Estimating the Mixing Time of Ergodic Markov Chains

G. Wolfer, A. Kontorovich
geo.wolfer@gmail.com, lkontor@gmail.com

February 1, 2019

Abstract

We address the problem of estimating the mixing time t_{mix} of an arbitrary ergodic finite Markov chain from a single trajectory of length m. The reversible case was addressed by Hsu et al. [2017], who left the general case as an open problem. In the reversible case, the analysis is greatly facilitated by the fact that the Markov operator is self-adjoint, and Weyl's inequality allows for a dimension-free perturbation analysis of the empirical eigenvalues. As Hsu et al. point out, in the absence of reversibility (and hence, the non-symmetry of the pair probabilities matrix), the existing perturbation analysis has a worst-case exponential dependence on the number of states d. Furthermore, even if an eigenvalue perturbation analysis with better dependence on d were available, in the non-reversible case the connection between the spectral gap and the mixing time is not nearly as straightforward as in the reversible case. Our key insight is to estimate the *pseudo-spectral gap* instead, which allows us to overcome the loss of self-adjointness and to achieve a polynomial dependence on d and the minimal stationary probability $\pi_\star$. Additionally, in the reversible case, we obtain simultaneous nearly (up to logarithmic factors) minimax rates in t_{mix} and precision ε, closing a gap in Hsu et al., who treated ε as constant in the lower bounds. Finally, we construct fully empirical confidence intervals for the pseudo-spectral gap, which shrink to zero at a rate of roughly $1/\sqrt{m}$, and improve the state of the art in even the reversible case.

1 Introduction

We address the problem of estimating the mixing time t_{mix} of a Markov chain from a single trajectory of observations. Approaching the problem from a minimax perspective, we construct point estimates as well as fully empirical confidence intervals for t_{mix} (defined in (4.5)) of an unknown ergodic finite state time homogeneous Markov chain.

It is a classical result [Levin et al., 2009] that the mixing time of an ergodic and reversible Markov chain is controlled by its *absolute spectral gap* $\gamma_\star$ and minimum stationary probability $\pi_\star$:

$$\left(\frac{1}{\gamma_\star} - 1\right)\ln 2 \leq t_{\mathsf{mix}} \leq \frac{\ln\left(4/\pi_\star\right)}{\gamma_\star}, \tag{1.1}$$

which Hsu et al. [2015] leverage to estimate t_{mix} in the reversible case. For non-reversible Markov chains, the relationship between the spectrum and the mixing time is not nearly as straightforward. Any eigenvalue $\lambda \neq 1$ provides a lower bound on the mixing time [Levin et al., 2009],

$$\left(\frac{1}{1-|\lambda|} - 1\right)\ln 2 \leq t_{\mathsf{mix}}, \tag{1.2}$$

and upper bounds may be obtained in terms of the spectral gap of the *multiplicative reversiblization* [Fill, 1991, Montenegro and Tetali, 2006],

$$t_{\mathsf{mix}} \leq \mathcal{O}\left(\frac{1}{\gamma\left(\boldsymbol{M}^\dagger \boldsymbol{M}\right)} \ln\left(\sqrt{\frac{1-\pi_\star}{\pi_\star}}\right)\right). \tag{1.3}$$

Unfortunately, the latter estimate is far from sharp[1]. A more delicate quantity, the *pseudo-spectral gap*, γ_{ps} (formally defined in (4.8)) was introduced by Paulin [2015], who showed that for ergodic $\boldsymbol{M}$,

$$\frac{1}{2\gamma_{\mathsf{ps}}} \leq t_{\mathsf{mix}} \leq \frac{1}{\gamma_{\mathsf{ps}}} \left(\ln \frac{1}{\pi_\star} + 2\ln 2 + 1 \right). \tag{1.4}$$

Thus, γ_{ps} plays a role analogous to that of $\gamma_\star$ in the reversible case, in that it controls the mixing time of non-reversible chains from above and below — and will be our main quantity of interest throughout the paper.

2 Main results

Here we give an informal overview of our main results, all of which pertain to an unknown d-state ergodic time homogeneous Markov chain $\boldsymbol{M}$ with mixing time t_{mix} and minimum stationary probability $\pi_\star$. The formal statements are deferred to Section 5. *Sample complexity* refers to the trajectory length of observations drawn from $\boldsymbol{M}$.

1. We determine the minimax sample complexity of estimating $\pi_\star$ to within a relative error of ε to be $\tilde{\Theta}\left(\frac{t_{\mathsf{mix}}}{\varepsilon^2 \pi_\star}\right)$. This improves the state of the art even in the reversible case.
2. We upper bound the sample complexity of estimating the *pseudo-spectral gap* γ_{ps} of any ergodic $\boldsymbol{M}$ to within a relative error of ε by $\tilde{\mathcal{O}}\left(\frac{t_{\mathsf{mix}}^2 \max\{t_{\mathsf{mix}}, \mathcal{C}(\boldsymbol{M})\}}{\pi_\star \varepsilon^2}\right)$, where $1 \leq \mathcal{C}(\boldsymbol{M}) \leq d/\pi_\star$ captures a notion of how far $\boldsymbol{M}$ is from being doubly-stochastic.
3. We lower bound the sample complexity of estimating t_{mix} by $\Omega\left(\frac{t_{\mathsf{mix}} d}{\varepsilon^2}\right)$, which holds for both the reversible and non-reversible cases. This shows that our upper bound is sharp in ε, up to logarithmic factors.
4. We construct fully empirical confidence intervals for $\pi_\star$ and γ_{ps} without assuming reversibility.
5. Finally, our analysis narrows the width of the confidence intervals and improves their computational cost as compared to the state of the art.

3 Related work

Our work is largely motivated by PAC-type learning problems with dependent data. Many results from statistical learning and empirical process theory have been extended to dependent data with sufficiently rapid mixing [Yu, 1994, Gamarnik, 2003, Karandikar and Vidyasagar, 2002, Mohri and Rostamizadeh, 2008, 2009, Steinwart and Christmann, 2009, Steinwart et al., 2009, Shalizi and Kontorovich, 2013, Wolfer and Kontorovich, 2019]. These have been used to provide generalization guarantees that depend on the possibly unknown mixing properties of the process. In the Markovian setting, the relevant quantity is usually the mixing time, and therefore empirical estimates of this quantity yield corresponding data-dependent generalization bounds.

Other applications include MCMC diagnostics for non-reversible chains, which have recently gained interest through accelerations methods. Chains generated by the classical Metropolis-Hastings are reversible, which is instrumental in analyzing the stationary distribution. However,

[1] Consider the Markov chain on $\{1,2,3\}$ with the transition probability matrix $\boldsymbol{M} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1/2 & 0 & 1/2 \end{pmatrix}$, which is rapidly mixing despite having $\gamma(\boldsymbol{M}^\dagger \boldsymbol{M}) = 0$ [Montenegro and Tetali, 2006].

non-reversible chains may enjoy better mixing properties as well as improved asymptotic variance. For theoretical and experimental results in this direction, see Hildebrand [1997], Chen et al. [1999], Diaconis et al. [2000], Sun et al. [2010], Turitsyn et al. [2011], Vucelja [2014], Suwa and Todo [2010], Turitsyn et al. [2011], Vucelja [2014], Diaconis et al. [2000], Neal [2004], Sun et al. [2010], Chen and Hwang [2013].

The problem of obtaining empirical estimates on the mixing time (with confidence) was first addressed in Hsu et al. [2015] for the reversible setting, by reducing the task to one of estimating the absolute spectral gap and minimum stationary probability. Hsu et al. gave a point estimator for the absolute spectral gap, up to fixed relative error, with sample complexity between $\Omega\left(\frac{d}{\gamma_\star} + \frac{1}{\pi_\star}\right)$ and $\tilde{\mathcal{O}}\left(\frac{1}{\gamma_\star^3 \pi_\star}\right)$. Following up, Levin and Peres [2016] sharpened the upper bound of $\tilde{\mathcal{O}}\left(\frac{1}{\gamma_\star \pi_\star}\right)$, again leveraging properties of the absolute spectral gap of the unknown chain. Additionally, Hsu et al. [2015] presented a point estimator for $\pi_\star$ with sufficient sample complexity $\tilde{\mathcal{O}}\left(\frac{1}{\pi_\star \gamma_\star}\right)$. The state of the art, as well as our improvements, are summarized in Table 1.

More recently, also in the reversible setting, algorithmic improvements in a different sampling model were obtained [Qin et al., 2017, Combes and Touati, 2018], for which new estimation techniques, derived from the *power iteration method* were introduced; these focus on computational rather than statistical efficiency. Our stringent one-sequence sampling model is at the root of many of the technical challenges encountered; allowing, for example, access to a *restart mechanism* of the chain from any given state considerably simplifies the estimation problem [Bhattacharya and Valiant, 2015, Batu et al., 2000, 2013].

Estimated quantity	Hsu et al. [2015]	Levin and Peres [2016]	Current work
$\pi_\star$ (rev. case)	$\tilde{\mathcal{O}}\left(\frac{1}{\pi_\star \gamma_\star \varepsilon^2}\right)$	-	$\tilde{\Omega}\left(\frac{1}{\pi_\star \gamma_\star \varepsilon^2}\right)$ Th. 5.2
$\pi_\star$ (non-rev. case)	-	-	$\tilde{\Theta}\left(\frac{1}{\pi_\star \gamma_{\mathsf{ps}} \varepsilon^2}\right)$ Th. 5.1, Th. 5.2
$\gamma_\star$ (rev. case)	$\tilde{\mathcal{O}}\left(\frac{1}{\pi_\star \gamma_\star^3 \varepsilon^2}\right), \tilde{\Omega}\left(\frac{d}{\gamma_\star} + \frac{1}{\pi_\star}\right)$	$\tilde{\mathcal{O}}\left(\frac{1}{\pi_\star \gamma_\star \varepsilon^2}\right)$	$\tilde{\Omega}\left(\frac{d}{\gamma_\star \varepsilon^2}\right)$ Th. 5.4
γ_{ps} (non-rev. case)	-	-	$\tilde{\mathcal{O}}\left(\frac{\max\{\gamma_{\mathsf{ps}}^{-1}, \mathcal{C}(\boldsymbol{M})\}}{\pi_\star \varepsilon^2 \gamma_{\mathsf{ps}}^2}\right), \tilde{\Omega}\left(\frac{d}{\gamma_{\mathsf{ps}} \varepsilon^2}\right)$ Th. 5.3,5.4

Table 1: Comparison with existing results in the literature.

4 Notation and definitions

We define $[d] \triangleq \{1, \ldots, d\}$, denote the simplex of all distributions over $[d]$ by Δ_d, and the collection of all $d \times d$ row-stochastic matrices by $\mathcal{M}_d$. $\mathbb{N}$ will refer to $\{1, 2, 3, \ldots\}$, and in particular $0 \notin \mathbb{N}$. For $\boldsymbol{\mu} \in \Delta_d$, we will write either $\boldsymbol{\mu}(i)$ or μ_i, as dictated by convenience. All vectors are rows unless indicated otherwise. A Markov chain on d states being entirely specified by an initial distribution $\boldsymbol{\mu} \in \Delta_d$ and a row-stochastic transition matrix $\boldsymbol{M} \in \mathcal{M}_d$, we identify the chain with the pair $(\boldsymbol{M}, \boldsymbol{\mu})$. Namely, by $(X_1, \ldots, X_m) \sim (\boldsymbol{M}, \boldsymbol{\mu})$, we mean that

$$\mathbf{P}\left((X_1, \ldots, X_m) = (x_1, \ldots, x_m)\right) = \boldsymbol{\mu}(x_1) \prod_{t=1}^{m-1} \boldsymbol{M}(x_t, x_{t+1}). \tag{4.1}$$

We write $\mathbf{P}_{\boldsymbol{M},\mu}(\cdot)$ to denote probabilities over sequences induced by the Markov chain $(\boldsymbol{M}, \boldsymbol{\mu})$, and omit the subscript when it is clear from context.

Skipped chains and associated random variables. For a Markov chain $X_1, \dots X_m \sim (\boldsymbol{M}, \boldsymbol{\mu})$, for any $k \in \mathbb{N}$ and $r \in \{0, \dots, k-1\}$ we can define the k-skipped r-offset Markov chain,

$$X_{r+k}, X_{r+2k}, \dots, X_{r+tk}, \dots, X_{r+\lfloor (m-r)/k \rfloor k} \sim (\boldsymbol{M}^k, \boldsymbol{\mu}\boldsymbol{M}^r),$$

and we will write it $X_1^{(k,r)}, \dots, X_{\lfloor (m-r)/k \rfloor}^{(k,r)}$, or more simply $X_1^{(k)}, \dots, X_{\lfloor m/k \rfloor}^{(k)}$ when $r = 0$. The main random quantities in use throughout this work are now defined for clarity. For two states i and j, a skipping rate k, and an offset r, we define the *number of visits to state i* to be

$$N_i^{(k,r)} \triangleq \left| \left\{ 1 \leq t \leq \lfloor (m-r)/k \rfloor - 1 : X_t^{(k,r)} = i \right\} \right|, \quad N_i^{(k)} \triangleq N_i^{(k,0)} \tag{4.2}$$

and the *number of transitions from i to j* to be

$$N_{ij}^{(k,r)} \triangleq \left| \left\{ 1 \leq t \leq \lfloor (m-r)/k \rfloor - 1 : X_t^{(k,r)} = i, X_{t+1}^{(k,r)} = j \right\} \right|, \quad N_{ij}^{(k)} \triangleq N_{ij}^{(k,0)}. \tag{4.3}$$

We will also use the shorthand notation

$$N_i \triangleq N_i^{(1)},\ N_{ij} \triangleq N_{ij}^{(1)},\ N_{\max}^{(k)} \triangleq \max_{i \in [d]} N_i^{(k)},\ N_{\min}^{(k)} \triangleq \min_{i \in [d]} N_i^{(k)},\ N_{\max} \triangleq N_{\max}^{(1)},\ N_{\min} \triangleq N_{\min}^{(1)}. \tag{4.4}$$

Stationarity. The Markov chain $(\boldsymbol{M}, \boldsymbol{\mu})$ is *stationary* if $\boldsymbol{\mu} = \boldsymbol{\mu}\boldsymbol{M}$ (i.e. $\boldsymbol{\mu}$ is a left-eigenvector associated to the eigenvalue 1). Unless noted otherwise, $\boldsymbol{\pi}$ is assumed to be a stationary distribution of the Markov chain in context. We also define $\boldsymbol{D_\pi} \triangleq \operatorname{diag}(\boldsymbol{\pi})$, the diagonal matrix whose entries correspond to the stationary distribution, i.e. $\boldsymbol{\pi} = \boldsymbol{1} \cdot \boldsymbol{D_\pi}$, with $\boldsymbol{1} = (1, \dots, 1)$.

Ergodicity. The Markov chain $(\boldsymbol{M}, \boldsymbol{\mu})$ is *ergodic* if $\boldsymbol{M}^k > 0$ (entry-wise positive) for some $k \geq 1$. If $\boldsymbol{M}$ is ergodic, it has a unique stationary distribution $\boldsymbol{\pi}$ and moreover $\pi_\star > 0$, where $\pi_\star = \min_{i \in [d]} \pi_i$ is called the *minimum stationary probability.* We henceforth only consider ergodic chains.

Mixing time. When the chain is ergodic, we can define its mixing time as the number of steps it requires to converge to its stationary distribution within a constant precision (traditionally taken to be 1/4):

$$t_{\mathsf{mix}} \triangleq \min_{t \in \mathbb{N}} \left\{ \sup_{\boldsymbol{\mu} \in \Delta_d} \left\| \boldsymbol{\mu}\boldsymbol{M}^{t-1} - \boldsymbol{\pi} \right\|_{\mathsf{TV}} \leq \frac{1}{4} \right\}. \tag{4.5}$$

Reversibility. A *reversible* $\boldsymbol{M} \in \mathcal{M}_d$ satisfies *detailed balance* for some distribution $\boldsymbol{\mu}$: for all $i, j \in [d]$, $\mu_i \boldsymbol{M}(i,j) = \mu_j \boldsymbol{M}(j,i)$ — in which case $\boldsymbol{\mu}$ is necessarily the unique stationary distribution. The eigenvalues of a reversible $\boldsymbol{M}$ lie in $(-1, 1]$, and these may be ordered (counting multiplicities): $1 = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d$. The *spectral gap* is defined to be $\gamma = 1 - \lambda_2(\boldsymbol{M})$, the *absolute spectral gap* is

$$\gamma_\star = 1 - \max\left\{ \lambda_2(\boldsymbol{M}), |\lambda_d(\boldsymbol{M})| \right\}. \tag{4.6}$$

It follows from above that whenever $\boldsymbol{M}$ is *lazy* ($\forall i \in [d], \boldsymbol{M}(i,i) \geq 1/2$), all eigenvalues of $\boldsymbol{M}$ are positive and $\gamma = \gamma_\star$. The matrix $\boldsymbol{D_\pi}\boldsymbol{M}$ consists of the *doublet probabilities* associated with $\boldsymbol{M}$ and is symmetric when $\boldsymbol{M}$ is reversible. The *rescaled transition matrix*,

$$\boldsymbol{L} = \boldsymbol{D_\pi}^{1/2} \boldsymbol{M} \boldsymbol{D_\pi}^{-1/2}, \tag{4.7}$$

is also symmetric for reversible $\boldsymbol{M}$ (but not in general). Since $\boldsymbol{L}$ and $\boldsymbol{M}$ are similar matrices, their eigenvalue systems are identical.

Non-reversibility. Chains that do not satisfy the detailed balanced equations are said to be non-reversible. In this case, the eigenvalues may be complex, and the transition matrix may not be diagonalizable, even over $\mathbb{C}$. Paulin [2015] defines the *pseudo-spectral gap* by

$$\gamma_{\mathsf{ps}} \triangleq \max_{k \in \mathbb{N}} \left\{ \frac{\gamma((\boldsymbol{M}^\dagger)^k \boldsymbol{M}^k)}{k} \right\}, \tag{4.8}$$

where $\boldsymbol{M}^\dagger$ is the *time reversal* of $\boldsymbol{M}$, given by $\boldsymbol{M}^\dagger(i,j) \triangleq \boldsymbol{\pi}(j)\boldsymbol{M}(j,i)/\boldsymbol{\pi}(i)$; the expression $\boldsymbol{M}^\dagger \boldsymbol{M}$ is called the *multiplicative reversiblization* of $\boldsymbol{M}$. The chain $\boldsymbol{M}^\dagger \boldsymbol{M}$ is alway reversible, and its eigenvalues are all real and non-negative [Fill, 1991]. We also denote by k_{ps} the smallest positive integer such that $\gamma_{\mathsf{ps}} = \frac{\gamma\left((\boldsymbol{M}^\dagger)^{k_{\mathsf{ps}}} \boldsymbol{M}^{k_{\mathsf{ps}}}\right)}{k_{\mathsf{ps}}}$; this is the power of $\boldsymbol{M}$ for which the multiplicative reversiblization achieves[2] its pseudo-spectral gap.

Norms and metrics. We use the standard ℓ_1, ℓ_2 norms $\|z\|_p = \left(\sum_{i\in[d]} |z_i|\right)^{1/p}$; in the context of distributions (and up to a convention-dependent factor of 2), $p = 1$ corresponds to the total variation norm. For $A \in \mathbb{R}^{d\times d}$, define the spectral radius $\rho(A)$ to be the largest absolute value of the eigenvalues of A, and recall the following operator norms for real matrices,

$$\|A\|_\infty = \max_{i\in[d]} \sum_{j\in[d]} |A(i,j)|, \quad \|A\|_1 = \max_{j\in[d]} \sum_{i\in[d]} |A(i,j)|, \quad \|A\|_2 = \sqrt{\rho(A^\intercal A)}. \tag{4.9}$$

We denote by $\langle\cdot,\cdot\rangle_{\boldsymbol{\pi}}$ the inner product on $\mathbb{R}^d$ defined by $\langle \mathbf{f}, \mathbf{g}\rangle_{\boldsymbol{\pi}} \triangleq \sum_{i\in[d]} \mathbf{f}(i)\mathbf{g}(i)\boldsymbol{\pi}(i)$, and write $\|\cdot\|_{2,\boldsymbol{\pi}}$ for its associated norm; $\ell^2(\boldsymbol{\pi})$ is the resulting Hilbert space. To any $(\boldsymbol{M}, \boldsymbol{\mu})$, we also associate

$$\|\boldsymbol{\mu}/\boldsymbol{\pi}\|_{2,\boldsymbol{\pi}}^2 \triangleq \sum_{i\in[d]} \mu_i^2/\pi_i, \tag{4.10}$$

which provides a notion of "distance from stationarity" and satisfies $\|\boldsymbol{\mu}/\boldsymbol{\pi}\|_{2,\boldsymbol{\pi}} \le 1/\pi_\star$. For two $[d]$-supported distributions $\boldsymbol{D} = (p_1,\ldots,p_d)$ and $\boldsymbol{D}' = (q_1,\ldots,q_d)$, we also define respectively the Hellinger distance and the KL divergence,

$$H(\boldsymbol{D},\boldsymbol{D}') = \frac{1}{\sqrt{2}}\sqrt{\sum_{i\in[d]} (\sqrt{p_i} - \sqrt{q_i})^2} \quad D_{\mathrm{KL}}\left(\boldsymbol{D}\middle|\middle|\boldsymbol{D}'\right) = \sum_{i\in[d]} p_i \ln\left(\frac{p_i}{q_i}\right). \tag{4.11}$$

Asymptotic notation. We use standard $\mathcal{O}(\cdot)$, $\Omega(\cdot)$ and $\Theta(\cdot)$ order-of-magnitude notation, as well as their tilde variants $\tilde{\mathcal{O}}(\cdot)$, $\tilde{\Omega}(\cdot)$, $\tilde{\Theta}(\cdot)$ where lower-order log factors of any variables are suppressed. Note that as we are giving finite sample fully empirical bounds with explicit multiplicative constants, the logarithm base is relevant and we write $\ln(\cdot)$ for the natural logarithm of base e.

5 Formal statement of the results

The results in this section are all summarized in terms of learning the parameters of interest up to some ε relative error and confidence δ. For absolute error point estimation learning bounds, we refer the reader to Section B.

Theorem 5.1 (Minimum stationary probability estimation upper bound)**.** *There exists an estimator $\hat{\pi}_\star$ which, for all $0 < \varepsilon < 2$, $0 < \delta < 1$, satisfies the following. If $\hat{\pi}_\star$ receives as input a sequence $\boldsymbol{X} = (X_1,\ldots,X_m)$ of length at least m_{UB} drawn according to an unknown d-state*

[2] Note that for ergodic chains the pseudo-spectral is always achieved for finite k so that k_{ps} is properly defined. Indeed, as $1/(2\gamma_{\mathsf{ps}}) \le t_{\mathsf{mix}} < \infty$ [Paulin, 2015] and directly from its definition $k_{\mathsf{ps}} \le 1/\gamma_{\mathsf{ps}}$.

