
Identity Testing of Reversible Markov Chains

Sela Fried

Department of Computer Science
Ben-Gurion University of the Negev

Geoffrey Wolfer

JSPS International Research Fellow
Department of Computer and Information Sciences
Tokyo University of Agriculture and Technology

Abstract

We consider the problem of identity testing of Markov chain transition matrices based on a single trajectory of observations under the distance notion introduced by Daskalakis et al. (2018a) and further analyzed by Cherapanamjeri and Bartlett (2019). Both works made the restrictive assumption that the Markov chains under consideration are symmetric. In this work we relax the symmetry assumption and show that it is possible to perform identity testing under the much weaker assumption of reversibility, provided that the stationary distributions of the reference and of the unknown Markov chains are close under a distance notion related to the separation distance. Additionally, we provide intuition on the distance notion of Daskalakis et al. (2018a) by showing how it behaves under several natural operations. In particular, we address some of their open questions.

1 INTRODUCTION

Efficiently distinguishing whether an unknown stochastic process is identical to a reference one or at least ε -far from it under some notion of distance is a fundamental problem of the field of *property testing*. Although the sample complexity in the iid case under the total variation distance is known to be of $\Theta(\sqrt{d}/\varepsilon^2)$, where d is the support size (see Waggoner (2015) for a summary), the Markovian case remains far from being settled. In this setting one has the transition matrix of some reference Markov chain and needs to decide, with prescribed confidence, based

on a single long trajectory sampled from an unknown Markov chain which started from an arbitrary state, whether the trajectory was sampled from the reference Markov chain or from a Markov chain that is at least ε -far from it, with respect to some notion of distance.

Identity testing in the Markovian setting was applied by Daskalakis et al. (2017) on the problem of testing whether an observed card shuffling is performed according to a certain riffle shuffle model. More recently, Matsui et al. (2022) used identity testing of Markov chains to analyse COVID-19 evolution.

2 RELATED WORK AND MOTIVATION

Prior work on the problem of identity testing of Markov chains within the property testing framework has so far focused on two main notions of distance: On one hand, Daskalakis et al. (2018a), motivated by Kazakos (1978), considered a distance notion defined as the spectral radius of the entry-wise geometric mean of the transition matrices. Restricting their analysis to symmetric Markov chains, they obtained an upper bound of $\tilde{O}(d/\varepsilon + h)$ on the sample complexity where d is the state space size, ε is the proximity parameter of the test, h depends on the hitting time of the reference Markov chain and the tilde notation hides polylogarithmic factors in all quantities being used. They also proved a lower bound of $\Omega(d/\varepsilon)$ and conjectured that this is the true sample complexity. Subsequently and still under the symmetry assumption, Cherapanamjeri and Bartlett (2019) used sparsest-cut techniques to obtain an upper bound of $\tilde{O}(d/\varepsilon^4)$ proving thereby that the sample complexity is independent of the hitting time.

In parallel, Wolfer and Kontorovich (2020) considered a distance between Markov chains that relies on the infinity norm over stochastic matrices. They showed that, under this distance, the broader class of ergodic Markov chains may be addressed and proved an upper bound on the sample complexity of $\tilde{O}((\sqrt{d} + t_{\text{rel}})/\pi_*)$

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

where t_{rel} and π_* are, respectively, the relaxation time and the minimum stationary probability of the reference Markov chain. Subsequently, Chan et al. (2021) characterized the sample complexity of the problem in terms of what they refer to as the k -cover time which enabled them to generalize to all irreducible Markov chains. At the same time, Fried and Wolfer (2021) too extended the results of Wolfer and Kontorovich (2020) to irreducible Markov chains by moving to the α -lazy versions of the Markov chains under consideration.

3 MAIN CONTRIBUTION

In this work we relax the symmetry assumption of Daskalakis et al. (2018a) and Cherapanamjeri and Bartlett (2019), which poses a considerable limitation in terms of applicability. In the words of Chan et al. (2021, p. 3), “However, identity testing under this distance only works for symmetric Markov chains, which is a quite restricted sub-family of Markov chains [...]. Thus we do not study learning and testing problems under this distance.”

Our main insight is that the methods of the aforementioned works are not limited to symmetric Markov chains and that a finer analysis allows to perform identity testing of Markov chains belonging to the much larger and natural class of reversible Markov chains. This generalization significantly enhances the applicability of the distance notion since the class of reversible Markov chains corresponds to random walks on networks (e.g. Levin and Peres (2017, Section 9.1)).

A consequence of the symmetry assumption of Daskalakis et al. (2018a) and Cherapanamjeri and Bartlett (2019) is that the stationary distributions of the Markov chains under consideration are uniform. In particular, despite the uncertainty regarding the unknown Markov chain, we do know that its stationary distribution is equal to the stationary distribution of the reference Markov chain. In the generalization to reversible Markov chains this property is lost and that poses a challenge. Clearly, one possibility is to assume that the reference and the unknown Markov chains share the same (but now arbitrary) stationary distribution. However, this assumption can be difficult to verify empirically since stationary distributions could be arbitrarily close while being different. We show that with a more careful analysis and only at a constant price in terms of the sample complexity, it is possible to perform identity testing on reversible Markov chains whose stationary distributions are not too far apart from each other.

Our main result is stated in the following theorem. It uses the distance notion between Markov chains of Daskalakis et al. (2018a) (cf. Definition 4.2) and of a

distance notion between probability distributions that is closely related to the separation distance (cf. Definition 4.3):

Theorem 3.1. *Suppose we have the transition matrix \bar{P} of an irreducible and reversible Markov chain with minimum stationary probability π_* and let $\varepsilon \in (0, 1)$. There is a polynomial time algorithm which given access to a single trajectory of length $\tilde{O}(1/\varepsilon^4\pi_*)$ from an unknown irreducible and reversible Markov chain P with a stationary distribution π that satisfies*

$$\left\| \frac{\pi}{\bar{\pi}} - \mathbf{1} \right\|_{\infty} \leq \frac{\varepsilon}{2},$$

distinguishes between

$$P = \bar{P} \quad \text{and} \quad \text{Distance}(P, \bar{P}) \geq \varepsilon,$$

with probability of success at least $3/5$.

Remark 3.2. We make the following observations:

1. When the Markov chains are symmetric, $\pi_* = 1/d$ and our upper bound of $\tilde{O}(1/\varepsilon^4\pi_*)$ takes the form of $\tilde{O}(d/\varepsilon^4)$, restoring the result of Cherapanamjeri and Bartlett (2019, Theorem 10). Thus, our generalization comes with no additional cost in terms of the sample complexity.
2. Similar to the symmetric case, the upper bound of $\tilde{O}(1/\varepsilon^4\pi_*)$ involves only efficiently computable parameters of the reference Markov chain and, in particular, does not depend on the unknown Markov chain.
3. Having an upper bound on the mixing time of the unknown Markov chain significantly simplifies the problem. Indeed, suppose the mixing times of both the reference as well as the unknown Markov chain are upper bounded by T . By thinning the process and keeping pairs of observations (X_t, X_{t+1}) separated roughly by $\tilde{O}(T)$ steps, a trajectory of length $\tilde{O}(Td/\varepsilon)$ yields $\tilde{O}(d/\varepsilon)$ almost independent pairs. This suffices to perform identity testing on the edge measures with respect to the Hellinger distance which dominates the one in Definition 4.2, allowing to solve the original problem by reduction. In particular, under this assumption, the requirement that the stationary distributions are close is redundant. Additionally, this shows that the dependency on π_* in our bound is not strictly necessary in this case. This work is concerned with the more challenging setting where the mixing time of the unknown Markov chain is unknown, precluding this approach.
4. Both Daskalakis et al. (2018a) and Cherapanamjeri and Bartlett (2019) make use of the hitting

time to bound the amount of time needed to visit every state “enough” times, which is, in the symmetric case, the same for all states. In contrast, in the reversible case, “enough” depends on the weight of each state in the stationary distribution and it is not clear how the hitting time could be used in this case. We avoid the use of the hitting time by applying the concentration bounds of Paulin (2015).

5. The algorithm of Cherapanamjeri and Bartlett (2019) has a polynomial time complexity and this is unchanged in our generalization.

4 PRELIMINARIES

For $d \in \mathbb{N}$ we denote $[d] = \{1, 2, \dots, d\}$ and write Δ_d for the set of all probability distributions over $[d]$. Vectors will be written as row-vectors and for $x \in \mathbb{R}^d$ and $i \in [d]$ we write $x(i)$ for the i th entry of x . Similarly, if P is a square matrix of size d and $i, j \in [d]$, we write $P(i, j)$ for the entry at the i th row and the j th column of P . We refer to transition matrices of Markov chains simply as Markov chains.

Irreducible and reversible Markov chains. Denote by \mathcal{M}_{irr} (resp. $\mathcal{M}_{\text{irr,rev}}, \mathcal{M}_{\text{erg,rev}}$) the set of all irreducible (resp. irreducible and reversible, ergodic and reversible) Markov chains on the state space $[d]$. Let $P \in \mathcal{M}_{\text{irr}}, \mu \in \Delta_d, m \in \mathbb{N}$ and $i_1, \dots, i_m \in [d]$. By $(X_t)_{t \in [m]} \sim (P, \mu)$ we mean

$$\mathbb{P}((X_1, \dots, X_m) = (i_1, \dots, i_m)) = \mu(i_1) \prod_{t=1}^{m-1} P(i_t, i_{t+1}).$$

For an entry-wise positive vector $\pi \in \mathbb{R}^d$ let $\langle \cdot, \cdot \rangle_\pi$ be the inner product on \mathbb{R}^d given by

$$\langle x, y \rangle_\pi = \sum_{i \in [d]} x(i)y(i)\pi(i), \quad \forall x, y \in \mathbb{R}^d$$

(cf. (Levin and Peres, 2017, p. 153)) and let $\|\cdot\|_{2,\pi}$ be the induced norm. We write $L_2(\pi) := (\mathbb{R}^d, \langle \cdot, \cdot \rangle_\pi)$ for the corresponding Hilbert space. For $P \in \mathcal{M}_{\text{irr}}$ with stationary probability π , the *edge measure* Q is defined by $Q = \text{diag}(\pi)P$, where $\text{diag}(\pi)$ is the diagonal matrix whose entries correspond to π (cf. Levin and Peres (2017, p. 88)). Define the *time reversal* of P by $P^* = \text{diag}(\pi)^{-1}P^T \text{diag}(\pi)$ (e.g. Levin and Peres (2017, 1.33)) and the *multiplicative reversibilization* of P by $P^\dagger = P^*P$ (e.g. Fill (1991, 2.2)). Finally, the *spectral gap* γ of $P \in \mathcal{M}_{\text{irr,rev}}$ is defined by $\gamma = 1 - \lambda_2$ where λ_2 is the second largest eigenvalue of P (cf. Levin and Peres (2017, p. 154)).

