Proceedings of Machine Learning Research vol 98:1–27, 2019 30th International Conference on Algorithmic Learning Theory

Minimax Learning of Ergodic Markov Chains

Geoffrey Wolfer GEOFFREY@POST.BGU.AC.IL
Computer Science Department
Ben-Gurion University of the Negev
Beer Sheva, Israel

Aryeh Kontorovich KARYEH@CS.BGU.AC.IL
Computer Science Department
Ben-Gurion University of the Negev
Beer Sheva, Israel

Editors: Aurélien Garivier and Satyen Kale

Abstract

We compute the finite-sample minimax (modulo logarithmic factors) sample complexity of learning the parameters of a finite Markov chain from a single long sequence of states. Our error metric is a natural variant of total variation. The sample complexity necessarily depends on the spectral gap and minimal stationary probability of the unknown chain, for which there are known finite-sample estimators with fully empirical confidence intervals. To our knowledge, this is the first PAC-type result with nearly matching (up to logarithmic factors) upper and lower bounds for learning, in any metric, in the context of Markov chains.

Keywords: ergodic Markov chain, learning, minimax

1. Introduction

Approximately recovering the parameters of a discrete distribution is a classical problem in computer science and statistics (see, e.g., Han et al. (2015); Kamath et al. (2015); Orlitsky and Suresh (2015) and the references therein). Total variation (TV) is a natural and well-motivated choice of approximation metric (Devroye and Lugosi, 2001), and the metric we use throughout the paper will be derived from TV (but see Waggoner (2015) for results on other ℓ_p norms). The minimax sample complexity for obtaining an ε-approximation to the unknown distribution in TV is well-known to be $\Theta(d/\varepsilon^2)$, where d is the support size (see, e.g., Anthony and Bartlett (1999)).

This paper deals with learning the transition probability parameters of a finite Markov chain in the minimax setting. The Markov case is much less well-understood than the iid one. The main additional complexity introduced by the Markov case on top of the iid one is that not only the state space size d and the precision parameter ε, but also the chain's mixing properties must be taken into account.

Our contribution. Up to logarithmic factors, we compute (apparently the first, in *any* metric) finite sample PAC-type minimax sample complexity for the learning problem in the Markovian setting, which seeks to recover, from a single long run of an unknown Markov chain, the values of its transition matrix up to a tolerance of ε in a certain natural TV-based metric $|||\cdot|||$ we define below. We obtain upper and lower bounds on the sample complexity (sequence length) in terms of ε, the number of states, the stationary distribution, and the spectral gap of the Markov chain.

2. Main results

Our definitions and notation are mostly standard, and are given in Section 3. Since the focus of this paper is on statistical rather than computational complexity, we defer the (straightforward) analysis of the runtime of our learner to the Appendix, Section A.

2.1. Minimax learning results

Theorem 1 (Learning sample complexity upper bound) *There exists an (ε, δ)-learner $\mathcal{L}$ (provided in Algorithm 1), which, for all $0 < \varepsilon < 2$, $0 < \delta < 1$, satisfies the following. If $\mathcal{L}$ 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 $(\boldsymbol{M}, \boldsymbol{\mu})$, then it outputs $\hat{\boldsymbol{M}} = \mathcal{L}(d, \boldsymbol{X})$ such that*

$$\left|\left|\left|\boldsymbol{M} - \hat{\boldsymbol{M}}\right|\right|\right| < \varepsilon$$

holds with probability at least $1 - \delta$. The sample complexity is upper-bounded by

$$m_{\text{UB}} \quad := c \cdot \max\left\{ \frac{1}{\gamma_{\mathsf{ps}}\pi_\star} \log\left(\frac{d\sqrt{\Pi_{\boldsymbol{\mu}}}}{\delta}\right), \frac{d}{\varepsilon^2 \pi_\star} \log\left(\frac{d}{\delta}\right) \right\} = \tilde{O}\left(\max\left\{ \frac{d}{\varepsilon^2\pi_\star}, \frac{1}{\gamma_{\mathsf{ps}}\pi_\star} \right\} \right),$$

where c is a universal constant, γ_{ps} is the pseudo-spectral gap *(3.4) of $\boldsymbol{M}$, $\pi_\star$ the minimum stationary probability (3.1) of $\boldsymbol{M}$, and $\Pi_{\boldsymbol{\mu}} \leq 1/\pi_\star$ is defined in (3.2).*

The proof shows that for reversible $\boldsymbol{M}$, the bound also holds with the spectral gap (3.3) in place of the pseudo-spectral gap.

Theorem 2 (Learning sample complexity lower bound) *For every $0 < \varepsilon < 1/32$, $0 < \gamma_{\mathsf{ps}} < 1/8$, and $d = 6k$, $k \geq 2$, there exists a d-state Markov chain $\boldsymbol{M}$ with pseudo-spectral gap γ_{ps} and stationary distribution $\boldsymbol{\pi}$ such that every $(\varepsilon, 1/10)$-learner 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*

$$m_{\text{LB}} := \Omega\left(\max\left\{ \frac{d}{\varepsilon^2\pi_\star}, \frac{\log d}{\gamma_{\mathsf{ps}}\pi_\star} \right\} \right),$$

where $\gamma_{\mathsf{ps}}, \pi_\star$ are as in Theorem 1.

The proof of Theorem 2 actually yields a bit more than claimed in the statement. For any $\pi_\star \in (0, 1/d]$, a Markov chain $\boldsymbol{M}$ can be constructed that achieves the $\frac{d}{\varepsilon^2\pi_\star}$ component of the bound. Additionally, the $\frac{1}{\gamma_{\mathsf{ps}}\pi_\star}$ component is achievable by a class of *reversible* Markov chains with spectral gap $\gamma = \Theta(\gamma_{\mathsf{ps}})$, and uniform stationary distribution.

Although the sample complexity m_{UB} depends on the pseudo-spectral gap γ_{ps} and minimal stationary probability $\pi_\star$ of the unknown chain, these can be efficiently estimated with finite-sample data-dependent confidence intervals from a single trajectory both in the reversible (Hsu et al., 2017), and even in the non-reversible case (Wolfer and Kontorovich, 2019). The form of the lower bound m_{LB} indicates that in some regimes, estimating the pseudo-spectral gap up to constant multiplicative error, which requires $\tilde{\Omega}\left(\frac{1}{\gamma_{\mathsf{ps}}\pi_\star}\right)$, is as difficult as learning the entire transition matrix (for our choice of metric $|||\cdot|||$). We stress that our learner only requires ergodicity (and not, say, reversibility) to work.

Our results also indicate that the transition matrix may be estimated to precision $\varepsilon = \sqrt{\gamma_{\mathsf{ps}} d}$ with sample complexity $\tilde{O}\left(\frac{1}{\gamma_{\mathsf{ps}}\pi_\star}\right)$, which is already relevant for slowly mixing Markov chains. For this level of precision in the reversible case, in light of Hsu et al. (2017), one also obtains estimates on γ and $\pi_\star$ with no increase in sample complexity.

Finally, even though the upper bound formally depends on the unknown (and, in our setting, not learnable) initial distribution $\boldsymbol{\mu}$, we note that (i) this dependence is logarithmic and (ii) an upper bound on $\Pi_{\boldsymbol{\mu}}$ in terms of the learnable quantity $\pi_\star$ is available.

2.2. Overview of techniques

The upper bound for learning in Theorem 1 is achieved by the mildly smoothed maximum-likelihood estimator given in Algorithm 1. If the stationary distribution is bounded away from 0, the chain will visit each state a constant fraction of the total sequence length. Exponential concentration (controlled by the spectral gap) provides high-probability confidence intervals about the expectations. A technical complication is that the empirical distribution of the transitions out of a state i, conditional on the number of visits N_i to that state, is not binomial but actually rather complicated — this is due to the fact that the sequence length is fixed and so a large value of N_i "crowds out" other observations. We overcome this via a matrix version of Freedman's inequality. The factor $\Pi_{\boldsymbol{\mu}}$ in the bounds quantifies the price one pays for not assuming (as we do not) stationarity of the unknown Markov chain.

Our chief technical contribution is in establishing the sample complexity lower bounds for the Markov chain learning problem. We do this by constructing two different independent lower bounds.

The lower bound in $\tilde{\Omega}\left(\frac{1}{\gamma_{\mathsf{ps}}\pi_\star}\right)$ is derived successively by a covering argument and a classical reduction scheme to a collection of testing problems using a class of Markov chains we construct, with a carefully controlled spectral gap.[1] The latter can be estimated via Cheeger's inequality, which gives sharp upper bounds but suboptimal lower bounds (Lemma 11). To get the correct order of magnitude, we use a contraction-based argument. The Dobrushin contraction coefficient $\kappa(\boldsymbol{M})$, defined in (3.6), is in general a much cruder indicator of the mixing rate than the spectral gap γ, defined in (3.3). Indeed, $1 - \gamma \leq \kappa$ holds for all reversible $\boldsymbol{M}$ (Brémaud, 1999, pp. 237-238), and for some ergodic $\boldsymbol{M}$, we have $\kappa(\boldsymbol{M}) = 1$ (in which case it yields no information, since the latter holds for non-ergodic $\boldsymbol{M}$ as well). This is in fact the case for the families of Markov chains we construct in the course of proving Theorem 2. Fortunately, in both cases, even though $\kappa(\boldsymbol{M}) = 1$, it turns out that $\kappa(\boldsymbol{M}^2) < 1$, and coupled with the contraction property (3.7), our bound on $\kappa(\boldsymbol{M}^2)$ actually yields an optimal estimate of γ_{ps}. Although the calculation of $\kappa(\boldsymbol{M}^2)$ in Lemma 9 is computationally intensive, the contraction coefficient is, in general, more amenable to analysis than the eigenvalues directly, and hence this technique may be of independent interest.

The lower bound in $\Omega\left(\frac{d}{\varepsilon^2\pi_\star}\right)$ arises from the idea that learning the whole transition is at least as hard as learning the conditional distribution of one of its states. From here, we design a class of matrices where one state is both hard to reach and difficult to learn, by constructing mixture of

1. The family of chains used in the lower bound of Hsu et al. (2017) does not suffice for our purposes; a considerably richer family is needed (see Remark 5).

indistinguishable distributions for that particular state, and indexed by a large subset of the binary hypercube. We express the statistical distance between words of length m distributed according to different matrices of this class in terms of $\pi_\star$ and the KL divergence between the conditional distributions of the hard-to-reach state, by taking advantage of the structure of the class, and invoke an argument from Tsybakov to conclude ours.

2.3. Related work

Our Markov chain learning setup is a natural extension of the PAC distribution learning model of Kearns et al. (1994). Despite the plethora of literature on estimating Markov transition matrices (see, e.g., Billingsley (1961); Craig and Sendi (2002); Welton and Ades (2005)) we were not able to locate any rigorous finite-sample PAC-type results.

The minimax problem has recently received some attention, and Hao et al. (2018) have, in parallel to us, shown the first minimax learning bounds, in expectation, for the problem of learning the transition matrix $\boldsymbol{M}$ of a Markov chain under a certain class of divergences. The authors consider the case where $\min_{i,j} \boldsymbol{M}(i,j) \geq \alpha > 0$, essentially showing that for some family of smooth f-divergences, the expected risk is $\Theta\left(\frac{df''(1)}{m\pi_\star}\right)$. The metric used in this paper is based on TV, which corresponds to the f-divergence induced by $f(t) = \frac{1}{2}\left|t-1\right|$, which is not differentiable at $t = 1$. The results of Hao et al. and the present paper are complementary and not directly comparable. We do note that (i) their guarantees are in expectation rather than with high-confidence, (ii) our TV-based metric is not covered by their smooth f-divergence family, and most important (iii) their notion of mixing is related to contraction as opposed to the spectral gap. In particular the α-minorization assumption implies (but is not implied by) a bound of $\kappa \leq 1 - d\alpha$ on the Dobrushin contraction coefficient (defined in (3.6); see Kontorovich (2007, Lemma 2.2.2) for the latter claim). Thus, the family of α-minorized Markov chains is strictly contained in the family of contracting chains, which in turn is a strict subset of the ergodic chains we consider.

3. Definitions and notation

We define $[d] := \{1, \ldots, d\}$ and use m to denote the size of the sample received by the Markov learner. The simplex of all distributions over $[d]$ will be denoted by Δ_d, and the collection of all $d \times d$ row-stochastic matrices by $\mathcal{M}_d$. 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. We assume familiarity with basic Markov chain concepts (see, e.g., Kemeny and Snell (1976); Levin et al. (2009)). A Markov chain $(\boldsymbol{M}, \boldsymbol{\mu})$ on d states is specified by an initial distribution $\boldsymbol{\mu} \in \Delta_d$ and a transition matrix $\boldsymbol{M} \in \mathcal{M}_d$ in the usual way. Namely, by $(X_1, \ldots, X_m) \sim (\boldsymbol{M}, \boldsymbol{\mu})$, we mean that

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

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.

