

Hypothesis testing for Markovian models with random time observations

Flavia Barsotti¹, Anne Philippe² and Paul Rochet²

Abstract

The aim of this paper is to propose a methodology for testing general hypothesis in a Markovian setting with random sampling. A discrete Markov chain X is observed at random time intervals τ_k , assumed to be iid with unknown distribution μ . Two test procedures are investigated. The first one is devoted to testing if the transition matrix P of the Markov chain X satisfies specific affine constraints, covering a wide range of situations such as symmetry or sparsity. The second procedure is a goodness-of-fit test on the distribution μ , which reveals to be consistent under mild assumptions even though the time gaps are not observed. The theoretical results are supported by a Monte Carlo simulation study to show the performance and robustness of the proposed methodologies on specific numerical examples.

Keywords: Asymptotic Tests of Statistical Hypotheses ; Markov Chain ; Random Sampling.

1 Introduction

It has been widely recognized that discrete Markov chains are a powerful probabilistic tool to study many real phenomena in different application fields. Interests from economics, finance, insurance and also medical research are just few examples. From a theoretical point of view, alternative technical frameworks have been defined and developed in the literature to perform statistical inference in Markovian models: multiple Markov chains [2, 4], hidden Markov processes [5, 6, 13], random walks on graphs [10] are well known mathematical settings used to describe the evolution of real events.

In this paper, we investigate hypothesis testing issues in a Markov model with random sampling. In the spirit of [3], a discrete homogenous Markov chain $(X_m)_{m \in \mathbb{N}}$ is observed at random times so that the only available observations consist in a sub-sequence of the initial process. Hereafter we denote by $(Y_k)_{k \in \mathbb{N}}$ the observed process. The time gaps τ_k (i.e. the number of jumps) between two consecutive observations are assumed non-negative, independent and identically distributed from an unknown distribution μ . A statistical methodology is proposed to test specific hypothesis on the transition matrix P or the distribution μ having observed (Y_1, \dots, Y_n) .

¹Risk Methodologies, Group Financial Risks, Group Risk Management, UniCredit S.p.A, 20154 Milano
The views expressed in this paper are those of the author and should not be attributed to UniCredit Group or to the author as representative or employee of UniCredit Group.

²Université de Nantes, Laboratoire de mathématiques Jean Leray

Our interest is to answer the following questions: can we test some specific hypothesis on the transition matrix P of the initial chain X when neither the time gaps τ_k nor their distribution μ are known? Additionally, can we perform a goodness-of-fit test on the distribution μ of the time gaps in this framework?

As explained in [3], a typical application of this setting occurs when considering a continuous time Markov chain $Z = (Z_t)_{t>0}$ observed at a discrete time grid $t_1 < \dots < t_n$. In this situation, X represents the jump process of Z and $(\tau_k)_{k \in \mathbb{N}}$ denotes the number of jumps occurring between two consecutive observations, $Y_k := Z_{t_k}$. If the discrete time grid is chosen independently from the chain, the time gaps τ_k are independent random variables *unknown* to the practitioner. Their distribution μ is Poisson if Z is observed at regularly spaced time intervals, but one can easily imagine a different distribution μ if t_1, \dots, t_n are affected by undesired random effects. This framework can describe many real phenomena: chemical reactions [1], financial markets [12], waiting lines in queuing theory [9], medical studies [7].

In [3], the sparsity of P reveals to be a crucial assumption for the proposed estimation methodology. In this paper, we extend the statistical setting defined in [3] by relaxing the sparsity assumption and working under a more general framework where the information about P can be expressed in the form of affine constraints. It is then possible to incorporate some prior information on P such as symmetry, reflexivity or even simply the knowledge of some of its entries. The sparsity hypothesis considered in [3] can be retrieved as a particular case.

In our model, the observations Y_1, \dots, Y_n are the realizations of an homogenous Markov chain, with a transition probability Q that can be written as an analytic function of P . As in [3], neither the transition matrix P nor the distribution μ are known and the time gaps τ_k between two consecutive observations Y_{k-1}, Y_k are unavailable. Our aim is now to perform hypothesis tests on both the transition kernel P and the time gaps' distribution μ . Theoretical results are separated in two consecutive steps. As a first step, we describe how to build an hypothesis test on P to study if the transition kernel P satisfies additional affine constraints. This framework can be used to test a wide range of hypothesis such as symmetry, sparsity or even to test a particular value of P or some of its entries. As a second step, we propose a procedure to test specific values for the distribution of time gaps μ , considering an hypothesis of the form $\mathcal{H}_0 : \mu = \mu_0$, given some preliminary information on P . Both tests are asymptotically exact and rely on the asymptotic normality of the empirical transition matrix \hat{Q} . A Monte Carlo simulation analysis is performed in order to highlight the performance of the proposed test.

The paper is organized as follows: Section 2 provides a detailed description of the statistical problem and discusses the identifiability issues encountered in this framework. Section 3 defines an asymptotically exact hypothesis test for the transition kernel P . Section 4 shows the theoretical framework to build an hypothesis test on the distribution μ . Section 5 supports the analysis of Sections 3 and 4 with numerical results from a Monte Carlo simulation study. Proofs are postponed in the Appendix.

2 The problem

We consider an irreducible homogenous Markov chain $X = (X_m)_{m \in \mathbb{N}}$ with finite state space $\mathcal{E} = \{1, \dots, N\}$, $N \geq 3$ and transition matrix P . We assume that X is observed at random times T_1, \dots, T_n so that the only available observations consist in the sub-sequence $Y_k := X_{T_k}$, $k = 1, \dots, n$. The numbers of jumps $\tau_k := T_{k-1} - T_k$ between two observations Y_k are assumed to be iid random variables with distribution μ on \mathbb{N} and independent from X . In this setting, the resulting process $Y = (Y_k)_{k \in \mathbb{N}}$ remains Markovian (see [3]) with transition matrix

$$Q := G_\mu(P) = \sum_{\ell \geq 0} P^\ell \mu(\ell), \quad (1)$$

where $G_\mu : [-1, 1] \rightarrow \mathbb{R}$ denotes the generator function of μ . We assume that some information is available on P in the form of an affine constraint

$$Ap = b$$

where $p = \text{vec}(P) = (P_{11}, \dots, P_{N1}, \dots, P_{1N}, \dots, P_{NN})^\top$ is the vectorization of P , $A \in \mathbb{R}^{r \times N^2}$ is a known full ranked matrix and $b \in \mathbb{R}^r$ a known vector. This affine constraint defines the model

$$\mathcal{M} := \{m \in \mathbb{R}^{N^2} : Am = b\}, \quad (2)$$

that is, the set of admissible values for $p = \text{vec}(P)$. Each additional affine condition satisfied by P reduces by one the dimension of the model. Therefore, the rank r of A indicates the information available on P . In the sequel, we denote by d the dimension of the model \mathcal{M} , given by $d = \dim(\mathcal{M}) = N^2 - r$.

Many different forms of information on the initial chain X_m can be expressed by affine constraints, as described in the following examples.

- a) $P\mathbf{1} = \mathbf{1}$, where $\mathbf{1} = (1, \dots, 1)^\top$, is an important information that can be expressed as an affine constraint on P by

$$(\mathbf{1}^\top \otimes \mathbf{I}) \text{vec}(P) = \mathbf{1},$$

where \otimes stands for the Kronecker product.

