

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

BY DANIEL HSU*, ARYEH KONTOROVICH†, DAVID A. LEVIN‡, YUVAL PERES§, CSABA SZEPESVÁRI¶ AND GEOFFREY WOLFER†

Columbia University*, Ben-Gurion University†, University of Oregon‡, Microsoft Research§ and University of Alberta/DeepMind¶

The spectral gap γ_* of a finite, ergodic and reversible Markov chain is an important parameter measuring the asymptotic rate of convergence. In applications, the transition matrix \mathbf{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 γ_* from this data. Let π be the stationary distribution of \mathbf{P} , and $\pi_* = \min_x \pi(x)$. We show that if n is at least $\frac{1}{\gamma_* \pi_*}$ times a logarithmic correction, then γ_* can be estimated to within a multiplicative factor with high probability. When π is uniform on d states, this nearly matches a lower bound of $\frac{d}{\gamma_*}$ steps required for precise estimation of γ_* . 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 $t_{\text{relax}} = 1/\gamma_*$, 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 $(X_t)_{t=1,2,\dots}$ be an irreducible, aperiodic time-homogeneous Markov chain on a finite state space $[d] := \{1, 2, \dots, d\}$ with transition matrix \mathbf{P} . Under this assumption, the chain converges to its unique stationary distribution $\pi = (\pi_i)_{i=1}^d$ regardless of the initial state distribution \mathbf{q} :

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

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:

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

Received November 2017; revised July 2018.

MSC2010 subject classifications. 60J10, 62M05, 62M99.

Key words and phrases. Markov chains, mixing time, spectral gap, empirical confidence interval.

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

This problem is motivated by the numerous scientific applications and machine learning tasks in which the quantity of interest is the mean $\pi(f) = \sum_i \pi_i f(i)$ for some function f of the states of a Markov chain. This is the setting of the celebrated Markov chain Monte Carlo (MCMC) paradigm (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 γ_\star and stationary distribution minorized by $\pi_\star := \min_{i \in [d]} \pi_i$. As is well known (see, e.g., Levin, Peres and Wilmer (2009), Theorems 12.3 and 12.4),

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

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

(1) In Section 3.1, we show that in some problems $n = \Omega(d/\gamma_\star + 1/\pi_\star)$ observations are necessary for any procedure to guarantee constant multiplicative accuracy in estimating γ_\star (Theorems 3.1 and 3.2). Essentially, in some problems a *majority of the states* may need to be visited about $1/\gamma_\star$ 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 γ_\star , based on 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 π_\star . This establishes the feasibility of *estimating* the mixing time in this setting, and the dependence on π_\star and γ_\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 π_\star and γ_\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 π_\star and γ_\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, 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 $\pi_n(f) := \frac{1}{n} \sum_{t=1}^n f(X_t)$ converges almost surely to $\pi(f)$ (Meyn and Tweedie (1993)), while the central limit theorem tells us that as $n \rightarrow \infty$, the distribution of the deviation $\sqrt{n}(\pi_n(f) - \pi(f))$ will be normal with mean zero and asymptotic variance $\lim_{n \rightarrow \infty} n \text{Var}(\pi_n(f))$ (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 nonasymptotic behavior, which is often needed for the prudent evaluation of a method or even its algorithmic design (Flegal and Jones (2011), Gyori and Paulin (2014), Kontoyiannis, Lastras-Montaño and Meyn (2006)). To address this need, numerous works have developed Chernoff-type bounds on $\Pr(|\pi_n(f) - \pi(f)| > \epsilon)$, thus providing valuable tools for nonasymptotic probabilistic analysis (Gillman (1998), Kontorovich and Weiss (2014), Kontoyiannis, Lastras-Montaño and Meyn (2006), León and Perron (2004), Paulin (2015)). These probability bounds are larger than the corresponding bounds for independent and identically distributed (i.i.d.) data due to the temporal dependence; intuitively, for the Markov chain to yield a fresh draw $X_{t'}$ that behaves as if it was independent of X_t , one must wait $\Theta(t_{\text{mix}})$ time steps. Note that the bounds generally depend on distribution-specific properties of the Markov chain (e.g., \mathbf{P} , t_{mix} , γ_*), which are often unknown a priori in practice. Consequently, much effort has been put toward estimating these unknown quantities, especially in the context of MCMC diagnostics, in order to provide data-dependent assessments of estimation accuracy (e.g., Atchadé (2016), Flegal and Jones (2011), Garren and Smith (2000), Gyori and Paulin (2014), Jones and Hobert (2001)). 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 (e.g.) 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., Gamarnik (2003), Karandikar and Vidyasagar (2002), Mohri and Rostamizadeh (2008), Steinwart and Christmann (2009), Steinwart, Hush and Scovel (2009), Yu (1994)), 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 et al. (2000, 2013), 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. *Notation.* 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 nonnegative 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 a natural logarithm, and we use $\log(\cdot)$ for a logarithm with an arbitrary constant base > 1 that does not matter in the given context. Boldface symbols are used for vectors and matrices (e.g., \mathbf{v}, \mathbf{M}), and their entries are referenced by subindexing (e.g., $v_i, M_{i,j}$). For a vector \mathbf{v} , $\|\mathbf{v}\|$ denotes its Euclidean norm; for a matrix \mathbf{M} , $\|\mathbf{M}\|$ denotes its spectral norm. We use $\text{Diag}(\mathbf{v})$ to denote the diagonal matrix whose (i, i) th entry is v_i . The probability simplex is denoted by $\Delta^{d-1} = \{\mathbf{p} \in [0, 1]^d : \sum_{i=1}^d p_i = 1\}$, and we regard vectors in Δ^{d-1} as row vectors.

2.2. *Setting.* Let $\mathbf{P} \in (\Delta^{d-1})^d \subset [0, 1]^{d \times d}$ be a $d \times d$ row-stochastic matrix for an ergodic (i.e., irreducible and aperiodic) Markov chain. This implies there is a unique stationary distribution $\boldsymbol{\pi} \in \Delta^{d-1}$ with $\pi_i > 0$ for all $i \in [d]$ (Levin, Peres and Wilmer (2009), Corollary 1.17). We also assume that \mathbf{P} is *reversible* (with respect to $\boldsymbol{\pi}$):

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

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

Define the matrices

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

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

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

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

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

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

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

The goal is to estimate π_\star and γ_\star from the length n sample path $(X_t)_{t \in [n]}$, and also to construct confidence intervals that π_\star and γ_\star with high probability; in particular, the construction of the intervals should be fully empirical and not depend on any unobservable quantities, including π_\star and γ_\star themselves. As mentioned in the Introduction, it is well known that the *mixing time* of the Markov chain t_{mix} (defined in equation (1)) is bounded in terms of π_\star and γ_\star , as shown in equation (2). Moreover, convergence rates for empirical processes on Markov chain sequences are also often given in terms of mixing coefficients that can ultimately be bounded in terms of π_\star and γ_\star (as we will show in the proof of our first result). Therefore, valid confidence intervals for π_\star and γ_\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 γ_\star . By equation (2), the multiplicative accuracy lower bound for γ_\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 γ_\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 $n \leq 1/(4\bar{\pi})$ from a Markov chain starting from any desired initial state distribution. There exists a two-state ergodic and reversible Markov chain distribution with spectral gap $\gamma_\star \geq 1/2$ and minimum stationary probability $\pi_\star \geq \bar{\pi}$ such that*

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

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

THEOREM 3.2. *There is an absolute constant $c > 0$ such that the following holds. Pick any positive integer $d \geq 10$ 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 $n < cd/\bar{\gamma}_\star$ 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 $\pi_\star \geq 1/(2d)$ such that*

$$\Pr[|\hat{\gamma}_\star - \gamma_\star| \geq \bar{\gamma}_\star/2] \geq 1/50.$$

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 γ_\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 equation (2) to trap the mixing time.

