

Predictive characterization of mixtures of Markov chains

SANDRA FORTINI* and SONIA PETRONE**

Department of Decision Sciences, Bocconi University, Via Roentgen 1, 20136 Milano, Italy.
*E-mail: *sandra.fortini@unibocconi.it; **sonia.petrone@unibocconi.it*

Predictive constructions are a powerful way of characterizing the probability laws of stochastic processes with certain forms of invariance, such as exchangeability or Markov exchangeability. When de Finetti-like representation theorems are available, the predictive characterization implicitly defines the prior distribution, starting from assumptions on the observables; moreover, it often helps in designing efficient computational strategies. In this paper we give necessary and sufficient conditions on the sequence of predictive distributions such that they characterize a Markov exchangeable probability law for a discrete valued process \mathbf{X} . Under recurrence, Markov exchangeable processes are mixtures of Markov chains. Our predictive conditions are in some sense minimal sufficient conditions for Markov exchangeability; we also provide predictive conditions for recurrence. We illustrate their application in relevant examples from the literature and in novel constructions.

Keywords: Bayesian inference; edge reinforced random walks; Markov exchangeability; predictive distributions; recurrence; reinforced processes

1. Introduction

Predictive characterization of the probability law of a stochastic process is a fundamental problem in probability and statistics. Informally, this means characterizing the probability law P of a process $(X_n, n \geq 1)$ through the sequence of predictive distributions $(P_n, n \geq 1)$, such that X_1 has distribution P_1 and $X_{n+1}|X_1, \dots, X_n$ has distribution P_n for $n \geq 1$. The sequence (P_n) characterizes a probability measure P for the stochastic process (X_n) under general assumptions, by the Ionescu–Tulcea theorem. The problem of interest is to determine the conditions under which it characterizes a law P with some given properties, and in particular, some specific invariance property. Necessary and sufficient conditions under which the sequence of predictive distributions (P_n) characterizes an exchangeable P are given in [15]. In the present paper, we give necessary and sufficient conditions for the sequence (P_n) to characterize a P which is partially exchangeable in the sense of Diaconis and Freedman [10] – or, using the terminology of Zaman [43] and Zabell [42], Markov exchangeable.

In Bayesian statistics, predictive characterizations have fundamental and practical relevance. Prediction is often the main goal of statistical analysis and, from a Bayesian perspective, the predictive approach seems natural. Even in the context of independent replicates of an experiment, probabilistic dependence is introduced through the assumption of exchangeability and prediction is naturally solved through the conditional distributions of future results given the observed facts. Indeed, according to de Finetti [7], a statistical model is just a link of the probabilistic chain that

leads from past to future events. Thus, at least in principle, models and priors on non-observable parameters can and should be induced by probability assertions on the observable X_n , such as exchangeability and predictive structures. The predictive characterization of prior distributions is a long studied problem in Bayesian statistics. Dirichlet conjugate priors for exchangeable categorical sequences have been characterized by Zabell [41] based on Johnson's sufficiency postulate. Diaconis and Ylvisaker [12] characterize conjugate priors for the natural exponential family through predictive conditions. Powerful predictive constructions have also been given in Bayesian nonparametrics. The predictive characterization of the Dirichlet process in terms of Pólya sequences [5] clarifies the relationship with random partitions in combinatorics and population biology [13]. Walker and Muliere [39] characterize neutral to the right processes through an extension of Johnson's sufficiency postulate. The general class of species sampling priors [30], which includes the Dirichlet process and the two parameter Poisson–Dirichlet process [29], is characterized in terms of the predictive distributions. Zabell [42] extends the characterization of Dirichlet conjugate priors through Johnson's sufficiency postulate to Markov exchangeable sequences. Reinforced processes [6,27,28] play an important role in predictive constructions of exchangeable and Markov exchangeable sequences; references include [24,38]. Diaconis and Rolles [11] provide a predictive characterization of conjugate priors for the transition matrix of a reversible Markov chain through edge reinforced random walks on a graph. Developments for variable order reversible Markov chains are found in [2] and [3]. Connections with the theory of vertex-reinforced jump processes are studied in [33]. Beyond the foundational issues, predictive characterizations are powerful tools in hierarchical modeling of symmetry structures and as generating algorithms which can be exploited for computational purposes, as recent developments at the interface between statistics and machine learning show. See, for example, the predictive construction of Markov exchangeable processes with countable unknown state space by Beal, Ghahramani and Rasmussen [4] (see also [35] and [16]), which has a wide application in hierarchical clustering and infinite hidden Markov models; or the predictive construction of the Indian buffet process [17] for latent features allocation, whose de Finetti-like representation has been later provided in [36]. Refer to [34] for an overview and further references.

We recall the main concepts and provide some preliminary results in Section 2. We first review two characterizations of mixtures of Markov chains, in terms of Markov exchangeability of the process and of partial exchangeability of the matrix of successor states. The main point of this section is to revisit main basic results in a predictive approach and relate the prior to the predictive distributions. For exchangeable sequences, the prior measure is the limit law of the sequence of predictive distributions; we formalize an analogous result for recurrent Markov exchangeable sequences in Section 2.3. This result further enhances the interest for predictive constructions and, therefore, for sufficient predictive conditions for Markov exchangeability.

Addressing the latter question is the main theoretical contribution of the paper, presented in Section 3. We give necessary and sufficient conditions on the sequence of predictive rules (P_n) under which they characterize a Markov exchangeable P for a discrete valued process $(X_n, n \geq 1)$. Furthermore, we give some predictive conditions for recurrence in Section 4. Under recurrence, a Markov exchangeable sequence is a mixture of recurrent Markov chains and the predictive structure characterizes the mixing distribution; that is, in a subjective, Bayesian approach, the prior distribution on the unknown transition matrix. Thus, our results are useful for verifying if a predictive scheme characterizes a prior for Bayesian inference on Markov chains.

In Section 5, we illustrate the results through a novel predictive construction, defined as an edge reinforced random walk on a colored graph. To some extent, the proposed scheme is a generalization of the edge reinforced random walks of [11], where colors allow the reinforcement of edges even when they are not crossed. The motivating idea is to introduce forms of probabilistic dependence or constraints on the random transition matrix through the reinforcement of both edges and colors. The predictive conditions given in the previous sections are used to establish when the proposed predictive scheme characterizes a Markov exchangeable process. Several proposals in the literature can be recovered as special cases of colored edge reinforced random walks, which, in this sense, offer a unifying framework. An advantage of the predictive approach is that it encourages prior distributions that are closed under sampling, as we illustrate in Section 5.1. We provide examples of colored edge reinforced random walks for which the prior distribution can incorporate information and constraints on the transition matrix, and can be easily updated given the data. Finally, we discuss extensions through the introduction of latent variables which may complicate the analytic computations, but can be easily simulated.

Some complements to these results and detailed proofs are provided in the [Appendix](#).

2. Overview and preliminary results

2.1. Markov exchangeability

Let S be a finite or countable set, containing at least two elements, and $\mathbf{X} = (X_n, n \geq 0)$ be a discrete-time stochastic process taking values in S and starting at a specific state $x_0 \in S$. The process \mathbf{X} is a mixture of Markov chains if there exists a probability law μ on the set \mathcal{P} of stochastic matrices on $S \times S$ endowed with the topology of element-wise convergence and the corresponding Borel sigma-algebra, such that, for every $\mathbf{x} = (x_1, \dots, x_n)$,

$$p(\mathbf{x}) = \int_{\mathcal{P}} \prod_{k=1}^n \tilde{p}_{x_{k-1}, x_k} \mu(d\tilde{p}), \quad (1)$$

where $p(\mathbf{x}) = P(\mathbf{X}_{1:n} = \mathbf{x})$ with $\mathbf{X}_{1:n} = (X_1, \dots, X_n)$. Equivalently, \mathbf{X} is a mixture of Markov chains if there exists a \mathcal{P} -valued random element \tilde{P} such that, conditionally on \tilde{P} , \mathbf{X} is a Markov chain with transition matrix \tilde{P} . In Bayesian inference for Markov chains, the probability law $\mu(\cdot)$ plays the role of the prior distribution on the unknown transition matrix.

In a predictive approach, interest is in understanding what probabilistic assumptions on the observable process \mathbf{X} imply (1), that is, the existence of a prior distribution on the non-observable matrix \tilde{P} . A main result by Diaconis and Freedman [10] shows that, for recurrent processes, such assumption is Markov exchangeability. Recurrence here means that $P(X_n = x_0 \text{ infinitely often}) = 1$. Markov exchangeability is defined as invariance under a certain kind of symmetry. Two finite strings \mathbf{z} and \mathbf{z}' are equivalent, written $\mathbf{z} \sim \mathbf{z}'$, if \mathbf{z} and \mathbf{z}' have the same first element and exhibit the same number of transitions from i to j , for every pair of states i and j . The process \mathbf{X} is Markov exchangeable if $(x_0, \mathbf{x}) \sim (x_0, \mathbf{x}')$ implies $p(\mathbf{x}) = p(\mathbf{x}')$. It can be proved ([10], Lemma 5) that, if $\mathbf{z}' \sim \mathbf{z}$, then \mathbf{z} and \mathbf{z}' have the same length, end at the same state and visit each state the same number of times. Hence, Markov exchangeability can be seen

as invariance with respect to a class of permutations, namely those permutations that do not alter the number of transitions between any two states. This motivates the original term of *partial exchangeability* used by Diaconis and Freedman.

Markov exchangeability is a necessary condition for the process to be representable as a mixture of Markov chains; under recurrence, it is also sufficient. The proof of this result is based on the exchangeability of the sequence of x_0 -blocks. A x_0 -block is defined as a finite sequence of states that begins at x_0 and contains no further instances of x_0 . Recurrence of the process \mathbf{X} ensures that the sequence of x_0 -blocks is infinite. Under recurrence, Markov exchangeability implies exchangeability of the sequence of x_0 -blocks. Because a Markov exchangeable sequence with independent and identically distributed x_0 -blocks is a Markov chain, it follows that \mathbf{X} is a mixture of Markov chains ([10], Proposition 15 and Theorem 7). Thus, a recurrent Markov exchangeable process is a mixture of recurrent Markov chains.

An important point is that, if \mathbf{X} is a mixture of recurrent Markov chains for some specified μ , one can determine μ from \mathbf{X} . Let $\mathbf{T}_{i,j}(x_0, \mathbf{x})$ denote the number of transitions from i to j in the finite sequence (x_0, \mathbf{x}) , $\mathbf{T}(x_0, \mathbf{x}) = [\mathbf{T}_{i,j}(x_0, \mathbf{x})]_{i,j \in S}$ be the matrix of transition counts, $\mathbf{T}_i(x_0, \mathbf{x})$ be the i th row of $\mathbf{T}(x_0, \mathbf{x})$ and $\mathbf{T}_{i,\cdot}(x_0, \mathbf{x}) = \sum_j \mathbf{T}_{i,j}(x_0, \mathbf{x})$. Define $\hat{\mathbf{T}}_{i,j}(x_0, \mathbf{x}) = \mathbf{T}_{i,j}(x_0, \mathbf{x})/\mathbf{T}_{i,\cdot}(x_0, \mathbf{x})$ if $\mathbf{T}_{i,\cdot}(x_0, \mathbf{x}) > 0$ and $\hat{\mathbf{T}}_{i,j}(x_0, \mathbf{x}) = \mathbf{T}_{i,j}(x_0, \mathbf{x}) = 0$ if $\mathbf{T}_{i,\cdot}(x_0, \mathbf{x}) = 0$. Then $\hat{\mathbf{T}}_{i,j}(x_0, \mathbf{X}_{1:n})$ converges to $\tilde{P}_{i,j}$ almost surely (a.s.) as $n \rightarrow \infty$ ([10], Section 4). The limit matrix \tilde{P} may not be a stochastic matrix on $S \times S$, as the sum of the elements in a row corresponding to a state that is not visited is zero. In fact, \tilde{P} is a stochastic matrix on the random set of the visited states, say $A_{\tilde{P}}$. In order for the limit matrix \tilde{P} to be a stochastic matrix (almost surely), we can use a conventional enlargement of the state space, as in [14]. Let us introduce an additional state ∂ and denote by S^* the enlarged space $S \cup \{\partial\}$. Then set $\hat{\mathbf{T}}_{i,\partial} = 1 - \sum_{j \in S} \hat{\mathbf{T}}_{i,j}$ and $\hat{\mathbf{T}}_{\partial,\partial} = 1$. The enlarged matrix $[\hat{\mathbf{T}}_{i,j}]_{i,j \in S^*}$ converges pointwise almost surely to a stochastic matrix \tilde{P} on $S^* \times S^*$. The mixing measure μ in (1) is uniquely determined as the probability law of \tilde{P} . We have understood that the probability measure P on S^∞ has been extended to $(S^*)^\infty$, and \mathcal{P} represents the class of stochastic matrices on S^* . For the sake of simplicity, we keep the same notation, and keep denoting by \mathbf{X} the coordinate process on $(S^*)^\infty$; the distinction is clear from the context.

2.2. Partial exchangeability of successor states

A different characterization of mixtures of recurrent Markov chains, hinted in [8] and [42], is developed by Fortini *et al.* [14], in terms of partial exchangeability of the matrix of successor states. The n th successor of a state $i \in S$ is the state that follows the n th visit to i . More formally, for every $i \in S^*$, let $\tau_n(i)$ be the time of the n th visit to i , with the proviso that $\tau_n(i) = \infty$ if state i is not visited n times. The n th successor state of i is defined as $V_{i,n} = X_{\tau_n(i)+1}$ if $\tau_n(i) < \infty$ and $V_{i,n} = \partial$ otherwise. The successor matrix associated to \mathbf{X} is then defined as the array $\mathbf{V} = [V_{i,n}, i \in S^*, n \geq 1]$. If the process \mathbf{X} is recurrent and Markov exchangeable, then it is also strongly recurrent [14]; that is, if a state i is visited, it is visited infinitely often, almost surely. Thus, the rows of the matrix \mathbf{V} are infinite sequences, the i th row being a sequence in S^∞ if i is visited, or a sequence equal to $(\partial, \partial, \dots)$ if i is not visited, or if $i = \partial$. Fortini *et al.* [14]

show that the process \mathbf{X} is recurrent and Markov exchangeable if and only if the successor matrix \mathbf{V} is partially exchangeable by rows, in the sense of de Finetti, that is, there exists a stochastic matrix \tilde{P} on S^* such that, conditionally on \tilde{P} , the random variables $(V_{i,n}, i \in S^*, n \geq 1)$ are independent with $P(V_{i,n} = j | \tilde{P}) = \tilde{P}_{i,j}$. By the properties of partially exchangeable sequences, the random matrix \tilde{P} is determined as

$$\tilde{P}_{i,j} = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \delta_{V_{i,k}}(\{j\}) \quad \text{a.s.}, \tag{2}$$

where δ is the degenerate measure defined by $\delta_a(\{a\}) = 1$. This gives an interpretation of the prior μ in (1) as the limit law of the empirical distributions of the successor states. Another characterization can be given in terms of predictive distributions, as we discuss in the next subsection.

