

Minimax Estimation of the Pseudo-Spectral Gap at Arbitrary Relative Precision

Abstract

We study estimation of the pseudo-spectral gap of a finite ergodic Markov chain from a single trajectory. For a known state space of size $d \geq 2$, stationary minimum π_* , and pseudo-spectral gap γ_{ps} , we construct an estimator with relative error $\varepsilon \in (0, 1/2)$ and confidence $1 - \delta$ using $O((\pi_* \gamma_{\text{ps}} \varepsilon^2)^{-1} \log(d/(\delta \pi_* \gamma_{\text{ps}})))$ transitions. The guarantee holds from every initial distribution, and the estimator requires no knowledge of the stationary law or gap. This extends the constant-factor inverse-gap bound of Wolfer and Kontorovich (Ann. Appl. Probab., 2024) to arbitrary relative precision, improving their general-precision dependence from γ_{ps}^{-3} to γ_{ps}^{-1} . On classes with $\pi_* \geq p$ and $\gamma_{\text{ps}} \geq g$, where $0 < p \leq 1/d$ and $0 < g \leq 1/8$, we prove a matching lower bound $\Omega((pg\varepsilon^2)^{-1})$ at constant confidence, even for reversible chains with a common known stationary law. The main estimate compares a bipartite Laplacian with its empirical counterpart relative to the Laplacian plus a regularization term. Positive semidefinite edge matrices allow the martingale variance to be controlled by occupation counts. Separate incoming and outgoing normalization ensures empirical contractivity, so the lagged estimates can be combined without additional inverse-gap factors.

Note. This paper was generated entirely by AI, including MiMo, using an automated research pipeline developed by Chenghua Liu and Hanyu Li.

1 Introduction

Concentration bounds for Markov observations depend on the relaxation of the underlying chain. When its transition matrix is unknown, the relaxation parameter must itself be estimated from dependent data. For nonreversible chains, Paulin’s pseudo-spectral gap controls concentration through the singular-value contraction of powers of the transition operator [4]. Estimates of this quantity also enter empirical PAC-Bayes bounds for Markov data [3]. The statistical difficulty is to estimate a potentially small spectral deficit while accounting for the unknown stationary law that defines the relevant inner product.

For reversible chains, Hsu et al. [2] obtained inverse-gap sample complexity for relative estimation of the absolute spectral gap, and constructed fully data-dependent confidence intervals. Wolfer and Kontorovich [6] extended gap estimation beyond reversibility by using multiplicative reversibilizations of powers. Their later work established an upper bound of $\tilde{O}((\pi_* \gamma_{\text{ps}})^{-1})$ for a fixed multiplicative factor [7, Theorem 2.2], but the bound for arbitrary relative error ε remained $\tilde{O}((\pi_* \gamma_{\text{ps}}^3 \varepsilon^2)^{-1})$ [7, Theorem 2.3]. Here the tilde suppresses logarithmic factors. Their Open Problem 2.1 asks to close the remaining precision and stationary-mass dependence between the upper and lower bounds.

We obtain the inverse-gap rate at every relative precision $0 < \varepsilon < 1/2$, with only one logarithmic factor. A complementary two-point construction proves the joint lower bound $\Omega((pg\varepsilon^2)^{-1})$ on classes with stationary minimum at least p and pseudo-spectral gap at least g , for $g \leq 1/8$. The earlier gap-estimation lower bound in [6, Theorem 4] is proportional to $d/(\gamma_{\text{ps}} \varepsilon^2)$, apart from its confidence factor; it does not supply the required $1/p$ dependence when $p \ll 1/d$. Our construction keeps the stationary distribution fixed and known. Thus the three factors p^{-1} , g^{-1} , and ε^{-2} arise together even without the need to estimate that distribution.

The proof retains the lagged-operator framework of [6] and the reversible-dilation viewpoint of [7]. The new estimate is a relative perturbation bound for the associated bipartite Laplacian. At lag k , this Laplacian has second-smallest eigenvalue $h_k = 1 - \sigma_2(D_\pi^{1/2} P^k D_\pi^{-1/2})$. Regularizing at scale $k\gamma_{\text{ps}}$ yields an eigenvalue error proportional to $\varepsilon k\gamma_{\text{ps}}$, rather than requiring the same small additive operator error at every lag. Dividing by k then gives the required gap accuracy. The regularization is used only in the proof, not in the estimator.

Two properties make this comparison effective. Each observed transition contributes a positive semidefinite edge matrix; conditional centering gives a matrix martingale whose predictable variation is controlled by occupation counts. In addition, normalizing incoming and outgoing counts separately makes every empirical operator a contraction. A candidate at lag k is therefore at most $1/k$, so lags beyond the inverse gap need no concentration estimate. For the remaining lags, thinning reduces the sample length but increases the pseudo-spectral gap by the same order of magnitude.

