

Minimax Testing of Identity to a Reference Ergodic Markov Chain

Geoffrey Wolfer
geoffrey@post.bgu.ac.il

Aryeh Kontorovich
karyeh@cs.bgu.ac.il

January 31, 2019

Abstract

We exhibit an efficient procedure for testing, based on a single long state sequence, whether an unknown Markov chain is identical to or ε -far from a given reference chain. We obtain nearly matching (up to logarithmic factors) upper and lower sample complexity bounds for our notion of distance, which is based on total variation. Perhaps surprisingly, we discover that the sample complexity depends solely on the properties of the known reference chain and does not involve the unknown chain at all, which is not even assumed to be ergodic.

1 Introduction

Distinguishing whether an unknown distribution \mathbf{D} is identical to a reference one $\bar{\mathbf{D}}$ or is ε -far from it in total variation (TV) is a special case of statistical *property testing*. For the iid case, it is known that a sample of size $\tilde{\Theta}(\sqrt{d}/\varepsilon^2)$, where d is the support size, is both sufficient and necessary [Batu et al., 2001, Valiant and Valiant, 2017]. This is in contradistinction to the corresponding learning problem, with the considerably higher sample complexity of $\tilde{\Theta}(d/\varepsilon^2)$ (see, e.g., Anthony and Bartlett [1999], Waggoner [2015], Kontorovich and Pinelis [2019+]). The Markovian setting has so far received no attention in the property testing framework, with the notable exception of the recent work of Daskalakis et al. [2017], which we discuss in greater detail in Section 2. The latter work “initiate[d] the study of Markov chain testing”, but imposed the stringent constraint of being symmetric on both the reference and unknown chains. In this paper, we only require ergodicity of the reference chain, and make no assumptions on the unknown one — from which the tester receives a single long trajectory of observations — other than it having d states.

Our contribution. We prove nearly matching upper and lower bounds on the sample complexity of the testing problem in terms of the accuracy ε and the number of states d , as well as parameters derived from the stationary distribution and mixing. We discover that for testing, only the reference chain affects the sample complexity, and no assumptions (including ergodicity) need be made on the unknown one. In particular, we exhibit an efficient testing procedure, which, given a Markovian sequence of length

$$m = \tilde{\mathcal{O}} \left(\frac{1}{\pi_\star} \max \left\{ \frac{\sqrt{d}}{\varepsilon^2}, t_{\text{mix}} \right\} \right), \quad (1.1)$$

correctly identifies the unknown chain with high probability, where π_\star and t_{mix} are, respectively, the minimum stationary probability and mixing time of the *reference* chain. We also derive an instance-specific version of the previous bound: the factor $\frac{\sqrt{d}}{\pi_\star}$ in (1.1) can be replaced with

the potentially much smaller $\max_{i \in [d]} \left\{ \pi(i)^{-1} \|\overline{\mathbf{M}}(i, \cdot)\|_{2/3} \right\}$, defined in (4.3). Additionally, we construct two separate worst case lower bounds of $\Omega(dt_{\text{mix}})$ and $\Omega\left(\frac{\sqrt{d}}{\pi_* \varepsilon^2}\right)$, exhibiting a regime for which our testing procedure is unimprovable.

2 Related work

We consider *distribution testing* in the *property testing* (as opposed to the more classical statistical *hypothesis testing*¹) framework — a research program initiated by Batu et al. [2000]. The special case of iid uniformity testing was addressed (for various metrics) by Goldreich and Ron [2011], Paninski [2008]. Extensions to iid identity testing for arbitrary finite distributions were then obtained [Goldreich, 2016, Diakonikolas et al., 2016], including the instance-optimal tester of Valiant and Valiant [2017], who showed that \sqrt{d} may be replaced with $\|\overline{\mathbf{D}}\|_{2/3}$, the (2/3)-pseudo-norm of the reference distribution.

To our knowledge, Daskalakis et al. [2017] were the first to consider the testing problem for Markov chains (see references therein for previous works addressing goodness-of-fit testing under Markov dependence). Their model is based on the pseudo-distance $\text{Dist}_{\sqrt{\cdot}}(\cdot, \cdot)$ defined first by Kazakos [1978] as

$$\text{Dist}_{\sqrt{\cdot}}(\mathbf{M}, \mathbf{M}') = 1 - \rho\left([\mathbf{M}, \mathbf{M}']_{\sqrt{\cdot}}\right), \quad (2.1)$$

where $\left([\mathbf{M}, \mathbf{M}']_{\sqrt{\cdot}}\right)_{(i,j)} = \sqrt{\mathbf{M}(i,j)\mathbf{M}'(i,j)}$ is the term-wise geometric mean of the transition kernels and ρ is the largest eigenvalue in magnitude. This pseudo-distance has the property of vanishing on pairs of chains sharing an identical connected component. Daskalakis et al.’s sample complexity upper bound of $\tilde{O}(\text{HitT}_{\overline{\mathbf{M}}} + d/\varepsilon)$ requires knowledge of the hitting time of the reference chain, while their lower bound $\Omega(d/\varepsilon)$ involves no quantities related to the mixing rate at all. The authors conjectured that for $\text{Dist}_{\sqrt{\cdot}}(\cdot)$, the correct sample complexity is $\Theta(d/\varepsilon)$ — i.e., independent of the mixing properties of the chain.

The present paper compares favorably with Daskalakis et al. in that the latter requires both the reference and the unknown chains to be symmetric (and, a fortiori, reversible) as well as ergodic. We only require ergodicity of the reference chain and assume nothing about the unknown one. Additionally, we obtain nearly sharp sample complexity bounds in terms of the reference chain’s mixing properties. Finally, our metric $\|\cdot\|$ dominates the pseudo-metric $\text{Dist}_{\sqrt{\cdot}}(\cdot)$, and hence their identity testing problem is reducible to ours (see Lemma C.1 in the Appendix).

We note that the corresponding PAC-type learning problem for Markov chains was only recently introduced in Hao et al. [2018], Wolfer and Kontorovich [2019]. This paper uses the same notion of distance as the latter work, and exhibits a quadratic reduction in the bound’s dependence on d , which was calculated therein to be $m = \tilde{\Theta}\left(\frac{1}{\pi_*} \max\left\{\frac{d}{\varepsilon^2}, t_{\text{mix}}\right\}\right)$.

3 Definitions and notation

We define $[d] \triangleq \{1, \dots, d\}$, denote the simplex of all distributions over $[d]$ by Δ_d , and the collection of all $d \times d$ row-stochastic matrices by \mathcal{M}_d . 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. For $n \in \mathbb{N}$, and any $\mathbf{D} \in \Delta_d$ we also consider its n -fold *product* $\mathbf{D}^{\otimes n}$, i.e. $(X_1, \dots, X_n) \sim \mathbf{D}^{\otimes n}$ is a shorthand for $(X_i)_{i \in [n]}$ being all mutually independent, and such that $\forall i \in [n], X_i \sim \mathbf{D}$. A Markov chain on d states being entirely specified by an initial distribution $\boldsymbol{\mu} \in \Delta_d$ and a row-stochastic transition matrix $\mathbf{M} \in \mathcal{M}_d$, we identify the chain with the pair $(\mathbf{M}, \boldsymbol{\mu})$. Namely, by

¹A recent result in this vein is Barsotti et al. [2016]; see also references therein.

$(X_1, \dots, X_m) \sim (\mathbf{M}, \boldsymbol{\mu})$, we mean that

$$\mathbf{P}((X_1, \dots, X_m) = (x_1, \dots, x_m)) = \boldsymbol{\mu}(x_1) \prod_{t=1}^{m-1} \mathbf{M}(x_t, x_{t+1}). \quad (3.1)$$

We write $\mathbf{P}_{\mathbf{M}, \boldsymbol{\mu}}(\cdot)$ to denote probabilities over sequences induced by the Markov chain $(\mathbf{M}, \boldsymbol{\mu})$, and omit the subscript when it is clear from context. Taking the null hypothesis to be that $\mathbf{M} = \overline{\mathbf{M}}$ (i.e., the chain being tested is identical to the reference one), $\mathbf{P}_0(\cdot)$ will denote probability in the completeness case, and $\mathbf{P}_1(\cdot)$ in the soundness case. The Markov chain $(\mathbf{M}, \boldsymbol{\mu})$ is *stationary* if $\boldsymbol{\mu} = \boldsymbol{\pi}$ for $\boldsymbol{\pi} = \boldsymbol{\pi}\mathbf{M}$, and *ergodic* if $\mathbf{M}^k > 0$ (entry-wise positive) for some $k \geq 1$. If \mathbf{M} is ergodic, it has a unique stationary distribution $\boldsymbol{\pi}$ and moreover the *minimum stationary probability* $\pi_\star > 0$, where

$$\pi_\star \triangleq \min_{i \in [d]} \boldsymbol{\pi}(i). \quad (3.2)$$

Unless noted otherwise, $\boldsymbol{\pi}$ is assumed to be the stationary distribution of the Markov chain in context. The *mixing time* of a chain is defined as the number of steps necessary for its state distribution to be sufficiently close to the stationary one (traditionally taken to be within $1/4$):

$$t_{\text{mix}} \triangleq \min_{t \in \mathbb{N}} \left\{ \sup_{\boldsymbol{\mu} \in \Delta_d} \|\boldsymbol{\mu}\mathbf{M}^{t-1} - \boldsymbol{\pi}\|_{\text{TV}} \leq \frac{1}{4} \right\}. \quad (3.3)$$

We use the standard ℓ_1 norm $\|z\|_1 = \sum_{i \in [d]} |z_i|$, 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

$$\|A\| \triangleq \max_{i \in [d]} \|A(i, \cdot)\|_1 = \max_{i \in [d]} \sum_{j \in [d]} |A(i, j)|. \quad (3.4)$$

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

Definition 3.1 *An (ε, δ) -identity tester \mathcal{T} for Markov chains with sample complexity function $m_0(\cdot)$ is an algorithm that takes as input a reference Markov chain $(\overline{\mathbf{M}}, \overline{\boldsymbol{\mu}})$ and $\mathbf{X} = (X_1, \dots, X_m)$ drawn from some unknown Markov chain $(\mathbf{M}, \boldsymbol{\mu})$, and outputs $\mathcal{T} = \mathcal{T}(d, \varepsilon, \delta, \overline{\mathbf{M}}, \overline{\boldsymbol{\mu}}, \mathbf{X}) \in \{0, 1\}$ such that for $m \geq m_0(d, \varepsilon, \delta, \overline{\mathbf{M}}, \overline{\boldsymbol{\mu}})$, both $\mathbf{M} = \overline{\mathbf{M}} \Rightarrow \mathcal{T} = 0$ and $\|\mathbf{M} - \overline{\mathbf{M}}\| > \varepsilon \Rightarrow \mathcal{T} = 1$ hold with probability at least $1 - \delta$. (The probability is over the draw of \mathbf{X} and any internal randomness of the tester.)*

4 Formal results

Since the focus of this paper is on statistical rather than computational complexity, we defer the (straightforward) analysis of the runtimes of our tester to the Appendix, Section B.2.

