

## BAYESIAN NONPARAMETRIC ANALYSIS OF REVERSIBLE MARKOV CHAINS

BY SERGIO BACALLADO<sup>§,\*</sup> STEFANO FAVARO<sup>¶,||,†</sup>  
AND LORENZO TRIPPA<sup>\*\*,\ddagger</sup>

*Stanford University*<sup>\*</sup>, *University of Torino*<sup>†</sup>  
*and Harvard School of Public Health*<sup>‡</sup>

We introduce a three-parameter random walk with reinforcement, called the  $(\theta, \alpha, \beta)$  scheme, which generalizes the linearly edge reinforced random walk to uncountable spaces. The parameter  $\beta$  smoothly tunes the  $(\theta, \alpha, \beta)$  scheme between this edge reinforced random walk and the classical exchangeable two-parameter Hoppe urn scheme, while the parameters  $\alpha$  and  $\theta$  modulate how many states are typically visited. Resorting to de Finetti's theorem for Markov chains, we use the  $(\theta, \alpha, \beta)$  scheme to define a nonparametric prior for Bayesian analysis of reversible Markov chains. The prior is applied in Bayesian nonparametric inference for species sampling problems with data generated from a reversible Markov chain with an unknown transition kernel. As a real example, we analyze data from molecular dynamics simulations of protein folding.

**1. Introduction.** The problem that motivated our study is the analysis of benchtop and computer experiments that produce dynamical data associated with the structural fluctuations of a protein in water. Frequently, the physical laws that govern these dynamics are time-reversible. Therefore, a stochastic model for the experiment should also be reversible. Reversible Markov models in particular have become widespread in the field of molecular dynamics (19). Modeling with reversible Markov chains is also natural in a number of other disciplines.

We consider the setting in which a scientist has a sequence of states  $X_1, \dots, X_n$  sampled from a reversible Markov chain. We propose a Bayesian

---

<sup>§</sup>This work was funded by grants NIH-R01-GM062868 and NSF-DMS-0900700.

<sup>¶</sup>This work was supported by the European Research Council (ERC) through StG N-BNP 306406.

<sup>||</sup>Also affiliated to Collegio Carlo Alberto, Moncalieri, Italy.

<sup>\*\*</sup>Also affiliated to the Dana-Farber Cancer Institute, Dept. of Biostatistics and Computational Biology.

*AMS 2000 subject classifications:* Primary 62M02; secondary 62C10.

*Keywords and phrases:* Reversibility, Mixtures of Markov chains, Reinforced random walks, Bayesian nonparametrics, Species sampling, Two-parameter Hoppe urn, Molecular dynamics.


model for a reversible Markov chain driven by an unknown transition kernel. Problems one can deal with using our model include (i) predicting how soon the process will return to a specific state of interest and (ii) predicting the number of states not yet explored by  $X_1, \dots, X_n$  that appear in the next  $m$  transitions  $X_{n+1}, \dots, X_{n+m}$ . More generally, the model can be used to predict any characteristic of the future trajectory of the process. Problems (i) and (ii) are of great interest in the analysis of computer experiments on protein dynamics.

Diaconis and Rolles (9) introduced a conjugate prior for Bayesian analysis of reversible Markov chains. This prior is defined via de Finetti's theorem for Markov chains (8). The predictive distribution is that of a linearly edge-reinforced random walk (ERRW) on an undirected graph (7). Much is known about the asymptotic properties of this process (15), its uniqueness (23), and its recurrence on infinite graphs (18, and references therein). Fortini, Petrone and Bacallado recently discussed other examples of Markov chain priors constructed through representation theorems (13; 1).

Our construction can be viewed as an extension of the ERRW defined on an infinite space. The prediction for the next state visited by the process is not solely a function of the number of transitions observed in and out of the last state. In effect, transition probabilities out of different states share statistical strength. This will become relevant in applications where many states are observed, especially for those states that occur rarely.

A major goal in our application is the prediction of the number of states that the Markov chain has not yet visited that will appear in the next  $m$  transitions. More generally, scientists are interested in predicting aspects of the protein dynamics that may be strongly correlated with the rate of discovery of unobserved states, for instance, the variability of the time needed to reach a conformation of interest  $y$ , starting from a specific state  $x$ . Predictive distributions for such attributes are useful in deciding whether one should continue a costly experiment to obtain substantial additional information on a kinetic property of interest.

Estimating the probability of discovering new species is a long-standing problem in statistics (5). Most contributions in the literature assume that observations, for example species of fish captured in a lake, can be modeled as independent and identically distributed random variables with an unknown discrete distribution. In this setting, several Bayesian nonparametric models have been studied (16; 17; 11). Here we assume that species, in our case protein conformational states, are sampled from a reversible Markov chain. To the best of our knowledge, this is the first Bayesian analysis of species sampling in this setting.

FIG 1. Diagram of the  $(\theta, \alpha, \beta)$  scheme and special cases.

We can now outline the article. Section 2 introduces the species sampling model, which we call the  $(\theta, \alpha, \beta)$  scheme. The process specializes to the ERRW, a Markov exchangeable scheme, and to the two-parameter Hoppe urn, a classical exchangeable scheme which gives rise to the Pitman-Yor process and the two-parameter Poisson-Dirichlet distribution (20; 21). As illustrated in Figure 1, the parameter  $\beta$  smoothly tunes the model between these two special cases. Section 3 shows that the  $(\theta, \alpha, \beta)$  scheme can be represented as a mixture of reversible Markov chains. This allows us to use its de Finetti measure as a prior for Bayesian analysis. Section 4 shows that our scheme is a projection of a conjugate prior for a random walk on a multigraph. This representation is then used to prove that our model has full weak support. Section 5 provides a sufficientness characterization of the proposed scheme. This result is strictly related to the characterizations of the ERRW and the two-parameter Hoppe urn discussed in (23) and (28) respectively. In Section 6, an expression for the law of the  $(\theta, \alpha, \beta)$  scheme is derived and this result is used in Section 7 to define algorithms for posterior simulation. Section 8 applies our model to the analysis of two molecular dynamics datasets. We evaluate the predictive performance of the model by splitting the data into training and validating datasets. Section 9 concludes with a discussion of remaining challenges.




FIG 2. Three kinds of transition in the  $(\theta, \alpha, \beta)$  scheme. The blue arrow represents the transition between two states in  $\mathcal{X}$ , while the red arrows represent the path of an auxiliary random walk with reinforcement. The edges that have positive weight  $g$  before the transition are drawn in black, and in each case, we mark the reinforcements of  $g$  produced by the transition. Self-transitions follow a slightly different reinforcement scheme formalized in Definition 2.1.

**2. The  $(\theta, \alpha, \beta)$  scheme.** The  $(\theta, \alpha, \beta)$  scheme is a stochastic process  $(X_i)_{i \in \mathbb{N}}$  on a Polish measurable space  $(\mathcal{X}, \mathcal{F})$  equipped with a diffuse (i.e. without point masses) probability measure  $\mu$ . We construct the law of the process using an auxiliary random walk with reinforcement on the extended space  $\mathcal{X}_+ := \mathcal{X} \cup \{\zeta\}$ . The auxiliary process classifies each transition  $X_i \rightarrow X_{i+1}$  into three categories listed in Figure 2 and defines latent variables  $(U_i)_{i \in \mathbb{N}}$ , taking values in  $\{a, b, c\}$ , that capture each transitions' category. In this section we first provide a formal definition of the  $(\theta, \alpha, \beta)$  scheme and then briefly describe the latent process.

The law of the  $(\theta, \alpha, \beta)$  scheme is specified by a weighted undirected graph  $g$  with vertices in  $\mathcal{X}_+$ . This graph can be formalized as a symmetric function  $g : \mathcal{X}_+ \times \mathcal{X}_+ \rightarrow [0, \infty)$ , where  $g(x, y)$  is the weight of an undirected edge with vertices  $x$  and  $y$ . We require that the set  $\mathcal{S} := \{x \in \mathcal{X}; g(\zeta, x) > 0\}$  is countable,  $\sum_{x \in \mathcal{S}} g(\zeta, x) < \infty$ , and that the set of edges  $\{(x, y) \in \mathcal{X}^2; g(x, y) > 0\}$  is a finite subset of  $\mathcal{S}^2$ . The graph will be sequentially reinforced after each transition of the  $(\theta, \alpha, \beta)$  scheme. In the following definition, we assume the initial state  $X_1$  is deterministic and contained in  $\mathcal{S}$ .

**DEFINITION 2.1.** *The  $(\theta, \alpha, \beta)$  scheme,  $(X_i)_{i \in \mathbb{N}}$ , has parameters  $\theta \geq 0$ ,  $\alpha \in [0, 1)$ , and  $\beta \in [0, 1]$ . The parameter  $\theta$  is equal to the initial weight*

$g(\zeta, \zeta)$ . Suppose we have sampled  $X_1, \dots, X_i$ , where  $i \geq 1$ . Then, given  $X_i$  and the reinforced graph  $g$ , we sample a

(a) **ERRW-like transition** to  $X_{i+1}$  with probability

$$\frac{g(X_i, X_{i+1})}{\sum_{x \in \mathcal{X}_+} g(X_i, x)},$$

and make the following reinforcement:

$$g(X_i, X_{i+1}) \rightarrow g(X_i, X_{i+1}) + 1 + \mathbf{1}_{X_i=X_{i+1}};$$

(b) **mediated transition without discovery** to  $X_{i+1}$  with probability

$$\frac{g(X_i, \zeta)}{\sum_{x \in \mathcal{X}_+} g(X_i, x)} \times \frac{g(\zeta, X_{i+1}) + \beta \times \mathbf{1}_{X_i=X_{i+1}}}{\beta + \sum_{x \in \mathcal{X}_+} g(\zeta, x)},$$

and make the following reinforcements:

$$g(X_i, X_{i+1}) \rightarrow g(X_i, X_{i+1}) + (1 - \beta) \times (1 + \mathbf{1}_{X_i=X_{i+1}}),$$

$$g(X_i, \zeta) \rightarrow g(X_i, \zeta) + \beta,$$

$$g(X_{i+1}, \zeta) \rightarrow g(X_{i+1}, \zeta) + \beta;$$

(c) **mediated transition with discovery** to a new state  $X_{i+1} \sim \mu$  with probability

$$\frac{g(X_i, \zeta)}{\sum_{x \in \mathcal{X}_+} g(X_i, x)} \times \frac{g(\zeta, \zeta)}{\beta + \sum_{x \in \mathcal{X}_+} g(\zeta, x)},$$

and make the following reinforcements:

$$g(X_i, X_{i+1}) \rightarrow g(X_i, X_{i+1}) + (1 - \beta),$$

$$g(X_i, \zeta) \rightarrow g(X_i, \zeta) + \beta,$$

$$g(X_{i+1}, \zeta) \rightarrow g(X_{i+1}, \zeta) + (1 - \alpha)\beta,$$

$$g(\zeta, \zeta) \rightarrow g(\zeta, \zeta) + \alpha\beta.$$

In several applications one may prefer to set the initial  $g$  to zero everywhere except for  $g(\zeta, \zeta) = \theta$  and  $g(X_1, \zeta)$ , which can be made infinitesimally small. This reduces difficulties associated with the model specification and does not affect the main properties of the model discussed in this article. In some cases, we will relax the assumption that  $X_1$  is deterministic, by specifying a distribution, say  $\mu$ , for  $X_1$  and choosing a positive value for  $g(X_1, \zeta)$ . In any case, the conditional distributions  $p(X_n | X_1, \dots, X_{n-1})$  are dictated by the reinforced scheme in Definition 2.1.

