

BAYESIAN ANALYSIS OF VARIABLE-ORDER, REVERSIBLE MARKOV CHAINS¹

BY SERGIO BACALLADO

Stanford University

We define a conjugate prior for the reversible Markov chain of order r . The prior arises from a partially exchangeable reinforced random walk, in the same way that the Beta distribution arises from the exchangeable Polya urn. An extension to variable-order Markov chains is also derived. We show the utility of this prior in testing the order and estimating the parameters of a reversible Markov model.

1. Introduction. Reversible Markov chains are central to a number of fields. They underlie problems in applied probability like card-shuffling and queueing networks [1, 13] and pervade computational statistics through the many variants of Markov chain Monte Carlo; in physics, they are natural stochastic models for time-reversible dynamics. However, the notion of reversibility in stochastic processes with memory is not as widely discussed, and statistical problems like testing the order of a reversible process remain a challenge.

We define a conjugate prior for higher-order, reversible Markov chains, which extends a prior for reversible Markov chains by Diaconis and Rolles [10]. We begin by defining reversibility in a more general setting and motivating the significance of higher-order processes. In Section 2, we present two graphical representations for an order- r , reversible Markov chain, which are used in Section 3 to derive the conjugate prior via a random walk with reinforcement. We dedicate Section 4 to variable-order Markov chains, a family of models that avoids the curse of dimensionality associated with higher-order Markov chains, proving essential in certain applications. Finally in Section 5, we discuss properties of the prior pertaining to Bayesian analysis. In examples, we test the extent of memory of a lumped Markov chain and discretized molecular dynamics trajectories, and compare the posterior inferences of different models.

DEFINITION 1.1. A stochastic process $X = X_n, n \in \mathbb{N}$, with distribution P is called *reversible*, if for any $m > n > 0$,

$$P(X_1, X_2, \dots, X_n) = P(X_{m-1}, X_{m-2}, \dots, X_{m-n}).$$

Received May 2010; revised September 2010.

¹Supported by a Smith Stanford Graduate Fellowship and Grants NIH-R01-GM062868 and NSF DMS-09-00700.

MSC2010 subject classifications. Primary 62M02; secondary 62C10.

Key words and phrases. Reversibility, reinforced random walks, variable-order Markov chains, Bayesian analysis, conjugate priors.

It is not difficult to show that reversibility implies stationarity [13]; if stationarity is given, the above condition need only be checked for $m = n + 1$. Now suppose X is an order- r , irreducible Markov chain taking values in a finite set \mathcal{X} . We will also apply the term *reversible* to this process when the stationary chain satisfies the reversibility condition.

PROPOSITION 1.2. *Let P be the stationary law of the order- r Markov chain X . If $P(X_1, \dots, X_{r+1}) = P(X_{r+1}, \dots, X_1)$, then the Markov chain is reversible.*

PROOF. It is not difficult to check that the hypothesis together with stationarity imply $P(X_1, \dots, X_n) = P(X_n, \dots, X_1)$ for any $n < r + 1$. For any $n > r + 1$:

$$\begin{aligned} P(X_1, \dots, X_n) &= P(X_1, \dots, X_{r+1}) \prod_{i=r+2}^n P(X_i | X_{i-r}, \dots, X_{i-1}) \\ &= P(X_1, \dots, X_{r+1}) \frac{P(X_2, \dots, X_{r+2})}{P(X_2, \dots, X_{r+1})} \dots \frac{P(X_{n-r}, \dots, X_n)}{P(X_{n-r}, \dots, X_{n-1})} \\ &= P(X_n, \dots, X_{n-r}) \frac{P(X_{n-1}, \dots, X_{n-r-1})}{P(X_{n-1}, \dots, X_{n-r})} \dots \frac{P(X_{r+1}, \dots, X_1)}{P(X_{r+1}, \dots, X_2)} \\ &= P(X_n, \dots, X_1), \end{aligned}$$

where we have used the Markov property, stationarity and the hypothesis. \square

As a first remark, note that $X_n, n \in \mathbb{N}$, can be represented as a first-order Markov chain $V_n, n \in \mathbb{N}$, taking values in the space of sequences \mathcal{X}^r . However, the reversibility of X does not imply the reversibility of its first-order representation; therefore, the analysis of higher-order reversible Markov chains requires novel techniques. In the following sections, we often use the first-order representation $V_n, n \in \mathbb{N}$, referring to it nonetheless as an order- r Markov chain and using the notion of reversibility associated with the order- r Markov chain.

Secondly, we recall that Kolmogorov’s criterion is another necessary and sufficient condition for the reversibility of a Markov chain, which only depends on the conditional transition probabilities [13]. Its equivalence to Definition 1.1 in the higher-order case is proven in the [Appendix](#). Kolmogorov’s criterion requires that the probability of traversing any cycle in either direction is the same. Accordingly, a reversible Markov chain can be interpreted as a process with no net circulation in space.

Reversibility is preserved under certain transformations. For example, let $X_n, n \in \mathbb{N}$, be a stationary, reversible Markov chain and consider a finitely valued function, $f(X_n), n \in \mathbb{N}$. It is easy to check that this process is stationary and reversible, even though it may not be a Markov chain of any finite order. Functions or projections of reversible Markov chains appear under different guises in physics

and other fields, and in many cases the effects of memory subside with time, motivating the use of finite order models. The problems of determining the order and estimating the parameters of Markov models have been studied extensively; here, we address these problems with the constraint of reversibility.