Theorem 4.1 (Identity testing sample complexity upper bound) *There exists an (ε, δ) -identity tester \mathcal{T} (provided in the Appendix, Section B.1, Algorithm 1), which, for all $0 < \varepsilon < 2$, $0 < \delta < 1$, satisfies the following. If \mathcal{T} receives as input a d -state “reference” ergodic Markov chain $(\overline{\mathbf{M}}, \overline{\boldsymbol{\mu}})$, as well as a sequence $\mathbf{X} = (X_1, \dots, X_m)$ of length at least m_{UB} , drawn according to an unknown chain \mathbf{M} (starting from an arbitrary state), then it outputs $\mathcal{T} = \mathcal{T}(d, \varepsilon, \delta, \overline{\mathbf{M}}, \overline{\boldsymbol{\mu}}, \mathbf{X}) \in \{0, 1\}$ such that*

$$\mathbf{M} = \overline{\mathbf{M}} \implies \mathcal{T} = 0 \quad \text{and} \quad \|\mathbf{M} - \overline{\mathbf{M}}\| > \varepsilon \implies \mathcal{T} = 1$$

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

$$m_{\text{UB}} := \frac{C_{\text{MC}}}{\pi_{\star}} \max \left\{ \frac{\sqrt{d}}{\varepsilon^2} \ln \left(\frac{d^{3/2}}{\delta \varepsilon^2} \right), t_{\text{mix}} \ln \left(\frac{d\sqrt{1/\pi_{\star}}}{\delta} \right) \right\} = \tilde{\mathcal{O}} \left(\frac{1}{\pi_{\star}} \max \left\{ \frac{\sqrt{d}}{\varepsilon^2}, t_{\text{mix}} \right\} \right), \quad (4.1)$$

where C_{MC} is a universal constant, and t_{mix} and π_{\star} are respectively the mixing time (Eq. 3.3) and the minimum stationary probability (Eq. 3.2) of $\bar{\mathbf{M}}$.

Remark 4.1 An important feature of Theorem 4.1 is that the sample complexity only depends on the (efficiently computable, see Section B.2) properties of the known reference chain. No assumptions, such as symmetry (as in Daskalakis et al. [2017]) or even ergodicity, are made on the unknown Markov chain, and none of its properties appear in the bound.

Remark 4.2 Our results indicate that in the regime where the $\frac{t_{\text{mix}}}{\pi_{\star}}$ term is not dominant, the use of optimized identity iid testers as subroutines confers an $\tilde{\mathcal{O}}(\sqrt{d})$ -fold improvement over the naive testing-by-learning strategy. In particular, one can always take $\varepsilon = \tilde{\mathcal{O}}\left(\sqrt{t_{\text{mix}}^{-1}d^{1/4}}\right)$, since settling for a worse precision provides no gains in sample complexity.

Remark 4.3 We note but do not pursue the fact that the logarithmic dependencies on δ in our upper bound could be improved via a combination of the techniques of Diakonikolas et al. [2017] and the reduction to uniformity testing of Goldreich [2016].

Theorem 4.2 (Instance-specific identity testing sample complexity upper bound) *There exists an (ε, δ) -identity tester \mathcal{T} , which, for all $0 < \varepsilon < 2$, $0 < \delta < 1$, satisfies the following. If \mathcal{T} receives as input a d -state “reference” ergodic Markov chain $(\bar{\mathbf{M}}, \bar{\boldsymbol{\mu}})$, as well as a sequence $\mathbf{X} = (X_1, \dots, X_m)$ of length at least m_{UB} , drawn according to an unknown chain \mathbf{M} (starting from an arbitrary state), then it outputs $\mathcal{T} = \mathcal{T}(d, \varepsilon, \delta, \bar{\mathbf{M}}, \bar{\boldsymbol{\mu}}, \mathbf{X}) \in \{0, 1\}$ such that*

$$\mathbf{M} = \bar{\mathbf{M}} \implies \mathcal{T} = 0 \quad \text{and} \quad \|\mathbf{M} - \bar{\mathbf{M}}\| > \varepsilon \implies \mathcal{T} = 1$$

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

$$m_{\text{UB}} := \tilde{\mathcal{O}} \left(\max \left\{ \frac{\|\bar{\mathbf{M}}\|_{\pi, 2/3}}{\varepsilon^2}, \frac{t_{\text{mix}}}{\pi_{\star}} \right\} \right), \quad (4.2)$$

where t_{mix} and π_{\star} are as in Theorem 4.1, and

$$\|\bar{\mathbf{M}}\|_{\pi, 2/3} \triangleq \max_{i \in [d]} \left\{ \frac{\left(\sum_{j \in [d]} \bar{\mathbf{M}}(i, j)^{2/3} \right)^{3/2}}{\boldsymbol{\pi}(i)} \right\}. \quad (4.3)$$

Remark 4.4 Since we always have $\|\bar{\mathbf{M}}\|_{\pi, 2/3} \leq \sqrt{d}/\pi_{\star}$, the instance-specific bound is always at least as sharp as the worst-case one in Theorem 4.1. It may, however, be considerably sharper, as illustrated by a simple random walk on a d -vertex, Δ -regular graph [Levin et al., 2009, Section 1.4], for which the instance-specific bound is $\tilde{\mathcal{O}}\left(d \max \left\{ \frac{\sqrt{\Delta}}{\varepsilon^2}, t_{\text{mix}} \right\}\right)$ — a savings of roughly \sqrt{d} .

Theorem 4.3 (Identity testing sample complexity lower bound) *For every $0 < \varepsilon < 1/8$, $t_{\text{mix}} \geq 50$, and $d = 6k$, $k \geq 2$, there exists a d -state Markov chain $\bar{\mathbf{M}}$ with mixing time t_{mix} and stationary distribution $\boldsymbol{\pi}$ such that every $(\varepsilon, 1/10)$ -identity tester for reference chain*

$\overline{\mathbf{M}}$ must require in the worst case a sequence $\mathbf{X} = (X_1, \dots, X_m)$ drawn from the unknown chain \mathbf{M} of length at least

$$m_{\text{LB}} := \tilde{\Omega} \left(\max \left\{ \frac{\sqrt{d}}{\varepsilon^2 \pi_\star}, dt_{\text{mix}} \right\} \right),$$

where $t_{\text{mix}}, \pi_\star$ are as in Theorem 4.1.

Remark 4.5 As the proof shows, for any $0 < \pi_\star < 1/(2(d+1))$, a testing problem can be constructed that achieves the $\frac{\sqrt{d}}{\varepsilon^2 \pi_\star}$ component of the lower bound. Moreover, for doubly-stochastic $\overline{\mathbf{M}}$, we have $\pi_\star = 1/d$, which shows that the upper bound cannot be improved in all parameters simultaneously.

5 Overview of techniques

For both upper and lower bounds, we survey existing techniques, describe their limitations vis-à-vis our problem, and highlight the key technical challenges as well as our solutions for overcoming these.

5.1 Upper bounds

Naïve approach: testing-by-learning. We mention this approach mainly to establish a baseline comparison. Wolfer and Kontorovich [2019] showed that in order to (ε, δ) -learn an unknown d -state ergodic Markov chain $(\mathbf{M}, \boldsymbol{\mu})$ under the $\|\cdot\|$ distance, a single trajectory of length $m = \tilde{O} \left(\frac{1}{\pi_\star} \max \left\{ \frac{d}{\varepsilon^2}, t_{\text{mix}} \right\} \right)$ is sufficient. It follows that one can test identity with sample complexity

$$m = \tilde{O} \left(\max \left\{ \frac{1}{\pi_\star(\mathbf{M})} \max \left\{ \frac{d}{\varepsilon^2}, t_{\text{mix}}(\mathbf{M}) \right\}, \frac{1}{\pi_\star(\overline{\mathbf{M}})} \max \left\{ \frac{d}{\varepsilon^2}, t_{\text{mix}}(\overline{\mathbf{M}}) \right\} \right\} \right). \quad (5.1)$$

This naïve bound, aside from being much looser than bounds provided in the present paper, has the additional drawback of depending on the unknown \mathbf{M} and, in particular, being completely uninformative when the latter is not ergodic.

Reduction to iid testing. Our upper bound in Theorem 4.1 is achieved via the stratagem of invoking an existing iid distribution identity tester as a black box (this is also the general approach of Daskalakis et al. [2017]). Intuitively, given the reference chain $\overline{\mathbf{M}}$, we can compute its stationary distribution $\boldsymbol{\pi}$ and thus know roughly how many visits to expect in each state. Further, computing the mixing time t_{mix} gives us confidence intervals about these expected visits (similar to Wolfer and Kontorovich [2019], via the concentration bounds of Paulin [2015]). Hence, if a chain fails to visit each state a “reasonable” number of times, our tester in Algorithm 1 rejects it. Otherwise, given that state i has been visited as expected, we can apply an iid identity tester to its conditional distribution. The unknown Markov chain passes the identity test if every state’s conditional distribution passes its corresponding iid test.

A central technical challenge in executing this stratagem is the fact that conditioning on the number of visits introduces dependencies on the sample, thereby breaking the Markov property. Getting around this difficulty required a novel argument involving a union bound over the possible numbers of visits, which incurred an additional logarithmic factor in the bound. Finally, invoking the tester of Valiant and Valiant [2017] as a black box, it is straightforward to sharpen the worst-case bound in Theorem 4.1 to the instance-specific one in Theorem 4.2.

5.2 Lower bounds

A lower bound of $\Omega(d/\varepsilon^2)$ is immediate via a reduction from the testing problem of Daskalakis et al. [2017] to ours (see Remark C.1). Although our construction for obtaining the sharper lower bound of $\Omega(\sqrt{d}/(\pi_*\varepsilon^2))$ shares some conceptual features with the constructions in Hao et al. [2018], Wolfer and Kontorovich [2019], a considerably more delicate analysis is required here. Indeed, the technique of tensorizing the KL divergence, instrumental in the lower bound of Wolfer and Kontorovich, would yield (at best) a sub-optimal estimate of $\Omega(1/(\pi_*\varepsilon^2))$ in our setting. Intuitively, bounding TV via KL divergence is too crude for our purposes. Instead, we take the approach of reducing the problem, via a covering argument, to one of iid testing, and construct a family of Markov chains whose structure allows us to recover the Markov property even after conditioning on the number of visits to a certain “special” state. The main contribution for this argument is the decoupling technique of Lemma A.4. The second lower bound is based on the construction of Wolfer and Kontorovich, for which the mixing time and accuracy of the test can both be controlled independently. Curiously, the aforementioned argument cannot be invoked verbatim for our problem, and so we introduce here the twist of considering half-covers of the chains (Lemma A.5), concluding the argument with a two-point technique. This adaptation shaves a logarithmic factor off the corresponding learning problem.