We can now describe the latent reinforced process in order to simplify the interpretation of the  $(\theta, \alpha, \beta)$  scheme. To sample a transition  $X_i \rightarrow X_{i+1}$ , we first take one step in the auxiliary random walk from  $X_i$ . If we land on a state  $x \in \mathcal{X}$  (panel *a* in Figure 2), we set  $X_{i+1} = x$  and  $U_i = a$ . If we land on  $\zeta$ , we sample another step of the random walk from  $\zeta$ . Once more, if we land on some  $x' \in \mathcal{X}$  (panel *b* in Figure 2), we set  $X_{i+1} = x'$  and  $U_i = b$ . Otherwise (panel *c* in Figure 2), we sample a new state  $X_{i+1}$  from  $\mu$  and set  $U_i = c$ .

REMARK 2.1. *Assume the initial graph  $g$  is null everywhere except for  $g(\zeta, \zeta) = \theta$  and  $g(\zeta, X_1) = 1$ . If  $\beta = 1$ ,  $\alpha = 0$  and  $X_1 \sim \mu$ , then the process  $(X_i)_{i \in \mathbb{N}}$  is a Blackwell-MacQueen urn (4) with base distribution  $\mu$  and concentration parameter  $\theta/2$ . In different words, the process is exchangeable and its directing random measure is the Dirichlet process (12).*

REMARK 2.2. *Under the assumptions in Remark 2.1, by setting  $\alpha > 0$  and  $g(\zeta, X_1) = 1 - \alpha$ , the process  $(X_i)_{i \in \mathbb{N}}$  is equal to an urn scheme introduced by Engen (10), known in the species sampling context as the two-parameter Hoppe urn (see Appendix B (2)). This exchangeable process has been studied extensively by Pitman and Yor (20; 21). Its directing random measure is the Pitman-Yor process (14) with base distribution  $\mu$ , concentration parameter  $(\theta - \alpha)/2$ , and discount parameter  $\alpha/2$ ; the sorted masses of this random measure have the two-parameter Poisson-Dirichlet distribution. Note that the discount parameter of the Pitman-Yor process can be chosen from the unit interval  $[0, 1)$ , while in our construction  $\alpha/2 < 0.5$ .*

REMARK 2.3. *When  $\beta = \theta = 0$ , edges connected to  $\zeta$  cannot be reinforced, and the  $(\theta, \alpha, \beta)$  scheme specializes to the ERRW on  $\mathcal{S}$ . See (7).*

REMARK 2.4. *Definition 2.1 brings to mind the two-parameter hierarchical Dirichlet Process hidden Markov model (HDP-HMM) (25) and its associated species sampling scheme, the two-parameter Chinese Restaurant Franchise. This process has been used for Bayesian modeling of Markov chains on infinite spaces. The predictive distribution can be viewed as a  $(\theta, \alpha, \beta)$  scheme in which the underlying infinite graph has directed edges. However, the  $(\theta, \alpha, \beta)$  scheme is not a special case of this model and has no equivalent hierarchical construction. This connection is explained in more detail in Appendix B (2).*

The influence of each parameter in the  $(\theta, \alpha, \beta)$  scheme can be described as follows. The parameter  $\beta$  determines the Markov character of the model;

as it approaches 1, the process becomes exchangeable. The parameter  $\theta$  is related to the concentration parameter of the two-parameter Hoppe urn, which controls the mode of the number of states visited in a given number of steps. The parameter  $\alpha$  is related to the discount factor in the two-parameter Hoppe urn, which controls the distribution of frequencies of different states. It is worth noting that  $\beta$  also controls the number of states visited, which increases markedly as  $\beta$  is made larger.

The recurrence of the ERRW on infinite graphs is far from trivial, especially for locally connected graphs (18, and references therein). However, it is not difficult to prove that the  $(\theta, \alpha, \beta)$  scheme a.s. returns infinitely often to all visited states, the state  $\zeta$  is visited infinitely often a.s. and, if  $\theta > 0$ , the edge  $g(\zeta, \zeta)$  is crossed infinitely often. This notion of recurrence is stated in the next proposition.

**PROPOSITION 2.1.** *The  $(\theta, \alpha, \beta)$  scheme is recurrent, i.e. the event  $\{\sum_{j>i} \mathbf{1}_{X_i=X_j} > 0\}$  has probability 1 for every integer  $i$ . When  $\theta > 0$  or when the set  $\mathcal{S}$  is infinite, the number of states visited is infinite almost surely.*

**3. de Finetti representation of the  $(\theta, \alpha, \beta)$  scheme.** Diaconis and Freedman defined a special notion of partial exchangeability to prove a version of de Finetti’s theorem for Markov chains (8).

**DEFINITION 3.1.** *A stochastic process on a countable space  $\mathcal{X}$  is Markov exchangeable if the probability of observing a path  $x_1, \dots, x_n$  is only a function of  $x_1$  and the transition counts  $C(x, y) := |\{x_i = x, x_{i+1} = y; 1 \leq i < n\}|$  for all  $x, y \in \mathcal{X}$ .*

**THEOREM 3.1 (Diaconis and Freedman).** *A process is Markov exchangeable and returns to every state visited infinitely often, if and only if it is a mixture of recurrent Markov chains.*

The  $(\theta, \alpha, \beta)$  scheme takes values in an uncountable space  $\mathcal{X}$ , which precludes a direct application of Theorem 3.1. We will state a more general notion of Markov exchangeability and use it to prove a de Finetti style representation for  $(X_i)_{i \in \mathbb{N}}$ . We use the notion of  $x$ -block defined in Diaconis and Freedman (8); given a recurrent trajectory  $x_1, x_2, \dots$ , the  $i$ -th  $x$ -block, for any state  $x$  appearing in  $x_1, x_2, \dots$ , is the finite subsequence that starts with the  $i$ -th occurrence of  $x$  and ends before the  $(i + 1)$ -th occurrence.

The process  $(X_i)_{i \in \mathbb{N}}$  visits every state in  $\mathcal{S}$  infinitely often; in addition, it will discover new species in  $\mathcal{X}$  in steps of the third kind in Figure 2. But the

new species are sampled independently from  $\mu$ , which motivates expressing  $(X_i)_{i \in \mathbb{N}}$  as a function of two independent processes on the same probability space:  $(Z_i)_{i \in \mathbb{N}}$  which represents the sequence where new species are labeled in order of appearance, and  $(T_i)_{i \in \mathbb{N}}$  which represents the  $\mathcal{X}$ -valued locations of each species. These are formally defined in the sequel.

Let  $\mathcal{Z} = \mathcal{X} \sqcup \mathbb{N}$  and let  $d$  be a function that maps a sequence in  $\mathcal{Z} \supset \mathcal{S}$  to a sequence in the disjoint union  $\mathcal{S} \sqcup \mathbb{N}$ . Each element of the sequence in  $\mathcal{S}$  is mapped to itself, and those not in  $\mathcal{S}$  are mapped to the order in which they appear in the sequence. Hence, the range of  $d$  consists of sequences where every state  $j \in \mathbb{N}$  may only appear after all states  $1, 2, \dots, j-1$  have appeared at least once. For example, if  $\mathcal{X}$  is the unit interval and  $\mathcal{S} = \{0.1, 0.2, 0.3\}$ , then

$$d : (0.1, 7, 4, 0.3, 7, 6, 4, 4) \mapsto (0.1, 1, 2, 0.3, 1, 3, 2, 2).$$

Define  $(Z_i)_{i \in \mathbb{N}} := d((X_i)_{i \in \mathbb{N}})$ , and let  $(T_i)_{i \in \mathbb{N}}$  be a sequence of independent random variables from  $\mu$ , with  $(T_i)_{i \in \mathbb{N}}$  independent from  $(Z_i)_{i \in \mathbb{N}}$ . Then,

$$(X_i)_{i \in \mathbb{N}} \stackrel{d}{=} (\bar{X}_i)_{i \in \mathbb{N}}, \quad \text{where we define} \quad \bar{X}_i := \begin{cases} Z_i & \text{if } Z_i \in \mathcal{S}, \\ T_{Z_i} & \text{otherwise.} \end{cases}$$

**PROPOSITION 3.1.** *Take two sequences  $x_1, \dots, x_n$  and  $x'_1, \dots, x'_n$  in  $\mathcal{S} \sqcup \mathbb{N}$  that are fixed points of  $d$ . Suppose one can map  $x_1, \dots, x_n$  to  $x'_1, \dots, x'_n$  by a transposition of two blocks in  $x_1, \dots, x_n$  which both begin in  $x \in \mathcal{S} \sqcup \mathbb{N}$  and end in  $y \in \mathcal{S} \sqcup \mathbb{N}$ , followed by an application of the mapping  $d$ . Then,*

$$p(Z_1 = x_1, \dots, Z_n = x_n) = p(Z_1 = x'_1, \dots, Z_n = x'_n).$$

**EXAMPLE 3.1.** *Assume again  $\mathcal{S} = \{0.1, 0.2, 0.3\}$ . If we set*

$$\begin{aligned} (x_1, \dots, x_n) &= (0.1, 1, 2, 0.3, 3, 2, 4, 0.3), \quad \text{and} \\ (x'_1, \dots, x'_n) &= (0.1, 1, 2, 3, 0.3, 4, 2, 0.3), \end{aligned}$$

*then, by transposing two blocks in  $(x_1, \dots, x_n)$  that start from 2 and finish in 0.3, we obtain the vector  $(0.1, 1, 2, 4, 0.3, 3, 2, 0.3)$ . Moreover,*

$$d(0.1, 1, 2, 4, 0.3, 3, 2, 0.3) = (0.1, 1, 2, 3, 0.3, 4, 2, 0.3) = (x'_1, \dots, x'_n).$$

*Proposition 3.1 then implies that the two probabilities,  $p(Z_1 = x_1, \dots, Z_n = x_n)$  and  $p(Z_1 = x'_1, \dots, Z_n = x'_n)$ , are identical.*

REMARK 3.1. *Note that if the process only visits states in  $\mathcal{S}$ , as is the case when  $\theta = 0$ , the statement of Proposition 3.1 is equivalent to Markov exchangeability (cf. Proposition 27 in 8). This fact, together with Proposition 2.1 is enough to show by a straightforward application of Theorem 3.1 that the  $(\theta, \alpha, \beta)$  scheme with  $\theta = 0$  is a mixture of recurrent Markov chains on  $\mathcal{S}$ .*

Equipped with this notion of Markov exchangeability for the species sampling sequence  $(Z_i)_{i \in \mathbb{N}}$ , we show that  $(X_i)_{i \in \mathbb{N}}$  can be represented as a mixture of Markov chains.