The Markov chain $(\boldsymbol{M}, \boldsymbol{\mu})$ is *stationary* if $\boldsymbol{\mu} = \boldsymbol{\pi}$ for $\boldsymbol{\pi} = \boldsymbol{\pi}\boldsymbol{M}$, and *ergodic* if $\boldsymbol{M}^k > 0$ entrywise for some $k \geq 1$. If $\boldsymbol{M}$ is ergodic, it has a unique stationary distribution $\boldsymbol{\pi}$ and moreover

$\pi_\star > 0$, where

$$\pi_\star = \min_{i\in[d]} \boldsymbol{\pi}(i). \tag{3.1}$$

Unless noted otherwise, $\boldsymbol{\pi}$ is assumed to be the stationary distribution of the Markov chain in context. To any Markov chain $(\boldsymbol{M}, \boldsymbol{\mu})$, we associate

$$\Pi_{\boldsymbol{\mu}} := \sum_{i\in[d]} \boldsymbol{\mu}(i)^2/\boldsymbol{\pi}(i), \tag{3.2}$$

which is always $\Pi_{\boldsymbol{\mu}} \le 1/\pi_\star$.

A *reversible* $\boldsymbol{M} \in \mathcal{M}_d$ satisfies *detailed balance* for some distribution $\boldsymbol{\mu}$: for all $i, j \in [d]$, $\boldsymbol{\mu}(i)\boldsymbol{M}(i,j) = \boldsymbol{\mu}(j)\boldsymbol{M}(j,i)$ — 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): $1 = \lambda_1 \ge \lambda_2 \ge \ldots \ge \lambda_d$. The *spectral gap* is

$$\gamma = \gamma(\boldsymbol{M}) = 1 - \lambda_2(\boldsymbol{M}). \tag{3.3}$$

Paulin (2015) defines the *pseudo-spectral gap* by

$$\gamma_{\mathsf{ps}} := \max_{k\ge 1} \left\{ \frac{\gamma((\boldsymbol{M}^\star)^k \boldsymbol{M}^k)}{k} \right\}, \tag{3.4}$$

where $\boldsymbol{M}^\star$ is the *time reversal* of $\boldsymbol{M}$, given by $\boldsymbol{M}^\star(i,j) := \boldsymbol{\pi}(j)\boldsymbol{M}(j,i)/\boldsymbol{\pi}(i)$; the expression $\boldsymbol{M}^\star\boldsymbol{M}$ is called the *multiplicative reversiblization* of $\boldsymbol{M}$.

We use the standard ℓ_1 norm $\|z\| = \sum_{i\in[d]} |z_i|$, which, in the context of distributions (and up to a convention-dependent factor of 2) corresponds to the total variation norm. For $A \in \mathbb{R}^{d\times d}$, define

$$|||A||| := \max_{i\in[d]} \|A(i,\cdot)\|_1 = \max_{i\in[d]} \sum_{j\in[d]} |A(i,j)| \tag{3.5}$$

(we note, but do not further exploit, that $|||\cdot|||$ corresponds to the $\ell_\infty \to \ell_\infty$ operator norm (Horn and Johnson, 1985)). For any $\boldsymbol{M} \in \mathcal{M}_d$, define its *Dobrushin contraction* coefficient

$$\kappa(\boldsymbol{M}) = \frac{1}{2} \max_{i,j\in[d]} \|\boldsymbol{M}(i,\cdot) - \boldsymbol{M}(j,\cdot)\|_1 ; \tag{3.6}$$

this quantity is also associated with Döblin's name. The term "contraction" refers to the property

$$\left\|(\boldsymbol{\mu} - \boldsymbol{\mu}')\boldsymbol{M}\right\|_1 \le \kappa(\boldsymbol{M}) \left\|\boldsymbol{\mu} - \boldsymbol{\mu}'\right\|_1, \qquad \boldsymbol{\mu}, \boldsymbol{\mu}' \in \Delta_d, \tag{3.7}$$

which was observed by Markov (1906, §5).

Finally, we use standard $O(\cdot)$, $\Omega(\cdot)$ and $\Theta(\cdot)$ order-of-magnitude notation, as well as their tilde variants $\tilde{O}(\cdot)$, $\tilde{\Omega}(\cdot)$, $\tilde{\Theta}(\cdot)$ where lower-order log factors are suppressed.

Definition 3 *An* (ε, δ)-learner $\mathcal{L}$ *for Markov chains with sample complexity function* $m_0(\cdot)$ *is an algorithm that takes as input* $\boldsymbol{X} = (X_1, \ldots, X_m)$ *drawn from some unknown Markov chain* $(\boldsymbol{M}, \boldsymbol{\mu})$*, and outputs* $\hat{\boldsymbol{M}} = \mathcal{L}(d, \boldsymbol{X})$ *such that* $m \ge m_0(\varepsilon, \delta, \boldsymbol{M}, \boldsymbol{\mu}) \Rightarrow \left|\left|\left|\boldsymbol{M} - \hat{\boldsymbol{M}}\right|\right|\right| < \varepsilon$ *holds with probability at least* $1 - \delta$.

The probability is over the draw of $\boldsymbol{X}$ and any internal randomness of the learner. Note that by Theorem 2, the learner's sample complexity must necessarily depend on the properties of the unknown Markov chain.

4. Proofs

Proof [of Theorem 1] We proceed to analyze Theorem 1, and in particular, the random variable $\hat{\boldsymbol{M}}(i,j) = \dfrac{N_{ij}}{N_i}$ it constructs, where

$$N_i = \left|\{t \in [m-1] : X_t = i\}\right|, N_{ij} = \left|\{t \in [m-1] : X_t = i, X_{t+1} = j\}\right|.$$

To do so, we make use of an adaptation of Freedman's inequality (Freedman, 1975) to random matrices (Tropp et al., 2011), which has been reported for convenience in the appendix as Theorem 12. Define the row vector sequence $\boldsymbol{Y}$ for a fixed i by

$$\boldsymbol{Y}_0 = 0, \boldsymbol{Y}_t = \frac{1}{\sqrt{2}}\Big(\mathbf{1}\left[X_{t-1} = i\right]\left(\mathbf{1}\left[X_t = j\right] - \boldsymbol{M}(i,j)\right)\Big)_{j \in [d]},$$

and notice that $\sum_{t=1}^{m} \boldsymbol{Y}_t = \frac{1}{\sqrt{2}}\left(N_{i1} - N_i \boldsymbol{M}(i,1), \dots, N_{id} - N_i \boldsymbol{M}(i,d)\right)$. We also have from the Markov property that $\underset{\boldsymbol{M},\boldsymbol{\mu}}{\mathbf{E}}\left[\boldsymbol{Y}_t \,|\, \boldsymbol{Y}_{t-1}\right] = \mathbf{0}$, so that $\boldsymbol{Y}_t$ defines a vector-valued martingale difference, and immediately,

$$\begin{aligned}
\boldsymbol{Y}_t \boldsymbol{Y}_t^\intercal = \|\boldsymbol{Y}_t\|_2^2 &= \sum_{j=1}^{d}\left(\frac{1}{\sqrt{2}}\mathbf{1}\left[X_{t-1} = i\right]\left(\mathbf{1}\left[X_t = j\right] - \boldsymbol{M}(i,j)\right)\right)^2 \\
&= \frac{1}{2}\mathbf{1}\left[X_{t-1} = i\right]\sum_{j=1}^{d}\mathbf{1}\left[X_t = j\right] + \boldsymbol{M}(i,j)^2 - 2 \cdot \mathbf{1}\left[X_t = j\right]\boldsymbol{M}(i,j) \\
&= \frac{1}{2}\mathbf{1}\left[X_{t-1} = i\right]\Big(1 + \|\boldsymbol{M}(i,\cdot)\|_2^2 - 2\boldsymbol{M}(i,X_t)\Big) \leq \mathbf{1}\left[X_{t-1} = i\right],
\end{aligned} \tag{4.1}$$

so that $\boldsymbol{W}_{col,m} := \sum_{t=1}^{m} \mathbf{E}\left[\boldsymbol{Y}_t \boldsymbol{Y}_t^\intercal \,|\, \mathcal{F}_{t-1}\right] \leq \sum_{t=1}^{m} \mathbf{1}\left[X_{t-1} = i\right] = N_i$, and $\|\boldsymbol{W}_{col,m}\|_2 \leq N_i$ as $\boldsymbol{W}_{col,m}$ is a real valued random variable. Construct now the $d \times d$ matrix $\boldsymbol{Y}_t^\intercal \boldsymbol{Y}_t$,

$$\boldsymbol{Y}_t^\intercal \boldsymbol{Y}_t = \frac{\mathbf{1}\left[X_{t-1} = i\right]}{2}\begin{pmatrix}
Z_{t,i,1,1} & Z_{t,i,1,2} & \cdots & Z_{t,i,1,d} \\
Z_{t,i,2,1} & Z_{t,i,2,2} & \cdots & Z_{t,i,2,d} \\
\vdots & \vdots & \vdots & \vdots \\
\vdots & Z_{t,i,j,k} & \vdots & \vdots \\
\vdots & \vdots & \vdots & \vdots \\
Z_{t,i,d,1} & Z_{t,i,d,2} & \cdots & Z_{t,i,d,d}
\end{pmatrix}, \tag{4.2}$$

with $Z_{t,i,j,k} = (\mathbf{1}\left[X_t = j\right] - \boldsymbol{M}(i,j))(\mathbf{1}\left[X_t = k\right] - \boldsymbol{M}(i,k))$. Computing the row sums and column sums of this matrix in absolute value,

$$\sum_{k=1}^{d} |Z_{t,i,j,k}| = |\mathbf{1}\left[X_t = j\right] - \boldsymbol{M}(i,j)| \sum_{k=1}^{d} |\mathbf{1}\left[X_t = k\right] - \boldsymbol{M}(i,k)| \leq \sum_{k=1}^{d} \mathbf{1}\left[X_t = k\right] + \sum_{k=1}^{d} \boldsymbol{M}(i,k) = 2 \tag{4.3}$$

and similarly, $\sum_{j=1}^{d} |Z_{t,i,j,k}| \leq 2$. From Hölder's inequality, $\|\boldsymbol{Y}_t^\intercal \boldsymbol{Y}_t\|_2 \leq \sqrt{\|\boldsymbol{Y}_t^\intercal \boldsymbol{Y}_t\|_1 \cdot \|\boldsymbol{Y}_t^\intercal \boldsymbol{Y}_t\|_\infty} \leq \mathbf{1}\left[X_{t-1} = i\right]$, and from the sub-additivity of the norm and Jensen's inequality $\|\boldsymbol{W}_{row,m}\|_2 := \left\|\sum_{t=1}^{m} \mathbf{E}\left[\boldsymbol{Y}_t^\intercal \boldsymbol{Y}_t \,|\, \mathcal{F}_{t-1}\right]\right\|_2 \leq \sum_{t=1}^{m} \mathbf{E}\left[\|\boldsymbol{Y}_t^\intercal \boldsymbol{Y}_t\|_2 \,|\, \mathcal{F}_{t-1}\right]$, it follows that

$$\|\boldsymbol{W}_{row,m}\|_2 \leq \sum_{t=1}^{m} \mathbf{E}\left[\sqrt{\frac{\mathbf{1}\left[X_{t-1} = i\right] \cdot 2}{2} \cdot \frac{\mathbf{1}\left[X_{t-1} = i\right] \cdot 2}{2}} \,|\, \mathcal{F}_{t-1}\right] \leq N_i.$$

Now decomposing the error probability of the learner, while choosing an arbitrary value $n_i \in \mathbb{N}$ for the desired number of visits to each state,

$$\begin{aligned}\mathop{\mathbf{P}}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left|\left|\left|\boldsymbol{M} - \hat{\boldsymbol{M}}\right|\right|\right| > \varepsilon\right) \leq \sum_{i=1}^{d} \mathop{\mathbf{P}}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\hat{\boldsymbol{M}}(i,\cdot) - \boldsymbol{M}(i,\cdot)\right\|_1 > \varepsilon \text{ and } N_i \in [n_i, 3n_i]\right) \\ + \mathop{\mathbf{P}}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\{\exists i \in [d] : N_i \notin [n_i, 3n_i]\}\right).\end{aligned} \tag{4.4}$$