We now specify the model and results. Let P be a row-stochastic transition matrix on the known state space $[d] = \{1, \dots, d\}$, where $d \geq 2$. Throughout, ergodic means irreducible and aperiodic. Write its stationary distribution as a row vector π , and put $D_\pi = \text{diag}(\pi)$, $\Pi = \mathbf{1}\pi$, and $\pi_\star = \min_i \pi_i$. Here $\mathbf{1}$ denotes the all-ones column vector. Functions on $[d]$ are column vectors, with $\|f\|_\pi^2 = \sum_i \pi_i f_i^2$; the corresponding operator norm is also denoted by $\|\cdot\|_\pi$. The time reversal is $P^* = D_\pi^{-1} P^T D_\pi$. For integers $k \geq 1$, define

$$L_k = D_\pi^{1/2} P^k D_\pi^{-1/2}, \quad (1)$$

$$a_k = 1 - \sigma_2(L_k)^2 = 1 - \|P^k - \Pi\|_\pi^2, \quad (2)$$

$$\gamma = \gamma_{\text{ps}}(P) = \max_{k \geq 1} \frac{a_k}{k}. \quad (3)$$

Singular values are ordered decreasingly, with multiplicities. In particular, an additional unit singular value makes $a_k = 0$. Equivalently, a_k is the spectral gap of $(P^*)^k P^k$, so (3) is Paulin's pseudo-spectral gap [4], in the convention of [7]. Since $P^k \rightarrow \Pi$ in finite dimension, $a_k \rightarrow 1$. Also $0 \leq a_k/k \leq 1/k$, so the maximum is attained and $0 < \gamma \leq 1$.

We observe $X_0, \dots, X_n$, comprising n transitions. Relative error ε means $|\hat{\gamma} - \gamma| \leq \varepsilon\gamma$. All logarithms are natural. Unsubscripted matrix norms are Euclidean operator norms, and $A \preceq B$ means that $B - A$ is positive semidefinite. The one-state case has a known deterministic transition matrix and is excluded from the estimation problem.

Theorem 1 (Uniform upper bound). *There is a statistic $\hat{\gamma}_n$, defined in (7), such that, for every ergodic P , every initial distribution, $0 < \varepsilon < 1/2$, and $0 < \delta < 1$,*

$$n \geq \frac{C}{\pi_\star \gamma \varepsilon^2} \log \frac{8d}{\delta \pi_\star \gamma} \implies \mathbb{P}(|\hat{\gamma}_n - \gamma| > \varepsilon\gamma) \leq \delta. \quad (4)$$

One may take $C = 73728$. The statistic uses only the trajectory and the known state space; it does not use π , $\pi_\star$, γ , ε , or δ .

The sample-size guarantee depends on the true parameters, although the statistic does not. The theorem is a fixed-sample estimation result, not a fully empirical confidence interval or a certified rule for stopping the observation of the chain.

For the minimax statement, define

$$\mathcal{C}(d, p, g) = \{P : P \text{ is ergodic on } [d], \pi_\star(P) \geq p, \gamma_{\text{ps}}(P) \geq g\}.$$

Let $n^*(d, p, g, \varepsilon)$ be the smallest nonnegative integer n for which an estimator based on $X_0, \dots, X_n$ has relative error at most ε , with probability at least $2/3$, uniformly over $\mathcal{C}(d, p, g)$ and all deterministic initial states. Such an estimator may depend on d, p, g, ε .

Theorem 2 (Minimax bounds). *For $d \geq 2$, $0 < p \leq 1/d$, $0 < g \leq 1/8$, and $0 < \varepsilon < 1/2$, there are universal constants $c, C > 0$ such that*

$$\frac{c}{pg\varepsilon^2} \leq n^*(d, p, g, \varepsilon) \leq \frac{C}{pg\varepsilon^2} \log \frac{24d}{pg}. \quad (5)$$

One may take $c = 1/2048$. The lower bound holds even for reversible chains with a common known stationary distribution. The two hard chains have stationary minimum exactly p and pseudo-spectral gaps in $[g, 5g]$.

The restriction on g keeps the lower bound away from a degenerate endpoint. If $g \geq (1-\varepsilon)/(1+\varepsilon)$, the constant output $1-\varepsilon$ has the required relative accuracy for every $\gamma \in [g, 1]$, without any transitions. Indeed, $(1-\varepsilon)\gamma \leq 1-\varepsilon \leq (1+\varepsilon)\gamma$ throughout that interval. Thus a positive lower bound of the form in (5) cannot hold uniformly up to $g = 1$. Theorem 1 has no such restriction.

Section 2 defines the statistic and compares the gaps of thinned chains. Section 3 proves the relative Laplacian estimate and the upper bound, and Section 4 gives the lower bound. Appendix A provides a direct contraction-based alternative to the occupation-count inequality used in the main proof.

2 The estimator and its lag structure

For each $1 \leq k \leq n$, let $m_k = \lfloor n/k \rfloor$ and consider the skipped chain $Y_t = X_{tk}$, $0 \leq t \leq m_k$, whose transition matrix is P^k . Define its transition counts and degrees by

$$C_{ij}^{(k)} = \sum_{t=1}^{m_k} \mathbf{1}\{Y_{t-1} = i, Y_t = j\}, \quad r_i^{(k)} = \sum_j C_{ij}^{(k)}, \quad c_j^{(k)} = \sum_i C_{ij}^{(k)}.$$

When all incoming and outgoing degrees are positive, set

$$\hat{L}_k = \text{diag}(r^{(k)})^{-1/2} C^{(k)} \text{diag}(c^{(k)})^{-1/2}, \quad \hat{a}_k = 1 - \sigma_2(\hat{L}_k)^2. \quad (6)$$

If any degree is zero, set $\hat{a}_k = 0$. Our estimator is

$$\hat{\gamma}_n = \max_{1 \leq k \leq n} \frac{\hat{a}_k}{k} \quad (n \geq 1). \quad (7)$$

For completeness, set $\hat{\gamma}_0 = 0$.