PROPOSITION 3.2. *There exists a mixture of Markov chains  $(W_i)_{i \in \mathbb{N}}$ , taking values in  $\mathcal{S} \sqcup \mathbb{N}$ , such that, if we define*

$$\tilde{X}_i := \begin{cases} W_i & \text{if } W_i \in \mathcal{S}, \\ T_{W_i} & \text{otherwise,} \end{cases}$$

where  $(W_i)_{i \in \mathbb{N}}$  and  $(T_i)_{i \in \mathbb{N}} \sim \mu^{\mathbb{N}}$  are independent, then  $(\tilde{X}_i)_{i \in \mathbb{N}} \stackrel{d}{=} (X_i)_{i \in \mathbb{N}}$ . That is, for some measure  $\phi$  on  $(\mathcal{S} \sqcup \mathbb{N}) \times \mathcal{P}$ , where  $\mathcal{P}$  is the space of stochastic matrices on  $\mathcal{S} \sqcup \mathbb{N}$ , the distribution of  $(W_i)_{i \in \mathbb{N}}$ , can be represented as

$$p(W_1 = w_1, \dots, W_n = w_n) = \int_{\mathcal{P}} \prod_{i=1}^{n-1} P(w_i, w_{i+1}) \phi(w_1, dP).$$

Let  $\mathcal{P}_r \subset \mathcal{P}$  be the set of transition probability matrices for recurrent reversible Markov chains. We can show that this set has probability 1 under the de Finetti measure.

PROPOSITION 3.3.  $\phi((\mathcal{S} \sqcup \mathbb{N}) \times \mathcal{P}_r) = 1$ .

**4. The  $(\theta, \alpha, \beta)$  scheme with colors.** This section shows that the  $(\theta, \alpha, \beta)$  scheme can be interpreted as a Bayesian conjugate model for a random walk on a multigraph. This representation is used for showing that the  $(\theta, \alpha, \beta)$  scheme has large support, in a sense that will be made precise in the sequel. We also make a connection between the de Finetti measure of the ERRW on a finite graph and our model.

We start by defining a colored random walk on a weighted multigraph  $\tilde{g}$ . The vertices of the graph take values in  $\mathcal{X}$ , and we now allow there to be more than one edge between every pair of vertices. Every edge is associated to a distinct color in a set  $\mathcal{C}$ . We assign a weight  $\tilde{g}(\{x, y\}, c)$  to the edge connecting  $x$  and  $y$  with color  $c$ , requiring that  $\tilde{g}(x) := \sum_{y, c} \tilde{g}(\{x, y\}, c) < \infty$

for all  $x \in \mathcal{X}$ . A random walk on this graph is a process that starts from  $x_1 \in \mathcal{X}$ , and after arriving at some state  $x$ , traverses the edge  $(\{x, z\}, c)$  with probability  $\tilde{g}(\{x, z\}, c)/\tilde{g}(x)$ . Let  $p_{\tilde{g}}$  be the law of this process.

A Bayesian statistician observes a finite sequence of traversed colored edges and wants to predict the future trajectory of the colored random walk. We suggest how to use the  $(\theta, \alpha, \beta)$  scheme in this context. Informally, in a ERRW-like transition, we reinforce a single edge of a specific color, while in a mediated transition, we draw a new edge with a novel color.

The Bayesian model is a random sequence of colored edges  $(E_i)_{i \in \mathbb{N}}$ . We use  $C_i$  and  $\{X_i, X_{i+1}\}$  to denote the color and vertices of  $E_i$ . Let  $\mu$  and  $\gamma$  be non-atomic distributions over  $\mathcal{X}$  and  $\mathcal{C}$  respectively, and specify  $\theta$ ,  $\alpha$  and  $\beta$  as in the previous sections. Let  $X_1 = x_1$ ,  $C_1 \sim \gamma$  and

$$X_2 \mid X_1, C_1 \sim \frac{\beta(1-\alpha)}{\beta+\theta} \delta_{X_1} + \frac{\theta+\alpha\beta}{\beta+\theta} \mu.$$

The distribution of  $(X_1, X_2)$  corresponds to an initial graph with only 2 edges, with endpoints  $\{x_1, \zeta\}$  and  $\{\zeta, \zeta\}$ , weighted by  $-\alpha\beta$  and  $\theta + \alpha\beta$ , respectively. If the initial weighted graph  $g$  in a  $(\theta, \alpha, \beta)$  scheme is chosen as above, the reinforcement rules in Definition 2.1 produce a well-defined process even if the initial value of  $g(x_1, \zeta)$  is negative. After the first transition all edges will have nonnegative weights. This choice for the initial weighted graph will be used in the present section and Section 5.

After a path  $(E_1, \dots, E_n)$ , the probability of recrossing an edge  $E_j$ , with  $j \leq n$  and  $X_{n+1} \in \{X_j, X_{j+1}\}$ , is

$$p(E_{n+1} = E_j \mid E_1, \dots, E_n) = \frac{G_n(\{X_j, X_{j+1}\}, C_j)}{W_{X_{n+1}, n}}, \quad \text{where}$$

$$G_n(\{x, y\}, c) = \max\left(0, -\beta + \sum_{i \leq n} \mathbf{1}_{E_i = (\{x, y\}, c)}\right) 2^{\mathbf{1}_{x=y}},$$

$$W_{x, n} = -\alpha\beta + \sum_{i \leq n} (\mathbf{1}_{x \in \{X_i, X_{i+1}\}}) 2^{\mathbf{1}_{X_i = X_{i+1}}}.$$

In words, the probability of recrossing an edge is linear in the number of crossings. Let  $\mathcal{C}_n$  be the set of distinct colors in  $(C_1, \dots, C_n)$ . The conditional probability that  $C_{n+1} \notin \mathcal{C}_n$ ,

$$p(C_{n+1} \notin \mathcal{C}_n \mid E_1, \dots, E_n) = \frac{B_{X_{n+1}, n}}{W_{X_{n+1}, n}}, \quad \text{where}$$

$$B_{x, n} = -\alpha\beta + \beta \sum_{i \leq n} (\mathbf{1}_{x \in \{X_i, X_{i+1}\}}) (\mathbf{1}_{C_i \notin \mathcal{C}_{i-1}}) 2^{\mathbf{1}_{X_i = X_{i+1}}},$$

is linear in the number of distinct colored edges adjacent to  $X_{n+1}$ . The probability that  $X_{n+2} = y$ , for any vertex  $y \in \{X_1, \dots, X_n\}$ , conditional on  $C_{n+1} \notin \mathcal{C}_n$ , is

$$(4.1) \quad p(X_{n+2} = y \mid E_1, \dots, E_n, C_{n+1} \notin \mathcal{C}_n) = \frac{\beta \mathbf{1}_{X_{n+1}=y} + B_{y,n}}{2\beta \times |\mathcal{C}_n| + \theta + \beta},$$

which depends linearly on the number of distinct colored edges adjacent to  $y$ . Finally,

$$\begin{aligned} X_{n+2} \mid C_{n+1} \notin \mathcal{C}_n, X_{n+2} \notin \{X_1, \dots, X_{n+1}\}, E_1, \dots, E_n &\sim \mu, \quad \text{and} \\ C_{n+1} \mid X_{n+2}, E_1, \dots, E_n, C_{n+1} \notin \mathcal{C}_n &\sim \gamma. \end{aligned}$$

The following property is a direct consequence of this definition.

**PROPOSITION 4.1.** *The sequence  $(X_i)_{i \in \mathbb{N}}$  of  $\mathcal{X}$ -valued states visited by  $(E_i)_{i \in \mathbb{N}}$  is identical in distribution to a  $(\theta, \alpha, \beta)$  scheme initiated at  $x_1$ , with  $g$  everywhere null except at  $g(x_1, \zeta) = -\beta\alpha$  and  $g(\zeta, \zeta) = \theta + \alpha\beta$ .*

Furthermore, given a  $(\theta, \alpha, \beta)$  scheme with an arbitrary initial graph  $g$ , it is possible to construct a colored  $(\theta, \alpha, \beta)$  scheme with a closed-form predictive distribution such that the equality stated in Proposition 4.1 holds. This would require changing the definition of  $(E_i)_{i \in \mathbb{N}}$  in a way that preserves the reinforcement scheme.

**PROPOSITION 4.2.** *Let  $e_1, \dots, e_n$  be a colored path, with  $e_i = (\{x_i, x_{i+1}\}, c_i)$ , and let  $\Lambda(e_1, \dots, e_n)$  be the probability of the event*

$$\bigcap_{e_i=e_j} \{E_i = E_j\} \bigcap_{e_i \neq e_j} \{E_i \neq E_j\} \bigcap_{x_i=x_j} \{X_i = X_j\} \bigcap_{x_i \neq x_j} \{X_i \neq X_j\}.$$

*Suppose  $e_{\sigma(1)}, \dots, e_{\sigma(n)}$ , for some permutation  $\sigma$ , is also a colored path starting at  $x_1$ . Then,*

$$(4.2) \quad \Lambda(e_1, \dots, e_n) = \Lambda(e_{\sigma_1}, \dots, e_{\sigma_n}).$$

This result, related to Proposition 3.1, establishes a probabilistic symmetry between paths that can be mapped to each other by permuting the order of edges crossed, and applying certain automorphisms to  $\mathcal{X}$  and  $\mathcal{C}$ . Proposition 4.2 gives rise to the following de Finetti representation.

**PROPOSITION 4.3.** *There exists a mixture of colored random walks on weighted multigraphs  $(\{W_i, W_{i+1}\}, J_i)_{i \in \mathbb{N}}$ , with  $W_i \in \mathbb{N}$  and  $J_i \in \mathbb{N}$ , and*

independent processes  $(T_i)_{i \in \mathbb{N}} \sim \mu^{\mathbb{N}}$  and  $(V_i)_{i \in \mathbb{N}} \sim \gamma^{\mathbb{N}}$ , such that the sequence of colored edges  $(\{\tilde{X}_i, \tilde{X}_{i+1}\}, \tilde{C}_i)_{i \in \mathbb{N}}$ , defined by

$$(\tilde{X}_i, \tilde{C}_i) = \begin{cases} (x_1, V_{J_i}) & \text{if } W_i = 1, \\ (T_{W_i}, V_{J_i}) & \text{if } W_i > 1, \end{cases}$$

is identical in distribution to  $(E_i)_{i \in \mathbb{N}}$ .

The proof of Proposition 4.3 constructs the discrete process  $(W_i, J_i)_{i \in \mathbb{N}}$ , which will now be used to show that the  $(\theta, \alpha, \beta)$  scheme has large support. The law of  $(W_i, J_i)_{i \in \mathbb{N}}$  can be mapped bijectively to the exchangeable law of the sequence of  $x_1$ -blocks in the process, which in the present section and the next are defined as sequences of labeled edges. Therefore, the de Finetti measure of  $(W_i, J_i)_{i \in \mathbb{N}}$  uniquely identifies the de Finetti measure of the  $x_1$ -blocks, and vice versa. We denote the random distribution of the  $x_1$ -blocks  $\eta$ . Note that the space of  $x_1$ -blocks is discrete because  $W_i$  and  $J_i$  are integer-valued.

We show that the law of  $\eta$  has full weak support. On the basis of Proposition 4.1, we then conclude that the  $x_1$ -block de Finetti measure induced by the  $(\theta, \alpha, \beta)$  scheme also has full weak support. The next proposition is proven for  $(E_i)_{i \in \mathbb{N}}$  as defined in this section; the result can be extended to any analogous reinforced process corresponding to a specific  $(\theta, \alpha, \beta)$  scheme.