Since $\left\|\hat{\boldsymbol{M}}(i,\cdot) - \boldsymbol{M}(i,\cdot)\right\|_1 > \varepsilon \implies \left\|\hat{\boldsymbol{M}}(i,\cdot) - \boldsymbol{M}(i,\cdot)\right\|_2 > \frac{\varepsilon}{\sqrt{d}}$, we have

$$\begin{aligned}&\mathop{\mathbf{P}}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\hat{\boldsymbol{M}}(i,\cdot) - \boldsymbol{M}(i,\cdot)\right\|_1 > \varepsilon \text{ and } N_i \in [n_i, 3n_i]\right) \\ &\leq \mathop{\mathbf{P}}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\hat{\boldsymbol{M}}(i,\cdot) - \boldsymbol{M}(i,\cdot)\right\|_2 > \frac{\varepsilon}{\sqrt{d}} \text{ and } N_i \in [n_i, 3n_i]\right) \\ &= \mathop{\mathbf{P}}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\frac{\sqrt{2}\,\|\sum_{t=1}^{m} \boldsymbol{Y}_t\|_2}{N_i} > \frac{\varepsilon}{\sqrt{d}} \text{ and } N_i \in [n_i, 3n_i]\right) \\ &\leq \mathop{\mathbf{P}}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\sum_{t=1}^{m} \boldsymbol{Y}_t\right\|_2 > \frac{\varepsilon}{\sqrt{2d}} n_i \text{ and } N_i \leq 3n_i\right) \\ &\leq \mathop{\mathbf{P}}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\sum_{t=1}^{m} \boldsymbol{Y}_t\right\|_2 > \frac{\varepsilon}{\sqrt{2d}} n_i \text{ and } \max\left\{\|\boldsymbol{W}_{row,m}\|_2, \|\boldsymbol{W}_{col,m}\|_2\right\} \leq 3n_i\right) \\ &\leq (d+1) \cdot \exp\left(-\frac{\varepsilon^2 n_i^2}{2d(3n_i + \varepsilon n_i/(3\sqrt{2d}))}\right) \text{ (Theorem 12)} \\ &\leq 2d \cdot \exp\left(-\frac{\varepsilon^2 n_i}{8d}\right),\end{aligned} \tag{4.5}$$

and setting $n_i = \frac{m\pi_i}{2}$, it follows that for all $i \in [d]$, $\mathop{\mathbf{P}}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left\|\hat{\boldsymbol{M}}(i,\cdot) - \boldsymbol{M}(i,\cdot)\right\|_1 > \varepsilon \text{ and } N_i \in [n_i, 3n_i]\right) \leq 2d \cdot \exp\left(-\frac{\varepsilon^2 m\pi_\star}{16d}\right)$, and finally, from Lemma 6 (stated and proven in the appendix), which follows easily from Paulin (2015), it is possible to control the number of visits to states, such that for m larger than $\frac{16d}{\varepsilon^2\pi_\star} \ln\left(\frac{4d^2}{\delta}\right)$, and $m \geq \frac{112}{\gamma_{\mathsf{ps}}\pi_\star} \log\left(\frac{2d\sqrt{\Pi_{\boldsymbol{\mu}}}}{\delta}\right)$ we have that $\mathop{\mathbf{P}}_{\boldsymbol{M},\boldsymbol{\mu}}\left(\left|\left|\left|\boldsymbol{M} - \hat{\boldsymbol{M}}\right|\right|\right| > \varepsilon\right) \leq \delta$, and the upper bound is proven.

■

Remark 4 *Note that one can derive an upper bound* $\tilde{O}\left(\frac{\max\left\{1/\varepsilon^2, 1/\gamma_{\mathsf{ps}}\right\}}{\pi_\star}\right)$ *for the problem with respect to the max norm* $\left\|\boldsymbol{M} - \hat{\boldsymbol{M}}\right\|_{\mathsf{MAX}} = \max_{(i,j)\in[d]^2}\left|\boldsymbol{M}(i,j) - \hat{\boldsymbol{M}}(i,j)\right|$*, by studying the entry-wise martingales and invoking the scalar version of Freedman's inequality (Freedman, 1975). Similarly, since for* $p \in [1,2)$*, it is the case that* $\{\|x\|_2 > \varepsilon\} \implies \left\{\|x\|_p > d^{1/p-1/2}\cdot\varepsilon\right\}$*, we can derive the more general upper bound* $\tilde{O}\left(\frac{\max\left\{d^{2/p-1}/\varepsilon^2, 1/\gamma_{\mathsf{ps}}\right\}}{\pi_\star}\right)$ *for the problem with respect to the norm* $\left\|\left\|\boldsymbol{M} - \hat{\boldsymbol{M}}\right\|\right\|_p = \max_{i\in[d]}\left\|\boldsymbol{M}(i,\cdot) - \hat{\boldsymbol{M}}(i,\cdot)\right\|_p$*.*

Proof [of Theorem 2 (part 1): learning lower bound $\Omega\left(\frac{d}{\varepsilon^2\pi_\star}\right)$]

Let $0 < \varepsilon < 1/32$, and $\mathcal{M}_{d,\gamma_{\mathsf{ps}},\pi_\star}$ be 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}. The quantity we wish to lower bound is the minimax risk for the learning problem :

$$\mathcal{R}_m = \inf_{\hat{\boldsymbol{M}}}\sup_{\boldsymbol{M}} \mathbf{P}_{\boldsymbol{M}}\left(\left\|\left\|\boldsymbol{M} - \hat{\boldsymbol{M}}\right\|\right\| > \varepsilon\right), \tag{4.6}$$

where the inf is taken over all learners and the sup over $\mathcal{M}_{d,\gamma_{\mathsf{ps}},\pi_\star}$. Suppose for simplicity of the analysis that we consider Markov chains of $d+1$ states instead of d, and that d is even. A slight modification of the proofs covers the odd case. We define the following class of Markov chains parametrized by a given distribution $\mathbf{p} \in \Delta_{d+1}$, where the conditional distribution defined at each state of the chain is always $\mathbf{p}$ with $p_{d+1} = p_\star$ and $p_k = \frac{1-p_\star}{d}$ for $k \in [d]$, with $p_\star < \frac{1}{d+2}$, except for state $d+1$, where it is only required that it has a loop of probability $p_\star$ to itself.

$$\mathcal{G}_{\mathbf{p}} = \left\{\boldsymbol{M}_{\boldsymbol{\eta}} = \begin{pmatrix} p_1 & \dots & p_d & p_\star \\ \vdots & \vdots & \vdots & \vdots \\ p_1 & \dots & p_d & p_\star \\ \eta_1 & \dots & \eta_d & p_\star \end{pmatrix} : \boldsymbol{\eta} = (\eta_1,\dots,\eta_d,p_\star) \in \Delta_{d+1}\right\}. \tag{4.7}$$

Remark: a family of Markov chains very similar to $\mathcal{G}_{\mathbf{p}}$ was independently considered by Hao et al. (2018) for proving their lower bound.
It is easy to see that the stationary distribution $\boldsymbol{\pi}$ of an element of $\mathcal{G}_{\mathbf{p}}$ indexed by $\boldsymbol{\eta}$ is

$$\pi_k = \frac{(1-p_\star)^2}{d} + \eta_k p_\star, \text{ for } k \in [d], \qquad \pi_{d+1} = p_\star \tag{4.8}$$

For $m \geq 4$, $\boldsymbol{\eta} = (\eta_1,\dots,\eta_d,p_\star) \in \Delta_{d+1}$ and $(X_1,\dots,X_m) \sim (\boldsymbol{M}_{\boldsymbol{\eta}}, \mathbf{p})$, set $N_i = |\{t \in [m] : X_t = i\}|$ the number of visits to the ith state. Focusing on the $(d+1)$th state, since $\forall i \in [d+1], \boldsymbol{M}_{\boldsymbol{\eta}}(i, d+1) = p_\star$, it is immediate that $N_{d+1} \sim \mathrm{Binomial}(m, p_\star)$. Introduce the subset of Markov chains in $\mathcal{G}_{\mathbf{p}}$ such that

$$\boldsymbol{\eta}(\boldsymbol{\sigma}) = \left(\frac{1-p_\star+16\sigma_1\varepsilon}{d}, \frac{1-p_\star-16\sigma_1\varepsilon}{d}, \dots, \frac{1-p_\star+16\sigma_{\frac{d}{2}}\varepsilon}{d}, \frac{1-p_\star-16\sigma_{\frac{d}{2}}\varepsilon}{d}, p_\star\right),$$

where $\boldsymbol{\sigma} = \left(\sigma_1,\dots,\sigma_{\frac{d}{2}}\right) \in \{-1,1\}^{\frac{d}{2}}$. Also define $\boldsymbol{M}_0$ with $\boldsymbol{\eta}_0 = \left(\frac{1-p_\star}{d},\dots,\frac{1-p_\star}{d}, p_\star\right)$. A direct computation yields that for $\boldsymbol{\sigma} \neq \boldsymbol{\sigma}'$, $\|\boldsymbol{M}_{\boldsymbol{\sigma}} - \boldsymbol{M}_{\boldsymbol{\sigma}'}\|_1 = \frac{32\varepsilon}{d} d_H(\boldsymbol{\sigma},\boldsymbol{\sigma}')$, where d_H is the Hamming distance. From the Varshamov-Gilbert lemma, we know that $\exists \Sigma \subset \{-1, 1\}^{d/2}$, $|\Sigma| \geq 2^{d/16}$, such that $\forall (\boldsymbol{\sigma}, \boldsymbol{\sigma}') \in \Sigma$ with $\boldsymbol{\sigma} \neq \boldsymbol{\sigma}'$, $d_H(\boldsymbol{\sigma}, \boldsymbol{\sigma}') \geq \frac{d}{16}$. Restricting our problem to this set Σ, and finally noticing that $\forall \boldsymbol{\sigma} \in \Sigma, \|\boldsymbol{M}_{\boldsymbol{\sigma}} - \boldsymbol{M}_0\|_1 = 16\varepsilon > 2\varepsilon$, from Tsybakov's method (Tsybakov, 2009) applied to our problem,

$$\mathcal{R}_m \geq \frac{1}{2}\left(1 - \frac{\frac{4}{2^{\frac{d}{16}}} \sum_{\boldsymbol{\sigma} \in \Sigma} D_{KL}\left(\boldsymbol{M}_{\boldsymbol{\sigma}}^m || \boldsymbol{M}_0^m\right)}{\log 2^{\frac{d}{16}}}\right), \tag{4.9}$$

where we wrote $D_{KL}\left(\boldsymbol{M}_{\boldsymbol{\sigma}}^m || \boldsymbol{M}_0^m\right)$ to be the KL divergence between the two distributions of words of length m from each of the Markov chains. Leveraging a tensorization property of the KL divergence, and as by construction, the only discrepancy occurs when visiting the $(d+1)$th state, Lemma 7 shows that

$$D_{KL}\left(\boldsymbol{M}_{\boldsymbol{\sigma}}^m || \boldsymbol{M}_0^m\right) \leq p_\star m D_{KL}\left(\boldsymbol{\eta}(\boldsymbol{\sigma}) || \boldsymbol{\eta}_0\right), \tag{4.10}$$

following up with a straightforward computation,

$$\begin{aligned} D_{KL}\left(\boldsymbol{\eta}(\boldsymbol{\sigma}) || \boldsymbol{\eta}_0\right) &= \frac{d}{2}\left(\frac{1 - p_\star + 16\varepsilon}{d}\right) \ln\left(\frac{\frac{1-p_\star+16\varepsilon}{d}}{\frac{1-p_\star}{d}}\right) + \frac{d}{2}\left(\frac{1 - p_\star - 16\varepsilon}{d}\right) \ln\left(\frac{\frac{1-p_\star-16\varepsilon}{d}}{\frac{1-p_\star}{d}}\right) \\ &\leq 128\varepsilon^2, \end{aligned} \tag{4.11}$$

and finally combining (4.6), (4.10) and (4.11), we get $\mathcal{R}_m \geq \frac{1}{2}\left(1 - \frac{512\varepsilon^2 m p_\star}{\frac{d}{16}\ln 2}\right)$. Further noticing that for the considered range of ε and for $p_\star < \frac{1}{d+2}$, it is always the case that $\pi_\star = p_\star$, so that for $m \leq \frac{d(1-2\delta)\ln 2}{8192\varepsilon^2 \pi_\star}$, $\mathcal{R}_m \geq \delta$. ■

Proof [of Theorem 2 (part 2): learning lower bound $\Omega\left(\frac{1}{\gamma_{\mathsf{ps}}\pi_\star}\right)$]

We treat $0 < \varepsilon \leq 1/8$ and $d = 6k$, $k \geq 2$ as fixed. For $\eta \in (0, 1/48)$ and $\boldsymbol{\tau} \in \{0,1\}^{d/3}$, define the block matrix

$$\boldsymbol{M}_{\eta,\boldsymbol{\tau}} = \begin{pmatrix} C_\eta & R_{\boldsymbol{\tau}} \\ R_{\boldsymbol{\tau}}^\intercal & L_{\boldsymbol{\tau}} \end{pmatrix},$$

where $C_\eta \in \mathbb{R}^{d/3 \times d/3}$, $L_{\boldsymbol{\tau}} \in \mathbb{R}^{2d/3 \times 2d/3}$, and $R_{\boldsymbol{\tau}} \in \mathbb{R}^{d/3 \times 2d/3}$ are given by

$$L_{\boldsymbol{\tau}} = \frac{1}{8}\operatorname{diag}\left(7 - 4\tau_1\varepsilon, 7 + 4\tau_1\varepsilon, \dots, 7 - 4\tau_{d/3}\varepsilon, 7 + 4\tau_{d/3}\varepsilon\right),$$

$$C_\eta = \begin{pmatrix} \frac{3}{4} - \eta & \frac{\eta}{d/3-1} & \cdots & \frac{\eta}{d/3-1} \\ \frac{\eta}{d/3-1} & \frac{3}{4} - \eta & \ddots & \vdots \\ \vdots & \ddots & \ddots & \frac{\eta}{d/3-1} \\ \frac{\eta}{d/3-1} & \cdots & \frac{\eta}{d/3-1} & \frac{3}{4} - \eta \end{pmatrix},$$

$$R_{\boldsymbol{\tau}} = \frac{1}{8}\begin{pmatrix} 1+4\tau_1\varepsilon & 1-4\tau_1\varepsilon & 0 & \dots & \dots & \dots & 0 \\ 0 & 0 & 1+4\tau_2\varepsilon & 1-4\tau_2\varepsilon & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \dots & \dots & \dots & 0 & 1+4\tau_{d/3}\varepsilon & 1-4\tau_{d/3}\varepsilon \end{pmatrix}.$$

Holding η fixed, define the collection

$$\mathcal{H}_\eta = \left\{ \boldsymbol{M}_{\eta,\boldsymbol{\tau}} : \boldsymbol{\tau} \in \{0,1\}^{d/3} \right\} \tag{4.12}$$

of Markov matrices. Denote by $\boldsymbol{M}_{\eta,\mathbf{0}} \in \mathcal{H}_\eta$ the element corresponding to $\boldsymbol{\tau} = \mathbf{0}$. Note that every $\boldsymbol{M} \in \mathcal{H}_\eta$ is ergodic and reversible, and its unique stationary distribution is uniform.