2.3. Predictive properties

The directing measure of an infinite exchangeable sequence can be characterized as the limit of the sequence of the predictive distributions $P(X_{n+1} \in \cdot | X_1, \dots, X_n)$, which converges weakly almost surely to a random distribution F such that the $X_i | F$ are a random sample from F ; see [1], page 60. Similar properties hold for recurrent Markov exchangeable processes. In this case, from the characterization in terms of partial exchangeability of the successor matrix \mathbf{V} , discussed in the previous section, it follows that, almost surely

$$\begin{aligned} & \lim_{n \rightarrow \infty} P(V_{i,n} = j | V_{i,1}, \dots, V_{i,n-1}) \\ &= \lim_{n \rightarrow \infty} P(V_{i,n} = j | V_{i,1}, \dots, V_{i,n-1}, V_{l,k}, k = 1, 2, \dots, l \in S, l \neq i) = \tilde{P}_{i,j}, \quad j \in S. \end{aligned}$$

The predictive rules above refer to the successors states. In terms of the sequence \mathbf{X} , the following result based on stopping times holds. Recall that $\tau_n(i)$ is the time of the n th visit to state i , and let $\mathcal{F}_{\tau_n(i)}$ be the sigma-algebra of the events until the n th visit of \mathbf{X} to state i . Let $X_{\tau_n(i)+1} = X_{k+1}$ if $\tau_n(i) = k$, and $X_{\tau_n(i)+1} = \partial$ if $\tau_n(i) = \infty$.

Theorem 1. *Let \mathbf{X} be a mixture of recurrent Markov chains with random transition matrix \tilde{P} . Then, for every $i \in S^*$,*

$$\lim_{n \rightarrow \infty} P(X_{\tau_n(i)+1} = j | \mathcal{F}_{\tau_n(i)}) = \tilde{P}_{i,j} \quad \text{a.s.}$$

Proof. The result is immediate for $i = \partial$. Consider $i \in S$. Denote by $\mathcal{F}_{\tilde{P}}$ the sigma-algebra generated by \tilde{P} , and by $\mathcal{F}_{\tau_n(i)} \vee \mathcal{F}_{\tilde{P}}$ the sigma algebra generated by $\mathcal{F}_{\tau_n(i)} \cup \mathcal{F}_{\tilde{P}}$. Then

$$P(X_{\tau_n(i)+1} = j | \mathcal{F}_{\tau_n(i)}) = P(V_{i,n} = j | \mathcal{F}_{\tau_n(i)}) = E(P(V_{i,n} = j | \mathcal{F}_{\tau_n(i)} \vee \mathcal{F}_{\tilde{P}}) | \mathcal{F}_{\tau_n(i)}).$$

Partial exchangeability of the successors matrix \mathbf{V} implies that $V_{i,n}$ is conditionally independent of the other successor states $(V_{i,k}, k < n; V_{j,l}, j \in S^*, j \neq i, l \geq 1)$, given \tilde{P} . As formally proved

in Lemma 4 in the Appendix, $\mathcal{F}_{\tau_n(i)}$ is included in the sigma algebra generated by $(V_{i,k}, k < n; V_{j,l}, j \in S^*, j \neq i, l \geq 1)$; therefore, $V_{i,n}$ is also conditionally independent of $\mathcal{F}_{\tau_n(i)}$, given \tilde{P} . Hence

$$P(X_{\tau_n(i)+1} = j | \mathcal{F}_{\tau_n(i)}) = E(\tilde{P}_{i,j} | \mathcal{F}_{\tau_n(i)}),$$

which converges to $E(\tilde{P}_{i,j} | \mathcal{F}_{\tau_\infty(i)})$, where $\mathcal{F}_{\tau_\infty(i)} = \bigvee_n \mathcal{F}_{\tau_n(i)}$ is the sigma-algebra generated by $\bigcup_{n=1}^\infty \mathcal{F}_{\tau_n(i)}$. Since $V_{i,k}$ is measurable with respect to $\mathcal{F}_{\tau_\infty(i)}$ for all $k \geq 1$, then $\tilde{P}_{i,j} = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \delta_{V_{i,k}}(\{j\})$ is $\mathcal{F}_{\tau_\infty(i)}$ -measurable, too. Thus, $E(\tilde{P}_{i,j} | \mathcal{F}_{\tau_\infty(i)}) = \tilde{P}_{i,j}$ almost surely and the proof is complete. \square

A mixture of recurrent Markov chains has a random transition matrix with independent rows if and only if the probability of observing a transition from i to j depends only on the past transition from state i .

Corollary 1. *Let \mathbf{X} be a mixture of recurrent Markov chains with random transition matrix \tilde{P} . The rows of \tilde{P} are stochastically independent if and only if, for all $n \geq 1$, X_{n+1} is conditionally independent of $\mathbf{X}_{1:n}$ given $(X_n, \mathbf{T}_{X_n}(x_0, \mathbf{X}_{1:n}))$.*

Although this fact appears to be known in the literature (hints are in [42]; see also, for example, [4,24]), for completeness we provide a proof in Appendix (Corollary 3), based on Theorem 1. Informally, the result follows from partial exchangeability of the matrix \mathbf{V} of successors states. For any i and j , the probability of a transition to j , given past observations ending at i , is equal to the probability that $V_{i,n_i+1} = j$ conditionally on the sequence of successors, where n_i is the number of past successors of state i . If such probability only depends on i and on the transitions from i , the probability distribution of V_{i,n_i+1} , given the other successor states, only depends on the successors of i ; thus, the rows of \mathbf{V} are independent. By (2), it follows that \tilde{P} has independent rows. Conversely, if \tilde{P} has independent rows, then partial exchangeability by rows of \mathbf{V} reduces to independence and internal exchangeability of the rows, therefore the probability of a transition from i to j only depends on transitions from i .

3. Predictive characterization of Markov exchangeability

The basic question studied in this work is: when does a sequence of predictive rules characterize a Markov exchangeable process? Before stating the results, let us introduce some simplifying notation. We continue to denote finite sequences of elements in S by bold letters, (e.g., $\mathbf{x}, \mathbf{y}, \dots$), while non-bold letters (e.g., x, y, \dots) denote single elements of S . Unless otherwise specified, a string can coincide with the empty string, denoted by \emptyset . The predictive probabilities are denoted by $p(\mathbf{y}|x_0, \mathbf{x})$. Hence, $p(\mathbf{y}|x_0, \mathbf{x}) = p(\mathbf{x}, \mathbf{y})/p(\mathbf{x})$ if $p(\mathbf{x}) \neq 0$ and $p(\mathbf{y}|x_0, \mathbf{x})$ is defined arbitrarily if $p(\mathbf{x}) = 0$. A string (x_0, \mathbf{x}, i) should be interpreted as a string of any length starting at x_0 and ending at i ; this includes the string of length one: (x_0) , when $i = x_0$.

Markov exchangeability clearly implies the following predictive properties:

- (A) $p(\mathbf{y}|x_0, \mathbf{x}, i) = p(\mathbf{y}|x_0, \mathbf{x}', i)$ for every \mathbf{y} and every $i, \mathbf{x}, \mathbf{x}'$ such that $(x_0, \mathbf{x}, i) \sim (x_0, \mathbf{x}', i)$;
- (B) $p(\mathbf{y}|x_0, \mathbf{x}, i) = p(\mathbf{y}'|x_0, \mathbf{x}, i)$ for every i, \mathbf{x} and every \mathbf{y}, \mathbf{y}' such that $(i, \mathbf{y}) \sim (i, \mathbf{y}')$.

Condition (B) is sufficient for Markov exchangeability; in fact, condition (B) for $(x_0, \mathbf{x}, i) = x_0$ is the definition of Markov exchangeability. Thus, verifying (B) is not easier than verifying Markov exchangeability directly from the joint distribution. Our aim is to show that a predictive condition weaker than (B) and a simpler form of condition (A) are jointly sufficient for Markov exchangeability.

Condition (A) is the joint predictive sufficiency of the last state and the transition counts: for every $n \geq 1$, $(X_{n+1}, X_{n+2}, \dots)$ are conditionally independent of $\mathbf{X}_{1:n}$, given $(X_n, \mathbf{T}(x_0, \mathbf{X}_{1:n}))$. This is because $\mathbf{T}(x_0, \mathbf{x}_{1:n}) = \mathbf{T}(x_0, \mathbf{x}'_{1:n})$ and $x_n = x'_n$ imply $(x_0, x_1, \dots, x_n, x_{n+1}, \dots, x_{n+k}) \sim (x_0, x'_1, \dots, x'_n, x_{n+1}, \dots, x_{n+k})$. This property is the analogous, for Markov exchangeable processes, of predictive sufficiency of the empirical distribution for exchangeable sequences. Clearly, a predictive sufficient statistic is also one-step-ahead predictive sufficient, that is, for any n , X_{n+1} is conditionally independent of $\mathbf{X}_{1:n}$, given the statistic. Predictive sufficiency and one-step-ahead predictive sufficiency are different conditions, in general. But in the specific case where the statistic has the form $(X_n, \mathbf{T}(x_0, \mathbf{X}_{1:n}))$, they turn out to be equivalent.

Proposition 1. *Let \mathbf{X} be an S -valued process starting at x_0 . Then, for any n , $(X_n, \mathbf{T}(x_0, \mathbf{X}_{1:n}))$ is predictive sufficient if and only if it is one-step-ahead predictive sufficient.*

Proof. In order to show the non-obvious implication, it is sufficient to write the conditional distribution of X_{n+1}, \dots, X_{n+k} , given X_1, \dots, X_n , as a product of conditional distributions and notice that $(x_0, \mathbf{x}_{1:n}) \sim (x_0, \mathbf{x}'_{1:n})$ implies $x_n = x'_n$ and $(x_0, x_1, \dots, x_n, x_{n+1}, \dots, x_{n+j}) \sim (x_0, x'_1, \dots, x'_n, x_{n+1}, \dots, x_{n+j})$ for every $j = 1, \dots, k$. □

Thus, we can replace condition (A) with one-step-ahead predictive sufficiency of $(X_n, \mathbf{T}(x_0, \mathbf{X}_{1:n}))$. However, such condition alone does not imply Markov exchangeability. As a counterexample, consider $S = \{0, 1, 2, \dots\}$, $x_0 = 0$ and for $n \geq 1$, $p(j|x_0, x_1, \dots, x_{n-1}) = 1/n$ for $j = 1, \dots, n$ and zero otherwise. Since $n = 1 + \sum_{i,j} \mathbf{T}_{i,j}(x_0, x_1, \dots, x_{n-1})$, the vector $(X_n, \mathbf{T}(x_0, \mathbf{X}_{1:n}))$ is one-step-ahead predictive sufficient. However, \mathbf{X} is not Markov exchangeable. For example, $p(1, 3, 1, 2, 3) = 0$ while $p(1, 2, 3, 1, 3) \neq 0$.

On the other hand, one-step-ahead predictive sufficiency of the transition counts, together with a weaker form of condition (B), implies Markov exchangeability. This is shown in the following theorem. We denote by $\{\mathbf{x}\}$ the set of distinct elements in \mathbf{x} .

Theorem 2. *Let $\mathbf{X} = (X_n, n \geq 0)$ be an S -valued process such that $X_0 = x_0$. Then \mathbf{X} is Markov exchangeable if and only if both the following conditions hold:*

- (a) $p(y|x_0, \mathbf{x}, i) = p(y|x_0, \mathbf{x}', i)$ for every $y \in S$ and every $i, \mathbf{x}, \mathbf{x}'$ such that $(x_0, \mathbf{x}, i) \sim (x_0, \mathbf{x}', i)$;
- (b) $p(\mathbf{y}|x_0, \mathbf{x}, i) = p(\mathbf{y}'|x_0, \mathbf{x}, i)$ for every $i, \mathbf{x}, \mathbf{y}$ and \mathbf{y}' such that $\mathbf{y} = (\mathbf{u}, \mathbf{w}, i, \mathbf{v}, \mathbf{w}, i)$ and $\mathbf{y}' = (\mathbf{v}, \mathbf{w}, i, \mathbf{u}, \mathbf{w}, i)$ with $\{i\}, \{\mathbf{u}\}, \{\mathbf{v}\}, \{\mathbf{w}\}$ disjoint.

The proof makes use of some lemmas, given below. Before proceeding, some remarks are in order.

Remark 1. Under condition (a), condition (b) is equivalent to simultaneously having:

- (bi) $p(\mathbf{u}, i, i|x_0, \mathbf{x}, i) = p(i, \mathbf{u}, i|x_0, \mathbf{x}, i)$ for every \mathbf{u}, i with $\{i\}, \{\mathbf{u}\}$ disjoint;
- (bii) $p(\mathbf{u}, i, \mathbf{v}, i|x_0, \mathbf{x}, i) = p(\mathbf{v}, i, \mathbf{u}, i|x_0, \mathbf{x}, i)$ for every $\mathbf{u}, \mathbf{v}, i$ with $\{i\}, \{\mathbf{u}\}, \{\mathbf{v}\}$ disjoint;
- (biii) $p(\mathbf{u}, j, \mathbf{w}, i, \mathbf{v}, j|x_0, \mathbf{x}, i) = p(\mathbf{v}, j, \mathbf{w}, i, \mathbf{u}, j|x_0, \mathbf{x}, i)$ for every $\mathbf{u}, \mathbf{v}, \mathbf{w}, i, j$ with $\{i\}, \{\mathbf{u}\}, \{\mathbf{v}\}, \{j, \mathbf{w}\}$ disjoint.

Conditions (bi) and (bii) are obtained from (b) by setting $\mathbf{w} = \mathbf{v} = \emptyset$ and $\mathbf{w} = \emptyset$, respectively; (biii) is recovered by substituting \mathbf{w} with (j, \mathbf{w}) and canceling $p(\mathbf{w}|x_0, \mathbf{x}, i, \mathbf{u}, j, \mathbf{w}, i, \mathbf{v}, j)$ from the left-hand side and $p(\mathbf{w}|x_0, \mathbf{x}, i, \mathbf{v}, j, \mathbf{w}, i, \mathbf{u}, j)$ from the right-hand side. Similarly for showing that (bi)–(biii) imply (b).

Remark 2. For a recurrent process \mathbf{X} , when conditions (a) and (b) of Theorem 2 hold for any $\mathbf{u}, \mathbf{v}, \mathbf{w}$ not including i , they imply exchangeability of the i -blocks, by Theorem 3.1 in [15]. For a general process, if conditions (bi)–(biii) hold for any \mathbf{u}, \mathbf{v} and \mathbf{w} not including i , they are equivalent to invariance under block-switch transformations, which in turn is equivalent to Markov exchangeability ([10], Section 4). Theorem 2 says that, under condition (a), it is enough to check block-switch invariance on a subset of disjoint strings.

Before proving Theorem 2, we introduce some convenient notation. The length of a string \mathbf{z} is denoted by $|\mathbf{z}|$. Given two finite strings \mathbf{z} and \mathbf{z}' , we say that \mathbf{z} is shorter than \mathbf{z}' and write $\mathbf{z} \leq \mathbf{z}'$, if there exists \mathbf{x} such that $(\mathbf{x}, \mathbf{z}) \sim \mathbf{z}'$; if $\mathbf{x} \neq \emptyset$, we say that \mathbf{z} is strictly shorter than \mathbf{z}' and write $\mathbf{z} < \mathbf{z}'$. Given a class \mathcal{C} of nonempty S -valued finite strings, an element \mathbf{z}^* is called minimal in \mathcal{C} if $\mathbf{z}^* \in \mathcal{C}$ and there is no $\mathbf{z} \in \mathcal{C}$ that is strictly shorter than \mathbf{z}^* .