6 Proofs

6.1 Proof of Theorem 4.1

For $\mathbf{X} \sim (M, \mu)$, we define our identity tester $\mathcal{T}(\mathbf{X})$ in terms of d sub-testers $\mathcal{T}^{(i)}(\mathbf{X})$, $i \in [d]$, whose definition we defer until further in the analysis. Intuitively, each $\mathcal{T}^{(i)}$ requires “reasonable” number of visits to i :

$$\mathcal{T}(\mathbf{X}) \triangleq \mathbf{1} \left\{ \exists i \in [d], \mathcal{T}^{(i)}(\mathbf{X}) = 1 \text{ or } |N_i - m\pi(i)| > (1/2)m\pi(i) \right\}, \quad (6.1)$$

where $N_i = \sum_{t=1}^{m-1} \mathbf{1}\{X_t = i\}$ is the number of visits to state i (not counting the final state at time m). A pseudo-code implementation of the test is given in the Appendix, Section B.1, Algorithm 1.

Completeness. We begin with an application of the union bound,

$$\mathbf{P}_0(\mathcal{T}(\mathbf{X}) = 1) \leq \sum_{i \in [d]} \mathbf{P}_0\left(\mathcal{T}^{(i)}(\mathbf{X}) = 1\right) + \sum_{i \in [d]} \mathbf{P}_0(|N_i - m\pi(i)| > (1/2)m\pi(i)). \quad (6.2)$$

For any $i \in [d]$, we decompose

$$\mathbb{N} = \{n_i : n_i < \lfloor (1/2)m\pi(i) \rfloor\} \cup \{\lfloor (1/2)m\pi(i) \rfloor, \dots, \lceil (3/2)m\pi(i) \rceil\} \cup \{n_i : n_i > \lceil (3/2)m\pi(i) \rceil\}, \quad (6.3)$$

and obtain from the law of total probability,

$$\begin{aligned} \mathbf{P}_0(\mathcal{T}(\mathbf{X}) = 1) &\leq \sum_{i \in [d]} \sum_{n_i = \lfloor (1/2)m\pi(i) \rfloor}^{\lceil (3/2)m\pi(i) \rceil} \mathbf{P}_0\left(\mathcal{T}^{(i)}(\mathbf{X}) = 1 \text{ and } N_i = n_i\right) \\ &\quad + 2 \sum_{i \in [d]} \mathbf{P}_0(|N_i - m\pi(i)| > (1/2)m\pi(i)). \end{aligned} \quad (6.4)$$

The error term $2 \sum_{i \in [d]} \mathbf{P}_0(|N_i - m\pi(i)| > (1/2)m\pi(i))$ can be made smaller than $\delta/3$ for $m \geq C \frac{t_{\text{mix}}}{\pi_*} \left(\frac{d\sqrt{1/\pi_*}}{\delta} \right)$, where $C > 0$ is a universal constant, using the Bernstein-type concentration

inequalities of Paulin [2015] as made explicit in Wolfer and Kontorovich [2019, Lemma 5]. Observe that no properties of the unknown chain were invoked in this deduction. The sub-testers have the general form $\mathcal{T}^{(i)}(\mathbf{X}) = \mathbf{1}\{S^{(i)}(\mathbf{X}) > t^{(i)}(\mathbf{X})\}$, where both the statistic $S^{(i)}$ and the threshold $t^{(i)}$ may be random, (e.g., can depend on the number of visits to the state i). Consider the subset of such tests, which can be expressed as $S^{(i)}(\mathbf{Y}_1^{N_i})$, where $\mathbf{Y}_1^{N_i} = (Y_1, \dots, Y_{N_i})$ are mutually independent and each is distributed according to $\mathbf{M}(i, \cdot)$, and $t^{(i)}(\mathbf{X})$ may be expressed as $t^{(i)}(N_i)$, depending only on the number of visits to state i . Using our previous decomposition (6.3), for any $i \in [d]$ and any $\lfloor (1/2)m\pi(i) \rfloor \leq n_i \leq \lceil (3/2)m\pi(i) \rceil$,

$$\left\{ \mathcal{T}^{(i)}(\mathbf{X}) = 1 \text{ and } N_i = n_i \right\} = \left\{ S^{(i)}(\mathbf{Y}_1^{N_i}) > t^{(i)}(N_i) \text{ and } N_i = n_i \right\} \subset \left\{ S^{(i)}(\mathbf{Y}_1^{n_i}) > t^{(i)}(n_i) \right\}. \quad (6.5)$$

We recall some standard results.

Lemma 6.1 (Fixed confidence iid $(\varepsilon, 2/5)$ -testing to identity (e.g. Waggoner, 2015))

Let $\varepsilon > 0$ and $d \in \mathbb{N}$. There exists a universal constant C_{IID} and a tester \mathcal{T}_{FIX} , such that for any reference distribution $\bar{\mathbf{D}} \in \Delta_d$, and any unknown distribution $\mathbf{D} \in \Delta_d$, for a sample of size $m_{\text{IID}, 2/5} \geq C_{\text{IID}} \frac{\sqrt{d}}{\varepsilon^2}$ drawn iid from \mathbf{D} , \mathcal{T}_{FIX} can distinguish between the cases $\mathbf{D} \equiv \bar{\mathbf{D}}$ and $\|\mathbf{D} - \bar{\mathbf{D}}\|_{\text{TV}} > \varepsilon$ with probability $3/5$.

Lemma 6.2 (BPP amplification) Given any $(\varepsilon, 2/5)$ -identity tester \mathcal{T} for the iid case with sample complexity $m_{\text{IID}, 2/5}$, and any $0 < \delta < 2/5$, we can construct (via a majority vote) an amplified tester \mathcal{T}_{BPP} such that for $m \geq 18 \ln\left(\frac{2}{\delta}\right) m_{\text{IID}, 2/5}$, \mathcal{T}_{BPP} can distinguish the cases $\mathbf{D} \equiv \bar{\mathbf{D}}$ and $\|\mathbf{D} - \bar{\mathbf{D}}\|_{\text{TV}} > \varepsilon$ with confidence $1 - \delta$.

Remark 6.1 Note that some iid tests, such as the uniformity test of Paninski [2008], work in specific parameter regimes and break down when the sample size becomes too large; hence we only consider tests for which a larger sample size always translates into better guarantees.

Applying (6.5) to (6.4),

$$\begin{aligned} \mathbf{P}_0(\mathcal{T}(\mathbf{X}) = 1) &\leq \sum_{i \in [d]} (m\pi(i) + 2) \mathbf{P}_0\left(\mathcal{T}^{(i)}(\mathbf{Y}_1, \dots, \mathbf{Y}_{\lfloor (1/2)m\pi(i) \rfloor}) = 1\right) \\ &\quad + 2 \sum_{i \in [d]} \mathbf{P}_0(|N_i - m\pi(i)| > (1/2)m\pi(i)). \end{aligned} \quad (6.6)$$

Notice that for $m \geq \frac{4}{\pi_*}$, we always have $m\pi(i) + 2 \leq \frac{3}{2}m\pi(i)$ and $\lfloor \frac{1}{2}m\pi(i) \rfloor \geq \frac{1}{4}m\pi(i)$. We will use the elementary fact that

$$w \geq 2(u + v \ln 2v - v) \implies w \geq u + v \ln w, \quad u, v > 0.$$

Setting $w = (1/4)m\pi(i)$, $v = 18C_{\text{IID}} \frac{\sqrt{d}}{\varepsilon^2}$, $u = v \ln \frac{36d}{\delta}$, we have

$$\begin{aligned} (1/4)m\pi(i) &\geq 36C_{\text{IID}} \frac{\sqrt{d}}{\varepsilon^2} \left[\ln \left(\frac{36d}{\delta} \cdot \frac{18C_{\text{IID}} \sqrt{d}}{\varepsilon^2} \right) - 1 \right] \\ &= 2v \left[\ln \left(\frac{36d}{\delta} v \right) - 1 \right] = 2[u + v \ln 2v - v], \end{aligned} \quad (6.7)$$

which implies that $(1/4)m\pi(i) \geq 18C_{\text{IID}} \frac{\sqrt{d}}{\varepsilon^2} \ln \left(\frac{9dm\pi(i)}{\delta} \right)$. Combined with Lemmas 6.1 and 6.2, this yields

$$\sum_{i \in [d]} ((3/2)m\pi(i)) \mathbf{P}_0\left(\mathcal{T}^{(i)}(\mathbf{Y}_1, \dots, \mathbf{Y}_{\lfloor (1/2)m\pi(i) \rfloor}) = 1\right) \leq \frac{\delta}{3}, \quad (6.8)$$

whence $\mathbf{P}_0(\mathcal{T}(\mathbf{X}) = 1) \leq \frac{2\delta}{3}$.

Soundness. Suppose $i_0 \in [d]$ is such that $\|\overline{\mathbf{M}}(i_0, \cdot) - \mathbf{M}(i_0, \cdot)\|_1 > \varepsilon$. Then, by inclusion of events, we have

$$\begin{aligned} \mathbf{P}_1(\mathcal{T}(\mathbf{X}) = 0) &\leq \mathbf{P}_1\left(\mathcal{T}^{(i_0)}(\mathbf{X}) = 0 \text{ and } |N_{i_0} - m\pi(i_0)| \leq (1/2)m\pi(i_0)\right) \\ &\leq \sum_{n_{i_0}=\lfloor(1/2)m\pi(i_0)\rfloor}^{\lceil(3/2)m\pi(i_0)\rceil} \mathbf{P}_1\left(\mathcal{T}^{(i_0)}(\mathbf{X}) = 0 \text{ and } N_{i_0} = n_{i_0}\right) \\ &\leq ((3/2)m\pi(i_0)) \mathbf{P}_1\left(\mathcal{T}^{(i_0)}(\mathbf{Y}_1, \dots, \mathbf{Y}_{\lfloor(1/2)m\pi(i_0)\rfloor}) = 0\right). \end{aligned} \quad (6.9)$$

Similarly to the completeness argument, in order to guarantee $\mathbf{P}_1(\mathcal{T}(\mathbf{X}) = 0) \leq \frac{\delta}{3}$, it suffices that

$$(1/4)m\pi(i_0) \geq 36C_{\text{IID}} \frac{\sqrt{d}}{\varepsilon^2} \left[\ln \left(\frac{36d}{\delta} \cdot \frac{18C_{\text{IID}}\sqrt{d}}{\varepsilon^2} \right) - 1 \right]. \quad (6.10)$$

Finally, enforcing this requirement on the sample size uniformly over all states completes the proof. \square

6.2 Proof of Theorem 4.2

This claim follows immediately from the analysis of the iid instance-optimal tester [Valiant and Valiant, 2017], which is invoked to test the conditional distributions of each state.

