(A)-\gamma_\star(A)| \le \varepsilon\,. \tag{32}$$

The proof of Proposition 7.1 is based on the following lemma.

Lemma 7.2. *Fix $n \geq n_1(\varepsilon/\sqrt{2}; \delta, \gamma_\star)$. If $a\gamma_\star \leq 1$, then*

$$\Pr(|\gamma_\star(a) - \hat{\gamma}_\star(a)| \leq \varepsilon) > 1 - \delta.$$

Proof. Recall the bound $M(n; \delta, \gamma_\star)$ on the right-hand side of Eq. (5). If $\gamma_\star(a) \geq \gamma_\star a/2$, then

$$M(n/a; \delta, \gamma_\star(a)) \leq \sqrt{2}M(n; a\delta, \gamma_\star) \leq \sqrt{2}M(n; \delta, \gamma_\star) \leq \sqrt{2} \cdot \frac{\varepsilon}{\sqrt{2}} = \varepsilon\,,$$

and the lemma follows from applying Theorem 3.3 to the $\boldsymbol{P}^a$-chain. We now show that $\gamma_\star(a) \geq \gamma_\star a/2$. A Taylor expansion of $(1 - \gamma_\star)^a$ implies that there exists $\xi \in [0, \gamma_\star] \subseteq [0, 1/a]$ such that

$$\gamma_\star(a) = 1 - (1 - \gamma_\star)^a = \gamma_\star a - \frac{a(a-1)(1-\xi)^{a-2}\gamma_\star^2}{2} \geq \frac{\gamma_\star a}{2}\,.$$

(We have used the hypothesis $a\gamma_\star \leq 1$ in the inequality.) □

Proof of Proposition 7.1. Define the events $G(a; \varepsilon) := \{|\gamma_\star(a) - \hat{\gamma}_\star(a)| \leq \varepsilon\}$, and $G = G(\varepsilon) := \bigcap_{k=0}^{K_{\gamma_\star}} G(2^k; \varepsilon)$. If $k \leq K_{\gamma_\star}$, then $\gamma_\star 2^k \leq \gamma_\star 2^{\log_2(1/\gamma_\star)} \leq 1$ and Lemma 7.2 implies that

$$\Pr(G^c) \leq \sum_{k=0}^{K_{\gamma_\star}} \Pr(G(2^k; \varepsilon)^c) \leq (K_{\gamma_\star} + 1) \cdot \frac{\delta}{K_{\gamma_\star} + 1} = \delta\,.$$

On G, if $\gamma_\star \geq 1/2$, then $|\hat{\gamma}_\star - \gamma_\star| \leq 0.01$, and consequently $\hat{\gamma}_\star \geq 0.49 > 0.31$. In this case, $A = 1$ on G.

On the event G, if the algorithm has not terminated by step $k - 1$, then the following hold:

(1) If $\gamma_\star(2^k) \leq 0.30$, then the algorithm does not terminate at step k.
(2) If $\gamma_\star(2^k) > 0.32$, then the algorithm terminates at step k.

Also, assuming $\gamma_\star \leq 1/2$,

$$\gamma_\star(2^{K_{\gamma_\star}}) \geq 1 - (1 - \gamma_\star)^{\frac{1}{2\gamma_\star}} \geq 1 - e^{-1/2} \geq 0.39\,,$$

so the algorithm always terminates before $k = K_{\gamma_\star}$ on G and thus (32) holds on G.

Finally, on G, if $A > 1$, then $\gamma_\star(A/2) \leq 0.32$, whence

$$\gamma_\star(A) = 1 - (1 - \gamma_\star(A/2))^2 \leq 1 - (0.68)^2 < 0.54\,.$$

If $\gamma_\star < 1/2$ and $A = 1$, then $\gamma_\star(A) = \gamma_\star \leq 1/2$. □

We now prove Theorem 3.4.