The use of both degree vectors ensures contractivity on every sample path. For any nonnegative matrix C with positive row sums r and column sums c , let $\hat{L} = \text{diag}(r)^{-1/2} C \text{diag}(c)^{-1/2}$. Weighted Cauchy–Schwarz gives

$$\|\hat{L}u\|_2^2 = \sum_i \frac{1}{r_i} \left(\sum_j C_{ij} \frac{u_j}{\sqrt{c_j}} \right)^2 \leq \sum_{i,j} C_{ij} \frac{u_j^2}{c_j} = \|u\|_2^2.$$

Moreover, $\hat{L}\sqrt{c} = \sqrt{r}$ and $\|\sqrt{c}\|_2 = \|\sqrt{r}\|_2$, where square roots are componentwise. Hence $\sigma_1(\hat{L}) = 1$ and all its singular values lie in $[0, 1]$. In particular,

$$0 \leq \hat{a}_k \leq 1 \quad (8)$$

for every data set, including the zero-degree case.

Consequently, lag k contributes at most $1/k$ to (7). This gives an exact way to terminate the lag search: after processing lag k , stop if the largest candidate so far is at least $1/(k+1)$. Every remaining candidate is bounded by that value, so the result is unchanged. This rule concerns computation on a fixed trajectory, not the number of observations required. Our bounds concern sample complexity rather than an optimized running time.

The estimator uses the exact transformation $a = h(2-h)$ between $a = 1 - \sigma_2^2$ and $h = 1 - \sigma_2$. Maximizing h_k/k instead would estimate only a quantity within a factor of two of γ_{ps} . The normalization is also consequential: incoming and outgoing counts differ at the endpoints of a trajectory. For the trajectory $(1, 1, 2, 2)$,

$$C = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \quad r = (2, 1), \quad c = (1, 2).$$

Using r on both sides gives

$$B = \text{diag}(r)^{-1/2} C \text{diag}(r)^{-1/2} = \begin{pmatrix} 1/2 & 1/\sqrt{2} \\ 0 & 1 \end{pmatrix}.$$

For the second coordinate vector e_2 , $\|Be_2\|_2^2 = 3/2$, so B is not a contraction. Separate incoming and outgoing normalization avoids this endpoint effect.

Thinning shortens the trajectory but strengthens contraction. The next lemma quantifies this balance, using the submultiplicativity approach of [7, Section 2.1].

Lemma 1 (Gaps of powers). *For every integer $k \geq 1$,*

$$\frac{1}{4} \min\{k\gamma_{\text{ps}}(P), 1\} \leq \gamma_{\text{ps}}(P^k) \leq \min\{k\gamma_{\text{ps}}(P), 1\}. \quad (9)$$

Moreover, $a_t > 1/2$ for every integer $t \geq 1/\gamma_{\text{ps}}(P)$.

Proof. Write $\gamma = \gamma_{\text{ps}}(P)$ and $b_t = \|P^t - \Pi\|_\pi$ for $t \geq 0$. Jensen's inequality and stationarity give

$$\|Pf\|_\pi^2 \leq \sum_{i,j} \pi_i P_{ij} f_j^2 = \|f\|_\pi^2.$$

Since Π is the orthogonal projection onto the constant functions, $P^t - \Pi = P^t(I - \Pi)$ has norm at most one. Thus $0 \leq b_t \leq 1$. Also $P\Pi = \Pi P = \Pi$, so $(P^r - \Pi)(P^s - \Pi) = P^{r+s} - \Pi$, and consequently $b_{r+s} \leq b_r b_s$.

Choose a lag k_0 attaining (3) and put $a = k_0\gamma \in (0, 1]$. If $a = 1$, then $b_{k_0} = 0$, so $b_t = 0$ for every $t \geq k_0 = 1/\gamma$. Otherwise, for an integer $t \geq 1/\gamma$,

$$b_t^2 \leq (1-a)^{\lfloor t/k_0 \rfloor} \leq (1-a)^{\lfloor 1/a \rfloor} < \frac{1}{2}.$$

For the strict inequality, let $q = \lfloor 1/a \rfloor$. Then $a > 1/(q+1)$, so $(1-a)^q < (q/(q+1))^q \leq 1/2$; the last step follows by expanding $(1+1/q)^q \geq 2$. This proves the assertion about $a_t = 1 - b_t^2$.

If $k\gamma \leq 1$, take $j = \lceil 1/(k\gamma) \rceil$. Then $kj \geq 1/\gamma$, and hence

$$\gamma_{\text{ps}}(P^k) \geq \frac{a_{kj}}{j} > \frac{1}{2j} \geq \frac{k\gamma}{2(1+k\gamma)} \geq \frac{k\gamma}{4}.$$

If $k\gamma \geq 1$, the bound $a_k > 1/2$ gives the lower estimate directly. Finally, $\gamma_{\text{ps}}(P^k) = \max_{j \geq 1} a_{kj}/j \leq k\gamma$, and every pseudo-spectral gap is at most one. This proves the upper estimate. $\square$

The scalar power identity for the absolute spectral gap of a reversible chain does not extend to the pseudo-spectral gap of a general chain. For example, consider the four-state shift-register chain

$$P = \frac{1}{2} \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix}.$$

It is doubly stochastic, $P^2 = \Pi$, and its singular values are $1, 1, 0, 0$. Therefore $\gamma_{\text{ps}}(P) = 1/2$ and $\gamma_{\text{ps}}(P^2) = 1$, whereas $1 - (1 - \gamma_{\text{ps}}(P))^2 = 3/4$. The two-sided comparison in Lemma 1, rather than such an identity, is what the statistical argument needs.