Define the random matrix $\widehat{\mathbf{M}} \in [0, 1]^{d \times d}$ 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

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

and let

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

be its symmetrized version. Let $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \dots \geq \hat{\lambda}_d$ be the eigenvalues of $\text{Sym}(\widehat{\mathbf{L}})$. Our estimator of the minimum stationary probability π_\star is $\hat{\pi}_\star := \min_{i \in [d]} \hat{\pi}_i$, and

our estimator of the spectral gap γ_\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 $\hat{\pi}$ 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 γ_\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 $1 - \delta$,*

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

and

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

Theorem 3.3 implies that the sequence lengths sufficient to estimate π_\star and γ_\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{M} and $\hat{\pi}$ 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) nonstationary Markov chain, we cannot use standard large deviation bounds; moreover applying Chernoff-type bounds for Markov chains to each entry of \widehat{M} would result in a significantly worse sequence length requirement, roughly a factor of d larger. Instead, we adapt probability tail bounds for sums of independent random matrices (Tropp (2015)) to our non-i.i.d. 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 $\pi^{(1)}$. 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 equation (5) to show that in fact, to obtain any prescribed multiplicative accuracy, $\tilde{O}(1/(\pi_\star \gamma_\star))$ steps suffice to estimate γ_\star . The idea is to apply the estimator $\hat{\gamma}_\star$ from equation (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 γ_\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 equation (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 γ_\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 γ_\star is not known (of course), we use a doubling trick to find a suitable value of a .

The estimator is defined as follows. 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 γ_\star^{-1} , π_\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 γ_\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 equation (37). The proof of Theorem 3.4 is given in Section 7. The result shows that to estimate both π_\star and γ_\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 γ_\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 γ_\star and π_\star . This is because the deviation terms themselves depend inversely both on γ_\star and π_\star , and hence can never rule out 0 (or an arbitrarily small positive value) as a possibility for γ_\star or π_\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 $\hat{\pi}_\star$.

¹Using Theorem 3.3, it is possible to trap γ_\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.

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

We adopt a different strategy for estimating π based on the group inverse $\hat{\mathbf{A}}^\#$ of $\hat{\mathbf{A}} = \mathbf{I} - \hat{\mathbf{P}}$. Recall that the group inverse of a square matrix \mathbf{M} , a special case of the Drazin inverse, is the unique matrix $\mathbf{M}^\#$ satisfying $\mathbf{M}\mathbf{M}^\#\mathbf{M} = \mathbf{M}$, $\mathbf{M}^\#\mathbf{M}\mathbf{M}^\# = \mathbf{M}^\#$ and $\mathbf{M}^\#\mathbf{M} = \mathbf{M}\mathbf{M}^\#$. In our case, where $\hat{\mathbf{P}}$ defines an ergodic chain (due to the use of the smoothed estimates), the group inverse $\hat{\mathbf{A}}^\#$ can be computed at the cost of inverting an $(d-1) \times (d-1)$ matrix (Meyer (1975), Theorem 5.2). Finally, given $\hat{\mathbf{A}}^\#$, the unique stationary distribution $\hat{\pi}$ of $\hat{\mathbf{P}}$ can be read out from the last row of $\hat{\mathbf{A}}^\#$ (Meyer (1975), Theorem 5.3), and $\hat{\pi}$ can be regarded as an estimate of the stationary distribution π of \mathbf{P} . This way of estimating π decouples the bound on the estimation error from the mixing time. Indeed, the sensitivity of $\hat{\pi}$ to $\hat{\mathbf{P}}$ is also controlled by the group inverse through

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

A perturbation bound from Cho and Meyer (2001) grants the inequality

$$(6) \quad \max_{i \in [d]} |\hat{\pi}_i - \pi_i| \leq \hat{\kappa} \max_{i \in [d]} \sum_{j \in [d]} |\hat{P}_{i,j} - P_{i,j}|.$$

(In fact, with $\mathbf{A} := \mathbf{I} - \mathbf{P}$ and

$$\kappa := \frac{1}{2} \max\{A_{j,j}^\# - \min\{A_{i,j}^\# : i \in [d]\} : j \in [d]\},$$

the inequality in (6) also holds with κ in place of $\hat{\kappa}$. The quantity κ appears in our main result below.)

Algorithm 1 Confidence intervals

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

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

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

$$N_{i,j} := |\{t \in [n - 1] : (X_t, X_{t+1}) = (i, j)\}|, \quad (i, j) \in [d]^2;$$

$$\widehat{P}_{i,j} := \frac{N_{i,j} + 1/d}{N_i + 1}, \quad (i, j) \in [d]^2.$$

- 2: Let $\widehat{A}^\#$ be the group inverse of $\widehat{A} := \mathbf{I} - \widehat{P}$.
- 3: Let $\widehat{\pi} \in \Delta^{d-1}$ be the unique stationary distribution for \widehat{P} .
- 4: Compute eigenvalues $\widehat{\lambda}_1 \geq \widehat{\lambda}_2 \geq \dots \geq \widehat{\lambda}_d$ of $\text{Sym}(\widehat{L})$, where

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

5: Spectral gap estimate:

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

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

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

7: Relative sensitivity of π :

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

8: Bounds for $\max_{i \in [d]} |\widehat{\pi}_i - \pi_i|$ and $\max \cup_{i \in [d]} \{|\sqrt{\pi_i/\widehat{\pi}_i} - 1|, |\sqrt{\widehat{\pi}_i/\pi_i} - 1|\}$:

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

9: Bounds for $|\widehat{\gamma}_\star - \gamma_\star|$:

$$\widehat{w} := 2\widehat{\rho} + \widehat{\rho}^2 + (1 + 2\widehat{\rho} + \widehat{\rho}^2) \|\widetilde{B}\|,$$

where

$$\widetilde{B}_{i,j} := \sqrt{\frac{\widehat{\pi}_i}{\widehat{\pi}_j}} \widehat{B}_{i,j}.$$

We can now follow the strategy based on estimating L alluded to above. Using $\hat{\pi}$ and \hat{P} , we construct the plug-in estimate \hat{L} of 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 $\hat{\pi}$ and π , viewing P as the perturbation of \hat{P} ; and also to relate $\hat{\gamma}_\star$ and γ_\star , viewing L as a perturbation of $\text{Sym}(\hat{L})$. 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 P .