Markov chain $(\boldsymbol{M}, \boldsymbol{\mu})$ *with minimal stationary probability* $\pi_\star$ *and pseudo-spectral gap* γ_{ps}*, then* $|\hat{\pi}_\star - \pi_\star| < \varepsilon \pi_\star$ *holds with probability at least* $1-\delta$*. The sample complexity is upper-bounded by*

$$m_{\mathrm{UB}} := \frac{C_\star}{\gamma_{\mathsf{ps}} \varepsilon^2 \pi_\star} \ln \left(\frac{d}{\delta} \sqrt{\pi_\star^{-1}} \right), \tag{5.1}$$

where $C_\star$ *is a universal constant.*

From this theorem, and the fact that for reversible Markov chains the absolute and pseudo-spectral gaps are within a constant multiplicative factor (see Lemma C.1), we immediately recover the point estimation upper bound for the minimum stationary probability provided by Hsu et al. [2015] for reversible chains.

Theorem 5.2 (Minimum stationary probability estimation lower bound)**.** *Let* $d \in \mathbb{N}, d \geq 4$*. For every* $0 < \varepsilon < 1/2, 0 < \pi_\star < \frac{1}{d}$ *and* $\pi_\star < \gamma_{\mathsf{ps}} < 1$*, there exists a* d*-state Markov chain* $\boldsymbol{M}$ *with pseudo-spectral gap* γ_{ps} *and minimum stationary probability* $\pi_\star$ *such that any estimator must require a sequence* $\boldsymbol{X} = (X_1, \ldots, X_m)$ *drawn from the unknown* $\boldsymbol{M}$ *of length at least*

$$m_{\mathrm{LB}} := \Omega \left(\frac{\ln \left(\delta^{-1} \right)}{\gamma_{\mathsf{ps}} \varepsilon^2 \pi_\star} \right), \tag{5.2}$$

for $|\hat{\pi}_\star - \pi_\star| < \varepsilon \pi_\star$ *to hold with probability* $\geq 1 - \delta$*.*

The upper and lower bounds in (5.1) and (5.2) match up to logarithmic factors — and continue to hold for the reversible case (Lemma C.1).

Theorem 5.3 (Pseudo-spectral gap estimation upper bound)**.** *There exists an estimator* $\widehat{\gamma}_{\mathsf{ps}}$ *which, for all* $0 < \varepsilon < 2$*,* $0 < \delta < 1$*, satisfies the following. If* $\widehat{\gamma}_{\mathsf{ps}}$ *receives as input a sequence* $\boldsymbol{X} = (X_1, \ldots, X_m)$ *of length at least* m_{UB} *drawn according to an unknown* d*-state Markov chain* $(\boldsymbol{M}, \boldsymbol{\mu})$ *with minimal stationary probability* $\pi_\star$ *and pseudo-spectral gap* γ_{ps}*, then* $|\widehat{\gamma}_{\mathsf{ps}} - \gamma_{\mathsf{ps}}| < \varepsilon \gamma_{\mathsf{ps}}$ *holds with probability at least* $1-\delta$*. The sample complexity is upper-bounded by*

$$m_{\mathrm{UB}} := \frac{C_{\mathsf{ps}}}{\pi_\star \varepsilon^2 \gamma_{\mathsf{ps}}^2} \max \left\{ \frac{1}{\gamma_{\mathsf{ps}}}, \mathcal{C}(\boldsymbol{M}) \right\} \ln \left(\frac{d \sqrt{\pi_\star^{-1}}}{\varepsilon^2 \delta} \right), \tag{5.3}$$

where $\mathcal{C}(\boldsymbol{M}) \leq \|\boldsymbol{M}\|_{\boldsymbol{\pi}} \min \{d, \|\boldsymbol{M}\|_{\boldsymbol{\pi}}\}$*,* $\|\boldsymbol{M}\|_{\boldsymbol{\pi}} \triangleq \max_{(i,j) \in [d]^2} \left\{ \frac{\pi_i}{\pi_j} \right\}$*, and* C_{ps} *is a universal constant.*

The quantity $\mathcal{C}(\boldsymbol{M})$ can be thought of as a measure of "distance to double stochasticity" of the chain, or a measure of non-uniformity of its stationary distribution; for doubly-stochastic chains, we have $\mathcal{C}(\boldsymbol{M}) = 1$. The proof of the theorem provides a more delicate yet less tractable expression for $\mathcal{C}(\boldsymbol{M})$, given in (B.34).

Theorem 5.4 (Pseudo-spectral gap estimation lower bound)**.** *Let* $d \in \mathbb{N}, d \geq 4$*. For every* $0 < \varepsilon < 1/4, 0 < \delta < 1/d, 0 < \gamma_{\mathsf{ps}} < 1/8$*, there exists a* d*-state Markov chains* $\boldsymbol{M}$ *with pseudo-spectral gap* γ_{ps} *such that every estimator* $\widehat{\gamma}_{\mathsf{ps}}$ *must require in the worst case a sequence* $\boldsymbol{X} = (X_1, \ldots, X_m)$ *drawn from the unknown* $\boldsymbol{M}$ *of length at least*

$$m_{\mathrm{LB}} := \Omega \left(\frac{d}{\gamma_{\mathsf{ps}} \varepsilon^2} \ln \left(\frac{1}{d\delta} \right) \right),$$

in order for $|\widehat{\gamma}_{\mathsf{ps}} - \gamma_{\mathsf{ps}}| < \varepsilon \gamma_{\mathsf{ps}}$ *to hold with probability at least* $1-\delta$*.*

The proof of the above result also yields the lower bound of $\tilde{\Omega}\left(\frac{d}{\gamma_\star \varepsilon^2}\right)$ is the reversible case, and closes the minimax estimation gap, also in terms of the precision parameter ε, at least for doubly-stochastic Markov chains (for which $\pi_\star = 1/d$), matching the upper bound of Hsu et al. [2017] up to logarithmic terms. For a comparison with existing results in the literature, see Table 1.

The minimax sample complexity in γ_{ps} remains to be pinned down as our current bounds exhibit a gap. Closing the gap appears not to be amenable to the techniques used to close the corresponding gap in the reversible case (see Remark B.2), and we leave this as an open problem.

6 Empirical procedure

The full procedure for estimating the pseudo-spectral gap is described below. For clarity, the computation of the confidence intervals is not made explicit in the pseudo-code and the reader is referred to Theorem A.1 for their expression.

Function `PseudoSpectralGap`$(d, (X_1, \dots, X_m), K)$**:**

- $\widehat{\gamma}_{\mathsf{ps}} \leftarrow 0$
- **for** $k \leftarrow 1$ **to** K **do**
  - $\tilde{\gamma}_k^\dagger \leftarrow 0$
  - **for** $r \leftarrow 0$ **to** $k-1$ **do**
    - $\hat{\gamma}_{k,r}^\dagger \leftarrow$ `SpectralGapMultRev`$(d, \alpha, (X_{r+k}, X_{r+2k}, \dots, X_{r+\lfloor (m-r)/k \rfloor k}))$
    - $\tilde{\gamma}_k^\dagger \leftarrow \tilde{\gamma}_k^\dagger + \hat{\gamma}_{k,r}^\dagger$
  - **end**
  - $\tilde{\gamma}_k^\dagger \leftarrow \tilde{\gamma}_k^\dagger / k$
  - **if** $\tilde{\gamma}_k^\dagger / k > \widehat{\gamma}_{\mathsf{ps}}$ **then**
    - $\widehat{\gamma}_{\mathsf{ps}} \leftarrow \tilde{\gamma}_k^\dagger / k$
  - **end**
- **end**
- **return** $\widehat{\gamma}_{\mathsf{ps}}$

Function `SpectralGapMultRev`$(d, \alpha, (X_1, \dots, X_n))$**:**

- $\mathbf{N} \leftarrow [d\alpha]_d$
- $\mathbf{T} \leftarrow [\alpha]_{d \times d}$
- **for** $t \leftarrow 1$ **to** $n-1$ **do**
  - $\mathbf{N}[X_t] \leftarrow \mathbf{N}[X_t] + 1$
  - $\mathbf{T}[X_t, X_{t+1}] \leftarrow \mathbf{T}[X_t, X_{t+1}] + 1$
- **end**
- $\mathbf{S} \leftarrow \sqrt{1/\mathbf{N}}$
- $\mathbf{F} \leftarrow$ `MatrixMultiply`$(\mathbf{S}^\intercal, \mathbf{S})/n$
- $\mathbf{D} \leftarrow$ `Diag`$(1/\mathbf{N})$
- $\mathbf{G} \leftarrow$ `MatrixMultiply`$(\mathbf{S}, \mathbf{T}^\intercal, \mathbf{D}, \mathbf{T}, \mathbf{S})$
- $\hat{\gamma}^\dagger \leftarrow 1 -$ `SpectralRadius`$(\mathbf{G} - \mathbf{F})$
- **return** $\hat{\gamma}^\dagger$

Algorithm 1: The estimation procedure outputting $\widehat{\gamma}_{\mathsf{ps}}$

We construct fully empirical confidence intervals whose non-asymptotic form is deferred to Theorem A.1, and the asymptotic behavior is summarized as follows:

Theorem 6.1 (Confidence intervals, asymptotic behavior)**.** *In the non-reversible case, the* *interval widths asymptotically behave as*

$$\sqrt{m}\left|\hat{\pi}_\star - \pi_\star\right| = \tilde{\mathcal{O}}\left(\frac{\sqrt{d}}{\gamma_{\mathsf{ps}}\sqrt{\pi_\star}}\right), \quad \sqrt{m}\left|\widehat{\gamma}_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}}\right| = \tilde{\mathcal{O}}\left(\frac{1}{K} + \sqrt{\frac{d}{\pi_\star}}\left(\sqrt{d}\left\|\boldsymbol{M}\right\|_{\boldsymbol{\pi}} + \frac{1}{\gamma_{\mathsf{ps}}\pi_\star}\right)\right), \tag{6.1}$$

and in the reversible case, they asymptotically behave as

$$\sqrt{m}\left|\hat{\pi}_\star - \pi_\star\right| = \tilde{\mathcal{O}}\left(\frac{\sqrt{d}}{\gamma_\star\sqrt{\pi_\star}}\right), \qquad \sqrt{m}\left|\widehat{\gamma}_\star - \gamma_\star\right| = \tilde{\mathcal{O}}\left(\sqrt{\frac{d}{\pi_\star}}\left(\sqrt{d}\left\|\boldsymbol{M}\right\|_{\boldsymbol{\pi}} + \frac{1}{\gamma_\star\pi_\star}\right)\right). \tag{6.2}$$

See Remark B.3 for a discussion of how Theorem 6.1 improves the state of the art.

7 Proof sketches

In this section we, provide proof sketches and explain the basic intuition. These are fully fleshed out in Section B.

7.1 Point estimation

7.1.1 Proof sketch of Theorem 5.1

We take the natural candidate $\hat{\pi}_\star \triangleq \frac{1}{m}\min_{i\in[d]}\left|t \in [m] : X_t = i\right|$ as our estimator. The proof follows along the lines of its counterpart in Hsu et al. [2015] for the reversible case, with the exception that it makes use of a more general concentration inequality from Paulin [2015] that also applies to non-reversible Markov chains.

7.1.2 Proof sketch of Theorem 5.2

To prove the claim, we first focus on the estimation problem with absolute error, in the regime $2\varepsilon < \pi_\star < \gamma_{\mathsf{ps}}$. We construct the following star-shaped class of *reversible* Markov chain where a single "hub" state is special, while the remaining "spoke" can only transition to themselves or to the hub. For $d \in \mathbb{N}, d \geq 4$,

$$\mathcal{S}_d = \left\{\boldsymbol{S}_\alpha(\boldsymbol{D}) : 0 < \alpha < 1, \boldsymbol{D} = (p_1, \dots, p_d) \in \Delta_d\right\} \text{ where } \boldsymbol{S}_\alpha(\boldsymbol{D}) = \begin{pmatrix} \alpha & (1-\alpha)p_1 & \cdots & (1-\alpha)p_d \\ \alpha & 1-\alpha & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \alpha & 0 & \dots & 1-\alpha \end{pmatrix}. \tag{7.1}$$

Any chain in $\mathcal{S}_d$ is readily computed to have stationary distribution $(\alpha, (1-\alpha)p_1, \dots, (1-\alpha)p_d)$ and pseudo-spectral gap $\gamma_{\mathsf{ps}}(\boldsymbol{S}_\alpha(\boldsymbol{D})) = \Theta(\alpha)$. Consider $\boldsymbol{D}, \boldsymbol{D}_\varepsilon \in \Delta_d$ defined by

$$\boldsymbol{D} := \left(\beta, \beta, \frac{1-2\beta}{d-2}, \dots, \frac{1-2\beta}{d-2}\right), \boldsymbol{D}_\varepsilon := \left(\beta + 2\varepsilon, \beta - 2\varepsilon, \frac{1-2\beta}{d-2}, \dots, \frac{1-2\beta}{d-2}\right), \tag{7.2}$$

with $2\varepsilon < \beta < 1/d$. Consider further the two (stationary) Markov chains $\boldsymbol{S}_\alpha, \boldsymbol{S}_{\alpha,\varepsilon} \in \mathcal{S}_d$ indexed, respectively, by $\boldsymbol{D}$ and $\boldsymbol{D}_\varepsilon$. Then $\left|\pi_\star(\boldsymbol{S}_\alpha) - \pi_\star(\boldsymbol{S}_{\alpha,\varepsilon})\right| = \Omega(\varepsilon)$, and for all m, we exploit the structure of $\mathcal{S}_d$ to derive a tensorization property of the KL divergence between trajectories $\boldsymbol{X}_1^m \sim \boldsymbol{S}_\alpha$ and $\boldsymbol{Y}_1^m \sim \boldsymbol{S}_{\alpha,\varepsilon}$:

$$D_{\mathrm{KL}}\left(\boldsymbol{X}_1^m || \boldsymbol{Y}_1^m\right) \leq \alpha m D_{\mathrm{KL}}\left(\boldsymbol{D}_\varepsilon || \boldsymbol{D}\right).$$

The argument is concluded with a direct computation of $D_{\mathrm{KL}}\left(\boldsymbol{D}_\varepsilon || \boldsymbol{D}\right) = \mathcal{O}\left(\frac{\varepsilon^2}{\beta}\right)$ and a KL version of Le Cam's two point method.

7.1.3 Proof sketch of Theorem 5.3

We first solve the estimation problem with respect to the absolute error. The quantity we wish to estimate is a maximum over the integers, while an empirical procedure can only consider a finite search space $[K] \subset \mathbb{N}$. To highlight the dependence of our estimator on the choice of K, we write

$$\gamma_k^\dagger \triangleq \gamma\left(\left(\boldsymbol{M}^\dagger\right)^k \boldsymbol{M}^k\right), \quad \widehat{\gamma}_k^\dagger \triangleq \gamma\left(\left(\widehat{\boldsymbol{L}}^{(k)}\right)^\intercal \widehat{\boldsymbol{L}}^{(k)}\right), \quad \gamma_{\mathsf{ps}[K]} \triangleq \max_{k \in [K]} \left\{ \frac{\gamma_k^\dagger}{k} \right\}, \tag{7.3}$$

where $\widehat{\boldsymbol{L}}^{(k)}$ the empirical version of the rescaled transition matrix (4.7) associated with the k-skipped chain. We denote by $\widehat{\gamma}_{\mathsf{ps}[K]}$ the empirical estimator, chosen in our case to be

$$\widehat{\gamma}_{\mathsf{ps}[K]} \triangleq \max_{k \in [K]} \left\{ \frac{\widehat{\gamma}_k^\dagger}{k} \right\}. \tag{7.4}$$

It is easily seen that $\left|\gamma_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}}\right| \leq \frac{1}{K}$, and so it suffices to consider $K = \lceil 2/\varepsilon \rceil$, which yields

$$\mathbf{P}\left(\left|\widehat{\gamma}_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}}\right| > \varepsilon\right) \leq \sum_{k=1}^{\lceil \frac{2}{\varepsilon} \rceil} \mathbf{P}\left(\left|\gamma_k^\dagger - \widehat{\gamma}_k^\dagger\right| > \frac{k\varepsilon}{2}\right). \tag{7.5}$$

Via a standard application of Weyl's eigenvalue perturbation inequality, Perron-Frobenius theory, and properties of similar matrices, we obtain:

$$\left|\widehat{\gamma}_k^\dagger - \gamma_k^\dagger\right| \leq \left\|(\widehat{\boldsymbol{L}}^{(k)})^\intercal \widehat{\boldsymbol{L}}^{(k)} - (\boldsymbol{L}^k)^\intercal \boldsymbol{L}^k\right\|_2 \leq 2\left\|\boldsymbol{L}^k - \widehat{\boldsymbol{L}}^{(k)}\right\|_2. \tag{7.6}$$

We continue by decomposing $\widehat{\boldsymbol{L}}^{(k)} - \boldsymbol{L}^k$,

$$\begin{aligned} \widehat{\boldsymbol{L}}^{(k)} - \boldsymbol{L}^k &= \mathcal{E}_{\boldsymbol{M}}^{(k)} + \mathcal{E}_{\boldsymbol{\pi},1}^{(k)} \boldsymbol{L}^k + \boldsymbol{L}^k \mathcal{E}_{\boldsymbol{\pi},2}^{(k)} + \mathcal{E}_{\boldsymbol{\pi},1}^{(k)} \boldsymbol{L}^k \mathcal{E}_{\boldsymbol{\pi},2}^{(k)} \\ \text{where } \mathcal{E}_{\boldsymbol{M}}^{(k)} &= \left(\widehat{\boldsymbol{D}}_{\boldsymbol{\pi}}^{(k)}\right)^{1/2} \left(\widehat{\boldsymbol{M}}^{(k)} - \boldsymbol{M}^k\right) \left(\widehat{\boldsymbol{D}}_{\boldsymbol{\pi}}^{(k)}\right)^{-1/2}, \\ \mathcal{E}_{\boldsymbol{\pi},1}^{(k)} &= \left(\widehat{\boldsymbol{D}}_{\boldsymbol{\pi}}^{(k)}\right)^{1/2} \boldsymbol{D}_{\boldsymbol{\pi}}^{-1/2} - \boldsymbol{I}, \qquad \mathcal{E}_{\boldsymbol{\pi},2}^{(k)} = \boldsymbol{D}_{\boldsymbol{\pi}}^{1/2} \left(\widehat{\boldsymbol{D}}_{\boldsymbol{\pi}}^{(k)}\right)^{-1/2} - \boldsymbol{I} \end{aligned} \tag{7.7}$$

so that

$$\left\|\widehat{\boldsymbol{L}}^{(k)} - \boldsymbol{L}^k\right\|_2 \leq \left\|\mathcal{E}_{\boldsymbol{M}}^{(k)}\right\|_2 + 2\left\|\mathcal{E}_{\boldsymbol{\pi}}^{(k)}\right\|_2 + \left\|\mathcal{E}_{\boldsymbol{\pi}}^{(k)}\right\|_2^2, \text{ where } \left\|\mathcal{E}_{\boldsymbol{\pi}}^{(k)}\right\|_2 \triangleq \max\left\{\left\|\mathcal{E}_{\boldsymbol{\pi},1}^{(k)}\right\|_2, \left\|\mathcal{E}_{\boldsymbol{\pi},2}^{(k)}\right\|_2\right\}. \tag{7.8}$$

Now we must upper bound $\left\|\mathcal{E}_{\boldsymbol{M}}^{(k)}\right\|_2$, $\left\|\mathcal{E}_{\boldsymbol{\pi}}^{(k)}\right\|_2$, and the "bad" event that a k-skipped chain did not visit every state a "reasonable" amount of times (in which case, the estimator might not even be properly defined). The most challenging quantity to bound is $\left\|\mathcal{E}_{\boldsymbol{M}}^{(k)}\right\|_2$, which we achieve via a row-martingale process. After carefully controlling its second-order induced row and column processes, we invoke a matrix martingale version of Freedman's inequality [Tropp, 2011], concluding that $\left\|\mathcal{E}_{\boldsymbol{M}}^{(k)}\right\|_2 = \mathcal{O}(k\varepsilon)$ with high confidence for a trajectory of length $m = \tilde{\Omega}\left(\frac{\left\|\boldsymbol{M}^k\right\|_1 \|\boldsymbol{M}\|_{\boldsymbol{\pi}}}{k\varepsilon^2}\right)$. The bound $\left\|\mathcal{E}_{\boldsymbol{\pi}}^{(k)}\right\|_2 = \mathcal{O}(k\varepsilon)$ follows from Hsu et al. [2017, Section 6.3] and Theorem C.1. The proof is concluded with an absolute-to-relative error conversion — i.e., replacing ε by $\varepsilon\gamma_{\mathsf{ps}}$.

7.1.4 Proof sketch of Theorem 5.4

For $0 < \alpha < \frac{1}{8}$ and $d \geq 4$, we define the following family of symmetric (and hence reversible) stochastic $d \times d$ matrices:

$$\boldsymbol{M}(\alpha) = \begin{pmatrix} 1-\alpha & \frac{\alpha}{d-1} & \cdots & \cdots & \frac{\alpha}{d-1} \\ \frac{\alpha}{d-1} & 1/2 - \frac{\alpha}{d-1} & \frac{1}{2(d-2)} & \cdots & \frac{1}{2(d-2)} \\ \vdots & \frac{1}{2(d-2)} & \ddots & & \frac{1}{2(d-2)} \\ \vdots & \vdots & & & \vdots \\ \frac{\alpha}{d-1} & \frac{1}{2(d-2)} & \frac{1}{2(d-2)} & \cdots & 1/2 - \frac{\alpha}{d-1} \end{pmatrix}, \tag{7.9}$$

for which $\boldsymbol{\pi} = \frac{1}{d} \cdot \mathbf{1}$, and $\gamma_{\mathsf{ps}} = \Theta(\alpha)$. We then invoke a result of Kazakos [1978], reproduced in Lemma D.1, which provides with a method for recursively computing the Hellinger distance between two distributions over words of length m sampled from different Markov chains in terms of the entry-wise geometric mean of their transition matrices. The problem is then reduced to one of controlling a spectral radius. Writing for convenience

$$p = \frac{\sqrt{\alpha_0 \alpha_1}}{d-1}, q = \frac{1}{2(d-2)}, r = \sqrt{(1-\alpha_0)(1-\alpha_1)}, s = \sqrt{\left(1/2 - \frac{\alpha_0}{d-1}\right)\left(1/2 - \frac{\alpha_1}{d-1}\right)},$$

we compute the entry-wise geometric mean to be

$$[\boldsymbol{M}(\alpha_0), \boldsymbol{M}(\alpha_1)]_{\sqrt{}} = \begin{pmatrix} r & p & \cdots & \cdots & p \\ p & s & q & \cdots & q \\ \vdots & q & \ddots & & q \\ \vdots & \vdots & & & \vdots \\ p & q & q & \cdots & s \end{pmatrix}. \tag{7.10}$$

Observing that the rank of this matrix is less than 2 and employing some careful analysis, we are able to bound the spectral radius ρ of $[\boldsymbol{M}(\alpha_0), \boldsymbol{M}(\alpha_1)]_{\sqrt{}}$ from below by $\rho \geq 1 - 6\frac{\alpha\varepsilon^2}{d-1}$. A Hellinger version of Le Cam's two-point method concludes the proof.