3 Relative Laplacian concentration

For a transition matrix Q , the normalized bipartite Laplacian has second-smallest eigenvalue $h = 1 - \sigma_2(D_\pi^{1/2} Q D_\pi^{-1/2})$. We will estimate h to accuracy proportional to $\eta(h + \tau)$, where $\tau > 0$ is a regularization parameter. This permits a useful bound even when $h = 0$. Occupation counts control the degrees and the drift of the edge process, while a matrix martingale inequality controls the conditionally centered fluctuation.

We first record the two concentration inputs. Paulin's Bernstein inequality [4, Theorem 3.11, equation (3.25)], in the corrected version specified in the bibliography, states that for a stationary ergodic chain Q with pseudo-spectral gap $\Gamma > 0$, a nonconstant real function f , and $u > 0$,

$$\mathbb{P}_\pi \left(\left| \sum_{t=0}^{m-1} f(Y_t) - m\pi f \right| \geq u \right) \leq 2 \exp \left(- \frac{u^2 \Gamma}{8(m + 1/\Gamma) \text{Var}_\pi(f) + 20u \|f - \pi f\|_\infty} \right).$$

Here $\pi f = \sum_i \pi_i f_i$ and $\text{Var}_\pi(f) = \sum_i \pi_i (f_i - \pi f)^2$. For $N_i = \sum_{t=0}^{m-1} \mathbf{1}\{Y_t = i\}$, $m\Gamma \geq 1$, and $0 < \eta \leq 1$, this implies

$$\mathbb{P}_\pi(|N_i - m\pi_i| > \eta m \pi_i) \leq 2 \exp \left(- \frac{\eta^2 m \pi_i \Gamma}{36} \right). \quad (10)$$

Indeed, use $f = \mathbf{1}\{\cdot = i\}$ and $u = \eta m \pi_i$; then $\text{Var}_\pi(f) \leq \pi_i$, $\|f - \pi f\|_\infty \leq 1$, and the denominator in the Bernstein bound is at most $36m\pi_i$. For any initial distribution μ and any trajectory event E ,

$$\mathbb{P}_\mu(E) = \sum_i \mu_i \mathbb{P}_i(E) \leq \pi_\star^{-1} \sum_i \pi_i \mathbb{P}_i(E) = \pi_\star^{-1} \mathbb{P}_\pi(E), \quad (11)$$

since $\mu_i \leq \pi_i/\pi_\star$.

The second input is the self-adjoint matrix Freedman inequality [5, Theorem 1.2]. Let Z_t be $D \times D$ self-adjoint martingale differences with $\|Z_t\| \leq R$, and put $S_m = \sum_{t=1}^m Z_t$ and $V_m = \sum_{t=1}^m \mathbb{E}[Z_t^2 | \mathcal{F}_{t-1}]$. For $u, v > 0$,

$$\mathbb{P}\{\|S_m\| \geq u, \|V_m\| \leq v\} \leq 2D \exp \left(- \frac{u^2}{2(v + Ru/3)} \right). \quad (12)$$

This follows by applying the maximal one-sided theorem to S_t and $-S_t$, taking a union bound, and restricting to time m . The predictable quadratic variation is unchanged by the sign. We use the joint-event form (12): the bound on V_m will hold on a high-probability event, rather than on every trajectory.

Proposition 1 (Regularized relative Laplacian concentration). *Consider m transitions from an ergodic chain Q on $[d]$, from any initial distribution. Let its stationary law be π , with $p = \min_i \pi_i$, and let $\Gamma = \gamma_{\text{ps}}(Q)$. Write $L = D_\pi^{1/2} Q D_\pi^{-1/2}$ and $h = 1 - \sigma_2(L)$. Form $\hat{L}$ from the transition counts as in (6), and write $\hat{h} = 1 - \sigma_2(\hat{L})$ when all degrees are positive. If $\tau > 0$, $0 < \eta \leq 1/4$, and $m\Gamma \geq 1$, then, except on an event of probability at most*

$$\frac{4d}{p} \exp\left(-\frac{mp\Gamma\eta^2}{36}\right) + 4d \exp\left(-\frac{3mp\tau\eta^2}{16}\right), \quad (13)$$

all degrees are positive and

$$|\hat{h} - h| \leq \frac{\eta(3h + \tau)}{1 - \eta}. \quad (14)$$

Proof. Let r_i and c_i be the outgoing and incoming counts. By (10)–(11) and a union bound, the event

$$\mathcal{E} = \{(1 - \eta)\pi_i \leq r_i/m, c_i/m \leq (1 + \eta)\pi_i \text{ for every } i\} \quad (15)$$

has failure probability bounded by the first summand in (13). Both degree vectors consist of sums of m state indicators; under stationarity, the arrival sequence is stationary as well. In particular, all degrees are positive on $\mathcal{E}$.