- b) In [7], the author investigates the progression of a disease using a Markov chain observed at unequal time intervals. For this particular application, the transition matrix of the chain is expected to be triangular: in this case the information can be represented with $N(N-1)/2$ linear conditions on P .
- c) If the initial chain X cannot remain at the same state on two consecutive times, the transition matrix P is known to have zero diagonal. This information can be expressed in the form of N linear constraints on P .
- d) More generally, if one or several state transitions are known to be impossible in the initial chain $(X_m)_{m \in \mathbb{N}}$, the nullity of the corresponding entries can be expressed as a set of linear conditions on p . The number of constraints is equal to the number of known zeros in P . This situation is treated in [3].

- e) A symmetric transition matrix P corresponds to $N(N - 1)/2$ linear constraints on $p = \text{vec}(P)$.
- f) P is doubly stochastic, one can include the condition $P^\top \mathbf{1} = \mathbf{1}$ as $N - 1$ additional affine constraints.
- g) X is a reversible Markov chain, its transition matrix P satisfies the linear conditions $\pi_i P_{ij} = \pi_j P_{ji}$ for all $i, j = 1, \dots, N$, where $\pi = (\pi_1, \dots, \pi_N)$ is the invariant measure. Although π is unknown in practice, it can be estimated directly from the observations since it is also the invariant measure of Q .
- h) The knowledge of some entries of P being equal, or equal to a known value $c \in [0, 1]$ can be expected in certain practical situations (e.g. the transition probability distribution from a given state i is uniform). These are of course particular examples of affine constraints on P .

The positivity of the entries of P appears to be more difficult to handle than affine constraints. Nevertheless, the positivity is somewhat less informative as it does not reduce the dimension of the model. In this paper, we choose to neglect this information on P for simplicity.

To avoid critical situations, we assume that $(X_m)_{m \in \mathbb{N}}$ is an aperiodic Markov chain, in which case the transition kernel P has a unique invariant distribution $\pi = (\pi_1, \dots, \pi_N)$, where π_i is positive for all $i = 1, \dots, N$. From the relation $Q = G_\mu(P)$, we know that, like P , Q is aperiodic recurrent and they share the same invariant distribution. Thus, π can be estimated consistently from the observations Y_k , along with Q by the empirical estimators

$$\hat{\pi}_i = \frac{1}{n} \sum_{k=1}^n \mathbb{1}\{Y_k = i\} \quad , \quad \hat{Q}_{ij} = \frac{\sum_{k=1}^{n-1} \mathbb{1}\{Y_k = i, Y_{k+1} = j\}}{\sum_{k=1}^{n-1} \mathbb{1}\{Y_k = i\}} \quad , \quad (3)$$

for all $i, j = 1, \dots, N$, provided that the state i has been observed at least once. It is known (see for instance [11]) that \hat{Q} is asymptotically unbiased and $\hat{q} = \text{vec}(\hat{Q})$ is asymptotically Gaussian,

$$\sqrt{n}(\hat{q} - q) \xrightarrow{d} \mathcal{N}(0, \Sigma).$$

The matrix $\Sigma \in \mathbb{R}^{N^2 \times N^2}$ can be deduced from the asymptotic behavior of the covariances

$$\forall i, j, k, \ell = 1, \dots, N, \quad \lim_{n \rightarrow \infty} n \text{cov}(\hat{Q}_{ij}, \hat{Q}_{k\ell}) = \begin{cases} Q_{ij}(1 - Q_{ij})/\pi_i & \text{if } (i, j) = (k, \ell), \\ -Q_{ij}Q_{i\ell}/\pi_i & \text{if } i = k, j \neq \ell, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

Remark that Σ is a singular matrix since $\hat{q} - q$ can only take values in a linear subspace of $\mathbb{R}^{N^2 \times N^2}$.

Since the distribution of the time gaps τ_k is unknown, the main information available on P is that it commutes with Q . Thus, given the model \mathcal{M} , the possible values for p can be restricted to $\mathcal{M} \cap \text{Com}(Q)$, where $\text{Com}(Q)$ is the commutant of Q . As in [3], we consider the commutation operator $m \mapsto \Delta(Q)m = \text{vec}(MQ - QM)$ where

$$\Delta(Q) = \text{I} \otimes Q - Q^\top \otimes \text{I}.$$

Of course, since Q is unknown, we use the empirical estimator \hat{Q} to retrieve some information on P .

3 Hypothesis tests on P

Given that the transition matrix P of the initial Markov process $(X_m)_{m \in \mathbb{N}}$ belongs to some model \mathcal{M} of the form (2), we want to test additional affine conditions on $p = \text{vec}(P)$, via the null hypothesis $\mathcal{H}_0 : A_0 p = b_0$, where $A_0 \in \mathbb{R}^{k \times N^2}$ is a full ranked matrix, $b_0 \in \mathbb{R}^k$ and $k \geq 1$ is the number of new conditions to test. To avoid technical issues, the new constraints must be compatible with the model, i.e. there exists at least one element $m \in \mathcal{M}$ such that $A_0 m = b_0$. Moreover, we assume without loss of generality that the new constraints are not redundant, in the sense that each row of A_0 brings a new information on P . Formally, this means that the model under \mathcal{H}_0 ,

$$\mathcal{M}_0 := \{m \in \mathbb{R}^{N^2} : Am = b, A_0 m = b_0\} \subset \mathcal{M},$$

is of dimension $d-k$ where we recall $d = \dim(\mathcal{M})$. Let $\phi_1, \dots, \phi_{d-k}$ be a basis of $\ker(A) \cap \ker(A_0)$, extended by $\phi_{d-k+1}, \dots, \phi_d$ to form a basis of $\ker(A)$. The initial model \mathcal{M} and the constrained model \mathcal{M}_0 can be expressed as

$$\mathcal{M} = \{p_0 + \Phi\beta : \beta \in \mathbb{R}^d\} \quad \text{and} \quad \mathcal{M}_0 = \{p_0 + \Phi_0\gamma : \gamma \in \mathbb{R}^{d-k}\},$$

setting $\Phi = (\phi_1, \dots, \phi_d) \in \mathbb{R}^{N^2 \times d}$, $\Phi_0 = (\phi_1, \dots, \phi_{d-k}) \in \mathbb{R}^{N^2 \times (d-k)}$ and with p_0 an element of \mathcal{M}_0 . Define the linear sets

$$\begin{aligned} E &= \text{Im}(\Delta(Q)\Phi) = \{\Delta(Q)m : m \in \ker(A)\} \\ E_0 &= \text{Im}(\Delta(Q)\Phi_0) = \{\Delta(Q)m : m \in \ker(A) \cap \ker(A_0)\} \\ F &= E \cap E_0^\perp, \end{aligned}$$

where E_0^\perp stands for the orthogonal complement of E_0 in \mathbb{R}^{N^2} . Note that since $E_0 \subseteq E$, we have $E = E_0 \oplus F$. Moreover, we shall denote by Π_E , Π_{E_0} and Π_F the orthogonal projectors in \mathbb{R}^{N^2} onto E , E_0 and F respectively. We make the following assumptions.

A1. The problem is identifiable under \mathcal{H}_0 , i.e. $\text{Com}(Q) \cap \ker(A) \cap \ker(A_0) = \{0\}$.