7.2 Empirical confidence intervals

Non-reversible setting. The complete algorithmic procedure is described in Algorithm 1. Formally, for a sample path $(X_1, \dots, X_m)$, a fixed k and a smoothing parameter α, we construct the estimator

$$\widehat{\gamma}^{\dagger}_{k,\alpha}(X_1, \dots, X_m) \triangleq \frac{1}{k} \sum_{r=0}^{k-1} \widehat{\gamma}^{\dagger}_{k,r,\alpha}(X_t^{(k,r)}, 1 \leq t \leq \lfloor (m-r)/k \rfloor), \tag{7.11}$$

where $\widehat{\gamma}^{\dagger}_{k,r,\alpha}$ is an estimator for the spectral gap of the multiplicative reversiblization of $\boldsymbol{M}^k$, which was constructed by observing the k-skipped r-offset Markov chain $X_t^{(k,r)}, 1 \leq t \leq \lfloor (m - r)/k \rfloor$, and applying Laplace α-smoothing. We then notice that $\widehat{\gamma}^{\dagger}_{k,r,\alpha} = \gamma\left(\left(\widehat{\boldsymbol{L}}^{(k,r,\alpha)}\right)^{\intercal} \widehat{\boldsymbol{L}}^{(k,r,\alpha)}\right)$, where

$$\begin{aligned} \left(\widehat{\boldsymbol{L}}^{(k,r,\alpha)}\right)^{\intercal} \widehat{\boldsymbol{L}}^{(k,r,\alpha)} &= \left(\mathbf{D}_N^{(k,r,\alpha)}\right)^{-1/2} \left(\mathbf{N}^{(k,r,\alpha)}\right)^{\intercal} \left(\mathbf{D}_N^{(k,r,\alpha)}\right)^{-1} \mathbf{N}^{(k,r,\alpha)} \left(\mathbf{D}_N^{(k,r,\alpha)}\right)^{-1/2}, \\ \mathbf{N}^{(k,r,\alpha)} &\triangleq \left[N_{ij}^{(k,r)} + \alpha\right]_{(i,j)}, \ \mathbf{D}_N^{(k,r,\alpha)} \triangleq \operatorname{diag}\left(N_1^{(k,r)} + d\alpha, \dots, N_d^{(k,r)} + d\alpha\right). \end{aligned} \tag{7.12}$$

The derivation of the confidence intervals starts with an empirical version of the decomposition introduced for the point estimator. The subsequent analysis has two key components. The first is a perturbation bound for the stationary distribution in terms of the pseudo-spectral gap and the stability of the perturbation of matrix with respect to the $\|\cdot\|_\infty$ norm. More precisely, Lemma C.3 guarantees that

$$\|\widehat{\boldsymbol{\pi}} - \boldsymbol{\pi}\|_\infty \leq \tilde{\mathcal{O}}(1) \frac{1}{\gamma_{\mathsf{ps}}\left(\widehat{\boldsymbol{M}}\right)} \left\|\widehat{\boldsymbol{M}} - \boldsymbol{M}\right\|_\infty. \tag{7.13}$$

The second component (Lemma C.4) involves controlling the latter perturbation in terms of empirically observable quantities. In particular,

$$\left\|\widehat{\boldsymbol{M}} - \boldsymbol{M}\right\|_\infty \leq \tilde{\mathcal{O}}(1) \sqrt{\frac{d}{N_{\min}}} \tag{7.14}$$

holds with high probability — which is an empirical version of the result of Wolfer and Kontorovich [2019, Theorem 1], achieved by constructing and analyzing appropriate row-martingales.

Reversible setting. Our analysis also yields improvements over the state of the art estimation procedure in the reversible setting, where Hsu et al. [2015] used the absolute spectral gap $\gamma_\star$ of the additive reversiblization of the empirical transition matrix $\frac{\widehat{\boldsymbol{M}}^\dagger + \widehat{\boldsymbol{M}}}{2}$ as the estimator for the mixing time. Our analysis via row-martingales sharpens the confidence intervals roughly by a factor of $\mathcal{O}(\sqrt{d})$ over the previous method. The latter relied on entry-wise martingales together with the metric inequality $\|\mathbf{A}\|_\infty \leq d \max_{(i,j)\in[d]^2} |\mathbf{A}(i,j)|, \mathbf{A} \in \mathbb{R}^{d\times d}$. Additionally, we show that the computation complexity of the task can be reduced over non-trivial parameter regimes. We achieve this via iterative methods for computeing the second largest eigenvalue, and by replacing an expensive pseudo-inverse computation by the already-computed estimator for $\gamma_\star$ itself (Corollary C.3). These computational improvements do not degrade the asymptotic behavior of the confidence intervals.

Acknowledgments

We are thankful to Daniel Paulin for the insightful conversations.

References

S. Arora, E. Hazan, and S. Kale. Fast algorithms for approximate semidefinite programming using the multiplicative weights update method. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*, pages 339–348. IEEE, 2005.

T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White. Testing that distributions are close. In *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*, pages 259–269. IEEE, 2000.

T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White. Testing closeness of discrete distributions. *Journal of the ACM (JACM)*, 60(1):4, 2013.

B. Bhattacharya and G. Valiant. Testing closeness with unequal sized samples. In *Advances in Neural Information Processing Systems*, pages 2611–2619, 2015.

O. Bousquet, S. Boucheron, and G. Lugosi. Introduction to statistical learning theory. In *Advanced lectures on machine learning*, pages 169–207. Springer, 2004.

F. Chen, L. Lovász, and I. Pak. Lifting markov chains to speed up mixing. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 275–281. ACM, 1999.

T.-L. Chen and C.-R. Hwang. Accelerating reversible markov chains. *Statistics & Probability Letters*, 83(9):1956–1962, 2013.

G. E. Cho and C. D. Meyer. Comparison of perturbation bounds for the stationary distribution of a markov chain. *Linear Algebra and its Applications*, 335(1-3):137–150, 2001.

R. Combes and M. Touati. Computationally efficient estimation of the spectral gap of a markov chain. *arXiv preprint arXiv:1806.06047*, 2018.

C. Daskalakis, N. Dikkala, and N. Gravin. Testing symmetric markov chains from a single trajectory. *arXiv preprint arXiv:1704.06850*, 2017.

P. Diaconis, S. Holmes, and R. M. Neal. Analysis of a nonreversible markov chain sampler. *Annals of Applied Probability*, pages 726–752, 2000.

J. A. Fill. Eigenvalue bounds on convergence to stationarity for nonreversible markov chains, with an application to the exclusion process. *The annals of applied probability*, pages 62–87, 1991.

D. Gamarnik. Extension of the pac framework to finite and countable markov chains. *IEEE Transactions on Information Theory*, 49(1):338–345, 2003.

M. Hildebrand. Rates of convergence for a non-reversible markov chain sampler. *preprint*, 1997.

D. J. Hsu, A. Kontorovich, and C. Szepesvári. Mixing time estimation in reversible markov chains from a single sample path. In *Advances in neural information processing systems*, pages 1459–1467, 2015.

D. J. Hsu, A. Kontorovich, D. A. Levin, Y. Peres, and C. Szepesvári. Mixing time estimation in reversible markov chains from a single sample path. *CoRR*, abs/1708.07367, 2017. URL `http://arxiv.org/abs/1708.07367`.

S. Kaniel. Estimates for some computational techniques in linear algebra. *Mathematics of Computation*, 20(95):369–378, 1966.

R. L. Karandikar and M. Vidyasagar. Rates of uniform convergence of empirical means with mixing processes. *Statistics & probability letters*, 58(3):297–307, 2002.

D. Kazakos. The Bhattacharyya distance and detection between markov chains. *IEEE Transactions on Information Theory*, 24(6):747–754, 1978.

M. J. Kearns, Y. Mansour, D. Ron, R. Rubinfeld, R. E. Schapire, and L. Sellie. On the learnability of discrete distributions. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada*, pages 273–282, 1994. doi: 10.1145/195058.195155. URL `http://doi.acm.org/10.1145/195058.195155`.

J. Kuczyński and H. Woźniakowski. Estimating the largest eigenvalue by the power and lanczos algorithms with a random start. *SIAM journal on matrix analysis and applications*, 13(4): 1094–1122, 1992.

F. Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th international symposium on symbolic and algebraic computation*, pages 296–303. ACM, 2014.

D. A. Levin and Y. Peres. Estimating the spectral gap of a reversible markov chain from a short trajectory. *arXiv preprint arXiv:1612.05330*, 2016.

D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times, second edition.* American Mathematical Soc., 2009.

C. D. Meyer, Jr. The role of the group generalized inverse in the theory of finite markov chains. *Siam Review*, 17(3):443–464, 1975.

M. Mohri and A. Rostamizadeh. Stability bounds for non-iid processes. In *Advances in Neural Information Processing Systems*, pages 1025–1032, 2008.

M. Mohri and A. Rostamizadeh. Rademacher complexity bounds for non-iid processes. In *Advances in Neural Information Processing Systems*, pages 1097–1104, 2009.

R. Montenegro and P. Tetali. Mathematical aspects of mixing times in markov chains. *Foundations and Trends® in Theoretical Computer Science*, 1(3):237–354, 2006.

R. M. Neal. Improving asymptotic variance of mcmc estimators: Non-reversible chains are better. *arXiv preprint math/0407281*, 2004.

C. C. Paige. *The computation of eigenvalues and eigenvectors of very large sparse matrices.* PhD thesis, University of London, 1971.

D. Paulin. Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electronic Journal of Probability*, 20, 2015.

Q. Qin, J. P. Hobert, and K. Khare. Estimating the spectral gap of a trace-class markov operator. *arXiv preprint arXiv:1704.00850*, 2017.

Y. Saad. On the rates of convergence of the lanczos and the block-lanczos methods. *SIAM Journal on Numerical Analysis*, 17(5):687–706, 1980.

C. R. Shalizi and A. Kontorovich. Predictive PAC learning and process decompositions. In *Neural Information Processing Systems (NIPS)*, 2013.

I. Steinwart and A. Christmann. Fast learning from non-iid observations. In *Advances in neural information processing systems*, pages 1768–1776, 2009.

I. Steinwart, D. Hush, and C. Scovel. Learning from dependent observations. *Journal of Multivariate Analysis*, 100(1):175–194, 2009.

Y. Sun, J. Schmidhuber, and F. J. Gomez. Improving the asymptotic performance of markov chain monte-carlo by inserting vortices. In *Advances in Neural Information Processing Systems*, pages 2235–2243, 2010.

H. Suwa and S. Todo. Markov chain monte carlo method without detailed balance. *Physical review letters*, 105(12):120603, 2010.

J. Tropp. Freedman's inequality for matrix martingales. *Electronic Communications in Probability*, 16:262–270, 2011.

A. B. Tsybakov. Introduction to nonparametric estimation, 2009. URL `https://doi.org/10.1007/b13794`. Revised and extended from the 2004 French original, Translated by Vladimir Zaiats.

K. S. Turitsyn, M. Chertkov, and M. Vucelja. Irreversible monte carlo algorithms for efficient sampling. *Physica D: Nonlinear Phenomena*, 240(4-5):410–414, 2011.

M. Vucelja. Lifting–a nonreversible markov chain monte carlo algorithm. *arXiv preprint arXiv:1412.8762*, 2014.

B. Waggoner. L_p testing and learning of discrete distributions. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS 2015, Rehovot, Israel, January 11-13, 2015*, pages 347–356, 2015. doi: 10.1145/2688073.2688095. URL `http://doi.acm.org/10.1145/2688073.2688095`.

G. Wolfer and A. Kontorovich. Minimax learning of ergodic markov chains. In *Algorithmic Learning Theory, ALT*, 2019. URL `https://arxiv.org/abs/1809.05014`.

B. Yu. Rates of convergence for empirical processes of stationary mixing sequences. *The Annals of Probability*, pages 94–116, 1994.

A Confidence intervals

The main algorithmic procedure whose description is given at Algorithm 1 yields the following fully empirical confidence intervals.

Theorem A.1. *Let $C_{\mathcal{K}} \leq 192$ be a universal constant and define*

$$\tau_{\delta,m} = \inf \left\{ t > 0 : \left(1 + \lceil \ln(2m/t) \rceil_+\right)(d+1)e^{-t} \leq \delta \right\}.$$

Then $\tau_{\delta,m} = \mathcal{O}\left(\ln\left(\frac{d \ln m}{\delta}\right)\right)$ and with probability at least $1-\delta$,

$$\left|\widehat{\gamma}^{(\alpha)}_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}}\right| \leq \frac{1}{K} + 2 \max_{k \in [K]} \left\{ \frac{1}{k^2} \sum_{r=0}^{k-1} \left(\hat{a}^{(k,r,\alpha)} + 2\hat{c}^{(k,r,\alpha)} + \left(\hat{c}^{(k,r,\alpha)}\right)^2 \right) \right\},$$

$$\text{where} \begin{cases} \hat{a}^{(k,r,\alpha)} & = \sqrt{d} \frac{N_{\max}^{(k,r)} + d\alpha}{N_{\min}^{(k,r)} + d\alpha} \hat{d}^{(k,r,\alpha)} \\ \hat{b}^{(k,r,\alpha)} & = \frac{C_{\mathcal{K}}}{\gamma_{\mathsf{ps}}\left(\widehat{\boldsymbol{M}}^{(k,r,\alpha)}\right)} \ln\left(2\sqrt{\frac{2(\lfloor (m-r)/k \rfloor + d^2\alpha)}{N_{\min}^{(k,r)} + d\alpha}}\right) \hat{d}^{(k,r,\alpha)} \\ \hat{c}^{(k,r,\alpha)} & = \frac{1}{2} \max \bigcup_{i \in [d]} \left\{ \frac{\hat{b}^{(k,r,\alpha)}}{\frac{N_i^{(k,r)} + d\alpha}{\lfloor (m-r)/k \rfloor + d^2\alpha}}, \frac{\hat{b}^{(k,r,\alpha)}}{\left[\frac{N_i^{(k,r)} + d\alpha}{\lfloor (m-r)/k \rfloor + d^2\alpha} - \hat{b}^{(k,r,\alpha)}\right]_+} \right\} \\ \hat{d}^{(k,r,\alpha)} & = 4\tau_{\delta/d, \lfloor (m-r)/k \rfloor} \sqrt{\frac{d}{N_{\min}^{(k,r)} + d\alpha}} + \frac{2\alpha d}{N_{\min}^{(k,r)} + d\alpha} \end{cases}. \tag{A.1}$$

The interval widths asymptotically behave as

$$\sqrt{m}\left|\hat{\pi}_\star^{(\alpha)} - \pi_\star\right| = \tilde{\mathcal{O}}\left(\frac{\sqrt{d}}{\gamma_{\mathsf{ps}}\sqrt{\pi_\star}}\right), \quad \sqrt{m}\left|\widehat{\gamma}^{(\alpha)}_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}}\right| = \tilde{\mathcal{O}}\left(\frac{1}{K} + \sqrt{\frac{d}{\pi_\star}}\left(\sqrt{d}\left\|\boldsymbol{M}\right\|_{\boldsymbol{\pi}} + \frac{1}{\gamma_{\mathsf{ps}}\pi_\star}\right)\right). \tag{A.2}$$

In the reversible case,

$$\left|\widehat{\gamma}_\star^{(\alpha)} - \gamma_\star\right| \leq \hat{a}^{(\alpha)} + 2\hat{c}^{(\alpha)} + \left(\hat{c}^{(\alpha)}\right)^2 \qquad \left|\widehat{\pi}_\star^{(\alpha)} - \widehat{\pi}_\star\right| \leq \hat{b}^{(\alpha)},$$

$$\text{where} \begin{cases} \hat{a}^{(\alpha)} &= \sqrt{d}\frac{N_{\max}+d\alpha}{N_{\min}+d\alpha}\hat{d}^{(\alpha)} \\ \hat{b}^{(\alpha)} &= \frac{C_{\mathcal{K}}}{\widehat{\gamma}_\star^{(\alpha)}} \ln\left(2\sqrt{\frac{2(m+d^2\alpha)}{N_{\min}+d\alpha}}\right)\hat{d}^{(\alpha)} \\ \hat{c}^{(\alpha)} &= \frac{1}{2}\max\bigcup_{i\in[d]}\left\{\frac{\hat{b}^{(\alpha)}}{\frac{N_i+d\alpha}{m+d^2\alpha}}, \frac{\hat{b}^{(\alpha)}}{\left[\frac{N_i+d\alpha}{m+d^2\alpha} - \hat{b}^{(\alpha)}\right]_+}\right\} \\ \hat{d}^{(\alpha)} &= 4\tau_{\delta/d,\lfloor m\rfloor}\sqrt{\frac{d}{N_{\min}+d\alpha}} + \frac{2\alpha d}{N_{\min}+d\alpha} \end{cases}. \tag{A.3}$$

The interval widths asymptotically behave as

$$\sqrt{m}\left|\widehat{\pi}_\star^{(\alpha)} - \pi_\star\right| = \tilde{\mathcal{O}}\left(\frac{\sqrt{d}}{\gamma_\star\sqrt{\pi_\star}}\right), \qquad \sqrt{m}\left|\widehat{\gamma}_\star^{(\alpha)} - \gamma_\star\right| = \tilde{\mathcal{O}}\left(\sqrt{\frac{d}{\pi_\star}}\left(\sqrt{d}\left\|\boldsymbol{M}\right\|_{\boldsymbol{\pi}} + \frac{1}{\gamma_\star\pi_\star}\right)\right). \tag{A.4}$$

B Proofs

B.1 Minimum stationary probability

B.1.1 Proof of Theorem 5.1

Let $(X_1, \ldots, X_m) \sim (\boldsymbol{M}, \boldsymbol{\mu})$ be a d-state Markov chain, with pseudo-spectral gap γ_{ps} and stationary distribution $\boldsymbol{\pi}$ minorized by $\pi_\star$. Our estimator for $\pi_\star$ is defined as the minimum of the empirical stationary distribution, or more formally,

$$\widehat{\pi}_\star \triangleq \min_{i\in[d]} \widehat{\boldsymbol{\pi}}(i) = \min_{i\in[d]} \widehat{\boldsymbol{\pi}}(i)\frac{1}{m}\left|t \in [m] : X_t = i\right|. \tag{B.1}$$

Without loss of generality, suppose that $\pi_\star = \pi_1 \leq \pi_2 \leq \cdots \leq \pi_d$ (renumber states if needed). A Bernstein-type inequality [Paulin, 2015, Theorem 3.4], combined with Paulin [2015, Theorem 3.10], yields that for all $i \in [d]$ and $t > 0$,

$$\mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left|\widehat{\pi}_i - \pi_i\right| \geq t\right) \leq \sqrt{2\left\|\boldsymbol{\mu}/\boldsymbol{\pi}\right\|_{2,\boldsymbol{\pi}}}\exp\left(-\frac{t^2\gamma_{\mathsf{ps}}m}{16(1+1/(m\gamma_{\mathsf{ps}}))\pi_i(1-\pi_i)+40t}\right). \tag{B.2}$$

Taking $m > \frac{1}{\gamma_{\mathsf{ps}}}$ and putting

$$t_m = \frac{\log\left(\frac{d}{\delta}\sqrt{2\left\|\boldsymbol{\mu}/\boldsymbol{\pi}\right\|_{2,\boldsymbol{\pi}}}\right)}{\gamma_{\mathsf{ps}}m}, \qquad t = \sqrt{32\pi_i t_m} + 40t_m,$$

yields (via a union bound) $\mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\widehat{\boldsymbol{\pi}} - \boldsymbol{\pi}\right\|_\infty \geq t\right) \leq \delta$. We claim that

$$\forall i \in [d], \left|\widehat{\pi}_i - \pi_i\right| < \sqrt{32\pi_i t_m} + 40t_m \implies \left|\widehat{\pi}_\star - \pi_\star\right| < 8\sqrt{t_m\pi_\star} + 136t_m. \tag{B.3}$$

Indeed, Let $i_\star$ be such that $\widehat{\pi}_\star = \widehat{\pi}_{i_\star}$, and suppose that $\forall i \in [d] : \left|\widehat{\pi}_i - \pi_i\right| < \sqrt{32\pi_i t_m} + 40t_m$. Since $\widehat{\pi}_\star \leq \widehat{\pi}_1$, we have

$$\widehat{\pi}_\star - \pi_\star \leq \widehat{\pi}_1 - \pi_1 \leq \sqrt{32\pi_\star t_m} + 40t_m \leq \pi_\star + 48t_m, \tag{B.4}$$

where the last inequality follows from the AM-GM inequality. Furthermore, $a \leq b\sqrt{a} + c \implies a \leq b^2 + b\sqrt{c} + c$ [Bousquet et al., 2004] and

$$\pi_{i_\star} \leq \sqrt{32t_m}\sqrt{\pi_{i_\star}} + (\widehat{\pi}_{i_\star} + 40t_m). \tag{B.5}$$

Thus,

$$\begin{aligned}\pi_{i_\star} &\leq 32t_m + (\widehat{\pi}_\star + 40t_m) + \sqrt{32t_m}\sqrt{\widehat{\pi}_\star + 40t_m} = \widehat{\pi}_\star + 72t_m + \sqrt{32t_m(\widehat{\pi}_\star + 40t_m)} \\ &\leq \widehat{\pi}_\star + \sqrt{32t_m(2\pi_\star + 88t_m)} + 72t_m \leq \widehat{\pi}_\star + \sqrt{32 \cdot 88t_m^2} + \sqrt{64t_m\pi_\star} + 72t_m \\ &= \widehat{\pi}_\star + 16\sqrt{11}t_m + 8\sqrt{t_m\pi_\star} + 72t_m \leq \widehat{\pi}_\star + 8\sqrt{t_m\pi_\star} + 136t_m m\end{aligned} \tag{B.6}$$

and therefore,

$$\pi_\star - \widehat{\pi}_\star \leq \pi_{i_\star} - \widehat{\pi}_{i_\star} \leq 8\sqrt{t_m\pi_\star} + 136t_m, \tag{B.7}$$

whence

$$|\widehat{\pi}_\star - \pi_\star| \leq 8\sqrt{t_m\pi_\star} + 136t_m. \tag{B.8}$$

A direct computation shows that

$$m \geq \frac{16^2\pi_\star \log\left(\frac{d}{\delta}\sqrt{2\left\|\boldsymbol{\mu}/\boldsymbol{\pi}\right\|_{2,\boldsymbol{\pi}}}\right)}{\gamma_{\mathsf{ps}}\varepsilon^2} \implies 8\sqrt{t_m\pi_\star} \leq \frac{\varepsilon}{2} \tag{B.9}$$

and

$$m \geq \frac{2 \cdot 136 \log\left(\frac{d}{\delta}\sqrt{2\left\|\boldsymbol{\mu}/\boldsymbol{\pi}\right\|_{2,\boldsymbol{\pi}}}\right)}{\gamma_{\mathsf{ps}}\varepsilon} \implies 136t_m \leq \frac{\varepsilon}{2}, \tag{B.10}$$

so that for $m \geq \frac{\log\left(\frac{d}{\delta}\sqrt{2\left\|\boldsymbol{\mu}/\boldsymbol{\pi}\right\|_{2,\boldsymbol{\pi}}}\right)}{\gamma_{\mathsf{ps}}\varepsilon} \max\left\{272, \frac{256\pi_\star}{\varepsilon}\right\}$, we have $|\widehat{\pi}_\star - \pi_\star| < \varepsilon$ with probability at least $1 - \delta$. The theorem then follows by choosing the precision to be $\varepsilon\pi_\star$.

B.1.2 Proof of Theorem 5.2

We first prove the claim for absolute error in the regime $2\varepsilon < \pi_\star < \gamma_{\mathsf{ps}}$, and at the end obtain the claimed result via an absolute-to-relative conversion.