For an edge $i \rightarrow j$, define a vector in $\mathbb{R}^{2d}$ by

$$v_{ij} = \frac{e_i}{\sqrt{\pi_i}} - \frac{e_{d+j}}{\sqrt{\pi_j}},$$

where $e_1, \dots, e_{2d}$ are the coordinate vectors. The stationary mean edge matrix is

$$H = \sum_{i,j} \pi_i Q_{ij} v_{ij} v_{ij}^T = \begin{pmatrix} I & -L \\ -L^T & I \end{pmatrix} \succeq 0. \quad (16)$$

Its eigenvalues are $1 \pm \sigma_j(L)$, so its second smallest eigenvalue, counted with multiplicity, is h . Set

$$\bar{H} = \frac{1}{m} \sum_{t=1}^m v_{Y_{t-1}, Y_t} v_{Y_{t-1}, Y_t}^T, \quad A = (H + \tau I)^{-1/2}.$$

For each state i , put

$$F_i = \sum_j Q_{ij} v_{ij} v_{ij}^T, \quad M_i = A F_i A,$$

and define

$$W_t = A v_{Y_{t-1}, Y_t} v_{Y_{t-1}, Y_t}^T A, \quad Z_t = W_t - M_{Y_{t-1}}.$$

With $\mathcal{F}_t = \sigma(Y_0, \dots, Y_t)$, the Markov property gives $\mathbb{E}[W_t | \mathcal{F}_{t-1}] = M_{Y_{t-1}}$. Thus Z_t are self-adjoint martingale differences.

Both W_t and $M_{Y_{t-1}}$ lie between 0 and RI , where $R = 2/(p\tau)$, because $\|v_{ij}\|_2^2 \leq 2/p$ and $\|A\|^2 \leq 1/\tau$. Hence $\|Z_t\| \leq R$, and

$$\mathbb{E}[Z_t^2 | \mathcal{F}_{t-1}] = \mathbb{E}[W_t^2 | \mathcal{F}_{t-1}] - M_{Y_{t-1}}^2 \preceq R M_{Y_{t-1}}.$$

On $\mathcal{E}$, positivity gives the predictable-variation bound

$$V_m \preceq R \sum_i r_i M_i \preceq Rm(1 + \eta) \sum_i \pi_i M_i = Rm(1 + \eta) A H A \preceq Rm(1 + \eta) I. \quad (17)$$

Apply (12) with $D = 2d$, $u = \eta m$, and $v = Rm(1 + \eta)$. Since $\eta \leq 1/4$,

$$\frac{u^2}{2(v + Ru/3)} = \frac{mp\tau\eta^2}{4(1 + 4\eta/3)} \geq \frac{3mp\tau\eta^2}{16}.$$

The inclusion supplied by (17) therefore yields

$$\mathbb{P}\left(\mathcal{E} \cap \left\{\left\|\sum_{t=1}^m Z_t\right\| > \eta m\right\}\right) \leq 4d \exp\left(-\frac{3mp\tau\eta^2}{16}\right). \quad (18)$$

In particular, we have not conditioned the martingale on $\mathcal{E}$.

On $\mathcal{E} \cap \{\|\sum_t Z_t\| \leq \eta m\}$, conjugating the martingale bound by A^{-1} gives

$$-\eta(H + \tau I) \preceq \bar{H} - \frac{1}{m} \sum_i r_i F_i \preceq \eta(H + \tau I).$$

Since each F_i is positive semidefinite and $\sum_i \pi_i F_i = H$, the degree event also implies

$$-\eta H \preceq \frac{1}{m} \sum_i r_i F_i - H \preceq \eta H.$$

Adding these bounds, we obtain

$$(1 - 2\eta)H - \eta\tau I \preceq \bar{H} \preceq (1 + 2\eta)H + \eta\tau I. \quad (19)$$

We next transfer this estimate to the empirical normalization. Let

$$T = \text{diag}\left(\frac{r_1}{m\pi_1}, \dots, \frac{r_d}{m\pi_d}, \frac{c_1}{m\pi_1}, \dots, \frac{c_d}{m\pi_d}\right).$$

Directly from the definitions,

$$\hat{H} = T^{-1/2} \bar{H} T^{-1/2} = \begin{pmatrix} I & -\hat{L} \\ -\hat{L}^T & I \end{pmatrix}.$$

On $\mathcal{E}$, $(1 - \eta)I \preceq T \preceq (1 + \eta)I$. As $\bar{H} \succeq 0$, every nonzero x satisfies

$$\frac{x^T \bar{H} x}{(1 + \eta)\|x\|_2^2} \leq \frac{x^T \bar{H} x}{x^T T x} \leq \frac{x^T \bar{H} x}{(1 - \eta)\|x\|_2^2}.$$

The middle quotient is the Rayleigh quotient of $\hat{H}$ under the invertible substitution $y = T^{1/2}x$. The min-max principle, with eigenvalues λ_j ordered increasingly, now gives

$$\frac{\lambda_j(\bar{H})}{1 + \eta} \leq \lambda_j(\hat{H}) \leq \frac{\lambda_j(\bar{H})}{1 - \eta}.$$

Both H and $\hat{H}$ have their second smallest eigenvalue equal to the corresponding singular-value deficit, regardless of its multiplicity. Applying this comparison at $j = 2$ and using (19) yields

$$\frac{(1 - 2\eta)h - \eta\tau}{1 + \eta} \leq \hat{h} \leq \frac{(1 + 2\eta)h + \eta\tau}{1 - \eta}.$$

Subtracting h proves (14), also when the displayed lower bound is negative. Combining the degree and martingale failure probabilities completes the proof. $\square$

The proposition applies uniformly to the lags that can attain the true gap. At these lags, both the regularization and the pseudo-spectral gap of the skipped chain are of order $k\gamma$, compensating for the reduction in sample length.