Censored Markov chains. Let $(X_t)_{t \in \mathbb{N}}$ be a Markov chain with transition matrix P and let $\emptyset \neq S \subseteq [d]$. Consider the stochastic process $(X_{\tau_t})_{t \in \mathbb{N}}$ defined as follows: $\tau_1 = \inf\{i \in \mathbb{N} \mid X_i \in S\}$ and $\tau_{t+1} = \inf\{i \in \mathbb{N} \mid i > \tau_t, X_i \in S\}$ for every $t \in \mathbb{N}$. It is well known that $(X_{\tau_t})_{t \in \mathbb{N}}$ is a Markov chain. It is called the *chain induced on S* (Levin and Peres, 2017, Example 13.19) or the *watched chain on S* (Lévy, 1951, 1952, 1958) and we shall denote its transition matrix by $P_{\text{cen}}(S)$. There is an explicit description of $P_{\text{cen}}(S)$ in terms of certain submatrices of P that uses the following notation: Let $R, T \subseteq [d]$. We write $P_{R,T}$ for the matrix obtained from P by keeping only the rows and columns with indices in R and T , respectively. If $R = T$ we write P_R instead of $P_{R,R}$. With this notation it holds

$$P_{\text{cen}}(S) = P_S + \sum_{t=1}^{\infty} P_{S,[d] \setminus S} P_{[d] \setminus S}^t P_{[d] \setminus S, S}$$

(e.g. Kemeny et al. (2012, Lemma 6-6)). In addition, $P_{\text{cen}}(S)$ is irreducible (resp. ergodic) if P is irreducible (resp. ergodic) and has the stationary distribution π_S given by

$$\pi_S(i) = \frac{\pi(i)}{\pi(S)}, \quad \forall i \in S$$

where $\pi(S) = \sum_{i \in S} \pi(i)$ (e.g. (Zhao and Liu, 1996, Lemma 2)). The following lemma generalizes to reversible Markov chains the statement in Cherapanamjeri and Bartlett (2019, Lemma 14) according to which $P_{\text{cen}}(S)$ is symmetric if P is symmetric. Its proof is given in Appendix A.

Lemma 4.1. *Let $P \in \mathcal{M}_{\text{irr,rev}}$ and let $S \subseteq [d]$. Then $P_{\text{cen}}(S) \in \mathcal{M}_{\text{irr,rev}}$.*

Statistical distances. The Hellinger distance and the total variation distance between two distributions $p, q \in \Delta_d$ are given by

$$\begin{aligned} d_{\text{Hel}}^2(p, q) &= \frac{1}{2} \sum_{i \in [d]} \left(\sqrt{p(i)} - \sqrt{q(i)} \right)^2 \\ &= 1 - \sum_{i \in [d]} \sqrt{p(i)q(i)} \end{aligned}$$

and

$$d_{\text{TV}}(p, q) = \frac{1}{2} \sum_{i \in [d]} |p(i) - q(i)|,$$

respectively.

It is well known (e.g. Gibbs and Su (2002, p. 13)) that

$$d_{\text{Hel}}^2(p, q) \leq d_{\text{TV}}(p, q) \leq \sqrt{2}d_{\text{Hel}}(p, q). \quad (1)$$

For $P, \bar{P} \in \mathcal{M}_{\text{irr}}$ we denote by $P \circ \bar{P}$ the *Hadamard product of P and \bar{P}* (e.g. Horn and Johnson (2012, Definition 7.5.1)) and by \sqrt{P} the square matrix of size d such that

$$\sqrt{P}(i, j) = \sqrt{P(i, j)}, \quad \forall i, j \in [d].$$

For $p, q \in \Delta_d$, we define $p \circ q$ and \sqrt{p} analogously.

The following distance between two Markov chains was proposed by Daskalakis et al. (2018a, p. 4) who motivated it by noticing that (a) it vanishes if and only if the Markov chains share an identical essential communicating class and (b) it captures the ability to distinguish between the Markov chains based on a single long trajectory (Daskalakis et al., 2018a, Claims 1 and 2). The distance relies on the *spectral radius* $\rho(P)$ of a square matrix P defined by

$$\rho(P) = \max\{|\lambda| \mid \lambda \text{ is an eigenvalue of } P\}$$

(e.g. Horn and Johnson (2012)).

Definition 4.2. Let P and \bar{P} be two Markov chains. Define

$$\text{Distance}(P, \bar{P}) = 1 - \rho\left(\sqrt{P \circ \bar{P}}\right).$$

We will perform identity testing on reversible Markov chains whose stationary distributions are not too far from each other under the following distance notion, which is closely related to the *separation distance* (e.g. Levin and Peres (2017, 6.7) or Aldous and Diaconis (1987, p. 72)):

Definition 4.3. Let $\mu, \nu \in \Delta_d$ such that ν is entry-wise positive. Denote by $\mathbf{1}$ the vector in \mathbb{R}^d that has all entries equal to 1. Define

$$\left\| \frac{\mu}{\nu} - \mathbf{1} \right\|_{\infty} := \max_{i \in [d]} \left| \frac{\mu(i)}{\nu(i)} - 1 \right|.$$

5 STATE SPACE PARTITIONING AND COMPONENT ANALYSIS

The main insight of Cherapanamjeri and Bartlett (2019) that enables them to improve upon Daskalakis et al. (2018a) and discard the dependency on the hitting time is that in order to distinguish between two different Markov chains it is sufficient to analyse trajectories in subsets of states that are internally well connected (see Corollary 5.4 and the paragraph preceding it). To achieve this they devise a new algorithm for partitioning of the state space $[d]$. This algorithm, upon receiving a reference Markov chain and a tolerance parameter, returns a tuple (\mathcal{S}, T) where \mathcal{S} is a set of well connected subsets of $[d]$ (components) and

T is a subset of $[d]$ in which the Markov chain does not spend too much time. The subsets in \mathcal{S} , together with T , form a partition of $[d]$. The Markov chain identity testing is then reduced to iid identity testing of distributions induced by these components. In this section we introduce the machinery and prove its properties.

5.1 Spectral and probabilistic properties of the components

Recall that for $P \in \mathcal{M}_{\text{irr,rev}}$ and $\nu \in \Delta_d$ the matrix $\text{diag}(\nu)P$ encodes a distribution over $[d]^2$. If $\nu = \pi$, where π is the stationary distribution of P , the joint distribution $\text{diag}(\nu)P$ corresponds to the edge measure Q . To some component $S \subseteq [d]$ we wish to similarly associate a probability distribution over which we will later apply iid identity testing. To this end, we first consider $\text{diag}(\nu_S)P_S$. Second, we merge all the outgoing transitions from S , i.e., all $(i, j) \in S \times [d]$ such that $j \notin S$, into a single symbol denoted by ∞ . This is the idea behind the following

Definition 5.1. Let P be a Markov chain and let $\nu \in \Delta_d$. For $S \subseteq [d]$ such that $\nu(S) > 0$ we denote by $\text{Distribution}(S, P, \nu)$ the probability distribution on the set $S^2 \cup \{\infty\}$ defined as follows:

$$\begin{aligned} \text{Distribution}(S, P, \nu)(i, j) &= \frac{\nu(i)P(i, j)}{\nu(S)}, \quad \forall i, j \in S \text{ and} \\ \text{Distribution}(S, P, \nu)(\infty) &= \\ &= 1 - \sum_{i, j \in S} \text{Distribution}(S, P, \nu)(i, j). \end{aligned}$$

Let $P, \bar{P} \in \mathcal{M}_{\text{irr,rev}}$ with stationary distributions $\pi, \bar{\pi}$, respectively, where we assume that \bar{P} is given and that P is unknown. A key property is that a positive distance between P and \bar{P} results in a positive Hellinger distance between the probability distributions induced by the corresponding edge measures over a component S , provided that in \bar{P} there is enough weight on the transitions between states in S (cf. Cherapanamjeri and Bartlett (2019, Lemma 13)). But, since in the reversible case the stationary distributions are, in general, not uniform, it is not clear how to sample from the unknown edge measure. We solve this problem by sampling from the joint distribution $R = \text{diag}(\bar{\pi})P$. This approach is guaranteed to succeed if π is not too far from $\bar{\pi}$. This is the content of the following lemma. A sketch of its proof is given here while the full proof is given in Appendix B.

Lemma 5.2. Let $P, \bar{P} \in \mathcal{M}_{\text{irr,rev}}$ with stationary distributions $\pi, \bar{\pi}$, respectively. Let $\varepsilon \in (0, 1)$ be such that $\text{Distance}(P, \bar{P}) \geq \varepsilon$ and

$$\left\| \frac{\pi}{\bar{\pi}} - \mathbf{1} \right\|_{\infty} \leq \frac{\varepsilon}{2}.$$

Let $S \subseteq [d]$ such that

$$\sum_{i,j \in S} \text{Distribution}(S, \bar{P}, \bar{\pi})(i, j) \geq 1 - \frac{\varepsilon}{16}.$$

Then

$$d_{\text{Hel}}^2(\text{Distribution}(S, P, \bar{\pi}), \text{Distribution}(S, \bar{P}, \bar{\pi})) \geq \frac{\varepsilon^2}{128}.$$

Proof sketch. Let $\bar{Q} = \text{diag}(\bar{\pi})\bar{P}$ and $R = \text{diag}(\bar{\pi})P$. From the reversibility of P, \bar{P} it follows that $\sqrt{P \circ \bar{P}}$ is self-adjoint in $L_2(\sqrt{\pi \circ \bar{\pi}})$. Applying the Courant–Fischer principle (e.g. Helmberg (2008, p. 219)), taking u to be the characteristic function of S and using the AM-GM inequality we obtain

$$\begin{aligned} \rho\left(\sqrt{P \circ \bar{P}}\right) &= \max_{u \neq 0} \frac{\langle \sqrt{P \circ \bar{P}} u, u \rangle_{\sqrt{\pi \circ \bar{\pi}}}}{\|u\|_{2, \sqrt{\pi \circ \bar{\pi}}}} \\ &\geq \sum_{i,j \in S} \sqrt{\frac{R(i, j)}{\bar{\pi}(S)}} \sqrt{\frac{\bar{Q}(i, j)}{\bar{\pi}(S)}} \sqrt{\frac{\pi(i)}{\bar{\pi}(i)}} \sqrt{\frac{\bar{\pi}(S)}{\pi(S)}}. \end{aligned}$$

The assumption on $\frac{\pi}{\bar{\pi}}$ guarantees that $\sqrt{\frac{\pi(i)}{\bar{\pi}(i)}} \sqrt{\frac{\bar{\pi}(S)}{\pi(S)}}$ is not too small. The claim follows by distinguishing between two cases. \square

Connectivity (or conductance) measures how a Markov chain navigates among its states and is controlled by the Cheeger constant (cf. Levin and Peres (2017, p. 88)):

Definition 5.3. Let $P \in \mathcal{M}_{\text{irr,rev}}$ with stationary distribution π and edge measure Q . Let $I \subseteq [d]$ and $S \subsetneq I$. The *bottleneck ratio* of S in I is defined by

$$\Phi(P, S, I) = \frac{\sum_{i \in S, j \in I \setminus S} Q(i, j)}{\min\{\pi(S), \pi(I \setminus S)\}}$$

and the *Cheeger constant* of P is defined by

$$\Phi_*(P) = \min_{S \subsetneq [d]} \Phi(P, S, [d]).$$

We shall refer to a component $S \subseteq [d]$ for which the condition in the following corollary is satisfied with some $\alpha > 0$ as *well connected*. A well connected component allows to control the spectral gap of the induced censored Markov chain:

Corollary 5.4. Let $P \in \mathcal{M}_{\text{irr,rev}}$ with stationary distribution π . Let $S \subseteq [d]$ and $\alpha \geq 0$ be such that

$$\Phi(P, R, S) \geq \alpha, \quad \forall R \subsetneq S.$$

Then $\gamma \geq \frac{\alpha^2}{2}$ where γ is the spectral gap of $P_{\text{cen}}(S)$.