A graphical illustration[2] of this class of Markov chains is provided in Figure 1; in particular, every $\boldsymbol{M} \in \mathcal{H}_\eta$ consists of an "inner clique" (i.e., the states indexed by $\{1, \dots, d/3\}$) and "outer rim" (i.e., the states indexed by $\{d/3+1, \dots, d\}$).

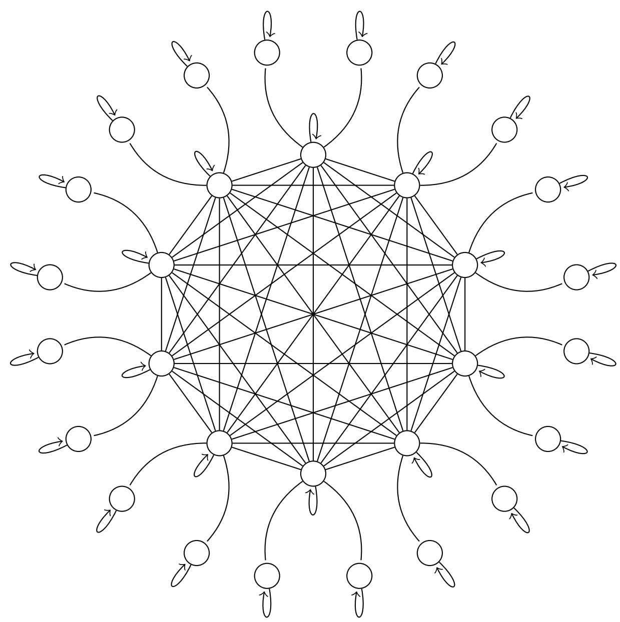

Figure 1: Generic topology of the $\mathcal{H}_\eta$ Markov chain class: every chain consists of an "inner clique" and an "outer rim".

Lemma 8 in the Appendix establishes a key property of the elements of $\mathcal{H}_\eta$: each $\boldsymbol{M}$ in this class satisfies

$$\gamma_{\mathsf{ps}}(\boldsymbol{M}) = \Theta(\eta). \tag{4.13}$$

Suppose that $\boldsymbol{X} = (X_1, \dots, X_m) \sim (\boldsymbol{M}_{\boldsymbol{\eta}}, \boldsymbol{\pi})$, where $\boldsymbol{M} \in \mathcal{H}_\eta$ and $\boldsymbol{\pi}$ is uniform. Define the random variable $T_{\text{CLIQ}} = T_{\text{CLIQ}}(\boldsymbol{M})$, to be the first time all of the states in the inner clique were

2. Additional figures are provided in Section C in the Appendix.

visited,

$$T_{\text{CLIQ}} = \inf \left\{ t \geq 1 : \left| \{X_1, \dots, X_t\} \cap [d/3] \right| = d/3 \right\}, \tag{4.14}$$

Lemma 10 in the Appendix gives a lower estimate on this quantity:

$$m \leq \frac{d}{20\eta} \ln \left(\frac{d}{3} \right) \quad \Longrightarrow \quad \mathbf{P}\left(T_{\text{CLIQ}} > m \right) \geq \tfrac{1}{5}. \tag{4.15}$$

Let $\mathcal{M}_{d,\gamma_{\mathsf{ps}},\pi_\star}$ be 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}. Writing $\boldsymbol{X} = (X_1, \dots, X_m)$, recall that the quantity we wish to lower bound is the minimax risk for the learning problem (it will be convenient to write $\varepsilon/2$ instead of ε, which only affects the constants):

$$\mathcal{R}_m = \inf_{\hat{\boldsymbol{M}}} \sup_{\boldsymbol{M}} \mathbf{P}_{\boldsymbol{M}} \left(\left|\left|\left| \boldsymbol{M} - \hat{\boldsymbol{M}} \right|\right|\right| > \frac{\varepsilon}{2} \right), \tag{4.16}$$

where the $\inf$ is taken over all learners and the $\sup$ over $\mathcal{M}_{d,\gamma_{\mathsf{ps}},\pi_\star}$. We employ the general reduction scheme of Tsybakov (2009, Chapter 2.2). The first step is to restrict the $\sup$ to the finite subset $\mathcal{H}_\eta \subsetneq \mathcal{M}_{d,\gamma_{\mathsf{ps}},\pi_\star}$.

$$\mathcal{R}_m \geq \inf_{\hat{\boldsymbol{M}}} \sup_{\boldsymbol{\tau}} \mathbf{P}_{\boldsymbol{M}_{\eta,\boldsymbol{\tau}}} \left(\left|\left|\left| \boldsymbol{M}_{\eta,\boldsymbol{\tau}} - \hat{\boldsymbol{M}} \right|\right|\right| > \frac{\varepsilon}{2} \right). \tag{4.17}$$

Define T_{CLIQ} as in (4.14). Then

$$\mathcal{R}_m \geq \inf_{\hat{\boldsymbol{M}}} \sup_{\boldsymbol{\tau}} \mathbf{P}_{\boldsymbol{M}_{\eta,\boldsymbol{\tau}}} \left(\left|\left|\left| \boldsymbol{M}_{\eta,\boldsymbol{\tau}} - \hat{\boldsymbol{M}} \right|\right|\right| > \varepsilon \,\middle|\, T_{\text{CLIQ}} > m \right) \mathbf{P}_{\boldsymbol{M}_{\eta,\boldsymbol{\tau}}} \left(T_{\text{CLIQ}} > m \right) \tag{4.18}$$

and Lemma 10 implies that for $m < \frac{d}{20\eta} \ln\left(\frac{d}{3}\right)$,

$$\mathcal{R}_m \geq \frac{1}{5} \inf_{\hat{\boldsymbol{M}}} \sup_{\boldsymbol{\tau}} \mathbf{P}_{\boldsymbol{M}_{\eta,\boldsymbol{\tau}}} \left(\left|\left|\left| \boldsymbol{M}_{\eta,\boldsymbol{\tau}} - \hat{\boldsymbol{M}} \right|\right|\right| > \varepsilon \,\middle|\, T_{\text{CLIQ}} > m \right). \tag{4.19}$$

Observe that all $\boldsymbol{\tau} \neq \boldsymbol{\tau}' \in \{0,1\}^{d/3}$ verify $|||\boldsymbol{M}_{\eta,\boldsymbol{\tau}} - \boldsymbol{M}_{\eta,\boldsymbol{\tau}'}||| = \varepsilon$. For any estimate $\hat{\boldsymbol{M}} = \hat{\boldsymbol{M}}(\boldsymbol{X})$, define

$$\boldsymbol{\tau}^\star(\boldsymbol{X}) = \operatorname*{argmin}_{\boldsymbol{\tau}} \left|\left|\left| \hat{\boldsymbol{M}} - \boldsymbol{M}_{\eta,\boldsymbol{\tau}} \right|\right|\right|.$$

Then for $\boldsymbol{\tau} \neq \boldsymbol{\tau}^\star(\boldsymbol{X})$, we have

$$\varepsilon = |||\boldsymbol{M}_{\eta,\boldsymbol{\tau}} - \boldsymbol{M}_{\eta,\boldsymbol{\tau}^\star}||| \leq \left|\left|\left| \boldsymbol{M}_{\eta,\boldsymbol{\tau}} - \hat{\boldsymbol{M}} \right|\right|\right| + \left|\left|\left| \hat{\boldsymbol{M}} - \boldsymbol{M}_{\eta,\boldsymbol{\tau}^\star} \right|\right|\right| \leq 2 \left|\left|\left| \boldsymbol{M}_{\eta,\boldsymbol{\tau}} - \hat{\boldsymbol{M}} \right|\right|\right|, \tag{4.20}$$

whence $\{\boldsymbol{\tau}^\star \neq \boldsymbol{\tau}\} \subset \left\{ \left|\left|\left| \boldsymbol{M}_{\eta,\boldsymbol{\tau}} - \hat{\boldsymbol{M}} \right|\right|\right| > \varepsilon/2 \right\}$ and

$$\mathcal{R}_m \geq \frac{1}{5} \inf_{\hat{\boldsymbol{M}}} \sup_{\boldsymbol{\tau}} \mathbf{P}_{\boldsymbol{M}_{\eta,\boldsymbol{\tau}}} \left(\boldsymbol{\tau}^\star \neq \boldsymbol{\tau} \,|\, T_{\text{CLIQ}} > m \right) = \frac{1}{5} \inf_{\hat{\boldsymbol{\tau}} : \boldsymbol{X} \mapsto \{0,1\}^{d/3}} \sup_{\boldsymbol{\tau}} \mathbf{P}_{\boldsymbol{M}_{\eta,\boldsymbol{\tau}}} \left(\hat{\boldsymbol{\tau}} \neq \boldsymbol{\tau} \,|\, T_{\text{CLIQ}} > m \right). \tag{4.21}$$

Since $T_{\text{CLIQ}} > m$ implies that $N_{i^\star} = 0$ for some $i^\star \in [d/3]$,

$$\mathcal{R}_m \geq \frac{1}{5} \inf_{\hat{\boldsymbol{\tau}}} \sup_{\boldsymbol{\tau}} \mathbf{P}_{\boldsymbol{M}_{\eta,\boldsymbol{\tau}}} \left(\hat{\tau}_{i^\star} \neq \tau_{i^\star} \,|\, N_{i^\star} = 0 \right). \tag{4.22}$$

There are as many $\boldsymbol{M} \in \mathcal{H}_\eta$ with $\tau_{i^\star} = 0$ as those with $\tau_{i^\star} = 1$, so if $\boldsymbol{M}$ is drawn uniformly at random and state $i^\star$ has not been visited, the learner can do no better than to make a random choice of $\hat{\tau}_{i^\star}$ (where $\hat{\boldsymbol{\tau}}$ determines $\hat{\boldsymbol{M}}$). More formally, writing $\boldsymbol{\tau}^{(i)} = (\tau_1, \dots, \tau_{i-1}, \tau_{i+1}, \dots, \tau_{d/3}) \in \{0,1\}^{d/3-1}$, the $\boldsymbol{\tau}$ vector without its ith coordinate, we can employ an Assouad-type of decomposition (Assouad, 1983; Yu, 1997):

$$\begin{aligned}
\mathcal{R}_m &\geq \frac{1}{5} \inf_{\hat{\boldsymbol{\tau}}} 2^{1-d/3} \sum_{\boldsymbol{\tau}^{(i)} \in \{0,1\}^{d/3-1}} \left[\frac{1}{2} \mathbf{P}_{\tau_i=0} \left(\hat{\tau}_i \neq \tau_i \,|\, N_i = 0 \right) + \frac{1}{2} \mathbf{P}_{\tau_i=1} \left(\hat{\tau}_i \neq \tau_i \,|\, N_i = 0 \right) \right] \\
&= \frac{2^{1-d/3}}{10} \sum_{\boldsymbol{\tau}^{(i)} \in \{0,1\}^{d/3-1}} \inf_{\hat{\boldsymbol{\tau}}} \left[\mathbf{P}_{\tau_i=0} \left(\hat{\tau}_i = 1 \,|\, N_i = 0 \right) + \mathbf{P}_{\tau_i=1} \left(\hat{\tau}_i = 0 \,|\, N_i = 0 \right) \right] \\
&= \frac{2^{1-d/3}}{10} \sum_{\boldsymbol{\tau}^{(i)} \in \{0,1\}^{d/3-1}} \left[1 - \left\| \mathbf{P}_{\tau_i=0} \left(\boldsymbol{X} = \cdot \,|\, N_i = 0 \right) + \mathbf{P}_{\tau_i=1} \left(\boldsymbol{X} = \cdot \,|\, N_i = 0 \right) \right\|_1 \right] \\
&= \frac{1}{10}.
\end{aligned} \tag{4.23}$$

Combined with Lemma 8, and inclusion of events, this implies lower bounds of $\Omega\left(\frac{d}{\gamma} \ln d\right)$ and $\Omega\left(\frac{d}{\gamma_{\mathsf{ps}}} \ln d\right)$ for the learning problem, which are tight for the case $\pi_\star = \frac{1}{d}$. ■

Remark 5 *Let us compare construction $\mathcal{H}_\eta$ to the family of Markov chains employed in the lower bound of Hsu et al. (2017):*

$$\boldsymbol{M}(i,j) = \begin{cases} 1 - \eta_i, & i = j \\ \frac{\eta_i}{d-1}, & \textit{else} \end{cases}, \tag{4.24}$$

where $\eta_i \in \{\eta, \eta'\}$ with $\eta' \approx \eta/2$. For our lower bound, $\mathcal{H}'_\eta$ has to be a ε-separated set under $|||\cdot|||$. In the construction of Hsu et al., the spectral gap γ and the separation distance ε are coupled, and using their family of Markov chains would lead to a lower bound of order $d/\gamma \approx d/\varepsilon$, which is inferior to $\Omega\left(\frac{\sqrt{d}}{\varepsilon^2 \pi_\star}\right)$. The free parameter η was key to our construction, which enabled us to decouple γ from ε.