Lemma 1. Let $\mathbf{x}_0 = (\mathbf{u}_0, j, j, \mathbf{v}_0, k, j)$ with $k \neq j$ and let $\mathcal{C} = \{\mathbf{x} \leq \mathbf{x}_0 : \mathbf{x} = (\mathbf{u}, j, j, \mathbf{v}, k, j)\}$. Then there exists a minimal element \mathbf{x}^* in \mathcal{C} , given by $\mathbf{x}^* = (j, j, \mathbf{v}^*, k, j)$ with $j \notin \{\mathbf{v}^*\}$.

Proof. The lemma is proved by contradiction. If there was no minimal element, we could find an infinite sequence $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots$ of strings in \mathcal{C} such that $\mathbf{x}_0 > \mathbf{x}_1 > \mathbf{x}_2 > \dots$. In that case, $|\mathbf{x}_0| > |\mathbf{x}_1| > |\mathbf{x}_2| > \dots$, which is impossible, since $|\mathbf{x}_0|$ is finite. Thus, a minimal element exists. Let such minimal element be $\mathbf{x}^* = (\mathbf{u}^*, j, j, \mathbf{v}^*, k, j)$. Then, \mathbf{u}^* has to be empty, as otherwise a shorter element of \mathcal{C} could be obtained by deleting \mathbf{u}^* . Furthermore, should \mathbf{v}^* contain j , we would have $\mathbf{x}^* = (j, j, \mathbf{v}_1^*, j, \mathbf{v}_2^*, k, j) \sim (j, \mathbf{v}_1^*, j, j, \mathbf{v}_2^*, k, j)$ and a shorter string could be obtained by deleting (j, \mathbf{v}_1^*) . \square

Lemma 2. Let j, k, k' be distinct elements of S and let $\mathbf{x} = (\mathbf{u}, k, j, \mathbf{v}, k', j)$ and $\mathbf{x}' = (\mathbf{u}', k', j, \mathbf{v}', k, j)$ be such that $\mathbf{x} \sim \mathbf{x}'$. Then, either $\{\mathbf{u}, k\} \cap \{j, \mathbf{v}, k'\} \neq \emptyset$ or $\{\mathbf{u}', k'\} \cap \{j, \mathbf{v}', k\} \neq \emptyset$ or both.

Proof. By contradiction: suppose that $j \notin \{\mathbf{u}\}, j \notin \{\mathbf{u}'\}, \{\mathbf{u}, k\} \cap \{\mathbf{v}, k'\} = \emptyset$ and $\{\mathbf{u}', k'\} \cap \{\mathbf{v}', k\} = \emptyset$. Then $k \notin \{\mathbf{v}\}$ and $k \notin \{\mathbf{u}'\}$. Let $\mathbf{u} = (u_1, \dots, u_m)$. We show by backward induction on s that $u_s \in \{\mathbf{v}', k\}$ for every $s = 1, \dots, m$. u_m is a predecessor of k in \mathbf{x} and therefore in \mathbf{x}' . Since $k \notin \{\mathbf{u}'\}$, then $u_m \in \{\mathbf{v}'\}$. Now, suppose that $u_s \in \{\mathbf{v}'\}$. Since u_{s-1} is a predecessor of u_s in \mathbf{x} , it is also a predecessor of u_s in \mathbf{x}' . Since $u_s \notin \{\mathbf{u}'\}$, then $u_{s-1} \in \{\mathbf{v}'\}$. It follows by induction that $u_s \in \{\mathbf{v}'\}$ for every s . Since $u_1 \in \{\mathbf{v}'\}$ and $\{\mathbf{u}'\} \cap \{\mathbf{v}'\} = \emptyset$, then $u_1 \notin \{\mathbf{u}'\}$. Hence, the first element of \mathbf{x}' is not u_1 , which contradicts $\mathbf{x} \sim \mathbf{x}'$. \square

Lemma 3. *Let j, k, k' be distinct elements of S and let $\mathbf{x}_0 = (\mathbf{u}_0, k, j, \mathbf{v}_0, k', j)$ with $\{\mathbf{u}_0, k\} \cap \{j, \mathbf{v}_0, k'\} \neq \emptyset$. Let $\mathcal{C} = \{\mathbf{x} \preceq \mathbf{x}_0 : \mathbf{x} = (\mathbf{u}, k, j, \mathbf{v}, k', j), \{\mathbf{u}, k\} \cap \{j, \mathbf{v}, k'\} \neq \emptyset\}$. Then, there exists a minimal element $\mathbf{x}^* = (\mathbf{u}^*, k, j, \mathbf{v}^*, k', j)$ in \mathcal{C} which is either of the form $\mathbf{x}^* = (j, \mathbf{u}, k, j, \mathbf{v}, k', j)$, with $\{j\}, \{k\}, \{k'\}, \{\mathbf{u}\}, \{\mathbf{v}\}$ disjoint, or $\mathbf{x}^* = (i, \mathbf{u}, k, j, \mathbf{w}, i, \mathbf{v}, k', j)$, with $\{i\}, \{k\}, \{k'\}, \{\mathbf{u}\}, \{\mathbf{v}\}, \{j, \mathbf{w}\}$ disjoint.*

The existence of a minimal element can be proved as for Lemma 1. The rest of the proof is by contradiction: if \mathbf{x}^* does not have the above structure, one can make a switch transformation by moving a piece of \mathbf{x}^* in front, and, by deleting such a piece, obtain a string in \mathcal{C} which is shorter than \mathbf{x}^* . A detailed proof is provided in the [Appendix](#).

Proof of Theorem 2. Markov exchangeability implies conditions (a) and (b). In order to prove that, under (a) and (b), $(x_0, \mathbf{z}) \sim (x_0, \mathbf{z}')$ implies $p(\mathbf{z}) = p(\mathbf{z}')$, we proceed by induction on $n = |\mathbf{z}|$. It is convenient to use the conditions (bi)–(biii). The thesis is true for $n = 1$. Suppose it is true for all sequences of length $k \leq n$ and let $\mathbf{z} = (z_1, \dots, z_{n+1})$, $\mathbf{z}' = (z'_1, \dots, z'_{n+1})$ be two strings such that $(x_0, \mathbf{z}) \sim (x_0, \mathbf{z}')$. Since equivalent strings end at the same state, $z_{n+1} = z'_{n+1}$. We treat separately the cases $z_n = z'_n$ and $z_n \neq z'_n$.

If $z_n = z'_n$, then $(x_0, z_1, \dots, z_n) \sim (x_0, z'_1, \dots, z'_n)$. By the induction hypothesis, $p(z_1, \dots, z_n) = p(z'_1, \dots, z'_n)$. Hence,

$$p(\mathbf{z}) = p(z_{n+1}|x_0, z_1, \dots, z_n)p(z_1, \dots, z_n) = p(z'_{n+1}|x_0, z'_1, \dots, z'_n)p(z'_1, \dots, z'_n) = p(\mathbf{z}'),$$

where we have used $z_{n+1} = z'_{n+1}$ and condition (a).

Suppose, now, that $z_n \neq z'_n$. Let $z_n = k$, $z'_n = k'$ and $z_{n+1} = z'_{n+1} = j$. It is enough to distinguish two cases: (I): $j = k'$ and (II): $j \neq k, k'$.

(I) Since (x_0, \mathbf{z}') contains the transition (j, j) and $(x_0, \mathbf{z}) \sim (x_0, \mathbf{z}')$, (x_0, \mathbf{z}) also contains (j, j) . Hence we can write $(x_0, \mathbf{z}) = (x_0, \mathbf{u}_0, j, j, \mathbf{v}_0, k, j)$ with $k \neq j$. The last expression should be intended as a string starting at x_0 , containing the transition (j, j) in some position and the transition (k, j) in the last position. Such strings include $(x_0, x_0, \mathbf{v}_0, k, x_0)$, for $j = x_0$. By Lemma 1, there exist \mathbf{x} and \mathbf{v}^* such that $(x_0, \mathbf{z}) \sim (x_0, \mathbf{x}, j, j, \mathbf{v}^*, k, j)$ with $j \notin \{\mathbf{v}^*\}$. Let $\mathbf{u} = (\mathbf{v}^*, k)$. Since $z_n = k$, again as in the first part of the proof, $p(\mathbf{z}) = p(\mathbf{x}, j, j, \mathbf{v}^*, k, j)$. Furthermore, $p(\mathbf{z}) = p(\mathbf{x}, j)p(j, \mathbf{u}, j|x_0, \mathbf{x}, j)$ and, by condition (bi), with $i = j$, $p(\mathbf{x}, j)p(j, \mathbf{u}, j|x_0, \mathbf{x}, j) = p(\mathbf{x}, j)p(\mathbf{u}, j, j|x_0, \mathbf{x}, j)$. Hence, $p(\mathbf{z}) = p(\mathbf{x}, j, \mathbf{u}, j, j)$. Since $z'_n = j$, reasoning as in the first part of the proof, $p(\mathbf{x}, j, \mathbf{u}, j, j) = p(\mathbf{z}')$. Thus, $p(\mathbf{z}) = p(\mathbf{z}')$.

(II) We can write $(x_0, \mathbf{z}) = (\mathbf{u}_0, k, j, \mathbf{v}_0, k', j)$ and $(x_0, \mathbf{z}') = (\mathbf{u}'_0, k', j, \mathbf{v}'_0, k, j)$. By Lemma 2, without loss of generality, we can suppose that $\{\mathbf{u}_0, k\} \cap \{j, \mathbf{v}_0, k'\} \neq \emptyset$. Let us consider the class $\mathcal{C} = \{(\mathbf{u}, k, j, \mathbf{v}, k', j) \preceq (x_0, \mathbf{z}) : \{\mathbf{u}, k\} \cap \{j, \mathbf{v}, k'\} \neq \emptyset\}$. By Lemma 3, a minimal element in \mathcal{C} exists and can be written either as $\mathbf{x}^* = (j, \mathbf{u}, k, j, \mathbf{v}, k', j)$, with $\{j\}, \{k\}, \{k'\}, \{\mathbf{u}\}, \{\mathbf{v}\}$ disjoint, or as $\mathbf{x}^* = (i, \mathbf{u}, k, j, \mathbf{w}, i, \mathbf{v}, k', j)$, with $\{i\}, \{k\}, \{k'\}, \{\mathbf{u}\}, \{\mathbf{v}\}, \{j, \mathbf{w}\}$ disjoint. In the first case, there exists \mathbf{x} such that $(x_0, \mathbf{z}) \sim (x_0, \mathbf{x}, j, \mathbf{u}, k, j, \mathbf{v}, k', j)$. Since $z_n = k'$, again as in the first part of the proof, $p(\mathbf{z}) = p(\mathbf{x}, j, \mathbf{u}, k, j, \mathbf{v}, k', j)$. Furthermore $p(\mathbf{x}, j, \mathbf{u}, k, j, \mathbf{v}, k', j) = p(\mathbf{x}, j)p(\mathbf{u}, k, j, \mathbf{v}, k', j|x_0, \mathbf{x}, j)$, which, by (bii) with $i = j$, is equal to

$$p(\mathbf{x}, j)p(\mathbf{v}, k', j, \mathbf{u}, k, j|x_0, \mathbf{x}, j) = p(\mathbf{x}, j, \mathbf{v}, k', j, \mathbf{u}, k, j).$$

Since $z'_n = k$, then $p(\mathbf{x}, j, \mathbf{v}, k', j, \mathbf{u}, k, j) = p(\mathbf{z}')$. If the minimal element is $\mathbf{x}^* = (i, \mathbf{u}, k, j, \mathbf{w}, i, \mathbf{v}, k', j)$ the proof can be obtained along the same steps by using property (biii). \square

Example 1 (Edge reinforced random walks). Edge reinforced random walks (ERRW) are predictive schemes that characterize conjugate priors for reversible Markov chains [6,11,31]. Consider a finite undirected graph G with vertex set V and edge set E (possibly including loops). All edges in E are given a strictly positive weight; at time zero, edge e has weight $\alpha_e > 0$. An edge-reinforced random walk on G with starting point $x_0 \in V$ is defined as follows. The process starts at x_0 at time 0. At each step, the random walker traverses an edge with probability proportional to its weight. Each time an edge in E , that is not a loop, is traversed, its weight is increased by 1. Each time a loop in E is traversed, its weight is increased by 2. Thus, the predictive probability of traversing edge $e = (i, j)$ is

$$p(j|x_0, \mathbf{x}, i) = \frac{\alpha_{(i,j)} + \mathbf{T}_{i,j}(x_0, \mathbf{x}, i) + \mathbf{T}_{j,i}(x_0, \mathbf{x}, i)}{\alpha_{(i,\cdot)} + \mathbf{T}_{i,\cdot}(x_0, \mathbf{x}, i) + \mathbf{T}_{\cdot,i}(x_0, \mathbf{x}, i)}, \quad (3)$$

where $\alpha_{(i,\cdot)} = \sum_k \alpha_{(i,k)}$ is the sum of the weights of the edges incident to i , and $\mathbf{T}_{i,\cdot} = \sum_{j'} \mathbf{T}_{i,j'}$ and $\mathbf{T}_{\cdot,i} = \sum_{j'} \mathbf{T}_{j',i}$ are the transitions from and to state i , respectively. ERRWs are known to generate Markov exchangeable processes. We notice that Markov exchangeability can be easily verified through Theorem 2. Condition (a) is immediate, as (3) depends on (x_0, \mathbf{x}, i) only through the transitions and the last state i . Condition (b) is satisfied if $p(\mathbf{y}|x_0, \mathbf{x}, i) = p(\mathbf{y}'|x_0, \mathbf{x}, i)$ for vectors \mathbf{y}, \mathbf{y}' of the form $\mathbf{y} = (\mathbf{u}, \mathbf{w}, i, \mathbf{v}, \mathbf{w}, i)$ and $\mathbf{y}' = (\mathbf{v}, \mathbf{w}, i, \mathbf{u}, \mathbf{w}, i)$ with $\{i\}, \{\mathbf{u}\}, \{\mathbf{v}\}, \{\mathbf{w}\}$ disjoint. Direct computation of the above conditional probabilities involves the product of terms of the form (3) recursively updated. The transition counts satisfy

$$\mathbf{T}_{u_{l-1}, u_l}(x_0, \mathbf{x}, i, \mathbf{v}, \mathbf{w}, i, u_1, \dots, u_{l-1}) = \mathbf{T}_{u_{l-1}, u_l}(x_0, \mathbf{x}, i, u_1, \dots, u_{l-1})$$

and similarly for $\mathbf{T}_{u_l, u_{l-1}}$ and for all the transition counts in the numerator of (3). Analogous equations hold for the denominators, with the exception of the terms involving i and \mathbf{w} , which satisfy

$$\mathbf{T}_{i,\cdot}(x_0, \mathbf{x}, i, \mathbf{u}, \mathbf{w}, i) = \mathbf{T}_{i,\cdot}(x_0, \mathbf{x}, i, \mathbf{v}, \mathbf{w}, i) = \mathbf{T}_{i,\cdot}(x_0, \mathbf{x}, i) + 1,$$

and similarly for $\mathbf{T}_{\cdot,i}(x_0, \mathbf{x}, i)$ and \mathbf{w} . The above equations imply condition (b).

If $(X_n, \mathbf{T}_{X_n}(x_0, \mathbf{X}_{1:n}))$ is one-step-ahead predictive sufficient, that is, $p(j|x_0, \mathbf{x}, i)$ is a function $\pi(j|\mathbf{T}_i, i)$ of the last element i and of the transition counts $\mathbf{T}_i = \mathbf{T}_i(x_0, \mathbf{X}_{1:n})$ from i , then the conditions for Markov exchangeability simplify greatly.