Consider a star-shaped class of Markov chains with a single "hub" connected to "spoke" states, each of which can only transition to itself and to the hub. Namely, we construct the family of $(d+1)$-state Markov chains for $d \in \mathbb{N}, d \geq 4$,

$$\mathcal{S}_d = \left\{\boldsymbol{S}_\alpha(\boldsymbol{D}) : 0 < \alpha < 1, \boldsymbol{D} = (p_1, \ldots, p_d) \in \Delta_d\right\}, \text{ where } \boldsymbol{S}_\alpha(\boldsymbol{D}) = \begin{pmatrix} \alpha & (1-\alpha)p_1 & \cdots & (1-\alpha)p_d \\ \alpha & 1-\alpha & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \alpha & 0 & \cdots & 1-\alpha \end{pmatrix}.$$

Notice that $\boldsymbol{\pi}(\boldsymbol{D}) = (\alpha, (1-\alpha)p_1, \ldots, (1-\alpha)p_d)$ is a stationary distribution for $\boldsymbol{S}_\alpha(\boldsymbol{D})$, that $\boldsymbol{S}_\alpha(\boldsymbol{D})$ is reversible and that the spectrum of $\boldsymbol{S}_\alpha(\boldsymbol{D})$ consists of $\lambda_0 = 1$, $\lambda_\alpha = 1 - \alpha$ (of multiplicity $d-1$), and $\lambda_d = 0$. Thus, the absolute spectral gap is $\gamma_\star(\boldsymbol{S}_\alpha(\boldsymbol{D})) = \alpha$, and the pseudo-spectral gap $\gamma_{\mathsf{ps}} = O(\alpha)$, by Lemma C.1. We apply Le Cam's two-point method as follows. Take $\boldsymbol{D}$ and $\boldsymbol{D}_\varepsilon$ in Δ_d defined as follows,

$$\boldsymbol{D} := \left(\beta, \beta, \frac{1-2\beta}{d-2}, \ldots, \frac{1-2\beta}{d-2}\right), \boldsymbol{D}_\varepsilon := \left(\beta + 2\varepsilon, \beta - 2\varepsilon, \frac{1-2\beta}{d-2}, \ldots, \frac{1-2\beta}{d-2}\right), \tag{B.11}$$

with $2\varepsilon < \beta < 1/d$. For readability, we will abbreviate in this section

$$\boldsymbol{S}_\alpha := \boldsymbol{S}_\alpha(\boldsymbol{D}), \boldsymbol{\pi} := \boldsymbol{\pi}(\boldsymbol{D}), \boldsymbol{S}_{\alpha,\varepsilon} := \boldsymbol{S}_\alpha(\boldsymbol{D}_\varepsilon), \boldsymbol{\pi}_\varepsilon := \boldsymbol{\pi}(\boldsymbol{D}_\varepsilon), \mathbf{P}(\cdot) := \mathbf{P}_{\boldsymbol{S}_\alpha,\pi}(\cdot), \mathbf{P}_\varepsilon(\cdot) := \mathbf{P}_{\boldsymbol{S}_{\alpha,\varepsilon},\pi_\varepsilon}(\cdot).$$

Consider the two stationary Markov chains $(\boldsymbol{S}_\alpha, \pi)$ and $(\boldsymbol{S}_{\alpha,\varepsilon}, \pi_\varepsilon)$ for $\alpha > \beta$. First notice that

$$\left| \min_{i \in [d+1]} \pi_\varepsilon(i) - \min_{i \in [d+1]} \pi(i) \right| = 2\varepsilon.$$

We now exhibit a tensorization property of the KL divergence between trajectories of length m sampled from the two chains. Let $\boldsymbol{X}_1^m \sim (\boldsymbol{S}_\alpha, \pi)$ and $\boldsymbol{Y}_1^m \sim (\boldsymbol{S}_{\alpha,\varepsilon}, \pi_\varepsilon)$, from the definition of the KL divergence and the Markov property,

$$\begin{aligned}
D_{\mathrm{KL}}\left(\boldsymbol{Y}_1^m || \boldsymbol{X}_1^m\right) &= \sum_{(z_1,\dots,z_m)\in[d]^m} \mathbf{P}_\varepsilon\left(\boldsymbol{Y}_1^m = z_1^m\right) \ln\left(\frac{\mathbf{P}_\varepsilon\left(\boldsymbol{Y}_1^m = z_1^m\right)}{\mathbf{P}\left(\boldsymbol{X}_1^m = z_1^m\right)}\right) \\
&= \underbrace{\sum_{(z_1,\dots,z_{m-1})\in[d]^{m-1}} \mathbf{P}_\varepsilon\left(\boldsymbol{Y}_1^{m-1} = z_1^{m-1}\right) \ln\left(\frac{\mathbf{P}_\varepsilon\left(\boldsymbol{Y}_1^{m-1} = z_1^{m-1}\right)}{\mathbf{P}\left(\boldsymbol{X}_1^{m-1} = z_1^{m-1}\right)}\right)}_{D_{\mathrm{KL}}\left(\boldsymbol{Y}_1^{m-1}||\boldsymbol{X}_1^{m-1}\right)} \underbrace{\sum_{z_m\in[d]} \boldsymbol{S}_{\alpha,\varepsilon}(z_{m-1}, z_m)}_{=1} \\
&+ \sum_{z_1^{m-2}\in[d]^{m-2}} \mathbf{P}_\varepsilon\left(\boldsymbol{Y}_1^{m-2} = z_1^{m-2}\right) \sum_{z_{m-1}\in[d]} \boldsymbol{S}_{\alpha,\varepsilon}(z_{m-2}, z_{m-1}) \underbrace{\sum_{z_m\in[d]} \boldsymbol{S}_{\alpha,\varepsilon}(z_{m-1}, z_m) \ln\left(\frac{\boldsymbol{S}_{\alpha,\varepsilon}(z_{m-1}, z_m)}{\boldsymbol{S}_\alpha(z_{m-1}, z_m)}\right)}_{=\mathbf{1}\{z_{m-1}=1\}D_{\mathrm{KL}}(\boldsymbol{D}_\varepsilon||\boldsymbol{D})},
\end{aligned} \tag{B.12}$$

and as by structural property of the chains of the class $\forall z_{m-2} \in [d], \boldsymbol{S}_{\alpha,\varepsilon}(z_{m-2}, 1) = \alpha$,

$$D_{\mathrm{KL}}\left(\boldsymbol{Y}_1^m || \boldsymbol{X}_1^m\right) = D_{\mathrm{KL}}\left(\boldsymbol{Y}_1^{m-1} || \boldsymbol{X}_1^{m-1}\right) + \alpha D_{\mathrm{KL}}\left(\boldsymbol{D}_\varepsilon || \boldsymbol{D}\right) \underbrace{\sum_{(z_1,\dots,z_{m-2})\in[d]^{m-2}} \mathbf{P}_\varepsilon\left(\boldsymbol{Y}_1^{m-2} = z_1^{m-2}\right)}_{=1}, \tag{B.13}$$

so by induction and stationarity, $D_{\mathrm{KL}}\left(\boldsymbol{Y}_1^m || \boldsymbol{X}_1^m\right) = m\alpha D_{\mathrm{KL}}\left(\boldsymbol{D}_\varepsilon || \boldsymbol{D}\right)$. It remains to compute the KL divergence between the two distributions,

$$D_{\mathrm{KL}}\left(\boldsymbol{D}_\varepsilon || \boldsymbol{D}\right) = (\beta + 2\varepsilon) \ln\left(1 + \frac{2\varepsilon}{\beta}\right) + (\beta - 2\varepsilon) \ln\left(1 + \frac{-2\varepsilon}{\beta}\right) \leq \frac{8\varepsilon^2}{\beta}. \tag{B.14}$$

Denote by $\mathcal{M}_{d,\gamma_{\mathsf{ps}},\pi_\star}$ the collection of all d-state Markov chains whose stationary distribution is minorized by $\pi_\star$ and whose pseudo-spectral gap is at least γ_{ps}, and define the minimax risk as

$$\mathcal{R}_m = \inf_{\hat{\pi}_\star} \sup_{\boldsymbol{M}\in\mathcal{M}_{d,\gamma_{\mathsf{ps}},\pi_\star}} \mathbf{P}_{\boldsymbol{M}}\left(|\hat{\pi}_\star - \pi_\star| > \varepsilon\right), \tag{B.15}$$

then from the KL divergence version of Le Cam's theorem [Tsybakov, 2009, Chapter 2],

$$\mathcal{R}_m \geq \frac{1}{4} \exp\left(-D_{\mathrm{KL}}\left(\boldsymbol{Y}_1^m || \boldsymbol{X}_1^m\right)\right) \geq \frac{1}{4} \exp\left(-\frac{8\alpha\varepsilon^2 m}{\beta}\right) \geq \frac{1}{4} \exp\left(-\frac{8\gamma_{\mathsf{ps}}\varepsilon^2 m}{\pi_\star}\right). \tag{B.16}$$

Hence, for $m \leq \frac{\pi_\star \ln\left(\frac{1}{4\delta}\right)}{8\gamma_{\mathsf{ps}}\varepsilon^2}$, we have $\mathcal{R}_m \geq \delta$, and so $m = \Omega\left(\frac{\pi_\star \ln\left(\frac{1}{\delta}\right)}{\gamma_{\mathsf{ps}}\varepsilon^2}\right) = \tilde{\Omega}\left(\frac{\pi_\star}{\gamma_{\mathsf{ps}}\varepsilon^2}\right)$ is a lower bound for the problem, in the $\gamma_{\mathsf{ps}} > \pi_\star$ regime, up to absolute error. When taking the accuracy to be $\varepsilon\pi_\star$ instead, the previous bound becomes $\tilde{\Omega}\left(\frac{1}{\gamma_{\mathsf{ps}}\pi_\star\varepsilon^2}\right)$, and the fact that the proof exclusively makes use of a reversible family confirms that the upper bound derived in Hsu et al. [2017] is minimax optimal up to a logarithmic a factor, in the parameters $\pi_\star, \gamma_{\mathsf{ps}}$ and ε.

B.2 Pseudo-spectral gap

B.2.1 Proof of Theorem 5.3

In this section, we analyze our point estimator for the pseudo-spectral gap.

Reduction to a maximum over a finite number of estimators. Recall the definitions from (7.3):

$$\gamma_k^\dagger \triangleq \gamma\left(\left(\boldsymbol{M}^\dagger\right)^k \boldsymbol{M}^k\right), \qquad \gamma_{\mathsf{ps}[K]} \triangleq \max_{k\in[K]}\left\{\frac{\gamma_k^\dagger}{k}\right\}.$$

It follows from the definition of $\boldsymbol{M}^\dagger$ that $\boldsymbol{\pi}$ is the stationary distribution of $\left(\boldsymbol{M}^\dagger\right)^k \boldsymbol{M}^k$ for all $k \in \mathbb{N}$. We denote by $\widehat{\gamma}_{\mathsf{ps}[K]}$ the empirical estimator for $\gamma_{\mathsf{ps}[K]}$:

$$\widehat{\gamma}_{\mathsf{ps}[K]} \triangleq \max_{k\in[K]}\left\{\frac{\widehat{\gamma}_k^\dagger(X_1,\dots,X_m)}{k}\right\}, \tag{B.17}$$

where $\widehat{\gamma}_k^\dagger(X_1,\dots,X_m)$ is an estimator for $\gamma_k^\dagger$ to be defined below. From the triangle inequality,

$$\mathbf{P}\left(\left|\widehat{\gamma}_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}}\right| > \varepsilon\right) \le \mathbf{P}\left(\left|\widehat{\gamma}_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}[K]}\right| + \left|\gamma_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}}\right| > \varepsilon\right). \tag{B.18}$$

By taking a maximum over a larger set, $\left|\gamma_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}}\right| = \gamma_{\mathsf{ps}} - \gamma_{\mathsf{ps}[K]} \le \max_{k\in\mathbb{N}\setminus[K]}\left\{\frac{\gamma_k^\dagger}{k}\right\}$, and since $\gamma_k^\dagger \le 1$ for all $k \in \mathbb{N}$, we have $\left|\gamma_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}}\right| \le \frac{1}{K}$. Thus, for $K \ge \frac{2}{\varepsilon}$,

$$\mathbf{P}\left(\left|\widehat{\gamma}_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}}\right| > \varepsilon\right) \le \mathbf{P}\left(\left|\widehat{\gamma}_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}[K]}\right| + \frac{1}{K} > \varepsilon\right) \le \mathbf{P}\left(\left|\widehat{\gamma}_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}[K]}\right| > \frac{\varepsilon}{2}\right), \tag{B.19}$$

and from another application of the triangle inequality, for $K \ge \frac{2}{\varepsilon}$,

$$\mathbf{P}\left(\left|\widehat{\gamma}_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}}\right| > \varepsilon\right) \le \mathbf{P}\left(\max_{k\in[K]}\left\{\frac{\left|\gamma_k^\dagger - \widehat{\gamma}_k^\dagger\right|}{k}\right\} > \frac{\varepsilon}{2}\right) \le \sum_{k=1}^{\lceil\frac{2}{\varepsilon}\rceil}\mathbf{P}\left(\left|\gamma_k^\dagger - \widehat{\gamma}_k^\dagger\right| > \frac{k\varepsilon}{2}\right). \tag{B.20}$$

Reduction to controlling spectral norms. Recall that $\gamma_k^\dagger = \gamma((\boldsymbol{M}^k)^\dagger \boldsymbol{M}^k)$ is the spectral gap of the multiplicative reversiblization of $\boldsymbol{M}^k$, the k-skipped Markov chain associated with $\boldsymbol{M}$. We now introduce natural estimators for $\boldsymbol{M}^k$ and $\boldsymbol{\pi}$,

$$\widehat{\boldsymbol{M}}^{(k)}(i,j) \triangleq \frac{N_{ij}^{(k)}}{N_i^{(k)}} \text{ and } \widehat{\boldsymbol{\pi}}^{(k)}(i) \triangleq \frac{N_i^{(k)}}{m}, \tag{B.21}$$

where $N_i^{(k)}$ and $N_{ij}^{(k)}$ are defined in (4.2, 4.3). It is readily verified that $\widehat{\boldsymbol{\pi}}^{(k)}$ is the stationary distribution of $\widehat{\boldsymbol{M}}^{(k)}$, and so $(\widehat{\boldsymbol{M}}^{(k)})^\dagger \widehat{\boldsymbol{M}}^{(k)}$ is a natural estimator for $(\boldsymbol{M}^k)^\dagger \boldsymbol{M}^k$. We also introduce $\widehat{\boldsymbol{L}}^{(k)} = (\widehat{\boldsymbol{D}}_{\boldsymbol{\pi}}^{(k)})^{1/2} \widehat{\boldsymbol{M}}^{(k)} (\widehat{\boldsymbol{D}}_{\boldsymbol{\pi}}^{(k)})^{-1/2}$, and will later show that the event $\min_{(k,i)\in[K]\times[d]} N_i^{(k)} > 0$ occurs with high probability for our sampling regime, in which case the above quantities will be well-defined, making smoothing unnecessary. Notice $(\boldsymbol{M}^k)^\dagger \boldsymbol{M}^k = \boldsymbol{D}_{\boldsymbol{\pi}}^{-1/2} (\boldsymbol{L}^k)^\intercal \boldsymbol{L}^k \boldsymbol{D}_{\boldsymbol{\pi}}^{1/2}$, and $(\boldsymbol{L}^k)^\intercal \boldsymbol{L}^k$ is symmetric, which makes Weyls' inequality applicable:

$$\left|\widehat{\gamma}_k^\dagger - \gamma_k^\dagger\right| \le \left\|(\widehat{\boldsymbol{L}}^{(k)})^\intercal \widehat{\boldsymbol{L}}^{(k)} - (\boldsymbol{L}^k)^\intercal \boldsymbol{L}^k\right\|_2.$$

By the triangle inequality and sub-multiplicativity of the spectral norm,

$$\left|\widehat{\gamma}_k^{\dagger} - \gamma_k^{\dagger}\right| \leq \left\|\left(\boldsymbol{L}^k\right)^{\intercal}\right\|_2 \left\|\widehat{\boldsymbol{L}}^{(k)} - \boldsymbol{L}^k\right\|_2 + \left\|\left(\widehat{\boldsymbol{L}}^{(k)}\right)^{\intercal} - \left(\boldsymbol{L}^k\right)^{\intercal}\right\|_2 \left\|\widehat{\boldsymbol{L}}^{(k)}\right\|_2 \leq 2\left\|\boldsymbol{L}^k - \widehat{\boldsymbol{L}}^{(k)}\right\|_2, \tag{B.22}$$

where for the second inequality we invoked the Perron-Frobenius theorem ($\sqrt{\boldsymbol{\pi}}$ is an eigenvector associated to eigenvalue 1 for $(\boldsymbol{L}^k)^{\intercal}\boldsymbol{L}^k$, $k \in [K]$ and $\left\|\boldsymbol{L}^k\right\|_2 = 1$; the same holds for its empirical version). We continue by decomposing $\widehat{\boldsymbol{L}}^{(k)} - \boldsymbol{L}^k$ into more manageable quantities,

$$\begin{aligned} \widehat{\boldsymbol{L}}^{(k)} - \boldsymbol{L}^k &= \mathcal{E}_{\boldsymbol{M}}^{(k)} + \mathcal{E}_{\boldsymbol{\pi},1}^{(k)}\boldsymbol{L}^k + \boldsymbol{L}^k\mathcal{E}_{\boldsymbol{\pi},2}^{(k)} + \mathcal{E}_{\boldsymbol{\pi},1}^{(k)}\boldsymbol{L}^k\mathcal{E}_{\boldsymbol{\pi},2}^{(k)} \\ \text{where } \mathcal{E}_{\boldsymbol{M}}^{(k)} &= \left(\widehat{\boldsymbol{D}}_{\boldsymbol{\pi}}^{(k)}\right)^{1/2}\left(\widehat{\boldsymbol{M}}^{(k)} - \boldsymbol{M}^k\right)\left(\widehat{\boldsymbol{D}}_{\boldsymbol{\pi}}^{(k)}\right)^{-1/2}, \\ \mathcal{E}_{\boldsymbol{\pi},1}^{(k)} &= \left(\widehat{\boldsymbol{D}}_{\boldsymbol{\pi}}^{(k)}\right)^{1/2}\boldsymbol{D}_{\boldsymbol{\pi}}^{-1/2} - \boldsymbol{I}, \qquad \mathcal{E}_{\boldsymbol{\pi},2}^{(k)} = \boldsymbol{D}_{\boldsymbol{\pi}}^{1/2}\left(\widehat{\boldsymbol{D}}_{\boldsymbol{\pi}}^{(k)}\right)^{-1/2} - \boldsymbol{I}. \end{aligned} \tag{B.23}$$

Writing $\left\|\mathcal{E}_{\boldsymbol{\pi}}^{(k)}\right\|_2 = \max\left\{\left\|\mathcal{E}_{\boldsymbol{\pi},1}^{(k)}\right\|_2, \left\|\mathcal{E}_{\boldsymbol{\pi},2}^{(k)}\right\|_2\right\}$, from the sub-multiplicativity of the spectral norm and another application of the Perron-Frobenius theorem,

$$\left\|\widehat{\boldsymbol{L}}^{(k)} - \boldsymbol{L}^k\right\|_2 \leq \left\|\mathcal{E}_{\boldsymbol{M}}^{(k)}\right\|_2 + 2\left\|\mathcal{E}_{\boldsymbol{\pi}}^{(k)}\right\|_2 + \left\|\mathcal{E}_{\boldsymbol{\pi}}^{(k)}\right\|_2^2. \tag{B.24}$$

Consider the event

$$\mathcal{K}_{1/2}^K = \left\{\max_{(k,i)\in[K]\times[d]}\left|N_i^{(k)} - \mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}\left[N_i^{(k)}\right]\right| \leq \frac{1}{2}\mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}\left[N_i^{(k)}\right]\right\}.$$

Applying the law of total probability,

$$\begin{aligned} \mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left|\widehat{\gamma}_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}}\right| > \varepsilon\right) \leq \mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\overline{\mathcal{K}_{1/2}^K}\right) + \sum_{k=1}^{\lceil \frac{2}{\varepsilon} \rceil}\Bigg(\mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\mathcal{E}_{\boldsymbol{M}}^{(k)}\right\|_2 > \frac{\varepsilon k}{8} \text{ and } \mathcal{K}_{1/2}^K\right) \\ + \mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(2\left\|\mathcal{E}_{\boldsymbol{\pi}}^{(k)}\right\|_2 + \left\|\mathcal{E}_{\boldsymbol{\pi}}^{(k)}\right\|_2^2 > \frac{\varepsilon k}{8}\right)\Bigg). \end{aligned} \tag{B.25}$$

A matrix martingale approach. Consider for $k \in [K]$, the following sequence of random matrices $\boldsymbol{Y}_1^{(k)} = \boldsymbol{0}, \boldsymbol{Y}_t^{(k)} = \left[\boldsymbol{1}\left\{X_{t-1}^{(k)} = i\right\}\left(\boldsymbol{1}\left\{X_t^{(k)} = j\right\} - \boldsymbol{M}^k(i,j)\right)\right]_{(i,j)\in[d]^2}$. Notice that $\sum_{t=1}^{\lfloor m/k \rfloor}\boldsymbol{Y}_t^{(k)} = \operatorname{diag}\left(N_1^{(k)}, \dots, N_d^{(k)}\right)\left(\widehat{\boldsymbol{M}}^{(k)} - \boldsymbol{M}^k\right)$. From the Markov property, for any $t \geq 1$, $\mathbf{E}_{t-1}\left[\boldsymbol{Y}_t^{(k)}\right] = \boldsymbol{0}$, proving that $\boldsymbol{Y}_t^{(k)}$ is a matrix martingale difference sequence, and we can write

$$\mathcal{E}_{\boldsymbol{M}}^{(k)} = \operatorname{diag}\left(\sqrt{N_1^{(k)}}, \dots, \sqrt{N_d^{(k)}}\right)^{-1}\left(\sum_{t=1}^{\lfloor m/k \rfloor}\boldsymbol{Y}_t^{(k)}\right)\operatorname{diag}\left(\sqrt{N_1^{(k)}}, \dots, \sqrt{N_d^{(k)}}\right)^{-1}. \tag{B.26}$$

At this point, we will invoke Freedman's inequality for matrices [Tropp, 2011, Corollary 1.3], stated here as Theorem D.1. A direct computation shows that for $t \geq 2$,

$$\left\|\boldsymbol{Y}_t^{(k)}\right\|_\infty = 2\left(1 - \boldsymbol{M}^k\left(X_{t-1}^{(k)}, X_t^{(k)}\right)\right) \text{ and } \left\|\boldsymbol{Y}_t^{(k)}\right\|_1 = \max_{j\in[d]}\left|\boldsymbol{1}\left\{X_t^{(k)} = j\right\} - \boldsymbol{M}^k\left(X_{t-1}^{(k)}, j\right)\right| \tag{B.27}$$

so that from Hölder's inequality for matrix norms, $\left\|\boldsymbol{Y}_t^{(k)}\right\|_2 \leq \sqrt{\left\|\boldsymbol{Y}_t^{(k)}\right\|_1 \left\|\boldsymbol{Y}_t^{(k)}\right\|_\infty} \leq \sqrt{2}$. The same trivially holds for $t = 1$. Similarly, using the Kronecker symbol δ_{ij}, we compute

$$\begin{aligned} \left[\boldsymbol{Y}_t^{(k)} \left(\boldsymbol{Y}_t^{(k)}\right)^\intercal\right]_{(i,j)} &= \delta_{ij} \mathbf{1}\left\{X_{t-1}^{(k)} = i\right\} \left(1 - 2\sum_{\ell=1}^{d} \boldsymbol{M}^k(i,\ell)\mathbf{1}\left\{X_t^{(k)} = \ell\right\} + \left\|\boldsymbol{M}^k(i,\cdot)\right\|_2^2\right) \\ \left[\left(\boldsymbol{Y}_t^{(k)}\right)^\intercal \boldsymbol{Y}_t^{(k)}\right]_{(i,j)} &= \sum_{\ell=1}^{d} \mathbf{1}\left\{X_{t-1}^{(k)} = \ell\right\} \left(\mathbf{1}\left\{X_t^{(k)} = i\right\}\mathbf{1}\left\{X_t^{(k)} = j\right\}\right. \\ &\quad \left. - \boldsymbol{M}^k(\ell,i)\mathbf{1}\left\{X_t^{(k)} = j\right\} - \boldsymbol{M}^k(\ell,j)\mathbf{1}\left\{X_t^{(k)} = i\right\} + \boldsymbol{M}^k(\ell,i)\boldsymbol{M}^k(\ell,j)\right). \end{aligned} \tag{B.28}$$