**PROPOSITION 4.4.** *Let  $\eta^\circ$  be the  $x_1$ -block distribution induced by a colored random walk on an arbitrary weighted multigraph with vertices and colors in  $\mathbb{N}$ , in which the sum of weights is finite,  $\sum_{x \in \mathbb{N}} \tilde{g}(x) < \infty$ . For every  $\epsilon > 0$ ,  $m \geq 1$  and any collection of bounded real functions  $f_1, \dots, f_m$  on the space of  $x_1$ -blocks,*

$$p(\eta \in U_{\epsilon, f_1, \dots, f_m}(\eta^\circ)) > 0, \quad \text{where}$$

$$U_{\epsilon, f_1, \dots, f_m}(\eta^\circ) = \left\{ \eta' : \left| \int f_i d\eta^\circ - \int f_i d\eta' \right| < \epsilon, \quad i = 1, \dots, m \right\}.$$

The process  $(E_i)_{i \in \mathbb{N}}$  also reveals a connection between the  $(\theta, \alpha, \beta)$  scheme and the ERRW. Recall that the colored  $x_1$ -blocks  $(H_i)_{i \in \mathbb{N}}$  in the Bayesian model are exchangeable and, by de Finetti's theorem, conditionally independent. Consider their posterior distribution given the subsequence  $(E_i)_{i \leq n}$ ,

and assume it has  $X_{n+1} = X_1$  and includes  $k$  colored  $x_1$ -blocks. These assumptions are only made to simplify the exposition. The limits

$$(4.3) \quad T_{(\{x,y\},c)} = \lim_{m \rightarrow \infty} \left( \frac{\sum_{i=n}^{n+m} \mathbf{1}_{(\{X_i, X_{i+1}\}, C_i) = (\{x,y\}, c)}}{\sqrt{\sum_{i=n}^{n+m} \mathbf{1}_{X_i=x}} \sqrt{\sum_{i=n}^{n+m} \mathbf{1}_{X_i=y}}} \right)$$

for every edge  $(\{x, y\}, c) \in \{E_1, \dots, E_n\}$  are functions of the directing random measure for the sequence of  $x_1$ -blocks. From these limits, one can obtain the probability in the directing random measure of any  $x_1$ -block formed with edges in  $\{E_1, \dots, E_n\}$ . Namely, given  $E_1, \dots, E_n$  and the tail  $\sigma$ -field of  $(E_i)_{i \in \mathbb{N}}$ , the probability of an  $x_1$ -block  $e_1, \dots, e_k$  is the product  $\prod_{i=1}^k T_{e_i}$ . We can now state the connection with the ERRW.

**PROPOSITION 4.5.** *There exists an ERRW on a multigraph with a finite number of edges, such that the joint posterior distribution of the random variables in Eq. 4.3, given  $E_1, \dots, E_n$ , is identical to the distribution of the same limits in the ERRW.*

Appendix C (2) contains a constructive proof of this proposition, in which one such ERRW, whose parameters depend on  $E_1, \dots, E_n$ , is defined.

**5. Sufficientness characterization.** This section provides a characterization of the colored  $(\theta, \alpha, \beta)$  scheme in terms of certain predictive sufficiencies or *sufficientness* conditions. The first characterization of this type, for the Pólya urn, was proven in an influential paper by W. E. Johnson (27) and has since been extended to other predictive schemes for discrete sequences such as the two-parameter Hoppe urn (28) and the edge-reinforced random walk (23). Our result is closely related to the work of these authors and uses similar proof techniques. In addition to its clear subjective motivation, the characterization elucidates connections between the  $(\theta, \alpha, \beta)$  scheme and other popular nonparametric Markov models (13).

Consider a random sequence of colored edges  $(\tilde{E}_i)_{i \in \mathbb{N}}$ , where  $\tilde{E}_i = (\{X_i, X_{i-1}\}, C_i)$ . Each color in  $(\tilde{E}_i)_{i \in \mathbb{N}}$  identifies an edge. We assume  $(\tilde{E}_i)_{i \in \mathbb{N}}$  is a mixture of random walks on weighted and colored multigraphs, in the sense of Proposition 4.3, which visits more than 2 vertices with probability 1. It will be shown that, if the predictive distribution of  $(\tilde{E}_i)_{i \in \mathbb{N}}$  satisfies certain conditions, then the process is a  $(\theta, \alpha, \beta)$  scheme.

Consider a path  $Z_n = (\tilde{E}_1, \dots, \tilde{E}_n)$  and define,

$$(i) \quad \kappa(Z_n) = |\{0 \leq i \leq n; X_i = X_n\}| + \frac{\mathbf{1}_{X_n \neq X_0}}{2},$$

- (ii)  $\kappa(e, Z_n) = |\{1 \leq i \leq n; \tilde{E}_i = e\}|$ , and
- (iii)  $\tau(Z_n) = \sum_v \eta(v, Z_n)/2$ , where,
- (iv)  $\eta(v, Z_n) = \sum_{i \leq n} (\mathbf{1}_{v \in \{X_i, X_{i+1}\}}) (\mathbf{1}_{C_i \notin \{C_1, \dots, C_{i-1}\}}) 2^{\mathbf{1}_{X_i = X_{i+1}}}$ .

These variables describe (i) how many times  $X_n$  has been visited and whether it coincides with  $X_0$ , (ii) how many times an edge  $e$  has been traversed, (iii) the number of observed colors and (iv) the degree of  $v \in \mathcal{X}$  in the multigraph constructed by all distinct colored edges in  $Z_n$ . In summary, they are easily interpretable. We also use  $\rho(Z_n)$  to denote the number of distinct  $\mathcal{X}$ -valued states in  $Z_n$ , and the indicator  $s(e)$ , which is equal to 1 if  $e$  is a loop and 0 otherwise.

We can now define sufficientness conditions for  $(\tilde{E}_i)_{i \in \mathbb{N}}$ . The process satisfies Condition 1 if there exist functions  $b_0$  and  $b_1$  such that for every  $e \in \{\tilde{E}_1, \dots, \tilde{E}_n\}$  incident on  $X_n$ ,

$$(5.1) \quad p(\tilde{E}_{i+1} = e | Z_n) = b_{s(e)}(\kappa(Z_n), \kappa(e, Z_n)) \in (0, 1).$$

In words, the probability of making a transition through an edge  $e = (\{X_n, v\}, c)$  in  $Z_n$  depends on the number of times the edge has been crossed and the number of visits to  $X_n$ . The process satisfies Condition 2 if there is a function  $g$  such that

$$(5.2) \quad p(\tilde{E}_{n+1} \notin \{\tilde{E}_1, \dots, \tilde{E}_n\} | Z_n) = g(\kappa(Z_n), \eta(X_n, Z_n)) \in (0, 1).$$

That is, the probability of a transition through a new edge is a function of the number of observed edges  $\eta(X_n, Z_n)$  incident on  $X_n$ , and the number of visits to  $X_n$ . Condition 3 requires that some function  $h$  satisfies, for every  $v \in \{X_0, \dots, X_n\}$ ,

$$(5.3) \quad \begin{aligned} p(X_{i+1} = v | Z_n, \tilde{E}_{n+1} \notin \{\tilde{E}_1, \dots, \tilde{E}_n\}) \\ = h(\tau(Z_n), \eta(v, Z_n) + \mathbf{1}_{X_n = v}) \in (0, 1). \end{aligned}$$

If a new edge will be traversed, then the conditional probability that the path will go to an already seen vertex depends solely on the number of edges out of said vertex and the overall number of observed edges. Finally, the process satisfies Condition 4 if there is a function  $q$ , such that

$$p(X_{i+1} \notin \{X_0, \dots, X_n\} | Z_n, \tilde{E}_{n+1} \notin \{\tilde{E}_1, \dots, \tilde{E}_n\})$$

$$(5.4) \quad = q(\tau(Z_n), \rho(Z_n)) \in (0, 1);$$

i.e. the conditional probability that the path will go to an unseen vertex is a function of the total number of edges and vertices.

The main result of this section can be divided into two lemmas, the first of which depends only on 3 of the conditions above.

LEMMA 5.1. *If the process  $(\tilde{E}_i)_{i \in \mathbb{N}}$  satisfies Conditions 1, 2, and 3, there exist  $\beta \in [0, 1)$  and  $\lambda \in [-\beta, \infty)$  such that,*

$$(5.5) \quad b_s(k, j) = \frac{(1+s)(j-\beta)}{\lambda+2k-2}, \quad \text{and} \quad g(k, t) = \frac{\lambda+\beta t}{\lambda+2k-2}.$$

LEMMA 5.2. *If the process  $(\tilde{E}_i)_{i \in \mathbb{N}}$  satisfies Conditions 1, 2, 3, and 4, there exist  $\alpha \in [0, 1)$  and  $\lambda' \in [-\alpha, \infty)$  such that,*

$$(5.6) \quad h(n, j) = \frac{j-\alpha}{\lambda'+2n+1-\alpha}, \quad \text{and} \quad q(n, t) = \frac{\lambda'+\alpha(t-1)}{\lambda'+2n+1-\alpha},$$

for  $n = 1, 2, \dots$ ,  $j = 1, \dots, 2n$  and  $t = 2, \dots, n+1$ .

The characterization of the  $(\theta, \alpha, \beta)$  scheme with colors follows from the two previous lemmas.

THEOREM 5.1. *If the process  $(\tilde{E}_i)_{i \in \mathbb{N}}$  satisfies Conditions 1, 2, 3, and 4, then there exist  $\alpha \in [0, 1)$ ,  $\beta \in [0, 1)$  and  $\theta > -2\alpha\beta$ , such that the conditional transition probabilities of the process  $(\tilde{E}_i)_{i \in \mathbb{N}}$ , given any path that visits more than 2 vertices, are equal to those in the  $(\theta, \alpha, \beta)$  scheme with colors  $(E_i)_{i \in \mathbb{N}}$ .*

**6. The law of the  $(\theta, \alpha, \beta)$  scheme.** We provide an expression for the law of the species sampling sequence  $(Z_i)_{i \in \mathbb{N}}$ , defined in Section 3. Recall that the process takes values on  $\mathcal{S} \sqcup \mathbb{N}$ , and consider a fixed path  $\mathbf{z}$ .

Let  $n_{xy}$  be the number of transitions in  $\mathbf{z}$  between  $x$  and  $y$  in either direction and  $n_x = \sum_{y \in \mathcal{S} \sqcup \mathbb{N}} n_{xy}$ . For each pair  $x, y \in \mathcal{S} \sqcup \mathbb{N}$ , we introduce  $k_{xy} \leq n_{xy}$  for the number of ERRW-like transitions (panel a, Figure 2) between  $x$  and  $y$  in either direction. Let  $\ell_x := \sum_{y \in \mathcal{S} \sqcup \mathbb{N}} (n_{xy} - k_{xy}) 2^{\mathbf{1}_{x=y}}$  be the number of times that the latent path traverses the edge  $(\zeta, x)$  in either direction, let  $\ell := \sum_{x \in \mathcal{S} \sqcup \mathbb{N}} \ell_x / 2$  be the number of mediated transitions, and let  $\ell'$  be the number of times that  $(\zeta, \zeta)$  is traversed. Note that  $\ell$ ,  $\ell_x$ , and  $\ell'$  are functions of  $\mathbf{z}$  and  $\mathbf{k} = \{k_{xy}; x, y \in \mathcal{S} \sqcup \mathbb{N}\}$ . We will also need  $g(x) = \sum_{y \in \mathcal{S}} g(x, y)$ , where  $g$  is the initial weighted graph.