Corollary 2. *Let \mathbf{X} be an S -valued stochastic process starting at x_0 and such that (X_n, \mathbf{T}_{X_n}) is one-step-ahead predictive sufficient. Then the process \mathbf{X} is Markov exchangeable if and only if, for every i, u, v ,*

$$\pi(u|\mathbf{T}_i, i)\pi(v|\mathbf{T}_i + \mathbf{e}_u, i) = \pi(v|\mathbf{T}_i, i)\pi(u|\mathbf{T}_i + \mathbf{e}_v, i), \quad (4)$$

where $\mathbf{T}_i + \mathbf{e}_k$ is the vector \mathbf{T}_i with the k th element incremented by one.

Proof. The proof is simple, applying Theorem 2. Condition (a) holds by assumption. We need to show that, under the assumptions of the corollary, (4) is equivalent to condition (b) of Theorem 2, that is, for any i , $\mathbf{u} = (u_1, \mathbf{u}')$, $\mathbf{v} = (v_1, \mathbf{v}')$, $\mathbf{w} = (w_1, \mathbf{w}')$ disjoint,

$$p(u_1, \mathbf{u}', w_1, \mathbf{w}', i, v_1, \mathbf{v}', w_1, \mathbf{w}' | x_0, \mathbf{x}, i) = p(v_1, \mathbf{v}', w_1, \mathbf{w}', i, u_1, \mathbf{u}', w_1, \mathbf{w}' | x_0, \mathbf{x}, i). \quad (5)$$

Writing both sides as products of conditional probabilities, and noticing that

$$\begin{aligned} p(\mathbf{u}' | x_0, \mathbf{x}, i, u_1) &= p(\mathbf{u}' | x_0, \mathbf{x}, i, \mathbf{v}, \mathbf{w}, i, u_1); \\ p(w_1 | x_0, \mathbf{x}, i, \mathbf{u}) &= p(w_1 | x_0, \mathbf{x}, i, \mathbf{v}, \mathbf{w}, i, \mathbf{u}); \\ p(\mathbf{w}' | x_0, \mathbf{x}, i, \mathbf{u}, w_1) &= p(\mathbf{w}' | x_0, \mathbf{x}, i, \mathbf{v}, w_1); \\ p(i | x_0, \mathbf{x}, i, \mathbf{u}, \mathbf{w}) &= p(i | x_0, \mathbf{x}, i, \mathbf{v}, \mathbf{w}); \\ p(\mathbf{v}' | x_0, \mathbf{x}, i, \mathbf{u}, \mathbf{w}, i, v_1) &= p(\mathbf{v}' | x_0, \mathbf{x}, i, v_1); \\ p(w_1 | x_0, \mathbf{x}, i, \mathbf{u}, \mathbf{w}, i, \mathbf{v}) &= p(w_1 | x_0, \mathbf{x}, i, \mathbf{v}); \\ p(\mathbf{w}' | x_0, \mathbf{x}, i, \mathbf{u}, \mathbf{w}, i, \mathbf{v}, w_1) &= p(\mathbf{w}' | x_0, \mathbf{x}, i, \mathbf{v}, \mathbf{w}, i, \mathbf{u}, w_1), \end{aligned}$$

by the predictive sufficiency assumptions, one can easily verify that (5) holds if and only if $p(u_1 | x_0, \mathbf{x}, i) p(v_1 | x_0, \mathbf{x}, i, \mathbf{u}, \mathbf{w}, i) = p(v_1 | x_0, \mathbf{x}, i) p(u_1 | x_0, \mathbf{x}, i, \mathbf{v}, \mathbf{w}, i)$, which, again by the predictive sufficiency assumptions, corresponds to (4). \square

An alternative proof, which however requires the additional assumption that \mathbf{X} is recurrent, can be given in terms of the successor states. If \mathbf{X} is recurrent, the matrix \mathbf{V} has rows of infinite length in S and, under the assumptions of Corollary 2, such rows are independent. Then, condition (4) is equivalent to exchangeability of the sequence of successor states of i , by Theorem 3.1 in [15], for any i ; therefore, it is equivalent to Markov exchangeability, as shown in Section 2.

Example 2 (Reinforced urn schemes). Let S be finite or countable. Consider the following predictive probabilities

$$p(j | x_0, \mathbf{x}, i) = \pi(j | \mathbf{T}_i, i) = \frac{\alpha_i q_i(j) + \mathbf{T}_{i,j}}{\alpha_i + \mathbf{T}_i}, \quad (6)$$

where $\sum_{j \in S} q_i(j) = 1$ and $\alpha_i > 0$ for any i . The predictive probability of observing state j is a weighted average of an initial weight $q_i(j)$ and the relative transition counts from the last element in the sample, $\mathbf{T}_{i,j} / \mathbf{T}_i$. This is a simple example of a predictive structure as considered in Corollary 2. It is immediate to verify that (4) holds; thus, by Corollary 2, the sequence of predictive distributions (6) characterizes a Markov exchangeable probability law for the process \mathbf{X} .

For a finite state space S , the predictive rule (6) has been derived by Zabell [42] from Johnson's sufficiency postulate: he shows that, if \mathbf{X} is recurrent and Markov exchangeable, and the predictive probability $p(j | x_0, \mathbf{x}, i)$ is a function of $i, j, \mathbf{T}_{i,j}(x_0, \mathbf{x}, i), \mathbf{T}_i \cdot (x_0, \mathbf{x}, i)$ for each j , then such function has to be of the linear form (6). Here, the linear structure (6) of the predictive probabilities is an assumption, while Markov exchangeability is deduced by Corollary 2. Muliere, Secchi

and Walker [24] construct the predictive rule (6) through a reinforced urn process. An extension of their construction, for a countable state space S , can be obtained through a reinforced Hoppe’s urn scheme. A Hoppe urn [18] is associated to each state, with urn i having initial number of black balls α_i and color distribution $q_i(\cdot)$ on S . The process starts at x_0 . At step n , a ball is drawn from urn x_{n-1} and, if colored, it is returned in the urn together with an additional ball of the same color; if black, a color is picked from the color distribution $q_{x_{n-1}}(\cdot)$ and a ball of the sampled color is added in the urn, together with the black ball. The process then moves to the urn associated with the color of the additional ball, and so on. X_n represents the color of the additional ball at the n th step. The predictive law of the process \mathbf{X} defined in such way is (6). If the process \mathbf{X} is recurrent, then it is a mixture of recurrent Markov chains, whose random transition matrix has independent rows. From the results discussed in Section 2, the i th row is the random limit of the sequence of predictive distributions of the successors of state i ; here, the draws from the urn associated to state i . These form an exchangeable sequence with predictive rule (6) which characterizes a Dirichlet prior distribution with parameters $(\alpha_i q_i(j), j \in S)$ when S is finite, or a Dirichlet process with parameter $\alpha q_i(\cdot)$, denoted $DP(\alpha q_i(\cdot))$, if S is countable. Extensions of the predictive rule (6), introducing dependence across the rows of the transition matrix, are given in [4,16,35], using hierarchical Hoppe’s urns.

A mixture of Markov chains with random transition matrix having independent Dirichlet rows has predictive rule (6). Thus, the Markov exchangeable process \mathbf{X} characterized by (6) can be represented as a mixture of Markov chains; but without recurrence such representation may not be unique. We provide predictive conditions for recurrence in the next section.

4. Predictive conditions for recurrence

A Markov exchangeable process \mathbf{X} is a mixture of processes of the following kinds ([10], page 124): (1) recurrent Markov chains, (2) processes starting with a string of transient states and continuing as recurrent Markov chains, (3) totally transient processes. It is of interest to have conditions under which \mathbf{X} is a mixture of Markov chains, thus excluding mixing of processes of kind 2 and of processes of kind 3 which are not Markov chains. Recurrence is a sufficient condition for restricting to mixtures of processes of kind 1. A recurrent process \mathbf{X} is Markov exchangeable if and only if it is a mixture of recurrent Markov chains. In this sense, recurrence is a simplifying assumption.

Proving recurrence of a general Markov exchangeable process is a difficult task, and results are mostly available for specific constructions. In a predictive approach, interest is in conditions on the predictive distributions that imply recurrence. We provide below a sufficient predictive condition for recurrence, for a general process. For a Markov exchangeable process, this condition is also necessary.

Theorem 3. *If, for any (x_1, x_2, \dots) in a set of probability one,*

$$\sum_{n=0}^{\infty} p(x_0|x_0, x_1, \dots, x_n) = \infty, \tag{7}$$

then \mathbf{X} is recurrent. Conversely, if \mathbf{X} is Markov exchangeable and recurrent, then (7) holds, almost surely.

Proof. We first prove that (7) implies recurrence. The proof is based on Lévy’s extension of the Borel Cantelli lemma. Let $A_n = \{X_n = x_0\}$ and \mathcal{F}_n the sigma-algebra generated by $\mathbf{X}_{1:n}$. Then $A_n \in \mathcal{F}_n$ and (7) can be rewritten as $\sum_{k=1}^\infty P(A_k|\mathcal{F}_{k-1}) = \infty$, almost surely. This implies (see, e.g., [40], Section 12.15)

$$\frac{\sum_{k=1}^n \mathbb{1}_{A_k}}{\sum_{k=1}^n P(A_k|\mathcal{F}_{k-1})} \rightarrow 1 \quad \text{a.s.}$$

as $n \rightarrow \infty$, where $\mathbb{1}_A$ is the indicator of event A . In turn, this entails $P(A_n \text{ infinitely often}) = 1$.

Now we prove the second assertion. If \mathbf{X} is recurrent and Markov exchangeable, it is a mixture of recurrent Markov chains. The process \mathbf{X} is also strongly recurrent, that is, $P(X_n = i \text{ infinitely often} | i \text{ is visited}) = 1$, for any state i . Furthermore, the set of elementary events (x_1, x_2, \dots) such that $\tilde{p}_{i,x_0} = 0$ for every visited i , has probability zero as, otherwise, x_0 would not be visited infinitely many times with probability one. It follows that, for any (x_1, x_2, \dots) in a set of probability one, there exists a state i , generally depending on (x_1, x_2, \dots) , that is visited infinitely many times, has positive transition probability \tilde{p}_{i,x_0} and satisfies $p(x_0|x_0, \dots, x_{\tau_n(i)}) \rightarrow \tilde{p}_{i,x_0} > 0$, where the last assertion follows from Theorem 1. Therefore,

$$\sum_{n=0}^\infty p(x_0|x_0, \dots, x_n) \geq \sum_{n=1}^\infty p(x_0|x_0, \dots, x_{\tau_n(i)}) = \infty. \quad \square$$

Theorem 3 offers a strategy to prove recurrence; in some cases, one can find a lower bound for $p(x_0|x_0, \dots, x_n)$ and easily prove that the series diverges.

Example 2 – Ctd. Consider again the reinforced Hoppe’s urns of Example 2 and suppose that $\alpha_i = \alpha$ and $q_i(\cdot) = q(\cdot)$ for every i (all urns have the same initial number of black balls and the same color distribution), with $q(x_0) > 0$. In this case, recurrence is an immediate consequence of Theorem 3, since $\sum_{n=0}^\infty p(x_0|x_0, x_1, \dots, x_n) \geq \sum_{n=0}^\infty \alpha q(x_0)/(\alpha + n) = \infty$. For this particular case, recurrence was verified in [16] through a different approach. The general case (different α_i and q_i) is still immediate if $\inf_i \alpha_i q_i(x_0) > 0$. A simple example is obtained when the state space is finite and $q_i(x_0) > 0$ for any i .

In the example above, if $q_i(x_0) = 0$ for some state i , recurrence is more difficult to prove. In fact, if there exists some $i \in S$ that is visited infinitely many times almost surely and such that $q_i(x_0) > 0$, we could still apply Theorem 3. However the existence of such i depends precisely on the recurrence property that we are exploring. The following theorem can be useful in this situation. Recall that $\mathcal{F}_{\tau_n(i)}$ is the sigma-field generated by all events until the n th visit of \mathbf{X} to state i .

Theorem 4. Given $i, j \in S$, if

$$\sum_{n=1}^{\infty} [P(X_{\tau_n(i)+1} = j | \mathcal{F}_{\tau_n(i)}) \mathbb{1}_{\{\tau_n(i) < \infty\}} + \mathbb{1}_{\{\tau_n(i) = \infty\}}] = \infty \tag{8}$$

almost surely, then $P((X_n = i \text{ for finitely many } n) \cup (X_n = j \text{ infinitely often})) = 1$.

Proof. For every $n \geq 1$, let $A_{n+1} = (\tau_n(i) = +\infty) \cup (X_{\tau_n(i)+1} = j)$ and $\mathcal{G}_n = \mathcal{F}_{\tau_n(i)}$. Then, $A_n \in \mathcal{G}_n$, and (8) can be written as $\sum_{n=1}^{\infty} P(A_{n+1} | \mathcal{G}_n) = \infty$, almost surely. By Lévy’s extension of the Borel–Cantelli lemma,

$$\frac{\sum_{k=2}^n \mathbb{1}_{A_k}}{\sum_{k=2}^n P(A_k | \mathcal{G}_{k-1})} \rightarrow 1$$

almost surely, as $n \rightarrow \infty$. This implies that $\sum_{k=2}^{\infty} \mathbb{1}_{A_k} = \infty$ a.s. The thesis follows. □

Theorem 4 implies that, if there exists a state i that is visited infinitely often, it is sufficient to check that (8) holds with $j = x_0$ to have recurrence.

Example 2 – Ctd. Consider the urn scheme of Example 2 with general weights $\alpha_i q_i(\cdot)$. Let Q denote the stochastic matrix with (i, j) th entry $Q_{i,j} = q_i(j)$. Consider two states i and j such that $q_i(j) > 0$. Then, almost surely,

$$\begin{aligned} & \sum_{n=1}^{\infty} [P(X_{\tau_n(i)+1} = j | \mathcal{F}_{\tau_n(i)}) \mathbb{1}_{\{\tau_n(i) < \infty\}} + \mathbb{1}_{\{\tau_n(i) = \infty\}}] \\ & \geq \sum_{n=1}^{\infty} \left[\frac{\alpha_i q_i(j)}{\alpha_i + n} \mathbb{1}_{\{\tau_n(i) < \infty\}} + \mathbb{1}_{\{\tau_n(i) = \infty\}} \right] = \infty. \end{aligned}$$

Hence, by Theorem 4, if the chain visits i infinitely many times, it also visits j infinitely many times. In the case when $q_i(j) = 0$, if state j can be reached from i in a finite number of steps, that is, there exist u_1, \dots, u_k such that $Q_{i,u_1} \cdots Q_{u_k,j} > 0$, we can use the same reasoning to show that condition (8) holds for each pair of those states. Thus, if j is accessible from i and i is recurrent, then j is recurrent as well. A recurrent state i certainly exists if the state space is finite. Thus, in the finite case, we can give a simple sufficient condition for recurrence: if the matrix Q is the transition matrix of an irreducible Markov chain, then \mathbf{X} is recurrent.

We should notice that these techniques are mostly useful for processes with a finite state space, where one can assume that at least one state is recurrent. Although they can also be helpful in fairly simple constructions with countable state space, such as in Example 2, proving recurrence in the infinite case is generally far more difficult than in the finite case and requires more sophisticated techniques.