Proof of Theorem 3.4. Let

$$n_0(\varepsilon; \delta, \gamma_\star, \pi_\star) = n_0(\varepsilon) := \frac{\mathcal{L}}{\pi_\star \gamma_\star \varepsilon^2}\,, \tag{33}$$

where

$$\mathcal{L} := 3 \cdot (16\sqrt{2})^2 \cdot \left(\log \frac{d(\lfloor \log_2(1/\gamma_\star) \rfloor + 1)}{\delta}\right) \cdot \left(\log \frac{3 \cdot (16\sqrt{2})^2 \cdot C^2(\lfloor \log_2(1/\gamma_\star) \rfloor + 1)}{\varepsilon^2 \pi_\star^2 \gamma_\star \delta}\right)\,, \tag{34}$$

and C is the constant in Eq. (5).

Fix $n > n_0(\varepsilon) = n_1(\varepsilon/(16\sqrt{2}); \delta_{\gamma_\star}, \gamma_\star)$. Let A and G be as defined in Proposition 7.1. Assume we are on the event $G = G(\varepsilon/16)$ for the rest of this proof.

Suppose first that $\gamma_\star < 1/2$. We have $0.30 < \gamma_\star(A) < 0.54$, and

$$|\hat{\gamma}_\star(A) - \gamma_\star(A)| \leq \frac{\varepsilon}{16} < 0.01\,,$$

so both $\gamma_\star(A)$ and $\hat{\gamma}_\star(A)$ are in $[0.29, 0.55]$, say.

Let $h(x) = 1-(1-x)^{1/A}$, so $\gamma_\star = h(\gamma_\star(A))$ and $\tilde{\gamma}_\star = h(\hat{\gamma}_\star(A))$. Since $(1-x)^{1/A} \leq 1 - x/A$, we have

$$\frac{1}{1-(1-x)^{1/A}} \leq \frac{A}{x}\,.$$

Consequently, on $[0.29, 0.55]$,

$$\left|\frac{d}{dx}\log h(x)\right| = \frac{\frac{1}{A}(1-x)^{1/A-1}}{1-(1-x)^{1/A}} \leq \frac{1}{A(1-x)}\frac{A}{x} = \frac{1}{(1-x)x} \leq \frac{1}{(0.45)(0.29)} < 8\,.$$

Thus, $|\frac{d}{dx}\log h(x)|$ is bounded (by 8) on $[0.29, 0.55]$. We have

$$|\log(h(\hat{\gamma}_\star(A))/\gamma_\star)| = |\log h(\gamma_\star(A)) - \log h(\hat{\gamma}_\star(A))| \leq 8|\gamma_\star(A) - \hat{\gamma}_\star(A)| \leq 8\frac{\varepsilon}{16} \leq \frac{\varepsilon}{2}\,.$$

Thus,

$$\frac{\tilde{\gamma}_\star}{\gamma_\star} = \frac{h(\hat{\gamma}_\star(A))}{\gamma_\star} \leq e^{\varepsilon/2} \leq 1+\varepsilon\,.$$

Similarly, $\frac{\gamma_\star}{h(\hat{\gamma}_\star(A))} \leq e^{\varepsilon/2}$, so

$$\frac{\tilde{\gamma}_\star}{\gamma_\star} = \frac{h(\hat{\gamma}_\star(A))}{\gamma_\star} \geq e^{-\varepsilon/2} \geq 1-\varepsilon\,.$$

Now instead suppose that $\gamma_\star \geq 1/2$. Then $A = 1$ on the event G, and

$$|\tilde{\gamma}_\star - \gamma_\star| < \frac{\varepsilon}{16}\,,$$

so

$$\left|\frac{\tilde{\gamma}_\star}{\gamma_\star} - 1\right| < \frac{\varepsilon}{16\gamma_\star} \leq \varepsilon\,. \qquad \square$$

8. Proof of Theorem 4.1

In this section, we derive Algorithm 1 and prove Theorem 4.1.

8.1. Estimators for $\boldsymbol{\pi}$ and $\gamma_\star$. The algorithm forms the estimator $\widehat{\boldsymbol{P}}$ of $\boldsymbol{P}$ 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 $\widehat{\boldsymbol{P}}$ are positive, and hence $\widehat{\boldsymbol{P}}$ is a transition probability matrix for an ergodic Markov chain. We let $\hat{\boldsymbol{\pi}}$ be the unique stationary distribution for $\widehat{\boldsymbol{P}}$. Using $\hat{\boldsymbol{\pi}}$, we form an estimator $\mathrm{Sym}(\widehat{\boldsymbol{L}})$ of $\boldsymbol{L}$ using:

$$\mathrm{Sym}(\widehat{\boldsymbol{L}}) := \frac{1}{2}(\widehat{\boldsymbol{L}} + \widehat{\boldsymbol{L}}^\top), \qquad \widehat{\boldsymbol{L}} := \mathrm{Diag}(\hat{\boldsymbol{\pi}})^{1/2}\widehat{\boldsymbol{P}}\,\mathrm{Diag}(\hat{\boldsymbol{\pi}})^{-1/2}.$$