Recall the random variables $N_{\min}^{(k)}$ and $N_{\max}^{(k)}$ defined in (4.4). As a consequence of (B.28), the two predictable quadratic variation processes $\boldsymbol{W}_{\mathrm{col},\lfloor m/k\rfloor}^{(k)}$ and $\boldsymbol{W}_{\mathrm{row},\lfloor m/k\rfloor}^{(k)}$ are well-defined:

$$\begin{aligned} \left[\boldsymbol{W}_{\mathrm{col},\lfloor m/k\rfloor}^{(k)}\right]_{(i,j)} &\triangleq \sum_{t=1}^{\lfloor m/k\rfloor} \mathbf{E}_{t-1}\left[\left[\boldsymbol{Y}_t^{(k)}\left(\boldsymbol{Y}_t^{(k)}\right)^\intercal\right]_{(i,j)}\right] = \delta_{ij} N_i^{(k)}\left(1 - \left\|\boldsymbol{M}^k(i,\cdot)\right\|_2^2\right) \\ \left[\boldsymbol{W}_{\mathrm{row},\lfloor m/k\rfloor}^{(k)}(i,j)\right]_{(i,j)} &\triangleq \sum_{t=1}^{\lfloor m/k\rfloor} \mathbf{E}_{t-1}\left[\left[\left(\boldsymbol{Y}_t^{(k)}\right)^\intercal \boldsymbol{Y}_t^{(k)}\right]_{(i,j)}\right] = \sum_{\ell=1}^{d} N_\ell^{(k)} \boldsymbol{M}^k(\ell,i)(\delta_{ij} - \boldsymbol{M}^k(\ell,j)) \end{aligned} \tag{B.29}$$

whence

$$\left\|\boldsymbol{W}_{\mathrm{col},\lfloor m/k\rfloor}^{(k)}\right\|_2 \leq N_{\max}^{(k)}, \qquad \left\|\boldsymbol{W}_{\mathrm{row},\lfloor m/k\rfloor}^{(k)}\right\|_2 \leq \frac{1}{4} N_{\max}^{(k)} \left\|\boldsymbol{M}^k\right\|_1, \tag{B.30}$$

where we used the fact that $\mathbf{A} \succcurlyeq \mathbf{B} \succcurlyeq \mathbf{0} \implies \|\mathbf{A}\|_2 \geq \|\mathbf{B}\|_2$ holds for all real valued matrices. Finally, as in the event where $\forall k \in [K], N_{\min}^{(k)} > 0$, $\mathcal{E}_{\boldsymbol{M}}^{(k)}$ is defined, and by sub-multiplicativity of the spectral norm and the fact that for $\mathbf{D} = \operatorname{diag}(v_1, \ldots, v_d)$ a diagonal matrix, $\|\mathbf{D}\|_2 = \max_{i\in[d]} v_i$,

$$\left\|\mathcal{E}_{\boldsymbol{M}}^{(k)}\right\|_2 \leq \frac{1}{N_{\min}^{(k)}} \left\|\sum_{t=1}^{\lfloor m/k\rfloor} \boldsymbol{Y}_t^{(k)}\right\|_2. \tag{B.31}$$

Since $\mathcal{K}_{1/2}^K \implies (N_{\min}^{(k)} \geq \frac{1}{2}\lfloor m/k\rfloor \pi_\star) \wedge (N_{\max}^{(k)} \leq \frac{3}{2}\lfloor m/k\rfloor \max_{i\in[d]}\{\pi_i\})$, we have

$$\begin{aligned} &\mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\mathcal{E}_{\boldsymbol{M}}^{(k)}\right\|_2 > \frac{k\varepsilon}{8} \text{ and } \mathcal{K}_{1/2}^K\right) \leq \mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\sum_{t=1}^{\lfloor m/k\rfloor} \boldsymbol{Y}_t^{(k)}\right\|_2 > \frac{k\varepsilon}{8}\frac{\lfloor m/k\rfloor \pi_\star}{2} \text{ and } N_{\max}^{(k)} \leq \frac{3}{2}\lfloor m/k\rfloor \max_{i\in[d]} \pi_i\right) \\ &\leq \mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\sum_{t=1}^{\lfloor m/k\rfloor} \boldsymbol{Y}_t\right\|_2 > \frac{(m-k)\varepsilon\pi_\star}{16} \text{ and } \max\left\{\left\|\boldsymbol{W}_{\mathrm{col},\lfloor m/k\rfloor}^{(k)}\right\|_2, \left\|\boldsymbol{W}_{\mathrm{row},\lfloor m/k\rfloor}^{(k)}\right\|_2\right\} \leq \frac{3m}{2k}\left\|\boldsymbol{M}^k\right\|_1 \max_{i\in[d]} \pi_i\right) \\ &\leq 2d \exp\left(-C\frac{m\varepsilon^2\pi_\star k}{\|\boldsymbol{M}\|_{\boldsymbol{\pi}} \left\|\boldsymbol{M}^k\right\|_1}\right) \text{ (Theorem D.1)}, \end{aligned} \tag{B.32}$$

and for $m \geq C\frac{\|\boldsymbol{M}^k\|_1 \|\boldsymbol{M}\|_{\boldsymbol{\pi}}}{k\varepsilon^2\pi_\star} \ln\left(\frac{2d}{\delta_{\boldsymbol{M}}^{(k)}}\right)$, the error probability is controlled by $\delta_{\boldsymbol{M}}^{(k)}$.

Finishing up. Since we have

$$\mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(2\left\|\mathcal{E}_{\boldsymbol{\pi}}^{(k)}\right\|_2+\left\|\mathcal{E}_{\boldsymbol{\pi}}^{(k)}\right\|_2^2>\frac{k\varepsilon}{8}\right)\leq \mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\mathcal{E}_{\boldsymbol{\pi}}^{(k)}\right\|_2>\frac{k\varepsilon}{32}\right)+\mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\mathcal{E}_{\boldsymbol{\pi}}^{(k)}\right\|_2>\frac{\sqrt{k\varepsilon}}{4}\right), \tag{B.33}$$

Theorem 5.1 together with Hsu et al. [2017, Section 6.3] imply that for $m \geq \frac{C}{\gamma_{\mathsf{ps}}\pi_\star\varepsilon^2}\ln\left(\frac{d}{\delta_{\boldsymbol{\pi}}^{(k)}}\sqrt{\pi_\star^{-1}}\right)$, this quantity is upper bounded by $\delta_{\boldsymbol{\pi}}^{(k)}$. Similarly, for $K \geq \frac{2}{\varepsilon}$, taking $m \geq \frac{C_{\mathcal{K}}}{\pi_\star\gamma_{\mathsf{ps}}\varepsilon^2}\ln\left(\frac{2\sqrt{2\pi_\star^{-1}}d}{\varepsilon\delta_{\mathcal{K}}}\right)$ and invoking Lemma C.2 gives $\mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\overline{\mathcal{K}}\right)\leq\delta_{\mathcal{K}}$.

Finally, choosing $\delta_{\mathcal{K}}:=\frac{\delta}{4}$, $\delta_{\boldsymbol{M}}^{(k)}:=\delta_{\boldsymbol{\pi}}^{(k)}:=\frac{\varepsilon\delta}{8}$, and taking the maximum of the three sample sizes,

$$m\geq\frac{C_{\mathsf{ps}}}{\pi_\star\varepsilon^2}\max\left\{\frac{1}{\gamma_{\mathsf{ps}}},\|\boldsymbol{M}\|_{\boldsymbol{\pi}}\max_{k\in\lceil 2/\varepsilon\rceil}\left\{\frac{\left\|\boldsymbol{M}^k\right\|_1}{k}\right\}\right\}\ln\left(\frac{d\sqrt{\pi_\star^{-1}}}{\varepsilon^2\delta}\right), C_{\mathsf{ps}}\in\mathbb{R}^+ \tag{B.34}$$

is sufficient to control the error up to absolute error ε and confidence $1-\delta$, which translates into an extra $\frac{1}{\gamma_{\mathsf{ps}}^2}$ factor for the relative error case.

Remark B.1. *Notice that for all $k\in\mathbb{N}$,*

$$\left\|\boldsymbol{M}^k\right\|_1=\max_{i\in[d]}\left\{\sum_{\ell=1}^d\boldsymbol{M}^k(\ell,i)\right\}\leq\frac{1}{\pi_\star}\max_{i\in[d]}\left\{\sum_{\ell=1}^d\pi_\ell\boldsymbol{M}^k(\ell,i)\right\}=\|\boldsymbol{M}\|_{\boldsymbol{\pi}},$$

and $\left\|\boldsymbol{M}^k\right\|_1\leq d$, so that $\frac{\left\|\boldsymbol{M}^k\right\|_1}{k}\leq\min\{d,\|\boldsymbol{M}\|_{\boldsymbol{\pi}}\}$; this justifies the appearance of $\mathcal{C}(\boldsymbol{M})=\|\boldsymbol{M}\|_{\boldsymbol{\pi}}\min\{d,\|\boldsymbol{M}\|_{\boldsymbol{\pi}}\}$ in the upper bound.

Remark B.2. *The dependence on $\gamma_\star$ was improved from $\gamma_\star^3$ in Hsu et al. [2015] to $\gamma_\star$ in Levin and Peres [2016], Hsu et al. [2017] via a "doubling trick" [Levin et al., 2009], which exploited the identity $\gamma_\star(\boldsymbol{M}^k)=1-(1-\gamma_\star(\boldsymbol{M}))^k$ for reversible Markov chains. The non-reversible analogue $\gamma((\boldsymbol{M}^\dagger)^k\boldsymbol{M}^k)=\gamma((\boldsymbol{M}^\dagger\boldsymbol{M})^k)$ in general fails, and we leave the question of the minimax dependence on γ_{ps} as an open problem.*

B.2.2 Proof of Theorem 5.4

Kazakos [1978] proposed an efficient method for recursively computing the Hellinger distance $H^2(\boldsymbol{X},\boldsymbol{Y})$ between two trajectories $\boldsymbol{X}=(X_1,\dots,X_m),\boldsymbol{Y}=(Y_1,\dots,Y_m)$ of length m sampled respectively from two Markov chains $(\boldsymbol{M}_0,\boldsymbol{\mu}_0)$ and $(\boldsymbol{M}_1,\boldsymbol{\mu}_1)$ in terms of the entry-wise geometric mean of their transition matrices and initial distributions, which we reproduce in Lemma D.1. Daskalakis et al. [2017, Proof of Claim 2] used this result to upper bound the Hellinger distance in terms of the spectral radius ρ for symmetric stationary Markov chains:

$$1-H^2(\boldsymbol{X},\boldsymbol{Y})\geq\frac{\rho^m}{d}. \tag{B.35}$$

For $0<\alpha<\frac{1}{8}$, consider the following family of symmetric stochastic matrices of size $d\geq 4$:

$$\boldsymbol{M}(\alpha)=\begin{pmatrix}1-\alpha & \frac{\alpha}{d-1} & \cdots & \cdots & \frac{\alpha}{d-1}\\ \frac{\alpha}{d-1} & 1/2-\frac{\alpha}{d-1} & \frac{1}{2(d-2)} & \cdots & \frac{1}{2(d-2)}\\ \vdots & \frac{1}{2(d-2)} & \ddots & & \frac{1}{2(d-2)}\\ \vdots & \vdots & & & \vdots\\ \frac{\alpha}{d-1} & \frac{1}{2(d-2)} & \frac{1}{2(d-2)} & \cdots & 1/2-\frac{\alpha}{d-1}\end{pmatrix}. \tag{B.36}$$

Being doubly-stochastic, all the chains described by this family are symmetric, reversible, and have $\boldsymbol{\pi} = (1/d, \dots, 1/d)$ as their stationary distribution. The eigenvalues of $\boldsymbol{M}(\alpha)$ are given by $\lambda_1 = 1, \lambda_{\alpha,1} = 1 - \frac{d}{d-1}\alpha, \lambda_{\alpha,2} = 1 - \left(\frac{d-1}{2(d-2)} + \frac{\alpha}{d-1}\right)$. Note that $\lambda_{\alpha,1} > \lambda_{\alpha,2}$ whenever $\alpha < \frac{d-1}{2(d-2)}$, and so $\alpha < \frac{1}{4} \implies \gamma(\boldsymbol{M}(\alpha)) = \frac{d}{d-1}\alpha$. Notice that although the constructed chains are not lazy, we still have $\lambda_{\alpha,1} > 0$ and $\lambda_{\alpha,2} > 0$, so that $\gamma_\star(\boldsymbol{M}(\alpha)) = \gamma(\boldsymbol{M}(\alpha))$ and is (by Lemma C.1) within a factor of 2 of $\gamma_{\mathsf{ps}}(\boldsymbol{M}(\alpha))$. We we proceed to compute $[\boldsymbol{M}(\alpha_0), \boldsymbol{M}(\alpha_1)]_{\surd} =$

$$\begin{pmatrix} \sqrt{(1-\alpha_0)(1-\alpha_1)} & \frac{\sqrt{\alpha_0\alpha_1}}{d-1} & \cdots & \cdots & \frac{\sqrt{\alpha_0\alpha_1}}{d-1} \\ \frac{\sqrt{\alpha_0\alpha_1}}{d-1} & \sqrt{\left(1/2 - \frac{\alpha_0}{d-1}\right)\left(1/2 - \frac{\alpha_1}{d-1}\right)} & \frac{1}{2(d-2)} & \cdots & \frac{1}{2(d-2)} \\ \vdots & \frac{1}{2(d-2)} & \ddots & & \frac{1}{2(d-2)} \\ \vdots & \vdots & & & \vdots \\ \frac{\sqrt{\alpha_0\alpha_1}}{d-1} & \frac{1}{2(d-2)} & \frac{1}{2(d-2)} & \cdots & \sqrt{\left(1/2 - \frac{\alpha_0}{d-1}\right)\left(1/2 - \frac{\alpha_1}{d-1}\right)} \end{pmatrix}. \tag{B.37}$$

Let $\boldsymbol{u} = (1, 0, \dots, 0)^\intercal$ and $\boldsymbol{v} = \frac{1}{\sqrt{d-1}}(0, 1, \dots, 1)^\intercal$ and put

$$p = \frac{\sqrt{\alpha_0\alpha_1}}{d-1}, q = \frac{1}{2(d-2)}, r = \sqrt{(1-\alpha_0)(1-\alpha_1)}, s = \sqrt{\left(1/2 - \frac{\alpha_0}{d-1}\right)\left(1/2 - \frac{\alpha_1}{d-1}\right)}.$$

Then

$$\begin{aligned} [\boldsymbol{M}(\alpha_0), \boldsymbol{M}(\alpha_1)]_{\surd} &= \begin{pmatrix} r & p & \cdots & \cdots & p \\ p & s & q & \cdots & q \\ \vdots & q & \ddots & & q \\ \vdots & \vdots & & & \vdots \\ p & q & q & \cdots & s \end{pmatrix} \\ &= (r - s + q)\boldsymbol{u}\boldsymbol{u}^\intercal + p\sqrt{d-1}(\boldsymbol{u}\boldsymbol{v}^\intercal + \boldsymbol{v}\boldsymbol{u}^\intercal) + (d-1)q\boldsymbol{v}\boldsymbol{v}^\intercal + (s-q)\boldsymbol{I}. \end{aligned} \tag{B.38}$$

So, $[\boldsymbol{M}(\alpha_0), \boldsymbol{M}(\alpha_1)]_{\surd} - (s-q)\boldsymbol{I}$ has rank ≤ 2 and its operator restriction to the subspace span $\{\boldsymbol{u}, \boldsymbol{v}\}$ has a matrix representation $\boldsymbol{R} = \begin{pmatrix} r - s + q & p\sqrt{d-1} \\ p\sqrt{d-1} & (d-1)q \end{pmatrix}$. For $d \geq 4$, and since $r - s + q > 0$, $\boldsymbol{R}$ is entry-wise positive. Hence

$$\rho := \rho\left([\boldsymbol{M}(\alpha_0), \boldsymbol{M}(\alpha_1)]_{\surd}\right) = s - q + \rho(\boldsymbol{R}), \tag{B.39}$$

with

$$\rho(\boldsymbol{R}) = \max\left\{\frac{\mathrm{Tr}(\boldsymbol{R})}{2} + \sqrt{\frac{\mathrm{Tr}^2(\boldsymbol{R})}{4} - |\boldsymbol{R}|}, \frac{\mathrm{Tr}(\boldsymbol{R})}{2} - \sqrt{\frac{\mathrm{Tr}^2(\boldsymbol{R})}{4} - |\boldsymbol{R}|}\right\}, \tag{B.40}$$

so that

$$\begin{aligned} \rho &= s - q + \frac{(r - s + dq) + \sqrt{[r - s - (d-2)q]^2 + 4(d-1)p^2}}{2} \\ &= \frac{(r + s + 1/2) + \sqrt{(r - s - 1/2)^2 + \frac{4\alpha_0\alpha_1}{d-1}}}{2}. \end{aligned} \tag{B.41}$$

Lemma B.4 shows that, for $d \geq 4, 0 < \alpha < 1/8, 0 < \varepsilon < 1/2$, and $\alpha_0 = \alpha(1-\varepsilon), \alpha_1 = \alpha(1+\varepsilon)$, we have

$$\rho = \rho\left([\boldsymbol{M}(\alpha_0), \boldsymbol{M}(\alpha_1)]_{\surd}\right) \geq 1 - 6\frac{\alpha\varepsilon^2}{d-1}.$$

The minimax risk for the problem of estimating the pseudo-spectral gap is defined as

$$\mathcal{R}_m^{\mathsf{ps}} \triangleq \inf_{\widehat{\gamma}_{\mathsf{ps}}} \sup_{(\boldsymbol{\mu}, \boldsymbol{M}) \in \mathcal{M}_{d, \gamma_{\mathsf{ps}}, \pi_\star}} \mathbf{P}_{\boldsymbol{M}, \boldsymbol{\mu}}\left(|\widehat{\gamma}_{\mathsf{ps}}(X_1, \dots, X_m) - \gamma_{\mathsf{ps}}(\boldsymbol{M})| > \varepsilon\right), \tag{B.42}$$

where the inf is taken over all measurable functions $\widehat{\gamma}_{\mathsf{ps}} : (X_1, \dots, X_m) \to (0, 1)$, and the sup over the set $\mathcal{M}_{d, \gamma_{\mathsf{ps}}, \pi_\star}$ of d-state Markov chains whose minimum stationary probability is $\pi_\star$, and of pseudo-spectral gap at least γ_{ps}. Using Le Cam's two point method [Tsybakov, 2009, Chapter 2],

$$\begin{aligned} \mathcal{R}_m^{\mathsf{ps}} &\geq \frac{1}{2}\left(1 - \sqrt{H^2\left(\boldsymbol{X} \sim \boldsymbol{M}(\alpha_0), \boldsymbol{Y} \sim \boldsymbol{M}(\alpha_1)\right)}\right) \geq \frac{1}{2}\left(1 - \sqrt{1 - \frac{\rho^m}{d}}\right) \geq \frac{\rho^m}{4d} \\ &\geq \frac{\exp\left(m \ln\left(1 - 6\frac{\varepsilon^2\alpha}{d}\right)\right)}{4d} \geq \frac{\exp\left(-9\frac{m\varepsilon^2\alpha}{d}\right)}{4d}, \end{aligned} \tag{B.43}$$

where the last inequality holds because $6\frac{\varepsilon^2\alpha}{d} \leq 3/8$ and $\ln(1-t) \geq -\frac{3}{2}t,\ t \in \left(0, \frac{1}{2}\right)$. Thus, for $\delta < \frac{1}{4d}$, a sample of size at least $m = \Omega\left(\frac{d}{9 \cdot \varepsilon^2 \alpha} \ln\left(\frac{1}{4d\delta}\right)\right)$ is necessary to achieve a confidence of $1 - \delta$.