Recurrence has been studied in depth for edge reinforced random walks. A phase transition in the recurrence and transience of ERRWs on an infinite binary tree has been shown by Pemantle

[26]. For ERRWs on *finite* graphs, Keane and Rolles ([19], Proposition 1) proved recurrence using techniques similar to those used here for proving Theorems 1 and 2. Exploiting recurrence, and developing a conjecture in [9], they also show that an ERRW can be represented as a random walk on a graph with random edge weights (a random walk on a random environment), whose distribution is the limit law of the fractions of time spent by the process on the edges. Recalling the representation of reversible Markov chains as random walks on undirected graphs (see, e.g., [11]), this result implies that ERRWs are mixtures of *reversible* Markov chains, and leads to the mixing distribution.

To extend the results to infinite graphs, more refined techniques are needed. Merkl and Rolles [21] proved that, for any *locally finite* graph, an ERRW is a mixture of reversible Markov chains, irrespectively whether it is recurrent or not. To show this, they first obtain more detailed properties of the random environment for a finite graph, then treat the infinite case by approximating the infinite (but locally finite) graph with a sequence of finite graphs. In the finite case, they study the predictive distribution of successor states of a state i , as considered here in Theorem 1. In this theorem, we show that such sequence of predictive distributions converges to the random transition $\hat{P}_{i,j}$, while [21] give more refined results on the limit law. By comparing the predictive distribution of successor states with a Polya urn, whose asymptotic behaviour is well understood, they obtain bounds on the tails of the distribution of the random transition probabilities $\hat{P}_{i,j}$. Such bounds are uniform in the size of the graph. Thus, they can be exploited in the approximation of the infinite graph by finite graphs, to show that the ERRW is a mixture of reversible Markov chains, by also taking into account that, for a reversible Markov chain, there is a relationship between the transition probabilities and the edge weights in its graph representation. Although the work by Merkl and Rolles [21] does not deal with recurrence explicitly, their results give a valuable contribution to the study of recurrence for the *infinite* case. A reversible irreducible Markov chain is recurrent if and only if, in its graph representation, the sum of the edge weights is finite. Thus, once it is proved that an ERRW can be represented as a random walk on a random environment, controlling the tails of the distribution of the random edge weights is a technique to show that they are strictly positive and summable, almost surely and, therefore, the process is recurrent. Proving such bounds is technically hard. Extending results in [20] and [32], Merkl and Rolles [22] study the asymptotic behavior of ERRWs on general multi-level ladders; in particular, they show that the edge weights decay exponentially in space, and prove recurrence. Developing on the results in [21], Merkl and Rolles [23] prove recurrence of ERRW on a large class of periodic graphs satisfying certain symmetry properties. A different technique based on a representation of ERRWs in terms of vertex-reinforced jump process is used by Sabot and Tarres [33], who show that ERRWs on \mathbb{Z}^d are strongly recurrent for any d , for large reinforcement, under mild conditions.

Although reversible Markov chains have specific features, some of these techniques, possibly the approximation by finite graphs to obtain detailed results on the limit law of the predictive distribution of successor states, could be helpful for other Markov exchangeable processes.

5. Colored edge reinforced random walks

In this section, we provide an illustration of the previous results through a novel predictive construction. To some extent, the predictive scheme proposed here is a generalization of ERRWs,

with colored edges. Introducing colors allows to reinforce groups of edges, even if they are not crossed. Colors' reinforcement could be exploited to express restrictions or global properties of the process. We use Theorem 2 to study the conditions under which this predictive scheme characterizes a Markov exchangeable process, and provide examples which satisfy such conditions. These include processes found in the literature as well as novel characterizations.

Informally, the process is described as a random walk on a colored graph where edges and colors are reinforced when crossed. Consider a directed graph with vertices in a finite or countable set S and with edge set E . Each directed edge (i, j) in E is given a color $c(i, j)$ from a set of colors C and a weight $\beta_{i,j} > 0$. Each color c is assigned an initial weight $\alpha_c > 0$. For each vertex i , let $C(i)$ denote the set of colors of all the edges starting at i and let $\alpha_{C(i)} = \sum_{c \in C(i)} \alpha_c$ be the overall weight of all colors from i . For each color c , let E_c denote the set of all edges of color c in the graph, and $\beta_{i,E_c} = \sum_{j:(i,j) \in E_c} \beta_{i,j}$ be the overall weight of all edges of color c starting from i . We assume $\alpha_{C(i)} < \infty$ and $\beta_{i,E_c} < \infty$ for every i and every c .

Let x_0 be the starting point of the random walk. At step one, a color c is selected among the colors in $C(x_0)$, with probability $\alpha_c/\alpha_{C(x_0)}$. Then, a directed edge (x_0, x_1) is selected among the edges of color c from x_0 , with probability $\beta_{x_0,x_1}/\beta_{x_0,E_c}$. The process moves to $X_1 = x_1$ and both the weight of color c and of edge (x_0, x_1) are incremented by one. The walk is repeated in the same way, this time starting from x_1 and with the new weights; and so on for the following steps.

Let \mathbf{X} denote the process generated in this manner, and let $\mathbf{T}(x_0, \mathbf{X}_{1:m})$ be its transitions count process, with $\mathbf{T}_{i,j}$ denoting the number of transitions across edge (i, j) in $(x_0, \mathbf{X}_{1:m})$. The resulting predictive probabilities for the process \mathbf{X} are as follows. For $(i, y) \in E_c$,

$$p(y|x_0, \mathbf{x}, i) = \frac{\alpha_c + \mathbf{T}_{E_c}(x_0, \mathbf{x}, i)}{\alpha_{C(i)} + \mathbf{T}_{C(i)}(x_0, \mathbf{x}, i)} \frac{\beta_{i,y} + \mathbf{T}_{i,y}(x_0, \mathbf{x}, i)}{\beta_{i,E_c} + \mathbf{T}_{i,E_c}(x_0, \mathbf{x}, i)}, \tag{9}$$

where, for each color c , $\mathbf{T}_{E_c} = \sum_{(u,v) \in E_c} \mathbf{T}_{u,v}$ is the number of transitions over edges of color c , and for each vertex i , $\mathbf{T}_{C(i)} = \sum_{c \in C(i)} \mathbf{T}_{E_c}$ and $\mathbf{T}_{i,E_c} = \sum_{u:(i,u) \in E_c} \mathbf{T}_{i,u}$. Notice that, if there is only one edge of color c from i , the second factor in (9) is equal to one, thus the predictive probabilities reduce to the color updating. Analogously, if all edges from i have the same color, the first factor in (9) is one and the predictive probabilities reduce to the edge's weight updating. Also, one could allow reinforcements different from one, or no reinforcement at all, and this could be done both for colors and for edges.

Proposition 2. *The sequence of predictive rules (9) defines a Markov exchangeable process if and only if*

$$\prod_{l=2}^m [\alpha_{C(y_{l-1})} + \mathbf{T}_{C(y_{l-1})}(x_0, \mathbf{x}, i, y_1, \dots, y_{l-1})] = \prod_{l=2}^m [\alpha_{C(y'_{l-1})} + \mathbf{T}_{C(y'_{l-1})}(x_0, \mathbf{x}, i, y'_1, \dots, y'_{l-1})] \tag{10}$$

for every i, \mathbf{x} and $\mathbf{y} = (y_1, \dots, y_m)$, $\mathbf{y}' = (y'_1, \dots, y'_m)$ such that $\mathbf{y} = (\mathbf{u}, \mathbf{w}, i, \mathbf{v}, \mathbf{w}, i)$ and $\mathbf{y}' = (\mathbf{v}, \mathbf{w}, i, \mathbf{u}, \mathbf{w}, i)$, with $\{\mathbf{u}\}, \{\mathbf{v}\}, \{\mathbf{w}\}, \{i\}$ disjoint.

Proof. The sequence of predictive rules (9) defines a Markov exchangeable process if and only if it satisfies conditions (a) and (b) of Theorem 2. Condition (a) is immediate, as (9) depends on (x_0, \mathbf{x}, i) only through the transitions and the last state i . Direct computation of $p(\mathbf{y}|x_0, \mathbf{x}, i)$ involves the product of terms of the form (9) recursively updated, and one can easily check that condition (b) is satisfied if, for vectors \mathbf{y}, \mathbf{y}' of the form $\mathbf{y} = (\mathbf{u}, \mathbf{w}, i, \mathbf{v}, \mathbf{w}, i)$ and $\mathbf{y}' = (\mathbf{v}, \mathbf{w}, i, \mathbf{u}, \mathbf{w}, i)$ with $\{i\}, \{\mathbf{u}\}, \{\mathbf{v}\}, \{\mathbf{w}\}$ disjoint, the following equivalences hold:

$$\frac{\prod_{l=1}^m [\alpha_{c(y_{l-1}, y_l)} + \mathbf{T}_{E_{c(y_{l-1}, y_l)}}(x_0, \dots, y_{l-1})]}{\prod_{l=1}^m [\alpha_{C(y_{l-1})} + \mathbf{T}_{C(y_{l-1})}(x_0, \dots, y_{l-1})]} \tag{11}$$

$$= \frac{\prod_{l=1}^m [\alpha_{c(y'_{l-1}, y'_l)} + \mathbf{T}_{E_{c(y'_{l-1}, y'_l)}}(x_0, \dots, y'_{l-1})]}{\prod_{l=1}^m [\alpha_{C(y'_{l-1})} + \mathbf{T}_{C(y'_{l-1})}(x_0, \dots, y'_{l-1})]};$$

$$\frac{\prod_{l=1}^m [\beta_{y_{l-1}, y_l} + \mathbf{T}_{y_{l-1}, y_l}(x_0, \dots, y_{l-1})]}{\prod_{l=1}^m [\beta_{y_{l-1}, E_{c(y_{l-1}, y_l)}} + \mathbf{T}_{y_{l-1}, E_{c(y_{l-1}, y_l)}}(x_0, \dots, y_{l-1})]} \tag{12}$$

$$= \frac{\prod_{l=1}^m [\beta_{y'_{l-1}, y'_l} + \mathbf{T}_{y'_{l-1}, y'_l}(x_0, \dots, y'_{l-1})]}{\prod_{l=1}^m [\beta_{y'_{l-1}, E_{c(y'_{l-1}, y'_l)}} + \mathbf{T}_{y'_{l-1}, E_{c(y'_{l-1}, y'_l)}}(x_0, \dots, y'_{l-1})]},$$

where $y_0 = y'_0 = i$. The factors on both sides of (12) only depend on the last state and the transitions from it. Similarly to the proof of Corollary 2, one can see that (12) always holds. As for condition (11), the numerators on the left hand side only depend on the number of times colors $c(y_{l-1}, y_l)$ are visited in $(x_0, \mathbf{x}, i, \mathbf{y})$, which remain unchanged in $(x_0, \mathbf{x}, i, \mathbf{y}')$. Thus, condition (b) of Theorem 2 is satisfied if the denominators on both sides of (11) are equal, that is equivalent to (10). Since (12) is always true, condition (10) is also necessary. \square

Condition (10) has to be checked case by case, depending on the structure of the graph and of the reinforcement. Reinforced Hoppe’s urns and ERRWs are special cases of the predictive scheme (9), for which (10) holds.

Reinforced Hoppe’s urn schemes discussed in Example 2 are a particular case of (9) for which all edges in the graph have the same color. Then, the predictive rule (9) reduces to (6); for which condition (10) is immediate. Some extensions will be given in Section 5.1.

ERRWs are defined for undirected graphs, but they can be framed in our scheme by assigning to every pair of directed edges (i, j) and (j, i) the same color, with a different color for each different pair, and augmenting the graph by associating to each vertex i an auxiliary vertex, say i^* , to represent loops (i, i) as the pair of directed edges $(i, i^*), (i^*, i)$. Then, when edge (i, j) is crossed, both (i, j) and (j, i) are reinforced by one, and loops are reinforced by 2. In this case, the predictive rule (9) reduces to the predictive rule that characterizes ERRWs. Markov exchangeability can be easily verified through condition (10), as it only involves disjoint vectors $\mathbf{v}, \mathbf{w}, \mathbf{w}, i$, and, in this case, colors between the vertices of disjoint vectors are all distinct. Thus, for any element u_l of \mathbf{u}

$$\mathbf{T}_{C(u_l)}(x_0, \mathbf{x}, i, u_1, \dots, u_l) = \mathbf{T}_{C(u_l)}(x_0, \mathbf{x}, i, \mathbf{v}, \mathbf{w}, i, u_1, \dots, u_l),$$

and similarly for \mathbf{v} . For vertex i , the terms involved in (10) are $\mathbf{T}_{C(i)}(x_0, \mathbf{x}, i)$ and $\mathbf{T}_{C(i)}(x_0, \mathbf{x}, i, \mathbf{u}, \mathbf{w}, i) = \mathbf{T}_{C(i)}(x_0, \mathbf{x}, i, \mathbf{v}, \mathbf{w}, i) = \mathbf{T}_{C(i)}(x_0, \mathbf{x}, i) + 2$. Similarly for \mathbf{w} .

5.1. Colored edge reinforced random walks with partitioned colors

Suppose that the graph's colors are partitioned into groups $\{C_1, \dots, C_N\}$, $N \leq \infty$, such that, for all i , there exist $m \leq N$ such that $C(i) = C_m$. In other words, for every pair of vertices i and j , either $C(i) = C(j)$ or $C(i) \cap C(j) = \emptyset$. As a consequence, the left hand side of (10) only depends on the transitions through edges with colors in the sets C_1, \dots, C_N in (i, \mathbf{y}) , which are the same for (i, \mathbf{y}') if $\mathbf{y} = (\mathbf{u}, \mathbf{w}, i, \mathbf{v}, \mathbf{w}, i)$ and $\mathbf{y}' = (\mathbf{v}, \mathbf{w}, i, \mathbf{u}, \mathbf{w}, i)$. Thus, (10) holds and the process \mathbf{X} is Markov exchangeable.

The process is recurrent if $\inf_i \alpha_{c(i, x_0)} > 0$ and $\inf_i \beta_{i, x_0} > 0$, by Theorem 3. Recurrence holds under a milder restriction if the state space is finite. Let $Q = [Q_{i, j}]$ be the matrix of normalized edge weights $Q_{i, j} = \beta_{i, j} / \sum_{j'} \beta_{i, j'}$, with $Q_{i, j} = 0$ if the edge (i, j) is not present in the graph.

Proposition 3. *Let \mathbf{X} be a colored edges reinforced random walk with partitioned set of colors, finite state space and irreducible weight matrix Q . Then \mathbf{X} is recurrent.*

The proof is given in the Appendix. One first shows that, in this case, if a set of colors C_m is visited infinitely often, then every color in C_m is visited infinitely often, almost surely; and if a set of edges E_c is visited infinitely often, every edge in E_c is visited infinitely often, almost surely.