Acknowledgments

This research was supported by the Lynn and William Frankel Center for Computer Science at Ben-Gurion University and by the Israel Science Foundation. We are thankful to John Lafferty for bringing this problem to our attention and numerous insightful conversations. We also thank the anonymous reviewers for their valuable comments.

References

Martin Anthony and Peter L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, Cambridge, 1999. ISBN 0-521-57353-X. doi: 10.1017/CBO9780511624216. URL http://dx.doi.org/10.1017/CBO9780511624216.

Patrice Assouad. Deux remarques sur l'estimation. *Comptes rendus des séances de l'Académie des sciences. Série 1, Mathématique*, 296(23):1021–1024, 1983.

Patrick Billingsley. Statistical methods in markov chains. *The Annals of Mathematical Statistics*, 32(1):12–40, 1961. ISSN 00034851. URL http://www.jstor.org/stable/2237603.

Pierre Brémaud. *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer-Verlag, 1999.

Bruce A. Craig and Peter P. Sendi. Estimation of the transition matrix of a discrete-time Markov chain. *Health economics*, 11(1):33–42, 2002.

Luc Devroye and Gábor Lugosi. *Combinatorial methods in density estimation*. Springer Series in Statistics. Springer-Verlag, New York, 2001. ISBN 0-387-95117-2. doi: 10.1007/978-1-4613-0125-7. URL http://dx.doi.org/10.1007/978-1-4613-0125-7.

David A Freedman. On tail probabilities for martingales. *the Annals of Probability*, pages 100–118, 1975.

Yanjun Han, Jiantao Jiao, and Tsachy Weissman. Minimax estimation of discrete distributions under ℓ_1 loss. *IEEE Transactions on Information Theory*, 61(11):6343–6354, Nov 2015. ISSN 0018-9448. doi: 10.1109/TIT.2015.2478816.

Yi Hao, Alon Orlitsky, and Venkatadheeraj Pichapati. On learning markov chains. In *Advances in Neural Information Processing Systems*, pages 646–655, 2018.

Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.

Daniel J. Hsu, Aryeh Kontorovich, David A. Levin, Yuval Peres, and Csaba 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.

Sudeep Kamath, Alon Orlitsky, Dheeraj Pichapati, and Ananda Theertha Suresh. On learning distributions from their samples. In *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, pages 1066–1100, 2015. URL http://jmlr.org/proceedings/papers/v40/Kamath15.html.

Michael J. Kearns, Yishay Mansour, Dana Ron, Ronitt Rubinfeld, Robert E. Schapire, and Linda 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.

John G. Kemeny and J. Laurie Snell. *Finite Markov chains*. Springer-Verlag, New York, 1976. Reprinting of the 1960 original, Undergraduate Texts in Mathematics.

Aryeh (Leonid) Kontorovich. *Measure Concentration of Strongly Mixing Processes with Applications*. PhD thesis, Carnegie Mellon University, 2007.

David Asher Levin, Yuval Peres, and Elizabeth Lee Wilmer. *Markov chains and mixing times, second edition*. American Mathematical Soc., 2009.

Andrei A. Markov. Extension of the law of large numbers to dependent quantities. *Izvestiia Fiz.-Matem. Obsch. Kazan Univ.*, 15:135–156, 1906.

Alon Orlitsky and Ananda Theertha Suresh. Competitive distribution estimation: Why is good-turing good. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 2143–2151, 2015. URL http://papers.nips.cc/paper/5762-competitive-distribution-estimation-why-is-good-turing-good.

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

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

Alexandre 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.

Bo 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.

Nicky Welton and A. E. Ades. Estimation of Markov chain transition probabilities and rates from fully and partially observed data: uncertainty propagation, evidence synthesis, and model calibration. *Medical Decision Making*, 25(6):633–645, 2005.

Geoffrey Wolfer and Aryeh Kontorovich. Estimating the mixing time of ergodic markov chains. arXiv preprint arXiv:1902.01224, 2019. URL https://arxiv.org/abs/1902.01224.

Bin Yu. Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam*, pages 423–435. Springer, 1997.

Appendix A. Algorithm

Algorithm 1 The learner $\mathcal{L}$

Input: $d, (X_1, \ldots, X_m)$
Output: $\hat{\boldsymbol{M}}$
$\hat{\boldsymbol{M}} \leftarrow \mathbf{0} \in \mathbb{R}^{d \times d}$
$\mathsf{Visits} \leftarrow \mathbf{0} \in \mathbb{R}^{d}$
for $t \leftarrow 1$ **to** m **do**
  $\mathsf{Visits}(X_t) \leftarrow \mathsf{Visits}(X_t) + 1$
end
for $i \leftarrow 1$ **to** d **do**
  if $\mathsf{Visits}(i) > 0$ **then**
    for $t \leftarrow 1$ **to** $m - 1$ **do**
      $\hat{\boldsymbol{M}}(i,j) \leftarrow \hat{\boldsymbol{M}}(i,j) + \mathbf{1}\left[X_t = i\right] \cdot \mathbf{1}\left[X_{t+1} = j\right] / \mathsf{Visits}(i)$
    end
  else
    $\hat{\boldsymbol{M}}(i,\cdot) = (1/d, 1/d, \ldots, 1/d)$
  end
end
return $\hat{\boldsymbol{M}}$

Constructing $\hat{\boldsymbol{M}}$ has time complexity $O(d(d+m))$.

Appendix B. Auxiliary lemmas and reported theorems for the proofs of Theorem 1 and Theorem 2

Lemma 6 *Let $(\boldsymbol{M}, \boldsymbol{\mu})$ an ergodic d-state Markov chain together with its initial distribution, with stationary distribution $\boldsymbol{\pi}$, pseudo spectral gap γ_{ps} and minimum stationary probability $\pi_\star$. For $m \geq \frac{112}{\gamma_{\mathsf{ps}} \pi_\star} \log \left(\frac{2d\sqrt{\Pi_{\boldsymbol{\mu}}}}{\delta} \right)$, $\underset{\boldsymbol{M},\boldsymbol{\mu}}{\mathbf{P}} \left(\left\{ \exists i \in [d] : N_i \notin \left[\frac{1}{2} m\pi_i, \frac{3}{2} m\pi_i \right] \right\} \right) \leq \frac{\delta}{2}$, where N_i is the number of visits to state i.*

Proof

Invoking Paulin (2015, Proposition 3.14, Theorem 3.8, Theorem 3.10), for any walk of length m,

$$\mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}} \left(|N_i - m\pi_i| > \frac{1}{2} m\pi_i \right) \leq \Pi_{\boldsymbol{\mu}}^{1/2} \cdot \mathbf{P}_{\boldsymbol{M},\boldsymbol{\pi}} \left(|N_i - m\pi_i| > \frac{1}{2} m\pi_i \right)^{1/2}. \tag{B.1}$$

Then for reversible $\boldsymbol{M}$ with spectral gap γ,

$$\mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}} \left(N_i \notin \left[\frac{1}{2} m\pi_i, \frac{3}{2} m\pi_i \right] \right) \leq \sqrt{\Pi_{\boldsymbol{\mu}}} \exp \left(- \frac{\gamma \left(\frac{1}{2} m\pi_i \right)^2}{2(4m\pi_i(1-\pi_i) + 10\frac{1}{2} m\pi_i)} \right) \tag{B.2}$$

and for general $\boldsymbol{M}$ with pseudo-spectral gap γ_{ps},

$$\mathbf{P}_{\boldsymbol{M},\boldsymbol{\mu}}\left(N_i \notin \left[\frac{1}{2}m\pi_i, \frac{3}{2}m\pi_i\right]\right) \leq \sqrt{\Pi_{\boldsymbol{\mu}}}\exp\left(-\frac{\gamma_{\mathsf{ps}}\left(\frac{1}{2}m\pi_i\right)^2}{2(8(m+1/\gamma_{\mathsf{ps}})\pi_i(1-\pi_i)+20\frac{1}{2}m\pi_i)}\right). \tag{B.3}$$

Hence, from a direct computation and an application of the union bound, as long as $m \geq \frac{112}{\pi_\star \gamma_{\mathsf{ps}}} \ln\left(\frac{2d\sqrt{\Pi_{\boldsymbol{\mu}}}}{\delta}\right)$, $\underset{\boldsymbol{M},\boldsymbol{\mu}}{\mathbf{P}}\left(N_i \notin [\frac{1}{2}m\pi_i, \frac{3}{2}m\pi_i]\right) \leq \frac{\delta}{2d}$, whence the lemma in the non-reversible case. The similar result with the spectral gap in lieu of the pseudo-spectral gap for reversible chains can be proven the exact same way from Theorem B.2. ■

Lemma 7 *For two Markov chains $\boldsymbol{M}_1$ and $\boldsymbol{M}_2$ of the class $\mathcal{G}_{\mathbf{p}}$ defined at (4.7) indexed respectively by $\boldsymbol{\eta}_1$ and $\boldsymbol{\eta}_2$, it is the case that*

$$D_{KL}\left(\boldsymbol{M}_1^m || \boldsymbol{M}_2^m\right) \leq p_\star m D_{KL}\left(\boldsymbol{\eta}_1 || \boldsymbol{\eta}_2\right), \tag{B.4}$$

Proof *Recall that from the tensorization property of the KL divergence,*

$$D_{KL}\left(\boldsymbol{M}_1^m || \boldsymbol{M}_2^m\right) = \sum_{t=1}^m \underset{X_1,\dots X_{t-1}}{\mathbf{E}}\left[D_{KL}\left(X_t \sim \boldsymbol{M}_{1\,|\,X_1,\dots,X_{t-1}} \middle|\middle| X_t \sim \boldsymbol{M}_{2\,|\,X_1,\dots,X_{t-1}}\right)\right], \tag{B.5}$$

so that successively,

$$\begin{aligned}
&D_{KL}\left(\boldsymbol{M}_1^m || \boldsymbol{M}_2^m\right) \\
&= \sum_{t=1}^m \underset{X_1,\dots X_{t-1}}{\mathbf{E}}\left[D_{KL}\left(X_t \sim \boldsymbol{M}_{1\,|\,X_{t-1}} \middle|\middle| X_t \sim \boldsymbol{M}_{2\,|\,X_{t-1}}\right)\right] \\
&= \sum_{t=1}^m \underset{X_1,\dots X_{t-2}}{\mathbf{E}}\left[\underset{X_{t-1}}{\mathbf{E}}\left[D_{KL}\left(X_t \sim \boldsymbol{M}_{1\,|\,X_{t-1}} \middle|\middle| X_t \sim \boldsymbol{M}_{2\,|\,X_{t-1}}\right) \mid X_1 = x_1, \dots, X_{t-2} = x_{t-2}\right]\right]
\end{aligned} \tag{B.6}$$

$$\begin{aligned}
&= \sum_{t=1}^m \underset{X_1,\dots X_{t-2}}{\mathbf{E}}\left[\sum_{x_{t-1}\in[d+1]}\left(D_{KL}\left(X_t \sim \boldsymbol{M}_{1\,|\,x_{t-1}} \middle|\middle| X_t \sim \boldsymbol{M}_{2\,|\,x_{t-1}}\right)\right) \underset{\boldsymbol{M}_1}{\mathbf{P}}\left(X_{t-1} = x_{t-1} \,|\, X_{t-2} = x_{t-2}\right)\right] \\
&= p_\star \sum_{t=1}^m \underset{X_1,\dots X_{t-2}}{\mathbf{E}}\left[D_{KL}\left(X_t \sim \boldsymbol{M}_{1\,|\,X_{t-1}=d+1} \middle|\middle| X_t \sim \boldsymbol{M}_{2\,|\,X_{t-1}=d+1}\right)\right] \\
&= p_\star \sum_{t=1}^m D_{KL}\left(\boldsymbol{M}_1(d+1,\cdot) || \boldsymbol{M}_2(d+1,\cdot)\right) \\
&= p_\star m D_{KL}\left(\boldsymbol{M}_1(d+1,\cdot) || \boldsymbol{M}_2(d+1,\cdot)\right)
\end{aligned} \tag{B.7}$$

and, $D_{KL}\left(\boldsymbol{M}_1^m || \boldsymbol{M}_2^m\right) \leq p_\star m D_{KL}\left(\boldsymbol{\eta}_1 || \boldsymbol{\eta}_2\right)$. ■

Lemma 8 *For all $\boldsymbol{M} \in \mathcal{H}_\eta$ [defined in (4.12)], we have $\gamma = \Theta(\eta)$ and $\gamma_{\mathsf{ps}} = \Theta(\eta)$.*

Remark. For the lower bound in Theorem 2, we only need $\gamma, \gamma_{\mathsf{ps}} \in \Omega(\eta)$. The other estimate follows from the Cheeger constant computation in Lemma 11, which has been deferred to Section D.