Lemma 6.3 (iid instance optimal $(\varepsilon, 2/5)$ -testing to identity [Valiant and Valiant, 2017])

Let $\varepsilon > 0$ and $d \in \mathbb{N}$. There exists a universal constant C_{IO} , such that for any reference distribution $\overline{\mathbf{D}} \in \Delta_d$, there exists a tester $\mathcal{T}_{\text{FIX,IO}}$ such that for any unknown distribution $\mathbf{D} \in \Delta_d$, for a sample of size $m_{\text{IO},2/5} \geq C_{\text{IO}} \frac{\|\overline{\mathbf{D}}\|_{2/3}}{\varepsilon^2}$ drawn iid from \mathbf{D} , $\mathcal{T}_{\text{FIX,IO}}$ can distinguish between the cases $\mathbf{D} \equiv \overline{\mathbf{D}}$ and $\|\mathbf{D} - \overline{\mathbf{D}}\|_{\text{TV}} > \varepsilon$ with probability $3/5$.

We then replace the condition in (6.7) by the requirement that, for $i \in [d]$,

$$(1/4)m\pi(i) \geq 36C_{\text{IO}} \frac{\|\overline{\mathbf{M}}(i, \cdot)\|_{2/3}}{\varepsilon^2} \left[\ln \left(\frac{36d}{\delta} \cdot \frac{18C_{\text{IO}} \|\overline{\mathbf{M}}(i, \cdot)\|_{2/3}}{\varepsilon^2} \right) - 1 \right], \quad (6.11)$$

whence the theorem. \square

6.3 Proof of Theorem 4.3

The metric domination result in Lemma C.1 immediately implies a lower bound of $\Omega(d/\varepsilon^2)$ (see Remark C.1). We now construct two independent and more delicate lower bounds of $\Omega\left(\frac{\sqrt{d}}{\varepsilon^2\pi_\star}\right)$ and $\Omega(dt_{\text{mix}})$. Let $\mathcal{M}_{d,t_{\text{mix}},\pi_\star}$ be the collection of all d -state Markov chains whose stationary distribution is minorized by π_\star and whose mixing time is at most t_{mix} . Our goal is to lower bound the minimax risk, defined by

$$\mathcal{R}_m = \inf_{\mathcal{T}} \sup_{\overline{\mathbf{M}}, \mathbf{M}} [\mathbf{P}_0(\mathcal{T} = 1) + \mathbf{P}_1(\mathcal{T} = 0)], \quad (6.12)$$

where the inf is over all testing procedures $\mathcal{T} : (X_1, \dots, X_m) \mapsto \{0, 1\}$, and the sup is over all $\overline{\mathbf{M}}, \mathbf{M} \in \mathcal{M}_{d,t_{\text{mix}},\pi_\star}$ such that $\|\overline{\mathbf{M}} - \mathbf{M}\| > \varepsilon$.

Lower bound $\Omega\left(\frac{\sqrt{d}}{\varepsilon^2 p_\star}\right)$. The analysis is simplified by considering $(d+1)$ -state Markov chains with d even; an obvious modification of the proof handles the case of odd d . Fix $0 < p_\star \leq 1/(2(d+1))$ and $0 < \varepsilon < 1/2$, and define $\mathbf{p} \in \Delta_{d+1}$ by

$$\mathbf{p}(d+1) = p_\star \text{ and } \mathbf{p}(i) = (1 - p_\star)/d, \quad i \in [d]. \quad (6.13)$$

Define the collection of $(d+1)$ -state Markov chain transitions matrices,

$$\mathcal{G}_{p_\star} = \left\{ \mathbf{M}_\eta = \begin{pmatrix} \mathbf{p}(1) & \dots & \mathbf{p}(d) & p_\star \\ \vdots & \vdots & \vdots & \vdots \\ \mathbf{p}(1) & \dots & \mathbf{p}(d) & p_\star \\ \boldsymbol{\eta}(1) & \dots & \boldsymbol{\eta}(d) & 0 \end{pmatrix} : \boldsymbol{\eta} \in \Delta_d \right\}. \quad (6.14)$$

The stationary distribution $\boldsymbol{\pi}$ of a chain of this family is given by

$$\boldsymbol{\pi}(i) = \frac{\mathbf{p}(i) + p_\star \boldsymbol{\eta}(i)}{1 + p_\star}, i \neq d+1, \quad \boldsymbol{\pi}(d+1) = \frac{p_\star}{1 + p_\star}, \quad (6.15)$$

and for $p_\star < \frac{d\boldsymbol{\eta}(i)}{d+1}, \forall i \in [d]$, we have $\boldsymbol{\pi}_\star = \boldsymbol{\pi}(d+1)$. We define the n th hitting time for state i as the random variable $H_i^{(n)} := \inf \{t \geq 1 : \sum_{s=1}^t \mathbf{1}\{X_s = i\} = n\}$; in words, this is the first time t at which state i has been visited n times. Suppose that $\mathbf{X} = (X_1, \dots, X_m) \sim (\mathbf{M}, \mathbf{p})$ for some $\mathbf{M} \in \mathcal{G}_{p_\star}$. For any $n \in \mathbb{N}$, the n th hitting time $H_{d+1}^{(n)}$ to state $d+1$, stochastically dominates² the random variable $\sum_{s=1}^n R_s$, where each R_s is an independent copy distributed as Geometric(p_\star). To see this, consider a similar chain where the value 0 in the last row is replaced with p_\star with $\boldsymbol{\eta}$ appropriately re-normalized; clearly, the modification can only make it easier to reach state $d+1$. Continuing, we compute $\mathbf{E}[\sum_{s=1}^n R_s] = n/p_\star$ and $\mathbf{Var}[\sum_{s=1}^n R_s] = n(1 - p_\star)/p_\star^2 \leq n/p_\star^2$. The Paley-Zygmund inequality implies that for $m < n/(2p_\star)$,

$$\begin{aligned} \mathbf{P}\left(H_{d+1}^{(n)} > m\right) &\geq \mathbf{P}\left(\sum_{s=1}^n R_s > m\right) \geq \mathbf{P}\left(\sum_{s=1}^n R_s > \frac{1}{2}\mathbf{E}\left[\sum_{s=1}^n R_s\right]\right) \\ &\geq \left(1 + \frac{4\mathbf{Var}[\sum_{s=1}^n R_s]}{\mathbf{E}[\sum_{s=1}^n R_s]^2}\right)^{-1} \geq 1 - \frac{1}{1 + n/4} \geq \frac{1}{5}. \end{aligned} \quad (6.16)$$

Define the random variable $N_{d+1} = \sum_{t=1}^m \mathbf{1}\{X_t = d+1\}$, i.e. the number of visits to state $d+1$, and consider a reference Markov chain $\overline{\mathbf{M}} := \mathbf{M}_{\bar{\boldsymbol{\eta}}} \in \mathcal{G}_{p_\star}$, where $\bar{\boldsymbol{\eta}} = \mathbf{U}_d$ is the $[d]$ -supported uniform distribution. Restricting the problem to a subset of the family \mathcal{G}_{p_\star} satisfying the ε -separation condition only makes it easier for the tester, as does taking any mixture \mathbf{M}_Σ of chains of this class in lieu of the sup in (6.12). More specifically, we choose

$$\mathbf{M}_\Sigma := \frac{1}{2^{d/2}} \sum_{\boldsymbol{\sigma} \in \{-1, 1\}^{d/2}} \mathbf{M}_\boldsymbol{\sigma}, \quad (6.17)$$

where $\mathbf{M}_\boldsymbol{\sigma}(i, \cdot) = \mathbf{p}$ for $i \in [d]$, and

$$\mathbf{M}_\boldsymbol{\sigma}(d+1, \cdot) = \mathbf{D}_\boldsymbol{\sigma} = \left(\frac{1 + \sigma_1 \varepsilon}{d}, \frac{1 - \sigma_1 \varepsilon}{d}, \dots, \frac{1 + \sigma_{d/2} \varepsilon}{d}, \frac{1 - \sigma_{d/2} \varepsilon}{d}\right). \quad (6.18)$$

By construction, for all $\boldsymbol{\sigma} \in \{-1, 1\}^{d/2}$, we have $\|\mathbf{M}_\boldsymbol{\sigma} - \overline{\mathbf{M}}\| = \|\mathbf{D}_\boldsymbol{\sigma} - \mathbf{U}_d\|_1 = \varepsilon$. We start both chains with distribution \mathbf{p} , defined in (6.13). Now (6.16) implies that for any $n \in \mathbb{N}$ and $m < n/(2p_\star)$, any testing procedure $\mathcal{T} : (X_1, \dots, X_m) \mapsto \{0, 1\}$ verifies

$$\mathbf{P}_{\overline{\mathbf{M}}}(\mathcal{T} = 1) + \mathbf{P}_{\mathbf{M}_\Sigma}(\mathcal{T} = 0) \geq \frac{1}{5} \left(\mathbf{P}_{\overline{\mathbf{M}}}(\mathcal{T} = 1 | N_{d+1} \leq n) + \mathbf{P}_{\mathbf{M}_\Sigma}(\mathcal{T} = 0 | N_{d+1} \leq n) \right). \quad (6.19)$$

²A random variable A stochastically dominates B if $\mathbf{P}(A \geq x) \geq \mathbf{P}(B \geq x)$ for all $x \in \mathbb{R}$.

It follows from Le Cam [2012, Chapter 16, Section 4] that

$$\mathcal{R}_m \geq \frac{1}{5} \left(1 - \left\| \mathbf{P}_{\mathbf{M}_\Sigma}(\mathbf{X}|N_{d+1} \leq n) - \mathbf{P}_{\overline{\mathbf{M}}}(\mathbf{X}|N_{d+1} \leq n) \right\|_{\text{TV}} \right), \quad (6.20)$$

and so it remains to upper bound a total variation distance. For any $\mathbf{M}, \mathbf{M}' \in \mathcal{G}_{p_\star}$, the statistics of the induced state sequence only differ in the visits to state $d+1$.