It follows from Proposition 3 that a colored edge reinforced random walk \mathbf{X} with partitioned set of colors, finite state space and irreducible weight matrix Q is a mixture of recurrent Markov chains. The prior is the limit probability law of the predictive distributions, as discussed in Section 3. Consider a state i with $C(i) = C_m$ for some $m \in \{1, \dots, N\}$. From expression (9), in the notation of Theorem 1, $P(X_{\tau_n(i)+1} = j | \mathcal{F}_{\tau_n(i)})$ is the product of two terms, the predictive probability of choosing the color, say c , of the edge (i, j) , times the predictive probability, given $\mathcal{F}_{\tau_n(i)}$ and the color c chosen, of picking the edge (i, j) among the edges of color c from i . One can easily see that the two terms have the expression of the predictive probabilities of having next color c and next state j for two independent exchangeable sequences, respectively: an exchangeable sequence of colors in C_m , with directing measure $(\tilde{P}_m(c), c \in C_m)$ having Dirichlet distribution with parameters $(\alpha(c), c \in C_m)$, and an exchangeable sequence of states in $A_{i, c} = \{j' : c(i, j') = c\}$, with directing measure $(\tilde{P}(j|i, c), j \in A_{i, c})$ having Dirichlet distribution with parameters $(\beta_{i, j}, j \in A_{i, c})$, independently of \tilde{P}_m . Because, for exchangeable sequences, the predictive distributions converge almost surely to the directing measure, it follows that $P(X_{\tau_n(i)+1} = j | \mathcal{F}_{\tau_n(i)})$ converges almost surely to

$$\tilde{P}_{i, j} = \tilde{P}_m(c(i, j)) \tilde{P}(j|i, c(i, j)).$$

Thus, for a state i with $C(i) = C_m$, the i th row of the random transition matrix, as a measure on S , is given by

$$\tilde{P}_i(\cdot) = \sum_{c \in C_m} \tilde{P}_m(c) \sum_{y \in A_{i, c}} \tilde{P}(y|i, c) \delta_y(\cdot), \tag{13}$$

where $(\tilde{P}_m(c), c \in C_m) \sim \text{Dirichlet}(\alpha(c), c \in C_m)$, $(\tilde{P}(j|i, c), j \in A_{i,c}) \sim \text{Dirichlet}(\beta_{i,j}, j \in A_{i,c})$ for any $c \in C_m$, and \tilde{P}_m and the $\tilde{P}(\cdot|i, c)$ are independent. In other words, the colors of the edges from i determine a partition of the sample space, and the transition probability $\tilde{P}_i(\cdot)$ from i first picks a region $A_{i,c}$ in the partition, with probability $\tilde{P}_m(c)$, then selects a state $y \in A_{i,c}$, with probability $\tilde{P}(y|i, c)$. The rows of the random transition matrix are in general dependent, through the common component $\tilde{P}_m(c)$. The same results hold for infinite state space if the process \mathbf{X} is recurrent, with the understanding that, if C_m or $A_{i,c}$ are countable, the Dirichlet distributions above become the appropriate Dirichlet processes.

Regarding the transition probability law \tilde{P}_i as a probability measure on $C(i) \times S$, the prior reveals some analogies with the Enriched Dirichlet process [37] for a bivariate random distribution, which arises as a nonparametric extension of the generalized Dirichlet distribution.

Let $\mu(\cdot; x_0, \alpha, \beta)$ be the prior distribution defined by (13) for a colored ERRW with weights α, β , and let \mathcal{D} denote the family of such prior distributions for different starting values and weights. The family \mathcal{D} is closed under sampling, that is, the posterior distribution still belongs to \mathcal{D} . Moreover, there exists a closed-form mapping to the posterior parameters (sometimes called *functional conjugacy*; see the discussion in [25]): if \mathbf{X} has prior $\mu(\cdot; x_0, \alpha, \beta)$, then the posterior distribution, given $\mathbf{X}_{1:n} = \mathbf{x}_{1:n}$, is $\mu(\cdot; x_n, \alpha(x_0, \mathbf{x}_{1:n}), \beta(x_0, \mathbf{x}_{1:n}))$, where

$$\alpha(x_0, \mathbf{x}_{1:n}) = \alpha_c + \mathbf{T}_{E_c}(x_0, \mathbf{x}_{1:n}), \quad \beta(x_0, \mathbf{x}_{1:n})_{i,j} = \beta_{i,j} + \mathbf{T}_{i,j}(x_0, \mathbf{x}_{1:n}). \tag{14}$$

This implies that the posterior distribution belongs to \mathcal{D} , thus \mathcal{D} is closed under sampling. To prove (14), notice that the conditional distribution of $(X_{n+1}, X_{n+2}, \dots)$, given $\mathbf{X}_{1:n} = \mathbf{x}_{1:n}$, is the distribution of the same colored edge reinforced random walk, but with initial state x_n and weights $\alpha(x_0, \mathbf{x}_{1:n}), \beta(x_0, \mathbf{x}_{1:n})$.

The following examples provide further insights into the nature of the process.

Example 3 (Independent enriched Dirichlet rows). Suppose that the set of colors associated to distinct vertices are all different, that is, $C(i) \cap C(j) = \emptyset$ for $i \neq j$. In this case, no probabilistic dependence is induced through the predictive distributions and the resulting random transition matrix has independent rows. As the Dirichlet process, the prior distribution on the rows is closed under sampling, but allows more flexibility in having the choice of two scale parameters, rather than just one. For example, suppose that the graph represents a physical network whose nodes are partitioned in local nets A_1, \dots, A_k , with \mathbf{X} describing some flow of information through the network, and assume one wishes to express the prior information that, from a node i , many local networks A_m are visited, but only a few states inside each local network tend to be visited. Such prior information could not be expressed by a Dirichlet process, but, by the clustering properties of Dirichlet processes, it could be incorporated in the prior (13) by choosing a large value of the normalizing constant $\alpha_{C(i)}$ and small values for β_{i,E_c} .

Example 4 (Analytic constraints). The case where some vertices share the same group of colors is somehow opposite to the previous example. The predictive probabilities (9) imply analytic constraints on the random transition matrix. Indeed, from (13), we obtain

$$\sum_{y \in A_{i,c}} \tilde{P}_{i,y} = \sum_{y \in A_{i',c}} \tilde{P}_{i',y} = \tilde{P}_m(c). \tag{15}$$

Thus, sums, by row, of probabilities of transitions along edges of the same color are constant. As a simple example, consider a graph with loops. Let the loops be colored in red and all other edges be colored in blue. The predictive distributions induce a mixture of random walks with equal transition probabilities on loops.

Example 5 (Colored edge reinforced random walks with dummy states). One way of introducing dependence among the rows of the random transition matrix, without strictly imposing constraints as those of expression (15), is to augment the graph G with auxiliary dummy states and edges, as follows. Consider a graph G , that can have partitioned colors or not. For an edge (i, j) in G , let us introduce a new vertex i^* between i and j , by adding a state i^* together with the edges (i, i^*) and (i^*, j) . The construction is repeated, adding one or more dummy vertices for each edge (i, j) in a set I . Let $I_{i,j}^*$ denote the set of dummy states between i and j ; I^* the set of dummy states and $S^* = S \cup I^*$.

If the augmented graph G^* has partitioned colors, then an ERRW \mathbf{X}^* on G^* , with weights α^*, β^* , and starting at $x_0 \in S$, is Markov exchangeable. Moreover, if recurrent, \mathbf{X}^* is a mixture of Markov chains with prior distribution $\mu^*(\cdot; x_0, \alpha^*, \beta^*)$ on the random transition matrix \tilde{P}^* .

Now, let \mathbf{X} be the process defined by deleting the dummy states from the paths of \mathbf{X}^* . The process \mathbf{X} is well defined and its probability distribution can be recovered from the law of \mathbf{X}^* , as

$$p(\mathbf{x}) = P(\mathbf{X}_{1:n} = \mathbf{x}) = \sum_{(x_0, \mathbf{x}^*) \in A(x_0, \mathbf{x})} P(\mathbf{X}_{1:n+m}^* = \mathbf{x}^*) = \sum_{(x_0, \mathbf{x}^*) \in A(x_0, \mathbf{x})} p^*(\mathbf{x}^*), \quad (16)$$

where $(n + m)$ is the length of \mathbf{x}^* and the sum is taken over the set $A(x_0, \mathbf{x})$ of all sequences (x_0, \mathbf{x}^*) consistent with (x_0, \mathbf{x}) ; that is, sequences that start and end as (x_0, \mathbf{x}) and lead to (x_0, \mathbf{x}) when deleting the dummy states. The process \mathbf{X} is Markov exchangeable. If $(x_0, \mathbf{x}) \sim (x_0, \mathbf{x}')$, it is fairly simple to see that the elements of $A(x_0, \mathbf{x})$ and $A(x_0, \mathbf{x}')$ are pairwise equivalent. Thus, $p(\mathbf{x}) = p(\mathbf{x}')$.

Moreover, if \mathbf{X}^* is recurrent, then \mathbf{X} is recurrent, as well, because $x_0 \in S$. In this case, \mathbf{X} is a mixture of recurrent Markov chains, and the random transition probabilities $\tilde{P}_{i,j}$ can be obtained by the transformation

$$\tilde{P}_{i,j} = \tilde{P}_{i,j}^* + \sum_{i^* \in I_{i,j}^*} \tilde{P}_{i,i^*}^*$$

of the random transition probabilities $[\tilde{P}_{i,j}^*] \sim \mu^*(\cdot; x_0, \alpha, \beta)$ for the process \mathbf{X}^* . The prior distribution for \tilde{P} redistributes the masses assigned by μ^* according to the dummy states that have been introduced, and generally expresses correlation across the rows of the transition matrix.

As an illustration, consider a monochromatic finite graph G , of color c_1 , with k vertices, including loops. As discussed before, a colored ERRW on G with edge weights $\beta_{i,j}$ is a Hoppe reinforced urn process with $\alpha_i q_i(j) = \beta_{i,j}$. When recurrent, such process is a mixture of Markov chains, with independent Dirichlet($\beta_{i,1}, \dots, \beta_{i,k}$) distributions on the rows of the transition matrix. Let us augment the graph by adding, for every vertex i , a dummy vertex i^* along with edges (i, i^*) and (i^*, i) . We color all edges (i, i^*) with a single color $c_2 \neq c_1$, and the edges (i^*, i) with a color c_3 different from c_1 and c_2 . This results in partitioned colors on the extended graph

G^* , the partition being $(C_1 = (c_1, c_2), C_2 = (c_3))$. Let \mathbf{X}^* be a recurrent colored ERRW on the augmented graph, and \mathbf{X} defined from \mathbf{X}^* as discussed above. Then, \mathbf{X} is a mixture of Markov chains with random transition probabilities:

$$\begin{aligned} \tilde{P}_{i,j} &= \tilde{P}_{i,j}^* = \tilde{P}_1^*(c_1)\tilde{P}^*(j|i, c_1) && \text{for } j \neq i, \\ \tilde{P}_{i,i} &= \tilde{P}_{i,i}^* + \tilde{P}_{i,i^*}^* = \tilde{P}_1^*(c_1)\tilde{P}^*(i|i, c_1) + (1 - \tilde{P}_1^*(c_1)). \end{aligned}$$

The i th row of the transition matrix \tilde{P} is a mixture of the probability distribution $(\tilde{P}^*(j|i, c_1), j \in S) \sim \text{Dirichlet}(\beta_{i,1}^*, \dots, \beta_{i,k}^*)$, the same distribution as for a Hoppe reinforced urn process, and a degenerate distribution on i , with mixing weight $\tilde{P}_1^*(c_1) \sim \text{Beta}(\alpha_{c_1}^*, \alpha_{c_2}^*)$, independent of the $\tilde{P}^*(j|i, c_1)$. The probability of a loop is a weighted average of $\tilde{P}^*(i|i, c_1)$ and $\tilde{P}^*(i^*|i, c_2) = 1$, thus it is higher than for the corresponding reinforced urn process. The rows of the transition matrix are correlated, due to the common random variable $\tilde{P}_1^*(c_1)$, but no strict equality is imposed, differently from Example 4.

As a further example, dummy states can be used to express prior inequalities between the elements of the random transition matrix. Suppose that, in the original graph G , edges (i, j) and (i', j') , j and j' distinct or not, have the same color c_1 and are the only edges from i and from i' with color c_1 . A dummy state i^* is added between i and j , together with edges (i, i^*) and (i^*, j) . The edge (i, i^*) is given a color $c_2 \neq c_1$. If necessary, other dummy vertices are added, until the extended graph has partitioned colors, never including edges of color c_1 from i or from i' , nor dummy vertices between i' and j' . The processes \mathbf{X}^* and \mathbf{X} are constructed as above. The random transition probabilities for the process \mathbf{X} satisfy

$$\tilde{P}_{i,j} = \tilde{P}_1^*(c_1) + \tilde{P}_1^*(c_2)\tilde{P}^*(j|i, c_2); \quad \tilde{P}_{i',j'} = \tilde{P}_1^*(c_1).$$

Hence, $\tilde{P}_{i,j} > \tilde{P}_{i',j'}$.

The construction by dummy states can be quite flexible in expressing prior beliefs. Notice that it is always possible to augment a graph G until the extended graph has partitioned colors, hence, the proposed construction can be used to define a Markov exchangeable reinforced random walk on any graph G . However, while the prior distribution μ^* for the process \mathbf{X}^* is closed under sampling, this is no longer true for the prior μ for \mathbf{X} . In principle, computing the posterior distribution $\mu(\cdot|\mathbf{x})$ on the random transition matrix \tilde{P} for \mathbf{X} , given observations $\mathbf{x} = \mathbf{x}_{1:n}$, remains simple, since

$$\mu(\cdot|\mathbf{x}) = \sum_{(x_0, \mathbf{x}^*) \in A(x_0, \mathbf{x})} \mu^*(\cdot; x_n, \boldsymbol{\alpha}^*(x_0, \mathbf{x}^*), \boldsymbol{\beta}^*(x_0, \mathbf{x}^*)) p_{X^*|X}(\mathbf{x}^*|\mathbf{x}), \tag{17}$$

where the conditional probability $p_{X^*|X}(\mathbf{x}^*|\mathbf{x}) = p^*(\mathbf{x}^*)/p(\mathbf{x})$ can be easily computed from p^* and using (16). The summation in (16) can actually be rearranged into a smaller number of terms. The set $A(x_0, \mathbf{x})$ can be partitioned in classes of equivalent vectors $A_{\mathbf{m}} = \{(x_0, \mathbf{x}^*) \in A(x_0, \mathbf{x}^*) : \mathbf{m}(x_0, \mathbf{x}^*) = \mathbf{m}\}$, where $\mathbf{m}(x_0, \mathbf{x}^*) = \mathbf{m}(x_0, \mathbf{x}^*) = (m_{i^*}(x_0, \mathbf{x}^*), i^* \in I^*)$ and $m_{i^*}(x_0, \mathbf{x}^*)$ denotes the number of visits to the dummy state i^* in (x_0, \mathbf{x}^*) . For any \mathbf{m} , all the strings in $A_{\mathbf{m}}$ are equivalent and, therefore, have the same probability. A simple combinatorial argument shows

that the cardinality of $A_{\mathbf{m}}$ is

$$N_{\mathbf{m}} = \prod_{(i,j) \in I} \frac{\mathbf{T}_{i,j}(x_0, \mathbf{x})!}{(\mathbf{T}_{i,j}(x_0, \mathbf{x}) - m_{i,j}^*)! \prod_{i^* \in I_{i,j}^*} m_{i^*}^*!},$$

where $m_{i,j}^* = \sum_{i^* \in I_{i,j}^*} m_{i^*}^*$. Thus,

$$p(\mathbf{x}) = \sum_{\mathbf{m}} p^*(\mathbf{x}_{\mathbf{m}}) N_{\mathbf{m}}, \tag{18}$$

where $(x_0, \mathbf{x}_{\mathbf{m}})$ is any string in $A_{\mathbf{m}}$. Although reduced with respect to (16), the number of terms in (18) still explodes for long paths, or when the number of dummy states is large. Yet, simulation techniques can be used to sample from $p_{X^*|X}(\mathbf{x}^*|\mathbf{x})$ and to compute a Monte Carlo approximation of (17). To simulate from $p_{X^*|X}$, one has to simulate the possibly missing passages through dummy states. Notice that a string (x_0, \mathbf{x}^*) in $A(x_0, \mathbf{x})$ can be described in terms of the successors states: each state $x_i \notin I$ in the observed sample (x_0, \mathbf{x}) has a known successor state in \mathbf{x}^* ; while for each observed pair of consecutive values (x_i, j) , where $x_i \in I$, the successor of x_i in \mathbf{x}^* can be any state in the set $\{j, I_{x_i,j}^*\}$. By construction, the successor of a dummy state $i^* \in I_{i,j}^*$ is always equal to j . Thus, generating \mathbf{x}^* from $p_{X^*|X}(\mathbf{x}^*|\mathbf{x})$ is equivalent to generating the sequence of unknown successor states \mathbf{V}^* from the conditional distribution $P(\mathbf{V}^* = \mathbf{v}^* | \mathbf{V} = \mathbf{v}, \mathbf{V}^* \in B)$, given the known successors \mathbf{V} and the appropriate constraints on the remaining ones, denoted by $\mathbf{V}^* \in B$. The simplest Gibbs sampling scheme consists in generating the \mathbf{V}^* one at the time from their full conditional distributions. Assume that we want to generate a successor V_i of state i , corresponding to consecutive observations (i, j) in (x_0, \mathbf{x}) . By partial exchangeability of the matrix of successor states, one can permute the successors so that V_i is the last successor of state i . Then, its full conditional distribution selects a state k in $\{j, I_{i,j}^*\}$ with probability proportional to

$$\frac{\alpha_{c(i,k)}^* + \mathbf{T}_{c(i,k)}}{\alpha_{C(i)}^* + \mathbf{T}_{C(i)}} \frac{\beta_{i,k}^* + \mathbf{T}_{i,k}}{\beta_{i,E_c(i,k)}^* + \mathbf{T}_{i,E_c(i,k)}},$$

where the transitions counts are computed on all the other successor states.