Given  $\mathbf{k}$  and  $\mathbf{z}$ , we know the number of transitions out of  $\zeta$  and out of  $x \in \mathcal{S} \sqcup \mathbb{N}$  in the latent path. Each transition adds a factor to the denominator

of the probability of a latent path, which increase by a fixed amount,  $2\beta$  or  $2$ , between occurrences. Similarly, given  $\mathbf{k}$ , we know the number of times that  $(\zeta, x)$  is traversed; each transition adds a factor in the numerator, and these factors are sequentially reinforced by a fixed amount  $\beta$ . Finally,  $(\zeta, \zeta)$  is traversed  $\ell'$  times, and this contributes a factor  $\theta(\theta + \alpha\beta) \cdots (\theta + [\ell' - 1]\alpha\beta)$  to the numerator of the probability of the latent path.

We can write  $p(\mathbf{z}) = \sum_{\mathbf{k}} p(\mathbf{z}, \mathbf{k})$ , where  $p(\mathbf{z}, \mathbf{k})$  is the total probability of all latent paths consistent with  $(\mathbf{z}, \mathbf{k})$ . Taking into account that the factors listed in the previous paragraph are common to all latent paths with a given  $(\mathbf{z}, \mathbf{k})$ , we obtain

$$p(\mathbf{z}, \mathbf{k}) = \frac{F(\mathbf{z}, \mathbf{k})(\theta)^{\ell' \uparrow \alpha\beta} \prod_{x: n_x > 0} \left( g(x, \zeta) + \beta(1 - \alpha)\mathbf{1}_{x \in \mathbb{N}} \right)_{\ell_x - \mathbf{1}_{x \in \mathbb{N}} \uparrow \beta}}{\left( g(\zeta) + \beta \right)_{\ell \uparrow 2\beta} \left( g(z_1) \right)_{\lfloor (n_{z_1} + 1)/2 \rfloor \uparrow 2} \prod_{\substack{x: n_x > 0 \\ x \neq z_1}} \left( g(x) + 1 - \alpha\beta\mathbf{1}_{x \in \mathbb{N}} \right)_{\lfloor n_x/2 \rfloor \uparrow 2}},$$

where we use Pitman's notation for factorial powers

$$(r)_{n \uparrow q} := r(r + q)(r + 2q) \cdots (r + (n - 1)q).$$

The function  $F(\mathbf{z}, \mathbf{k})$  is a sum with as many terms as the possible latent paths consistent with  $(\mathbf{z}, \mathbf{k})$ . The term corresponding to a specific latent path is the product of those factors that appear in the numerator of the latent path probability and correspond to ERRW-like transitions. For every pair of states  $x, y \in \mathcal{S} \sqcup \mathbb{N}$ , there are  $k_{xy}$  factors, but their sequential reinforcement depends on the order in which  $k_{xy}$  ERRW-like and  $(n_{xy} - k_{xy})$  mediated transitions appear in a specific latent path. Summing these factors over all possible orders, one pair of states at a time, we can factorize  $F(\mathbf{z}, \mathbf{k})$ ;

$$F(\mathbf{z}, \mathbf{k}) = \prod_{x, y: n_{xy} > 0} 2^{k_{xx}\mathbf{1}_{x=y}} f_{e_{xy}, \beta}(n_{xy} - \mathbf{1}_{g(x, y)=0}, k_{xy}),$$

where

$$f_{e, \beta}(n, k) = \sum_{u \in \{0, 1\}^n, \|u\|_1 = k} \prod_{j=1}^n \left( e + (1 - \beta)(j - 1) + \beta \sum_{\ell < j} u_\ell \right)^{u_j}$$

and

$$e_{xy} := \begin{cases} g(x, y) & \text{if } g(x, y) > 0, \\ 1 - \beta & \text{if } g(x, y) = 0. \end{cases}$$

PROPOSITION 6.1. *The function  $f_{e,\beta}$  satisfies the following recursion for all  $0 < k < n$ ,*

$$(6.1) \quad f_{e,\beta}(n, k) = f_{e,\beta}(n-1, k) + f_{e,\beta}(n-1, k-1)[e-1 + \beta k + (1-\beta)n],$$

where we set, for all  $n \geq 0$ ,

$$f_{e,\beta}(n, 0) = 1 \quad \text{and} \quad f_{e,\beta}(n, n) = (e)_{n \uparrow 1}.$$

The recursive representation allows one to compute  $p(\mathbf{z}, \mathbf{k})$  quickly. In order to obtain the values of  $f_{e,\beta}(n, k)$  for every  $n < \tilde{n}$ , where  $\tilde{n}$  is an arbitrarily selected integer, and  $k < n$  it is sufficient to solve 6.1 fewer than  $\tilde{n}^2$  times.

In the next proposition, we provide a closed-form solution for  $f_{e,\beta}$  in terms of the generalized Lah numbers, a well-known triangular array (6).

DEFINITION 6.1. *Let  $(t)_{n, V_0}$  be the generalized factorial of  $t$  of order  $n$  and increments  $V_0 = (v_j)_{j \geq 0}$ , namely*

$$(t)_{n, V_0} = (t - v_0)(t - v_1) \cdots (t - v_{n-1})$$

with  $(t)_{0, V_0} := 1$ . The generalized Lah numbers  $C(n, k, V_0, W_0)$  (sometimes referred to as generalized Stirling numbers), are defined by

$$(6.2) \quad C(n, k, V_0, W_0) = \begin{cases} 1 & \text{if } k = n = 0, \\ (w_0)_{n, V_0} & \text{if } n > 0 \text{ and } k = 0, \\ 0 & \text{if } k > n, \\ \sum_{j=0}^k \frac{(w_j)_{n, V_0}}{(w_j)_{j, W_0} (w_j)_{k-j, W_{j+1}}} & \text{if } 0 < k \leq n, \end{cases}$$

where  $W_i = (w_j)_{j \geq i}$ .

PROPOSITION 6.2. *For any  $n \geq 1$  and  $0 < k < n$  the function  $f_{e,\beta}$  coincides with*

$$f_{e,\beta}(n, k) = (1 - \beta)^k C(n, n - k, V_0^{(e,\beta)}, W_0^{(\beta)}),$$

where

$$(6.3) \quad V_0^{(e,\beta)} = (v_j)_{j \geq 0} := \left( -\frac{e+j}{1-\beta} \right)_{j \geq 0}$$

and

$$(6.4) \quad W_0^{(\beta)} = (w_j)_{j \geq 0} := \left( j - \frac{j}{1-\beta} \right)_{j \geq 0}.$$

**7. Posterior Simulations.** In this section we introduce a Gibbs algorithm for performing Bayesian inference with the  $(\theta, \alpha, \beta)$  scheme given the trajectory of a reversible Markov chain  $X_1, \dots, X_n$ . On the basis of the almost conjugate structure of the prior model described in the previous sections we only need to sample the latent variables  $\mathbf{k}$  conditionally on the data. Recall that the latent variables  $\mathbf{k}$  express what fraction of the transitions in  $X_1, \dots, X_n$  are ERRW-like transitions (cf. Figure 2).

We want to sample from  $p(\mathbf{k} \mid X_1, \dots, X_n)$  or equivalently from

$$p(\mathbf{k} \mid \mathbf{z}) \propto p(\mathbf{k}, \mathbf{z}).$$

Recall that  $\ell$  and  $\ell_x$  are functions of  $(\mathbf{k}, \mathbf{z})$  and that  $\mathbf{z} = d(X_1, \dots, X_n)$ . For simplicity, and without loss of generality, we consider the case where initially  $g(X_1, \zeta)$  is infinitesimal,  $g(\zeta, \zeta) = \theta$ , and  $g(\cdot, \cdot) = 0$  otherwise. The count  $\ell'$  is a function of  $\mathbf{z}$  and therefore

$$\begin{aligned} p(\mathbf{k} \mid \mathbf{z}) &\propto F(\mathbf{z}, \mathbf{k}) \frac{\prod_{x:n_x > 0} (\beta(1-\alpha)^{\mathbf{1}_{x \neq X_1}})_{\ell_x - 1 \uparrow \beta}}{(\theta + \beta)_{\ell \uparrow 2\beta}} \\ &= F(\mathbf{z}, \mathbf{k}) \frac{\prod_{x:n_x > 0} (\beta(1-\alpha)^{\mathbf{1}_{x \neq X_1}})_{\ell_x - 1 \uparrow \beta}}{(\theta + \ell'\beta)_{2\ell - \ell' \uparrow \beta}} \frac{(\theta + \ell'\beta)_{2\ell - \ell' \uparrow \beta}}{(\theta + \beta)_{\ell \uparrow 2\beta}} \\ &\propto F(\mathbf{z}, \mathbf{k}) \frac{\prod_{x:n_x > 0} (\beta(1-\alpha)^{\mathbf{1}_{x \neq X_1}})_{\ell_x - 1 \uparrow \beta}}{(\theta + \ell'\beta)_{2\ell - \ell' \uparrow \beta}} \frac{(\theta)_{2\ell \uparrow \beta}}{(\theta + \beta)_{\ell \uparrow 2\beta}} \\ &\propto F(\mathbf{z}, \mathbf{k}) \frac{\prod_{x:n_x > 0} (\beta(1-\alpha)^{\mathbf{1}_{x \neq X_1}})_{\ell_x - 1 \uparrow \beta}}{(\theta + \ell'\beta + \beta)_{2\ell - \ell' - 1 \uparrow \beta}} (\theta)_{\ell \uparrow 2\beta}. \end{aligned}$$

If

$$G \sim \text{Gamma}(\text{scale} = 1, \text{shape} = \theta/(2\beta)) \quad \text{and}$$

$$D := (D_{X_1}, D_1, D_2, \dots, D_{\ell'+1}) \sim \text{Dirichlet}\left(1, 1 - \alpha, 1 - \alpha, \dots, 1 - \alpha, \frac{\theta}{\beta} + \ell'\alpha\right),$$

then, we can write  $p(\mathbf{k} \mid \mathbf{z}) \propto \mathbb{E}_{G,D} [\psi(\mathbf{k}, G, D)]$ , where,

$$\begin{aligned} \psi(\mathbf{k}, G, D) &= \\ &\left( \prod_x \frac{1}{D_x} \right) \left( \prod_{x,y} 2^{k_{xy} \mathbf{1}_{x=y}} f_{1-\beta, \beta}(n_{xy} - 1, k_{xy}) (D_x D_y 2\beta G)^{n_{xy} - k_{xy}} \right). \end{aligned}$$

In other words, if we consider the joint distribution of three variables,  $\mathbf{k}^* = \{k_{xy}^*; x, y \in \mathcal{S} \sqcup \mathbb{N}, n_{xy} > 0\}$ ,  $D^* = (D_{X_1}^*, D_1^*, D_2^*, \dots, D_{\ell'+1}^*)$  and  $G^*$ ,

$$p(\mathbf{k}^*, D^*, G^*) \propto p_G(G^*) p_D(D^*) \psi(\mathbf{k}^*, G^*, D^*),$$

where  $p_G$  and  $p_D$  are the distributions of  $G$  and  $D$ , then the marginal law of  $\mathbf{k}^*$  coincides with  $p(\mathbf{k}|\mathbf{z})$ . We note that sampling from  $p(\mathbf{k}^* | D^*, G^*)$  is simple, because the variables  $k_{xy}^*$  are conditionally independent, and that sampling from  $p(D^*, G^*|\mathbf{k}^*)$  is straightforward. The random variables  $D^*$  and  $G^*$  conditionally on  $\mathbf{k}^*$  are independent with Dirichlet and Gamma distributions.