The most computationally expensive step in Algorithm 1 is the computation of the group inverse $\hat{A}^\#$, 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 (nonrandom) positive sequence $(\varepsilon_{\theta,n})_{n \geq 1}$ parameterized by θ , we say “ $Y_n = O(\varepsilon_{\theta,n})$ holds almost surely as $n \rightarrow \infty$ ” if there is some universal constant $C > 0$ such that for all θ , $\limsup_{n \rightarrow \infty} Y_n / \varepsilon_{\theta,n} \leq C$ holds almost surely.

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, π the unique stationary distribution, and $\pi_\star > 0$ the minimum stationary probability. Then, on an event of probability at least $1 - \delta$,*

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

Moreover,

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

almost surely as $n \rightarrow \infty$.

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 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 (Cho and Meyer (2001), Meyer (1975)) and eigenvalue perturbation theory (Stewart and Sun (1990)) are used to validate the empirical bounds on π_i and γ_\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 γ_\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 P, π, 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

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

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

It is easy to combine Theorems 3.3 and 4.1 to yield intervals whose widths shrink at least as fast as both the nonempirical intervals from Theorem 3.3 and the empirical intervals from Theorem 4.1. Specifically, determine lower bounds on π_\star and γ_\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 π_\star and γ_\star in the deviation bounds in equation (5) from Theorem 3.3. This yields a new interval centered around the estimate of γ_\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 γ_\star , and for sufficiently large n ; the width shrinks at the rate given in equation (5). We can similarly construct an empirical confidence interval for π_\star using equation (4), which is valid on the same $1 - 2\delta$ probability event.² 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 \hat{U} and \hat{V} described above for π_\star and γ_\star , respectively, satisfy $\pi_\star \in \hat{U}$ and $\gamma_\star \in \hat{V}$ with probability at least $1 - 2\delta$. Furthermore, $|\hat{U}| = O\left(\sqrt{\frac{\pi_\star \log \frac{d}{\pi_\star \delta}}{\gamma_\star n}}\right)$ and $|\hat{V}| = O\left(\min\left\{\sqrt{\frac{\log \frac{d}{\delta} \cdot \log(n)}{\pi_\star \gamma_\star n}}, \hat{w}\right\}\right)$ almost surely as $n \rightarrow \infty$, where \hat{w} is the width from Algorithm 1.*

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

- 1: **if** $n = 2^k$ for an integer k , **then**
- 2: 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 π_\star and γ_\star .

²For the π_\star interval, we only plug-in lower bounds on π_\star and γ_\star only where these quantities appear as $1/\pi_\star$ and $1/\gamma_\star$ in equation (4). It is then possible to “solve” for observable bounds on π_\star . See Section 9 for details.

- 3: Stop if, for each interval, the interval width divided by the lower bound on estimated quantity falls below ϵ .
- 4: **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:

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

Each Markov chain is ergodic and reversible; their stationary distributions are, respectively, $\pi^{(1)} = (1 - \bar{\pi}, \bar{\pi})$ and $\pi^{(2)} = (1/(1 + 2\bar{\pi}), 2\bar{\pi}/(1 + 2\bar{\pi}))$. We have $\pi_\star \geq \bar{\pi}$ 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.* For $d \geq 10$, and for simplicity of the analysis d even (a slight modification of the proof covers the odd case), we consider d -state Markov chains of the following form:

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

for some $\epsilon_1, \epsilon_2, \dots, \epsilon_d \in (0, 1)$. Such a chain is ergodic and reversible, and its unique stationary distribution π satisfies

$$\pi_i = \frac{1/\epsilon_i}{\sum_{j=1}^d 1/\epsilon_j}.$$

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

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

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

Following a half-covering argument similar to [Wolfer and Kontorovich \(2019b\)](#), in order to guarantee $|\hat{\gamma}_\star - \gamma_\star| < \bar{\gamma}/2$, it must be possible to distinguish $\mathbf{P}^{(0)}$ from $\mathbf{P}^{(i)}$, where $i \in [d]$ is drawn uniformly at random. But $\mathbf{P}^{(0)}$ is identical to $\mathbf{P}^{(i)}$ except for the transition probabilities from state i . Therefore, regardless of the initial state, and from symmetry considerations, the sample path must visit at least half of the states in order to distinguish $\mathbf{P}^{(0)}$ from $\mathbf{P}^{(i)}$, $i \in [d]$ with constant probability of success. For any of the $d + 1$ Markov chains above, the earliest time in which a sample path visits at least $d/2$ states stochastically dominates a generalized coupon half-collection time $T_{1/2} = 1 + \sum_{i=1}^{d/2-1} T_i$, where T_i is the number of steps required to see the $(i + 1)$ th distinct state in the sample path beyond the first i . The random variables $T_1, T_2, \dots, T_{d/2-1}$ are independent, and are geometrically distributed, $T_i \sim \text{Geom}(\varepsilon - (i - 1)\varepsilon/(d - 1))$. We have that

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

Therefore,

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

where $H_{d-1} = 1 + 1/2 + 1/3 + \dots + 1/(d - 1)$. For $d \geq 10$, it is the case that $H_{d-1} - H_{d/2} \geq 2/5$, and by the Paley–Zygmund inequality,

$$\Pr\left(T_{1/2} > \frac{1}{3}\mathbb{E}[T_{1/2}]\right) \geq \frac{1}{1 + \frac{\text{var}(T_{1/2})}{(1-1/3)^2\mathbb{E}[T_{1/2}]^2}} \geq \frac{1}{1 + \frac{(\frac{d-1}{\varepsilon})^2 \frac{\pi^2}{6}}{(4/9)(\frac{d-1}{\varepsilon}(2/5))^2}} \geq \frac{1}{25}.$$

Since $n < cd/\bar{\gamma} \leq \mathbb{E}[T_{1/2}]/3$ (for an appropriate absolute constant c), with probability at least $1/25$, the sample path does not visit at least half of the d states. By symmetry, with probability $1/2$ over the draw of i , the state that differs was not visited, the distribution of trajectories of length n conditioned on this event are identical, and one cannot do better than to choose an hypothesis at random.

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 equation (4) in Theorem 3.3.

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

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

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

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

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

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

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

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

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

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

Now assume this $1 - \delta$ probability event holds; it remains to prove that equation (9) also holds in this event. Without loss of generality, we assume that $\pi_\star = \pi_1 \leq \pi_2 \leq \dots \leq \pi_d$. Let $j \in [d]$ be such that $\hat{\pi}_\star = \hat{\pi}_j$. By equation (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,

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

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

6.2. *Accuracy of $\hat{\gamma}_\star$.* Let us now turn to proving equation (5), that is, the bound on the error of the spectral gap estimate $\hat{\gamma}_\star$. Recall the definitions of $\widehat{\mathbf{L}}$ and $\text{Sym}(\widehat{\mathbf{L}})$:

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

The accuracy of $\hat{\gamma}_\star$ is based on the accuracy of $\text{Sym}(\widehat{\mathbf{L}})$ in approximating \mathbf{L} via Weyl’s inequality:

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

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

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

Therefore, we can deduce equation (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*

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

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

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

where

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

We briefly describe how to obtain the bound on $|\hat{\gamma}_\star - \gamma_\star|$ that appears in equation (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 equation (10), so by Lemma 6.2, we have $|\hat{\gamma}_\star - \gamma_\star| \leq \|\widehat{\mathbf{L}} - \mathbf{L}\| \leq C(\sqrt{\varepsilon} + \varepsilon + \varepsilon^2) \leq C'\sqrt{\varepsilon}$.