Proof. Let Q be the edge measure of P and suppose $R \subsetneq S$ is such that $\pi(R) \leq \frac{\pi(S)}{2}$. Since $P_{\text{cen}}(S)(i, j) \geq P(i, j)$ for every $i, j \in S$, we obtain

$$\begin{aligned} \Phi(P_{\text{cen}}(S), R, S) &= \frac{\sum_{i \in R, j \in S \setminus R} \frac{\pi(i)}{\bar{\pi}(S)} P_{\text{cen}}(S)(i, j)}{\pi(R)/\pi(S)} \\ &\geq \frac{\sum_{i \in R, j \in S \setminus R} Q(i, j)}{\pi(R)} \\ &= \Phi(P, R, S) \geq \alpha. \end{aligned}$$

Thus, $\Phi_*(P_{\text{cen}}(S)) \geq \alpha$. By Cheeger’s inequality (Sinclair and Jerrum, 1989, Lemma 3.3), the spectral gap γ of $P_{\text{cen}}(S)$ satisfies $\gamma \geq \Phi_*(P_{\text{cen}}(S))^2/2$. \square

Recall that the state space partitioning algorithm, in addition to the well connected components, returns a subset $T \subseteq [d]$ comprising of states that belong to no component. The following lemma bounds the largest eigenvalue of the submatrix P_T . Its proof is a variation of the proof of (Sinclair and Jerrum, 1989, Lemma 3.3) and is given in Appendix C. The bound is then used in Lemma 5.6 to upper bound the time the Markov chain spends in T . The proof of Lemma 5.6 is given in Appendix D.

Lemma 5.5. Let $P \in \mathcal{M}_{\text{irr,rev}}$ with stationary distribution π and edge measure Q . Let $T \subsetneq [d]$ and $\alpha \geq 0$ be such that

$$\frac{\sum_{i \in R, j \in [d] \setminus R} Q(i, j)}{\pi(R)} \geq \alpha, \quad \forall R \subseteq T.$$

Let λ denote the largest eigenvalue of P_T . Then

$$\lambda \leq 1 - \frac{\alpha^2}{2}.$$

Lemma 5.6. Let $P \in \mathcal{M}_{\text{irr,rev}}$ with stationary distribution π and edge measure Q . Let $T \subsetneq [d]$ and $\alpha > 0$ be such that

$$\frac{\sum_{i \in R, j \in [d] \setminus R} Q(i, j)}{\pi(R)} \geq \alpha, \quad \forall R \subseteq T.$$

Denote $(\pi_T)_* = \min_{i \in T} \{\pi(i)\}$. Let $\delta \in (0, 1)$ and suppose $m \geq \Omega\left(\frac{\log \frac{1}{(\pi_T)_*} \log \frac{1}{\delta}}{\alpha^2}\right)$. Then

$$\mathbb{P}\left(\sum_{t=1}^m \mathbf{1}\{X_t \notin T\} \geq \Omega\left(\frac{m\alpha^2}{\log \frac{1}{(\pi_T)_*}}\right)\right) \geq 1 - \delta.$$

5.2 Markov chain partitioning algorithm

In Cherapanamjeri and Bartlett (2019) an algorithm for partitioning the state space $[d]$ is devised which,

upon receiving a reference Markov chain and a tolerance parameter, returns two objects: A set \mathcal{S} of subsets of $[d]$ in which the iid identity tester will be applied and $T \subseteq [d]$ in which the Markov chain does not spend too much time. The properties of these two objects are given in the following theorem. The algorithm (and its analysis) needs to be modified in order to be applicable beyond the symmetric case. In the rest of the section we give only the details whose modification was more involved. The reader is referred to Cherapanamjeri and Bartlett (2019, Appendix A) for the complete description and analysis of the algorithm.

Theorem 5.7. *Let $P \in \mathcal{M}_{\text{irr,rev}}$ with stationary distribution π and edge measure Q . Let $\beta \in (0, 1)$. There exists an algorithm that returns a tuple (\mathcal{S}, T) such that $S \subseteq [d], \forall S \in \mathcal{S}, T \subseteq [d]$ and $\mathcal{S} \cup \{T\}$ is a partition of $[d]$. Furthermore, for every $S \in \mathcal{S}$ it holds:*

- (1) $\frac{\sum_{i,j \in S} Q(i,j)}{\pi(S)} \geq 1 - \beta$.
- (2) $\Phi(P, R, S) \geq \Omega\left(\frac{\beta}{\log^2 d}\right), \forall R \subsetneq S$.
- (3) $\frac{\sum_{i \in R, j \in [d] \setminus R} Q(i,j)}{\pi(R)} \geq \Omega\left(\frac{\beta}{\log d}\right), \forall R \subseteq T$.

Definition 5.8. Let $S \subsetneq [d]$. The cut metric associated with S is defined by

$$\delta_S(i, j) = \begin{cases} 0 & \text{if } i, j \in S \text{ or } i, j \in [d] \setminus S \\ 1 & \text{otherwise.} \end{cases}$$

Definition 5.9. Let $P \in \mathcal{M}_{\text{irr,rev}}$ with stationary distribution π and edge measure Q . Let $I \subseteq [d]$ and $T \subsetneq I$. The sparsest cut with component constraints (SPCCC) is defined as

$$S^* = \arg \min_{T \subseteq S \subsetneq I} \frac{\sum_{i,j \in I} Q(i, j) \delta_S(i, j)}{\sum_{i,j \in I} \pi(i) \pi(j) \delta_S(i, j)}.$$

The corresponding linear programming relaxation is given by

$$\begin{aligned} & \min \sum_{i,j \in I} Q(i, j) \delta_{ij} \text{ such that} \\ & \delta_{ii} = 0, \quad \forall i \in I \\ & \delta_{ij} \leq \delta_{ik} + \delta_{kj}, \quad \forall i, j, k \in I \\ & \sum_{i,j \in I} \pi(i) \pi(j) \delta_{ij} = 1 \\ & \delta_{ij} \geq 0 \\ & \delta_{ij} = 0, \quad \forall i, j \in T \\ & \delta_{ik} = \delta_{jk}, \quad \forall i, j \in T, k \in I. \end{aligned} \tag{2}$$

Theorem 5.10. *Given an instance of the SPCCC problem there exists a polynomial time algorithm FindComp that returns $S' \subsetneq I$ such that $S' \cap T = \emptyset$ and*

$$\begin{aligned} & \frac{\sum_{i,j \in I} Q(i, j) \delta_{S'}(i, j)}{\sum_{i,j \in I} \pi(i) \pi(j) \delta_{S'}(i, j)} \\ & \leq O(\log d) \min_{T \subseteq S \subsetneq I} \frac{\sum_{i,j \in I} Q(i, j) \delta_S(i, j)}{\sum_{i,j \in I} \pi(i) \pi(j) \delta_S(i, j)}. \end{aligned} \tag{3}$$

Proof. The proof of Linial et al. (1995, Theorem 4.1) literally extends to our case (the symmetry of $Q(i, j)$ due to reversibility and the symmetry of $\pi(i) \pi(j)$ are crucial) and establishes (3) for some $S' \subsetneq I$. We shall provide only the details necessary to prove the claim regarding the intersection: First, notice that it suffices to prove that either $S' \cap T = \emptyset$ or $(I \setminus S') \cap T = \emptyset$ since the left hand-side of (3) is the same for both S' and $I \setminus S'$ and we let the algorithm return the set satisfying the intersection condition. Now, denote by δ the minimizing metric that is found by the linear program above. By Bourgain's metric-embedding theorem (e.g. Linial et al. (1995, Corollary 3.4)), there are $x_1, \dots, x_d \in \mathbb{R}^m$ where $m = O(\log^2 d)$ such that for every $i, j \in I$ it holds

$$\Omega\left(\frac{1}{\log d}\right) \delta_{ij} \leq \|x_i - x_j\|_1 \leq \delta_{ij}.$$

It follows from this, together with the constraint (2), that $x_i = x_j$ for every $i, j \in T$. Recall that by the proof of Linial et al. (1995, Theorem 4.1), $S' = \{i \in I \mid x_i(r) = 1\}$ for some $r \in [m]$. Conclude that either $T \subseteq S'$ or $T \subseteq I \setminus S'$. \square

Corollary 5.11. *In the setting of Theorem 5.10, for S' returned by the FindComp algorithm it holds*

$$\Phi(P, S', I) \leq O(\log d) \min_{T \subseteq S \subsetneq I} \Phi(P, S, I).$$

Proof. Let $S \subsetneq I$. It is easy to see that

$$\frac{\sum_{i,j \in I} Q(i, j) \delta_S(i, j)}{\sum_{i,j \in I} \pi(i) \pi(j) \delta_S(i, j)} = \frac{\sum_{i \in S, j \in I \setminus S} Q(i, j)}{\pi(S) \pi(I \setminus S)}. \tag{4}$$

Now, assume that $\frac{\pi(I)}{2} \leq \pi(S) \leq \pi(I)$. Then

$$1 \leq \frac{\pi(I) \min\{\pi(S), \pi(I \setminus S)\}}{\pi(S) \pi(I \setminus S)} \leq 2.$$

By symmetry, this holds also if $\frac{\pi(I)}{2} \leq \pi(I \setminus S) \leq \pi(I)$. Conclude that

$$\begin{aligned} & \frac{\pi(I) \sum_{i \in S, j \in I \setminus S} Q(i, j)}{2\pi(S) \pi(I \setminus S)} \leq \Phi(P, S, I) \\ & \leq \frac{\pi(I) \sum_{i \in S, j \in I \setminus S} Q(i, j)}{\pi(S) \pi(I \setminus S)}. \end{aligned} \tag{5}$$

The claim follows by combining (3), (4) and (5) together. \square

5.3 Sample generation

Both Daskalakis et al. (2018a) and Cherapanamjeri and Bartlett (2019) reduce the Markov chain identity testing to identity testing of distributions in the iid case and use a black box tester which we will refer to as `iidTester` (cf. (Daskalakis et al., 2018b, Algorithm 1)). The algorithm accepts the arguments $(Y_t)_{t \in [m]}, \bar{p}, \varepsilon, \delta$ and returns either 0 or 1. The guarantees on the tester’s performance as well as the meaning of each of its arguments is given in the following lemma, which is used in Lemma 5.16.

Lemma 5.12. *Suppose we have the description of a probability distribution $\bar{p} \in \Delta_d$ and let $\varepsilon, \delta \in (0, 1)$. Then there is a tester called `iidTester` such that if it is given $m \geq \Omega\left(\frac{\sqrt{d} \log \frac{1}{\delta}}{\varepsilon^2}\right)$ iid samples $(Y_t)_{t \in [m]}$ from an unknown probability distribution $p \in \Delta_d$ the following holds with probability at least $1 - \delta$:*

$$\text{iidTester}((Y_t)_{t \in [m]}, \bar{p}, \varepsilon, \delta) = \begin{cases} 0 & \text{if } p = \bar{p} \\ 1 & \text{if } d_{\text{Hel}}(p, \bar{p}) \geq \varepsilon. \end{cases}$$

The iid sampling process from the (Markovian) observed trajectory comprises of two stages: First, sample from $[d]$ according to the stationary distribution $\bar{\pi}$ of the reference Markov chain. Second, record the transitions from these states in the trajectory (cf. Daskalakis et al. (2018a, pp. 12-13)). This process is Cherapanamjeri and Bartlett (2019, Algorithm 1) and will be referred to as `iidGenerator`. Notice that when it is applied on $S \subseteq [d]$, instead of sampling uniformly and iid from S , which corresponds to the symmetric case, it needs to sample from S according to the distribution $\bar{\pi}_S$. We elaborate on the `iidGenerator` algorithm in Appendix H.