A2. The dimension of E is maximal over the set of stochastic matrices with the same support as Q :

$$\dim(E) = \text{rank}(\Delta(Q)\Phi) = \max_{\substack{Q' : Q' \mathbf{1} = \mathbf{1} \\ \text{supp}(Q') \subseteq \text{supp}(Q)}} \text{rank}(\Delta(Q')\Phi).$$

The equivalence, stated in **A1**, between the identifiability and the matrix $\Delta(Q)\Phi_0$ being of full rank is mentioned in [3]. Here, the identifiability under \mathcal{H}_0 means that P is fully characterized by the condition $p \in \mathcal{M}_0$ and the fact that P and Q commute. Theorem 3.2 in [3] ensures the existence of a consistent estimator of p as soon as the hypothesis \mathcal{H}_0 is true. Remark that unlike in [3], we do not require the initial model \mathcal{M} to be identifiable, although if it is the case, the two assumptions always hold.

Though the assumption **A2** is quite technical, it is very mild, since the set of matrices Q' for which the rank of $\Delta(Q')\Phi$ is maximal is a dense open set. The argument is similar to the one used in the proof of Lemma 2.1 in [3].

Hereafter we denote by $\chi_G^2(W)$ the centered generalized chi-squared distribution

$$\chi_G^2(W) \stackrel{d}{=} \epsilon^\top W \epsilon,$$

where ϵ is a standard Gaussian vector in \mathbb{R}^{N^2} and $W \in \mathbb{R}^{N^2 \times N^2}$. Remark that the above distribution is simply a Dirac mass at zero if W is the null matrix.

Theorem 3.1 *Suppose that **A1** and **A2** hold. Under $\mathcal{H}_0 : A_0 p = b_0$, as $n \rightarrow \infty$ the statistic*

$$S = n \left(\inf_{m \in \mathcal{M}_0} \|\Delta(\hat{Q})m\|^2 - \inf_{m \in \mathcal{M}} \|\Delta(\hat{Q})m\|^2 \right)$$

converges in distribution to a $\chi_G^2(W)$ as $n \rightarrow \infty$ where

$$W := \Pi_F \Delta(P) \Sigma \Delta(P)^\top \Pi_F \quad (5)$$

The proof is given in Appendix.

This result does not rule out the possibility that W is the null matrix (the limit distribution is 0 in this case), however, it is of no use to perform the test in this situation.

Theorem 3.1 alone is not sufficient to build a test for \mathcal{H}_0 since the limit distribution is unknown in practice. Nevertheless, W can be estimated consistently under \mathcal{H}_0 using estimates of Π_F , P and Σ .

- An estimator of Σ can be easily obtained by using the plug in principle. We replace in (4) the parameters π and Q by the empirical versions $\hat{\pi}$ and \hat{Q} defined in (3). Since π_i is positive for all $i = 1, \dots, N$, the resulting estimator $\hat{\Sigma}$ is consistent.
- In [3], the authors show that a consistent estimator of P can be obtained if the problem is identifiable. In our framework, the identifiability under \mathcal{H}_0 assumed in **A1** reduces to

$$\{P\} = \mathcal{M}_0 \cap \text{Com}(Q).$$

From Lemma 6.1 in [3], this condition is actually equivalent to $\Delta(Q)\Phi_0$ being of full rank, in which case

$$\hat{p} = \left(\mathbb{I} - \Phi_0 (\Phi_0^\top \Delta(\hat{Q})^\top \Delta(\hat{Q}) \Phi_0)^{-1} \Phi_0^\top \Delta(\hat{Q})^\top \Delta(\hat{Q}) \right) p_0 \quad (6)$$

is a consistent estimator of p under \mathcal{H}_0 .

- In the proof of Theorem 3.1 (see Appendix) an estimator of Π_F is obtained by considering the linear space

$$\hat{F} = \text{Im}(\Delta(\hat{Q})\Phi) \cap \text{Im}(\Delta(\hat{Q})\Phi_0)^\perp.$$

Under **A1** and **A2**, we prove that the resulting projector $\Pi_{\hat{F}}$ converges in probability to Π_F as $n \rightarrow \infty$.

By plugging the estimates $\hat{\Sigma}$, \hat{P} and $\Pi_{\hat{F}}$ in (5), we build

$$\hat{W} := \Pi_{\hat{F}} \Delta(\hat{P}) \hat{\Sigma} \Delta(\hat{P})^\top \Pi_{\hat{F}}.$$

which is a consistent estimate of W under \mathcal{H}_0 . As a result, the limit distribution of the statistic S can be approximated by $\chi_G^2(\hat{W})$. This leads to the following result.

Corollary 3.2 *Assume that **A1** and **A2** hold. Let $\hat{u}_{1-\alpha}$ denote the quantile of order $1 - \alpha$ of the $\chi_G^2(\hat{W})$ distribution for $\alpha \in (0, 1)$. If $W \neq 0$, we have under $\mathcal{H}_0 : A_0 p = b_0$,*

$$\lim_{n \rightarrow \infty} \mathbb{P}(S > \hat{u}_{1-\alpha}) = \alpha.$$

Moreover, if $\forall p' \in \mathcal{M}_0 : \Delta(Q)p' \neq 0$, then $\lim_{n \rightarrow \infty} \mathbb{P}(S > \hat{u}_{1-\alpha}) = 1$.

The proof is given in Appendix.

Because the quantiles of the generalized χ^2 distribution are not available in an analytic form, they are approximated by Monte-Carlo methods, following [8].

4 Test on μ

The estimation of the transition kernel P is based on the property that P and Q commute, which is due to the existence of an analytic function $G_\mu(\cdot)$ such that $Q = G_\mu(P)$ defined in (1). However, no particular interest has been given so far to the actual value of $G_\mu(\cdot)$, since the knowledge of the distribution μ is not needed to build the test procedure on \hat{P} .

We emphasize that in our framework it is in general not possible to exactly recover the distribution μ , since different distributions μ and ν may lead to the same image $G_\nu(P) = G_\mu(P)$. Nevertheless, questions regarding the number of jumps between consecutive observations, or their distribution μ can naturally arise in practical cases. Suppose for instance that the observations Y_i come from a continuous process $(X_t)_{t \geq 0}$ observed at regular time intervals. Assuming that nothing other than the Y_i 's is known, one might be interested in checking if the underlying process $(X_t)_{t \geq 0}$ is a continuous time Markov chain. This hypothesis relies on the distribution of the time intervals between two consecutive jumps of X_t . If the process $(X_t)_{t \geq 0}$ is in fact Markovian, the times between jumps are drawn from an exponential distribution, resulting in Poisson random variables τ_k . In such a case, testing this hypothesis on μ reduces to verify if G_μ is the generator function of a Poisson variable, that is, if $G_\mu(t) = e^{\lambda(t-1)}$ for some $\lambda > 0$.

In this section we consider hypotheses of the form $\mathcal{H}_0 : \mu = \mu_0$. We assume that the model is identifiable, in which case we can build a consistent estimator of p as follows

$$\hat{p} = \left(\mathbf{I} - \Phi (\Phi^\top \Delta(\hat{Q})^\top \Delta(\hat{Q}) \Phi)^{-1} \Phi^\top \Delta(\hat{Q})^\top \Delta(\hat{Q}) \right) p_0.$$

From Theorem 3.2 in [3], we know that \hat{p} satisfies

$$\sqrt{n} (\hat{p} - p) = \sqrt{n} B(\hat{q} - q) + o_P(1), \quad (7)$$

where $B = \Phi [\Phi^\top \Delta(Q)^\top \Delta(Q) \Phi]^{-1} \Phi^\top \Delta(Q)^\top \Delta(P)$.