At this point, we would like to invoke an iid testing lower bound — but are cautioned against doing so naively, as conditioning on the number of visits to a state breaks the Markov property. Instead, in Lemmas A.2, A.3 and A.4 we develop a decoupling technique, which yields

$$\left\| \mathbf{P}_{\mathbf{M}_\Sigma}(\mathbf{X}|N_{d+1} \leq n) - \mathbf{P}_{\overline{\mathbf{M}}}(\mathbf{X}|N_{d+1} \leq n) \right\|_{\text{TV}} \leq \left\| \mathbf{D}_\Sigma^{\otimes n} - \mathbf{U}_d^{\otimes n} \right\|_{\text{TV}}. \quad (6.21)$$

We shall make use of Paninski [2008, Theorem 4], which states:

$$\left\| \mathbf{D}_\Sigma^{\otimes n} - \mathbf{U}_d^{\otimes n} \right\|_{\text{TV}} \leq \sqrt{\exp\left(\frac{n^2 \varepsilon^4}{d}\right)}. \quad (6.22)$$

It follows that

$$\mathcal{R}_m \geq \frac{1}{5} \left(1 - \frac{1}{2} \sqrt{\exp\left(\frac{n^2 \varepsilon^4}{d}\right) - 1} \right). \quad (6.23)$$

Finally, for the mixture of chains and parameter regime in question, we have $\boldsymbol{\pi}(d+1) = \pi_\star$ and $p_\star/2 \leq \pi_\star \leq p_\star$, so that for $\delta < 1/10$, $m < \frac{n}{2p_\star}$ and $n \leq \frac{\sqrt{d}}{\varepsilon^2} \sqrt{\ln\left(\frac{4}{C^2}(1-5\delta)^2\right)}$, it follows that $\mathcal{R}_m \geq \delta$. This implies a lower bound of $m = \Omega\left(\frac{\sqrt{d}}{\varepsilon^2 \pi_\star}\right)$ for the testing problem.

Lower bound $\Omega(dt_{\text{mix}})$. Let us recall the construction of Wolfer and Kontorovich [2019]. Taking $0 < \varepsilon \leq 1/8$ and $d = 6k$, $k \geq 2$ fixed, $0 < \eta < 1/48$ and $\boldsymbol{\tau} \in \{0, 1\}^{d/3}$, we define the block matrix

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

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

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

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

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

Holding η fixed, define the collection

$$\mathcal{H}_\eta = \left\{ \mathbf{M}_{\eta, \boldsymbol{\tau}} : \boldsymbol{\tau} \in \{0, 1\}^{d/3} \right\} \quad (6.24)$$

of ergodic and symmetric stochastic matrices. Suppose that $\mathbf{X} = (X_1, \dots, X_m) \sim (\mathbf{M}, \boldsymbol{\mu})$, where $\mathbf{M} \in \mathcal{H}_\eta$, and $\boldsymbol{\mu}$ is the uniform distribution over the inner clique nodes, indexed by $\{1, \dots, d/3\}$. Define the random variable $T_{\text{CLIQ}/2}$ to be the first time some half of the states in the inner clique were visited,

$$T_{\text{CLIQ}/2} = \inf \{t \geq 1 : |\{X_1, \dots, X_t\} \cap [d/3]| = d/6\}. \quad (6.25)$$

Lemma A.5 lower bounds the half cover time:

$$m \leq \frac{d}{120\eta} \implies \mathbf{P}(T_{\text{CLIQ}/2} > m) \geq \frac{1}{5}, \quad (6.26)$$

while Wolfer and Kontorovich [2019, Lemma 6] establishes the key property that any element \mathbf{M} of \mathcal{H}_η satisfies

$$t_{\text{mix}}(\mathbf{M}) = \tilde{\Theta}(1/\eta). \quad (6.27)$$

Let us fix some $i_\star \in [d]$, choose as reference $\overline{\mathbf{M}} := \mathbf{M}_{\eta, \mathbf{0}}$ and as an alternative hypothesis $\mathbf{M} := \mathbf{M}_{\eta, \boldsymbol{\tau}}$, with $\tau_i = \mathbf{1}\{i = i_\star\}$. Take both chains to have the uniform distribution $\boldsymbol{\mu}$ over the clique nodes as their initial one. It is easily verified that $\|\overline{\mathbf{M}} - \mathbf{M}\| = \varepsilon$, so that

$$\mathcal{R}_m \geq \inf_{\mathcal{T}} [\mathbf{P}_0(\mathcal{T} = 1 | T_{\text{CLIQ}/2} > m) \mathbf{P}_0(T_{\text{CLIQ}/2} > m) + \mathbf{P}_1(\mathcal{T} = 0 | T_{\text{CLIQ}/2} > m) \mathbf{P}_1(T_{\text{CLIQ}/2} > m)]. \quad (6.28)$$

Further, for $m < \frac{d}{120\eta}$, we have

$$\mathcal{R}_m \geq \frac{1}{5} \inf_{\mathcal{T}} [\mathbf{P}_0(\mathcal{T} = 1 | T_{\text{CLIQ}/2} > m) + \mathbf{P}_1(\mathcal{T} = 0 | T_{\text{CLIQ}/2} > m)]. \quad (6.29)$$

Since $\mathbf{P}(X|Y) \geq \mathbf{P}(X|Y, Z) \mathbf{P}(Z|Y)$, we have

$$\mathbf{P}_0(\mathcal{T} = 1 | T_{\text{CLIQ}/2} > m) \geq \mathbf{P}_0(\mathcal{T} = 1 | T_{\text{CLIQ}/2} > m, N_{i_\star} = 0) \mathbf{P}_0(N_{i_\star} = 0 | T_{\text{CLIQ}/2} > m). \quad (6.30)$$

Additionally, the symmetry of our reference chain implies that $\mathbf{P}_0(N_{i_\star} = 0 | T_{\text{CLIQ}/2} > m) \geq 1/2$. It follows, via an analogous argument that $\mathbf{P}_1(N_{i_\star} = 0 | T_{\text{CLIQ}/2} > m) \geq 1/2$, so that

$$\mathcal{R}_m \geq \frac{1}{10} \inf_{\mathcal{T}} [\mathbf{P}_0(\mathcal{T} = 1 | T_{\text{CLIQ}/2} > m, N_{i_\star} = 0) + \mathbf{P}_1(\mathcal{T} = 0 | T_{\text{CLIQ}/2} > m, N_{i_\star} = 0)]. \quad (6.31)$$

By Le Cam's theorem [Le Cam, 2012, Chapter 16, Section 4],

$$\mathcal{R}_m = \frac{1}{10} \left[1 - \|\mathbf{P}_0(\mathbf{X} | T_{\text{CLIQ}/2} > m, N_{i_\star} = 0) - \mathbf{P}_1(\mathbf{X} | T_{\text{CLIQ}/2} > m, N_{i_\star} = 0)\|_{\text{TV}} \right]. \quad (6.32)$$

Other than state i_\star and its connected outer nodes, the reference chain $\mathbf{M}_{\eta, \mathbf{0}}$ and the alternative chain $\mathbf{M}_{\eta, \boldsymbol{\tau}}$ are identical. Conditional on $N_{i_\star} = 0$, the outer states connected to i_\star were never visited, since these are only connected to the rest of the chain via i_\star and our choice of the initial distribution $\boldsymbol{\mu}$ constrains the initial state to the inner clique. Thus, the two distributions over sequences conditioned on $N_{i_\star} = 0$ are identical, causing the term $\|\mathbf{P}_0(\cdot) - \mathbf{P}_1(\cdot)\|_{\text{TV}}$ in (6.32) to vanish:

$$\mathcal{R}_m \geq \frac{1}{10}, \quad (6.33)$$

which proves a sample complexity lower bound of $\tilde{\Omega}(dt_{\text{mix}})$. Since our family of Markov chains has uniform stationary distribution ($\pi_\star = 1/d$), this further proves that the dependence on π_\star in our bound is in general not improvable. \square

Acknowledgments

We are thankful to John Lafferty for bringing this problem to our attention and numerous insightful conversations.

References

- M. Anthony and P. 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>.
- F. Barsotti, A. Philippe, and P. Rochet. Hypothesis testing for markovian models with random time observations. *Journal of Statistical Planning and Inference*, 173:87–98, 2016.
- T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White. Testing that distributions are close. In *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA*, pages 259–269, 2000. doi: 10.1109/SFCS.2000.892113. URL <https://doi.org/10.1109/SFCS.2000.892113>.
- T. Batu, E. Fischer, L. Fortnow, R. Kumar, R. Rubinfeld, and P. White. Testing random variables for independence and identity. In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, pages 442–451. IEEE, 2001.
- C. Daskalakis, N. Dikkala, and N. Gravin. Testing symmetric markov chains from a single trajectory. *arXiv preprint arXiv:1704.06850*, 2017.
- I. Diakonikolas, T. Gouleakis, J. Peebles, and E. Price. Collision-based testers are optimal for uniformity and closeness. *arXiv preprint arXiv:1611.03579*, 2016.
- I. Diakonikolas, T. Gouleakis, J. Peebles, and E. Price. Sample-optimal identity testing with high probability. *CoRR*, abs/1708.02728, 2017. URL <http://arxiv.org/abs/1708.02728>.
- J. A. Fill. Eigenvalue bounds on convergence to stationarity for nonreversible markov chains, with an application to the exclusion process. *The annals of applied probability*, pages 62–87, 1991.
- O. Goldreich. The uniform distribution is complete with respect to testing identity to a fixed distribution. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 23, page 1, 2016.
- O. Goldreich and D. Ron. On testing expansion in bounded-degree graphs. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, pages 68–75. Springer, 2011.
- Y. Hao, A. Orlitsky, and V. Pichapati. On learning markov chains. *arXiv preprint arXiv:1810.11754*, 2018.
- D. Kazakos. The Bhattacharyya distance and detection between markov chains. *IEEE Transactions on Information Theory*, 24(6):747–754, 1978.
- A. Kontorovich and I. Pinelis. Exact lower bounds for the agnostic probably-approximately-correct (PAC) machine learning model. *Ann. Statist.*, 2019+.
- L. Le Cam. *Asymptotic methods in statistical decision theory*. Springer Science & Business Media, 2012.