Let $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \cdots \geq \hat{\lambda}_d$ be the eigenvalues of $\mathrm{Sym}(\widehat{\boldsymbol{L}})$ (and in fact, we have $1 = \hat{\lambda}_1 > \hat{\lambda}_2$ and $\hat{\lambda}_d > -1$). The algorithm estimates the spectral gap $\gamma_\star$ using

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

8.2. **Empirical bounds for $\boldsymbol{P}$.** We make use of a simple corollary of Freedman's inequality for martingales (Freedman 1975, Theorem 1.6).

Theorem 8.1 (Freedman's inequality)**.** *Let $(Y_t)_{t\in\mathbb{N}}$ be a bounded martingale difference sequence with respect to the filtration $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \mathcal{F}_2 \subset \cdots$; assume for some $b > 0$, $|Y_t| \leq b$ almost surely for all $t \in \mathbb{N}$. Let $V_k := \sum_{t=1}^k \mathbb{E}\left(Y_t^2 | \mathcal{F}_{t-1}\right)$ and $S_k := \sum_{t=1}^k Y_t$ for $k \in \mathbb{N}$. For all $s, v > 0$,*

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

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

Observe that in Theorem 8.1, 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

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

since $h(\sqrt{2z} + z/3)/z \geq 1$ for all $z > 0$ (see, e.g., Audibert, Munos, and Szepesvári (2009, proof of Lemma 5)).

Corollary 8.2. *Under the same setting as Theorem 8.1, for any $n \geq 1$, $x > 0$, and $c > 1$,*

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

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

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

where the final inequality uses Theorem 8.1. 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 8.3. *The following holds for any constant $c > 1$ with probability at least $1-\delta$: for all $(i,j) \in [d]^2$,*

$$|\widehat{P}_{i,j} - P_{i,j}| \leq \sqrt{\left(\frac{N_i}{N_i + d\alpha}\right) \frac{2cP_{i,j}(1-P_{i,j})\tau_{n,\delta}}{N_i + d\alpha}} + \frac{(4/3)\tau_{n,\delta}}{N_i + d\alpha} + \frac{|\alpha - d\alpha P_{i,j}|}{N_i + d\alpha}, \tag{35}$$

where

$$\tau_{n,\delta} := \inf\left\{t \geq 0 : 2d^2 \left(1 + \lceil \log_c(2n/t) \rceil_+\right) e^{-t} \leq \delta\right\} = O\left(\log\left(\frac{d\log(n)}{\delta}\right)\right). \tag{36}$$

Proof. Let $\mathcal{F}_t$ be the σ-field generated by $X_1, X_2, \ldots, X_t$. 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 $t \geq 2$,

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

Therefore, by Corollary 8.2 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 (35) 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. (35), we can numerically maximize $|\widehat{P}_{i,j} - P_{i,j}|$ subject to the constraint in Eq. (35) (viewing $P_{i,j}$ as the optimization variable). Let $B^*_{i,j}$ be this maximum value, so we have

$$P_{i,j} \in \left[\widehat{P}_{i,j} - B^*_{i,j},\ \widehat{P}_{i,j} + B^*_{i,j}\right]$$

in the same event where Eq. (35) 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 (Audibert, Munos, and Szepesvári 2009). Specifically, with $c := 1.1$ (an arbitrary choice), we compute

$$\widehat{B}_{i,j} := \left(\sqrt{\frac{c\tau_{n,\delta}}{2N_i}} + \sqrt{\frac{c\tau_{n,\delta}}{2N_i} + \sqrt{\frac{2c\widehat{P}_{i,j}(1-\widehat{P}_{i,j})\tau_{n,\delta}}{N_i}} + \frac{(4/3)\tau_{n,\delta} + |\alpha - d\alpha\widehat{P}_{i,j}|}{N_i}}\right)^2 \tag{37}$$

for each $(i,j) \in [d]^2$, where $\tau_{n,\delta}$ is defined in Eq. (36). We show in Lemma 8.4 that

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

again, in the same event where Eq. (35) holds. The observable bound in Eq. (37) is not too far from the unobservable bound in Eq. (35).