The procedure to perform a test on the hypothesis $\mathcal{H}_0 : \mu = \mu_0$ relies on the comparison between \hat{Q} and $G_\mu(\hat{P})$. Since $Q = G_\mu(P)$, the consistency of \hat{P} implies that $\hat{Q} - G_\mu(\hat{P})$ converges in probability to zero. Moreover, we can derive the asymptotic distribution of $g_\mu(\hat{P}) := \text{vec}(G_\mu(\hat{P}))$, leading to the following result.

Proposition 4.1 *If the time gaps distribution μ satisfies the moment condition $\mathbb{E}_\mu(\tau) < \infty$, then the matrix $\Gamma = \sum_{k \geq 1} (\sum_{j=1}^k (P^{j-1})^\top \otimes P^{k-j}) \mu(k)$ is well defined and*

$$\sqrt{n}(\hat{q} - g_\mu(\hat{P})) \xrightarrow{d} \mathcal{N}\left(0, (\mathbf{I} - \Gamma B) \Sigma (\mathbf{I} - \Gamma B)^\top\right).$$

The proof is given in Appendix.

Consequently we get, under \mathcal{H}_0 ,

$$n \|\hat{q} - g_{\mu_0}(\hat{P})\|^2 \xrightarrow{d} \chi_G^2((\mathbf{I} - \Gamma B) \Sigma (\mathbf{I} - \Gamma B)^\top).$$

This result is not sufficient to build the statistical test on the distribution μ . So, we proceed similarly as before by using consistent estimates of Γ , B and Σ to replace the unavailable true values in order to approximate the limit variance. Then, the resulting test statistic has critical region of the form $\{n \|\hat{q} - g_{\mu_0}(\hat{P})\|^2 > \tilde{u}_{1-\alpha}\}$ where $\tilde{u}_{1-\alpha}$ is the quantile of the $\chi_G^2((\mathbf{I} - \hat{\Gamma} \hat{B}) \hat{\Sigma} (\mathbf{I} - \hat{\Gamma} \hat{B})^\top)$ distribution.

5 Computational study

This section is devoted to a Monte Carlo simulation analysis for three hypothesis tests proposed in this paper. We study two examples. The first one evaluates the performance of the test procedure in a situation where X is a random walk. The second example deals with identifiability conditions related to symmetry and bistochasticity.

5.1 Example 1

In this example, the initial Markov chain X behaves like a reflected random walk. The chain takes values in some finite space having $N = 10$ states. Our ultimate aim is to test the

assumption $P = P_0$, where P_0 is given by

$$P_0 = \begin{bmatrix} 0 & 1 & 0 & \dots & \\ 0.5 & 0 & 0.5 & \ddots & \vdots \\ 0 & \ddots & \ddots & \ddots & 0 \\ \vdots & \ddots & 0.5 & 0 & 0.5 \\ \dots & 0 & 1 & 0 & \end{bmatrix}. \quad (8)$$

Assuming that nothing is known on the transition kernel P , we divide the hypothesis tests on P into three steps. First, we want to test if the support of P is contained in that of P_0 , i.e. if only the upper and lower diagonals of P have non-zero entries. Secondly, we investigate the test procedure for the null hypothesis $P = P_0$, under the assumption that the support of P is restricted to that of P_0 . Finally, a third test is carried out for hypothesis on the distribution μ .

We perform the three tests for four sample sizes: $n = 200$, $n = 500$, $n = 1000$ and $n = 2000$. Under the null hypothesis, we simulate the Markov chain $X = (X_m)_{m \in \mathbb{N}}$ with transition matrix P_0 and the times gaps τ_k are drawn from a Poisson distribution $\mu = \mathcal{P}(1)$. Thus, the available observations are $Y_k := X_{T_k}$ where $T_k = \sum_{j=1}^k \tau_j$ for $k = 1, \dots, n$. We recall that neither the initial Markov chain X nor the time gaps τ_k are observed and the distribution μ of the τ_k is unknown to the practitioner. The nominal significance level is fixed equal to $\alpha = 5\%$ for all numerical experiments.

Test 1.

We wish to test if the support of P is contained in that of P_0 , i.e. if only the upper and lower diagonals of P have non-zero entries. The hypothesis is

$$\mathcal{H}_0^{(1)} : \text{supp}(P) \subseteq \text{supp}(P_0),$$

and the model $\mathcal{M}^{(1)}$ is maximal, i.e. it contains all stochastic matrices: $\mathcal{M}^{(1)} = \{\text{vec}(M) : M\mathbf{1} = \mathbf{1}\}$. By Monte-Carlo method, the significance level is estimated for the four sample sizes using 10^4 replications. Results are gathered in Table 4.

n	200	500	1000	2000
$\hat{\alpha}$	0.131 (0.0034)	0.064 (0.0024)	0.059 (0.0024)	0.056 (0.0023)

Table 1: Test for hypothesis $\mathcal{H}_0^{(1)} : \text{supp}(P) \subseteq \text{supp}(P_0)$ in the model $\mathcal{M}^{(1)} = \{\text{vec}(M) : M\mathbf{1} = \mathbf{1}\}$. The table contains summary statistics of Monte Carlo simulation results obtained with $P = P_0$ given in Eq. (8), for the four different sample sizes n and for $\mu = \mathcal{P}(1)$. The probability to reject $\mathcal{H}_0^{(1)}$ is estimated for a nominal size $\alpha = 0.05$, based on 10^4 replications. Standard errors are given in brackets.

Table 4 shows that the estimated significance level converges to the nominal size 5% as n increases. It seems that the test is accurate as soon as the sample sizes is greater than 500.

To evaluate the power of the test, we generate a Markov chain X with transition matrix $P_t = tP_0 + (1-t)C$ where C is a stochastic matrix with full support drawn randomly beforehand (the same value of C is used during the whole simulation study). We estimate the power for different values of t varying in $(0, 1)$ with 0.1 increments. The empirical reject frequencies under the alternative hypothesis are given in Figure 1.