- F. Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th international symposium on symbolic and algebraic computation*, pages 296–303. ACM, 2014.
- D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov chains and mixing times, second edition*. American Mathematical Soc., 2009.
- C. Meyer. *Matrix analysis and applied linear algebra*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2000. ISBN 0-89871-454-0. URL <https://doi.org/10.1137/1.9780898719512>. With 1 CD-ROM (Windows, Macintosh and UNIX) and a solutions manual (iv+171 pp.).
- L. Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Transactions on Information Theory*, 54(10):4750–4755, 2008.
- D. Paulin. Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electronic Journal of Probability*, 20, 2015.
- G. Valiant and P. Valiant. An automatic inequality prover and instance optimal identity testing. *SIAM Journal on Computing*, 46(1):429–455, 2017.
- B. Waggoner. L_p testing and learning of discrete distributions. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS 2015, Rehovot, Israel, January 11-13, 2015*, pages 347–356, 2015. doi: 10.1145/2688073.2688095. URL <http://doi.acm.org/10.1145/2688073.2688095>.
- G. Wolfer and A. Kontorovich. Minimax learning of ergodic markov chains. In *Algorithmic Learning Theory, ALT*, 2019. URL <https://arxiv.org/abs/1809.05014>.

A Auxiliary lemmas


The following standard combinatorial fact will be useful.

Lemma A.1 *Let $(m, n) \in \mathbb{N}^2$ such that $m + 1 \geq 2n$. Then there are $\binom{m-n+1}{n}$ ways of selecting n non-consecutive integers from $[m]$.*

Lemma A.2 *For any $M \in \mathcal{G}_{p_\star}$ defined in (6.14), started with initial distribution \mathbf{p} defined in (6.13),*

$$P(m, n, p_\star) \triangleq \mathbf{P}_{M, \mathbf{p}}(N_{d+1} = n) = \begin{cases} (1 - p_\star)^n & \text{if } n = 0 \\ p_\star^n (1 - p_\star)^{m-2n} \left[\binom{m-n+1}{n} - \binom{m-n}{n-1} p_\star \right] & \text{if } 1 \leq n \leq \frac{m+1}{2} \\ 0 & \text{if } n > \frac{m+1}{2} \end{cases}.$$

Proof: For any $M \in \mathcal{G}_{p_\star}$, we construct the following associated two-state Markov chain, with initial distribution $(1 - p_\star, p_\star)$, where all states $i \in [d]$ are merged into a single state, which we call $\overline{d+1}$, while state $d+1$ is kept distinct. Observe that this two-state Markov chain is the same for all $M \in \mathcal{G}_{p_\star}$, regardless of $\boldsymbol{\eta}$, and that the probability distribution of the number of visits to state $d+1$, when sampling from M , is the same as when sampling from this newly constructed chain.



Let $m \geq 1$. The case where $n = 0$ is trivial as it corresponds to n failures to reach the state $d + 1$, and there is only one such path. When $n > \frac{m+1}{2}$, there is no path of length m that contains n visits to state $d + 1$, as any visit to this state almost surely cannot be directly followed by another visit to this same state. It remains to analyze the final case where $1 \leq n \leq \frac{m+1}{2}$. Take (x_1, \dots, x_m) to be a sample path in which the state $d + 1$ was visited n times. We consider two sub-cases.

The last state in the sample path is $d + 1$: In the case where $x_m = d + 1$, note that also necessarily $x_{m-1} = \overline{d + 1}$. The $n - 1$ previous visits to state $d + 1$ were followed by a probability 1 transition to state $\overline{d + 1}$, and the remaining transitions have value $1 - p_*$, so that $\mathbf{P}(\mathbf{X} = \mathbf{x}) = p_*^{n-1} 1^{n-1} (1 - p_*)^{m-(n-1)-(n-1)-1} p_* = p_*^n (1 - p_*)^{m-2n+1}$. Since the last two states in the sequence are fixed and known, $x_m = d + 1$, $x_{m-1} = \overline{d + 1}$, counting the number of such paths amounts to counting the number of subsets of $m - 2$ of size $n - 1$ such that no two elements are consecutive, i.e. $\binom{(m-2)-(n-1)+1}{n-1} = \binom{m-n}{n-1}$ (Lemma A.1).

The last state in the sample path is $\overline{d + 1}$: By reasoning similar to above, such paths have probability $p_*^n (1 - p_*)^{m-2n}$. To count such paths, consider all possible subsets of m of size n such that no two elements are consecutive, and subtract the count of paths in the other case where the last state was $d + 1$. There are then $\binom{m-n+1}{n} - \binom{m-n}{n-1} = \binom{m-n}{n}$ such paths.

It follows that

$$\begin{aligned} P(m, n, p_*) &= \binom{m-n}{n-1} p_*^n (1 - p_*)^{m-2n+1} + \binom{m-n}{n} p_*^n (1 - p_*)^{m-2n} \\ &= p_*^n (1 - p_*)^{m-2n} \left[\binom{m-n+1}{n} - \binom{m-n}{n-1} p_* \right]. \end{aligned} \quad (\text{A.1})$$

□

Lemma A.3 Let $\mathbf{M}_1, \mathbf{M}_2 \in \mathcal{G}_{p_*}$, defined in (6.14), and start both chains with initial distribution \mathbf{p} defined in (6.13). For arbitrary $(m, n) \in \mathbb{N}^2$ such that $m + 1 \geq 2n$, let $N_{d+1} = \sum_{t=1}^m \mathbf{1}\{X_t = d + 1\}$ be the number of visits to state $(d + 1)$. Then, for trajectories $\mathbf{X} = (X_1, \dots, X_m)$ sampled from either chain, we have

$$\begin{aligned} &\|\mathbf{P}_{\mathbf{M}_1}(\mathbf{X} \mid N_{d+1} \leq n) - \mathbf{P}_{\mathbf{M}_2}(\mathbf{X} \mid N_{d+1} \leq n)\|_{\text{TV}} \\ &\leq \|\mathbf{P}_{\mathbf{M}_1}(\mathbf{X} \mid N_{d+1} = n) - \mathbf{P}_{\mathbf{M}_2}(\mathbf{X} \mid N_{d+1} = n)\|_{\text{TV}}. \end{aligned} \quad (\text{A.2})$$

Proof: Partitioning over all possible number of visits to $d + 1$ for \mathbf{M}_1 ,

$$\begin{aligned} \mathbf{P}_{\mathbf{M}_1}(\mathbf{X} \mid N_{d+1} \leq n) &= \sum_{k=0}^{\infty} \mathbf{P}_{\mathbf{M}_1}(\mathbf{X} \mid N_{d+1} \leq n, N_{d+1} = k) \mathbf{P}_{\mathbf{M}_1}(N_{d+1} = k \mid N_{d+1} \leq n) \\ &= \sum_{k=0}^n \mathbf{P}_{\mathbf{M}_1}(\mathbf{X} \mid N_{d+1} = k) \mathbf{P}_{\mathbf{M}_1}(N_{d+1} = k \mid N_{d+1} \leq n) \\ &= \sum_{k=0}^n \mathbf{P}_{\mathbf{M}_1}(\mathbf{X} \mid N_{d+1} = k) \frac{\overbrace{\mathbf{P}_{\mathbf{M}_1}(N_{d+1} \leq n \mid N_{d+1} = k)}^{=1} \mathbf{P}_{\mathbf{M}_1}(N_{d+1} = k)}{\mathbf{P}_{\mathbf{M}_1}(N_{d+1} \leq n)} \end{aligned} \quad (\text{A.3})$$

From Lemma A.2, we have

$$\begin{aligned} \mathbf{P}_{\mathbf{M}_1}(N_{d+1} = k) &= \mathbf{P}_{\mathbf{M}_2}(N_{d+1} = k) = P(m, k, p_*) \\ \mathbf{P}_{\mathbf{M}_1}(N_{d+1} \leq n) &= \mathbf{P}_{\mathbf{M}_2}(N_{d+1} \leq n) = \sum_{s=0}^n P(m, s, p_*), \end{aligned} \quad (\text{A.4})$$

and subsequently,

$$\begin{aligned}
& \|\mathbf{P}_{M_1}(\mathbf{X} \mid N_{d+1} \leq n) - \mathbf{P}_{M_2}(\mathbf{X} \mid N_{d+1} \leq n)\|_{\text{TV}} \\
&= \left\| \sum_{k=0}^n \left(\mathbf{P}_{M_1}(\mathbf{X} \mid N_{d+1} = k) - \mathbf{P}_{M_2}(\mathbf{X} \mid N_{d+1} = k) \right) \frac{P(m, k, p_\star)}{\sum_{s=0}^n P(m, s, p_\star)} \right\|_{\text{TV}} \\
&\leq \sum_{k=0}^n \|\mathbf{P}_{M_1}(\mathbf{X} \mid N_{d+1} = k) - \mathbf{P}_{M_2}(\mathbf{X} \mid N_{d+1} = k)\|_{\text{TV}} \frac{P(m, k, p_\star)}{\sum_{s=0}^n P(m, s, p_\star)} \\
&\leq \max_{k \in \{0, \dots, n\}} \|\mathbf{P}_{M_1}(\mathbf{X} \mid N_{d+1} = k) - \mathbf{P}_{M_2}(\mathbf{X} \mid N_{d+1} = k)\|_{\text{TV}} \\
&\leq \|\mathbf{P}_{M_1}(\mathbf{X} \mid N_{d+1} = n) - \mathbf{P}_{M_2}(\mathbf{X} \mid N_{d+1} = n)\|_{\text{TV}}.
\end{aligned} \tag{A.5}$$

□

The following lemma shows that for the family \mathcal{G}_{p_\star} of chains constructed in (6.14), conditioned on the number of visits to state $d+1$, it is possible to control the total variation between two trajectories drawn from two chains of the class in terms of the total variation between product distributions.