Lemma 8.4. *In the same $1-\delta$ event as from Lemma 8.3, we have $P_{i,j} \in [\widehat{P}_{i,j} - \widehat{B}_{i,j}, \widehat{P}_{i,j} + \widehat{B}_{i,j}]$ for all $(i,j) \in [d]^2$, where $\widehat{B}_{i,j}$ is defined in Eq. (37).*

Proof. Recall that in the $1-\delta$ probability event from Lemma 8.3, we have for all $(i,j) \in [d]^2$,

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

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

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

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

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

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

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

$$\widehat{\boldsymbol{A}}\widehat{\boldsymbol{A}}^{\#}\widehat{\boldsymbol{A}} = \widehat{\boldsymbol{A}}, \quad \widehat{\boldsymbol{A}}^{\#}\widehat{\boldsymbol{A}}\widehat{\boldsymbol{A}}^{\#} = \widehat{\boldsymbol{A}}^{\#}, \quad \widehat{\boldsymbol{A}}^{\#}\widehat{\boldsymbol{A}} = \widehat{\boldsymbol{A}}\widehat{\boldsymbol{A}}^{\#}.$$

The matrix $\widehat{\boldsymbol{A}}^{\#}$, which is well defined no matter what transition probability matrix $\widehat{\boldsymbol{P}}$ we start with (Meyer Jr. 1975), is a central quantity that captures many properties of the ergodic Markov chain with transition matrix $\widehat{\boldsymbol{P}}$ (Meyer Jr. 1975). We denote the (i,j)-th entry of $\widehat{\boldsymbol{A}}^{\#}$ by $\widehat{\boldsymbol{A}}^{\#}_{i,j}$. Define

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

Analogously define

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

We now use the following perturbation bound from Cho and Meyer (2001, Section 3.3) (derived from Haviv and Van der Heyden (1984) and Kirkland, Neumann, and Shader (1998)).

Lemma 8.5 (Haviv and Van der Heyden 1984; Kirkland, Neumann, and Shader 1998)**.** *If* $|\widehat{P}_{i,j} - P_{i,j}| \leq \widehat{B}_{i,j}$ *for each* $(i,j) \in [d]^2$*, then*

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

This establishes the validity of the confidence intervals for the π_i in the same event from Lemma 8.3.

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 8.6. *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 8.5, 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 $i \in [d]$,

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

□

8.4. **Empirical bounds for $\boldsymbol{L}$.** By Weyl's inequality and the triangle inequality,

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

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

Lemma 8.7. *If* $|\widehat{P}_{i,j} - P_{i,j}| \leq \widehat{B}_{i,j}$ *for each* $(i,j) \in [d]^2$ *and* $\max\{|\hat{\pi}_i - \pi_i| : i \in [d]\} \leq \hat{b}$*, then*

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

where

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

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

$$\boldsymbol{L} - \widehat{\boldsymbol{L}} = \boldsymbol{\mathcal{E}}_{\boldsymbol{P}} + \boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},1}\widehat{\boldsymbol{L}} + \widehat{\boldsymbol{L}}\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},2} + \boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},1}\boldsymbol{\mathcal{E}}_{\boldsymbol{P}} + \boldsymbol{\mathcal{E}}_{\boldsymbol{P}}\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},2} + \boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},1}\widehat{\boldsymbol{L}}\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},2} + \boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},1}\boldsymbol{\mathcal{E}}_{\boldsymbol{P}}\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},2}$$

where

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

Therefore

$$\begin{aligned}\|\boldsymbol{L} - \widehat{\boldsymbol{L}}\| \le\ & \|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},1}\| + \|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},2}\| + \|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},1}\|\|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},2}\|\\ &+ (1 + \|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},1}\| + \|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},2}\| + \|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},1}\|\|\boldsymbol{\mathcal{E}}_{\boldsymbol{\pi},2}\|)\,\|\boldsymbol{\mathcal{E}}_{\boldsymbol{P}}\|.\end{aligned}$$

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

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

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

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

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

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

which can be bounded using Lemma 8.6. □

This establishes the validity of the confidence interval for $\gamma_\star$ in the same event from Lemma 8.3.