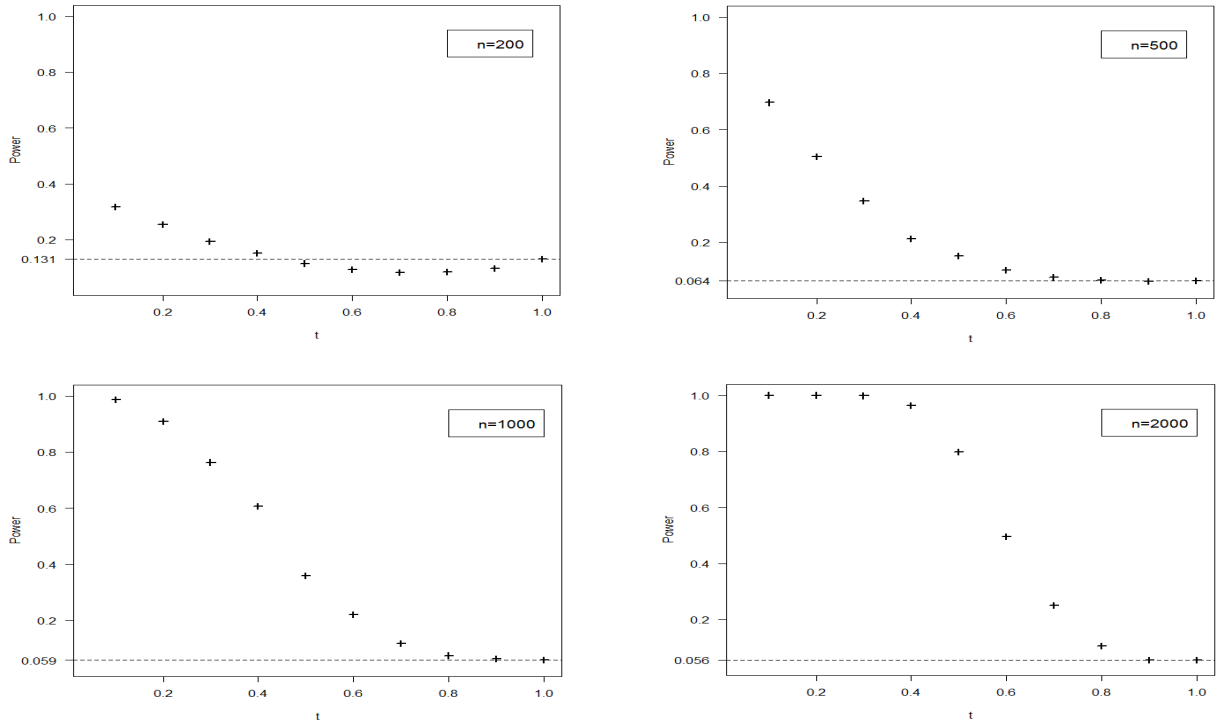


Figure 1: *Test of $\mathcal{H}_0^{(1)} : \text{supp}(P) \subseteq \text{supp}(P_0)$ in the model $\mathcal{M}^{(1)} = \{\text{vec}(M) : M\mathbf{1} = \mathbf{1}\}$.* The graphs represent the estimated power of the test under the alternative $P = tP_0 + (1-t)C$, for the four sample sizes $n = 200$ (top-left), $n = 500$ (top-right), $n = 1000$ (bottom-left) and $n = 2000$ (bottom-right). The power is estimated for a nominal size $\alpha = 0.05$, based on 10^4 replications.

For small sample sizes, e.g. $n = 200$, the estimated type I error probability (≈ 0.13 corresponding to $t = 1$) is far from the nominal value $\alpha = 0.05$. The test is not powerful in this case as we can see that the power remains smaller than 0.4 for $t \in (0, 1)$ and increases in a neighborhood of 1, which indicates a bias. Nevertheless, the theoretical results are verified for sample sizes as large as $n = 500$ and the test is no longer biased. Both the estimated powers and significance level appear to converge to the expected values as n grows.

Test 2.

We investigate the test procedure for the hypothesis

$$\mathcal{H}_0^{(2)} : P = P_0,$$

assuming that the support of P is restricted to that of P_0 . Note that in this case, the first and last row of P are known since they contain only one non-zero element. The model is given by

$$\mathcal{M}^{(2)} = \{ \text{vec}(M) : \text{supp}(M) \subseteq \text{supp}(P_0) \}.$$

The estimated significance levels are given in Table 2.

n	200	500	1000	2000
$\hat{\alpha}$	0.255 (0.0044)	0.104 (0.0031)	0.075 (0.0026)	0.057 (0.0023)

Table 2: *Test for the hypothesis $\mathcal{H}_0^{(2)} : P = P_0$ in the model $\mathcal{M}^{(2)} = \{ \text{vec}(M) : \text{supp}(M) \subseteq \text{supp}(P_0) \}$.* The table contains summary statistics of Monte Carlo simulation results obtained with $P = P_0$ given in Eq. (8), for the four considered sample sizes and for $\mu \sim \mathcal{P}(1)$. The probability of rejecting $\mathcal{H}_0^{(1)}$ is estimated for a significance level $\alpha = 0.05$, based on 10^4 replications.

The convergence seems slower in this case compared to the previous test procedure, with an estimated significance level not quite in a two standard-deviation range of objective value for $n = 2000$. The slow convergence is partly explained by the approximation of the limit distribution by the generalized chi-squared $\chi_G^2(\hat{W})$ obtained with the estimated matrix \hat{W} . Nevertheless, the convergence in distribution of the test statistic to the $\chi_G^2(W)$ is clearly observed in Figure 2.




Figure 2: *Histograms of the 10^4 realizations of the test statistic for the test of $\mathcal{H}_0^{(2)} : P = P_0$ in the model $\mathcal{M}^{(2)} = \{ \text{vec}(M) : \text{supp}(M) \subseteq \text{supp}(P_0) \}$.* The histograms are obtained for $n = 500$ (left), $n = 1000$ (center) and $n = 2000$ (right) with the density of the limit distribution $\chi_G^2(W)$ added in straight line.

Test 3.

We now focus on the distribution μ of the random sampling $(\tau_k)_k$. We want to test the hypothesis

$$\mathcal{H}_0^{(3)} : \mu = \mu_0 = \mathcal{P}(1),$$

in the model $\mathcal{M}^{(2)}$. This means that we perform the test to see if μ is a Poisson distribution with parameter 1, knowing that the support of P is contained in the upper and lower diagonals. This model satisfies the identifiability assumption **A1**.

We estimate the probability of rejecting the null hypothesis when the time gaps τ_k are distributed from a Poisson distribution $\mathcal{P}(\lambda)$, with the parameter λ varying in $(0.5, 1.5)$ with 0.1 increments. Table 6 provides the empirical rejection frequency under the null hypothesis $\lambda = 1$. Figure 3 gives the estimated power of the test as function of the parameter λ .

n	200	500	1000	2000
$\hat{\alpha}$	0.269 (0.0044)	0.081 (0.0027)	0.056 (0.0023)	0.054 (0.0023)

Table 3: Test for the hypothesis $\mathcal{H}_0^{(3)} : \mu = \mathcal{P}(1)$ in the model $\mathcal{M}^{(2)} = \{\text{vec}(M) : \text{supp}(M) \subseteq \text{supp}(P_0)\}$. The table contains summary statistics of Monte Carlo simulation results obtained with $P = P_0$ for the four sample sizes under the null hypothesis. The nominal size of the test is $\alpha = 0.05$, based on 10^4 replications.