Lemma B.1. *For $d \geq 4,\ 0 < \alpha_0, \alpha_1 < 1/4$, and $s = \sqrt{\left(1/2 - \frac{\alpha_0}{d-1}\right)\left(1/2 - \frac{\alpha_1}{d-1}\right)}$, we have*

$$\frac{1}{2}\left(1 - \frac{\alpha_0 + \alpha_1}{d-1} - 2\left(\frac{\alpha_0 - \alpha_1}{d-1}\right)^2\right) \leq s \leq \frac{1}{2}\left(1 - \frac{\alpha_0 + \alpha_1}{d-1}\right). \tag{B.44}$$

Proof. We begin by showing that $s \leq \frac{1}{2}\left(1 - \frac{\alpha_0+\alpha_1}{d-1}\right)$. By the AM-GM inequality, $4\alpha_0\alpha_1 \leq (\alpha_0 + \alpha_1)^2$, whence

$$\begin{aligned} \left(\sqrt{1 - \frac{2\alpha_0}{d-1}}\sqrt{1 - \frac{2\alpha_1}{d-1}}\right)^2 &= 1 - \frac{2(\alpha_0 + \alpha_1)}{d-1} + \frac{4\alpha_0\alpha_1}{(d-1)^2} \\ &\leq 1 - \frac{2(\alpha_0 + \alpha_1)}{d-1} + \left(\frac{\alpha_0 + \alpha_1}{d-1}\right)^2 \\ &= \left(1 - \frac{\alpha_0 + \alpha_1}{d-1}\right)^2, \end{aligned} \tag{B.45}$$

which, together with $\frac{\alpha_0+\alpha_1}{d-1} \leq 1$, proves the upper bound. For the lower bound, observe that

$$\begin{aligned} \left(1 - \frac{\alpha_0 + \alpha_1}{d-1} - 2\left(\frac{\alpha_0 - \alpha_1}{d-1}\right)^2\right)^2 &= \left(\sqrt{1 - \frac{2\alpha_0}{d-1}}\sqrt{1 - \frac{2\alpha_1}{d-1}}\right)^2 \\ &+ \left(\frac{\alpha_0 - \alpha_1}{d-1}\right)^2\left[\frac{4(\alpha_0 + \alpha_1)}{d-1} + \frac{4(\alpha_0 - \alpha_1)^2}{(d-1)^2} - 3\right]. \end{aligned} \tag{B.46}$$

Now for $d \geq 4,\ 0 < \alpha_0, \alpha_1 < 1/4$, we have $\frac{4(\alpha_0+\alpha_1)}{d-1} + 4\left(\frac{\alpha_0-\alpha_1}{d-1}\right)^2 < 3$ and the lemma is proved. □

Lemma B.2. *For $d \geq 4, 0 < \alpha < 1/8, 0 < \varepsilon < 1/2$ and $\alpha_0 = \alpha(1-\varepsilon), \alpha_1 = \alpha(1+\varepsilon), r = \sqrt{(1-\alpha_0)(1-\alpha_1)}$, we have*

$$\sqrt{(r-1)^2 + \frac{4\alpha_0\alpha_1 + (\alpha_0+\alpha_1)(r-1)}{d-1} + \frac{\alpha_0\alpha_1}{(d-1)^2}} \geq (1-r) + \left[\frac{4\alpha_0\alpha_1}{1-r} - \frac{3}{2}(\alpha_0+\alpha_1)\right]\frac{1}{d-1}. \tag{B.47}$$

Proof. Squaring (B.47), the claim will follow immediately from

$$2(\alpha_0+\alpha_1)(1-r) + \frac{\alpha_0\alpha_1}{d-1} \geq \left(\frac{4\alpha_0\alpha_1}{1-r} - \frac{3}{2}(\alpha_0+\alpha_1)\right)^2 \frac{1}{d-1} + 4\alpha_0\alpha_1. \tag{B.48}$$

Fix α_0, α_1 and define

$$\varphi(t) = 2(\alpha_0+\alpha_1)(1-r) - 4\alpha_0\alpha_1 + t\left[\alpha_0\alpha_1 - \left(\frac{4\alpha_0\alpha_1}{1-r} - \frac{3}{2}(\alpha_0+\alpha_1)\right)^2\right], \qquad t \in (0, +\infty).$$

The sign of $\frac{d\varphi(t)}{dt}$ is the same as that of $\alpha_0\alpha_1 - \left(\frac{4\alpha_0\alpha_1}{1-r} - \frac{3}{2}(\alpha_0+\alpha_1)\right)^2$. Since $\varepsilon < \frac{1}{2}$, it is straightforward to verify that $4(1-\varepsilon)(1+\varepsilon) - 3 \leq \sqrt{(1+\varepsilon)(1-\varepsilon)}$. Additionally, $r^2 = 1 - 2\alpha + \alpha^2(1+\varepsilon)(1-\varepsilon) \leq (1-\alpha)^2$, so that $\frac{1}{1-r} \leq \frac{1}{\alpha}$, and $\frac{4(1-\varepsilon)(1+\varepsilon)\alpha}{1-r} - 3 \leq \sqrt{(1+\varepsilon)(1-\varepsilon)}$. Squaring this inequality yields, for our range of parameters, $\frac{d\varphi(t)}{dt} \geq 0$, and so φ is minimized at $t = 0$, and $\varphi(t) \geq \varphi(0) = 2(\alpha_0+\alpha_1)(1-r) - 4\alpha_0\alpha_1$.

By the AM-GM inequality, $\sqrt{\alpha_0\alpha_1} \leq \frac{\alpha_0+\alpha_1}{2} \implies (1-\alpha_0)(1-\alpha_1) \leq \left(1 - \frac{\alpha_0+\alpha_1}{2}\right)^2$, and so $r \leq 1 - \frac{\alpha_0+\alpha_1}{2}$, and another application of AM-GM yields $1 - r \geq 2\frac{\alpha_0\alpha_1}{\alpha_0+\alpha_1}$. We conclude that $\varphi(0) \geq 0$, which proves the claim. □

Lemma B.3. *For $d \geq 4, 0 < \alpha < 1/8, 0 < \varepsilon < 1/2$ and $\alpha_0 = \alpha(1-\varepsilon), \alpha_1 = \alpha(1+\varepsilon), r = \sqrt{(1-\alpha_0)(1-\alpha_1)}$, we have*

$$\alpha_0 + \alpha_1 - 2\frac{\alpha_0\alpha_1}{1-r} \leq 4\varepsilon^2\alpha. \tag{B.49}$$

Proof. Observe that $\alpha(\alpha\varepsilon^2 + 2) \leq 1$ holds for our assumed range of parameters, which implies $2(1-\alpha) - \varepsilon^2\alpha^2 \geq 1$, and further

$$(1 - \alpha - \alpha^2\varepsilon^2)^2 = (1-\alpha)^2 - \alpha^2\varepsilon^2[2(1-\alpha) - \varepsilon^2\alpha^2] \leq (1-\alpha)^2 - \alpha^2\varepsilon^2 = r^2. \tag{B.50}$$

As a consequence, $\frac{1}{1-r} \geq \frac{1}{\alpha(1+\alpha\varepsilon^2)} \geq \frac{1-\alpha\varepsilon^2}{\alpha}$, and

$$\alpha_0 + \alpha_1 - 2\frac{\alpha_0\alpha_1}{1-r} \leq 2\alpha - 2\alpha(1+\varepsilon)(1-\varepsilon)(1-\alpha\varepsilon^2) = 2\alpha\varepsilon^2(1 + \alpha - \alpha\varepsilon^2) \leq 4\alpha\varepsilon^2. \tag{B.51}$$

□

Lemma B.4. *For $d \geq 4, 0 < \alpha < 1/8, 0 < \varepsilon < 1/2$ ρ as defined in* (B.41)*, and $\alpha_0 = \alpha(1-\varepsilon), \alpha_1 = \alpha(1+\varepsilon), r = \sqrt{(1-\alpha_0)(1-\alpha_1)}$, we have*

$$\rho \geq 1 - \left[(\alpha_0+\alpha_1) - \frac{2\alpha_0\alpha_1}{1-r}\right]\frac{1}{d-1} - \frac{1}{2}\left(\frac{\alpha_0-\alpha_1}{d-1}\right)^2 \geq 1 - 6\frac{\alpha\varepsilon^2}{d-1}. \tag{B.52}$$

Proof.

$$
\begin{aligned}
2\rho &= (r+s+1/2) + \left((r-1/2)^2 + s^2 - 2(r-1/2)s + \frac{4\alpha_0\alpha_1}{d-1}\right)^{1/2} \\
&\geq (r+s+1/2) + \left((r-1/2)^2 + \left(1/2 - \frac{\alpha_0}{d-1}\right)\left(1/2 - \frac{\alpha_1}{d-1}\right)\right. \\
&\left. - 2(r-1/2)\left(1 - \frac{\alpha_0+\alpha_1}{d-1}\right) + \frac{4\alpha_0\alpha_1}{d-1}\right)^{1/2} \\
&= (r+s+1/2) + \left((r-1)^2 + \frac{4\alpha_0\alpha_1 + (\alpha_0+\alpha_1)(r-1)}{d-1} + \frac{\alpha_0\alpha_1}{(d-1)^2}\right)^{1/2},
\end{aligned}
\tag{B.53}
$$

where the inequality is due to Lemma B.1. Invoking Lemmas B.1, and B.2, we have

$$
\begin{aligned}
2\rho &\geq r + \frac{1}{2}\left(1 - \frac{\alpha_0+\alpha_1}{d-1} - 2\left(\frac{\alpha_0-\alpha_1}{d-1}\right)^2\right) + 1/2 + (1-r) + \left[\frac{4\alpha_0\alpha_1}{1-r} - \frac{3}{2}(\alpha_0+\alpha_1)\right]\frac{1}{d-1}, \\
\rho &\geq 1 - \left[(\alpha_0+\alpha_1) - \frac{2\alpha_0\alpha_1}{1-r}\right]\frac{1}{d-1} - \frac{1}{2}\left(\frac{\alpha_0-\alpha_1}{d-1}\right)^2.
\end{aligned}
\tag{B.54}
$$

Finally, Lemma B.3 implies a lower bound on ρ:

$$
\rho \geq 1 - 4\frac{\alpha\varepsilon^2}{d-1} - 4\frac{\alpha^2\varepsilon^2}{2(d-1)^2} \geq 1 - 6\frac{\alpha\varepsilon^2}{d-1}. \tag{B.55}
$$

□

B.3 Analysis of the empirical procedure

Before delving into the proofs, we provide some preliminary motivation and analysis behind the main estimation procedure in Algorithm 1.

Although we only have access to a single path from $\boldsymbol{M}$, by considering skipped chains with offsets $r \in \{0, \dots, k-1\}$ we effectively obtain k different paths from which to estimate $\boldsymbol{M}^k$. Define the averaged estimator

$$
\widetilde{\gamma}^{\dagger}_{k,\alpha}(X_1, \dots, X_m) \triangleq \frac{1}{k}\sum_{r=0}^{k-1} \widehat{\gamma}^{\dagger}_{k,r,\alpha}\left(X_t^{(k,r)}, 1 \leq t \leq \lfloor (m-r)/k \rfloor\right), \tag{B.56}
$$

where $\widehat{\gamma}^{\dagger}_{k,r,\alpha}$ is the spectral gap of the multiplicative reversiblization of the α-smoothed empirical transition matrix constructed from $X_t^{(k,r)}$:

$$
\widehat{\boldsymbol{M}}^{(k,r,\alpha)}(i,j) \triangleq \frac{N_{ij}^{(k,r)} + \alpha}{N_i^{(k,r)} + d\alpha}, \qquad \widehat{\pi}_i^{(k,r,\alpha)} \triangleq \frac{N_i^{(k,r)} + d\alpha}{\lfloor (m-r)/k \rfloor + d^2\alpha}. \tag{B.57}
$$

Notice that $\widehat{\pi}_i^{(k,r,\alpha)}$ requires more aggressive smoothing than for the transition matrix in order to ensure stationarity. We now derive the smoothed empirical form of (B.72) for a generic skipping rate k and offset r:

$$
\left[\left(\widehat{\boldsymbol{L}}^{(k,r,\alpha)}\right)^{\intercal}\widehat{\boldsymbol{L}}^{(k,r,\alpha)}\right]_{(i,j)} = \frac{1}{\sqrt{\left(N_i^{(k,r)} + d\alpha\right)\left(N_j^{(k,r)} + d\alpha\right)}}\sum_{\ell=1}^{d}\frac{\left(N_{\ell i}^{(k,r)} + \alpha\right)\left(N_{\ell j}^{(k,r)} + \alpha\right)}{N_\ell^{(k,r)} + d\alpha}, \tag{B.58}
$$

where $N_i^{(k,r)}, N_{ij}^{(k,r)}$ are defined in (4.2, 4.3). The expression can be alternatively be written in its vectorized form:

$$\left(\widehat{\boldsymbol{L}}^{(k,r,\alpha)}\right)^{\mathsf{T}} \widehat{\boldsymbol{L}}^{(k,r,\alpha)} = \left(\mathbf{D}_N^{(k,r,\alpha)}\right)^{-1/2} \left(\mathbf{N}^{(k,r,\alpha)}\right)^{\mathsf{T}} \left(\mathbf{D}_N^{(k,r,\alpha)}\right)^{-1} \mathbf{N}^{(k,r,\alpha)} \left(\mathbf{D}_N^{(k,r,\alpha)}\right)^{-1/2}$$
$$\text{with } \mathbf{N}^{(k,r,\alpha)} \triangleq \left[N_{ij}^{(k,r)} + \alpha\right]_{(i,j)}, \ \mathbf{D}_N^{(k,r,\alpha)} \triangleq \operatorname{diag}\left(N_1^{(k,r)} + d\alpha, \dots, N_d^{(k,r)} + d\alpha\right). \tag{B.59}$$

B.3.1 Proof of Theorems 6.1 and A.1

We begin with a decomposition very similar to the one employed in Section B.2.1. From the definition of $\widetilde{\gamma}^{\dagger}_{k,\alpha}$ in (B.56), it follows that for all $K \in \mathbb{N}$, we have

$$\left|\widehat{\gamma}^{(\alpha)}_{\mathsf{ps}[K]} - \gamma_{\mathsf{ps}}\right| \leq \frac{1}{K} + 2 \max_{k \in [K]} \left\{ \frac{1}{k^2} \sum_{r=0}^{k-1} \left(\left\| \mathcal{E}_{\boldsymbol{M}}^{(k,r,\alpha)} \right\|_2 + 2 \left\| \mathcal{E}_{\boldsymbol{\pi}}^{(k,r,\alpha)} \right\|_2 + \left\| \mathcal{E}_{\boldsymbol{\pi}}^{(k,r,\alpha)} \right\|_2^2 \right) \right\},$$
$$\text{where } \mathcal{E}_{\boldsymbol{M}}^{(k,r,\alpha)} \triangleq \left(\widehat{\boldsymbol{D}}_{\boldsymbol{\pi}}^{(k,r,\alpha)}\right)^{1/2} \left(\widehat{\boldsymbol{M}}^{(k,r,\alpha)} - \boldsymbol{M}^k\right) \left(\widehat{\boldsymbol{D}}_{\boldsymbol{\pi}}^{(k,r,\alpha)}\right)^{-1/2},$$
$$\mathcal{E}_{\boldsymbol{\pi},1}^{(k,r,\alpha)} = \left(\widehat{\boldsymbol{D}}_{\boldsymbol{\pi}}^{(k,r,\alpha)}\right)^{1/2} \boldsymbol{D}_{\boldsymbol{\pi}}^{-1/2} - \boldsymbol{I}, \qquad \mathcal{E}_{\boldsymbol{\pi},2}^{(k,r,\alpha)} = \boldsymbol{D}_{\boldsymbol{\pi}}^{1/2} \left(\widehat{\boldsymbol{D}}_{\boldsymbol{\pi}}^{(k,r,\alpha)}\right)^{-1/2} - \boldsymbol{I}, \tag{B.60}$$

and $\widehat{\boldsymbol{M}}^{(k,r,\alpha)}, \widehat{\boldsymbol{D}}_{\boldsymbol{\pi}}^{(k,r,\alpha)}$ are the α-smoothed estimators for $\boldsymbol{M}^k$ and $\boldsymbol{D}_{\boldsymbol{\pi}}$ constructed from the sample path $X_t^{(k,r)}$.

Cho and Meyer [2001, Section 3.3] prove the perturbation bound

$$\|\widehat{\boldsymbol{\pi}} - \boldsymbol{\pi}\|_\infty \leq \hat{\kappa} \left\|\widehat{\boldsymbol{M}} - \boldsymbol{M}\right\|_\infty,$$

where $\hat{\kappa} \triangleq \frac{1}{2} \max_{j \in [d]} \left\{ \widehat{\boldsymbol{A}}^{\#}_{j,j} - \min_{i \in [d]} \left\{ \widehat{\boldsymbol{A}}^{\#}_{i,j} \right\} \right\}$ and $\widehat{\boldsymbol{A}}^{\#}$ is the empirical (Drazin) *group inverse* of $\widehat{\boldsymbol{M}}$ [Meyer, 1975]. Since computing $\hat{\kappa}$ via matrix inversion is computationally expensive, we show in Lemma C.3 that

$$\left\|\widehat{\boldsymbol{\pi}}^{(k,r,\alpha)} - \boldsymbol{\pi}\right\|_\infty \leq \frac{C_{\mathcal{K}}}{\gamma_{\mathsf{ps}}\left(\widehat{\boldsymbol{M}}^{(k,r,\alpha)}\right)} \ln\left(2\sqrt{\frac{2(\lfloor (m-r)/k \rfloor + d^2\alpha)}{N_{\min}^{(k,r)} + d\alpha}}\right) \left\|\widehat{\boldsymbol{M}}^{(k,r,\alpha)} - \boldsymbol{M}^k\right\|_\infty, \tag{B.61}$$

where $C_{\mathcal{K}}$ is the universal constant from Lemma C.2. From the triangle inequality,

$$\left\|\widehat{\boldsymbol{M}}^{(k,r,\alpha)}(i,\cdot) - \boldsymbol{M}^k(i,\cdot)\right\|_1 \leq \frac{N_i^{(k,r)}}{N_i^{(k,r)} + d\alpha} \left\|\widehat{\boldsymbol{M}}^{(k,r)}(i,\cdot) - \boldsymbol{M}^k(i,\cdot)\right\|_1 + \frac{\alpha d \left\|(1/d)\cdot \mathbf{1} - \boldsymbol{M}^k(i,\cdot)\right\|_1}{N_i^{(k,r)} + d\alpha}, \tag{B.62}$$

and writing

$$\hat{d}^{(k,r,\alpha)} \triangleq 4\tau_{\delta/d, \lfloor (m-r)/k \rfloor} \sqrt{\frac{d}{N_{\min}^{(k,r)} + d\alpha}} + \frac{2\alpha d}{N_{\min}^{(k,r)} + d\alpha},$$
$$\hat{b}^{(k,r,\alpha)} \triangleq \frac{C_{\mathcal{K}}}{\gamma_{\mathsf{ps}}\left(\widehat{\boldsymbol{M}}^{(k,r,\alpha)}\right)} \ln\left(2\sqrt{\frac{2(\lfloor (m-r)/k \rfloor + d^2\alpha)}{N_{\min}^{(k,r)} + d\alpha}}\right) \hat{d}^{(k,r,\alpha)}, \tag{B.63}$$

it follows via Lemma C.4 that $\left\| \widehat{\boldsymbol{\pi}}^{(k,r,\alpha)} - \boldsymbol{\pi} \right\|_\infty \leq \hat{b}^{(k,r,\alpha)}$.

Define $\hat{c}^{(k,r,\alpha)} \triangleq \frac{1}{2} \max \bigcup_{i \in [d]} \left\{ \frac{\hat{b}^{(k,r,\alpha)}}{\widehat{\boldsymbol{\pi}}^{(k,r,\alpha)}(i)}, \frac{\hat{b}^{(k,r,\alpha)}}{\left[\widehat{\boldsymbol{\pi}}^{(k,r,\alpha)}(i) - \hat{b}^{(k,r,\alpha)}\right]_+} \right\}$, and recall from Hsu et al. [2017], that $\left\| \widehat{\boldsymbol{\pi}}^{(k,r,\alpha)} - \boldsymbol{\pi} \right\|_\infty \leq \hat{b}^{(k,r,\alpha)} \implies \left\| \mathcal{E}_{\boldsymbol{\pi}}^{(k,r,\alpha)} \right\|_2 \leq \hat{c}^{(k,r,\alpha)}$. By sub-multiplicativity of the spectral norm and norm properties of diagonal matrices, we have

$$\left\| \mathcal{E}_{\boldsymbol{M}}^{(k,r,\alpha)} \right\|_2 \leq \frac{N_{\max}^{(k,r)} + d\alpha}{\lfloor (m-r)/k \rfloor + d^2\alpha} \sqrt{d} \left\| \widehat{\boldsymbol{M}}^{(k,r,\alpha)} - \boldsymbol{M}^k \right\|_\infty \frac{\lfloor (m-r)/k \rfloor + d^2\alpha}{N_{\min}^{(k,r)} + d\alpha}. \tag{B.64}$$

Putting

$$\hat{a}^{(k,r,\alpha)} \triangleq \sqrt{d} \frac{N_{\max}^{(k,r)} + d\alpha}{N_{\min}^{(k,r)} + d\alpha} \hat{d}^{(k,r,\alpha)}, \tag{B.65}$$

we have that $\left\| \mathcal{E}_{\boldsymbol{M}}^{(k,r,\alpha)} \right\|_2 \leq \hat{a}^{(k,r,\alpha)}$ holds with probability at least $1 - \delta$.

Turning to the simpler reversible case, consider the α-smoothed version of the estimator $\frac{1}{2}\left(\widehat{\boldsymbol{M}}^\dagger + \widehat{\boldsymbol{M}}\right)$, where the (i,j)th entry is $\frac{N_{ij} + N_{ji}}{2N_i}$. Reversibility allows us to apply Weyl's inequality:

$$\left| \lambda_i \left(\frac{1}{2} \left(\widehat{\boldsymbol{M}}^\dagger + \widehat{\boldsymbol{M}} \right) \right) - \lambda_i (\boldsymbol{M}) \right| = \left| \lambda_i \left(\frac{1}{2} \left(\widehat{\boldsymbol{L}}^\intercal + \widehat{\boldsymbol{L}} \right) \right) - \lambda_i (\boldsymbol{L}) \right| \leq \left\| \widehat{\boldsymbol{L}} - \boldsymbol{L} \right\|_2, \qquad i \in [d]. \tag{B.66}$$

The interval widths can now be deduced from Corollary C.2 and Corollary C.3.

B.3.2 Asymptotic interval widths

Non-reversible setting. The definition of the pseudo-spectral gap implies that $\gamma_{\mathsf{ps}}(\boldsymbol{M}^k) \geq k\gamma_{\mathsf{ps}}(\boldsymbol{M})$, and assuming a smoothing parameter $\alpha < 1/d$, a straightforward computation (ignoring logarithmic factors) yields

$$\sqrt{m}\hat{a}^{(k,r,\alpha)} = \tilde{\mathcal{O}}\left(\frac{\sqrt{k} \|\boldsymbol{M}\|_{\boldsymbol{\pi}} d}{\sqrt{\pi_\star}} \right), \sqrt{m}\hat{b}^{(k,r,\alpha)} \leq \tilde{\mathcal{O}}\left(\frac{\sqrt{d}}{\sqrt{k}\gamma_{\mathsf{ps}}\sqrt{\pi_\star}} \right), \sqrt{m}\hat{c}^{(k,r,\alpha)} \leq \tilde{\mathcal{O}}\left(\frac{\sqrt{d}}{\sqrt{k}\gamma_{\mathsf{ps}}\pi_\star^{3/2}} \right). \tag{B.67}$$

It follows that

$$\sqrt{m}\left| \hat{\pi}_\star^{(\alpha)} - \pi_\star \right| = \tilde{\mathcal{O}}\left(\frac{\sqrt{d}}{\gamma_{\mathsf{ps}}\sqrt{\pi_\star}} \right), \quad \sqrt{m}\left| \widehat{\gamma}_{\mathsf{ps}[K]}^{(\alpha)} - \gamma_{\mathsf{ps}} \right| = \tilde{\mathcal{O}}\left(\frac{1}{K} + \sqrt{\frac{d}{\pi_\star}} \left(\sqrt{d} \|\boldsymbol{M}\|_{\boldsymbol{\pi}} + \frac{1}{\gamma_{\mathsf{ps}}\pi_\star} \right) \right). \tag{B.68}$$

Reversible setting. Here,

$$\sqrt{m}\hat{a}^{(\alpha)} = \tilde{\mathcal{O}}\left(\frac{\|\boldsymbol{M}\|_{\boldsymbol{\pi}} d}{\sqrt{\pi_\star}} \right), \qquad \sqrt{m}\hat{b}^{(\alpha)} = \tilde{\mathcal{O}}\left(\frac{\sqrt{d}}{\gamma_\star\sqrt{\pi_\star}} \right), \qquad \sqrt{m}\hat{c}^{(\alpha)} = \tilde{\mathcal{O}}\left(\frac{\sqrt{d}}{\gamma_\star\pi_\star^{3/2}} \right), \tag{B.69}$$

so that

$$\sqrt{m}\left| \hat{\pi}_\star^{(\alpha)} - \pi_\star \right| = \tilde{\mathcal{O}}\left(\frac{\sqrt{d}}{\gamma_\star\sqrt{\pi_\star}} \right), \qquad \sqrt{m}\left| \widehat{\gamma}_\star^{(\alpha)} - \gamma_\star \right| = \tilde{\mathcal{O}}\left(\sqrt{\frac{d}{\pi_\star}} \left(\sqrt{d} \|\boldsymbol{M}\|_{\boldsymbol{\pi}} + \frac{1}{\gamma_\star\pi_\star} \right) \right). \tag{B.70}$$

Remark B.3. *Let us compare the intervals obtained in Hsu et al. [2017, Theorem 4.1] for the reversible case with ours. For estimating $\pi_\star$, we obtain an improvement of a factor of $\tilde{\mathcal{O}}\left(\sqrt{d}\right)$, and for γ_{ps}, we witness a similar improvement when $t_{\mathsf{mix}} = \Omega\left(\sqrt{d}\max_{i\in[d]}\boldsymbol{\pi}(i)\right)$ — i.e. in the case where the chain is not rapidly mixing. When comparing to Hsu et al. [2017, Theorem 4.2], which intersects the point-estimate intervals with the empirical ones, our results are asymptotically equivalent. However, this is a somewhat misleading comparison, since the aforementioned intersection, being asymptotic in nature, is oblivious to the rate decay of the empirical intervals.*

B.3.3 Computational complexity in the reversible setting

The best current time complexity of multiplying (and inverting, and diagonalizing) $d\times d$ matrices is $\mathcal{O}(d^\omega)$ where $2 \leq \omega \leq 2.3728639$ [Le Gall, 2014].