Following Daskalakis et al. (2018a), Cherapanamjeri and Bartlett (2019, Lemma 19) use a bound involving the hitting time to guarantee enough visits of each state in a component. It is not immediately clear how this approach generalizes to the reversible case. We, instead, exploit the concentration bounds of Paulin (2015) that rely on the spectral gap. This comes at no additional cost and arguably simplifies the analysis.

The following two lemmas give guarantees for the success of this iid sampling process. Their proofs are given in Appendix E and F.

Lemma 5.13. *Let $P \in \mathcal{M}_{\text{erg,rev}}$ with stationary distribution π and spectral gap γ . Let $\mu \in \Delta_d$ and $\delta \in (0, 1)$. Assume $(X_t)_{t \in \mathbb{N}} \sim (P, \mu)$. For $m \in \mathbb{N}$ and $i \in [d]$ denote by $N_m(i) = \sum_{t \in [m]} 1\{X_t = i\}$ the*

number of occurrences of state i in $(X_t)_{t \in [m]}$. Suppose $m \geq \Omega\left(\frac{\log \frac{1}{\delta \pi_}}{\pi_* \gamma}\right)$. Then*

$$\mathbb{P}\left(N_m(i) \geq \frac{\pi(i)}{2} m, \forall i \in [d]\right) \geq 1 - \delta.$$

Lemma 5.14. *Let $m \in \mathbb{N}$ and suppose $(Y_t)_{t \in [m]}$ are iid sampled according to a distribution $p \in \Delta_d$. Let $v \in \{0, 1, \dots, m\}^d$ be the histogram of $(Y_t)_{t \in [m]}$ and let $\delta \in (0, 1)$. Assume $m \geq \Omega\left(\frac{\log \frac{d}{\delta}}{p_*}\right)$. Then*

$$\mathbb{P}(v(i) \leq 2mp(i), \forall i \in [d]) \geq 1 - \delta.$$

In Lemma 5.16 it is assumed that there is access to enough iid samples from $\text{Distribution}(S, P, \bar{\pi})$. The following lemma upper bounds the number of visits to S that guarantees this. Due to the way the identity tester of the Markov chains is defined, such a guarantee is only needed in the case that $P = \bar{P}$.

Lemma 5.15. *Let $P \in \mathcal{M}_{\text{erg,rev}}$ with stationary distribution π . Let $\varepsilon \in (0, 1)$ and let (\mathcal{S}, T) be the tuple returned by the state space partitioning algorithm (Theorem 5.7) applied on P and $\beta = \varepsilon$. Let $\mu \in \Delta_d$ and suppose $(X_t)_{t \in \mathbb{N}} \sim (P, \mu)$. Then it is possible, with probability at least $\frac{9}{10}$, to generate $\frac{m}{4}$ iid samples from $\text{Distribution}(S, P, \pi)$, for every $S \in \mathcal{S}$ that is visited at least $m \geq \Omega\left(\frac{\log^4 d \log \frac{d}{(\pi_S)_*}}{\varepsilon^2 (\pi_S)_*}\right)$ times.*

Proof. Let $S \in \mathcal{S}$. By Theorem 5.7 (2),

$$\Phi(P, R, S) \geq \Omega\left(\frac{\varepsilon}{\log^2 d}\right), \quad \forall R \not\subseteq S.$$

By Corollary 5.4, $\gamma \geq \Omega\left(\frac{\varepsilon^2}{\log^4 d}\right)$ where γ is the spectral gap of $P_{\text{cen}}(S)$. Thus, by Lemma 5.13 applied on $P_{\text{cen}}(S)$ and $\delta = \frac{1}{20d}$, in a trajectory of length $m \geq \Omega\left(\frac{\log^4 d \log \frac{20d^2}{(\pi_S)_*}}{\varepsilon^2 (\pi_S)_*}\right)$ each state $i \in S$ is visited at least $\frac{\pi(i)}{2\pi(S)} m$ times, with probability at least $1 - \frac{1}{20d}$. By Lemma 5.14, applied on $p = \pi_S$ and $\delta = \frac{1}{20d}$, if m satisfies the above inequality, then

$$\mathbb{P}\left(v(i) \leq 2m \frac{\pi(i)}{\pi(S)}, \forall i \in S\right) \geq 1 - \frac{1}{20d}.$$

Intersecting the above two events, we conclude that, with probability at least $1 - \frac{1}{10d}$, we can to generate $\frac{m}{4}$ iid samples from $\text{Distribution}(S, P, \pi)$. By the union bound, with probability at least $\frac{9}{10}$, this is true for every $S \in \mathcal{S}$ that is visited at least m times. \square

5.4 Reduction to the iid case

The following lemma shows that in order to perform identity testing of Markov chains it suffices to perform identity testing in the iid case over any component in \mathcal{S} that is visited enough times:

Lemma 5.16. *Let $P, \bar{P} \in \mathcal{M}_{\text{irr,rev}}$ with stationary distributions $\pi, \bar{\pi}$, respectively, and let $\varepsilon \in (0, 1)$. Assume that*

$$\left\| \frac{\pi}{\bar{\pi}} - \mathbf{1} \right\|_{\infty} \leq \frac{\varepsilon}{2}.$$

Let (\mathcal{S}, T) be the tuple returned by the state space partitioning algorithm (Theorem 5.7) applied on \bar{P} and $\beta = \frac{\varepsilon}{16}$. Let $\mu \in \Delta_d$ and suppose $(X_t)_{t \in \mathbb{N}} \sim (P, \mu)$. Let $S \in \mathcal{S}$ and assume that we have $m \geq \Omega\left(\frac{|S| \log d}{\varepsilon^2}\right)$ iid samples $(Y_t)_{t \in [m]}$ from $\text{Distribution}(S, P, \bar{\pi})$. Then, with probability at least $\frac{9}{10}$,

$$\begin{aligned} \text{iidTester} \left((Y_t)_{t \in [m]}, \text{Distribution}(S, \bar{P}, \bar{\pi}), \frac{\varepsilon^2}{128}, \frac{1}{10d} \right) \\ = \begin{cases} 0 & \text{if } P = \bar{P} \\ 1 & \text{if } \text{Distance}(P, \bar{P}) \geq \varepsilon. \end{cases} \end{aligned}$$

Proof. Let \bar{Q} be the edge measure of \bar{P} . By Theorem 5.7 (1),

$$\begin{aligned} \sum_{i,j \in S} \text{Distribution}(S, \bar{P}, \bar{\pi})(i, j) &= \frac{\sum_{i,j \in S} \bar{Q}(i, j)}{\pi(S)} \\ &\geq 1 - \frac{\varepsilon}{16}. \end{aligned}$$

By Lemma 5.2,

$$d_{\text{Hel}}^2(\text{Distribution}(S, P, \bar{\pi}), \text{Distribution}(S, \bar{P}, \bar{\pi})) \geq \frac{\varepsilon^2}{128}.$$

Since $|\mathcal{S}| \leq d$, the claim follows from a union bound and the guarantees of Lemma 5.12. \square

6 PROOF OF THEOREM 3.1

To avoid periodicity issues, we exploit the fact (Proposition 7.1 (i)) that the distance between two Markov chains does not change much if the Markov chains are replaced by their respective α -lazy versions, as long as α is not too big. Thus, we may assume that P, \bar{P} are ergodic, at the cost of replacing ε with $\frac{\varepsilon}{2}$. Let (\mathcal{S}, T) be the tuple returned by the state space partitioning algorithm (Theorem 5.7) applied on P and $\beta = \varepsilon$. First assume that $P = \bar{P}$. Let Q be the edge measure of P . By Theorem 5.7 (3),

$$\frac{\sum_{i \in R, j \in [d] \setminus R} Q(i, j)}{\pi(R)} \geq \Omega\left(\frac{\varepsilon}{\log d}\right), \quad \forall R \subseteq T.$$

Applying Lemma 5.6 with $\alpha = \frac{\varepsilon}{\log d}$ and $\delta = \frac{1}{10}$ we have that if $m \geq \Omega\left(\frac{\log^6 d \log \frac{1}{\pi_*} \log \frac{d}{\pi_*}}{\varepsilon^4 \pi_*}\right)$ then

$$\mathbb{P}\left(\sum_{t=1}^m \mathbf{1}\{X_t \in \bigcup_{S \in \mathcal{S}} S\} \geq \Omega\left(\frac{\log^4 d \log \frac{d}{\pi_*}}{\varepsilon^2 \pi_*}\right)\right) \geq \frac{9}{10}.$$

It holds

$$\sum_{S \in \mathcal{S}} \frac{1}{(\pi_S)_*} = \sum_{S \in \mathcal{S}} \frac{\pi(S)}{\min_{i \in S} \{\pi(i)\}} \leq \sum_{S \in \mathcal{S}} \frac{\pi(S)}{\pi_*} = \frac{1}{\pi_*}.$$

Thus, by the strong form of the pigeonhole principle (e.g. Brualdi (1977, Theorem 2.2.1)), at least one $S \in \mathcal{S}$ is visited more than $\Omega\left(\frac{\log^4 d \log \frac{d}{\pi_*}}{\varepsilon^2 (\pi_S)_*}\right)$ times. By Lemma 5.15, with probability at least $\frac{9}{10}$, we may sample

$$\Omega\left(\frac{\log^4 d \log \frac{d}{\pi_*}}{\varepsilon^2 (\pi_S)_*}\right) \geq \Omega\left(\frac{|S| \log d}{\varepsilon^2}\right)$$

iid samples from

$$\text{Distribution}(S, P, \bar{\pi}) = \text{Distribution}(S, \bar{P}, \bar{\pi}).$$

Applying Lemma 5.16 we identify that $P = \bar{P}$ with probability of success at least $\frac{9}{10}$. We conclude that, in this case, the tester succeeds with probability at least $\frac{7}{10}$.

Assume now that $\text{Distance}(P, \bar{P}) \geq \varepsilon$. In this case the algorithm fails only if the iid sampling succeeds in at least one component $S \in \mathcal{S}$ but the iid tester fails on S . The probability that this happens is upper bounded by $\frac{1}{10}$ (Lemma 5.12). \square

7 SOME PROPERTIES OF THE DISTANCE

The purpose of this section is to develop intuition for the distance notion considered in this work (Definition 4.2). This will be done by establishing its behaviour under several natural operations, two of which were inspired by Daskalakis et al. (2018a, Open Questions 2 and 3). The proof of the following proposition is given in Appendix G.

Proposition 7.1. *Let $P, \bar{P} \in \mathcal{M}_{\text{irr}}$ with stationary distributions $\pi, \bar{\pi}$, time reversals P^*, \bar{P}^* and multiplicative reversibilizations $P^\dagger, \bar{P}^\dagger$, respectively. Let $\varepsilon \in (0, 1)$.*

- (i) *Assume that $\text{Distance}(P, \bar{P}) \geq \varepsilon$. Let $\alpha = \frac{\varepsilon^2}{2\sqrt{2}}$ and consider the α -lazy versions of P and \bar{P} :*

$$P' = \alpha I + (1 - \alpha) P,$$

$$\bar{P}' = \alpha I + (1 - \alpha) \bar{P},$$

respectively. Then $\text{Distance}(P', \bar{P}') \geq \frac{\varepsilon}{2}$.