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

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

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

where

$$\boldsymbol{\mathcal{E}}_\pi := \text{Diag}(\hat{\boldsymbol{\pi}})^{-1/2} \text{Diag}(\boldsymbol{\pi})^{1/2} - \mathbf{I} \quad \text{and}$$

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

Therefore,

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

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

$$(11) \quad \|\widehat{L} - L\| \leq \|\mathcal{E}_M\| + \|\mathcal{E}_M\|^2 + 3\|\mathcal{E}_\pi\|.$$

6.3. *A bound on $\|\mathcal{E}_\pi\|$.* Since \mathcal{E}_π is diagonal,

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

Assume that

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

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 equation (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 $\hat{\pi}$ 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 equation (12), then on this $1 - \delta$ probability event from Lemma 6.1, $\hat{\pi}$ is positive valued and

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

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

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

6.4.1. *Bounding $\|\mathbb{E}(\mathcal{E}_M)\|$: The price of a nonstationary 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*

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

Then

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

An immediate corollary of this result is that

$$(16) \quad \|\text{Diag}(\boldsymbol{\pi}^{(t)}) \text{Diag}(\boldsymbol{\pi})^{-1} - \mathbf{I}\| \leq \frac{(1 - \gamma_\star)^{t-1}}{\pi_\star}.$$

Now note that

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

and thus

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

Combining this, $\|\mathbf{L}\| \leq 1$ and equation (16), we get

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

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

THEOREM 6.4. *If*

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

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

$$\begin{aligned} & \|\mathcal{E}_M - \mathbb{E}(\mathcal{E}_M)\| \\ &= \|\text{Diag}(\boldsymbol{\pi})^{-1/2}(\widehat{\mathbf{M}} - \mathbb{E}[\widehat{\mathbf{M}}])\text{Diag}(\boldsymbol{\pi})^{-1/2}\| \\ (18) \quad & \leq \frac{4\lceil \frac{1}{\gamma_\star} \ln \frac{2(n-2)}{\pi_\star \delta} \rceil}{\pi_\star(n-1)} + \sqrt{\frac{4(d_P + 2) \ln \frac{4d}{\delta}}{\mu}} + \frac{2(\frac{1}{\pi_\star} + 2) \ln \frac{4d}{\delta}}{3\mu}, \end{aligned}$$

where

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

and μ , defined below, satisfies

$$\mu \geq \frac{n-1}{2(1 + \frac{1}{\gamma_\star} \ln \frac{2(n-2)}{\pi_\star \delta})} - 2.$$

The proof of Theorem 6.4 proceeds in several steps. First, we note that the matrix $\widehat{\mathbf{M}} - \mathbb{E}(\widehat{\mathbf{M}})$ is defined as a sum of dependent centered random matrices. We will use the blocking technique of 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 be bounded with the help of a Bernstein-type matrix tail inequality due to Tropp (2015), stated in Lemma 6.5. There is a tradeoff in choosing the block size a : larger blocks allow for faster decay of dependencies but induce a smaller effective sample size; this tradeoff is optimized in our choice of a in (27).

PROOF OF THEOREM 6.4. We divide $[n - 1]$ into contiguous blocks of time steps; each has size $a \leq n/3$ except possibly the first block, which has size between a and $2a - 1$. Formally, let $a' := a + ((n - 1) \bmod a) \leq 2a - 1$, and define

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

for $s = 1, 2, \dots$. Let μ_H (resp., μ_T) be the number of nonempty 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 $\bigcup_s H_s$ (resp., $\bigcup_s T_s$). We have

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

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

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

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

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

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

so

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

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

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

(note the $+1$ in the upper limit of t), so \mathbf{Y}_{s+1} is a time steps ahead of \mathbf{Y}_s . When a is sufficiently large, we will be able to effectively treat the random matrices \mathbf{Y}_s

as if they were independent. In the sequel, we shall always assume that the block length a satisfies

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

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

Define

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

Observe that

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

so

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

Define

$$(22) \quad \mathbf{Z}_s := \text{Diag}(\boldsymbol{\pi})^{-1/2} (\mathbf{Y}_s - \mathbb{E}(\mathbf{Y}_s)) \text{Diag}(\boldsymbol{\pi})^{-1/2}.$$

We apply the matrix tail inequality in Lemma 6.5 to the average of *independent* copies of the \mathbf{Z}_s 's. More precisely, we will apply the tail inequality to independent copies $\tilde{\mathbf{Z}}_s, s \in [\mu_H]$ of the random variables \mathbf{Z}_s and then relate the average of $\tilde{\mathbf{Z}}_s$ to that of \mathbf{Z}_s . To apply Lemma 6.5, it suffices to bound the spectral norms of \mathbf{Z}_s (almost surely), $\mathbb{E}(\mathbf{Z}_s \mathbf{Z}_s^\top)$ and $\mathbb{E}(\mathbf{Z}_s^\top \mathbf{Z}_s)$. These are furnished by Lemmas 6.6 and 6.7 below:

$$(23) \quad \|\mathbf{Z}_s\| \leq \frac{1}{\pi_\star} + 2,$$

$$(24) \quad \|\mathbb{E}(\mathbf{Z}_s \mathbf{Z}_s^\top)\|, \|\mathbb{E}(\mathbf{Z}_s^\top \mathbf{Z}_s)\| \leq 2 \max_{i \in [d]} \left(\sum_{j=1}^d \frac{P_{i,j}}{\pi_j} \right) + 4.$$

Armed with these estimates, we are ready to invoke Lemma 6.5. Let $\tilde{\mathbf{Z}}_s$ for $s \in [\mu_H]$ be independent copies of \mathbf{Z}_s for $s \in [\mu_H]$. Applying Lemma 6.5 to the average of these random matrices, we have

$$(25) \quad \mathbb{P}\left(\left\| \frac{1}{\mu_H} \sum_{s=1}^{\mu_H} \tilde{\mathbf{Z}}_s \right\| > \sqrt{\frac{4(d\mathbf{P} + 2) \ln \frac{4d}{\delta}}{\mu_H}} + \frac{2(\frac{1}{\pi_\star} + 2) \ln \frac{4d}{\delta}}{3\mu_H}\right) \leq \delta,$$

where

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

To bound the probability that $\|\sum_{s=1}^{\mu_H} \mathbf{Z}_s/\mu_H\|$ 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 \mathbf{Z}_s . Let \mathbb{P} denote the joint distribution of $(X^{(H_s)} : s \in [\mu_H])$; let \mathbb{P}_s be its marginal over $X^{(H_s)}$, and let $\mathbb{P}_{1:s+1}$ be its marginal over $(X^{(H_1)}, X^{(H_2)}, \dots, X^{(H_{s+1})})$. Let $\tilde{\mathbb{P}}$ be the product distribution formed from the marginals $\mathbb{P}_1, \mathbb{P}_2, \dots, \mathbb{P}_{\mu_H}$, so $\tilde{\mathbb{P}}$ governs the joint distribution of $(\tilde{\mathbf{Z}}_s : s \in [\mu_H])$. The result from [Yu \(1994\)](#), Corollary 2.7, implies for any event E ,