Appendix: Complements and proofs

A.1. Complements for Section 2

The following lemma is used for the proof of Theorem 1. For a state i , let $\mathcal{F}_{V_{i,n}^{\checkmark}}$ denote the sigma algebra generated by $(V_{i,k}, k < n; V_{j,l}, j \in S^*, j \neq i, l \geq 1)$.

Lemma 4. *Under the hypotheses of Theorem 1, $\mathcal{F}_{\tau_n(i)} \subseteq \mathcal{F}_{V_{i,n}^{\checkmark}}$.*

Proof. Let us consider $B \in \mathcal{F}_{\tau_n(i)}$ and show that $B \in \mathcal{F}_{V_{i,n}^\vee}$. One can express B as

$$B = \bigcup_{k=1}^{\infty} (B \cap (\tau_n(i) = k)) \cup (B \cap (\tau_n(i) = \infty)).$$

We show that the events: (i) $B \cap (\tau_n(i) = k)$, $k \geq 1$ and (ii) $B \cap (\tau_n(i) = \infty)$ belong to $\mathcal{F}_{V_{i,n}^\vee}$.

(i) Let \mathcal{F}_k be the sigma-field generated by $\mathbf{X}_{1:k}$. Since $B \cap (\tau_n(i) = k) \in \mathcal{F}_k$, there exists $B_{n,k}$ such that

$$B \cap (\tau_n(i) \leq k) = \bigcup_{\mathbf{x}_{1:k} \in B_{n,k}} (\mathbf{X}_{1:k} = \mathbf{x}_{1:k}).$$

It is proved in [14] that, for any $\mathbf{x}_{1:k}$, there exist $m, s_1, \dots, s_m, n_{s_1}, \dots, n_{s_m}$ and $v_{s_1,1}, \dots, v_{s_1,n_{s_1}}, \dots, v_{s_m,1}, \dots, v_{s_m,n_{s_m}}$ such that

$$\begin{aligned} (\mathbf{X}_{1:k} = \mathbf{x}_{1:k}) & \\ &= (V_{s_1,1} = v_{s_1,1}, \dots, V_{s_1,n_{s_1}} = v_{s_1,n_{s_1}}, \dots, V_{s_m,1} = v_{s_m,1}, \dots, V_{s_m,n_{s_m}} = v_{s_m,n_{s_m}}). \end{aligned} \tag{19}$$

For $\mathbf{x}_{1:k} \in B_{n,k}$, one has $(\mathbf{X}_{1:k} = \mathbf{x}_{1:k}) \subseteq (\tau_n(i) = k)$, thus $n_i = n - 1$; therefore, $B \cap (\tau_n(i) = k) \in \mathcal{F}_{V_{i,n}^\vee}$.

(ii) To prove that $B \cap (\tau_n(i) = \infty) \in \mathcal{F}_{V_{i,n}^\vee}$, let $\mathcal{C} = \{A \in \mathcal{F} : A \cap (\tau_n(i) = \infty) \in \mathcal{F}_{V_{i,n}^\vee}\}$. It is easy to verify that \mathcal{C} is a sigma-algebra. Furthermore, for every k and $\mathbf{x}_{1:k} \in S^k$, the event $(\mathbf{X}_{1:k} = \mathbf{x}_{1:k})$ belongs to \mathcal{C} . Indeed, if $m, s_1, \dots, s_m, n_{s_1}, \dots, n_{s_m}$ and $v_{s_1,1}, \dots, v_{s_1,n_{s_1}}, \dots, v_{s_m,1}, \dots, v_{s_m,n_{s_m}}$ satisfy (19), then $(\mathbf{X}_{1:k} = \mathbf{x}_{1:k}) \cap (\tau_n(i) = \infty) = \emptyset$ if $n_i \geq n$ or if $n_i = n - 1$ and $x_k = i$; while, for $n_i < n - 1$ or for $n_i = n - 1$, with $x_k \neq i$, one has

$$\begin{aligned} &(\mathbf{X}_{1:k} = \mathbf{x}_{1:k}) \cap (\tau_n(i) = \infty) \\ &= \left(\prod_{j=1}^m \prod_{l=1}^{n_{s_j}} (V_{s_j,l} = v_{s_j,l}) \right) \cap \left(\prod_{s_j \neq i} \prod_{l=n_{s_j}+1}^{\infty} (V_{s_j,l} \neq i) \right) \cap \left(\prod_{s \neq s_1, \dots, s_m} \prod_{l=1}^{\infty} (V_{s,l} \neq i) \right). \end{aligned}$$

Therefore, for every k and $\mathbf{x}_{1:k}$, the event $(\mathbf{X}_{1:k} = \mathbf{x}_{1:k})$ belongs to \mathcal{C} . It follows that $\mathcal{F}_k \subseteq \mathcal{C}$ for every k . Since \mathcal{C} is a sigma-algebra and includes $\bigcup_{k=1}^{\infty} \mathcal{F}_k$, it includes the sigma-algebra $\bigvee_k \mathcal{F}_k$ generated by $\bigcup_{k=1}^{\infty} \mathcal{F}_k$. Let $B \in \mathcal{F}_{\tau_n(i)}$. Since $\mathcal{F}_{\tau_n(i)} \subseteq \bigvee_k \mathcal{F}_k$, then $B \in \mathcal{C}$; therefore, $B \cap (\tau_n(i) = \infty) \in \mathcal{F}_{V_{i,n}^\vee}$. \square

The following corollary rephrases and proves Corollary 1.

Corollary 3. *Let \mathbf{X} be a mixture of recurrent Markov chains with random transition matrix \tilde{P} . The rows of \tilde{P} are stochastically independent if and only if*

$$P(X_{\tau_n(i)+1} = j | \mathbf{T}(x_0, \mathbf{X}_{1:\tau_n(i)})) = P(X_{\tau_n(i)+1} = j | \mathbf{T}_i(x_0, \mathbf{X}_{1:\tau_n(i)})) \quad a.s.$$

Furthermore, in this case,

$$P(X_{\tau_n(i)+1} = j | \mathbf{T}_i(x_0, \mathbf{X}_{1:\tau_n(i)}) = \mathbf{t}_i) = P\left(V_{i,n} = j \mid \sum_{l=1}^{n-1} \delta_{V_{i,l}}(k) = \mathbf{t}_{i,k} : k \in S\right)$$

almost surely with respect to the probability distribution of $\mathbf{T}_i(x_0, \mathbf{X}_{1:\tau_n(i)})$.

Proof. Suppose first that the rows of \tilde{P} are stochastically independent. By Lemma 4, $\mathcal{F}_{\tau_n(i)} \subseteq \mathcal{F}_{V_{i,n}}$ for $i \in S$, and by the independence assumption,

$$P(V_{i,n} = j | \mathcal{F}_{V_{i,n}}) = P(V_{i,n} = j | V_{i,1}, \dots, V_{i,n-1}).$$

Hence, the predictive probabilities from i satisfy

$$\begin{aligned} P(X_{\tau_n(i)+1} = j | \mathcal{F}_{\tau_n(i)}) &= E(P(V_{i,n} = j | \mathcal{F}_{V_{i,n}}) | \mathcal{F}_{\tau_n(i)}) \\ &= E(P(V_{i,n} = j | V_{i,1}, \dots, V_{i,n-1}) | \mathcal{F}_{\tau_n(i)}) \\ &= P(V_{i,n} = j | V_{i,1}, \dots, V_{i,n-1}). \end{aligned} \quad (20)$$

This proves that the transition probabilities from i depend on $\mathbf{X}_{1:\tau_n(i)}$ only through the transition counts from i .

Conversely, suppose that, for every j, n and $\mathbf{x}_{1:n}$, the predictive probability $p(j|x_0, \mathbf{x}_{1:n})$ is a function $\pi(j|x_n, \mathbf{T}_{x_n})$ of x_n and of the transition counts \mathbf{T}_{x_n} from x_n . Then $P(X_{\tau_n(i)+1} = j | \mathcal{F}_{\tau_n(i)})$ is a function of $V_{i,1}, \dots, V_{i,n-1}$. This implies that the rows of the array $(V_{i,n}, i \in S, n \geq 1)$ are stochastically independent. Since the random transition probability $\tilde{P}_{i,j}$ is the limit of the relative frequency of j in $V_{i,1}, V_{i,2}, \dots$, then the rows of \tilde{P} are stochastically independent as well.

To prove the last assertion, notice that, if the sequences $(V_{i,n}, n \geq 1)_{i \in S^*}$ are independent, then (20) follows, and by exchangeability of $(V_{i,n}, n \geq 1)$ we have the thesis. \square

A.2. Complements for Section 3

Proof of Lemma 3. The existence of a minimal element can be proved as for Lemma 1. We want to prove that the minimal element $\mathbf{x}^* = (\mathbf{u}^*, k, j, \mathbf{v}^*, k', j')$ satisfies:

- (1) $\{\mathbf{u}^*, k\} \cap \{j, \mathbf{v}^*, k'\}$ contains only one element, say i , that appears once in (\mathbf{u}^*, k) , as the first element, and once in (j, \mathbf{v}^*, k') ;
- (2) if $i \neq j$, then (\mathbf{v}^*, k') contains no j after i , and the strings before and after i in \mathbf{v}^* have no common elements.

Let us prove (1). Had (\mathbf{u}^*, k) and (j, \mathbf{v}^*, k') more than one common element, we could delete the first part of \mathbf{x}^* and obtain a string in \mathcal{C} that is shorter than \mathbf{x}^* . Let us denote by i the only common element. For the same reason, the state i has to appear once in (\mathbf{u}^*, k) , as its first element. For showing that i appears only once in (j, \mathbf{v}^*, k') , assume by contradiction that i appears at least

twice. Let us distinguish the following cases: $i = j$, $i = k$, $i = k'$ and $i \neq j, k, k'$. If $i = j$, then, by the above result, j is the first element of \mathbf{u}^* , say $\mathbf{u}^* = (j, \mathbf{u}_1^*)$; then \mathbf{v}^* cannot contain j , otherwise we could write

$$\mathbf{x}^* = (j, \mathbf{u}_1^*, k, j, \mathbf{v}_1^*, j, \mathbf{v}_2^*, k', j) \sim (j, \mathbf{v}_1^*, j, \mathbf{u}_1^*, k, j, \mathbf{v}_2^*, k', j)$$

and obtain a shorter string in \mathcal{C} by deleting (j, \mathbf{v}_1^*) . Similarly, for $i = k$, we could write $\mathbf{x}^* = (k, \mathbf{u}^*, k, j, \mathbf{v}_1^*, k, \mathbf{v}_2^*, k, \mathbf{v}_3^*, k', j) \sim (k, \mathbf{v}_2^*, k, \mathbf{u}^*, k, j, \mathbf{v}_1^*, k, \mathbf{v}_3^*, k', j)$ and obtain a shorter element of \mathcal{C} by deleting (k, \mathbf{v}_2^*) . For $i = k'$, we could set $\mathbf{x}^* = (k', \mathbf{u}^*, k, j, \mathbf{v}_1^*, k', \mathbf{v}_2^*, k', j) \sim (k', \mathbf{v}_2^*, k', \mathbf{u}^*, k, j, \mathbf{v}_1^*, k', j)$ and a shorter element could be obtained by deleting (k', \mathbf{v}_2^*) . For $i \neq j, k, k'$, we could write $\mathbf{x}^* = (i, \mathbf{u}^*, k, j, \mathbf{v}_1^*, i, \mathbf{v}_2^*, i, \mathbf{v}_3^*, k', j) \sim (i, \mathbf{v}_2^*, i, \mathbf{u}^*, k, j, \mathbf{v}_1^*, i, \mathbf{v}_3^*, k', j)$ and we could delete (i, \mathbf{v}_2^*) .