(ii) It holds

$$\text{Distance}(P^*, \bar{P}^*) = \text{Distance}(P, \bar{P}).$$

(iii) Assume that P, \bar{P} are reversible and that

$$\left\| \frac{\pi}{\bar{\pi}} - \mathbf{1} \right\|_{\infty} < \varepsilon.$$

Let $\alpha \in [0, 1]$ and denote $P(\alpha) = \alpha P + (1 - \alpha)\bar{P}$. Then

$$\begin{aligned} \text{Distance}(P, P(\alpha)) &\geq 1 - \sqrt{\alpha} - 2\sqrt{\frac{1 - \alpha}{1 - \varepsilon}} \\ &\quad + 2\sqrt{\frac{1 - \alpha}{1 - \varepsilon}} \text{Distance}(P, \bar{P}). \end{aligned}$$

(iv) Let $k \in \mathbb{N}$. Then

$$1 - (1 - \text{Distance}(P, \bar{P}))^k \geq \text{Distance}(P^k, \bar{P}^k).$$

(v)

$$\text{Distance}(P^k, \bar{P}^k) \xrightarrow[k \rightarrow \infty]{} d_{\text{Hel}}^2(\pi, \bar{\pi}).$$

(vi) It holds

$$\text{Distance}(P^\dagger, \bar{P}^\dagger) \leq 2\text{Distance}(P, \bar{P}).$$

(vii) There exist irreducible and reversible Markov chains that are arbitrarily close under the distance such that the Hellinger distance between their stationary distributions is bounded away from 0.

Remark 7.2. Property (i) exhibits the robustness of the distance under transitions to α -lazy versions. This allows us to handle arbitrary irreducible Markov chains, although we invoke the concentration bounds of Paulin (2015) which hold only for ergodic Markov chains.

Property (iii) addresses Daskalakis et al. (2018a, Open Questions 2) asking how the distance between two Markov chains changes when one substitutes one of them with a convex combination of both.

Property (iv) addresses Daskalakis et al. (2018a, Open Questions 3) asking how the distance between two Markov chains is related to the distance between the same Markov chains being observed only at intervals of size k .

Property (v) is related to the previous one and shows that with increasing k it becomes increasingly harder to distinguish between two Markov chains that have the same stationary distribution.

Reversible Markov chains enjoy favourable properties. One way to make an irreducible Markov reversible is

by moving to its multiplicative reversibilization. It is therefore natural to ask how the distance behaves under this operation. This is addressed in (vi).

Finally, property (vii) reflects the fact that if two reducible Markov chains share an identical essential communicating class, the distance between them is 0 (cf. Daskalakis et al. (2018a, Claim 1)).

8 CONCLUSION

In this work we have replaced the restrictive symmetry assumption made by Daskalakis et al. (2018a) and Cherapanamjeri and Bartlett (2019) with the more natural one of reversibility and showed that it is possible to perform identity testing in this class under the distance notion between Markov chains introduced by Daskalakis et al. (2018a), provided that the stationary distributions of the reference and of the unknown Markov chains are not too far from each other under a certain distance notion between probability distributions. In addition, we provided intuition regarding the distance notion between Markov chains by making statements on its behaviour under several natural operations. The next step in our research agenda in identity testing of Markov chains is to investigate the possibility of removing the assumption on the closeness of the stationary distributions of the reference and the unknown chain Markov chains.

Societal impact The problem of goodness-of-fit is fundamental in statistics. We do not foresee any new negative societal impact arising from this research that has not been well-understood by now.

Acknowledgments Sela Fried is supported by the Israel Science Foundation (ISF) through grant no. 1456/18 and by the European Research Council through grant no. 949707. Geoffrey Wolfer is supported by the Japan Society for the Promotion of Science. We thank the referees for their careful reading and valuable comments that helped us improve the presentation of this paper.

References

- D. Aldous and P. Diaconis. Strong uniform times and finite random walks. *Advances in Applied Mathematics*, 8(1):69–97, 1987.
- R. A. Brualdi. *Introductory combinatorics*. Pearson Education India, 1977.
- S. Chan, Q. Ding, and S. Li. Learning and testing irreducible Markov chains via the k -cover time. In *Algorithmic Learning Theory*, pages 458–480. PMLR, 2021.

- Y. Cherapanamjeri and P. L. Bartlett. Testing symmetric Markov chains without hitting. In *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 758–785. PMLR, 2019.
- C. Daskalakis, N. Dikkala, and N. Gravin. Testing from one sample: Is the casino really using a riffle shuffle? *ArXiv*, abs/1704.06850, 2017.
- C. Daskalakis, N. Dikkala, and N. Gravin. Testing symmetric Markov chains from a single trajectory. In *Conference On Learning Theory*, pages 385–409. PMLR, 2018a.
- C. Daskalakis, G. Kamath, and J. Wright. Which distribution distances are sublinearly testable? In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2747–2764. SIAM, 2018b.
- R. Drnovšek and A. Peperko. Inequalities for the Hadamard weighted geometric mean of positive kernel operators on Banach function spaces. *Positivity*, 10(4):613–626, 2006.
- D. P. Dubhashi and A. Panconesi. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009.
- 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.
- S. Fried and G. Wolfer. On the α -lazy version of Markov chains in estimation and testing problems. *arXiv preprint arXiv:2105.09536*, 2021.
- A. L. Gibbs and F. E. Su. On choosing and bounding probability metrics. *International statistical review*, 70(3):419–435, 2002.
- I. Gohberg and M. G. Kreĭn. *Introduction to the theory of linear nonselfadjoint operators*, volume 18. American Mathematical Soc., 1978.
- G. Helmberg. *Introduction to spectral theory in Hilbert space*. Courier Dover Publications, 2008.
- R. A. Horn and C. R. Johnson. *Matrix analysis*. Cambridge university press, 2012.
- D. Kazakos. The Bhattacharyya distance and detection between Markov chains. *IEEE Transactions on Information Theory*, 24(6):747–754, 1978.
- J. G. Kemeny, J. L. Snell, and A. W. Knapp. *Denumerable Markov chains: with a chapter of Markov random fields by David Griffeth*, volume 40. Springer Science & Business Media, 2012.
- D. A. Levin and Y. Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.
- P. Lévy. Systèmes Markoviens et stationnaires. cas dénombrable. *Annales scientifiques de l'École Normale Supérieure*, 3e série, 68:327–381, 1951.
- P. Lévy. Complément à l'étude des processus de Markoff. *Annales scientifiques de l'École Normale Supérieure*, 3e série, 69:203–212, 1952.
- P. Lévy. Processus Markoviens et stationnaires. cas dénombrable. *Annales de l'institut Henri Poincaré*, 16(1):7–25, 1958.
- N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.
- P. Massart. Concentration inequalities and model selection. Lectures from the 33rd summer school on probability theory held in Saint-Flour, July 6–23, 2003. *Lecture Notes in Mathematics*, 1896, 2007.
- T. Matsui, N. Azzaoui, and D. Murakami. Analysis of COVID-19 evolution based on testing closeness of sequential data. *Japanese Journal of Statistics and Data Science*, pages 1–18, 2022.
- P. J. Olver, C. Shakiban, and C. Shakiban. *Applied linear algebra*, volume 1. Springer, 2006.
- D. Paulin. Concentration inequalities for Markov chains by Marton couplings and spectral methods. *Electron. J. Probab.*, 20:32 pp., 2015. doi: 10.1214/EJP.v20-4039.
- A. Sinclair and M. Jerrum. Approximate counting, uniform generation and rapidly mixing Markov chains. *Information and Computation*, 82(1):93–133, 1989.
- B. Waggoner. Lp testing and learning of discrete distributions. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science*, pages 347–356, 2015.
- G. Wolfer and A. Kontorovich. Minimax testing of identity to a reference ergodic Markov chain. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 191–201. PMLR, 26–28 Aug 2020.
- X. Zhan. *Matrix theory*, volume 147. American Mathematical Soc., 2013.
- Y. Q. Zhao and D. Liu. The censored Markov chain and the best augmentation. *Journal of Applied Probability*, pages 623–629, 1996.

Supplementary Material: Identity Testing of Reversible Markov Chains

A PROOF OF LEMMA 4.1

Suppose P is reversible, i.e. P satisfies the detailed balance equation (e.g. Levin and Peres (2017, (1.30)):

$$Q(i, j) = Q(j, i), \quad \forall i, j \in [d].$$

Assume without loss that $S = [k]$ for some $k \in [d]$ and let $i, j \in [k]$. We have

$$\begin{aligned} \pi_S(i)P_{\text{cen}}(S)(i, j) &= \frac{\pi(i)}{\pi(S)} \left(P_S + \sum_{t=1}^{\infty} P_{S, [d] \setminus S} P_{[d] \setminus S}^t P_{[d] \setminus S, S} \right) (i, j) \\ &= \frac{1}{\pi(S)} \left(\pi(i)P_S(i, j) + \pi(i) \left(\sum_{t=1}^{\infty} P_{S, [d] \setminus S} P_{[d] \setminus S}^t P_{[d] \setminus S, S} \right) (i, j) \right) \\ &= \frac{1}{\pi(S)} \left(\pi(j)P_S(j, i) + \sum_{t=1}^{\infty} \left(\pi(i)P_{S, [d] \setminus S} P_{[d] \setminus S}^t P_{[d] \setminus S, S} \right) (i, j) \right). \end{aligned} \quad (6)$$

Now, let $t \in \mathbb{N}$. We have

$$\begin{aligned} \left(\pi(i)P_{S, [d] \setminus S} P_{[d] \setminus S}^t P_{[d] \setminus S, S} \right) (i, j) &= \sum_{l=k+1}^d \sum_{m=k+1}^d \pi(i)P(i, l)P^t(l, m)P(m, j) \\ &= \pi(j) \sum_{l=k+1}^d \sum_{m=k+1}^d P(l, i)P^t(m, l)P(j, m) \\ &= \left(\pi(j)P_{S, [d] \setminus S} P_{[d] \setminus S}^t P_{[d] \setminus S, S} \right) (j, i). \end{aligned}$$

Thus,

$$\begin{aligned} (6) &= \frac{1}{\pi(S)} \left(\pi(j)P_S(j, i) + \sum_{t=1}^{\infty} \left(\pi(j)P_{S, [d] \setminus S} P_{[d] \setminus S}^t P_{[d] \setminus S, S} \right) (j, i) \right) \\ &= \pi_S(j)P_{\text{cen}}(S)(j, i). \end{aligned}$$

□

B PROOF OF LEMMA 5.2

It is easy to see that $\sqrt{P \circ \bar{P}}$ is self-adjoint in $L_2(\sqrt{\pi \circ \bar{\pi}})$. Let $\mathbf{1}_S \in \mathbb{R}^d$ be given by

$$\mathbf{1}_S(i) = \begin{cases} 1 & i \in S \\ 0 & \text{otherwise,} \end{cases} \quad \forall i \in [d].$$