Proof of Theorem 1. Fix the chain and write $p = \pi_\star$ and $\gamma = \gamma_{\text{ps}}(P)$. Set $\eta = \varepsilon/16$ and $K = \lfloor 1/\gamma \rfloor$; these quantities are used only in the analysis. A maximizing lag k_0 in (3) satisfies $k_0\gamma = a_{k_0} \leq 1$, so $k_0 \leq K$.

For $1 \leq k \leq K$, the matrix P^k is ergodic with stationary law π . Apply Proposition 1 with $Q = P^k$, $m = m_k$, and $\tau = k\gamma$. Lemma 1 gives $\Gamma_k = \gamma_{\text{ps}}(P^k) \geq k\gamma/4$. If $n \geq 8/\gamma$, then

$$m_k \geq \frac{n}{2k}, \quad m_k \Gamma_k \geq \frac{n\gamma}{8} \geq 1, \quad m_k \tau \geq \frac{n\gamma}{2}.$$

Thus each lag has failure probability at most

$$\left(\frac{4d}{p} + 4d\right) \exp\left(-\frac{np\gamma\eta^2}{288}\right).$$

A union bound over at most $1/\gamma$ lags gives total failure probability at most

$$\frac{8d}{p\gamma} \exp\left(-\frac{np\gamma\varepsilon^2}{73728}\right). \quad (20)$$

The sample-size condition (4), with $C = 73728$, makes this at most δ and also implies $n \geq 8/\gamma$.

On the simultaneous good event, all degrees at these lags are positive. Write $h_k = 1 - \sigma_2(L_k)$ and $\hat{h}_k = 1 - \sigma_2(\hat{L}_k)$. Since $h_k \leq h_k(2 - h_k) = a_k \leq k\gamma$, (14) gives

$$|\hat{h}_k - h_k| \leq \frac{4\eta k\gamma}{1 - \eta}.$$

The function $x \mapsto x(2 - x)$ is 2-Lipschitz on $[0, 1]$, which contains both h_k and $\hat{h}_k$. Therefore

$$|\hat{a}_k - a_k| \leq \frac{8\eta k\gamma}{1 - \eta} \leq \varepsilon k\gamma, \quad 1 \leq k \leq K. \quad (21)$$

The last inequality uses $\eta = \varepsilon/16 < 1/32$. The candidate at k_0 is consequently at least $(1 - \varepsilon)\gamma$, while all candidates with $k \leq K$ are at most $(1 + \varepsilon)\gamma$. For $K < k \leq n$, the deterministic bound (8) gives $\hat{a}_k/k \leq 1/k < \gamma$. Taking the maximum proves the two-sided relative-error guarantee. $\square$

4 The minimax lower bound

The hard pair consists of a rare state coupled to a well-mixed reservoir. Varying the rate of exchange changes the gap without changing the stationary law. The transitions that distinguish nearby rates occur with frequency of order pg , which produces the three factors in the lower bound simultaneously.

Proof of Theorem 2. Fix $d \geq 2$ and $0 < p \leq 1/d$. Use the stationary law

$$\pi_1 = p, \quad \pi_i = \frac{1 - p}{d - 1} \quad (2 \leq i \leq d).$$

For $0 < s < 1$, define

$$\begin{aligned} P_s(1, 1) &= 1 - (1 - p)s, & P_s(1, j) &= \frac{(1 - p)s}{d - 1} & (j \geq 2), \\ P_s(i, 1) &= ps, & P_s(i, j) &= \frac{1 - ps}{d - 1} & (i, j \geq 2). \end{aligned}$$

All entries are positive, and every row sums to one. Across the rare-state cut, $\pi_1 P_s(1, j) = p(1 - p)s/(d - 1) = \pi_j P_s(j, 1)$; detailed balance within the reservoir follows from its equal masses and equal transition probabilities. Thus P_s is ergodic and reversible with stationary minimum exactly p .

On the two-dimensional space of functions constant on the reservoir, P_s acts as

$$\begin{pmatrix} 1 - (1 - p)s & (1 - p)s \\ ps & 1 - ps \end{pmatrix},$$

with eigenvalues 1 and $1 - s$. It annihilates functions that vanish at state 1 and sum to zero on the reservoir. Consequently its remaining $d - 2$ eigenvalues are zero, and reversibility implies

$$\gamma_{\text{ps}}(P_s) = \max_{k \geq 1} \frac{1 - (1 - s)^{2k}}{k} = 1 - (1 - s)^2 = s(2 - s). \quad (22)$$

For the maximizing claim, take $r = (1 - s)^2$ and use $1 - r^k = (1 - r) \sum_{j=0}^{k-1} r^j \leq k(1 - r)$, with equality at $k = 1$.