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

where

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

Here, $\|\cdot\|_{\text{tv}}$ denotes the total variation norm. The β -mixing coefficient of the stochastic process $\{X^{(H_s)}\}_{s \in [\mu_H]}$ ([Bradley \(2005\)](#)) can be seen to give an upper bound on $\beta(\mathbb{P})$: In the definition of $\beta(\mathbb{P})$, hidden in the total variation norm, the maximum is taken over a projection of the measure whose total variation norm is used in the definition of β -mixing. The cited result of [Yu](#) implies that the bound from equation (25) for $\|\sum_{s=1}^{\mu_H} \tilde{\mathbf{Z}}_s/\mu_H\|$ also holds for $\|\sum_{s=1}^{\mu_H} \mathbf{Z}_s/\mu_H\|$, except the probability bound increases from δ to $\delta + (\mu_H - 1)\beta(\mathbb{P})$:

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

By the triangle inequality,

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

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

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

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 $i \in [d]$,

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

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

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

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 equation (21).

We return to the decomposition from equation (19). We apply equation (26) to both the H_s blocks and the T_s blocks, and combine with equation (20) to obtain the following probabilistic bound. Pick any $\delta \in (0, 1)$, let the block length be

$$(27) \quad a := \lceil a_\delta \rceil = \left\lceil \frac{1}{\gamma_\star} \ln \frac{2(n - 2)}{\pi_\star \delta} \right\rceil,$$

so

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

Thus, for $n \geq 7 + \frac{6}{\gamma_\star} \ln \frac{2(n-2)}{\pi_\star \delta} \geq 3a$, the bound in (18) holds. \square

LEMMA 6.5 (Matrix Bernstein inequality (Tropp (2015), Theorem 6.1.1)). *Let $\mathbf{Q}_1, \mathbf{Q}_2, \dots, \mathbf{Q}_m$ be a sequence of independent, random $d_1 \times d_2$ matrices. Assume that $\mathbb{E}(\mathbf{Q}_i) = \mathbf{0}$ and $\|\mathbf{Q}_i\| \leq R$ for each $1 \leq i \leq m$. Let $\mathbf{S} = \sum_{i=1}^m \mathbf{Q}_i$ and let*

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

Then, for all $t \geq 0$,

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

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

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

LEMMA 6.6 (Range bound). *Let \mathbf{Z}_s be as defined in (22). Then*

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

PROOF. By the triangle inequality,

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

For the first term, we have

$$(28) \quad \|\text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbf{Y}_s \text{Diag}(\boldsymbol{\pi})^{-1/2}\| \leq \frac{1}{\pi_\star}.$$

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

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

Then, using equation (16),

$$(29) \quad \|\text{Diag}(\boldsymbol{\pi}^{(H_s)}) \text{Diag}(\boldsymbol{\pi})^{-1} - \mathbf{I}\| \leq \frac{(1 - \gamma_\star)^{a'+2(s-1)a}}{\pi_\star} \leq \frac{(1 - \gamma_\star)^a}{\pi_\star} \leq 1,$$

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

$$(30) \quad \|\text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbb{E}(\mathbf{Y}_s) \text{Diag}(\boldsymbol{\pi})^{-1/2}\| \leq 2$$

by the triangle inequality. Therefore, together with equation (28), we obtain the claimed range bound. \square

LEMMA 6.7 (Variance bound). *Let \mathbf{Z}_s be as defined in (22). Then*

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

PROOF. Observe that

$$\begin{aligned} &\mathbb{E}(\mathbf{Z}_s \mathbf{Z}_s^\top) \\ (31) \quad &= \frac{1}{a^2} \sum_{t \in H_s} \mathbb{E}(\text{Diag}(\boldsymbol{\pi})^{-1/2} \mathbf{e}_{X_t} \mathbf{e}_{X_{t+1}}^\top \\ &\quad \times \text{Diag}(\boldsymbol{\pi})^{-1} \mathbf{e}_{X_{t+1}} \mathbf{e}_{X_t}^\top \text{Diag}(\boldsymbol{\pi})^{-1/2}) \end{aligned}$$

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

The first sum, equation (31), easily simplifies to the diagonal matrix

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

For the second sum, equation (32), a symmetric matrix, consider

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

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

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

which simplifies to

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

The expectation is the same as that for the first term, equation (31).

Finally, the spectral norm of the third term, equation (33), is bounded using equation (30):

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

Therefore, by the triangle inequality, the bound $\pi_i^{(H)}/\pi_i \leq 2$ from equation (29), and simplifications,

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

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

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

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

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

again using the bound $\pi_i^{(H)}/\pi_i \leq 2$ from equation (29). \square

6.4.3. *The bound on $\|\mathcal{E}_M\|$.* Combining the probabilistic bound from above with the bound on the bias from equation (17), we obtain the following. Assuming the condition on n from equation (6.4), with probability at least $1 - 4\delta$,

$$\begin{aligned} (34) \quad \|\mathcal{E}_M\| &\leq \frac{1}{(n-1)\gamma_\star \pi_\star} + \frac{4 \lceil \frac{1}{\gamma_\star} \ln \frac{2(n-2)}{\pi_\star \delta} \rceil}{\pi_\star (n-1)} \\ &\quad + \sqrt{\frac{4(d_P + 2) \ln \frac{4d}{\delta}}{\mu}} + \frac{2(\frac{1}{\pi_\star} + 2) \ln \frac{4d}{\delta}}{3\mu} \leq C'(\sqrt{\varepsilon} + \varepsilon), \end{aligned}$$

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

6.5. *Overall error bound.* Observe that the assumption on the sequence length in equation (10) implies the conditions in equation (12) and equation (6.4) for a suitable choice of $C > 0$. With this assumption, there is a $1 - 5\delta$ probability event in which equations (8), (9), (34) hold; in particular, we have the bound on $\|\mathcal{E}_M\|$ from equation (34). In this event, the bound on $\|\mathcal{E}_\pi\|$ in equation (13) also holds, and the claimed bound on $\|\hat{\mathbf{L}} - \mathbf{L}\|$ follows from combining the bound in equation (11) with the bounds on $\|\mathcal{E}_\pi\|$ and $\|\mathcal{E}_M\|$:

$$\begin{aligned} \|\hat{\mathbf{L}} - \mathbf{L}\| &\leq \|\mathcal{E}_M\| + \|\mathcal{E}_M\|^2 + 3\|\mathcal{E}_\pi\| \\ &\leq 4C'(\sqrt{\varepsilon} + \varepsilon) + C'^2(\sqrt{\varepsilon} + \varepsilon)^2 \leq C(\sqrt{\varepsilon} + \varepsilon + \varepsilon^2), \end{aligned}$$

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 equation (5). Observe that

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

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

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

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

Define $K_{\gamma_\star} := \lceil \log_2(1/\gamma_\star) \rceil$ 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 $G(\varepsilon)$,

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

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

$$(35) \quad |\hat{\gamma}_\star(A) - \gamma_\star(A)| \leq \varepsilon.$$

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 equation (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 \mathbf{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.) \square

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 (35) holds on G .