Let $\bar{Q} = \text{diag}(\bar{\pi})\bar{P}$ be the edge measure and let $R = \text{diag}(\bar{\pi})P$. Then

$$\begin{aligned}
 \rho\left(\sqrt{P \circ \bar{P}}\right) &= \max_{u \neq 0} \frac{\langle \sqrt{P \circ \bar{P}}u, u \rangle_{\sqrt{\pi \circ \bar{\pi}}}}{\|u\|_{2, \sqrt{\pi \circ \bar{\pi}}}} \\
 &\stackrel{(i)}{\geq} \frac{\langle \sqrt{P \circ \bar{P}}\mathbf{1}_S, \mathbf{1}_S \rangle_{\sqrt{\pi \circ \bar{\pi}}}}{\|\mathbf{1}_S\|_{2, \sqrt{\pi \circ \bar{\pi}}}} \\
 &= \frac{\sum_{i,j \in S} \sqrt{R(i,j)} \sqrt{\bar{Q}(i,j)} \sqrt{\frac{\pi(i)}{\bar{\pi}(i)}}}{\sum_{i \in S} \sqrt{\pi(i)} \sqrt{\bar{\pi}(i)}} \\
 &= \frac{\sum_{i,j \in S} \sqrt{\frac{R(i,j)}{\bar{\pi}(S)}} \sqrt{\frac{\bar{Q}(i,j)}{\bar{\pi}(S)}} \sqrt{\frac{\pi(i)}{\bar{\pi}(i)}}}{\sum_{i \in S} \sqrt{\frac{\pi(i)}{\bar{\pi}(S)}} \sqrt{\frac{\bar{\pi}(i)}{\bar{\pi}(S)}} \sqrt{\frac{\pi(S)}{\bar{\pi}(S)}}} \\
 &\stackrel{(ii)}{\geq} \sum_{i,j \in S} \sqrt{\frac{R(i,j)}{\bar{\pi}(S)}} \sqrt{\frac{\bar{Q}(i,j)}{\bar{\pi}(S)}} \sqrt{\frac{\pi(i)}{\bar{\pi}(i)}} \sqrt{\frac{\bar{\pi}(S)}{\pi(S)}} \\
 &\stackrel{(iii)}{\geq} \sum_{i,j \in S} \sqrt{\frac{R(i,j)}{\bar{\pi}(S)}} \sqrt{\frac{\bar{Q}(i,j)}{\bar{\pi}(S)}} \underbrace{\sqrt{1 - \frac{\varepsilon}{2}} \sqrt{1 + \frac{\varepsilon}{2}}}_{\geq 1 - \frac{\varepsilon}{2}},
 \end{aligned}$$

where (i) is due to the Courant–Fischer principle (e.g. Helmberg (2008, p. 219)), (ii) stems from the AM-GM inequality as follows:

$$\sum_{i \in S} \sqrt{\frac{\pi(i)}{\bar{\pi}(S)}} \sqrt{\frac{\bar{\pi}(i)}{\bar{\pi}(S)}} \leq \frac{1}{2} \sum_{i \in S} \left(\frac{\pi(i)}{\bar{\pi}(S)} + \frac{\bar{\pi}(i)}{\bar{\pi}(S)} \right) = 1$$

and (iii) follows by definition of $\pi(S)$ and the assumption that $\|\frac{\pi}{\bar{\pi}} - \mathbf{1}\|_{\infty} \leq \frac{\varepsilon}{2}$:

$$\pi(S) = \sum_{i \in S} \pi(i) \leq \left(1 + \frac{\varepsilon}{2}\right) \sum_{i \in S} \bar{\pi}(i) = \left(1 + \frac{\varepsilon}{2}\right) \bar{\pi}(S).$$

By assumption, $\text{Distance}(P, \bar{P}) \geq \varepsilon$. Thus,

$$\varepsilon \leq 1 - \left(1 - \frac{\varepsilon}{2}\right) \sum_{i,j \in S} \sqrt{\frac{R(i,j)}{\bar{\pi}(S)}} \sqrt{\frac{\bar{Q}(i,j)}{\bar{\pi}(S)}}.$$

It follows

$$\frac{\varepsilon}{2} \leq \left(1 - \frac{\varepsilon}{2}\right) \left(1 - \sum_{i,j \in S} \sqrt{\frac{R(i,j)}{\bar{\pi}(S)}} \sqrt{\frac{\bar{Q}(i,j)}{\bar{\pi}(S)}}\right) \leq 1 - \sum_{i,j \in S} \sqrt{\frac{R(i,j)}{\bar{\pi}(S)}} \sqrt{\frac{\bar{Q}(i,j)}{\bar{\pi}(S)}}. \quad (7)$$

We distinguish between two cases: First assume that

$$\sum_{i,j \in S} \text{Distribution}(S, P, \bar{\pi})(i, j) \geq 1 - \frac{5\varepsilon}{16}. \quad (8)$$

In this case we have

$$\begin{aligned}
& d_{\text{Hel}}^2(\text{Distribution}(S, P, \bar{\pi}), \text{Distribution}(S, \bar{P}, \bar{\pi})) \\
& \geq \frac{1}{2} \sum_{i,j \in S} \left(\sqrt{\text{Distribution}(S, P, \bar{\pi})(i, j)} - \sqrt{\text{Distribution}(S, \bar{P}, \bar{\pi})(i, j)} \right)^2 \\
& = \frac{1}{2} \sum_{i,j \in S} \left(\sqrt{\frac{R(i, j)}{\bar{\pi}(S)}} - \sqrt{\frac{\bar{Q}(i, j)}{\bar{\pi}(S)}} \right)^2 \\
& = \frac{1}{2} \sum_{i,j \in S} \left(\frac{R(i, j)}{\bar{\pi}(S)} + \frac{\bar{Q}(i, j)}{\bar{\pi}(S)} - 2\sqrt{\frac{R(i, j)}{\bar{\pi}(S)}}\sqrt{\frac{\bar{Q}(i, j)}{\bar{\pi}(S)}} \right) \\
& \geq 1 - \frac{3\varepsilon}{16} - \sum_{i,j \in S} \sqrt{\frac{R(i, j)}{\bar{\pi}(S)}}\sqrt{\frac{\bar{Q}(i, j)}{\bar{\pi}(S)}} \\
& \geq \varepsilon - \frac{3\varepsilon}{16} = \frac{13\varepsilon}{16} \geq \frac{\varepsilon^2}{128}
\end{aligned}$$

where in the second inequality we used (8), (7) and that, by assumption,

$$\sum_{i,j \in S} \text{Distribution}(S, \bar{P}, \bar{\pi})(i, j) \geq 1 - \frac{\varepsilon}{16}.$$

Consider now the case

$$\sum_{i,j \in S} \text{Distribution}(S, P, \bar{\pi})(i, j) \leq 1 - \frac{5\varepsilon}{16}.$$

In this case we have

$$\begin{aligned}
d_{\text{TV}}(\text{Distribution}(S, P, \bar{\pi}), \text{Distribution}(S, \bar{P}, \bar{\pi})) & \geq \frac{1}{2} \sum_{i,j \in S} \left(\frac{\bar{Q}(i, j)}{\bar{\pi}(S)} - \frac{R(i, j)}{\bar{\pi}(S)} \right) \\
& \geq \frac{1}{2} \left(1 - \frac{\varepsilon}{16} - 1 + \frac{5\varepsilon}{16} \right) = \frac{\varepsilon}{16}.
\end{aligned}$$

By (1),

$$d_{\text{Hel}}^2(\text{Distribution}(S, P, \bar{\pi}), \text{Distribution}(S, \bar{P}, \bar{\pi})) \geq \frac{\varepsilon^2}{128}.$$

□

C PROOF OF LEMMA 5.5

Denote $m = |T|$ and assume without loss that $T = [m]$. There exists a non negative left Perron vector $u \in \mathbb{R}^m$ corresponding to λ which we may assume to be not descending, i.e.,

$$u(1) \leq u(2) \leq \dots \leq u(m).$$

By abuse of notation we denote by u the vector in \mathbb{R}^d obtained from u by extending it with $d - m$ zeros. Now, define $\hat{u} \in \mathbb{R}^d$ by $\hat{u}(i) = \frac{u(i)}{\pi(i)}, \forall i \in [d]$. It holds

$$\langle u(I - P), \hat{u} \rangle = (1 - \lambda) \langle u, \hat{u} \rangle. \tag{9}$$

The right-hand side of (9) equals

$$(1 - \lambda) \sum_{i \in [d]} \pi(i) \hat{u}(i)^2$$

while the left-hand side is bounded below by

$$\sum_{1 \leq i < j \leq d} Q(i, j) (\hat{u}(i) - \hat{u}(j))^2.$$

Thus,

$$1 - \lambda \geq \frac{\sum_{1 \leq i < j \leq d} Q(i, j) (\hat{u}(i) - \hat{u}(j))^2}{\sum_{i \in [d]} \pi(i) \hat{u}(i)^2}.$$

Now,

$$\sum_{1 \leq i < j \leq d} Q(i, j) (\hat{u}(i) + \hat{u}(j))^2 \leq 2 \sum_{i \in [d]} \pi(i) \hat{u}(i)^2.$$

It follows that

$$1 - \lambda \geq \frac{1}{2} \left(\frac{\sum_{1 \leq i < j \leq d} Q(i, j) (\hat{u}(i)^2 - \hat{u}(j)^2)}{\sum_{i \in [d]} \pi(i) \hat{u}(i)^2} \right)^2.$$

Now, for $1 \leq k \leq d - 1$ let $S_k = [k]$. Then

$$\begin{aligned} \sum_{1 \leq i < j \leq d} Q(i, j) (\hat{u}(i)^2 - \hat{u}(j)^2) &= \sum_{k=1}^m (\hat{u}(k)^2 - \hat{u}(k+1)^2) \sum_{i \in S_k, j \in \bar{S}_k} Q(i, j) \\ &\geq \alpha \sum_{k=1}^m (\hat{u}(k)^2 - \hat{u}(k+1)^2) \sum_{l=1}^k \pi(l) \\ &= \alpha \sum_{l=1}^m \pi(l) \sum_{k=l}^m (\hat{u}(k)^2 - \hat{u}(k+1)^2) \\ &= \alpha \sum_{l=1}^d \pi(l) \hat{u}(l)^2. \end{aligned}$$

□

D PROOF OF LEMMA 5.6

Let π_T be the vector obtained from π by keeping only the entries at indices belonging to T and let λ be the largest eigenvalue of P_T . By Lemma 5.5, $\lambda \leq 1 - \frac{\alpha^2}{2}$. Let $k = \frac{8 \log \frac{1}{(\pi_T)_*}}{\alpha^2}$ and let $j \in [k]$. Define

$$Y_j = 1\{\exists i \in [(j-1)k+1, jk] \mid X_i \notin T\}.$$

We distinguish between two cases:

1. $X_{(j-1)k+1} \notin T$: In this case, $\mathbb{P}(Y_j = 1 \mid X_1, \dots, X_{(j-1)k+1}) = 1$.
2. $X_{(j-1)k+1} = i \in T$: In this case, with $D = \text{diag}(\pi_T)$ where

$$\begin{aligned} \mathbb{P}(Y_j = 0 \mid X_1, \dots, X_{(j-1)k+1}) &= e_i^T P_T^{k-1} \mathbf{1} \\ &= e_i^T D^{-1} D P_T^{k-1} \mathbf{1} \\ &= \langle P_T^{k-1} \mathbf{1}, D^{-1} e_i \rangle_{\pi_T} \\ &= \sqrt{\frac{\pi(T)}{\pi(i)}} \left\langle P_T^{k-1} \frac{1}{\sqrt{\pi(T)}} \mathbf{1}, \sqrt{\pi(i)} D^{-1} e_i \right\rangle_{\pi_T}. \end{aligned} \quad (10)$$

Notice that $\left\| \frac{1}{\sqrt{\pi(T)}} \mathbf{1} \right\|_{\pi_T} = \left\| \sqrt{\pi(i)} D^{-1} e_i \right\|_{\pi_T} = 1$. By the Courant–Fischer principle (e.g. Helmberg (2008, p. 219)),

$$(10) \leq \frac{1}{\sqrt{(\pi_T)_*}} \left(1 - \frac{\alpha^2}{2} \right)^{k-1} \leq \frac{1}{2}$$

where in the first inequality we also used that $\sqrt{\frac{\pi(T)}{\pi(i)}} \leq \frac{1}{\sqrt{(\pi_T)_*}}$ and in the second the definition of k .