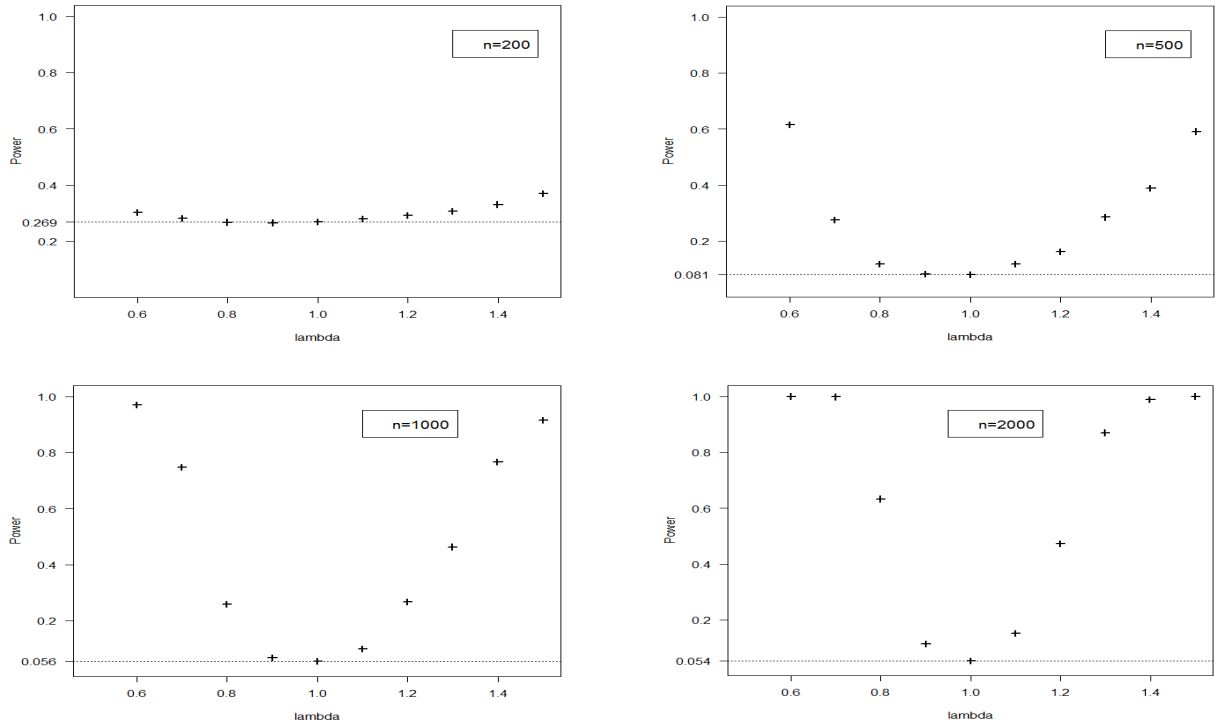


Figure 3: Test for the hypothesis $\mathcal{H}_0^{(3)} : \mu = \mathcal{P}(1)$ in the model $\mathcal{M}^{(2)} = \{\text{vec}(M) : \text{supp}(M) \subseteq \text{supp}(P_0)\}$. The graphs represent the estimated power under the alternative $\mu = \mathcal{P}(\lambda)$, for λ in $(0.5, 1.5)$ and for $n = 200$ (top-left), $n = 500$ (top-right), $n = 1000$ (bottom-left) and $n = 2000$ (bottom-right). The power is estimated for a nominal size $\alpha = 0.05$, based on 10^4 replications.

The power of the test is not satisfactory for $n = 200$ but it increases significantly with n . Although the time gaps are not observed, this simulation confirms that performing a goodness-

of-fit test on their distribution is still possible in this framework. In particular, the intensity of the process assuming a Poisson distribution can be recovered from the indirect observations. The hypothesis test turns out to be rather satisfactory provided the number of observations is sufficient.

5.2 Example 2

In this example, we consider a four-state Markov chain X with transition matrix P and we are interested in testing whether P is doubly-stochastic, and ultimately, symmetric. We assume X can only take 4 different values, the true transition matrix of X is chosen as

$$P = P_0 = \begin{bmatrix} 0.8 & 0 & 0.1 & 0.1 \\ 0.2 & 0.4 & 0.3 & 0.1 \\ 0 & 0.1 & 0.2 & 0.7 \\ 0 & 0.5 & 0.4 & 0.1 \end{bmatrix}. \quad (9)$$

Remark that P_0 is doubly-stochastic but not symmetric. Thus, testing if P is doubly-stochastic aims to measure the accuracy of the test under the null hypothesis while the test on the symmetry should evaluate the power. To avoid the technical considerations of the identifiability under the null hypothesis (assumption **A1**), we will assume that the diagonal entries of P are known. This condition ensures the problem to be identifiable since, while the bistochasticity certainly provides some information on P , it does not help making the problem identifiable. This is due to the fact that the bistochasticity property (and symmetry as a matter of fact) is preserved by the transformation $Q = G_\mu(P)$. In particular, this hypothesis could be tested directly on Q . Nevertheless, this example shows that the test can be carried out successfully on P in this context.

In this example, the initial model \mathcal{M} is of dimension 8 as it contains all stochastic matrices with fixed diagonal. The first hypothesis (bistochasticity) reduces the model dimension to 5 while the symmetry assumption only preserves 2 degrees of freedom. As previously, we perform the tests for four sample sizes: $n = 200$, $n = 500$, $n = 1000$ and $n = 2000$. Under the null hypothesis, we simulate the Markov chain $X = (X_m)_{m \in \mathbb{N}}$ with transition matrix P_0 and the times gaps τ_k are drawn from a Poisson distribution $\mu = \mathcal{P}(1)$. Only the sequence $Y = (X_{T_k})_{k=1, \dots, n}$, with $T_k = \sum_{j=1}^k \tau_j$, is observed and the distribution μ of the τ_k is unknown to the practitioner. The nominal significance level is fixed equal to $\alpha = 5\%$ for all numerical experiments.

Test 1.

We wish to test if P is bistochastic, knowing its diagonal entries. The initial model $\mathcal{M}^{(1)}$ and the model under $\mathcal{H}_0^{(1)}$ are given by

$$\begin{aligned} \mathcal{M}^{(1)} &: \{ \text{vec}(M) \in \mathbb{R}^{16} : M\mathbf{1} = \mathbf{1}, \text{diag}(M) = \text{diag}(P_0) \} \\ \mathcal{M}_0^{(1)} &: \{ \text{vec}(M) \in \mathbb{R}^{16} : M\mathbf{1} = M^\top \mathbf{1} = \mathbf{1}, \text{diag}(M) = \text{diag}(P_0) \} \end{aligned}$$

By Monte-Carlo method, the significance level is estimated for the four sample sizes using 10^4 replications. Results are gathered in Table 4.

n	200	500	1000	2000
$\hat{\alpha}$	0.050 (0.0022)	0.043 (0.0020)	0.045 (0.0020)	0.046 (0.0021)

Table 4: *Test for hypothesis $\mathcal{H}_0^{(1)}$: “ P is bistochastic” in the model $\mathcal{M}^{(1)} = \{\text{vec}(M) : M\mathbf{1} = \mathbf{1}, \text{diag}(M) = \text{diag}(P_0)\}$.* The table contains summary statistics of Monte Carlo simulation results obtained with $P = P_0$ given in Eq. (9), for the four different sample sizes n and for $\mu = \mathcal{P}(1)$. The probability to reject $\mathcal{H}_0^{(1)}$ is estimated for a nominal size $\alpha = 0.05$, based on 10^4 replications, with standard errors in brackets.

The estimated significance level converges to the nominal size 5% as n increases. The test is surprisingly accurate for a small sample size $n = 200$, although the extreme precision in this case is probably only coincidental. For larger sample sizes, the significance level seems to stabilize around its nominal value.

Test 2.