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

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

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

We now prove Theorem 3.4.

PROOF OF THEOREM 3.4. Let

$$(36) \quad n_0(\varepsilon; \delta, \gamma_\star, \pi_\star) = n_0(\varepsilon) := \frac{\mathcal{L}}{\pi_\star \gamma_\star \varepsilon^2},$$

where

$$(37) \quad \mathcal{L} := 3 \cdot (16\sqrt{2})^2 \cdot \left(\log \frac{d(\lfloor \log_2(1/\gamma_\star) \rfloor + 1)}{\delta} \right) \cdot \left(\log \frac{3 \cdot (16\sqrt{2})^2 \cdot C^2(\lfloor \log_2(1/\gamma_\star) \rfloor + 1)}{\varepsilon^2 \pi_\star^2 \gamma_\star \delta} \right),$$

and C is the constant in equation (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 $[0.29, 0.55]$,

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

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

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

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

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

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

8.1. *Estimators for π and γ_\star .* The algorithm forms the estimator $\hat{\mathbf{P}}$ of \mathbf{P} using Laplace smoothing:

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

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 $\hat{\mathbf{P}}$ are positive, and hence $\hat{\mathbf{P}}$ is a transition probability matrix for an ergodic Markov chain. We let $\hat{\pi}$ be the unique stationary distribution for $\hat{\mathbf{P}}$. Using $\hat{\pi}$, we form an estimator $\widehat{\mathbf{L}}$ of \mathbf{L} using:

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

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

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

8.2. *Empirical bounds for \mathbf{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 $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \mathcal{F}_2 \subset \dots$; assume for some $b > 0$, $|Y_t| \leq b$ almost surely for all $t \in \mathbb{N}$. Let $V_k := \sum_{t=1}^k \mathbb{E}(Y_t^2 | \mathcal{F}_{t-1})$ and $S_k := \sum_{t=1}^k Y_t$ for $k \in \mathbb{N}$. For all $s, v > 0$,*

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

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

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

since $h(\sqrt{2z} + z/3)/z \geq 1$ for all $z > 0$ (see, e.g., 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 $c > 1$,*

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

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

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

where the final inequality uses Theorem 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 $k \in [n]$. \square

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

$$(38) \quad \begin{aligned} |\widehat{P}_{i,j} - P_{i,j}| &\leq \sqrt{\left(\frac{N_i}{N_i + d\alpha}\right) \frac{2cP_{i,j}(1 - P_{i,j})\tau_{n,\delta}}{N_i + d\alpha}} \\ &\quad + \frac{(4/3)\tau_{n,\delta}}{N_i + d\alpha} + \frac{|\alpha - d\alpha P_{i,j}|}{N_i + d\alpha}, \end{aligned}$$

where

$$(39) \quad \tau_{n,\delta} := \inf\{t \geq 0 : 2d^2(1 + \lceil \log_c(2n/t) \rceil_+)e^{-t} \leq \delta\} = O\left(\log\left(\frac{d \log(n)}{\delta}\right)\right).$$

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

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

so that

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

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

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

and for $t \geq 2$,

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

Therefore, by Corollary 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 $(i, j) \in [d]^2$. \square

Equation (38) 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 equation (38), we can numerically maximize $|\widehat{P}_{i,j} - P_{i,j}|$ subject to the constraint in equation (38) (viewing $P_{i,j}$ as the optimization variable). Let $B_{i,j}^*$ be this maximum value, so we have

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

in the same event where equation (38) 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

$$(40) \quad \widehat{B}_{i,j} := \left(\sqrt{\frac{c\tau_{n,\delta}}{2N_i}} + \sqrt{\frac{c\tau_{n,\delta}}{2N_i} + \sqrt{\frac{2c\widehat{P}_{i,j}(1 - \widehat{P}_{i,j})\tau_{n,\delta}}{N_i} + \frac{(4/3)\tau_{n,\delta} + |\alpha - d\alpha\widehat{P}_{i,j}|}{N_i}}} \right)^2$$

for each $(i, j) \in [d]^2$, where $\tau_{n,\delta}$ is defined in equation (39). We show in Lemma 8.4 that

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

again, in the same event where equation (38) holds. The observable bound in equation (40) is not too far from the unobservable bound in equation (38).

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 equation (40).*

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

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

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

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

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

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

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

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

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

Analogously, define

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

We now use the following perturbation bound from Cho and Meyer (2001), Section 3.3 (derived from Haviv and Van der Heyden (1984), 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*

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

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 $\sqrt{\widehat{\pi}_i/\pi_i}$ and $\sqrt{\pi_i/\widehat{\pi}_i}$.

LEMMA 8.6. *If $\max\{|\widehat{\pi}_i - \pi_i| : i \in [d]\} \leq \widehat{b}$, then*

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

PROOF. By Lemma 8.5, we have for each $i \in [d]$,

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

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

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

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

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

8.4. *Empirical bounds for L.* By Weyl’s inequality and the triangle inequality,

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

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

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

$$\|\hat{L} - L\| \leq 2\hat{\rho} + \hat{\rho}^2 + (1 + 2\hat{\rho} + \hat{\rho}^2)\|\tilde{B}\|,$$

where \tilde{B} is the matrix with

$$\tilde{B}_{i,j} := \sqrt{\frac{\hat{\pi}_i}{\hat{\pi}_j}} \hat{B}_{i,j}$$

and

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

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

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

where

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

Therefore,

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

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

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

It follows that the spectral norm of \mathcal{E}_P can be bounded as

$$\|\mathcal{E}_P\| \leq \|\tilde{\mathbf{B}}\|.$$

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

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

which can be bounded using Lemma 8.6. \square

This establishes the validity of the confidence interval for γ_\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 $n \rightarrow \infty$,

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

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

Further, since

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

we thus have

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

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*

(41)
$$\mathbb{E}_i \tau_j \leq 2\left(t_{\text{mix}} + 8\frac{t_{\text{relax}}}{\pi_j}\right).$$

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 $t = t_{\text{mix}} + 8t_{\text{relax}}/\pi_j$, then