Combining the two cases we obtain, $\mathbb{P}(Y_j = 1 \mid X_1, \dots, X_{(j-1)k}) \geq \frac{1}{2}$. Thus,

$$\mathbb{P}\left(\sum_{i=1}^m \mathbf{1}\{X_i \notin T\} \geq \frac{m}{4k}\right) \geq \mathbb{P}\left(\sum_{i=1}^{m/k} Y_i \geq \frac{m}{4k}\right) \geq 1 - \delta$$

where in the second inequality we used an adaptation of Hoeffding's inequality stated and proved in the following lemma. \square

Lemma D.1. *Let $\alpha \geq 0$ and let $(B_t)_{t \in \mathbb{N}}$ be Bernoulli random variables, not necessarily independent, such that for every $t \in \mathbb{N}$ it holds*

$$\min_{b_1, \dots, b_{t-1} \in \{0,1\}} \mathbb{E}[B_t \mid B_1 = b_1, \dots, B_{t-1} = b_{t-1}] \geq \alpha.$$

Then

$$\mathbb{P}\left(\sum_{t=1}^n B_t \leq \frac{\alpha n}{2}\right) \leq e^{-\frac{\alpha^2 n}{2}}.$$

Proof. Let $n \in \mathbb{N}$ and $\lambda > 0$. By Markov's inequality,

$$\mathbb{P}\left(\sum_{t=1}^n B_t \leq \frac{\alpha n}{2}\right) \leq e^{\frac{\lambda \alpha n}{2}} \mathbb{E}\left[\prod_{t=1}^n e^{-\lambda B_t}\right].$$

Now,

$$\begin{aligned} \mathbb{E}\left[\prod_{t=1}^n e^{-\lambda B_t}\right] &= \sum_{b_1, \dots, b_n \in \{0,1\}} \prod_{t=1}^n e^{-\lambda b_t} \mathbb{P}(B_1 = b_1, \dots, B_n = b_n) \\ &= \sum_{b_1 \in \{0,1\}} e^{-\lambda b_1} \mathbb{P}(B_1 = b_1) \sum_{b_2 \in \{0,1\}} e^{-\lambda b_2} \mathbb{P}(B_2 = b_2 \mid B_1 = b_1) \cdots \\ &\quad \sum_{b_n \in \{0,1\}} e^{-\lambda b_n} \mathbb{P}(B_n = b_n \mid B_1 = b_1, \dots, B_{n-1} = b_{n-1}). \end{aligned}$$

For the last term in the above equality it holds

$$\sum_{b_n \in \{0,1\}} e^{-\lambda b_n} \mathbb{P}(B_n = b_n \mid B_1 = b_1, \dots, B_{n-1} = b_{n-1}) = \mathbb{E}_{B_n \mid B_1=b_1, \dots, B_{n-1}=b_{n-1}} [e^{-\lambda B_n}] \leq e^{-\lambda \alpha + \frac{\lambda^2}{8}}$$

where the inequality is due to Hoeffding's lemma (e.g. Massart (2007, Lemma 2.6)). Proceeding inductively, we obtain

$$\mathbb{E}\left[\prod_{t=1}^n e^{-\lambda B_t}\right] \leq \left(e^{-\lambda \alpha + \frac{\lambda^2}{8}}\right)^n.$$

Thus,

$$\mathbb{P}\left(\sum_{t=1}^n B_t \leq \frac{\alpha n}{2}\right) \leq \left(e^{-\lambda \alpha + \frac{\lambda^2}{4}}\right)^{n/2} \leq e^{-\frac{\alpha^2 n}{2}}.$$

where in the second inequality we used that $-\lambda \alpha + \frac{\lambda^2}{4}$ attains a minimum of $-\alpha^2$ at $\lambda = 2\alpha$. \square

E PROOF OF LEMMA 5.13

Let $m \in \mathbb{N}$. Consider first the stationary case, i.e., $\mu = \pi$. By Paulin (2015, Theorem 3.4),

$$\mathbb{P}\left(|N_m(i) - m\pi(i)| \geq \frac{1}{2}\pi(i)\right) \leq 2 \exp\left(-\frac{m\gamma\pi(i)}{36}\right).$$

Thus, there exists a universal constant c' such that if $m > \frac{c' \log \frac{d}{\delta}}{\pi(i)\gamma}$ then

$$\mathbb{P}\left(N_m(i) \notin \left(\frac{1}{2}m\pi(i), \frac{3}{2}m\pi(i)\right)\right) < \frac{\delta}{d}.$$

Now, in order to accommodate for non-stationary chains, by Paulin (2015, Proposition 3.10), we need to replace $\log \frac{d}{\delta}$ with

$$\log \frac{d\|\mu/\pi\|_{2,\pi}}{\delta} \leq \log \frac{d}{\delta\pi_\star} \leq 2 \log \frac{1}{\delta\pi_\star}.$$

Replacing $\pi(i)$ with π_\star and using the union bound proves the assertion. \square

F PROOF OF LEMMA 5.14

It holds

$$\begin{aligned} \mathbb{P}(\exists i \in [d] \text{ s.t. } v(i) > 2mp(i)) &\leq d \max_{i \in [d]} \mathbb{P}(v(i) > 2mp(i)) \\ &\leq d \exp\left(-\frac{m^2 p_\star^2}{2(mp_\star(1-p_\star) + \frac{mp_\star}{3})}\right) \\ &\leq d \exp\left(-\frac{mp_\star}{4}\right) < \delta \end{aligned}$$

where the second inequality is due to Bernstein's inequality (e.g. Dubhashi and Panconesi (2009, Theorem 1.2)). \square

G PROOF OF PROPOSITION 7.1

(i) Let $i, j \in [d]$. It holds

$$\begin{aligned} \left(\sqrt{P' \circ \bar{P}'}\right)(i, j) &= \sqrt{((1-\alpha)P(i, j) + \alpha 1\{i=j\})((1-\alpha)\bar{P}(i, j) + \alpha 1\{i=j\})} \\ &= \sqrt{(1-\alpha)^2 P(i, j)\bar{P}(i, j) + \alpha 1\{i=j\}((1-\alpha)P(i, j) + (1-\alpha)\bar{P}(i, j) + \alpha)} \\ &\leq \sqrt{P(i, j)\bar{P}(i, j)} + \sqrt{2\alpha} 1\{i=j\} \\ &= \left(\sqrt{P \circ \bar{P}} + \sqrt{2\alpha} I\right)(i, j). \end{aligned}$$

Now,

$$\rho\left(\sqrt{P' \circ \bar{P}'}\right) \leq \rho\left(\sqrt{P \circ \bar{P}} + \sqrt{2\alpha} I\right) = \rho\left(\sqrt{P \circ \bar{P}}\right) + \sqrt{2\alpha}$$

where the inequality is due to the monotonicity of the spectral radius (e.g. Horn and Johnson (2012, Theorem 8.1.18)). Thus,

$$\text{Distance}(P', \bar{P}') \geq \text{Distance}(P, \bar{P}) - \sqrt{2\alpha} = \frac{\varepsilon}{2}.$$

(ii) Using that the spectral radius is invariant under transposition and matrix similarity,

$$\begin{aligned}
 \rho\left(\sqrt{P^* \circ \bar{P}^*}\right) &= \rho\left(\sqrt{P^* \circ \bar{P}^{*T}}\right) \\
 &= \rho\left(\sqrt{D_\pi P D_{\pi^{-1}} \circ D_{\bar{\pi}} \bar{P} D_{\bar{\pi}^{-1}}}\right) \\
 &= \rho\left(D_{\sqrt{\pi \circ \bar{\pi}}} \sqrt{P \circ \bar{P}} D_{\sqrt{\pi \circ \bar{\pi}}}^{-1}\right) \\
 &= \rho\left(\sqrt{P \circ \bar{P}}\right).
 \end{aligned}$$

(iii) Let $i, j \in [d]$. Then

$$\sqrt{P \circ P(\alpha)}(i, j) = \sqrt{P \circ (\alpha P + (1 - \alpha)\bar{P})}(i, j) \leq \sqrt{\alpha}P(i, j) + \sqrt{1 - \alpha}\sqrt{P \circ \bar{P}}(i, j)$$

where the inequality is due to the subadditivity of the function $x \mapsto \sqrt{x}$. By Horn and Johnson (2012, Theorem 8.1.18),

$$\rho\left(\sqrt{P \circ P(\alpha)}\right) \leq \rho\left(\sqrt{\alpha}P + \sqrt{1 - \alpha}\sqrt{P \circ \bar{P}}\right).$$

Now, with $D = \text{diag}(\pi)^{\frac{1}{2}}$ and $\bar{D} = \text{diag}(\bar{\pi})^{\frac{1}{2}}$, both $D^{-1}PD$ and $\sqrt{(D\bar{D})^{-1}P \circ \bar{P}(D\bar{D})}$ are symmetric. Furthermore, for every $i, j \in [d]$ it holds

$$\begin{aligned}
 \left(D^{-1}\sqrt{P \circ \bar{P}}D\right)(i, j) &= \pi(j)^{-\frac{1}{2}}\sqrt{P(i, j)\bar{P}(i, j)}\pi(i)^{\frac{1}{2}} \\
 &= \sqrt{\pi(j)^{-\frac{1}{2}}P(i, j)\pi(i)^{\frac{1}{2}}\bar{\pi}(j)^{-\frac{1}{2}}\bar{P}(i, j)\bar{\pi}(i)^{\frac{1}{2}}}\sqrt{\frac{\pi(i)\bar{\pi}(j)}{\bar{\pi}(i)\pi(j)}} \\
 &\leq \sqrt[4]{\frac{1 + \varepsilon}{1 - \varepsilon}}\sqrt{(D\bar{D})^{-1}P \circ \bar{P}(D\bar{D})}(i, j).
 \end{aligned}$$

Thus,

$$\begin{aligned}
 \rho\left(\sqrt{\alpha}P + \sqrt{1 - \alpha}\sqrt{P \circ \bar{P}}\right) &= \rho\left(\sqrt{\alpha}D^{-1}PD + \sqrt{1 - \alpha}D^{-1}\sqrt{P \circ \bar{P}}D\right) \\
 &\leq \rho\left(\sqrt{\alpha}D^{-1}PD + \sqrt{1 - \alpha}\sqrt[4]{\frac{1 + \varepsilon}{1 - \varepsilon}}\sqrt{(D\bar{D})^{-1}P \circ \bar{P}(D\bar{D})}\right) \\
 &\leq \sqrt{\alpha}\rho(D^{-1}PD) + \sqrt{1 - \alpha}\sqrt[4]{\frac{1 + \varepsilon}{1 - \varepsilon}}\rho\left(D^{-1}\sqrt{P \circ \bar{P}}D\right) \\
 &= \sqrt{\alpha} + \sqrt{1 - \alpha}\sqrt[4]{\frac{1 + \varepsilon}{1 - \varepsilon}}\rho\left(\sqrt{P \circ \bar{P}}\right)
 \end{aligned}$$

where the first inequality is due to the monotonicity of the spectral radius (e.g. Horn and Johnson (2012, Theorem 8.1.18)) and the second inequality is due to the fact that the spectral radius is subadditive for symmetric matrices (e.g. Olver et al. (2006, Theorem 9.21 together with Exercise 9.2.37)).