Proof We focus our proof on the spectral gap, and will later show that the pseudo spectral gap is of the same order for our class of Markov matrices. We begin by computing the Cheeger constant (see, e.g., Levin et al. (2009, Chapter 7)) of the chains in $\mathcal{H}_\eta$. For $\boldsymbol{M} \in \mathcal{H}_\eta$, denote its Cheeger constant by

$$\Phi = \min_{S \subset [d], \boldsymbol{\pi}(S) \le \frac{1}{2}} \frac{\sum_{i \in S, j \in [d] \setminus S} \boldsymbol{\pi}(i) \boldsymbol{M}(i,j)}{\boldsymbol{\pi}(S)}, \tag{B.8}$$

and recall that for a lazy reversible Markov chain, Cheeger's inequality states that the spectral gap γ satisfies $\frac{\Phi^2}{2} \le \gamma \le 2\Phi$. From this inequality and Lemma 11, we have that $\gamma \le 6\eta$. It remains to prove the corresponding linear lower bound. From Levin et al. (2009, Theorem 12.5), we have

$$t_{\mathsf{mix}} \ge \ln 2 (1/\gamma - 1),$$

where

$$t_{\mathsf{mix}} = t_{\mathsf{mix}}(\boldsymbol{M}) = \min \left\{ t \in \mathbb{N} : \sup_{\boldsymbol{\mu} \in \Delta_d} \frac{1}{2} \left\| \boldsymbol{\mu} \boldsymbol{M}^t - \boldsymbol{\pi} \right\|_1 \le 1/4 \right\}.$$

From (3.7), we have that the Dobrushin coefficient $\kappa(\cdot)$ satisfies

$$\begin{aligned} \left\| \boldsymbol{\mu} \boldsymbol{M}^t - \boldsymbol{\pi} \right\|_1 = \left\| \boldsymbol{\mu} \boldsymbol{M}^t - \boldsymbol{\pi} \boldsymbol{M}^t \right\|_1 &= \left\| \boldsymbol{M}^t (\boldsymbol{\mu} - \boldsymbol{\pi}) \right\|_1 \\ &\le \kappa(\boldsymbol{M}^t) \left\| \boldsymbol{\mu} - \boldsymbol{\pi} \right\|_1 \\ &\le \kappa(\boldsymbol{M}^2)^{\lfloor t/2 \rfloor} \left\| \boldsymbol{\mu} - \boldsymbol{\pi} \right\|_1 \le \kappa(\boldsymbol{M}^2)^{\lfloor t/2 \rfloor}. \end{aligned} \tag{B.9}$$

Since $\kappa(\boldsymbol{M}^2) \le 1 - \frac{\eta}{16}$ from Lemma 9, it follows that $t_{\mathsf{mix}} \le \frac{64 \ln 2}{\eta}$, and hence

$$\gamma \ge \left(1 + \frac{t_{\mathsf{mix}}}{\ln 2}\right)^{-1} \ge \left(1 + \frac{64}{\eta}\right)^{-1} \ge \frac{\eta}{64}, \tag{B.10}$$

whence $\gamma = \Theta(\eta)$.

Now note that for a symmetric $\boldsymbol{M}$, $\boldsymbol{\pi}$ is the uniform distribution, $\boldsymbol{M}^\star = \boldsymbol{M}^\intercal = \boldsymbol{M}$, and $\gamma_{\mathsf{ps}} = \max_{k \ge 1} \left\{ \frac{\gamma(\boldsymbol{M}^{2k})}{k} \right\}$. Denoting by $1 = \lambda_1 > \lambda_2 \ge \dots \ge \lambda_d$ the eigenvalues of $\boldsymbol{M}$, we have that for all $i \in [d]$ and $k \ge 1$, λ_i^{2k} is an eigenvalue for $\boldsymbol{M}^{2k}$, and furthermore $1 = \lambda_1^{2k} > \lambda_2^{2k} \ge \dots \ge \lambda_d^{2k}$. We claim that

$$\gamma_{\mathsf{ps}} = \max_{k \ge 1} \frac{1 - \lambda_2^{2k}}{k} = 1 - \lambda_2^2$$

— that is, the maximum is achieved at $k = 1$. Indeed, $1 - \lambda_2^{2k} = (1 - \lambda_2^2)\left(\sum_{i=0}^{k-1} \lambda_2^i\right)$ and the latter sum is at most k since $\lambda_2 < 1$. As a result, $\gamma_{\mathsf{ps}}(\boldsymbol{M}) = 1 - \lambda_2^2 = 1 - (1 - \gamma(\boldsymbol{M}))^2 = \gamma(\boldsymbol{M})[2 - \gamma(\boldsymbol{M})]$ and

$$\gamma(\boldsymbol{M}) \leq \gamma_{\mathsf{ps}}(\boldsymbol{M}) \leq 2\gamma(\boldsymbol{M}), \tag{B.11}$$

which completes the proof. ■

Lemma 9 (Bounding the Dobrushin contraction coefficient) *For all $\boldsymbol{M} \in \mathcal{H}_\eta$,*

$$\kappa(\boldsymbol{M}^2) \leq 1 - \eta\left(\frac{1}{8} - \frac{\varepsilon}{2}\right)\left(1 + \frac{1}{d/3 - 1}\right) \leq 1 - \frac{\eta}{16}. \tag{B.12}$$

Proof

Special case: $\boldsymbol{M} = \boldsymbol{M}_{\eta,\mathbf{0}}$. Such an $\boldsymbol{M} \in \mathcal{H}_\eta$ corresponds to the case $\boldsymbol{\tau} = \mathbf{0}$, and in this case we claim that

$$\kappa(\boldsymbol{M}^2) = 1 - \frac{\eta}{8}\left(1 + \frac{1}{d/3 - 1}\right) \leq 1 - \frac{\eta}{8}. \tag{B.13}$$

We begin by computing $\boldsymbol{M}^2$:

$$\boldsymbol{M}^2 = \begin{pmatrix} C_{\eta,\mathbf{0}}^{(2)} & R_{\eta,\mathbf{0}}^{(2)} \\ R_{\eta,\mathbf{0}}^{(2)\intercal} & L_{\mathbf{0}}^{(2)} \end{pmatrix}, \tag{B.14}$$

where $C_{\eta,\mathbf{0}}^{(2)} \in \mathbb{R}^{d/3 \times d/3}$, $L_{\mathbf{0}}^{(2)} \in \mathbb{R}^{2d/3 \times 2d/3}$, and $R_{\eta,\mathbf{0}}^{(2)} \in \mathbb{R}^{d/3 \times 2d/3}$ are given by

$$C_{\eta,\mathbf{0}}^{(2)} = \begin{pmatrix} \alpha_1 & \alpha_6 & \dots & \alpha_6 \\ \alpha_6 & \alpha_1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \alpha_6 \\ \alpha_6 & \dots & \alpha_6 & \alpha_1 \end{pmatrix}, \tag{B.15}$$

$$R_{\eta,\mathbf{0}}^{(2)} = \begin{pmatrix} \alpha_3 & \alpha_3 & \alpha_4 & \dots & \dots & \dots & \alpha_4 \\ \alpha_4 & \alpha_4 & \alpha_3 & \alpha_3 & \alpha_4 & \dots & \alpha_4 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \alpha_4 & \dots & \dots & \dots & \alpha_4 & \alpha_3 & \alpha_3 \end{pmatrix}, \tag{B.16}$$

$$L_{\mathbf{0}}^{(2)} = \begin{pmatrix} S_{\mathbf{0}}^{(2)} & & & \mathbf{0} \\ & S_{\mathbf{0}}^{(2)} & & \\ & & \ddots & \\ \mathbf{0} & & & S_{\mathbf{0}}^{(2)} \end{pmatrix}, \qquad S_{\mathbf{0}}^{(2)} = \begin{pmatrix} \alpha_2 & \alpha_5 \\ \alpha_5 & \alpha_2 \end{pmatrix}, \tag{B.17}$$

and

$$\begin{aligned}
\alpha_1 &= (3/4-\eta)^2 + \frac{\eta^2}{d/3-1} + 1/32 \\
\alpha_2 &= 25/32 \\
\alpha_3 &= (1/8)(3/4-\eta) + 7/64 \\
\alpha_4 &= (1/8)\frac{\eta}{d/3-1} \\
\alpha_5 &= 1/64 \\
\alpha_6 &= 2(3/4-\eta)\frac{\eta}{d/3-1} + (d/3-2)\frac{\eta^2}{(d/3-1)^2}.
\end{aligned} \tag{B.18}$$

We observe that $\kappa(\boldsymbol{M}^2) = \max_{i\in[5]} \kappa_i$, where

$$\begin{aligned}
\kappa_1 &= \frac{1}{2}\left(2|\alpha_3-\alpha_4| + 2\alpha_2 + 2\alpha_5\right) = |\alpha_3-\alpha_4| + \alpha_2 + \alpha_5 \\
\kappa_2 &= \frac{1}{2}\left(2|\alpha_5-\alpha_2|\right) = |\alpha_5-\alpha_2| \\
\kappa_3 &= \frac{1}{2}\left((d/3-1)|\alpha_6-\alpha_4| + |\alpha_3-\alpha_1| + (2d/3-2)\alpha_4 + |\alpha_3-\alpha_2| + |\alpha_3-\alpha_5|\right) \\
\kappa_4 &= \frac{1}{2}\left((d/3-2)|\alpha_4-\alpha_6| + |\alpha_4-\alpha_1| + |\alpha_3-\alpha_6| + (2d/3-4)\alpha_4 + 2\alpha_3 + |\alpha_4-\alpha_5| + |\alpha_4-\alpha_2|\right) \\
\kappa_5 &= \frac{1}{2}\left(2|\alpha_6-\alpha_1| + 4|\alpha_4-\alpha_3|\right) = |\alpha_6-\alpha_1| + 2|\alpha_4-\alpha_3|.
\end{aligned} \tag{B.19}$$

We proceed to compute $\max \kappa_i$.

κ_1**:** Since $d = 6k$, $k \geq 2$ and $\eta < \frac{1}{48} < \frac{13}{16}$, we have

$$\begin{aligned}
\kappa_1 &= |\alpha_3-\alpha_4| + \alpha_2 + \alpha_5 \\
&= \frac{1}{8}\left|\frac{13}{8} - \eta\left(1 + \frac{1}{d/3-1}\right)\right| + \frac{51}{64} \\
&= 1 - \frac{\eta}{8}\left(1 + \frac{1}{d/3-1}\right).
\end{aligned} \tag{B.20}$$

$\kappa_2 < \kappa_1$**:** Since $\kappa_2 = |\alpha_5 - \alpha_2| = \frac{49}{64}$, and as $\eta < \frac{1}{48}$, it follows that $\kappa_2 < \kappa_1$.