We perform a test on the symmetry of P , considering the null hypothesis $\mathcal{H}_0^{(2)} : P = P^\top$. In this example, we assume that P is known to be bistochastic. We thus define

$$\mathcal{M}^{(2)} = \mathcal{M}_0^{(1)} = \{\text{vec}(M) : M\mathbf{1} = M^\top\mathbf{1} = \mathbf{1}, \text{diag}(M) = \text{diag}(P_0)\}.$$

The constrained model, i.e. the model under $\mathcal{H}_0^{(2)}$ is given by

$$\mathcal{M}_0^{(2)} = \{\text{vec}(M) : M\mathbf{1} = \mathbf{1}, M = M^\top\}.$$

Because the true value $P = P_0$ is not symmetric, this simulation study evaluates the power of the test in this particular framework. The estimated values are given in Table 2.

n	200	500	1000	2000
$\hat{\beta}$	0.413 (0.0049)	0.916 (0.0028)	0.995 (0.0007)	1 (0)

Table 5: *Test for the hypothesis $\mathcal{H}_0^{(2)} : P = P^\top$ in the model $\mathcal{M}^{(2)} = \{\text{vec}(M) : M\mathbf{1} = M^\top\mathbf{1} = \mathbf{1}, \text{diag}(M) = \text{diag}(P_0)\}$.* The table contains summary statistics of Monte Carlo simulation results obtained with $P = P_0$ given in Eq. (9), for the four considered sample sizes and for $\mu \sim \mathcal{P}(1)$. The power of the test is estimated for a significance level $\alpha = 0.05$, based on 10^4 replications.

The result are quite satisfactory for middle sample sizes with the test procedure rightfully ruling out the symmetry of P in more than 91% of the cases from $n = 500$, and more than 99% of the time for $n \geq 1000$.

Test 3.

We are now interested in knowing if the distribution μ of the random sampling $(\tau_k)_k$ is Poisson distributed. We are thus interested in the null hypothesis

$$\mathcal{H}_0^{(3)} : \mu = \mu_0 = \mathcal{P}(1).$$

We assume here that P is known to be bistochastic, choosing the model

$$\mathcal{M}^{(3)} = \mathcal{M}_0^{(1)} = \{\text{vec}(M) \in \mathbb{R}^{16} : M\mathbf{1} = M^\top \mathbf{1} = \mathbf{1}, \text{diag}(M) = \text{diag}(P_0)\}.$$

The identifiability assumption, necessary for this procedure, is thus verified. The results are very similar to that of the previous example. For instance, the significance level gets close to its nominal size $\alpha = 0.05$ for relatively small samples, as illustrated in Table 3.

n	200	500	1000	2000
$\hat{\alpha}$	0.056 (0.0023)	0.040 (0.0019)	0.0042 (0.0020)	0.048 (0.0021)

Table 6: Test for the hypothesis $\mathcal{H}_0^{(3)} : \mu = \mathcal{P}(1)$ in the model $\mathcal{M}^{(3)} = \{\text{vec}(M) : M\mathbf{1} = M^\top \mathbf{1} = \mathbf{1}, \text{diag}(M) = \text{diag}(P_0)\}$. The table contains summary statistics of Monte Carlo simulation results obtained with $P = P_0$ for the four sample sizes under the null hypothesis. The nominal size of the test is $\alpha = 0.05$, based on 10^4 replications.

The significance level is satisfactory for $n = 200$ but only seems to stabilize for greater sample sizes. Here again, this simulation confirms that, while the time gaps are not observed, we are still able to extract some information on their distribution.

6 Conclusion

We develop an hypothesis testing methodology on the transition matrix P and the distribution μ of time gaps in a Markovian model with random sampling. The original contribution consists in the construction of a testing procedure for affine hypothesis on P in this framework as well as a test on the distribution μ of the observation times. This setting requires additional information on P to make the model identifiable. The information is expressed in the form of affine constraints on P , thus extending the framework considered in [3] in which P is assumed to be sparse.

We show that the different forms of information on P , formalized via affine constraints, can be sufficient to recover the true distribution μ of the time gaps, or at least to build a consistent testing procedure on specific values of μ . The simulation study confirms that a goodness-of-fit test on μ can be carried out successfully despite the fact that the time gaps τ_k are not observed.

7 Appendix

Proof of Theorem 3.1. We use the fact that, for \mathcal{A} an affine subset of an Euclidean space,

$$\inf_{a \in \mathcal{A}} \|a\|^2 = \|(I - \Pi_{\mathcal{A}_{\text{lin}}})a_0\|^2,$$

for all $a_0 \in \mathcal{A}$, with $\Pi_{\mathcal{A}_{\text{lin}}}$ the orthogonal projector onto $\mathcal{A}_{\text{lin}} = \{a - a' : a, a' \in \mathcal{A}\}$. Define the linear sets $\hat{E} = \text{Im}(\Delta(\hat{Q})\Phi)$ and $\hat{E}_0 = \text{Im}(\Delta(\hat{Q})\Phi_0)$, we get

$$\inf_{m \in \mathcal{M}_0} \|\Delta(\hat{Q})m\|^2 - \inf_{m \in \mathcal{M}} \|\Delta(\hat{Q})m\|^2 = \|(I - \Pi_{\hat{E}_0})\Delta(\hat{Q})m_0\|^2 - \|(I - \Pi_{\hat{E}})\Delta(\hat{Q})m_0\|^2$$

for any m_0 in \mathcal{M}_0 and in particular, for $m_0 = p$ under \mathcal{H}_0 . Since $\hat{E}_0 \subseteq \hat{E}$, the Pythagorean theorem yields, under \mathcal{H}_0 ,

$$\inf_{m \in \mathcal{M}_0} \|\Delta(\hat{Q})m\|^2 - \inf_{m \in \mathcal{M}} \|\Delta(\hat{Q})m\|^2 = \|(\Pi_{\hat{E}} - \Pi_{\hat{E}_0})\Delta(\hat{Q})p\|^2 = \|\Pi_{\hat{F}}\Delta(\hat{Q})p\|^2,$$

setting $\hat{F} = \hat{E} \cap \hat{E}_0^\perp$. Write $\Pi_{\hat{F}}\Delta(\hat{Q})p = \Pi_F\Delta(\hat{Q})p + (\Pi_{\hat{F}} - \Pi_F)\Delta(\hat{Q})p$, we have

$$\sqrt{n} \Pi_F\Delta(\hat{Q})p = -\sqrt{n} \Pi_F\Delta(P)\hat{q} \xrightarrow{d} \mathcal{N}(0, \Pi_F\Delta(P)\Sigma\Delta(P)^\top \Pi_F),$$

yielding $n\|\Pi_F\Delta(\hat{Q})p\|^2 \xrightarrow{d} \chi^2(W)$. So, to prove the result, it remains to show that

$$n(\|\Pi_F\Delta(\hat{Q})p\|^2 - \|\Pi_{\hat{F}}\Delta(\hat{Q})p\|^2) = o_P(1).$$