Given $0 < g \leq 1/8$, choose

$$s_0 = 1 - \sqrt{1 - g}, \quad s_1 = (1 + 8\varepsilon)s_0, \quad g_i = \gamma_{\text{ps}}(P_{s_i}) \quad (i = 0, 1).$$

Since $s_0 \leq g$ and $s_0 = g/(2 - s_0)$, $s_0 \leq g/(2 - g) \leq 1/15$, and hence $s_1 < 1/3$. Formula (22) gives $g_0 = g$ and $g \leq g_1 \leq (1 + 8\varepsilon)g \leq 5g$. Furthermore,

$$\frac{g_1 - g_0}{g_0} = 8\varepsilon \frac{2 - (2 + 8\varepsilon)s_0}{2 - s_0} \geq \frac{32}{5}\varepsilon > \frac{2\varepsilon}{1 - \varepsilon}.$$

Here the numerator in the fraction is at least $2 - 6/15 = 8/5$, its denominator is at most 2, and $\varepsilon < 1/2$. Thus the two intervals of successful estimation are disjoint:

$$(1 + \varepsilon)g_0 < (1 - \varepsilon)g_1. \quad (23)$$

Start both chains at the fixed reservoir state $X_0 = 2$, and denote their trajectory laws by $P_i^{(n)}$. We use KL for relative entropy with natural logarithms and TV for total variation distance. In a reservoir row, the conditional distribution of the next state within the reservoir is the same under both models. The row divergence is therefore the divergence of Bernoulli laws with parameters ps_0 and ps_1 . At the rare state, the corresponding parameters are $(1 - p)s_0$ and $(1 - p)s_1$. For $0 < u < v < 1$, the inequality $\log x \leq x - 1$ gives

$$\text{KL}(\text{Ber}(u) \parallel \text{Ber}(v)) \leq u \left(\frac{u}{v} - 1 \right) + (1 - u) \left(\frac{1 - u}{1 - v} - 1 \right) = \frac{(u - v)^2}{v(1 - v)}.$$

In our setting, $v = (1 + 8\varepsilon)u \leq 1/3$, so

$$\text{KL}(\text{Ber}(u) \parallel \text{Ber}(v)) \leq \frac{64\varepsilon^2 u}{(1 + 8\varepsilon)(1 - v)} \leq 128\varepsilon^2 u. \quad (24)$$

Under P_{s_0} , the rare-state probability $q_t = \mathbb{P}_0(X_t = 1)$ satisfies $q_0 = 0$ and $q_{t+1} = ps_0 + (1 - s_0)q_t$. Solving this recurrence yields $q_t = p\{1 - (1 - s_0)^t\} \leq p$. The chain rule for relative entropy and (24) therefore imply

$$\begin{aligned} \text{KL}(\mathbf{P}_0^{(n)} \parallel \mathbf{P}_1^{(n)}) &= \sum_{t=0}^{n-1} \mathbb{E}_0[\text{KL}(P_{s_0}(X_t, \cdot) \parallel P_{s_1}(X_t, \cdot))] \\ &\leq 128s_0\varepsilon^2 \sum_{t=0}^{n-1} (q_t(1 - p) + (1 - q_t)p) \\ &\leq 256nps_0\varepsilon^2 \leq 256npg\varepsilon^2. \end{aligned} \tag{25}$$

If $n \leq 1/(2048pg\varepsilon^2)$, this divergence is at most $1/8$. Pinsker's inequality gives $\text{TV}(\mathbf{P}_0^{(n)}, \mathbf{P}_1^{(n)}) \leq 1/4$. Thus every test of these two models has the sum of its two error probabilities at least $1 - \text{TV} \geq 3/4$.

An estimator successful with probability at least $2/3$ under both models would contradict this bound: thresholding its output at any point strictly between the two quantities in (23) would give a test with error sum at most $2/3$. Both models belong to $\mathcal{C}(d, p, g)$ and have the same stationary law, so revealing that law cannot distinguish them. This proves the lower bound with $c = 1/2048$.

For the upper bound, apply Theorem 1 with $\delta = 1/3$. On $\mathcal{C}(d, p, g)$, the product of the actual stationary minimum and gap is at least pg , and $x \mapsto x^{-1} \log(24d/x)$ is decreasing for $0 < x \leq 1$. Taking the ceiling of the resulting sample requirement, and absorbing it into the universal constant, proves the upper bound in (5). $\square$

References

- [1] D. A. Freedman. On tail probabilities for martingales. *The Annals of Probability*, 3(1):100–118, 1975. <https://doi.org/10.1214/aop/1176996452>.
- [2] D. Hsu, A. Kontorovich, D. A. Levin, Y. Peres, C. Szepesvári, and G. Wolfer. Mixing time estimation in reversible Markov chains from a single sample path. *The Annals of Applied Probability*, 29(4):2439–2480, 2019. <https://doi.org/10.1214/18-AAP1457>.
- [3] V. Karagulyan and P. Alquier. Empirical PAC-Bayes bounds for Markov chains. In *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics, Proceedings of Machine Learning Research*, 300:361–369, 2026. <https://proceedings.mlr.press/v300/karagulyan26a.html>.
- [4] D. Paulin. Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electronic Journal of Probability*, 20, paper no. 79, 1–32, 2015. <https://doi.org/10.1214/EJP.v20-4039>. The concentration inequality cited here is Theorem 3.11, equation (3.25), of the corrected version, arXiv:1212.2015v5, 2018. <https://arxiv.org/abs/1212.2015v5>.
- [5] J. A. Tropp. Freedman's inequality for matrix martingales. *Electronic Communications in Probability*, 16:262–270, 2011. <https://doi.org/10.1214/ECP.v16-1624>.
- [6] G. Wolfer and A. Kontorovich. Estimating the mixing time of ergodic Markov chains. In *Proceedings of the 32nd Conference on Learning Theory, Proceedings of Machine Learning Research*, 99:3120–3159, 2019. <https://proceedings.mlr.press/v99/wolfer19a.html>.
- [7] G. Wolfer and A. Kontorovich. Improved estimation of relaxation time in nonreversible Markov chains. *The Annals of Applied Probability*, 34(1A):249–276, 2024. <https://doi.org/10.1214/23-AAP1963>. Theorem and open-problem numbers refer to arXiv:2209.00175v3, 2023. <https://arxiv.org/abs/2209.00175v3>.