Finally, we use these conditional distributions to construct a Gibbs sampler for  $p(\mathbf{k}^*, D^*, G^*)$ . In any Markov Chain Monte Carlo algorithm, it is important to ensure mixing. In Appendix D, we derive an exact sampler for  $p(\mathbf{k}|\mathbf{z})$  which uses a coupling of the Gibbs Markov chain just defined. The method is related to Coupling From The Past (22). We performed simulations with the exact sampler to check the convergence of the proposed Gibbs algorithm.

## 8. Analysis of molecular dynamics simulations.

8.1. *The species sampling problem.* Species sampling problems have a long history in ecological and biological studies. The aim is to determine the species composition of a population containing an unknown number of species when only a sample drawn from it is available.

A common statistical issue is how to estimate species richness, which can be quantified in different ways. For example, given an initial sample of size  $n$ , species richness might be quantified by the number of new species we expect to observe in an additional sample of size  $m$ . It can be alternatively evaluated in terms of the probability of discovering at the  $(n+m)$ -th draw a new species that does not appear across the previous  $(n+m-1)$  observations; this yields the discovery rate as a function of the size of an hypothetical additional sample. These estimates allow one to infer the coverage of a sample of size  $n+m$ , in other words, the relative abundance of distinct species observed in a sample of size  $n+m$ .

A review of the literature on this problem can be found in Bunge and Fitzpatrick (5). Lijoi et al. proposed a Bayesian nonparametric approach for evaluating species richness, considering a large class of exchangeable models, which include as special case the two-parameter Hoppe urn (16). See also Lijoi et al. (17) and Favaro et al. (11) for a practitioner-oriented illustration using Expressed Sequence Tag (EST) data obtained by sequencing cDNA libraries.

We illustrate the use of the  $(\theta, \alpha, \beta)$  scheme in species sampling problems.

In particular, we evaluate species richness in molecular dynamics simulations.

8.2. *Data.* The data we analyze come from a series of recent studies applying Markov models to protein molecular dynamics simulations (19, and references therein). These computer experiments produce time series of protein structures. The space of structures is discretized, such that two structures in a given state are geometrically similar; this yields a sequence of species which correspond to conformational states that the molecule adopts in water. We apply the  $(\theta, \alpha, \beta)$  scheme to perform predictive inference of this discrete time series.

We analyze two datasets. The first is a simulation of the alanine dipeptide, a very simple molecule. The dataset consists of 25,000 transitions, sampled every 2 picoseconds, in which 104 distinct states are observed. In this case the 50 most frequently observed states constitute 85% of the chain and each of the 104 observed states appears at least 12 times. The second dataset is a simulation of a more complex protein, the WW domain, performed in the supercomputer Anton (24). This example illustrates the complexity of the

Dataset 1: Alanine dipeptide ( $\theta = 25$ )

$\alpha \backslash \beta$	0.03	0.2	0.5	0.8	0.97
0.03	$-3212 \pm 10.4$	$-170 \pm 11.7$	$0 \pm 5.9$	$-982 \pm 7.4$	$-2828 \pm 12.3$
0.2	$-3263 \pm 4.0$	$-220 \pm 6.6$	$-28 \pm 8.6$	$-997 \pm 6.0$	$-2809 \pm 9.5$
0.5	$-3404 \pm 2.1$	$-333 \pm 4.1$	$-125 \pm 15.5$	$-1024 \pm 5.4$	$-2815 \pm 2.3$
0.8	$-3621 \pm 6.9$	$-525 \pm 6.9$	$-232 \pm 3.0$	$-1099 \pm 4.2$	$-2857 \pm 5.8$
0.97	$-3868 \pm 3.1$	$-763 \pm 4.6$	$-447 \pm 3.4$	$-1280 \pm 4.2$	$-2960 \pm 11.9$

Dataset 2: WW domain ( $\theta = 500$ )

$\alpha \backslash \beta$	0.03	0.2	0.5	0.8	0.97
0.03	$-14695 \pm 2.0$	$-5147 \pm 2.2$	$-1701 \pm 1.8$	$-361 \pm 4.0$	$-234 \pm 1.9$
0.2	$-15167 \pm 1.6$	$-5507 \pm 1.4$	$-1865 \pm 4.2$	$-329 \pm 3.5$	$-95 \pm 3.2$
0.5	$-16211 \pm 2.4$	$-6354 \pm 2.4$	$-2365 \pm 0.8$	$-482 \pm 5.6$	$0 \pm 1.2$
0.8	$-17943 \pm 1.6$	$-7893 \pm 1.8$	$-3542 \pm 3.0$	$-1120 \pm 7.6$	$-119 \pm 1.4$
0.97	$-20892 \pm 1.9$	$-10739 \pm 5.0$	$-6143 \pm 1.5$	$-3194 \pm 7.1$	$-964 \pm 0.9$

TABLE 1

The log-likelihood of the data,  $\log p(\mathbf{z})$ , for the  $(\theta, \alpha, \beta)$  scheme at the optimal value of  $\theta$  for a range of values of  $\alpha$  and  $\beta$ . In each table, every entry is shifted by a constant such that the largest entry equals 0.

technology and the large amount of resources required for simulating protein dynamics *in silico*. It also motivates the need for suitable statistical tools for the design and analysis of these experiments. In this dataset 1,410 distinct states are observed in 10,000 transitions, sampled every 20 nanoseconds. Many of the states are observed only a few times; in particular we have 991 states that have been observed fewer than 4 times and 547 states that appear only once.

8.3. *Prior Specification.* To apply the  $(\theta, \alpha, \beta)$  scheme it is necessary to tune the three parameters. We consider the initial weights  $g$  everywhere null except for  $g(\zeta, \zeta) = \theta$  and  $g(X_1, \zeta)$  infinitesimal. The parameters  $\theta$  and  $\alpha$  affect the probability of finding a novel state when the latent process reaches  $\zeta$ , while the parameter  $\beta$  tunes the degree of dependence between the random transition probabilities. We recall that in the extreme case of  $\beta = 1$  the sequence  $(X_i)_{i \in \mathbb{N}}$  is exchangeable and the random transition probabilities out of the observed states become identical.

We proceed by approximating the marginal likelihood of the data for set of parameters  $(\theta, \alpha, \beta)$ , where  $\theta \in \{1, 5, 10, 25, 50, 100, 300, 400, 500\}$ ,  $\alpha \in \{0.03, 0.2, 0.5, 0.8, 0.97\}$  and  $\beta \in \{0.03, 0.2, 0.5, 0.8, 0.97\}$ . We iteratively drew samples, under specific  $(\theta, \alpha, \beta)$  values, from the conditional distribution  $p(\mathbf{k}|\mathbf{z})$  using the Gibbs algorithm defined in the previous section. Note that

$$\eta := \sum_{\mathbf{k}} \frac{1}{p(\mathbf{z}, \mathbf{k})} p(\mathbf{k}|\mathbf{z}) = \sum_{\mathbf{k}} \frac{1}{p(\mathbf{z}, \mathbf{k})} \frac{p(\mathbf{z}, \mathbf{k})}{p(\mathbf{z})} = \frac{\prod_{x,y \in \mathbb{N} \sqcup \mathcal{S}} (n_{xy} - 1)}{p(\mathbf{z})},$$

where the last equality is obtained counting the possible values of  $\mathbf{k}$ . We compare the models on the basis of approximations  $\hat{p}(\mathbf{z})$  of the marginal probabilities  $p(\mathbf{z})$  across prior parameterizations. The samples we drew from  $p(\mathbf{k}|\mathbf{z})$  are used to compute Monte Carlo estimates  $\hat{\eta}$  of  $\eta$ . Recall that the probability  $p(\mathbf{k}, \mathbf{z})$  can be computed using the analytic expressions derived in Section 6. Using  $\hat{\eta}$  we compute the estimates  $\hat{p}(\mathbf{z}) := \prod_{x,y \in \mathbb{N} \sqcup \mathcal{S}} (n_{xy} - 1) / \hat{\eta}$  and obtain standard errors by Bootstrapping.

In Table 1, we report the logarithm of these estimates for each model, shifted by a constant such that the largest entry for each dataset is 0. We only show, due to limits of space, these results for the  $\theta$  values associated with the maxima of  $\hat{p}(\mathbf{z})$  across the considered parameterizations. The difference between two entries corresponds to a logarithmic Bayes factor between two models. The values in Table 1 indicate that in each dataset there is one model for which there is strong evidence against all others. This also holds when several values of  $\theta$  are considered. For each dataset, we have highlighted

the optimal parameters. The degenerate cases  $\alpha = 0$  and  $\beta = 1$  were also included in the comparisons but are not shown in Table 1. The difference in the marginal log-likelihood between models with  $\alpha = 0$  and  $\alpha = 0.03$  is negligible. On the other hand, shifting the parameter  $\beta$  from 0.97 to 1 in the optimal model for Dataset 2 decreased the log-likelihood by 7565, as this model is exchangeable and does not capture the Markovian nature of the data. These observations suggest that a fully Bayesian treatment with a hyper-prior over a grid of possible  $(\theta, \alpha, \beta)$  combinations would produce similar results.

Summarizing, the use of a three dimensional grid and the computation of Monte Carlo estimates allows one to effectively obtain a parsimonious approximation of the likelihood function that, in our case, supported selection of single parameterizations.

*8.4. Posterior Estimates.* The main results of our analysis are summarized in Figure 3. Conditional on each sample  $\mathbf{k} \sim p(\mathbf{k}|\mathbf{z})$ , generated under the selected  $(\theta, \alpha, \beta)$  parametrization, we simulated 20,000 future transitions using our predictive scheme. Once  $\mathbf{k}$  is conditionally sampled, the predictive simulations become straightforward with the reinforcement scheme. To provide a measure of species richness and the associated uncertainty, we histogram the number of new states discovered in our simulations in Figure 3. Only a few states are predicted to be found for dataset 1, while a large number of new states are predicted for dataset 2. This result is not surprising because the alanine dipeptide dataset has a limited number of rarely observed states, while in the WW domain data a significant number of states are observed once. This result also seems consistent with the selected values of  $\theta$  and  $\beta$  in these two experiments.

As previously mentioned the  $(\theta, \alpha, \beta)$  scheme is a Bayesian tool for predicting any characteristic of the future trajectories  $X_{n+1}, \dots, X_{n+m}$ . The bottom panels in Figure 3 show confidence bands for the predicted fractions of time that will be spent at the most frequently observed states in the next 20,000 transitions. Each box in the plots refers to a single state and shows the quartiles and the 10th and 90th percentiles of the predictive distribution; states are ordered according to their mean observed frequency. We only show these occupancies for the 20 most populated states, and below the dashed line, we show the total occupancy for states that do not appear in the original data. In the WW domain example, the simulation is expected to spend between 2.5% and 5% of the time at new states.