Lemma A.4 *Let $M_{\eta_1}, M_{\eta_2} \in \mathcal{G}_{p_\star}$ defined in (6.14), both started with initial distribution \mathbf{p} defined in (6.13). Then, for $1 \leq n \leq \frac{m+1}{2}$,*

$$\|\mathbf{P}_{M_{\eta_1}}(\mathbf{X} \mid N_{d+1} = n) - \mathbf{P}_{M_{\eta_2}}(\mathbf{X} \mid N_{d+1} = n)\|_{\text{TV}} \leq \|\boldsymbol{\eta}_1^{\otimes n} - \boldsymbol{\eta}_2^{\otimes n}\|_{\text{TV}}. \tag{A.6}$$

Proof: Total variation and ℓ_1 norm are equal up to a conventional factor of 2,

$$\begin{aligned}
& 2 \|\mathbf{P}_{M_{\eta_1}}(\mathbf{X} \mid N_{d+1} = n) - \mathbf{P}_{M_{\eta_2}}(\mathbf{X} \mid N_{d+1} = n)\|_{\text{TV}} \\
&= \sum_{\mathbf{x}=(x_1, \dots, x_m) \in [d+1]^m} |\mathbf{P}_{M_{\eta_1}}(\mathbf{X} = \mathbf{x} \mid N_{d+1} = n) - \mathbf{P}_{M_{\eta_2}}(\mathbf{X} = \mathbf{x} \mid N_{d+1} = n)|.
\end{aligned} \tag{A.7}$$

Notice now that

$$\begin{aligned}
\mathbf{P}_{M_{\eta_1}}(\mathbf{X} = \mathbf{x} \mid N_{d+1} = n) &= \frac{\mathbf{P}_{M_{\eta_1}}(N_{d+1} = n \mid \mathbf{X} = \mathbf{x}) \mathbf{P}_{M_{\eta_1}}(\mathbf{X} = \mathbf{x})}{\mathbf{P}_{M_{\eta_1}}(N_{d+1} = n)} \\
&= \frac{\mathbf{1}\{n_{d+1} = n\} \mathbf{P}_{M_{\eta_1}}(\mathbf{X} = \mathbf{x})}{\mathbf{P}_{M_{\eta_1}}(N_{d+1} = n)},
\end{aligned} \tag{A.8}$$

and similarly for M_{η_2} , so that

$$\begin{aligned}
& 2 \|\mathbf{P}_{M_{\eta_1}}(\mathbf{X} \mid N_{d+1} = n) - \mathbf{P}_{M_{\eta_2}}(\mathbf{X} \mid N_{d+1} = n)\|_{\text{TV}} \\
&= \sum_{\mathbf{x} \in [d+1]^m} \left| \frac{\mathbf{1}\{n_{d+1} = n\} \mathbf{P}_{M_{\eta_1}}(\mathbf{X} = \mathbf{x})}{\mathbf{P}_{M_{\eta_1}}(N_{d+1} = n)} - \frac{\mathbf{1}\{n_{d+1} = n\} \mathbf{P}_{M_{\eta_2}}(\mathbf{X} = \mathbf{x})}{\mathbf{P}_{M_{\eta_2}}(N_{d+1} = n)} \right|.
\end{aligned} \tag{A.9}$$

Invoking Lemma A.2, write

$$P(m, n, p_\star) = \mathbf{P}_{M_{\eta_1}}(N_{d+1} = n) = \mathbf{P}_{M_{\eta_2}}(N_{d+1} = n).$$

For $1 \leq n \leq \frac{m+1}{2}$,

$$\begin{aligned}
& 2 \left\| \mathbf{P}_{M_{\eta_1}}(\mathbf{X} | N_{d+1} = n) - \mathbf{P}_{M_{\eta_2}}(\mathbf{X} | N_{d+1} = n) \right\|_{\text{TV}} \\
&= \frac{1}{P(m, n, p_\star)} \sum_{\mathbf{x} \in [d+1]^m} \mathbf{1}\{n_{d+1} = n\} \left| \mathbf{P}_{M_{\eta_1}}(\mathbf{X} = \mathbf{x}) - \mathbf{P}_{M_{\eta_2}}(\mathbf{X} = \mathbf{x}) \right| \\
&= \frac{1}{P(m, n, p_\star)} \left(\sum_{\substack{\mathbf{x} \in [d+1]^m \\ n_{d+1} = n \\ x_m = d+1}} \left| \mathbf{P}_{M_{\eta_1}}(\mathbf{X} = \mathbf{x}) - \mathbf{P}_{M_{\eta_2}}(\mathbf{X} = \mathbf{x}) \right| \right. \\
&\quad \left. + \sum_{\substack{\mathbf{x} \in [d+1]^m \\ n_{d+1} = n \\ x_m \neq d+1}} \left| \mathbf{P}_{M_{\eta_1}}(\mathbf{X} = \mathbf{x}) - \mathbf{P}_{M_{\eta_2}}(\mathbf{X} = \mathbf{x}) \right| \right). \tag{A.10}
\end{aligned}$$

Recall that it is impossible to visit state $d+1$ twice in a row. Computing the first sum,

$$\begin{aligned}
& \sum_{\substack{\mathbf{x} \in [d+1]^m \\ n_{d+1} = n \\ x_m = d+1}} \left| \mathbf{P}_{M_{\eta_1}}(\mathbf{X} = \mathbf{x}) - \mathbf{P}_{M_{\eta_2}}(\mathbf{X} = \mathbf{x}) \right| \\
&= \sum_{\substack{S = (s_1, \dots, s_n) \\ S \subset [m] \\ s_n = m \\ i \neq j \Rightarrow |s_i - s_j| > 1}} \sum_{(x_{s_1}, \dots, x_{s_{n-1}}) \in [d+1]^{n-1}} d^{m-2n+1} p_\star^n \left(\frac{1-p_\star}{d} \right)^{m-2n+1} \left| \prod_{k=1}^{n-1} \eta_1(x_{s_k}) - \prod_{k=1}^{n-1} \eta_2(x_{s_k}) \right| \\
&= \binom{(m-2) - (n-1) + 1}{n-1} p_\star^n (1-p_\star)^{m-2n+1} \sum_{(x_{s_1}, \dots, x_{s_{n-1}}) \in [d]^{n-1}} \left| \prod_{k=1}^{n-1} \eta_1(x_{s_k}) - \prod_{k=1}^{n-1} \eta_2(x_{s_k}) \right| \\
&= \binom{m-n}{n-1} p_\star^n (1-p_\star)^{m-2n+1} 2 \left\| \boldsymbol{\eta}_1^{\otimes n-1} - \boldsymbol{\eta}_2^{\otimes n-1} \right\|_{\text{TV}}, \tag{A.11}
\end{aligned}$$

where the second inequality is from Lemma A.1. Similarly for the second sum,

$$\begin{aligned}
& \sum_{\substack{\mathbf{x} \in [d+1]^m \\ n_{d+1} = n \\ x_m \neq d+1}} \left| \mathbf{P}_{M_{\eta_1}}(\mathbf{X} = \mathbf{x}) - \mathbf{P}_{M_{\eta_2}}(\mathbf{X} = \mathbf{x}) \right| \\
&= \sum_{\substack{S = (s_1, \dots, s_n) \\ S \subset [m] \\ s_n \neq m \\ i \neq j \Rightarrow |s_i - s_j| > 1}} \sum_{(x_{s_1}, \dots, x_{s_n}) \in [d+1]^n} d^{m-2n} p_\star^n \left(\frac{1-p_\star}{d} \right)^{m-2n} \left| \prod_{k=1}^n \eta_1(x_{s_k}) - \prod_{k=1}^n \eta_2(x_{s_k}) \right| \\
&= \left(\binom{m-n+1}{n} - \binom{m-n}{n-1} \right) p_\star^n (1-p_\star)^{m-2n} \sum_{(x_{s_1}, \dots, x_{s_n}) \in [d]^n} \left| \prod_{k=1}^n \eta_1(x_{s_k}) - \prod_{k=1}^n \eta_2(x_{s_k}) \right| \\
&= \binom{m-n}{n} p_\star^n (1-p_\star)^{m-2n} 2 \left\| \boldsymbol{\eta}_1^{\otimes n} - \boldsymbol{\eta}_2^{\otimes n} \right\|_{\text{TV}}. \tag{A.12}
\end{aligned}$$

Hence,

$$\begin{aligned}
& 2 \left\| \mathbf{P}_{\mathbf{M}_{\eta_1}}(\mathbf{X} | N_{d+1} = n) - \mathbf{P}_{\mathbf{M}_{\eta_2}}(\mathbf{X} | N_{d+1} = n) \right\|_{\text{TV}} P(m, n, p_\star) \\
&= 2p_\star^n (1 - p_\star)^{m-2n} \left[\binom{m-n}{n-1} (1 - p_\star) \left\| \boldsymbol{\eta}_1^{\otimes n-1} - \boldsymbol{\eta}_2^{\otimes n-1} \right\|_{\text{TV}} + \binom{m-n}{n} \left\| \boldsymbol{\eta}_1^{\otimes n} - \boldsymbol{\eta}_2^{\otimes n} \right\|_{\text{TV}} \right] \\
&\leq 2p_\star^n \left\| \boldsymbol{\eta}_1^{\otimes n} - \boldsymbol{\eta}_2^{\otimes n} \right\|_{\text{TV}} (1 - p_\star)^{m-2n} \left[\binom{m-n}{n-1} (1 - p_\star) + \binom{m-n}{n} \right] \\
&= 2 \left\| \boldsymbol{\eta}_1^{\otimes n} - \boldsymbol{\eta}_2^{\otimes n} \right\|_{\text{TV}} P(m, n, p_\star) \quad (\text{Lemma A.1}).
\end{aligned} \tag{A.13}$$

□

Lemma A.5 (Cover time) For $\mathbf{M} \in \mathcal{H}_\eta$ [defined in (6.24)], the “half cover time” random variable $T_{\text{CLIQ}/2}$ [defined in (6.25)] satisfies

$$m \leq \frac{d}{120\eta} \implies \mathbf{P}(T_{\text{CLIQ}/2} > m) \geq \frac{1}{5}. \tag{A.14}$$

Proof: The proof pursues a strategy similar to Wolfer and Kontorovich [2019], which is adapted to “half” rather than “full” coverings. Let $\mathbf{M} \in \mathcal{H}_\eta$ and $\mathbf{M}_I \in \mathcal{M}_{d/3}$ be such that \mathbf{M}_I consists only in the inner clique of \mathbf{M} , and each outer rim state got absorbed into its unique inner clique neighbor:

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

By construction, it is clear that $T_{\text{CLIQ}/2}$ is almost surely greater than the half cover time of \mathbf{M}_I . The latter corresponds to a generalized coupon half collection time $U_{\text{COVER}/2} = 1 + \sum_{i=1}^{d/6-1} U_i$ where U_i is the time increment between the i th and the $(i+1)$ th unique visited state. Formally, if \mathbf{X} is a random walk according to \mathbf{M}_I (started from any state), then $U_1 = \min\{t > 1 : X_t \neq X_1\}$ and for $i > 1$,

$$U_i = \min\{t > 1 : X_t \notin \{X_1, \dots, X_{U_{i-1}}\}\} - U_{i-1}. \tag{A.15}$$