A A visit-count bound from contraction

The occupation estimate can alternatively be obtained directly from the contraction defining the pseudo-spectral gap. This argument uses thinning and scalar Freedman concentration [1]; it recovers the same polynomial sample complexity with additional logarithmic factors. It is not needed for the single-logarithm bound in Theorem 1.

Lemma 2. *Let Q be an ergodic chain on $[d]$, with stationary law π , minimum $p > 0$, and pseudo-spectral gap $\Gamma > 0$. Fix $0 < \eta \leq 1/4$, and set*

$$b = \left\lceil \frac{1 + 2 \log(4/(\eta p))}{\Gamma} \right\rceil.$$

For m transitions, let r_i and c_i be the outgoing and incoming counts. From any initial distribution, both degree vectors have relative error at most η with probability at least $1 - \delta$ whenever

$$m \geq \frac{72b}{p\eta^2} \log \frac{4db}{\delta}, \quad 0 < \delta < 1. \quad (26)$$

More precisely, if $m \geq 2b$ and $mp/(2b) \geq 2/\eta$, the failure probability is at most $4db \exp(-mp\eta^2/(72b))$.

Proof. Let k_0 attain Γ and put $a = k_0\Gamma \in (0, 1]$. The contraction and submultiplicativity used in Lemma 1 imply, for every integer $t \geq 0$,

$$\|Q^t - \Pi\|_\pi^2 \leq (1 - a)^{\lfloor t/k_0 \rfloor} \leq \exp(1 - \Gamma t). \quad (27)$$

For $0 < a < 1$, the last inequality follows from $(1 - a)^q \leq e^{-aq}$ and $a \lfloor t/k_0 \rfloor \geq \Gamma t - a \geq \Gamma t - 1$. For $a = 1$, the first bound is interpreted as one when its exponent is zero; if $t < k_0 = 1/\Gamma$, the last bound is at least one, and if $t \geq k_0$ the norm is zero. At $t = 0$, $Q^0 - \Pi = I - \Pi$ has norm one.

The similarity transform by $D_\pi^{1/2}$ gives the entrywise bound

$$\left| \frac{Q^t(i, j)}{\pi_j} - 1 \right| \leq \frac{\|Q^t - \Pi\|_\pi}{\sqrt{\pi_i \pi_j}} \leq p^{-1} \exp((1 - \Gamma t)/2).$$

By the choice of b , every $Q^b(i, j)$ lies between $(1 - \eta/4)\pi_j$ and $(1 + \eta/4)\pi_j$.

Consider first m observations $Y_0, \dots, Y_{m-1}$, and partition their indices into residue classes modulo b . Fix a state i and a residue containing ℓ observations. After setting aside its first indicator, center each subsequent indicator using its conditional mean in the filtration of this residue sequence. The Markov property identifies that mean with a transition probability of Q^b , hence places it within a relative $\eta/4$ of π_i . The centered sum S has increments bounded in absolute value by one and predictable variance at most $(\ell - 1)(1 + \eta/4)\pi_i \leq (1 + \eta/4)\ell\pi_i$. Scalar Freedman, equivalently the $D = 1$ case of (12), at $u = (\eta/4)\ell\pi_i$ gives

$$\mathbb{P}\{|S| > (\eta/4)\ell\pi_i\} \leq 2 \exp\left(-\frac{\eta^2 \ell \pi_i}{32(1 + \eta/4 + \eta/12)}\right) \leq 2 \exp(-\eta^2 \ell \pi_i/36).$$

The predictable-mean error is at most $(\eta/4)(\ell - 1)\pi_i$, and restoring the first indicator costs at most one. Outside the displayed exceptional event, the residue count therefore differs from $\ell\pi_i$ by at most $\eta\ell\pi_i$, provided $\ell\pi_i \geq 2/\eta$.

If $m \geq 2b$, each residue has $\ell \geq m/(2b)$. Under $mp/(2b) \geq 2/\eta$, summing the residue counts and taking a union bound over the b residues and d states yields

$$\mathbb{P}\{\exists i : |r_i - m\pi_i| > \eta m \pi_i\} \leq 2db \exp\left(-\frac{mp\eta^2}{72b}\right).$$

Apply the same argument to $Y_1, \dots, Y_m$ for incoming counts, and double the prefactor. No assumption on the initial distribution is used, because the first observation in each residue is controlled deterministically. Finally, (26) makes this failure probability at most δ and implies both $m \geq 2b$ and $mp/(2b) \geq 2/\eta$. $\square$

Since $b = O(\Gamma^{-1} \log(1/(p\eta)))$, the preceding lemma requires $\tilde{O}((p\Gamma\eta^2)^{-1})$ transitions. For a lag $k \leq 1/\gamma$, Lemma 1 gives $\Gamma = \gamma_{\text{ps}}(P^k) \geq k\gamma/4$, whereas $m_k \geq n/(2k)$ for $n \geq 2/\gamma$. Thus the polynomial factors of k cancel here as well. Replacing the occupation bound in Proposition 1 by Lemma 2, keeping its matrix martingale argument, and taking a union bound over these lags yields $n = \tilde{O}((p\gamma\epsilon^2)^{-1})$. Paulin's sharper inequality is what gives the single logarithm and the explicit constant in Theorem 1.