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

Thus, $\Pr_i(\tau_j > tk) \leq 2^{-k}$, so equation (41) follows. \square

PROOF OF LEMMA 8.8. It is established by Cho and Meyer (2001) that

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

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

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

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

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

9. Proof of Theorem 4.2. Let $\hat{\pi}_{\star, \text{lb}}$ and $\hat{\gamma}_{\star, \text{lb}}$ be the lower bounds on π_\star and γ_\star , respectively, computed from Algorithm 1. Let $\hat{\pi}_\star$ and $\hat{\gamma}_\star$ be the estimates of π_\star and γ_\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 $1 - 2\delta$,

$$(42) \quad |\hat{\pi}_\star - \pi_\star| \leq C \left(\sqrt{\frac{\pi_\star \log \frac{d}{\hat{\pi}_{\star, \text{lb}} \delta}}{\hat{\gamma}_{\star, \text{lb}} n}} + \frac{\log \frac{d}{\hat{\pi}_{\star, \text{lb}} \delta}}{\hat{\gamma}_{\star, \text{lb}} n} \right)$$

and

$$(43) \quad |\hat{\gamma}_\star - \gamma_\star| \leq C \left(\sqrt{\frac{\log \frac{d}{\delta} \cdot \log \frac{n}{\hat{\pi}_{\star, \text{lb}} \delta}}{\hat{\pi}_{\star, \text{lb}} \hat{\gamma}_{\star, \text{lb}} n}} \right).$$

The bound on $|\hat{\gamma}_\star - \gamma_\star|$ in equation (43)—call it \hat{w}' —is fully observable, and hence yields a confidence interval for γ_\star . The bound on $|\hat{\pi}_\star - \pi_\star|$ in equation (42) depends on π_\star , but from it one can derive

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

using the approach from the proof of Lemma 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,

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

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

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

and

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

This completes the proof of Theorem 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, $I'_n = E_n(\delta) \cap \bigcup_{(\gamma, \pi) \in E_n(\delta)} I_n(\gamma, \pi, \delta)$ 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 theoretical 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 π 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. Both of these questions, as well as an extension to nonreversible Markov chains, are addressed in [Wolfer and Kontorovich \(2019a\)](#).

For practical purposes, there is much room for improvement. One major deficiency of this work is that the confidence interval constructed by our procedure (Algorithm 1) appears to be quite conservative. This likely stems from the reliance on perturbation bounds that guard against worst-case perturbations but are loose in the “typical” cases. To illustrate this, we generated sample paths of varying lengths n from the Markov chain $\mathbf{P}^{(d)}$ described in the proof of Theorem 3.2, with $d = 3$ and $\gamma_\star = 0.49$, each starting from state 1. The following plot shows the widths of the γ_\star -interval constructed by Algorithm 1 in Figure 1. However, the actual coverage of each interval appears to be far greater than $1 - \delta$: in 1000 independent trials, the intervals *never* failed to contain γ_\star . Moreover, on each of 1000 independent sample paths of length $n = 1024$, we computed the spectral gap of the maximum likelihood estimate of \mathbf{P} . In 950 of these cases, this empirical estimate was within a factor of 1 ± 0.115 of the true spectral gap, and an interval of length 0.109 around the empirical estimate contained the true spectral gap.

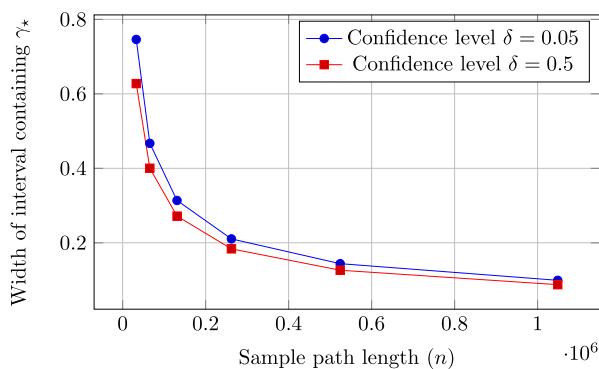


FIG. 1. Confidence interval width as function of sample path length.

Another deficiency is that the procedure is restricted to finite (and small) state spaces. 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. Another natural candidate is the exponential family of transition matrices considered by Hayashi and Watanabe (2016). 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. *Bernoulli* **22** 1808–1838. [MR3474834](#)
- AUDIBERT, J.-Y., MUNOS, R. and SZEPESVÁRI, C. (2009). Exploration-exploitation tradeoff using variance estimates in multi-armed bandits. *Theoret. Comput. Sci.* **410** 1876–1902. [MR2514714](#)
- BATU, T., FORTNOW, L., RUBINFELD, R., SMITH, W. D. and WHITE, P. (2000). Testing that distributions are close. In *41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000)* 259–269. IEEE Comput. Soc. Press, Los Alamitos, CA. [MR1931824](#)
- BATU, T., FORTNOW, L., RUBINFELD, R., SMITH, W. D. and WHITE, P. (2013). Testing closeness of discrete distributions. *J. ACM* **60** Art. 4, 25. [MR3033221](#)
- BENÍTEZ, J. and LIU, X. (2012). On the continuity of the group inverse. *Oper. Matrices* **6** 859–868. [MR3025040](#)
- BERNSTEIN, S. (1927). Sur l’extension du théorème limite du calcul des probabilités aux sommes de quantités dépendantes. *Math. Ann.* **97** 1–59. [MR1512353](#)
- BHATNAGAR, N., BOGDANOV, A. and MOSSEL, E. (2011). The computational complexity of estimating MCMC convergence time. In *Approximation, Randomization, and Combinatorial Optimization. Lecture Notes in Computer Science* **6845** 424–435. Springer, Heidelberg. [MR2863279](#)