The random variables $U_1, U_2, \dots, U_{d/6-1}$ are independent and $U_i \sim \text{Geometric}\left(\eta - \frac{(i-1)\eta}{d/3}\right)$, whence

$$\mathbf{E}[U_i] = \frac{d/3}{\eta(d/3 - i + 1)}, \quad \mathbf{Var}[U_i] = \frac{1 - \left(\eta - \frac{(i-1)\eta}{d/3}\right)}{\left(\eta - \frac{(i-1)\eta}{d/3}\right)^2} \tag{A.16}$$

and

$$\mathbf{E}[U_{\text{COVER}/2}] \geq 1 + \frac{d/3}{\eta}(\sigma_{d/3} - \sigma_{d/6}), \quad \mathbf{Var}[U_{\text{COVER}/2}] \leq \frac{(d/3)^2 \pi^2}{\eta^2 6}, \tag{A.17}$$

where $\sigma_d = \sum_{i=1}^d \frac{1}{i}$, and $\pi = 3.1416\dots$. Since $\ln(d+1) \leq \sigma_d \leq 1 + \ln d$, and for $d = 6k, k \geq 2$, we have $\sigma_d - \sigma_{d/2} \geq \ln 2$ it follows that

$$\mathbf{E}[U_{\text{COVER}/2}] \geq \frac{d \ln 2}{\eta 3}, \quad \mathbf{Var}[U_{\text{COVER}/2}] \leq \frac{d^2 \pi^2}{\eta^2 54}. \tag{A.18}$$

Invoking the Paley-Zygmund inequality with $\theta = 1 - \frac{\sqrt{15}}{6\ln 2}$, yields

$$\mathbf{P}(U_{\text{COVER}/2} > \theta \mathbf{E}[U_{\text{COVER}/2}]) \geq \left(1 + \frac{\text{Var}[U_{\text{COVER}/2}]}{(1 - \theta)^2 (\mathbf{E}[U_{\text{COVER}/2}])^2}\right)^{-1} \geq \frac{1}{5}, \quad (\text{A.19})$$

so that for $m \leq \frac{d}{120\eta}$ we have $\mathbf{P}(T_{\text{CLIQ}/2} > m) \geq \frac{1}{5}$. \square

B Algorithm

B.1 Pseudo-code implementation

```

Input:  $d, \varepsilon, \delta, \overline{\mathbf{M}}, \boldsymbol{\pi}, (X_1, \dots, X_m)$ 
Output: ACCEPT = 0 or REJECT = 1
Visits  $\leftarrow \mathbf{0} \in \mathbb{R}^d$ 
for  $t \leftarrow 1$  to  $m - 1$  do
  | Visits( $X_t$ )  $\leftarrow$  Visits( $X_t$ ) + 1
end
for  $i \leftarrow 1$  to  $d$  do
  | if  $|\text{Visits}(i) - m\boldsymbol{\pi}(i)| > (1/2)m\boldsymbol{\pi}(i)$  then
  | | return REJECT
  | end
end
for  $i \leftarrow 1$  to  $d$  do
  | Transitions  $\leftarrow \mathbf{0} \in \mathbb{R}^d$ 
  | for  $t \leftarrow 1$  to  $m - 1$  do
  | | Transitions( $j$ )  $\leftarrow$  Transitions( $j$ ) +  $\mathbf{1}\{X_t = i\} \mathbf{1}\{X_{t+1} = j\}$ 
  | end
  | if DistributionIIDIdentityTester( $\varepsilon, d, \overline{\mathbf{M}}(i, \cdot), \text{Transitions}$ ) is REJECT then
  | | return REJECT
  | end
end
return ACCEPT

```

Algorithm 1: The tester \mathcal{T}

B.2 Computational complexity

Time complexity of designing a test. In order to determine the size of the sample required in order to reach desired proximity ε and confidence $1 - \delta$, and to run the test algorithm, one must first compute $\boldsymbol{\pi}$ and t_{mix} . These two quantities are relative to the reference Markov chain, for which we assume to have a full description of the transition matrix $\overline{\mathbf{M}}$. Interestingly they can be computed *offline*, only once per test definition, no matter how many observed trajectories we will want to test.

The eigenproblem $\boldsymbol{\pi} \overline{\mathbf{M}} = \boldsymbol{\pi}$ can be solved in time $\tilde{O}(d^\omega)$, where $2 \leq \omega \leq 2.3728639$ [Le Gall, 2014]; this is a method for recovering $\boldsymbol{\pi}$.

In the reversible case, it is well known [Levin et al., 2009] that the mixing time of an ergodic Markov chain is controlled by its *absolute spectral gap* $\gamma_\star \triangleq 1 - \max\{\lambda_2, |\lambda_d|\}$, where $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d$ is the ordered spectrum of $\overline{\mathbf{M}}$, and minimum stationary probability π_\star :

$$\left(\frac{1}{\gamma_\star} - 1\right) \ln 2 \leq t_{\text{mix}} \leq \frac{\ln(4/\pi_\star)}{\gamma_\star}. \quad (\text{B.1})$$

The full eigen-decomposition used to obtain $\boldsymbol{\pi}$ of cost $\tilde{O}(d^\omega)$ also yields λ_\star .

In the non-reversible case, the relationship between the spectrum and the mixing time is not nearly as straightforward, and it is the *pseudo-spectral gap* [Paulin, 2015],

$$\gamma_{\text{ps}} \triangleq \max_{k \in \mathbb{N}} \left\{ \frac{\gamma((\overline{M}^\dagger)^k \overline{M}^k)}{k} \right\}, \quad (\text{B.2})$$

where \overline{M}^\dagger is the time reversal [Fill, 1991] of \overline{M} , that gives effective control:

$$\frac{1}{2\gamma_{\text{ps}}} \leq t_{\text{mix}} \leq \frac{1}{\gamma_{\text{ps}}} \left(\ln \frac{1}{\pi_\star} + 2 \ln 2 + 1 \right). \quad (\text{B.3})$$

For any $k \in \mathbb{N}$, we have $\frac{\gamma((\overline{M}^\dagger)^k \overline{M}^k)}{k} \leq 1/k$, and so γ_{ps} is computable to within an additive error of η in time $\mathcal{O}(d^3/\eta)$.

The instance-specific test requires computing the d values $\|\overline{M}(i, \cdot)\|_{2/3}$, which is feasible in time $\mathcal{O}(d^2)$.

Time and space complexity of performing a test. The first step is to verify whether each state has been visited a sufficient amount of times $n_i, i \in [d]$, which can be achieved in time $\mathcal{O}(m+d)$ and space $\mathcal{O}(d)$. The second step is to proceed with the state-wise testing strategy, whose runtime will depend on the black-box iid tester. For instance, the test of Valiant and Valiant [2017] can be performed in $\mathcal{O}(m+d)$ operations. Using this as the iid tester, we get an overall time complexity of $\mathcal{O}(d(d+m))$.

C Comparison with existing work

Let us compare our distance $\|\cdot\|$ to the one of Daskalakis et al. [2017]:

$$\text{Dist}_{\sqrt{\cdot}}(\mathbf{M}, \mathbf{M}') = 1 - \rho([\mathbf{M}, \mathbf{M}']_{\sqrt{\cdot}}), \quad (\text{C.1})$$

where $([\mathbf{M}, \mathbf{M}']_{\sqrt{\cdot}})(i, j) = \left[\sqrt{\mathbf{M}(i, j)\mathbf{M}'(i, j)} \right]$ and ρ is to the spectral radius.

Lemma C.1 *For all $\mathbf{M}, \mathbf{M}' \in \mathcal{M}_d$, $\|\mathbf{M} - \mathbf{M}'\| \geq 2\text{Dist}_{\sqrt{\cdot}}(\mathbf{M}, \mathbf{M}')$. Conversely, for $d \geq 4$, there exist $\mathbf{M}, \mathbf{M}' \in \mathcal{M}_d$ such that $\text{Dist}_{\sqrt{\cdot}}(\mathbf{M}, \mathbf{M}') = 0$ while $\|\mathbf{M} - \mathbf{M}'\| = 1$.*

Proof: For $\boldsymbol{\mu}, \boldsymbol{\nu} \in \Delta_d$, define the Hellinger distance

$$H(\boldsymbol{\mu}, \boldsymbol{\nu}) = \frac{1}{\sqrt{2}} \sqrt{\sum_{i=1}^d (\sqrt{\mu(i)} - \sqrt{\nu(i)})^2}. \quad (\text{C.2})$$

For arbitrary $\mathbf{M}, \mathbf{M}' \in \mathcal{M}_d$, a standard calculation yields

$$\begin{aligned} \|\mathbf{M} - \mathbf{M}'\| &= \max_{i \in [d]} \|\mathbf{M}(i, \cdot) - \mathbf{M}'(i, \cdot)\|_{\text{TV}} \geq 2 \max_{i \in [d]} H^2(\mathbf{M}(i, \cdot), \mathbf{M}'(i, \cdot)) \\ &= 2 \max_{i \in [d]} \left(1 - \sum_{j=1}^d \sqrt{\mathbf{M}(i, j)\mathbf{M}'(i, j)} \right) = 2 \left(1 - \min_{i \in [d]} \sum_{j=1}^d \sqrt{\mathbf{M}(i, j)\mathbf{M}'(i, j)} \right) \\ &\geq 2 \left(1 - \rho([\mathbf{M}, \mathbf{M}']_{\sqrt{\cdot}}) \right) = 2\text{Dist}_{\sqrt{\cdot}}(\mathbf{M}, \mathbf{M}'), \end{aligned} \quad (\text{C.3})$$

where the second inequality follows from the Perron-Frobenius theorem [Meyer, 2000, Chapter 8]. For the second claim, choose any $\theta \in [-1, 1]$ and put

$$\mathbf{M} = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 \\ \frac{1}{2} & \frac{1}{2} & 0 & 0 \\ 0 & 0 & \frac{1}{2} & \frac{1+\theta}{2} \\ 0 & 0 & \frac{1}{2} & \frac{1-\theta}{2} \end{pmatrix}, \quad \mathbf{M}' = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 \\ \frac{1}{2} & \frac{1}{2} & 0 & 0 \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix}. \quad (\text{C.4})$$

Since they have an identical connected component, $\text{Dist}_{\sqrt{\cdot}}(\mathbf{M}, \mathbf{M}') = 0$, whereas $\|\mathbf{M} - \mathbf{M}'\| = |\theta|$. \square

Remark C.1 *It follows that ε -identity testing with respect $\text{Dist}_{\sqrt{\cdot}}(\cdot, \cdot)$ reduces to 2ε -identity testing with respect to $\|\cdot\|$. In particular, the results of Daskalakis et al. [2017] immediately imply a lower bound of $\Omega(d/\varepsilon^2)$ for our testing problem.*