$\kappa_5 < \kappa_1$**:**

$$\begin{aligned}
|\alpha_6-\alpha_1| &= \left|2(3/4-\eta)\frac{\eta}{d/3-1} + (d/3-2)\frac{\eta^2}{(d/3-1)^2} - (3/4-\eta)^2 + \frac{\eta^2}{d/3-1} + 1/32\right| \\
&= \left|\frac{\eta(3(d-3)d - 2((d-6)d+18)\eta)}{2(d-3)^2} - \frac{17}{32}\right|.
\end{aligned} \tag{B.21}$$

Write $g(\eta, d) = \dfrac{\eta(3(d-3)d - 2((d-6)d+18)\eta)}{2(d-3)^2}$, and notice that for $d \geq 6$,

$$g(\eta, d) \leq \frac{\eta(3(d-3)d - 2((d-6)d+18)\eta)}{18}.$$

Now

$$\frac{\mathrm{d}}{\mathrm{d}\eta}\eta(3(d-3)d - 2((d-6)d+18)\eta) = 3(d-3)d - 4((d-6)d+18)\eta,$$

which is strictly positive as long as $\eta < \frac{3(d-3)d}{4((d-6)d+18)}$. But since $\eta < 1/48$, this holds for all $d \geq 6$, and thus $g(\cdot, d)$ is strictly increasing on $[0, 1/48]$. Further,

$$g(1/48, d) = \frac{71d^2 - 210d - 18}{2304(d-3)^2},$$

which is decreasing for $d \geq 6$, is majorized by $g(1/48, 6) = \frac{71}{1152} < \frac{17}{32}$, and so

$$|\alpha_6 - \alpha_1| = \frac{17}{32} - g(\eta, d). \tag{B.22}$$

Since $g(\eta, d) > 0$ for $\eta < \frac{1}{48} < \frac{3}{2}\left(\frac{(d-3)d}{(d-6)d+18}\right)$, it follows that $\alpha_2 + \alpha_5 > |\alpha_6 - \alpha_1| + |\alpha_4 - \alpha_3|$, which shows that. that $\kappa_5 < \kappa_1$.

$\kappa_3 < \kappa_1$: Let us compute

$$\begin{aligned}(d/3-1)|\alpha_6 - \alpha_4| &= (d/3-1)\left|2(3/4-\eta)\frac{\eta}{d/3-1} + (d/3-2)\frac{\eta^2}{(d/3-1)^2} - (1/8)\frac{\eta}{d/3-1}\right| \\ &= \eta\left|\frac{11}{8} - \eta\left(1 + \frac{1}{d/3-1}\right)\right| \\ &= \eta\left[\frac{11}{8} - \eta\left(1 + \frac{1}{d/3-1}\right)\right], \text{ since } \eta < \frac{1}{48} < \frac{11}{16},\end{aligned} \tag{B.23}$$

and

$$\begin{aligned}|\alpha_3 - \alpha_1| &= \left|(1/8)(3/4-\eta) + 7/64 - (3/4-\eta)^2 - \frac{\eta^2}{d/3-1} - 1/32\right| \\ &= \left|\frac{25}{64} - \frac{11}{8}\eta + \eta^2\left(1 + \frac{1}{d/3-1}\right)\right|.\end{aligned} \tag{B.24}$$

Noticing that for $d \geq 6$, we have $\frac{25}{64} - \frac{11}{8}\eta + \eta^2\left(1 + \frac{1}{d/3-1}\right) > \frac{25}{64} - \frac{11}{8}\eta$ and $\eta < \frac{1}{48} < \frac{25 \cdot 8}{64 \cdot 11} \approx 0.284$,

$$|\alpha_3 - \alpha_1| = \frac{25}{64} - \frac{11}{8}\eta + \eta^2\left(1 + \frac{1}{d/3-1}\right) \tag{B.25}$$

and

$$|\alpha_3 - \alpha_2| = |(1/8)(3/4 - \eta) + 7/64 - 25/32| = \frac{37}{64} + \frac{\eta}{8}. \tag{B.26}$$

Since $\eta < \frac{1}{48} < \frac{3}{2}$,

$$|\alpha_3 - \alpha_5| = |(1/8)(3/4 - \eta) + 7/64 - 1/64| = \frac{3}{16} - \frac{\eta}{8}, \tag{B.27}$$

and a direct computation shows that

$$\kappa_3 = \frac{1}{2}\left((d/3 - 1)|\alpha_6 - \alpha_4| + |\alpha_3 - \alpha_1| + (2d/3 - 2)\alpha_4 + |\alpha_3 - \alpha_2| + |\alpha_3 - \alpha_5|\right) = \frac{37}{64} + \frac{\eta}{8}, \tag{B.28}$$

whence $\kappa_3 < \kappa_1$.

$\kappa_4 < \kappa_1$: We compute

$$|\alpha_4 - \alpha_2| = \left|(1/8)\frac{\eta}{d/3 - 1} - 25/32\right| = \frac{25}{32} - \frac{\eta}{8(d/3 - 1)} \tag{B.29}$$

$$|\alpha_4 - \alpha_5| = \left|(1/8)\frac{\eta}{d/3 - 1} - 1/64\right| = \frac{1}{64} - \frac{\eta}{8(d/3 - 1)} \tag{B.30}$$

$$\begin{aligned} |\alpha_4 - \alpha_1| &= \left|(3/4 - \eta)^2 + \frac{\eta^2}{d/3 - 1} + 1/32 - (1/8)\frac{\eta}{d/3 - 1}\right| \\ &= \left|\frac{19}{32} - h(\eta, d)\right|, \end{aligned} \tag{B.31}$$

where $h(\eta, d) = \frac{\eta(1/8 - \eta)}{d/3 - 1} + \eta(3/2 - \eta)$. For η restricted to $[0, 1/48]$,

$h(\eta, \cdot)$ is positive and decreasing, and achieves its maximum at $d = 6$. Since $h(\eta, 6) = \eta(13/8 - 2\eta) < 19/32$, we have $|\alpha_4 - \alpha_1| = \frac{19}{32} - h(\eta, d)$. Further,

$$\begin{aligned} |\alpha_6 - \alpha_3| &= \left|2(3/4 - \eta)\frac{\eta}{d/3 - 1} + (d/3 - 2)\frac{\eta^2}{(d/3 - 1)^2} - (1/8)(3/4 - \eta) - 7/64\right| \\ &= \left|\frac{3}{2}\left(\frac{\eta}{d/3 - 1}\right) - \frac{\eta^2}{d/3 - 1} - \frac{\eta^2}{(d/3 - 1)^2} + \frac{\eta}{8} - \frac{13}{64}\right| \\ &=: \left|f(\eta, d) - \frac{13}{64}\right|. \end{aligned} \tag{B.32}$$

Now

$$\frac{\mathrm{d}}{\mathrm{d}\eta} f(\eta, d) = \frac{3}{2}\left(\frac{1}{d/3 - 1}\right) - \frac{2\eta}{d/3 - 1} - \frac{2\eta}{(d/3 - 1)^2} + \frac{1}{8},$$

which is positive as long as $\eta < \frac{3}{2}\dfrac{1 + \frac{d/3 - 1}{8}}{1 + \frac{1}{d/3 - 1}}$. The latter expression increases in d and is minimized at $d = 6$, so the sufficient condition becomes $\eta < \frac{27}{32}$, which holds for our range of η. It follows that

$f(\cdot, d)$ is increasing, and $f(1/48, d) = \frac{2d^2+59d-198}{768(d-3)^2}$, decreasing in d, is majorized by $f(1/48, 6) = \frac{19}{576} \approx 0.033 < 13/64 \approx 0.20$. Thus, we may omit the absolute value in (B.32):

$$|\alpha_6 - \alpha_3| = \frac{13}{64} - \frac{3}{2}\left(\frac{\eta}{d/3-1}\right) + \frac{\eta^2}{d/3-1} + \frac{\eta^2}{(d/3-1)^2} - \frac{\eta}{8}. \tag{B.33}$$

Putting $\tilde{\kappa_4} = (d/3-2)|\alpha_6 - \alpha_4| + (2d/3-4)\alpha_4 + 2\alpha_3 + |\alpha_4 - \alpha_5| + |\alpha_4 - \alpha_2|$, a routine but tedious computation yields

$$\tilde{\kappa_4} = \frac{9}{8}\eta + \eta^2\left(\frac{1}{(d/3-1)^2} - 1\right) + \eta\frac{2d/3-17}{8(d/3-1)} + \frac{77}{64}. \tag{B.34}$$

Combining this with the estimates above, we get

$$\begin{aligned} \kappa_4 &= \frac{1}{2}\left(\tilde{\kappa_4} + |\alpha_4 - \alpha_1| + |\alpha_3 - \alpha_6|\right) \\ &= 1 - \frac{\eta}{8}\left(1 + \frac{14}{d/3-1}\right) + \frac{\eta^2}{d/3-1}\left(1 + \frac{1}{d/3-1}\right) \end{aligned} \tag{B.35}$$

and so $\kappa_4 < \kappa_1$ as long as $\eta < \frac{13}{8}\left(1 + \frac{1}{d/3-1}\right)^{-1}$, which always holds for our range of η and d. This completes the proof of (B.13) — the special case where $\boldsymbol{\tau} = \mathbf{0}$.

General $\boldsymbol{\tau}$. The general case is proved along a very similar scheme, which we outline below. Start by computing $\boldsymbol{M}^2$:

$$\boldsymbol{M}^2_{\eta,\boldsymbol{\tau}} = \begin{pmatrix} C^{(2)}_{\eta,\boldsymbol{\tau}} & R^{(2)}_{\eta,\boldsymbol{\tau}} \\ R^{(2)\intercal}_{\eta,\boldsymbol{\tau}} & L^{(2)}_{\boldsymbol{\tau}} \end{pmatrix} \tag{B.36}$$

where $C^{(2)}_{\eta,\boldsymbol{\tau}} \in \mathbb{R}^{d/3\times d/3}$, $L^{(2)}_{\boldsymbol{\tau}} \in \mathbb{R}^{2d/3\times 2d/3}$, and $R^{(2)}_{\eta,\boldsymbol{\tau}} \in \mathbb{R}^{d/3\times 2d/3}$ are given by

$$C^{(2)}_{\eta,\boldsymbol{\tau}} = \begin{pmatrix} \alpha_1^{(1)} & \alpha_6 & \dots & \alpha_6 \\ \alpha_6 & \alpha_1^{(2)} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \alpha_6 \\ \alpha_6 & \dots & \alpha_6 & \alpha_1^{(d/3)} \end{pmatrix}, \tag{B.37}$$

$$R^{(2)}_{\eta,\boldsymbol{\tau}} = \begin{pmatrix} \alpha_3^{(d/3+1)} & \alpha_3^{(d/3+2)} & \alpha_4^{(d/3+3)} & \dots & \dots & \dots & \alpha_4^{(d)} \\ \alpha_4^{(d/3+1)} & \alpha_4^{(d/3+2)} & \alpha_3^{(d/3+3)} & \alpha_3^{(d/3+4)} & \alpha_4^{(d/3+5)} & \dots & \alpha_4^{(d)} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \alpha_4^{(d/3+1)} & \dots & \dots & \dots & \alpha_4^{(d-2)} & \alpha_3^{(d-1)} & \alpha_3^{(d)} \end{pmatrix}, \tag{B.38}$$

$$L^{(2)}_{\boldsymbol{\tau}} = \begin{pmatrix} S^{(2,1)}_{\boldsymbol{\tau}} & & & \mathbf{0} \\ & S^{(2,2)}_{\boldsymbol{\tau}} & & \\ & & \ddots & \\ \mathbf{0} & & & S^{(2,d/3)}_{\boldsymbol{\tau}} \end{pmatrix}, \qquad S^{(2,k)}_{\boldsymbol{\tau}} = \begin{pmatrix} \alpha_2^{(d/3+2k-1)} & \alpha_5^{(d/3+2k)} \\ \alpha_5^{(d/3+2k)} & \alpha_2^{(d/3+2k)} \end{pmatrix}, \tag{B.39}$$

where

$$\begin{aligned}
\alpha_1^{(i)} &= (3/4-\eta)^2 + \frac{\eta^2}{d/3-1} + 1/32 + \frac{\tau_i\varepsilon^2}{2} \\
\alpha_2^{(2k+1)} &= 25/32 + \frac{\tau_k\varepsilon}{2}\left(\varepsilon - \frac{3}{2}\right) \\
\alpha_2^{(2k+2)} &= 25/32 + \frac{\tau_k\varepsilon}{2}\left(\varepsilon + \frac{3}{2}\right) \\
\alpha_3^{(2k+1)} &= (1/8)(3/4-\eta) + 7/64 + \frac{\tau_k\varepsilon}{2}\left(\frac{3}{2}-\eta\right) - \frac{\tau_k\varepsilon^2}{4} \\
\alpha_3^{(2k+2)} &= (1/8)(3/4-\eta) + 7/64 - \frac{\tau_k\varepsilon}{2}\left(\frac{3}{2}-\eta\right) - \frac{\tau_k\varepsilon^2}{4} \\
\alpha_4^{(2k+1)} &= (1/8)\frac{\eta}{d/3-1} + \frac{\tau_k\varepsilon}{2}\frac{\eta}{d/3-1} \\
\alpha_4^{(2k+2)} &= (1/8)\frac{\eta}{d/3-1} - \frac{\tau_k\varepsilon}{2}\frac{\eta}{d/3-1} \\
\alpha_5^{(i)} &= 1/64 - \frac{\tau_i\varepsilon^2}{4} \\
\alpha_6 &= 2(3/4-\eta)\frac{\eta}{d/3-1} + (d/3-2)\frac{\eta^2}{(d/3-1)^2}.
\end{aligned} \tag{B.40}$$