(iv) By monotonicity of the spectral radius (e.g. Zhan (2013, Corollary 6.15)), it suffices to show that for each $i, j \in [d]$ it holds

$$\sqrt{P \circ \bar{P}}^k(i, j) \leq \sqrt{P^k \circ \bar{P}^k}(i, j).$$

We proceed by induction. The claim is obviously true for $k = 1$. Suppose it holds for $k \in \mathbb{N}$. We have

$$\begin{aligned} \sqrt{P \circ \bar{P}^{k+1}}(i, j) &= \sum_{l=1}^d \sqrt{P \circ \bar{P}}(i, l) \sqrt{P \circ \bar{P}^k}(l, j) \\ &\leq \sum_{l=1}^d \sqrt{P(i, l)P^k(l, j)} \sqrt{\bar{P}(i, l)\bar{P}^k(l, j)} \\ &\leq \sqrt{\sum_{l=1}^d P(i, l)P^k(l, j)} \sqrt{\sum_{l=1}^d \bar{P}(i, l)\bar{P}^k(l, j)} \\ &= \sqrt{P_1^{k+1} \circ P_2^{k+1}}(i, j) \end{aligned}$$

where we used the induction hypothesis in the first inequality and Cauchy-Schwarz in the second.

(v) It holds

$$P^k \xrightarrow[k \rightarrow \infty]{} \begin{pmatrix} - & \pi & - \\ & \vdots & - \\ - & \pi & - \end{pmatrix} \text{ and } \bar{P}^k \xrightarrow[k \rightarrow \infty]{} \begin{pmatrix} - & \bar{\pi} & - \\ & \vdots & - \\ - & \bar{\pi} & - \end{pmatrix}$$

where convergence is under the infinity norm over matrices. Thus,

$$\text{Distance}(P^k, \bar{P}^k) \xrightarrow[k \rightarrow \infty]{} d_{\text{Hel}}^2(\pi, \bar{\pi}).$$

(vi) It holds

$$\begin{aligned} \sqrt{P^\dagger \circ \bar{P}^\dagger} &= \sqrt{P^* P \circ \bar{P}^* \bar{P}} \\ &\geq \sqrt{P^* \circ \bar{P}^*} \sqrt{P \circ \bar{P}} \\ &= \sqrt{D_\pi^{-1} P^T D_\pi \circ D_{\bar{\pi}}^{-1} \bar{P}^T D_{\bar{\pi}}} \sqrt{P \circ \bar{P}} \\ &= D_{\sqrt{\pi \circ \bar{\pi}}}^{-1} \sqrt{P \circ \bar{P}}^T D_{\sqrt{\pi \circ \bar{\pi}}} \sqrt{P \circ \bar{P}} \\ &= \sqrt{P \circ \bar{P}}^* \sqrt{P \circ \bar{P}} \\ &= \sqrt{P \circ \bar{P}^\dagger} \end{aligned}$$

where the first inequality is due to Drnovšek and Peperko (2006, Theorem 2.2). It follows that

$$\rho\left(\sqrt{P^\dagger \circ \bar{P}^\dagger}\right) \geq \rho\left(\sqrt{P \circ \bar{P}^\dagger}\right) \geq \rho\left(\sqrt{P \circ \bar{P}}\right)^2$$

where the second inequality is due to Weyl (e.g. Gohberg and Kreĭn (1978, Lemma 3.3)).

(vii) For $\alpha \in (0, 1)$ define

$$P = \begin{pmatrix} 1 - \alpha & \alpha \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \text{ and } \bar{P} = \begin{pmatrix} 1 - \alpha & \alpha \\ \alpha & 1 - \alpha \end{pmatrix}.$$

Clearly, P, \bar{P} are irreducible and reversible and one verifies easily that

- (a) The stationary distributions of P, \bar{P} are $\pi = \left(\frac{\frac{1}{2}}{\frac{1}{2} + \alpha}, \frac{\alpha}{\frac{1}{2} + \alpha}\right)$, $\bar{\pi} = \left(\frac{1}{2}, \frac{1}{2}\right)$, respectively.
- (b) $1 - \sum_{i \in [2]} \sqrt{\pi(i)\bar{\pi}(i)} \xrightarrow[\alpha \rightarrow 0]{} 1 - \frac{1}{\sqrt{2}} > \frac{1}{4}$.
- (c) $\text{Distance}(P, \bar{P}) \xrightarrow[\alpha \rightarrow 0]{} 0$.

□

H ON THE IIDGENERATOR ALGORITHM

Let $P \in \mathcal{M}_{\text{irr}}$ with stationary distribution π and let $\mu \in \Delta_d$. Let $(X_t)_{t \in \mathbb{N}} \sim (P, \mu)$. For $i \in [d]$ we define the *first hitting time for i* to be

$$\tau_1^{(i)} = \inf \{t \geq 1 : X_t = i\}$$

(cf. Levin and Peres (2017, p. 11)) and for $1 < s \in \mathbb{N}$ we define the *sth hitting time for i* to be

$$\tau_s^{(i)} = \inf \left\{ t > \tau_{s-1}^{(i)} : X_t = i \right\}.$$

Following Daskalakis et al. (2018a, pp. 12-13), for $v \in \mathbb{N}^d$ we define the mapping $\Psi_v : \{0, 1, \dots, d\}^\infty \rightarrow \prod_{i=1}^d [d]^{v(i)}$ by

$$(X_t)_{t \in \mathbb{N}} \mapsto \left(\left(X_{\tau_t^{(1)}+1} \right)_{t \in [v(1)]}, \left(X_{\tau_t^{(2)}+1} \right)_{t \in [v(2)]}, \dots, \left(X_{\tau_t^{(d)}+1} \right)_{t \in [v(d)]} \right).$$

The map outputs, when given a trajectory sampled from P , for each $i \in [d]$, the first state that has been visited immediately after hitting i , for each of the first $v(i)$ visits to i . It is a consequence of the Markov property that all the coordinates of $\Psi_v((X)_{t \in \mathbb{N}})$ are independent and that for each $i \in [d]$, the i th entry of the d -tuple consists of a sample that is iid according to the conditional distribution defined by the i th row of P . That is, for $i \in [d]$ it holds

$$\left(X_{\tau_t^{(i)}+1} \right)_{t \in [v(i)]} \sim P(i, \cdot)^{\otimes v(i)}.$$

The mapping Ψ_v allows us to sample from the edge measure $Q = \text{diag}(\pi)P$ of P as described in the following two-stages procedure: Let $m \in \mathbb{N}$ be the desired size of the sample. First sample from m iid random variables $Z_1, \dots, Z_m \sim \pi^{\otimes m}$ and denote by v the corresponding histogram, i.e., $v(i) = \sum_{t=1}^m \mathbf{1}[Z_t = i]$, for every $i \in [d]$. Second, define $\Phi_\pi : [d]^\infty \rightarrow ([d] \times [d])^m$ by

$$(X_t)_{t \in \mathbb{N}} \mapsto (Y_k)_{k \in [m]},$$

where

$$Y_k = \left(Z_k, \left((\Psi_v(X_t)_{t \in \mathbb{N}})(Z_k) \left(1 + \sum_{j=1}^{k-1} \mathbf{1}[Z_j = Z_k] \right) \right) \right), \quad \forall k \in [m].$$

Then $(Y_k)_{k \in [m]} \sim Q^{\otimes m}$. For an infinite trajectory, the mapping Φ_π is well-defined almost surely since, by assumption, P is irreducible. However, applied on a finite trajectory $(X_t)_{t=1}^n$ for some $n \in \mathbb{N}$, the well-definedness of $\Phi_\pi((X_t)_{t=1}^n)$ is a random event that strongly depends on n and on the properties of the Markov chain P .

Notice that the procedure above may be performed with any probability distribution $\nu \in \Delta_d$ (instead of π). In this case the sampling is from $\text{diag}(\nu)P$ (instead of Q). Now, recall that in Lemma 5.16 it is assumed that one has access to $m \in \mathbb{N}$ iid samples from $\text{Distribution}(S, P, \bar{\pi})$ where $\emptyset \neq S \subseteq [d]$, $\bar{\pi}$ is the stationary distribution of the reference Markov chain and P is the transition matrix of the unknown Markov chain. To achieve this, we define $\nu \in \Delta_d$ as follows:

$$\nu(i) = \begin{cases} \frac{\pi(i)}{\pi(S)} & i \in S \\ 0 & \text{otherwise,} \end{cases} \quad \forall i \in [d] \quad (11)$$

and apply Φ_ν on a trajectory $(X_t)_{t \in \mathbb{N}}$. This gives $(Y_k)_{k \in [m]}$ in which we replace every Y_k whose second coordinate does not belong to S with the symbol ∞ . The pseudocode for the procedure of sampling from $\text{Distribution}(S, P, \bar{\pi})$ is given in Cherapanamjeri and Bartlett (2019, Algorithm 1). One only needs to replace the line

“ $v \leftarrow l$ samples from $\text{Uniform}(T)$ ”

with

“ $v \leftarrow l$ samples from ν ”

where ν is defined in (11).

I ALGORITHM FOR MARKOV CHAIN IDENTITY TESTING

Input: $d, \varepsilon, \bar{P}, (X_t)_{t \in [m]}$
Output: ACCEPT = 0 or REJECT = 1
 $(S, T) \leftarrow \text{Partition}([d], \Theta(\varepsilon))$
 $\bar{\pi} \leftarrow \text{StationaryDistribution}(\bar{P})$
for $S \in \mathcal{S}$ **do**
 $\nu \leftarrow [0]_d$
 for $i \in [d]$ **do**
 if $i \in S$ **then**
 $\nu[i] \leftarrow \bar{\pi}[i] / \sum_{j \in S} \bar{\pi}[j]$
 end
 end
 $(\bar{\pi}_S)_* \leftarrow \max(\nu)$
 $l \leftarrow \tilde{\mathcal{O}}(1/\varepsilon^2(\bar{\pi}_S)_*)$
 $Y \leftarrow \text{iidGenerator}((X_t)_{t \in [m]}, S, \nu, l)$
 if $Y \neq \text{FALSE}$ **then**
 return $\text{iidTester}(Y, \text{Distribution}(S, \bar{P}, \bar{\pi}), \Theta(\varepsilon^2), \Theta(1/d))$
 end
end
return REJECT

Remark I.1. 1. “ACCEPT” means that the trajectory $(X_t)_{t \in [m]}$ from the unknown Markov chain P suggests that $P = \bar{P}$ and “REJECT” means that probably $\text{Distance}(P, \bar{P}) \geq \varepsilon$.

2. The StationaryDistribution algorithm is any algorithm that upon receiving the transition matrix of an irreducible Markov chain returns its stationary distribution.
3. The Partition algorithm corresponds to Cherapanamjeri and Bartlett (2019, Algorithms 3 and 4) and its necessary modifications are described in Section 5.2.
4. The assignment to l follows Lemma 5.15.
5. The iidGenerator algorithm corresponds to Daskalakis et al. (2018b, Algorithm 1) and is described in full detail in Section H.
6. The iidTester algorithm corresponds to (Daskalakis et al., 2018b, Algorithm 1). See also Section 5.3 and Lemma 5.16.