8.5. **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 $n \to \infty$,

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

Here, we use the fact that $\hat{\kappa} \to \kappa$ as $n \to \infty$ since $\widehat{\boldsymbol{A}}^{\#} \to \boldsymbol{A}^{\#}$ as $\widehat{\boldsymbol{P}} \to \boldsymbol{P}$ (Li and Wei 2001; Benítez and X. Liu 2012).

Further, since

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

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 4.1. □

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

Lemma 8.8. $\kappa \leq \frac{1}{\gamma_\star} \min\{d, 8 + \ln(4/\pi_\star)\}$

Before proving this, we prove a lemma of independent interest.

Lemma 8.9. *Let τ_j be the first positive time that state j is visited by the Markov chain. Then*

$$\mathbb{E}_i \tau_j \leq 2 \left(t_{\text{mix}} + 8 \frac{t_{\text{relax}}}{\pi_j} \right) . \tag{38}$$

Proof. By taking f to be the indicator of state j in Theorem 12.19 of Levin, Peres, and Wilmer (2009), for any i, if $t = t_{\text{mix}} + 8 t_{\text{relax}}/\pi_j$, then

$$\Pr_i(\tau_j > t) \leq \frac{1}{2} .$$

Thus, $\Pr_i(\tau_j > tk) \leq 2^{-k}$, whence Eq. (38) follows. □

Proof of Lemma 8.8. It is established by Cho and Meyer (2001) that

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

(our κ is the κ_4 quantity from Cho and Meyer (2001)), and Seneta (1993) establishes

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

Since it is shown in Cho and Meyer (2001) that

$$\kappa = \frac{1}{2} \max_j \left[\max_{i \neq j} \mathbb{E}_i(\tau_j) \right] \pi_j ,$$

it follows from Lemma 8.9 that

$$\kappa \leq t_{\text{mix}} + 8 t_{\text{relax}} \leq t_{\text{relax}} (8 + \ln(4/\pi_\star)) .$$ □

9. Proof of Theorem 4.2

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 3.3. By a union bound, we have by Theorems 3.3 and 4.1 that with probability at least $1 - 2\delta$,

$$|\hat{\pi}_\star - \pi_\star| \leq C \left(\sqrt{\frac{\pi_\star \log \frac{d}{\hat{\pi}_{\star,\text{lb}} \delta}}{\hat{\gamma}_{\star,\text{lb}} n}} + \frac{\log \frac{d}{\hat{\pi}_{\star,\text{lb}} \delta}}{\hat{\gamma}_{\star,\text{lb}} n} \right) \tag{39}$$

and

$$|\hat{\gamma}_\star - \gamma_\star| \leq C \left(\sqrt{\frac{\log \frac{d}{\delta} \cdot \log \frac{n}{\hat{\pi}_{\star,\text{lb}} \delta}}{\hat{\pi}_{\star,\text{lb}} \hat{\gamma}_{\star,\text{lb}} n}} + \frac{\log \frac{d}{\delta} \cdot \log \frac{n}{\hat{\pi}_{\star,\text{lb}} \delta}}{\hat{\pi}_{\star,\text{lb}} \hat{\gamma}_{\star,\text{lb}} n} + \frac{\log \frac{1}{\hat{\gamma}_{\star,\text{lb}}}}{\hat{\gamma}_{\star,\text{lb}} n} \right) . \tag{40}$$

The bound on $|\hat{\gamma}_\star - \gamma_\star|$ in Eq. (40)—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. (39) depends on $\pi_\star$, but from it one can derive

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

using the approach from the proof of Lemma 8.4. 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 $n \to \infty$),

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

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 4.2. □

10. Discussion

The construction used in Theorem 4.2 applies more generally: Given a confidence interval of the form $I_n = I_n(\gamma_\star, \pi_\star, \delta)$ for some confidence level δ and a confidence set $E_n(\delta)$ for $(\gamma_\star, \pi_\star)$ for the same level, $I'_n = E_n(\delta) \cap \cup_{(\gamma,\pi) \in E_n(\delta)} I_n(\gamma, \pi, \delta)$ is a valid 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. The bootstrap estimator of Theorem 3.4 closes most of the gap when $\boldsymbol{\pi}$ is uniform. 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. To subvert our lower bounds, we must restrict attention to Markov chains with additional structure. Parametric classes, such as Markov chains with factored transition kernels with a few factors, are promising candidates 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 on the topic of fully empirical estimation of Markov chain parameters based on a single sample path.