Time complexity. In the reversible case, the time complexity of our algorithm is $\mathcal{O}\left(m + d^2 + \mathcal{C}_{\lambda_\star}\right)$, where $\mathcal{C}_{\lambda_\star}$ is the complexity of computing the second largest-magnitude eigenvalue of a symmetric matrix. For this task, we consider here the Lanczos algorithm. Let $\lambda_1, \dots, \lambda_d$ be the eigenvalues of a symmetric real matrix ordered by magnitude, and denote by $\lambda_1^{\text{LANCZOS}}$ the algorithm's approximation for λ_1. Then, for a stochastic matrix, it is known [Kaniel, 1966, Paige, 1971, Saad, 1980] that

$$|\lambda_1 - \lambda_1^{\text{LANCZOS}}| \leq C_{\text{LANCZOS}} R^{-2(n-1)}, \tag{B.71}$$

where C_{LANCZOS} a universal constant, n is the number of iterations (in practice often $n \ll d$), $R = 1 + 2r + 2\sqrt{r^2 + r}$, and $r = \frac{\lambda_1 - \lambda_2}{\lambda_2 - \lambda_d}$. So in order to attain additive accuracy η, it suffices iterate the method $n \geq 1 + \frac{1}{2}\frac{\ln\left(C\eta^{-1}\right)}{\ln(R)} = \mathcal{O}\left(\frac{\ln\left(\eta^{-1}\right)}{\ln(R)}\right)$ times.

A single iteration involves multiplying a vector by a matrix, incurring a cost of $\mathcal{O}\left(d^2\right)$, and so the full complexity of the Lanczos algorithm is $\mathcal{O}\left(m + d^2\frac{\ln\left(\eta^{-1}\right)}{\ln(R)}\right)$. More refined complexity analyses may be found in Kuczyński and Woźniakowski [1992], Arora et al. [2005]

The previous approach of Hsu et al. [2017] involved computing the Drazin inverse via matrix inversion, incurring a cost of $\mathcal{O}\left(m + d^\omega\right)$. Thus, our proposed computational method is faster over a non-trivial regime.

Space complexity. When estimating the absolute spectral gap of a reversible case, the space complexity remains $\mathcal{O}\left(d^2\right)$, as the full empirical transition matrix is being constructed.

B.3.4 Computational complexity in the non-reversible setting

Time complexity. For the non-reversible setting, we are estimating multiplicative reversiblizations of powers of chains which involves matrix multiplication. Our time complexity is $\mathcal{O}\left(K^2\left(m + d^\omega\right)\right)$.

Space complexity. In the non-reversible case, $\mathcal{O}\left(K^2\right)$ matrices are being constructed, so that the overall space complexity is $\mathcal{O}\left(K^2 d^2\right)$.

Reducing computation of a pseudo-spectral gap to the computation of spectral radii. We now show that, in order to compute $\gamma_{\mathsf{ps}[K]}$ it suffices to compute the value of K

spectral radii. For a Markov chain $\boldsymbol{M}$ with stationary distribution $\boldsymbol{\pi}$ and corresponding $\boldsymbol{L}$, we observe that $\gamma((\boldsymbol{M}^k)^\dagger \boldsymbol{M}^k) = \gamma((\boldsymbol{L}^k)^\intercal \boldsymbol{L}^k)$ and compute,

$$\left[\left(\boldsymbol{L}^k\right)^\intercal \boldsymbol{L}^k\right]_{(i,j)} = \sum_{\ell=1}^{d} \sqrt{\frac{\pi_\ell}{\pi_i}} \boldsymbol{M}(\ell,i)^k \sqrt{\frac{\pi_\ell}{\pi_j}} \boldsymbol{M}(\ell,j)^k, \qquad k \in \mathbb{N}. \tag{B.72}$$

Since $(\boldsymbol{L}^k)^\intercal \boldsymbol{L}^k$ is a symmetric positive semi-definite matrix, its eigenvalues may be ordered such that $\nu_1 \geq \nu_2 \geq \dots \nu_d \geq 0$. From (B.72) it follows that $\sqrt{\boldsymbol{\pi}}$ is a left eigenvector for eigenvalue 1, and by the Perron-Frobenius theorem, $\rho((\boldsymbol{L}^k)^\intercal \boldsymbol{L}^k) = 1$. By symmetry, we can express $(\boldsymbol{L}^k)^\intercal \boldsymbol{L}^k$ over an orthogonal left row eigen-basis $v_1, \dots, v_d$ with the associated eigenvalues $\nu_1, \dots, \nu_d$, where $(\nu_1, v_1) = (1, \sqrt{\boldsymbol{\pi}})$:

$$\left(\boldsymbol{L}^k\right)^\intercal \boldsymbol{L}^k = \sum_{i=1}^{d} \nu_i (v_i^\intercal v_i) \text{ hence } \rho\left(\left(\boldsymbol{L}^k\right)^\intercal \boldsymbol{L}^k - \sqrt{\boldsymbol{\pi}}^\intercal \sqrt{\boldsymbol{\pi}}\right) = \nu_2. \tag{B.73}$$

We have thus shown that at the cost of computing $\sqrt{\boldsymbol{\pi}}^\intercal \sqrt{\boldsymbol{\pi}}$ does not depend on $k \in [K]$, reducing the problem to computing the largest eigenvalue of a real symmetric matrix. The latter may be achieved via the Lanczos method, and is more efficient than the power iteration methods in the symmetric case. In summary, we have shown that

$$\gamma_{\mathsf{ps}[K]}(\boldsymbol{M}) = \max_{k \in [K]} \left\{ \frac{1 - \rho\left(\left(\boldsymbol{L}^k\right)^\intercal \boldsymbol{L}^k - \sqrt{\boldsymbol{\pi}}^\intercal \sqrt{\boldsymbol{\pi}}\right)}{k} \right\}. \tag{B.74}$$

Discussion. Recall also that for any $k \in \mathbb{N}$, $\frac{\gamma_k^\dagger}{k}$ is a lower bound for the pseudo-spectral gap, hence an upper one for the mixing time, so that any k would yield a usable and conservative value for it. In practice this implies that the procedure makes iterative improvements to its known value of t_{mix}.

C Additional auxiliary results

C.1 Pseudo-spectral gap of reversible chains.

An immediate consequence of (1.1) and (1.4) is that in the reversible case,

$$\frac{\gamma_\star}{2 \ln(4/\pi_\star)} \leq \gamma_{\mathsf{ps}} \leq \frac{\gamma_\star}{1 - \gamma_\star} \left(\frac{\ln(1/\pi_\star) + 1}{\ln 2} + 2 \right), \tag{C.1}$$

i.e. $\gamma_{\mathsf{ps}} = \tilde{\Theta}(\gamma_\star)$. Lemma C.1 below establishes the stronger fact that for reversible Markov chains, the absolute spectral gap and the pseudo-spectral gap are within a multiplicative factor of 2; moreover, the pseudo-spectral gap has a closed form in terms of the absolute spectral gap:

Lemma C.1. *For reversible $\boldsymbol{M}$, we have*

$$\gamma_\star(\boldsymbol{M}) \leq \gamma_{\mathsf{ps}}(\boldsymbol{M}) = \gamma_\star(\boldsymbol{M})[2 - \gamma_\star(\boldsymbol{M})] \leq 2\gamma_\star(\boldsymbol{M}). \tag{C.2}$$

Proof. Reversibility implies $\boldsymbol{M}^\dagger = \boldsymbol{M}$, and so $\gamma_{\mathsf{ps}} = \max_{k \geq 1} \left\{ \frac{\gamma(\boldsymbol{M}^{2k})}{k} \right\}$. Denoting by $1 = \lambda_1 > \lambda_2 \geq \dots \geq \lambda_d$ the eigenvalues of $\boldsymbol{M}$, we have that for all $i \in [d]$ and $k \geq 1$, λ_i^{2k} is an eigenvalue for $\boldsymbol{M}^{2k}$, and furthermore $\lambda_\star^{2k}$ with $\lambda_\star = \max\{\lambda_2(\boldsymbol{M}), |\lambda_d(\boldsymbol{M})|\}$ is necessarily the second largest. We claim that

$$\gamma_{\mathsf{ps}} = \max_{k \geq 1} \left\{ \frac{1 - \lambda_\star^{2k}}{k} \right\} = 1 - \lambda_\star^2 \tag{C.3}$$

— that is, the maximum is achieved at $k = 1$. Indeed, $1 - \lambda_\star^{2k} = (1 - \lambda_\star^2) \left(\sum_{i=0}^{k-1} \lambda_\star^i\right)$ and the latter sum is at most k since $\lambda_\star < 1$. As a result, $\gamma_{\mathsf{ps}}(\boldsymbol{M}) = 1 - \lambda_\star^2 = 1 - (1 - \gamma_\star(\boldsymbol{M}))^2 = \gamma_\star(\boldsymbol{M})[2 - \gamma_\star(\boldsymbol{M})]$ and also $\gamma_\star(\boldsymbol{M}) \leq \gamma_{\mathsf{ps}}(\boldsymbol{M}) \leq 2\gamma_\star(\boldsymbol{M})$. □

C.2 Controlling the number of visits to all k-skipped r-offset chains.

The following Lemma C.2 is instrumental in proving Theorem 5.3. It quantifies the trajectory length sufficient to guarantee that that all states of an ergodic Markov chain with skipping rate $k \in [K]$ have been visited approximately according to their expected value with high confidence. Recall that k_{ps}, defined immediately following (4.8), is the smallest positive integer such that $\gamma_{\mathsf{ps}} = \gamma\left(\left(\boldsymbol{M}^\dagger\right)^{k_{\mathsf{ps}}} \boldsymbol{M}^{k_{\mathsf{ps}}}\right)/k_{\mathsf{ps}}$.

Lemma C.2. *For $(d, K) \in \mathbb{N}^2$, let $X_1, \dots, X_m \sim (\boldsymbol{M}, \boldsymbol{\mu})$ be a d-state time homogeneous ergodic Markov chain with pseudo-spectral gap γ_{ps} and stationary distribution $\boldsymbol{\pi}$ minorized by $\pi_\star$. For $0 < \eta \leq 1$, consider the event*

$$\mathcal{K}_\eta^K = \left\{ \max_{(k,i)\in[K]\times[d]} \left| N_i^{(k)} - \mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}\left[N_i^{(k)}\right] \right| \leq \eta \mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}\left[N_i^{(k)}\right] \right\},$$

where $N_i^{(k)}$ is defined in (4.2). Then, for $m \geq \frac{C_{\mathcal{K}} K^2}{\pi_\star \eta^2 \gamma_{\mathsf{ps}}} \ln\left(\frac{\sqrt{2\pi_\star^{-1}} dK}{\delta}\right)$, we have that $\mathcal{K}_\eta^K$ holds with probability at least $1 - \delta$, where $C_{\mathcal{K}} \leq 192$ is a universal constant.

Proof. By two applications of the union bound, followed by Paulin [2015, Proposition 3.10], we have

$$\mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\overline{\mathcal{K}_\eta^K}\right) \leq \sqrt{\|\boldsymbol{\mu}/\boldsymbol{\pi}\|_{2,\boldsymbol{\pi}}} \sum_{k=1}^{K} \sum_{i=1}^{d} \mathbf{P}_{\boldsymbol{M},\boldsymbol{\pi}}^{1/2}\left[\left| N_i^{(k)} - \mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}\left[N_i^{(k)}\right]\right| > \mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}\left[N_i^{(k)}\right]\right]. \tag{C.4}$$

This accounts for non-stationary starting distributions, and it remains to upper bound each of the summands. Fixing k and i, define $\phi_t(X_t) = \mathbf{1}\{t \equiv 0 \pmod{k} \text{ and } X_t = i\}$, $t \in [m]$. Putting $N_i^{(k)} = \sum_{t=1}^m \phi_t(X_t)$, we obtain $\mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}[\phi_t] = \mathbf{1}\{t \equiv 0 \pmod{k}\} \pi_i$ from the ergodic theorem for stationary chains. Thus, for $j \in [d]$, we have $|\phi_t(j) - \mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}[\phi_t]| \leq 1$. We cannot directly apply Paulin [2015, Theorem 3.4] for reasons that we detail in Remark C.1, and instead derive in Theorem C.1 (using his methods) a concentration bound tailored to our needs.

Let $\tilde{\sigma}_r$ defined as in Theorem C.1. For any t, we compute the closed form for the variance of ϕ_t:

$$\mathbf{Var}_{\boldsymbol{M},\boldsymbol{\pi}}[\phi_t] = \mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}[\phi_t](1 - \mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}[\phi_t]) = \mathbf{1}\{t \equiv 0 \pmod{k}\} \pi_i(1 - \pi_i). \tag{C.5}$$

Thus, for $1 \leq r \leq k_{\mathsf{ps}}$, we have

$$\tilde{\sigma}_r = \pi_i(1 - \pi_i) \left|\{0 \leq s \leq \lfloor (m - r)/k_{\mathsf{ps}} \rfloor : r + s k_{\mathsf{ps}} \equiv 0 \pmod{k}\}\right| \leq \frac{\pi_i}{4k_{\mathsf{ps}}}(m + k_{\mathsf{ps}}). \tag{C.6}$$

Invoking Theorem C.1,

$$\begin{aligned}\mathbf{P}_{\boldsymbol{M},\boldsymbol{\pi}}\left(\left| N_i^{(k)} - \mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}\left[N_i^{(k)}\right]\right| > \eta \mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}\left[N_i^{(k)}\right]\right) &\leq 2\exp\left(-\frac{\lfloor m/k \rfloor^2 \pi_i \eta^2 \gamma_{\mathsf{ps}}}{2(m + k_{\mathsf{ps}}) + 20\lfloor m/k \rfloor \eta}\right) \\ &\leq 2\exp\left(-\frac{m\pi_i\eta^2\gamma_{\mathsf{ps}}}{96k^2}\right),\end{aligned} \tag{C.7}$$

where the second inequality relies on the fact that $m \geq \max\left\{2K, \frac{1}{\gamma_{\mathsf{ps}}}\right\}$, and $\gamma_{\mathsf{ps}} = \frac{\gamma^\dagger_{k_{\mathsf{ps}}}}{k_{\mathsf{ps}}} \leq {k_{\mathsf{ps}}}^{-1}$. It follows that

$$\mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\overline{\mathcal{K}^K_\eta}\right) \leq \sqrt{2\pi_\star^{-1}} \sum_{k=1}^{K} \sum_{i=1}^{d} \exp\left(-\frac{m\pi_i \eta^2 \gamma_{\mathsf{ps}}}{96k^2}\right) \leq \sqrt{2\pi_\star^{-1}} K d \exp\left(-\frac{m\pi_\star \eta^2 \gamma_{\mathsf{ps}}}{192K^2}\right). \quad \text{(C.8)}$$

□

The following corollary was also proven in Hsu et al. [2015] – from a different starting point – and is merely included here for completeness.

Theorem C.1 (Bernstein-type bound for non-reversible Markov chain.)**.** *For $d \in \mathbb{N}$, let $X_1, \dots, X_m \sim (\boldsymbol{M}, \boldsymbol{\pi})$ be a d-state stationary time homogeneous ergodic Markov chain with pseudo-spectral gap γ_{ps}. Suppose that $\phi_1, \dots \phi_m : [d] \to \mathbb{R}$ are such that for some $C > 0$ and all $j \in [d], t \in [m]$, we have $|\phi_t(j) - \mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}[\phi_t]| \leq C$. Define, for $1 \leq r \leq k_{\mathsf{ps}}$, the quantities*

$$\tilde{\sigma}_r \triangleq \sum_{s=0}^{\lfloor (m-r)/k_{\mathsf{ps}} \rfloor} \mathbf{Var}_{\boldsymbol{M},\boldsymbol{\pi}}\left[\phi_{r+sk_{\mathsf{ps}}}\right]$$

and

$$\tilde{\sigma} \triangleq k_{\mathsf{ps}} \max_{1 \leq r \leq k_{\mathsf{ps}}} \tilde{\sigma}_r,$$

where k_{ps} is the smallest positive integer such that $\gamma_{\mathsf{ps}} = \frac{\gamma^\dagger_{k_{\mathsf{ps}}}}{k_{\mathsf{ps}}}$. Then

$$\mathbf{P}_{\boldsymbol{M},\boldsymbol{\pi}}\left(\left|\sum_{t=t}^{m} \phi_t(X_t) - \mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}\left[\sum_{t=t}^{m} \phi_t(X_t)\right]\right| > \varepsilon\right) \leq 2\exp\left(-\frac{\varepsilon^2 \gamma_{\mathsf{ps}}}{8\tilde{\sigma} + 20\varepsilon C}\right). \quad \text{(C.9)}$$

Proof. Our proof proceeds along the lines of Paulin [2015, Theorem 3.4]. We begin by partitioning the sequence $(\phi_1(X_1), \dots, \phi_m(X_m))$ into k_{ps} skipped sub-sequences indexed by $1 \leq r \leq k_{\mathsf{ps}}$,

$$\left(\phi(\boldsymbol{X}_{r+sk_{\mathsf{ps}}})\right)_{0 \leq s \leq \lfloor (m-r)/k_{\mathsf{ps}} \rfloor}, \quad \text{(C.10)}$$

with

$$\left(\boldsymbol{X}_{r+sk_{\mathsf{ps}}}\right)_{0 \leq s \leq \lfloor (m-r)/k_{\mathsf{ps}} \rfloor} \sim (\boldsymbol{M}^{k_{\mathsf{ps}}}, \boldsymbol{\pi}\boldsymbol{M}^{r-1}) \quad \text{(C.11)}$$

from Jensen's inequality, for any distribution $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_{k_{\mathsf{ps}}}) \in \Delta_{k_{\mathsf{ps}}}$ and $\theta > 0$,

$$\mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}\left[\exp\left(\theta \sum_{t=1}^{m} \phi_t(X_t)\right)\right] \leq \sum_{r=1}^{k_{\mathsf{ps}}} \alpha_i \mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}\left[\exp\left(\frac{\theta}{\alpha_i} \sum_{s=0}^{\lfloor (m-r)/k_{\mathsf{ps}} \rfloor} \phi_{r+sk_{\mathsf{ps}}}(X_{r+sk_{\mathsf{ps}}})\right)\right]. \quad \text{(C.12)}$$

In particular for $i \in [k_{\mathsf{ps}}]$ and $\alpha_i = \frac{1}{k_{\mathsf{ps}}}$, it follows from Paulin [2015, (5.6), (5.7)] that

$$\begin{aligned}\mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}\left[\exp\left(\theta \sum_{t=t}^{m} \phi_t(X_t)\right)\right] &\leq \frac{1}{k_{\mathsf{ps}}} \sum_{r=1}^{k_{\mathsf{ps}}} \exp\left(\frac{2\theta^2 {k_{\mathsf{ps}}}^2 \tilde{\sigma}_r}{\gamma^\dagger_{k_{\mathsf{ps}}}} \left(1 - \frac{10\theta k_{\mathsf{ps}}}{\gamma^\dagger_{k_{\mathsf{ps}}}}\right)^{-1}\right) \\ &\leq \exp\left(\frac{2\theta^2 k_{\mathsf{ps}} \max_{1 \leq r \leq k_{\mathsf{ps}}} \tilde{\sigma}_r}{\gamma_{\mathsf{ps}}} \left(1 - \frac{10\theta}{\gamma_{\mathsf{ps}}}\right)^{-1}\right).\end{aligned} \quad \text{(C.13)}$$

From the estimate on its moment generating function, the random variable $\sum_{t=t}^{m} \phi_t(X_t)$ is sub-gamma with variance factor $\nu := \frac{4k_{\mathsf{ps}} \max_{1 \leq r \leq k_{\mathsf{ps}}} \tilde{\sigma}_r}{\gamma_{\mathsf{ps}}}$ and scale parameter $B := \frac{10}{\gamma_{\mathsf{ps}}}$, i.e.

$$\mathbf{E}_{\boldsymbol{M},\boldsymbol{\pi}}\left[\exp\left(\theta \sum_{t=t}^{m} \phi_t(X_t)\right)\right] \leq \exp\left(\frac{\nu\theta^2}{2(1-B\theta)}\right).$$

An application of Markov's inequality concludes the proof. □

Remark C.1. *Curiously, Paulin [2015, Theorem 3.4], in its most general form, does not seem to yield useful bounds for our particular problem. Indeed, one is faced with bounding the ratio* $\frac{R(m,k,k_{\mathsf{ps}})}{k_{\mathsf{ps}}}$ *where* $R(m,k,k_{\mathsf{ps}}) \triangleq \frac{\sum_{1 \leq r \leq k_{\mathsf{ps}}} \sqrt{\tilde{\sigma}_r}}{\min_{1 \leq r \leq k_{\mathsf{ps}}} \sqrt{\tilde{\sigma}_r}}$. *However, for our considered* $\phi_1, \dots, \phi_m$*, the latter quantity is computed to be*

$$R(m,k,k_{\mathsf{ps}}) = \frac{\sum_{1 \leq r \leq k_{\mathsf{ps}}} \sqrt{|\{0 \leq s \leq \lfloor (m-r)/k_{\mathsf{ps}} \rfloor : r + s k_{\mathsf{ps}} \equiv 0 \pmod{k}\}|}}{\min_{1 \leq r \leq k_{\mathsf{ps}}} \sqrt{|\{0 \leq s \leq \lfloor (m-r)/k_{\mathsf{ps}} \rfloor : r + s k_{\mathsf{ps}} \equiv 0 \pmod{k}\}|}},$$

and in particular, the the denominator might be zero. (For a concrete example, take $k = k_{\mathsf{ps}} = 5$*, and* $r = 1$*.) We hope that our observation is a step towards answering Paulin [2015, Remark 3.7].*

C.3 Empirical bounds for learning the stationary distribution of a chain.

The technical results of this section are at the core of three different improvements. First, to extend the perturbation results of Hsu et al. [2015] to the non-reversible setting, second, to improve the width of their confidence intervals for $\pi_\star$ roughly by a factor $\tilde{\mathcal{O}}(\sqrt{d})$ in the reversible case through the use of a vector-valued martingale technique, and third, to reduce the computation cost of their intervals from $\mathcal{O}(d^3)$ to $\mathcal{O}(d^2)$, by showing that we can trade the computation of a pseudo-inverse with the already computed estimator for the absolute spectral gap.