To assess the predictive performance of the model we split each dataset into a training set and a validation set. The rationale of this procedure




FIG 3. *Posterior simulations for two molecular dynamics datasets. Top: Histogram of the number of new species found in 600 simulations from the predictive distribution for  $X_{n+1}, \dots, X_{n+20,000}$ , given the data  $X_1, \dots, X_n$ . Bottom: Box plot of the fraction of time spent at each state in these simulations. Only the twenty most populated states are shown; below the dashed line, we show the fraction of time spent at states not observed in the dataset.*

is identical to routinely performed cross validations for i.i.d. data. In our setting, the training and validation sets are independent portions of a homogeneous Markov chain. The first part of the procedure, which uses only the training set, includes selection of the  $(\theta, \alpha, \beta)$  parameters and posterior computations. Then, we contrast Bayesian predictions to statistics of the validation set. Overall, this approach suggests that our model generates




FIG 4. Bayesian predictions for the number of new species in a validation set. The histograms show the number of species found in 600 simulations from the predictive distribution for  $X_{n+1}, \dots, X_{n+m}$ , where  $m$  is the length of the validation set and  $X_1, \dots, X_n$  is the training set. The blue line shows the mean of these samples. The red line shows the actual number of new species found in the validation set. Note that in the right panel, the lines overlap due to the small separation between them.

reliable predictions. Figure 4 shows histograms for the number of new species found in predictive simulations of equal length as the validation set. In each panel, the blue line is the Bayes estimate and the red line corresponds to the number of species that was actually discovered in the validation set. This approach also supports the inference reported with box plots in Figure 3. We repeated the computations for deriving the results in Figure 3 using only the training data, and considering a future trajectory equal in length to the validation data. In this case, 37 out of 42 of the true state occupancies in the validation set were contained in the 90% posterior confidence bands.

**9. Discussion.** We introduced a reinforced random walk with a simple predictive structure that can be represented as a mixture of reversible Markov chains. The model generalizes exchangeable and partially exchangeable sequences that have been extensively studied in the literature. Our nonparametric prior, the de Finetti measure of the  $(\theta, \alpha, \beta)$  scheme, can be viewed as a distribution over weighted graphs with a countable number of vertices in a possibly uncountable space  $\mathcal{X}$ . As is the case for other well known Bayesian nonparametric models such as the Dirichlet process (12), the hierarchical Dirichlet process (26) and the infinite hidden Markov model (3), it is possible to represent our model as a function of two independent components, a species sampling sequence  $(Z_i)_{i \in \mathbb{N}}$  and a process  $(T_i)_{i \in \mathbb{N}}$  which determines the species' locations. This property is fundamental in ap-

plications including Dirichlet process mixture models and the infinite hidden Markov model.

A natural extension of our model, not tackled here, is the definition of hidden reversible Markov models. A simple construction would consist of convolving our vertices with suitable density functions. We hope reversibility can be an advantageous assumption in relevant applications; in particular we think reversibility can be explored as a tool for the analysis of genomic data and time series from single-molecule biophysics experiments.

**Acknowledgements.** We are grateful to an Associate Editor and three Referees for their constructive comments and suggestions. We would like to thank DE Shaw Research (24) for providing the molecular dynamics simulations of the WW domain. The Markov models analyzed in Section 8 were generated by Kyle Beauchamp, using the methodology described in (19). We would also like to thank Persi Diaconis and Vijay Pande for helpful suggestions.

## SUPPLEMENTARY MATERIAL

### Appendices B, C, and D.

(doi: XXX, <http://lib.stat.cmu.edu/aos/XXX/XXX>; .pdf). Appendix B describes the two-parameter HDP-HMM in relation to the  $(\theta, \alpha, \beta)$  scheme. Appendix C contains all proofs from Sections 4, 5, and 6. Appendix D contains a derivation of the exact sampler mentioned in Section 7 using Coupling From the Past.

### References.

- [1] BACALLADO, S. (2011). Bayesian analysis of variable-order, reversible Markov chains. *Ann. Statist.* **39** 838–864.
- [2] BACALLADO, S., FAVARO, S. and TRIPPA, L. (2012). Supplement to: "Bayesian nonparametric analysis of reversible Markov chains".
- [3] BEAL, M., GHARAMANI, Z. and RASMUSSEN, C. (2002). The infinite hidden Markov model. *Adv. Neur. Inf. Proc. Sys.* **14** 577–584.
- [4] BLACKWELL, D. and MACQUEEN, J. B. (1973). Ferguson distributions via Pólya urn schemes. *Ann. Statist.* **1** 353–355.
- [5] BUNGE, J. and FITZPATRICK, M. (1993). Estimating the number of species: a review. *J. Amer. Statist. Assoc.* **88** 364–373.
- [6] COMTET, L. (1974). *Advanced Combinatorics: The art of finite and infinite expansions*. D. Reidel Pub. Co.
- [7] DIACONIS, P. (1988). Recent progress on de Finetti notions of exchangeability. In *Bayesian Statistics 3* (J. M. BERNARDO, M. H. DEGROOT, D. V. LINDLEY and A. F. M. SMITH, eds.) 111-125. Oxford University Press, New York.
- [8] DIACONIS, P. and FREEDMAN, D. (1980). de Finetti theorem for Markov chains. *Ann. Probab.* **8** 115-130.

- [9] DIACONIS, P. and ROLLES, S. W. W. (2006). Bayesian analysis for reversible Markov chains. *Ann. Statist.* **34** 1270–1292.
- [10] ENGEN, S. (1978). *Stochastic abundance models, with emphasis on biological communities and species diversity*. Chapman and Hall.
- [11] FAVARO, S., LIJOI, A., MENA, R. H. and PRÜNSTER, I. (2009). Bayesian nonparametric inference for species variety with a two-parameter Poisson–Dirichlet process prior. *J. R. Stat. Soc. Ser. B* **71** 993–1008.
- [12] FERGUSON, T. S. (1973). A Bayesian analysis of some nonparametric problems. *Ann. Statist.* **1** 209–230.
- [13] FORTINI, S. and PETRONE, S. (2012). Hierarchical reinforced urn processes. *Stat. Prob. Lett.* **82** 1521–1529.
- [14] ISHWARAN, H. and JAMES, L. F. (2001). Gibbs sampling methods for stick-breaking priors. *J. Amer. Statist. Assoc.* **96** 161–173.
- [15] KEANE, M. S. and ROLLES, S. W. W. (2000). Edge-reinforced random walk on finite graphs. In *Infinite dimensional stochastic analysis*. 217–234. R. Neth. Acad. Arts Sci., Amsterdam.
- [16] LIJOI, A., MENA, R. H. and PRÜNSTER, I. (2007). Bayesian nonparametric estimation of the probability of discovering new species. *Biometrika* **94** 769–786.
- [17] LIJOI, A., MENA, R. H. and PRÜNSTER, I. (2007). A Bayesian nonparametric method for prediction in EST analysis. *BMC Bioinformatics* **8** 339.
- [18] MERKL, F. and ROLLES, S. W. W. (2009). Recurrence of edge-reinforced random walk on a two-dimensional graph. *Ann. Probab.* **37** 1679–1714.
- [19] PANDE, V. S., BEAUCHAMP, K. and BOWMAN, G. R. (2010). Everything you wanted to know about Markov State Models but were afraid to ask. *Methods* **52** 99–105.
- [20] PITMAN, J. (1996). Some developments of the Blackwell–MacQueen urn scheme. In *Statistics, Probability and Game Theory; Papers in honor of David Blackwell*, (T. S. Ferguson, L. S. Shapley and M. J. B., eds.) **30** 245–267. Institute of Mathematical Statistics.
- [21] PITMAN, J. and YOR, M. (1997). The two-parameter Poisson–Dirichlet distribution derived from a stable subordinator. *Ann. Probab.* **25** 855–900.
- [22] PROPP, J. G. and WILSON, D. B. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures Algorithms* **9** 223–252.
- [23] ROLLES, S. W. W. (2003). How edge-reinforced random walk arises naturally. *Probab. Theory Related Fields* **126** 243–260.
- [24] SHAW, D. E. (2010). Atomic-level characterization of the structural dynamics of proteins. *Science* **330** 341.
- [25] TEH, Y. and JORDAN, M. (2010). Hierarchical Bayesian nonparametric models with applications. *Bayesian nonparametrics* 158–207.
- [26] TEH, Y., JORDAN, M., BEAL, M. and BLEI, D. (2006). Hierarchical dirichlet processes. *J. Amer. Statist. Assoc.* **101** 1566–1581.
- [27] ZABELL, S. L. (1982). W.E. Johnson’s sufficientness postulate. *Ann. Statist.* **10** 1091–1099.
- [28] ZABELL, S. (2005). The continuum of inductive methods revisited. In *Symmetry and its Discontents: Essays on the History of Inductive Probability*. New York: Cambridge University Press.

APPENDIX A: PROOFS FROM SECTIONS 2 AND 3

PROOF OF PROPOSITION 2.1. Consider the latent process on  $\mathcal{X}_+$ . The transition probability from  $x_1 \in \mathcal{X}_+$  to  $x_2 \in \mathcal{X}_+$  with  $g(x_1, x_2) > 0$  is of the form  $g(x_1, x_2) / \sum_{y \in \mathcal{X}_+} g(x_1, y)$ . Between successive visits to  $x_1$ , the denominator  $\sum_{y \in \mathcal{X}_+} g(x_1, y)$  is increased by at most 2, and the numerator may only increase. Assume that almost surely, the process visits  $x_1$  infinitely often. There exist  $c_2 \geq c_1 > 0$ , such that if  $A_{n,m}$  is the event that we do not traverse  $(x_1, x_2)$  between the  $n$ -th and  $m$ -th visits to  $x_1$ ,

$$p(A_{n,m}) \leq \prod_{i=n}^{m-1} \left[ 1 - \frac{c_1}{c_2 + 2(i-n)} \right],$$

which goes to 0 as  $m \rightarrow \infty$ . Therefore the edge  $(x_1, x_2)$  is a.s. traversed infinitely often. Thus, if  $x_1 \in \mathcal{X}_+$  is a.s. visited infinitely often, by induction the process a.s. returns infinitely often to all visited states. Suppose a state in  $\mathcal{X}$  is visited infinitely often a.s., then the process visits  $\zeta$  infinitely often by the previous argument. Otherwise, the process must visit an infinite number of states in  $\mathcal{X}$ , and since the set of pairs  $(x, y) \in \mathcal{X}^2$  with a positive initial weight  $g(x, y)$  is a finite subset of  $\mathcal{S}^2$ , we must go through  $\zeta$  an infinite number of times. We conclude that  $\zeta$  is visited infinitely often a.s. and therefore the process returns to every state visited infinitely often. If  $\theta > 0$ , then the edge  $(\zeta, \zeta)$  is crossed infinitely often, and we see an infinite number of distinct states.  $\square$

PROOF OF PROPOSITION 3.1. For  $1 \leq i < n$  there is a latent variable  $U_i \in \{a, b, c\}$  that determines in which of the three ways outlined in Figure 2 the transition  $X_i \rightarrow X_{i+1}$  proceeded. The probability of  $Z_1 = x_1, \dots, Z_n = x_n$  is the sum of its joint probability with every latent sequence  $U_1 = u_1, \dots, U_{n-1} = u_{n-1}$ . We will show that there is a one to one map  $L$  of the latent sequences such that, letting  $L(u_1, \dots, u_{n-1}) = u'_1, \dots, u'_{n-1}$ , one has

$$(A.1) \quad \begin{aligned} p(Z_1 = x_1, \dots, Z_n = x_n; U_1 = u_1, \dots, U_{n-1} = u_{n-1}) \\ = p(Z_1 = x'_1, \dots, Z_n = x'_n; U_1 = u'_1, \dots, U_{n-1} = u'_{n-1}). \end{aligned}$$