- BHATTACHARYA, B. B. and VALIANT, G. (2015). Testing closeness with unequal sized samples. In *Advances in Neural Information Processing Systems* 28 (C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama and R. Garnett, eds.) 2611–2619. Curran Associates, Red Hook, NY.
- BOUSQUET, O., BOUCHERON, S. and LUGOSI, G. (2004). Introduction to statistical learning theory. *Lecture Notes in Artificial Intelligence* **3176** 169–207.
- BRADLEY, R. C. (2005). Basic properties of strong mixing conditions. A survey and some open questions. *Probab. Surv.* **2** 107–144. Update of, and a supplement to, the 1986 original. [MR2178042](#)
- CHO, G. E. and MEYER, C. D. (2001). Comparison of perturbation bounds for the stationary distribution of a Markov chain. *Linear Algebra Appl.* **335** 137–150. [MR1850819](#)
- FLEGAL, J. M. and JONES, G. L. (2011). Implementing MCMC: Estimating with confidence. In *Handbook of Markov Chain Monte Carlo. Chapman & Hall/CRC Handb. Mod. Stat. Methods* 175–197. CRC Press, Boca Raton, FL. [MR2858449](#)
- FREEDMAN, D. A. (1975). On tail probabilities for martingales. *Ann. Probab.* **3** 100–118. [MR0380971](#)
- GAMARNIK, D. (2003). Extension of the PAC framework to finite and countable Markov chains. *IEEE Trans. Inform. Theory* **49** 338–345. [MR1966715](#)
- GARREN, S. T. and SMITH, R. L. (2000). Estimating the second largest eigenvalue of a Markov transition matrix. *Bernoulli* **6** 215–242. [MR1748720](#)
- GILLMAN, D. (1998). A Chernoff bound for random walks on expander graphs. *SIAM J. Comput.* **27** 1203–1220. [MR1621958](#)
- GYORI, B. M. and PAULIN, D. (2014). Non-asymptotic confidence intervals for MCMC in practice. Available at [arXiv:1212.2016](#).
- HAVIV, M. and VAN DER HEYDEN, L. (1984). Perturbation bounds for the stationary probabilities of a finite Markov chain. *Adv. in Appl. Probab.* **16** 804–818. [MR0766780](#)
- HAYASHI, M. and WATANABE, S. (2016). Information geometry approach to parameter estimation in Markov chains. *Ann. Statist.* **44** 1495–1535. [MR3519931](#)
- HSU, D., KONTOROVICH, A. and SZEPESVÁRI, C. (2015). Mixing time estimation in reversible Markov chains from a single sample path. In *Advances in Neural Information Processing Systems* 28 (C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama and R. Garnett, eds.). Curran Associates, Red Hook, NY.
- JONES, G. L. and HOBERT, J. P. (2001). Honest exploration of intractable probability distributions via Markov chain Monte Carlo. *Statist. Sci.* **16** 312–334. [MR1888447](#)
- KARANDIKAR, R. L. and VIDYASAGAR, M. (2002). Rates of uniform convergence of empirical means with mixing processes. *Statist. Probab. Lett.* **58** 297–307. [MR1921877](#)
- KIPNIS, C. and VARADHAN, S. R. S. (1986). Central limit theorem for additive functionals of reversible Markov processes and applications to simple exclusions. *Comm. Math. Phys.* **104** 1–19. [MR0834478](#)
- KIRKLAND, S. J., NEUMANN, M. and SHADER, B. L. (1998). Applications of Paz’s inequality to perturbation bounds for Markov chains. *Linear Algebra Appl.* **268** 183–196. [MR1480340](#)
- KONTOROVICH, A. and WEISS, R. (2014). Uniform Chernoff and Dvoretzky–Kiefer–Wolfowitz-type inequalities for Markov chains and related processes. *J. Appl. Probab.* **51** 1100–1113. [MR3301291](#)
- KONTOYIANNIS, I., LASTRAS-MONTAÑO, L. A. and MEYN, S. P. (2006). Exponential bounds and stopping rules for MCMC and general Markov chains. In *VALUETOOLS* 45.
- LEÓN, C. A. and PERRON, F. (2004). Optimal Hoeffding bounds for discrete reversible Markov chains. *Ann. Appl. Probab.* **14** 958–970. [MR2052909](#)
- LEVIN, D. A. and PERES, Y. (2016). Estimating the spectral gap of a reversible Markov chain from a short trajectory. Available at [arXiv:1612.05330](#).
- LEVIN, D. A., PERES, Y. and WILMER, E. L. (2009). *Markov Chains and Mixing Times*. Amer. Math. Soc., Providence, RI. [MR2466937](#)

- LI, X. and WEI, Y. (2001). An improvement on the perturbation of the group inverse and oblique projection. *Linear Algebra Appl.* **338** 53–66. [MR1860312](#)
- LIU, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. Springer Series in Statistics. Springer, New York. [MR1842342](#)
- MCDONALD, D. J., SHALIZI, C. R. and SCHERVISH, M. J. (2011). Estimating beta-mixing coefficients. In *International Conference on Artificial Intelligence and Statistics* 516–524.
- MEYER, C. D. JR. (1975). The role of the group generalized inverse in the theory of finite Markov chains. *SIAM Rev.* **17** 443–464. [MR0383538](#)
- MEYN, S. P. and TWEEDIE, R. L. (1993). *Markov Chains and Stochastic Stability*. Communications and Control Engineering Series. Springer, London. [MR1287609](#)
- MOHRI, M. and ROSTAMIZADEH, A. (2008). Stability bounds for non-iid processes. In *Advances in Neural Information Processing Systems* 20.
- MONTENEGRO, R. and TETALI, P. (2006). Mathematical aspects of mixing times in Markov chains. *Found. Trends Theor. Comput. Sci.* **1** x+121. [MR2341319](#)
- PAULIN, D. (2015). Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electron. J. Probab.* **20** no. 79, 32. [MR3383563](#)
- SENETA, E. (1993). Sensitivity of finite Markov chains under perturbation. *Statist. Probab. Lett.* **17** 163–168. [MR1223539](#)
- STEINWART, I. and CHRISTMANN, A. (2009). Fast learning from non-i.i.d. observations. In *Advances in Neural Information Processing Systems* 22.
- STEINWART, I., HUSH, D. and SCOVEL, C. (2009). Learning from dependent observations. *J. Multivariate Anal.* **100** 175–194. [MR2460486](#)
- STEWART, G. W. and SUN, J. G. (1990). *Matrix Perturbation Theory*. Computer Science and Scientific Computing. Academic Press, Boston, MA. [MR1061154](#)
- SUTTON, R. S. and BARTO, A. G. (1998). *Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning)*. MIT Press, Cambridge, MA.
- TROPP, J. A. (2015). An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning* **8** 1–230.
- WOLFER, G. and KONTOROVICH, A. (2019a). Estimating the mixing time of ergodic Markov chains. Available at [arXiv:1902.01224](#).
- WOLFER, G. and KONTOROVICH, A. (2019b). Minimax identity testing to ergodic Markov chains. Available at [arXiv:1902.00080](#).
- YU, B. (1994). Rates of convergence for empirical processes of stationary mixing sequences. *Ann. Probab.* **22** 94–116. [MR1258867](#)

D. HSU
COMPUTER SCIENCE DEPARTMENT
AND
DATA SCIENCE INSTITUTE
COLUMBIA UNIVERSITY
500 W 120TH STREET
MC 0401
NEW YORK, NEW YORK 10027
USA
E-MAIL: djhsu@cs.columbia.edu

D. A. LEVIN
DEPARTMENT OF MATHEMATICS
UNIVERSITY OF OREGON
FENTON HALL
EUGENE, OREGON 97403-1221
USA
E-MAIL: dlevin@uoregon.edu

A. KONTOROVICH
G. WOLFER
DEPARTMENT OF COMPUTER SCIENCE
BEN-GURION UNIVERSITY OF THE NEGEV
BEER-SHEVA, 84105
ISRAEL
E-MAIL: karyeh@cs.bgu.ac.il
geo.wolfer@gmail.com

Y. PERES
MICROSOFT RESEARCH
1 MICROSOFT WAY
REDMOND, WASHINGTON 98052
USA
E-MAIL: yuval@yuvalperes.com

C. SZEPESVÁRI
DEPARTMENT OF COMPUTING SCIENCE
UNIVERSITY OF ALBERTA
EDMONTON, ALBERTA T6G 2E8
CANADA
E-MAIL: csaba.szepesvari@gmail.com