Lemma C.3 (Perturbation bound for stationary distribution of ergodic Markov chain)**.** *Let* $\boldsymbol{M}_1$ *(resp.* $\boldsymbol{M}_2$*) be an ergodic Markov chain with stationary distribution* $\boldsymbol{\pi}_1$ *(resp.* $\boldsymbol{\pi}_2$*) minorized by* $\pi_\star(\boldsymbol{M}_1)$ *(resp.* $\pi_\star(\boldsymbol{M}_2)$*) and pseudo-spectral gap* $\gamma_{\mathsf{ps}}(\boldsymbol{M}_1)$ *(resp.* $\gamma_{\mathsf{ps}}(\boldsymbol{M}_2)$*). Then*

$$\|\boldsymbol{\pi}_1 - \boldsymbol{\pi}_2\|_\infty \leq \frac{C_{\mathcal{K}}}{\max\{\gamma_{\mathsf{ps}}(\boldsymbol{M}_1), \gamma_{\mathsf{ps}}(\boldsymbol{M}_2)\}} \ln\left(2\sqrt{\frac{2}{\max\{\pi_\star(\boldsymbol{M}_1), \pi_\star(\boldsymbol{M}_2)\}}}\right) \|\boldsymbol{M}_1 - \boldsymbol{M}_2\|_\infty, \tag{C.14}$$

where $C_{\mathcal{K}}$ *is the universal constant in Lemma C.2.*

Proof. The proof follows a similar argument as to the one at [Hsu et al., 2017, Lemma 8.9] except that (C.7), valid for non-reversible Markov chains, is used instead of Levin et al. [2009, Chapter 12]. Cho and Meyer [2001, Section 3.3] have shown established that $\|\boldsymbol{\pi}_1 - \boldsymbol{\pi}_2\|_\infty \leq \kappa_1 \|\boldsymbol{M}_1 - \boldsymbol{M}_2\|_\infty$, where $\kappa_1 = \frac{1}{2}\max_{j \in [d]} \{\pi_j \max_{i \neq j} \mathbf{E}_{\boldsymbol{M}_1, \delta_i}[\tau_j]\}$, $\boldsymbol{\delta}_i$ is the distribution whose support consists of state i, and τ_j is the first time state j is visited by the Markov chain. We proceed by with upper bounding $\mathbf{E}_{\boldsymbol{M}_1, \boldsymbol{\delta}_i}[\tau_j]$. Setting $\eta := k := 1$ in (C.7), we have, for $(m, i, j) \in \mathbb{N} \times [d]^2$, that

$$\mathbf{P}_{\boldsymbol{M}_1, \boldsymbol{\delta}_i}(\tau_j > m) = \mathbf{P}_{\boldsymbol{M}_1, \boldsymbol{\delta}_i}(N_j = 0) \leq \mathbf{P}_{\boldsymbol{M}_1, \boldsymbol{\delta}_i}(|N_j - \pi_j m| > \pi_j m). \tag{C.15}$$

Further, since $\|\boldsymbol{\delta}_i/\boldsymbol{\pi}_1\|_{2,\boldsymbol{\pi}_1} = \frac{1}{\boldsymbol{\pi}_1(i)}$ for $m \geq \frac{C_{\mathcal{K}}}{\boldsymbol{\pi}_1(j)\gamma_{\mathsf{ps}}(\boldsymbol{M}_1)} \ln\left(2\sqrt{\frac{2}{\boldsymbol{\pi}_1(i)}}\right)$, it follows that $\mathbf{P}_{\boldsymbol{M}_1,\boldsymbol{\delta}_i}(\tau_j > m) \leq \frac{1}{2}$. Hence,

$$\begin{aligned}\mathbf{E}_{\boldsymbol{M}_1,\boldsymbol{\delta}_i}[\tau_j] &= \sum_{\ell\geq 0}\mathbf{P}_{\boldsymbol{M}_1,\boldsymbol{\delta}_i}(\tau_j > \ell) = \sum_{0\leq \ell < m}\mathbf{P}_{\boldsymbol{M}_1,\boldsymbol{\delta}_i}(\tau_j > \ell) + \sum_{k\geq 1}\sum_{km\leq \ell<(k+1)m}\mathbf{P}_{\boldsymbol{M}_1,\boldsymbol{\delta}_i}(\tau_j > \ell) \\ &\leq m + \sum_{k\geq 1}\sum_{km\leq \ell<(k+1)m}\mathbf{P}_{\boldsymbol{M}_1,\boldsymbol{\delta}_i}(\tau_j > km).\end{aligned} \tag{C.16}$$

By the Markov property, $\mathbf{P}_{\boldsymbol{M}_1,\boldsymbol{\delta}_i}(\tau_j > m) \leq 1/2 \implies \mathbf{P}_{\boldsymbol{M}_1,\boldsymbol{\delta}_i}(\tau_j > km) \leq 2^{-k}$, $k \geq 1$, and so

$$\mathbf{E}_{\boldsymbol{M}_1,\boldsymbol{\delta}_i}[\tau_j] \leq m + m\sum_{k\geq 1}2^{-k} = m + m\frac{2^{-1} - \lim_{k\to\infty}2^{-k}}{1-1/2} = 2m, \tag{C.17}$$

which completes the proof. □

The following corollary is a natural extension of a similar result of Hsu et al. [2015] to matrix-valued martingales.

Corollary C.1 (to Theorem D.1)**.** *Consider a matrix martingale difference sequence $\{\boldsymbol{Y}_t : t = 1,2,3,\ldots\}$ with dimensions $d_1\times d_2$, such that $\|\boldsymbol{Y}_t\|_2 \leq 1$ almost surely for $t = 1,2,\ldots$. Then with $\|\boldsymbol{\Sigma}_m\|_2$ as defined in Theorem D.1, for all $\varepsilon \geq 0$ and $\sigma^2 > 0$,*

$$\mathbf{P}\left(\left\|\sum_{t=1}^{m}\boldsymbol{Y}_t\right\|_2 > \sqrt{2\|\boldsymbol{\Sigma}_m\|_2\,\tau_{\delta,m}} + \frac{5}{3}\tau_{\delta,m}\right) \leq \delta,$$

$$\text{where } \tau_{\delta,m} \triangleq \inf\left\{t > 0 : \left(1 + \lceil\ln(2m/t)\rceil_+\right)(d_1+d_2)e^{-t} \leq \delta\right\} = \mathcal{O}\left(\ln\left(\frac{(d_1+d_2)\ln m}{\delta}\right)\right). \tag{C.18}$$

Proof. Let $m \in \mathbb{N}$. From properties of the spectral norm (sub-additivity, sub-multiplicativity, invariance under transpose), an application of Jensen's inequality, and $\|\boldsymbol{Y}_t\|_2 \leq 1$, we have

$$\|\boldsymbol{W}_{col,k}\|_2 = \left\|\sum_{t=1}^{k}\mathbf{E}_{t-1}\left[\boldsymbol{Y}_t\boldsymbol{Y}_t^\intercal\right]\right\|_2 \leq \sum_{t=1}^{k}\mathbf{E}_{t-1}\|\boldsymbol{Y}_t\boldsymbol{Y}_t^\intercal\|_2 = \sum_{t=1}^{k}\mathbf{E}_{t-1}\|\boldsymbol{Y}_t\|_2^2 \leq k, \tag{C.19}$$

and similarly $\|\boldsymbol{W}_{\text{row},k}\|_2 \leq k$, concluding that $\|\boldsymbol{\Sigma}_k\|_2 \leq k$. Let $\varepsilon > 0$ and $\sigma > 0$. Define $\sigma_i^2 \triangleq \frac{e^i\varepsilon}{2}$ for $i \in \{0,1,\ldots,\lceil\ln(2m/\varepsilon)\rceil_+\}$, and $\sigma_{-1}^2 \triangleq -\infty$. Observing that $\sqrt{2\max\{\sigma_0^2, e\|\boldsymbol{\Sigma}_k\|_2\,\varepsilon\}} \leq \sqrt{2\|\boldsymbol{\Sigma}_k\|_2\,\varepsilon} + \varepsilon$, we have

$$\begin{aligned}&\mathbf{P}\left(\exists k\in[m], \left\|\sum_{t=1}^{k}\boldsymbol{Y}_t\right\|_2 > \sqrt{2\|\boldsymbol{\Sigma}_k\|_2\,\varepsilon} + \frac{5}{3}\varepsilon\right) \\ &\leq \sum_{i=0}^{\lceil\ln(2m/\varepsilon)\rceil_+}\mathbf{P}\left(\exists k\in[m], \left\|\sum_{t=1}^{k}\boldsymbol{Y}_t\right\|_2 > \sqrt{2\max\{\sigma_0^2, e\|\boldsymbol{\Sigma}_k\|_2\}\varepsilon} + \frac{2}{3}\varepsilon \text{ and } \|\boldsymbol{\Sigma}_k\|_2 \in (\sigma_{i-1}^2, \sigma_i^2]\right) \\ &\leq \sum_{i=0}^{\lceil\ln(2m/\varepsilon)\rceil_+}\mathbf{P}\left(\exists k\in[m], \left\|\sum_{t=1}^{k}\boldsymbol{Y}_t\right\|_2 > \sqrt{2\max\{\sigma_0^2, e\sigma_{i-1}^2\}\varepsilon} + \frac{2}{3}\varepsilon \text{ and } \|\boldsymbol{\Sigma}_k\|_2 \in (\sigma_{i-1}^2, \sigma_i^2]\right) \\ &\leq \sum_{i=0}^{\lceil\ln(2m/\varepsilon)\rceil_+}\mathbf{P}\left(\exists k\in[m], \left\|\sum_{t=1}^{k}\boldsymbol{Y}_t\right\|_2 > \sqrt{2\sigma_i^2\varepsilon} + \frac{2}{3}\varepsilon \text{ and } \|\boldsymbol{\Sigma}_k\|_2 \leq \sigma_i^2\right).\end{aligned} \tag{C.20}$$

Applying Theorem D.1 deviation size $\sqrt{2\sigma^2\varepsilon} + \varepsilon$ yields

$$\mathbf{P}\left(\exists k \in [m], \left\|\sum_{t=1}^{k} \boldsymbol{Y}_t\right\|_2 > \sqrt{2\left\|\boldsymbol{\Sigma}_k\right\|_2 \varepsilon} + \frac{5}{3}\varepsilon\right) \leq \left(1 + \lceil \ln(2m/\varepsilon) \rceil_+\right)(d_1 + d_2)e^{-\varepsilon}, \tag{C.21}$$

which concludes the proof. □

The following lemma gives an empirical high-confidence bound for the problem of learning an unknown Markov chain $\boldsymbol{M}$ with respect to the $\|\cdot\|_\infty$ norm.

Lemma C.4. *Let $X_1, \dots, X_m \sim (\boldsymbol{M}, \boldsymbol{\mu})$ a d-state Markov chain and $\widehat{\boldsymbol{M}}$ defined as in Section B.2.1. Then, with probability at least $1 - \delta$,*

$$\left\|\widehat{\boldsymbol{M}} - \boldsymbol{M}\right\|_\infty \leq 4\tau_{\delta/d,m}\sqrt{\frac{d}{N_{\min}}}, \tag{C.22}$$

where $\tau_{\delta,m} = \inf\left\{t > 0 : \left(1 + \lceil \ln(2m/t) \rceil_+\right)(d+1)e^{-t} \leq \delta\right\}$, and $N_{\min}$ is defined in (4.4).

Proof. For a fixed i, define the row vector sequence $\boldsymbol{Y}$ by

$$\boldsymbol{Y}_0 = 0, \boldsymbol{Y}_t = \frac{1}{\sqrt{2}}\left[\mathbf{1}\{X_{t-1} = i\}\left(\mathbf{1}\{X_t = j\} - \boldsymbol{M}(i,j)\right)\right]_j. \tag{C.23}$$

Notice that $\sum_{t=1}^m \boldsymbol{Y}_t = \left[N_{ij} - N_i\boldsymbol{M}(i,j)\right]_j$, and from the Markov property $\mathbf{E}_{t-1}[\boldsymbol{Y}_t] = \mathbf{0}$, so that $\boldsymbol{Y}_t$ defines a vector valued martingale difference. Let $\boldsymbol{W}_{\text{col},m}, \boldsymbol{W}_{\text{row},m}$ and $\|\boldsymbol{\Sigma}_m\|_2$ be as defined in Theorem D.1. Then

$$\begin{aligned}\boldsymbol{Y}_t\boldsymbol{Y}_t^\intercal = \|\boldsymbol{Y}_t\|_2^2 &= \frac{1}{2}\mathbf{1}\{X_{t-1} = i\}\left(1 + \|\boldsymbol{M}(i,\cdot)\|_2^2 - 2\boldsymbol{M}(i,X_t)\right) \leq \mathbf{1}\{X_{t-1} = i\} \\ \left[\boldsymbol{Y}_t^\intercal\boldsymbol{Y}_t\right]_{(j,\ell)} &= \frac{\mathbf{1}\{X_{t-1} = i\}}{2}\left[\mathbf{1}\{X_t = j\} - \boldsymbol{M}(i,j)\mathbf{1}\{X_t = \ell\} - \boldsymbol{M}(i,\ell)\right]_{(j,\ell)},\end{aligned} \tag{C.24}$$

so that $\boldsymbol{W}_{\text{col},m} \leq N_i$. Further, $\|\boldsymbol{Y}_t^\intercal\boldsymbol{Y}_t\|_2 \leq \sqrt{\|\boldsymbol{Y}_t^\intercal\boldsymbol{Y}_t\|_1 \|\boldsymbol{Y}_t^\intercal\boldsymbol{Y}_t\|_\infty}$ (Hölder's inequality), $\|\boldsymbol{Y}_t^\intercal\boldsymbol{Y}_t\|_2 \leq \mathbf{1}\{X_{t-1} = i\}$, and $\|\boldsymbol{\Sigma}_m\|_2 \leq N_i$ (sub-additivity of the norm and Jensen's inequality).

Corollary C.1 and another application of Hölder's inequality, yields, for all $i \in [d]$,

$$\begin{aligned}&\mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\widehat{\boldsymbol{M}}(i,\cdot) - \boldsymbol{M}(i,\cdot)\right\|_1 > \sqrt{\frac{2d\tau_{\delta/d,m}}{N_i}} + \frac{5\tau_{\delta/d,m}\sqrt{d}}{3N_i}\right) \\ &\leq \mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\widehat{\boldsymbol{M}}(i,\cdot) - \boldsymbol{M}(i,\cdot)\right\|_1 > \left(\sqrt{2\|\boldsymbol{\Sigma}_m\|_2\,\tau_{\delta/d,m}} + \frac{5}{3}\tau_{\delta/d,m}\right)\frac{\sqrt{d}}{N_i}\right) \\ &\leq \mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\sum_{t=1}^m \boldsymbol{Y}_t\right\|_2 > \sqrt{2\|\boldsymbol{\Sigma}_m\|_2\,\tau_{\delta/d,m}} + \frac{5}{3}\tau_{\delta/d,m}\right) \leq \frac{\delta}{d}.\end{aligned} \tag{C.25}$$

Finally, the observation that

$$\sqrt{\frac{2d\tau_{\delta/d,m}}{N_{\min}}} + \frac{5\tau_{\delta/d,m}\sqrt{d}}{3N_{\min}} \leq 4\tau_{\delta/d,m}\sqrt{\frac{d}{N_{\min}}} \tag{C.26}$$

completes the proof. □

Remark C.2. *In the discrete distribution learning model Kearns et al. [1994], Waggoner [2015], the minimax complexity for learning a $[d]$-supported distribution up to precision ε with high confidence is of order $m = \tilde{\Theta}\left(\frac{d}{\varepsilon^2}\right)$, and so the bound in Lemma C.4 is in a sense optimal, and can be thought of as an empirical version of the bounds derived in Wolfer and Kontorovich [2019].*

Corollary C.2 (to Lemma C.4)**.** *Let $X_1, \dots, X_m \sim (\boldsymbol{M}, \boldsymbol{\mu})$ a d-state Markov chain and $\widehat{\boldsymbol{M}}^\dagger$ is such that $\widehat{\boldsymbol{M}}^\dagger(i,j) \triangleq \frac{N_{ji}}{N_i}$. Then, with probability at least $1-\delta$,*

$$\left\| \widehat{\boldsymbol{M}}^\dagger - \boldsymbol{M}^\dagger \right\|_\infty \leq 4\tau_{\delta/d,m}\sqrt{\frac{d}{N_{\min}}}, \tag{C.27}$$

where $\boldsymbol{M}^\dagger$ is the time reversal of $\boldsymbol{M}$, $\tau_{\delta,m}$ is as in Lemma C.4, and $N_{\min}$ is defined in (4.4)*.*

Proof. The only change is to consider the time-reversed martingale

$$\boldsymbol{Y}_0 = 0, \boldsymbol{Y}_t = \frac{1}{\sqrt{2}}\left[\mathbf{1}\left\{X_t = i\right\}\left(\mathbf{1}\left\{X_{t-1} = j\right\} - \boldsymbol{M}^\dagger(i,j)\right)\right]_j \tag{C.28}$$

and mimic the proof of Lemma C.4. This yields the claim with probability at least $1-\delta$. □

The next lemma provides a perturbation bound for the stationary distribution of a reversible Markov chain in terms of the symmetrized estimator of the absolute spectral gap.

Corollary C.3. *Let $X_1, \dots, X_m \sim \boldsymbol{M}$ an ergodic reversible Markov chain with stationary distribution $\boldsymbol{\pi}$ minorized by $\pi_\star$ and absolute-spectral gap $\gamma_\star$, and let $\widehat{\boldsymbol{M}}, \widehat{\pi}$ the normalized-count estimators for $\boldsymbol{M}, \boldsymbol{\pi}$. Then*

$$\left\|\widehat{\boldsymbol{\pi}} - \boldsymbol{\pi}\right\|_\infty \leq \frac{C_{\mathcal{K}}}{\widehat{\gamma}_\star}\ln\left(2\sqrt{\frac{2m}{N_{\min}}}\right)\left(\left\|\widehat{\boldsymbol{M}} - \boldsymbol{M}\right\|_\infty + \left\|\widehat{\boldsymbol{M}}^\dagger - \boldsymbol{M}^\dagger\right\|_\infty\right), \tag{C.29}$$

where $C_{\mathcal{K}}$ is the universal constant in Lemma C.2, and $\widehat{\gamma}_\star \triangleq \gamma_\star\left(\frac{1}{2}\left(\widehat{\boldsymbol{M}} + \widehat{\boldsymbol{M}}^\dagger\right)\right)$, the absolute spectral gap of the additive reversiblization of $\widehat{\boldsymbol{M}}$.

Proof. As $\widehat{\boldsymbol{\pi}}$ is also the stationary distribution of $\frac{1}{2}\left(\widehat{\boldsymbol{M}} + \widehat{\boldsymbol{M}}^\dagger\right)$ we have the the perturbation bound $\left\|\widehat{\boldsymbol{\pi}} - \boldsymbol{\pi}\right\|_\infty \leq \hat{\kappa}\left\|\frac{1}{2}\left(\widehat{\boldsymbol{M}} + \widehat{\boldsymbol{M}}^\dagger\right) - \boldsymbol{M}\right\|_\infty$, where $\hat{\kappa} = \kappa\left(\frac{1}{2}\left(\widehat{\boldsymbol{M}} + \widehat{\boldsymbol{M}}^\dagger\right)\right)$. By reversibility of $\boldsymbol{M}$ and norm sub-additivity,

$$\left\|\widehat{\boldsymbol{\pi}} - \boldsymbol{\pi}\right\|_\infty \leq \hat{\kappa}\frac{1}{2}\left(\left\|\widehat{\boldsymbol{M}} - \boldsymbol{M}\right\|_\infty + \left\|\widehat{\boldsymbol{M}}^\dagger - \boldsymbol{M}^\dagger\right\|_\infty\right). \tag{C.30}$$

The proof is concluded by invoking Lemmas C.3 and C.1. □

D Results from the literature

Lemma D.1 (from Kazakos, 1978)**.** *Let $(\boldsymbol{M}_0, \boldsymbol{\mu}_0)$ and $(\boldsymbol{M}_1, \boldsymbol{\mu}_1)$ two d-state Markov chains, and define*

$$[\boldsymbol{\mu}_0, \boldsymbol{\mu}_1]_{\surd} \triangleq \left[\sqrt{\boldsymbol{\mu}_0(i)\boldsymbol{\mu}_1(i)}\right]_{i\in[d]} \text{ and } [\boldsymbol{M}_0, \boldsymbol{M}_1]_{\surd} \triangleq \left[\sqrt{\boldsymbol{M}_0(i,j)\boldsymbol{M}_1(i,j)}\right]_{(i,j)\in[d]^2}.$$

Then for $\boldsymbol{X} = (X_1, \dots, X_m), \boldsymbol{Y} = (Y_1, \dots, Y_m)$ two trajectories of length m sampled respectively from two $(\boldsymbol{M}_0, \boldsymbol{\mu}_0)$ and $(\boldsymbol{M}_1, \boldsymbol{\mu}_1)$, it holds that

$$1 - H^2(\boldsymbol{X}, \boldsymbol{Y}) = [\boldsymbol{\mu}_0, \boldsymbol{\mu}_1]_{\surd}^\intercal \cdot \left([\boldsymbol{M}_0, \boldsymbol{M}_1]_{\surd}\right)^m \cdot \mathbf{1}.$$

Theorem D.1 (Rectangular Matrix Freedman, Tropp [2011, Corollary 1.3])**.** *Consider a matrix martingale* $\{\boldsymbol{X}_t : t = 0, 1, 2, \ldots\}$ *whose values are matrices with dimension* $d_1 \times d_2$*, and let* $\{\boldsymbol{Y}_t : t = 1, 2, 3, \ldots\}$ *be the difference sequence. Assume that the difference sequence is uniformly bounded with respect to the spectral norm, i.e.* $\exists R > 0$*,*

$$\|\boldsymbol{Y}_t\|_2 \leq R \textit{ almost surely for } t = 1, 2, \ldots$$

Define two predictable quadratic variation processes for this martingale:

$$\boldsymbol{W}_{\text{col},m} \triangleq \sum_{t=1}^{m} \mathbf{E}_{t-1}\left[\boldsymbol{Y}_t \boldsymbol{Y}_t^{\intercal}\right] \textit{ and } \boldsymbol{W}_{\text{row},m} \triangleq \sum_{t=1}^{m} \mathbf{E}_{t-1}\left[\boldsymbol{Y}_t^{\intercal} \boldsymbol{Y}_t\right] \textit{ for } m = 1, 2, 3, \ldots, \tag{D.1}$$

and write $\|\boldsymbol{\Sigma}_m\|_2 = \max\left\{\|\boldsymbol{W}_{\text{col},m}\|_2, \|\boldsymbol{W}_{\text{row},m}\|_2\right\}$*. Then, for all* $\varepsilon \geq 0$ *and* $\sigma^2 > 0$*,*

$$\mathbf{P}\left(\max_{m \geq 0} \left\|\sum_{t=1}^{m} \boldsymbol{Y}_t\right\|_2 > \varepsilon \textit{ and } \|\boldsymbol{\Sigma}_m\|_2 \leq \sigma^2\right) \leq (d_1 + d_2) \exp\left(-\frac{\varepsilon^2/2}{\sigma^2 + R\varepsilon/3}\right) \tag{D.2}$$