Let us prove assertion (2). We distinguish the following cases: $i = k$; $i = k'$, and $i \neq k, k'$. Suppose, by contradiction, that (2) does not hold. Then:

If $i = k$, we could write $\mathbf{x}^* = (k, j, \mathbf{v}_1^*, k, \mathbf{v}_2^*, j, \mathbf{v}_3^*, k', j) \sim (k, \mathbf{v}_2^*, j, \mathbf{v}_1^*, k, j, \mathbf{v}_3^*, k', j)$ and obtain an element of \mathcal{C} that is shorter than \mathbf{x}^* by deleting (k, \mathbf{v}_2^*) ; furthermore, we could write $\mathbf{x}^* = (k, j, \mathbf{v}_1^*, w, \mathbf{v}_2^*, k, \mathbf{v}_3^*, w, \mathbf{v}_4^*, k', j) \sim (k, \mathbf{v}_3^*, w, \mathbf{v}_2^*, k, j, \mathbf{v}_1^*, w, \mathbf{v}_4^*, k', j)$ and obtain a shorter element by deleting (k, \mathbf{v}_3^*) ;

If $i = k'$, we could write $\mathbf{x}^* = (k', \mathbf{u}^*, k, j, \mathbf{v}_1^*, k', \mathbf{v}_2^*, j, \mathbf{v}_3^*, k', j) \sim (k', \mathbf{v}_2^*, j, \mathbf{v}_1^*, k', \mathbf{u}^*, k, j, \mathbf{v}_3^*, k', j)$ and $\mathbf{x}^* = (k', \mathbf{u}^*, k, j, \mathbf{v}_1^*, w, \mathbf{v}_2^*, k', \mathbf{v}_3^*, w, \mathbf{v}_4^*, k', j) \sim (k', \mathbf{v}_3^*, w, \mathbf{v}_2^*, k', \mathbf{u}^*, k, j, \mathbf{v}_1^*, w, \mathbf{v}_4^*, k', j)$, and obtain shorter elements in \mathcal{C} by deleting (k', \mathbf{v}_2^*) and (k', \mathbf{v}_3^*) , respectively;

If $i \neq k, k'$, we could write $\mathbf{x}^* = (i, \mathbf{u}^*, k, j, \mathbf{v}_1^*, i, \mathbf{v}_2^*, j, \mathbf{v}_3^*, k', j) \sim (i, \mathbf{v}_2^*, j, \mathbf{v}_1^*, i, \mathbf{u}^*, k, j, \mathbf{v}_3^*, k', j)$ and $\mathbf{x}^* = (i, \mathbf{u}^*, k, j, \mathbf{v}_1^*, w, \mathbf{v}_2^*, i, \mathbf{v}_3^*, w, \mathbf{v}_4^*, k', j) \sim (i, \mathbf{v}_3^*, w, \mathbf{v}_2^*, i, \mathbf{u}^*, k, j, \mathbf{v}_1^*, w, \mathbf{v}_4^*, k', j)$ and obtain shorter elements by deleting (i, \mathbf{v}_2^*) and (i, \mathbf{v}_3^*) , respectively. \square

A.3. Complements for Section 5

Proof of Proposition 3. Let \mathbf{X} be a colored edge reinforced random walk with partitioned colors $\{C_1, \dots, C_N\}$, finite state space and irreducible normalized weights matrix Q . We first show that:

(a) If a set of colors C_m is visited infinitely often, then every color $c \in C_m$ is visited infinitely often, almost surely.

(b) If a set of edges E_c is visited infinitely often, then every edge in E_c is visited infinitely often, almost surely.

Let us prove (a). For every $c \in C(i)$, let $A_{i,c} = \{j : c(i, j) = c\}$. We have

$$P(X_{n+1} \in A_{i,c} | \mathbf{X}_{1:n} = (\mathbf{x}, i)) = \frac{\alpha_c + \mathbf{T}_{E_c}(x_0, \mathbf{x}, i)}{\alpha_{C(i)} + \mathbf{T}_{C(i)}(x_0, \mathbf{x}, i)}.$$

Hence, if $c \in C_m$ and $\tau_n(C_m) = \inf\{t \geq \tau_{n-1}(C_m) : C(X_t) = C_m\}$ is the time of the n th visit to a state with color set equal to C_m , we have

$$P(c(X_{\tau_n(C_m)}, X_{\tau_n(C_m)+1}) = c | \tau_n(C_m) < \infty) \geq \frac{\alpha_c}{\alpha_{C_m} + n}.$$

It follows that

$$\sum_{n=1}^{\infty} [P(c(X_{\tau_n(C_m)}, X_{\tau_n(C_m)+1}) = c | \mathcal{F}_{\tau_n(C_m)}) \mathbb{1}_{\{\tau_n(C_m) < \infty\}} + \mathbb{1}_{\{\tau_n(C_m) = \infty\}}] = \infty \quad \text{a.s.}$$

Reasoning as in the proof of Theorem 4, we obtain that, if C_m is visited infinitely often, then every $c \in C_m$ is visited infinitely often, almost surely.

(b). Let $(i, y) \in E_c$ and $\mathbf{T}(x_0, \mathbf{x}) = \mathbf{T}$. Then

$$\begin{aligned} &P(X_{n+1} = y | \mathbf{X}_{1:n-1} = \mathbf{x}, c(X_n, X_{n+1}) = c) \\ &= \frac{\sum_{i:c(i,y)=c} P(X_{n+1} = y | \mathbf{X}_{1:n} = (\mathbf{x}, i)) P(X_n = i | \mathbf{X}_{1:n-1} = \mathbf{x})}{\sum_{i:c(i,y)=c} P(c(X_n, X_{n+1}) = c | \mathbf{X}_{1:n} = (\mathbf{x}, i)) P(X_n = i | \mathbf{X}_{1:n-1} = \mathbf{x})} \\ &\geq \frac{\min_{i:(i,y) \in E_c} P(X_{n+1} = y | \mathbf{X}_{1:n} = (\mathbf{x}, i))}{\max_{i:(i,y) \in E_c} P(c(X_n, X_{n+1}) = c | \mathbf{X}_{1:n} = (\mathbf{x}, i))} \\ &\geq \frac{\alpha_c + \mathbf{T}_{E_c}}{\alpha_c + \mathbf{T}_{E_c} + 1} \frac{\min_{i:(i,y) \in E_c} \beta_{i,y}}{\max_{i:(i,y) \in E_c} \beta_{i,y} + \mathbf{T}_{E_c}}. \end{aligned}$$

Let $\tau_n(c)$ be the n th time the chain visits E_c . Then, for every y such that $c(i, y) = c$ for some i ,

$$P(X_{\tau_n(c)+1} = y | \tau_n(c) < \infty) \geq \frac{\alpha_c + \mathbf{T}_{E_c}}{\alpha_c + \mathbf{T}_{E_c} + 1} \frac{\min_{i:(i,y) \in E_c} \beta_{i,y}}{\max_{i:(i,y) \in E_c} \beta_{i,y} + n}.$$

Hence,

$$\sum_{n=1}^{\infty} [P(X_{\tau_n(c)+1} = y | \mathcal{F}_{\tau_n(c)}) \mathbb{1}_{\{\tau_n(c) < \infty\}} + \mathbb{1}_{\{\tau_n(c) = \infty\}}] = \infty \quad \text{a.s.}$$

Again as in the proof of Theorem 4, we obtain that, if E_c is visited infinitely often, then every edge in E_c is visited infinitely often, almost surely.

To conclude the proof, let i and u be such that $Q_{i,u} > 0$. If i is visited infinitely often, then $C(i)$ is visited infinitely often; by (a), every E_c with $c \in C(i)$ is visited infinitely often and, by (b), every $y \in A_{i,c}$ is visited infinitely often. Hence, u is visited infinitely often. Since S is finite, for every path of the process there exists a state i that is visited infinitely often. Then, every state j , such that $Q_{i,u_1} Q_{u_1,u_2} \cdots Q_{u_n,j} > 0$ for some u_1, \dots, u_n and n , is visited infinitely often. Since Q is irreducible, for every pair i, j there exist such u_1, \dots, u_n . Hence, the process is recurrent. \square

Acknowledgments

We thank an anonymous referee for suggesting a shorter proof of Theorems 3 and 4, and all the reviewers and the Associate Editor for their thoughtful comments.

References

- [1] Aldous, D.J. (1985). Exchangeability and related topics. In *École D'été de Probabilités de Saint-Flour, XIII—1983. Lecture Notes in Math.* **1117** 1–198. Berlin: Springer. [MR0883646](#)
- [2] Bacallado, S. (2011). Bayesian analysis of variable-order, reversible Markov chains. *Ann. Statist.* **39** 838–864. [MR2816340](#)
- [3] Bacallado, S., Favaro, S. and Trippa, L. (2013). Bayesian nonparametric analysis of reversible Markov chains. *Ann. Statist.* **41** 870–896. [MR3099124](#)
- [4] Beal, M.J., Ghahramani, Z. and Rasmussen, C.E. (2002). The infinite hidden Markov model. In *Adv. Neural Inf. Process. Syst.* **14** 577–584. Cambridge, MA: MIT Press.
- [5] Blackwell, D. and MacQueen, J.B. (1973). Ferguson distributions via Pólya urn schemes. *Ann. Statist.* **1** 353–355. [MR0362614](#)
- [6] Coppersmith, D. and Diaconis, P. (1987). Random walk with reinforcement. Unpublished manuscript.
- [7] de Finetti, B. (1937). La prévision : Ses lois logiques, ses sources subjectives. *Ann. Inst. Henri Poincaré* **7** 1–68. [MR1508036](#)
- [8] de Finetti, B. (1959). La probabilità e la statistica nei rapporti con l'induzione secondo i diversi punti di vista. In *Centro Internazionale Matematico Estivo (CIME), Induzione e Statistica* 1–115. Roma: Cremonese. English translation in B. de Finetti (1972). *Probability, Induction and Statistics*, 147–227.
- [9] Diaconis, P. (1988). Recent progress on de Finetti's notions of exchangeability. In *Bayesian Statistics, 3 (Valencia, 1987)* (J.M. Bernardo, M.H. DeGroot, D.V. Lindley and A.F.M. Smith, eds.). *Oxford Sci. Publ.* 111–125. New York: Oxford Univ. Press. [MR1008047](#)
- [10] Diaconis, P. and Freedman, D. (1980). de Finetti's theorem for Markov chains. *Ann. Probab.* **8** 115–130. [MR0556418](#)
- [11] Diaconis, P. and Rolles, S.W.W. (2006). Bayesian analysis for reversible Markov chains. *Ann. Statist.* **34** 1270–1292. [MR2278358](#)
- [12] Diaconis, P. and Ylvisaker, D. (1979). Conjugate priors for exponential families. *Ann. Statist.* **7** 269–281. [MR0520238](#)
- [13] Ewens, W.J. (1972). The sampling theory of selectively neutral alleles. *Theor. Popul. Biol.* **3** 87–112. [MR0325177](#)
- [14] Fortini, S., Ladelli, L., Petris, G. and Regazzini, E. (2002). On mixtures of distributions of Markov chains. *Stochastic Process. Appl.* **100** 147–165. [MR1919611](#)
- [15] Fortini, S., Ladelli, L. and Regazzini, E. (2000). Exchangeability, predictive distributions and parametric models. *Sankhya, Ser. A* **62** 86–109. [MR1769738](#)
- [16] Fortini, S. and Petrone, S. (2012). Hierarchical reinforced urn processes. *Statist. Probab. Lett.* **82** 1521–1529. [MR2930656](#)
- [17] Griffiths, T.L. and Ghahramani, Z. (2006). Infinite latent feature models and the Indian buffet process. In *Adv. Neural Inf. Process. Syst.* **18** 475–482. Cambridge, MA: MIT Press.
- [18] Hoppe, F.M. (1984). Pólya-like urns and the Ewens' sampling formula. *J. Math. Biol.* **20** 91–94. [MR0758915](#)
- [19] Keane, M.S. and Rolles, S.W.W. (2000). Edge-reinforced random walk on finite graphs. In *Infinite Dimensional Stochastic Analysis (Amsterdam, 1999)* (P. Clement, F. den Hollander, J. van Neerven and B. de Pagter, eds.). *Verh. Afd. Natuurkd. 1. Reeks. K. Ned. Akad. Wet.* **52** 217–234. Amsterdam: R. Neth. Acad. Arts Sci. [MR1832379](#)
- [20] Merkl, F. and Rolles, S.W.W. (2005). Edge-reinforced random walk on a ladder. *Ann. Probab.* **33** 2051–2093. [MR2184091](#)
- [21] Merkl, F. and Rolles, S.W.W. (2007). A random environment for linearly edge-reinforced random walks on infinite graphs. *Probab. Theory Related Fields* **138** 157–176. [MR2288067](#)

- [22] Merkl, F. and Rolles, S.W.W. (2007). Asymptotic behavior of edge-reinforced random walks. *Ann. Probab.* **35** 115–140. [MR2303945](#)
- [23] Merkl, F. and Rolles, S.W.W. (2009). Recurrence of edge-reinforced random walk on a two-dimensional graph. *Ann. Probab.* **37** 1679–1714. [MR2561431](#)
- [24] Muliere, P., Secchi, P. and Walker, S.G. (2000). Urn schemes and reinforced random walks. *Stochastic Process. Appl.* **88** 59–78. [MR1761692](#)
- [25] Orbanz, P. (2009). Functional conjugacy in parametric Bayesian models. Technical report, Univ. Cambridge.
- [26] Pemantle, R. (1988). Phase transition in reinforced random walk and RWRE on trees. *Ann. Probab.* **16** 1229–1241. [MR0942765](#)
- [27] Pemantle, R. (2007). A survey of random processes with reinforcement. *Probab. Surv.* **4** 1–79. [MR2282181](#)
- [28] Pemantle, R.A. (1988). Random processes with reinforcement. Ph.D. thesis, Dept. Mathematics, Massachusetts Institute of Technology.
- [29] Perman, M., Pitman, J. and Yor, M. (1992). Size-biased sampling of Poisson point processes and excursions. *Probab. Theory Related Fields* **92** 21–39. [MR1156448](#)
- [30] Pitman, J. (1996). Some developments of the Blackwell–MacQueen urn scheme. In *Statistics, Probability and Game Theory. Institute of Mathematical Statistics Lecture Notes – Monograph Series 30* 245–267. Hayward, CA: IMS. [MR1481784](#)
- [31] Rolles, S.W.W. (2003). How edge-reinforced random walk arises naturally. *Probab. Theory Related Fields* **126** 243–260. [MR1990056](#)
- [32] Rolles, S.W.W. (2006). On the recurrence of edge-reinforced random walk on $\mathbb{Z} \times G$. *Probab. Theory Related Fields* **135** 216–264. [MR2218872](#)
- [33] Sabot, C. and Tarrès, P. (2015). Edge-reinforced random walk, vertex-reinforced jump process and the supersymmetric hyperbolic sigma model. *J. Eur. Math. Soc. (JEMS)* **17** 2353–2378. [MR3420510](#)
- [34] Teh, Y.W. and Jordan, M.I. (2010). Hierarchical Bayesian nonparametric models with applications. In *Bayesian Nonparametrics* (N.L. Hjort, C. Holmes, P. Muller and S.G. Walker, eds.). *Camb. Ser. Stat. Probab. Math.* 158–207. Cambridge: Cambridge Univ. Press. [MR2730663](#)
- [35] Teh, Y.W., Jordan, M.I., Beal, M.J. and Blei, D.M. (2006). Hierarchical Dirichlet processes. *J. Amer. Statist. Assoc.* **101** 1566–1581. [MR2279480](#)
- [36] Thibaux, R. and Jordan, M.I. (2007). Hierarchical beta processes and the Indian buffet process. In *JMLR Workshops and Proceedings: AISTATS 2007* **2** 564–571.
- [37] Wade, S., Mongelluzzo, S. and Petrone, S. (2011). An enriched conjugate prior for Bayesian nonparametric inference. *Bayesian Anal.* **6** 359–385. [MR2843536](#)
- [38] Walker, S. and Muliere, P. (1997). Beta-Stacy processes and a generalization of the Pólya-urn scheme. *Ann. Statist.* **25** 1762–1780. [MR1463574](#)
- [39] Walker, S. and Muliere, P. (1999). A characterization of a neutral to the right prior via an extension of Johnson’s sufficientness postulate. *Ann. Statist.* **27** 589–599. [MR1714716](#)
- [40] Williams, D. (1991). *Probability with Martingales*. Cambridge: Cambridge Univ. Press. [MR1155402](#)
- [41] Zabell, S.L. (1982). W. E. Johnson’s “sufficientness” postulate. *Ann. Statist.* **10** 1090–1099. [MR0673645](#)
- [42] Zabell, S.L. (1995). Characterizing Markov exchangeable sequences. *J. Theoret. Probab.* **8** 175–178. [MR1308676](#)
- [43] Zaman, A. (1984). Urn models for Markov exchangeability. *Ann. Probab.* **12** 223–229. [MR0723741](#)

Received May 2014 and revised April 2015