The proposition follows from this claim. Let  $t$  denote the transposition such that  $d(t(x_1, \dots, x_n)) = x'_1, \dots, x'_n$ . Define  $x''_1, \dots, x''_n = t(x_1, \dots, x_n)$ . The map  $L$  is defined so that for any  $u_1, \dots, u_{n-1}$  and  $j, k < n$ , satisfying

$$\mathbf{1}_{(x_j=x, x_{j+1}=y)} + \mathbf{1}_{(x_j=y, x_{j+1}=x)} > 0,$$

$$\mathbf{1}_{(x'_k=x, x''_{k+1}=y)} + \mathbf{1}_{(x'_k=y, x''_{k+1}=x)} > 0$$

for some  $x$  and  $y$ , if

$$\sum_{i=1}^j \mathbf{1}_{(x_i=x, x_{i+1}=y)} + \mathbf{1}_{(x_i=y, x_{i+1}=x)} = \sum_{i=1}^k \mathbf{1}_{(x'_i=x, x''_{i+1}=y)} + \mathbf{1}_{(x'_i=y, x''_{i+1}=x)}$$

then  $u_j = u'_k$ . Note that we can define the joint probability of  $(Z_i)_{i \leq n}$  and  $(U_i)_{i < n}$  through a reinforcement scheme identical the one defined in Section 2. Precisely, the probability of each transition and associated category is of the form

$$\frac{g(x_i, x_{i+1})}{g(x_i)} \mathbf{1}_{u_i=a}, \quad \frac{g(x_i, \zeta)g(\zeta, x_{i+1})}{g(x_i)g(\zeta)} \mathbf{1}_{u_i=b}, \quad \frac{g(x_i, \zeta)g(\zeta, \zeta)}{g(x_i)g(\zeta)} \mathbf{1}_{u_i=c},$$

where  $g(x) := \sum_{y \in \mathcal{X}_+} g(x, y)$ .

The factors  $g(\zeta)$ , which appear in the denominator when  $U_i \in \{b, c\}$ , are reinforced by  $2\beta$  between successive visits to  $\zeta$ . Therefore, their product only depends on the number of mediated transitions, which is invariant under  $L$ . Similarly, factors  $g(\zeta, \zeta)$  increase by  $\alpha\beta$  between successive occurrences; their product is identical when we compute the two sides of (A.1) because the number of mediated transitions with discovery remains identical. Also, the factors  $g(x)$  in the denominators increase by 2 between successive occurrences of the same  $\mathcal{X}$  state; their product is identical when we compute the two sides of (A.1) because the number of transitions out of any state  $x$  (or toward  $x$ ) remains identical. Finally, we need to prove the identity between

$$(A.2) \quad \prod_i \left( g(x_i, x_{i+1}) \mathbf{1}_{u_i=a} + g(x_i, \zeta)g(x_{i+1}, \zeta) \mathbf{1}_{u_i=b} + g(x_i, \zeta) \mathbf{1}_{u_i=c} \right)$$

and

$$(A.3) \quad \prod_i \left( g(x'_i, x'_{i+1}) \mathbf{1}_{u'_i=a} + g(x'_i, \zeta)g(x'_{i+1}, \zeta) \mathbf{1}_{u'_i=b} + g(x'_i, \zeta) \mathbf{1}_{u'_i=c} \right).$$

The identity between (A.2) and (A.3) follows by combining the definitions of  $d$  and  $L$  with the reinforcement mechanism. Specifically, the factors  $g(x, \zeta)$  and  $g(x', \zeta)$  are increased by  $\beta$  between successive occurrences. Since  $g(x, \zeta)$  appears as many times in the lefthand side of (A.1) as  $g(x', \zeta)$  does in the righthand side of (A.1), the product of these factors is identical in each case. The remaining factors  $g(x, y)$  may increase by different amounts between successive occurrences. Their product is a function of the subsequence of  $U_1, \dots, U_{n-1}$  with indices  $\{1 \leq i < n : \{Z_i, Z_{i+1}\} = \{x, y\}\}$ . By the definition of  $L$ , this subsequence is the same in the left and right hand sides of (A.1), which completes the proof of our claim.  $\square$

PROOF OF PROPOSITION 3.2. Let  $(X_i'')_{i \in \mathbb{N}}$  be a  $(\theta, \alpha, \beta)$  scheme. The process  $(X_i'')_{i \in \mathbb{N}}$  returns to  $X_1''$  infinitely often a.s. Let  $h_i$  be the  $i$ -th  $X_1''$ -block. Define  $(X_i)_{i \in \mathbb{N}} := (h_1, h_3, h_5, \dots)$  and  $(X_i')_{i \in \mathbb{N}} := (h_2, h_4, h_6, \dots)$ . Proposition 3.1 implies  $(X_i)_{i \in \mathbb{N}} \stackrel{d}{=} (X_i')_{i \in \mathbb{N}} \stackrel{d}{=} (X_i'')_{i \in \mathbb{N}}$ . Let  $\text{Fin}(x)$  be the last element of a vector  $x$ . Define

$$W_i := \lim_{m \rightarrow \infty} \text{Fin} \left( d(X_1', X_2', \dots, X_m', X_1, X_2, \dots, X_i) \right).$$

This limit exists a.s. because  $X_i$  is recurrent in  $(X_j'')_{j \in \mathbb{N}}$ ; therefore, the sequence of blocks that form  $(X_j')_{j \in \mathbb{N}}$  is conditionally i.i.d. from a distribution which a.s. assigns positive probability to blocks containing  $X_i$ , which implies  $(X_j')_{j \in \mathbb{N}}$  visits  $X_i$  after a finite time a.s., at which point the limit settles.

In Lemma A.1 we show that  $(W_i)_{i \in \mathbb{N}}$  is Markov exchangeable and recurrent. Therefore, by de Finetti's theorem for Markov chains (3.1), it is a mixture of Markov chains. Finally, by Lemma A.2, we obtain the representation claimed in the proposition.  $\square$

LEMMA A.1. *Without loss of generality, let  $\mathcal{X} = (0, 1)$ . The process  $(W_i)_{i \in \mathbb{N}}$  is Markov exchangeable and returns to every state in  $\mathcal{S} \sqcup \mathbb{N}$  infinitely often a.s.*

PROOF. The recurrence of  $(Z_i)_{i \in \mathbb{N}}$ , which is a consequence of Proposition 2.1, implies the recurrence of  $(W_i)_{i \in \mathbb{N}}$ . Thus, we have left to show Markov exchangeability.

The sequence  $W_1, \dots, W_n$  can be mapped through  $d$  to  $Z_1, \dots, Z_n$ , which is a species sampling sequence for the  $(\theta, \alpha, \beta)$  scheme. Take any sequence  $w_1, \dots, w_n$  and let  $z_1, \dots, z_n := d(w_1, \dots, w_n)$ . We have

$$(A.4) \quad \begin{aligned} p(W_1 = w_1, \dots, W_n = w_n) &= p(Z_1 = z_1, \dots, Z_n = z_n) \\ &\quad \times p(W_1 = w_1, \dots, W_n = w_n | Z_1 = z_1, \dots, Z_n = z_n). \end{aligned}$$

Consider any pair of sequences  $w_1, \dots, w_n$  and  $w'_1, \dots, w'_n$  related by a transposition of two blocks with identical initial and final states. Proposition 3.1 implies

$$p\left((Z_1, \dots, Z_n) = d(w_1, \dots, w_n)\right) = p\left((Z_1, \dots, Z_n) = d(w'_1, \dots, w'_n)\right).$$

We have left to show that the second factor on the right hand side of Eq. A.4 is identical for  $w_1, \dots, w_n$  and  $w'_1, \dots, w'_n$ . The identity of the conditional distribution of  $(X_i')_{i \in \mathbb{N}}$  given  $(Z_1, \dots, Z_n)$  equal to  $d(w_1, \dots, w_n)$  or equal to  $d(w'_1, \dots, w'_n)$  proves the lemma.  $\square$

LEMMA A.2. *The process  $(T_{W_i}\mathbf{1}_{W_i \notin \mathcal{S}} + W_i\mathbf{1}_{W_i \in \mathcal{S}})_{i \in \mathbb{N}}$  has the same distribution as  $(X_i)_{i \in \mathbb{N}}$ .*

PROOF. By definition  $(X_i)_{i \in \mathbb{N}} \stackrel{d}{=} (T_{Z_i}\mathbf{1}_{Z_i \notin \mathcal{S}} + Z_i\mathbf{1}_{Z_i \in \mathcal{S}})_{i \in \mathbb{N}}$ . Note that  $(T_i)_{i \in \mathbb{N}}$  is an i.i.d. sequence, independent from  $(Z_i)_{i \in \mathbb{N}}$ , and  $d((W_i)_{i \in \mathbb{N}}) = d((Z_i)_{i \in \mathbb{N}})$ . These facts imply that  $(T_{W_i}\mathbf{1}_{W_i \notin \mathcal{S}} + W_i\mathbf{1}_{W_i \in \mathcal{S}})_{i \in \mathbb{N}} \stackrel{d}{=} (T_{Z_i}\mathbf{1}_{Z_i \notin \mathcal{S}} + Z_i\mathbf{1}_{Z_i \in \mathcal{S}})_{i \in \mathbb{N}}$ .  $\square$

PROOF OF PROPOSITION 3.3. Let  $h_1, h_2, \dots$  be the  $X_1$ -blocks of the  $(\theta, \alpha, \beta)$  scheme. Consider a map  $s$  on the  $X_1$ -blocks' space; if  $a = (a_1, \dots, a_m)$  then  $s(a) = (a_1, a_m, a_{m-1}, \dots, a_2)$ . We can observe, following the same arguments used for proving Proposition 3.1, that for any  $X_1$ -block  $a$  and any integer  $n$

$$p(h_1 = a, h_2, \dots, h_n) = p(h_1 = s(a), h_2, \dots, h_n).$$

Let  $F$  be a random measure distributed according to the de Finetti measure of the  $X_1$ -blocks. The above expression and the equality

$$\lim_{n \rightarrow \infty} p(h_1 \in A \mid h_2, \dots, h_n) \stackrel{\text{a.s.}}{=} F(A),$$

where  $A$  is a generic measurable set, imply that a.s. the distance in total variation between  $F$  and  $F \circ s$  is null.  $\square$

UNIVERSITY OF TORINO  
CORSO UNIONE SOVIETICA 218/BIS  
TORINO, ITALY 10134  
E-MAIL: [stefano.favaro@unito.it](mailto:stefano.favaro@unito.it)

DANA-FARBER CANCER INSTITUTE  
450 BROOKLINE AVE. CLSB 11039  
BOSTON, MA 02215  
E-MAIL: [ltrippa@jimmy.harvard.edu](mailto:ltrippa@jimmy.harvard.edu)

CLARK CENTER, S296  
STANFORD UNIVERSITY  
STANFORD, CA 94305  
E-MAIL: [sergiob@stanford.edu](mailto:sergiob@stanford.edu)