References

Atchadé, Y. F. (2016). "Markov chain Monte Carlo confidence intervals". In: *Bernoulli* 22.3, pp. 1808–1838. ISSN: 1350-7265. DOI: 10.3150/15-BEJ712.

Audibert, J.-Y., R. Munos, and Cs. Szepesvári (2009). "Exploration-exploitation Tradeoff using Variance Estimates in Multi-armed Bandits". In: *Theoretical Computer Science* 410.19, pp. 1876–1902.

Batu, T., L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White (2000). "Testing that distributions are close". In: *FOCS*. IEEE, pp. 259–269.

– (2013). "Testing closeness of discrete distributions". In: *Journal of the ACM (JACM)* 60.1, 4:2–4:25.

Benítez, J. and X. Liu (2012). "On the continuity of the group inverse". In: *Operators and Matrices* 6.4, pp. 859–868.

Bernstein, S. (1927). "Sur l'extension du theoreme limite du calcul des probabilites aux sommes de quantites dependantes". In: *Mathematische Annalen* 97, pp. 1–59.

Bhatnagar, N., A. Bogdanov, and E. Mossel (2011). "The computational complexity of estimating MCMC convergence time". In: *RANDOM*. Springer, pp. 424–435.

Bhattacharya, B. B. and G. Valiant (2015). "Testing Closeness With Unequal Sized Samples". In: *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pp. 2611–2619. URL: http://papers.nips.cc/paper/5908-testing-closeness-with-unequal-sized-samples.

Bousquet, O., S. Boucheron, and G. Lugosi (2004). "Introduction to statistical learning theory". In: *Lecture Notes in Artificial Intelligence* 3176, pp. 169–207.

Cho, G. and C. Meyer (2001). "Comparison of perturbation bounds for the stationary distribution of a Markov chain". In: *Linear Algebra and its Applications* 335, pp. 137–150.

Flegal, J. M. and G. L. Jones (2011). "Implementing MCMC: estimating with confidence". In: *Handbook of Markov chain Monte Carlo*. Chapman & Hall/CRC, pp. 175–197.

Freedman, D. (1975). "On tail probabilities for martingales". In: *The Annals of Probability* 3.1, pp. 100–118.

Gamarnik, D. (2003). "Extension of the PAC framework to finite and countable Markov chains". In: *IEEE Transactions on Information Theory* 49.1, pp. 338–345.

Garren, S. T. and R. L. Smith (2000). "Estimating the second largest eigenvalue of a Markov transition matrix". In: *Bernoulli* 6, pp. 215–242.

Gillman, D. (1998). "A Chernoff bound for random walks on expander graphs". In: *SIAM Journal on Computing* 27.4, pp. 1203–1220.

Gyori, B. M. and D. Paulin (2014). "Non-asymptotic confidence intervals for MCMC in practice". arXiv:1212.2016.

Haviv, M. and L. Van der Heyden (1984). "Perturbation bounds for the stationary probabilities of a finite Markov chain". In: *Advances in Applied Probability* 16, pp. 804–818.

Hsu, D., A. Kontorovich, and Cs. Szepesvári (2015). "Mixing Time Estimation in Reversible Markov Chains from a Single Sample Path". In: *Advances in Neural Information Processing Systems 28*. Ed. by C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett. Curran Associates, Inc.

Jones, G. L. and J. P. Hobert (2001). "Honest Exploration of Intractable Probability Distributions via Markov Chain Monte Carlo". In: *Statist. Sci.* 16.4, pp. 312–

334. DOI: 10.1214/ss/1015346317. URL: http://dx.doi.org/10.1214/ss/1015346317.

Karandikar, R. L. and M. Vidyasagar (2002). "Rates of uniform convergence of empirical means with mixing processes". In: *Statistics and Probability Letters* 58.3, pp. 297–307.