Since $\sqrt{n}\|\Delta(\hat{Q})p\| = O_P(1)$, it suffices to show that $\Pi_{\hat{F}}$ converges in probability to Π_F in view of

$$\begin{aligned} n(\|\Pi_F\Delta(\hat{Q})p\|^2 - \|\Pi_{\hat{F}}\Delta(\hat{Q})p\|^2) &\leq n\|(\Pi_F - \Pi_{\hat{F}})\Delta(\hat{Q})p\| (\|\Pi_F\Delta(\hat{Q})p\| + \|\Pi_{\hat{F}}\Delta(\hat{Q})p\|) \\ &\leq \|\Pi_F - \Pi_{\hat{F}}\| \times 2n\|\Delta(\hat{Q})p\|^2, \end{aligned}$$

where $\|\cdot\|$ denotes the operator norm. Recall that $\Pi_{\hat{F}} = \Pi_{\hat{E}} - \Pi_{\hat{E}_0}$ with $\hat{E} = \text{Im}(\Delta(\hat{Q})\Phi)$ and $\hat{E}_0 = \text{Im}(\Delta(\hat{Q})\Phi_0)$. By **A1**, $\Delta(Q)\Phi_0$ is of full rank and

$$\Pi_{\hat{E}_0} = \Delta(\hat{Q})\Phi_0 \left(\Phi_0^\top \Delta(\hat{Q})^\top \Delta(\hat{Q})\Phi_0 \right)^{-1} \Phi_0^\top \Delta(\hat{Q})^\top$$

clearly converges in probability to Π_{E_0} . Moreover, we know by **A2** that $\dim(\hat{E}) = \text{Im}(\Delta(\hat{Q})\Phi) \leq \dim(E)$, so that the convergence of $\Delta(\hat{Q})\Phi$ to $\Delta(Q)\Phi$ is sufficient for $\Pi_{\hat{E}}$ to converge to Π_E . Thus, $\Pi_{\hat{F}}$ converges in probability to Π_F as $n \rightarrow \infty$ which ends the proof.

Proof of Corollary 3.2. By continuity of the map $(\epsilon, W) \mapsto \epsilon^\top W \epsilon$, Slutsky's lemma ensures the convergence in distribution of $\chi_G^2(\hat{W})$ to $\chi_G^2(W)$ and therefore, the convergence of the estimated quantile $\hat{u}_{1-\alpha}$ to the true quantile $u_{1-\alpha}$ of the $\chi_G^2(W)$ distribution. It follows that

$\lim_{n \rightarrow \infty} \mathbb{P}(S > \hat{u}_{1-\alpha}) = \mathbb{P}(S > u_{1-\alpha}) = \alpha$. Now, if $\forall p' \in \mathcal{M}_0 : \Delta(Q)p' \neq 0$, we show easily that $\inf_{m \in \mathcal{M}_0} \|\Delta(\hat{Q})m\|^2$ does not converge to zero as n grows to infinity, while we still have

$$\inf_{m \in \mathcal{M}} \|\Delta(\hat{Q})m\|^2 \leq \|\Delta(\hat{Q})p\|^2 \xrightarrow{\mathbb{P}} 0.$$

Hence, $S = n(\inf_{m \in \mathcal{M}_0} \|\Delta(\hat{Q})m\|^2 - \inf_{m \in \mathcal{M}} \|\Delta(\hat{Q})m\|^2)$ diverges in probability in this case, and thus the result follows.

Proof of Proposition 4.1. To show that the matrix Γ exists, it suffices to show that the series is normally convergent. We use that $\|((P^{j-1})^\top \otimes P^{k-j})\| \leq \|P^{j-1}\| \|P^{k-j}\| \leq 1$ to get

$$\sum_{k \geq 1} \left(\sum_{j=1}^k \|((P^{j-1})^\top \otimes P^{k-j})\| \right) \mu(k) \leq \sum_{k \geq 1} k \mu(k) = \mathbb{E}_\mu(\tau),$$

which is finite by assumption. We now compute the differential of g_μ at P . For $H \in \mathbb{R}^{N \times N}$ such that $\|H\| \leq 1$, we have

$$\begin{aligned} \lim_{t \rightarrow 0} \frac{1}{t} (G_\mu(P + tH) - G_\mu(P)) &= \lim_{t \rightarrow 0} \frac{1}{t} \sum_{k \geq 1} [(P + tH)^k - P^k] \mu(k) \\ &= \lim_{t \rightarrow 0} \frac{1}{t} \left[\sum_{k \geq 1} t \left(\sum_{j=1}^k P^{j-1} H P^{k-j} \right) \mu(k) + o(t) \right] \\ &= \sum_{k \geq 1} \left(\sum_{j=1}^k P^{k-j} H P^{j-1} \right) \mu(k). \end{aligned}$$

Applying the vectorization yields, in view of $\text{vec}(ABC) = (C^\top \otimes A) \text{vec}(B)$,

$$\begin{aligned} \lim_{t \rightarrow 0} \frac{1}{t} (g_\mu(P + tH) - g_\mu(P)) &= \sum_{k \geq 1} \left(\sum_{j=1}^k \text{vec}(P^{k-j} H P^{j-1}) \right) \mu(k) \\ &= \left[\sum_{k \geq 1} \left(\sum_{j=1}^k (P^{j-1})^\top \otimes P^{k-j} \right) \mu(k) \right] \text{vec}(H) \\ &= \Gamma \text{vec}(H). \end{aligned}$$

It follows from Cramer's theorem that,

$$\sqrt{n}(g_\mu(\hat{P}) - q) = \sqrt{n} \Gamma (\hat{p} - p) + o_P(1).$$

Using (7), we deduce that $\sqrt{n}(g_\mu(\hat{P}) - q) = \sqrt{n} \Gamma B(\hat{q} - q) + o_P(1)$, which leads to

$$\sqrt{n} (\hat{q} - g_\mu(\hat{P})) = \sqrt{n} (I - \Gamma B)(\hat{q} - q) + o_P(1),$$

and the result follows.

References

- [1] ANDERSON, D. F., AND KURTZ, T. G. Continuous time Markov chain models for chemical reaction networks. In *Design and Analysis of Biomolecular Circuits*. Springer, 2011, pp. 3–42.
- [2] ANDERSON, T. W., AND GOODMAN, L. A. Statistical inference about Markov chains. *Ann. Math. Statist.* 28 (1957), 89–110.
- [3] BARSOTTI, F., DE CASTRO, Y., ESPINASSE, T., AND ROCHET, P. Estimating the transition matrix of a Markov chain observed at random times. *Statistics & Probability Letters* 94 (2014), 98–105.
- [4] BARTLETT, M. S. The frequency goodness of fit test for probability chains. *Proc. Cambridge Philos. Soc.* 47 (1951), 86–95.
- [5] BAUM, L. E., AND PETRIE, T. Statistical inference for probabilistic functions of finite state Markov chains. *Ann. Math. Statist.* 37 (1966), 1554–1563.
- [6] CAPPÉ, O., MOULINES, E., AND RYDÉN, T. *Inference in hidden Markov models*, 2009.
- [7] CRAIG, B. A., AND SENDI, P. P. Estimation of the transition matrix of a discrete-time Markov chain. *Health Economics*, 11 (2002), 33–42.
- [8] DAVIES, R. B. Algorithm as 155: The distribution of a linear combination of χ^2 random variables. *Applied Statistics* (1980), 323–333.
- [9] GAVER JR, D. P. Imbedded Markov chain analysis of a waiting-line process in continuous time. *The Annals of Mathematical Statistics* (1959), 698–720.
- [10] GKANTSIDIS, C., MIHAIL, M., AND SABERI, A. Random walks in peer-to-peer networks. In *INFOCOM 2004. Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies* (2004), vol. 1, IEEE.
- [11] GUTTORP, P., AND MININ, V. N. *Stochastic modeling of scientific data*. CRC Press, 1995.
- [12] ISRAEL, R. B., ROSENTHAL, J. S., AND WEI, J. Z. Finding generators for Markov chains via empirical transition matrices, with applications to credit ratings. *Math. Finance* 11, 2 (2001), 245–265.
- [13] RABINER, L. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE* 77, 2 (1989), 257–286.