The case analysis, entirely analogous to our argument for $\boldsymbol{\tau} = \mathbf{0}$ (but with more cases to consider), yields

$$\kappa(\boldsymbol{M}_{\eta,\boldsymbol{\tau}}^2) = 1 - \eta\left(\frac{1}{8} - \frac{\varepsilon}{2}\right)\left(1 + \frac{1}{d/3-1}\right) \leq 1 - \frac{\eta}{16}.$$

■

Lemma 10 (Cover time) *For $\boldsymbol{M} \in \mathcal{H}_\eta$ [defined in (4.12)], the random variable $T_{\text{CLIQ}} = T_{\text{CLIQ}}(\boldsymbol{M})$ [defined in (4.14)] satisfies*

$$m \leq \frac{d}{20\eta}\ln\left(\frac{d}{3}\right) \implies \mathbf{P}\left(T_{\text{CLIQ}} > m\right) \geq \frac{1}{5} \tag{B.41}$$

Proof

Let $\boldsymbol{M} \in \mathcal{H}_\eta$ and $\boldsymbol{M}_I \in \mathcal{M}_{d/3}$ be such that $\boldsymbol{M}_I$ consists only in the inner clique of $\boldsymbol{M}$, and each outer rim state got absorbed into its unique inner clique neighbor:

$$\boldsymbol{M}_I = \begin{pmatrix} 1-\eta & \frac{\eta}{d/3-1} & \cdots & \frac{\eta}{d/3-1} \\ \frac{\eta}{d/3-1} & 1-\eta & \ddots & \vdots \\ \vdots & \ddots & \ddots & \frac{\eta}{d/3-1} \\ \frac{\eta}{d/3-1} & \cdots & \frac{\eta}{d/3-1} & 1-\eta \end{pmatrix}.$$

By construction, it is clear that $T_{\text{CLIQ}}(\boldsymbol{M})$ is almost surely greater than the cover time of $\boldsymbol{M}_I$. The latter corresponds to a generalized coupon collection time
$U = 1 + \sum_{i=1}^{d/3-1} U_i$ where U_i
is the time increment between the ith and the $(i+1)$th unique visited state.
Formally, if $\boldsymbol{X}$ is a random walk according to $\boldsymbol{M}_I$ (started from any state), then $U_1 = \min\{t > 1 : X_t \neq X_1\}$ and for $i > 1$,

$$U_i = \min\{t > 1 : X_t \notin \{X_1, \dots, X_{U_{i-1}}\}\} - U_{i-1}. \tag{B.42}$$

The random variables $U_1, U_2, \dots, U_{d/3-1}$ are independent and $U_i \sim \text{Geometric}\left(\eta - \dfrac{(i-1)\eta}{d/3}\right)$, whence

$$\mathbf{E}[U_i] = \frac{d/3}{\eta(d/3 - i + 1)}, \qquad \mathbf{Var}[U_i] = \frac{1 - \left(\eta - \dfrac{(i-1)\eta}{d/3}\right)}{\left(\eta - \dfrac{(i-1)\eta}{d/3}\right)^2} \tag{B.43}$$

and

$$\mathbf{E}[U] \geq 1 + \frac{d/3}{\eta}\sigma_{d/3-1}, \qquad \mathbf{Var}[U] \leq \frac{(d/3-1)^2}{\eta^2}\frac{\pi^2}{6} \tag{B.44}$$

where $\sigma_d = \sum_{i=1}^{d} \frac{1}{i}$, and $\pi = 3.1416\dots$.

Invoking the Paley-Zygmund inequality with $\theta = 1 - \frac{2\sqrt{2/3}}{\sigma_{d/3-1}}$ we have

$$\begin{aligned} \mathbf{P}\left(U > \theta\mathbf{E}[U]\right) &\geq \left(1 + \frac{\mathbf{Var}[U]}{(1-\theta)^2(\mathbf{E}[U])^2}\right)^{-1} \\ &\geq \left(1 + \frac{5}{3(1-\theta)^2\sigma_{d/3-1}^2}\right)^{-1} = \frac{1}{5} \end{aligned} \tag{B.45}$$

(since $\frac{\pi^2}{6} \leq \frac{5}{3}$). Further, $\sigma_{d/3-1} \geq \sigma_3 = 11/6$ implies

$$\theta\mathbf{E}[U] \geq \frac{3}{20} \cdot \frac{d/3}{\eta}\sigma_{d/3-1} \geq \frac{d}{20\eta}\ln\left(\frac{d}{3}\right),$$

and thus for $m \leq \frac{d}{20\eta}\ln\left(\frac{d}{3}\right)$, we have $\mathbf{P}\left(T_{\text{CLIQ}} > m\right) \geq \frac{1}{5}$.

■

Appendix C. Additional figures

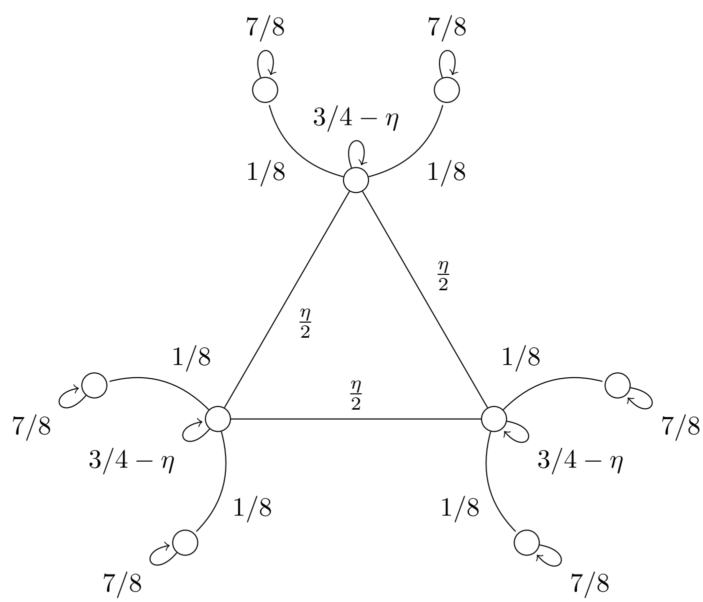

Figure 2: Graph representation of $\boldsymbol{M}_{\eta,\mathbf{0}}$ with $d = 9$.

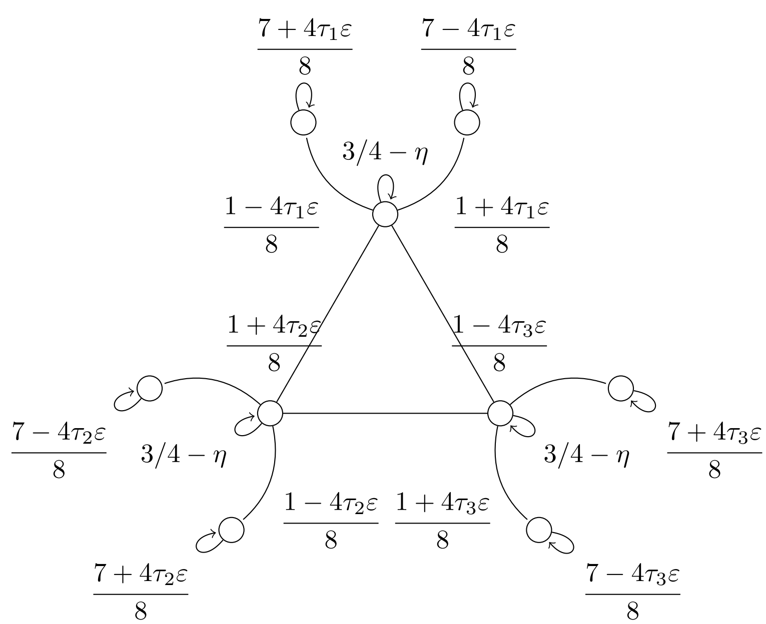

Figure 3: Graph representation of $\boldsymbol{M}_{\eta,\boldsymbol{\tau}}$ for $d = 9$ and general $\boldsymbol{\tau} \in \{0,1\}^3$.

Appendix D. Miscellaneous results

Lemma 11 (Computing the Cheeger constant) *For $0 < \eta < 1/48$ and $\boldsymbol{M} \in \mathcal{H}_\eta$, we have $\Phi(\boldsymbol{M}) = 3\eta$. [$\mathcal{H}_\eta$ and Φ are defined in (4.12) and (B.8).]*

Proof The proof proceeds in three steps. First we exhibit a set of states $S_{3\eta}$ that achieves the value 3η, which proves the upper estimate on Φ. We then argue that if an $S \subset [d]$ fails to satisfy some property, then it cannot achieve a value smaller than 3η in (B.8). Finally, we optimize over all S that *do* satisfy the latter property.

Step I. For $\boldsymbol{M}$ with uniform $\boldsymbol{\pi}$, (B.8) simplifies to

$$\Phi = \min_{S \subset [d], |S| \le \frac{d}{2}} |S| \sum_{i \in S, j \notin S} \boldsymbol{M}(i,j). \tag{D.1}$$

Referring to the construction of $\mathcal{H}_\eta$ [defined in (4.12) and illustrated in Figure 1], we partition the states into the *inner clique* $I = \{1, \ldots, d/3\}$ and its complement, the *outer rim* $J = \{d/3+1, \ldots, d\}$; these inherit the obvious connectedness properties from the Markov graph. Let $S_{3\eta} = \{i, j, k\}$ be such that $i \in I$, $j, k \in J$, and $\boldsymbol{M}(i,j) = \frac{1}{8} + \frac{1}{2}\tau_i\varepsilon$, $\boldsymbol{M}(i,k) = \frac{1}{8} - \frac{1}{2}\tau_i\varepsilon$ (that is, i is connected to j and k). As $|S_{3\eta}| = 3 \le \frac{d}{2}$, we have

$$\Phi \le |S_{3\eta}| \sum_{i \in S_{3\eta}, j \notin S_{3\eta}} \boldsymbol{M}(i,j) = 3 \cdot \left(\frac{d}{3} - 1\right) \cdot \left(\frac{\eta}{d/3 - 1}\right) = 3\eta.$$

Step II. For any $S \subset [d]$, suppose that $i \in S \cap I$ has two neighbors $j, k \in J$ such that at least one (say, j) is not in S. Since $\eta < 1/48$ and $\varepsilon < 1/8$, plugging such an S into the minimand in (D.1) yields

$$|S| \sum_{i \in S, j \notin S} \boldsymbol{M}(i,j) \ge 1/8 \pm \tau_i\varepsilon/2 > 3\eta. \tag{D.2}$$

An analogous argument shows that (D.2) also holds if $j \in S \cap J$ has a neighbor $i \in I \setminus S$. It follows that Φ is fully determined by the quantity $|S \cap I|$.

Step III. In light of Step II, we may rewrite the objective function of the optimization problem (D.1) as follows, where, for $\tilde{I} \subseteq I$, we write $\tilde{J}(\tilde{I})$ to denote the neighbors of $\tilde{I}$ in J:

$$\begin{aligned} \Phi &= \min_{\tilde{I} \subseteq I, |\tilde{I}| \le d/6} 3|\tilde{I}| \sum_{i \in \tilde{I} \cup \tilde{J}(\tilde{I}), j \notin \tilde{I} \cup \tilde{J}(\tilde{I})} \boldsymbol{M}(i,j) \\ &= \min_{k \in [d/6]} (3k)k \left(\frac{d}{3} - k\right)\left(\frac{\eta}{d/3 - 1}\right). \end{aligned}$$

Since $k \mapsto k^2(d/3 - k)$ is increasing on $[d/6]$, the the minimum is achieved at $k = 1$, which shows that $\Phi = 3\eta$. ■

Appendix E. Reported results from literature

Theorem 12 (Rectangular Matrix Freedman, (Tropp et al., 2011, Corollary 1.3) (weakened version)) *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:*

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

Define two predictable quadratic variation processes for this martingale:

$$\begin{aligned} \boldsymbol{W}_{col,m} &:= \sum_{t=1}^{m} \mathbf{E}\left[\boldsymbol{Y}_t \boldsymbol{Y}_t^{\intercal} \,|\, \mathcal{F}_{t-1}\right] \text{ and} \\ \boldsymbol{W}_{row,m} &:= \sum_{t=1}^{m} \mathbf{E}\left[\boldsymbol{Y}_t^{\intercal} \boldsymbol{Y}_t \,|\, \mathcal{F}_{t-1}\right] \text{ for } m = 1, 2, 3, \ldots. \end{aligned} \tag{E.1}$$

Then, for all $\varepsilon \geq 0$ *and* $\sigma^2 > 0$*,*

$$\mathbf{P}\left(\left\|\sum_{t=1}^{m} \boldsymbol{Y}_t\right\|_2 > \varepsilon \text{ and } \max\left\{\|\boldsymbol{W}_{row,m}\|_2, \|\boldsymbol{W}_{col,m}\|_2\right\} \leq \sigma^2\right) \leq (d_1 + d_2) \cdot \exp\left(-\frac{\varepsilon^2/2}{\sigma^2 + R\varepsilon/3}\right). \tag{E.2}$$