Kipnis, C. and S. R. S. Varadhan (1986). "Central limit theorem for additive functionals of reversible Markov processes and applications to simple exclusions". In: *Comm. Math. Phys.* 104.1, pp. 1–19. URL: http://projecteuclid.org/euclid.cmp/1104114929.

Kirkland, S., M. Neumann, and B. Shader (1998). "Applications of Paz's inequality to perturbation bounds for Markov chains". In: *Linear Algebra and its Applications* 268, pp. 183–196.

Kontorovich, A. and R. Weiss (2014). "Uniform Chernoff and Dvoretzky-Kiefer-Wolfowitz-type inequalities for Markov chains and related processes". In: *Journal of Applied Probability* 51.4, pp. 1100–1113.

Kontoyiannis, I., L. A. Lastras-Montaño, and S. P. Meyn (2006). "Exponential bounds and stopping rules for MCMC and general Markov chains". In: *VALUETOOLS*, p. 45.

León, C. A. and F. Perron (2004). "Optimal Hoeffding bounds for discrete reversible Markov chains". In: *Annals of Applied Probability*, pp. 958–970.

Levin, D. A. and Y. Peres (2016). *Estimating the Spectral Gap of a Reversible Markov Chain from a Short Trajectory*. arXiv: 1612.05330 [math.ST/stat.TH].

Levin, D. A., Y. Peres, and E. L. Wilmer (2009). *Markov chains and mixing times*. With a chapter by James G. Propp and David B. Wilson. American Mathematical Society, Providence, RI.

Li, X. and Y. Wei (2001). "An improvement on the perturbation of the group inverse and oblique projection". In: *Linear Algebra and its Applications* 338, pp. 53–66.

Liu, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. Springer Series in Statistics. Springer-Verlag.

McDonald, D., C. Shalizi, and M. Schervish (2011). "Estimating beta-mixing coefficients". In: *AISTATS*, pp. 516–524.

Meyer Jr., C. D. (1975). "The Role of the Group Generalized Inverse in the Theory of Finite Markov Chains". In: *SIAM Review* 17.3, pp. 443–464.

Meyn, S. P. and R. L. Tweedie (1993). *Markov Chains and Stochastic Stability*. Springer.

Mohri, M. and A. Rostamizadeh (2008). "Stability bounds for non-iid processes". In: *NIPS*.

Montenegro, R. and P. Tetali (2006). *Mathematical Aspects of Mixing Times in Markov Chains*. Now Publishers.

Paulin, D. (2015). "Concentration inequalities for Markov chains by Marton couplings and spectral methods". In: *Electronic Journal of Probability* 20, pp. 1–32.

Seneta, E. (1993). "Sensitivity of finite Markov chains under perturbation". In: *Statistics & Probability Letters* 17, pp. 163–168.

Steinwart, I. and A. Christmann (2009). "Fast Learning from Non-i.i.d. Observations". In: *NIPS*.

Steinwart, I., D. Hush, and C. Scovel (2009). "Learning from dependent observations". In: *Journal of Multivariate Analysis* 100.1, pp. 175–194.

Stewart, G. W. and J. Sun (1990). *Matrix perturbation theory*. Boston: Academic Press. ISBN: 0126702306.

Sutton, R. S. and A. G. Barto (1998). *Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning)*. A Bradford Book. ISBN: 9780262193986.

Tropp, J. (2015). "An Introduction to Matrix Concentration Inequalities". In: *Foundations and Trends in Machine Learning*.

Yu, B. (1994). "Rates of convergence for empirical processes of stationary mixing sequences". In: *The Annals of Probability* 22.1, pp. 94–116.

Computer Science Department, Columbia University, New York, NY 10027
E-mail address: djhsu@cs.columbia.edu

Ben-Gurion University
E-mail address: karyeh@cs.bgu.ac.il

Department of Mathematics, University of Oregon, Eugene, OR 97403-1220
E-mail address: dlevin@uoregon.edu

Microsoft Research
E-mail address: peres@microsoft.com

University of Alberta
E-mail address: csaba.szepesvari@gmail.com