2. Graphical representations of reversible Markov chains. For any sequence $u \in \mathcal{X}^s$, let u^* be its inverse, $A(u)$ the subsequence obtained by deleting its last element and $\Omega(u)$ the one obtained by deleting its first element. We call u_1, u_2, \dots, u_n with $u_i \in \mathcal{X}^s$ an *admissible path* if $\Omega(u_i) = A(u_{i+1})$ for all $1 \leq i < n$. The concatenation of these sequences without repeated overlaps is denoted $\overline{u_1 \cdots u_n} \in \mathcal{X}^{s+n-1}$.

The first representation we will consider is the circuit process of MacQueen [14]. Let a *circuit* be a periodic function on \mathcal{X} , and consider a class of positively weighted circuits \mathcal{C} (for an example, see Figure 1).

DEFINITION 2.1. A *circuit process* of order r is a Markov chain of the same order, where the transition probability from $u \in \mathcal{X}^r$ to any $v \in \mathcal{X}^r$ with $\Omega(u) = A(v)$ is given by

$$\frac{\sum_{\gamma \in \mathcal{C}} w_\gamma J_\gamma(\overline{uv})}{\sum_{\gamma \in \mathcal{C}} w_\gamma J_\gamma(u)},$$

where $w_\gamma > 0$ is the weight of circuit γ , and the function $J_\gamma(\cdot)$ counts the number of times that the circuit traverses a sequence in one period. In other words, in each step we move along some circuit in \mathcal{C} containing the current state with probability proportional to its weight. The process only visits states that appear in the circuits, for which transition probabilities are well defined.

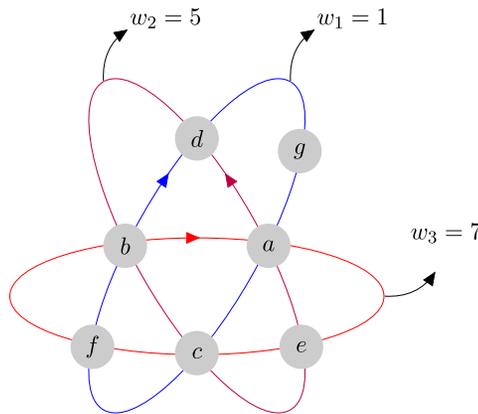


FIG. 1. A set of weighted circuits on the set $\mathcal{X} = \{a, b, c, d, e, f, g\}$. In a circuit process started at u in \mathcal{X}^r , we transition on some circuit that contains u with probability proportional to its weight.

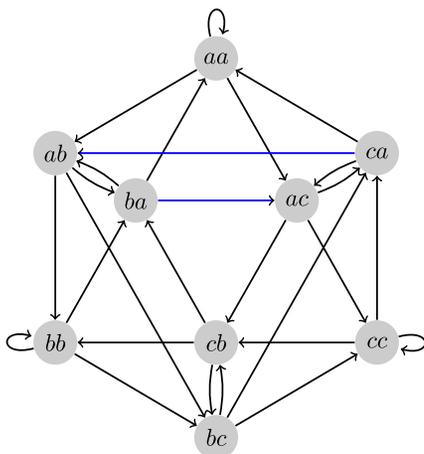


FIG. 2. A de Bruijn graph of order 2 on the state space $\mathcal{X} = \{a, b, c\}$. In a reversible random walk, the two highlighted edges have the same weight.

An irreducible order- r Markov chain with stationary law P_π is parametrized by $P_\pi(u)$ for all $u \in \mathcal{X}^{r+1}$. One can check that in a circuit process, this is just $P_\pi(u) = \sum_{\gamma \in \mathcal{C}} w_\gamma J_\gamma(u)$. MacQueen showed that any order- r Markov chain can be represented as a circuit process on a finite set \mathcal{C} , which is not unique [14]. This is true in particular when the chain is reversible.

We introduce a second graphical representation that is canonical, unlike the circuit process. Consider a *de Bruijn graph* on the vertices \mathcal{X}^r , which has a directed edge from u to v if and only if $\Omega(u) = \Lambda(v)$. That is, every path on the graph is an admissible path. For an example, see Figure 2. Assign a weight $k_{uv} \geq 0$ to each edge, and let k_u be the summed weights of edges departing from u . Furthermore, require that

- (1) $k_{uv} = k_{v^*u^*}$ for every edge uv ,
- (2) $k_u = k_{u^*}$ for all $u \in \mathcal{X}^r$ and
- (3) $\sum_{u \in \mathcal{X}^r} k_u = 1$.

DEFINITION 2.2. The *reversible random walk of order r* is a random walk on such a graph, with transition probabilities

$$p(v|u) = \frac{k_{uv}}{k_u}.$$

PROPOSITION 2.3. An irreducible, reversible random walk of order r represents a reversible Markov chain of the same order. Every irreducible, reversible order- r Markov chain is equivalent to a unique reversible random walk of order r .

PROOF. Let π be the stationary distribution of the random walk. To prove the first statement, we will first verify that $\pi(u) = k_u$ for all $u \in \mathcal{X}^r$. Let $p(u|v)$ be the transition probability from v to u in the random walk, and recall that $\Omega(u) = A(v)$ iff $\Omega(v^*) = A(u^*)$, then

$$\begin{aligned} \sum_{u \in \mathcal{X}^r} \pi(u)p(v|u) &= \sum_{\{u \in \mathcal{X}^r : \Omega(u)=A(v)\}} k_u \frac{k_{uv}}{k_u} \\ &= \sum_{\{u \in \mathcal{X}^r : \Omega(v^*)=A(u^*)\}} k_{v^*u^*} = k_{v^*} = k_v = \pi(v). \end{aligned}$$

Then, the stationary law P_π in the random walk of a path u, v is just

$$P_\pi(u, v) = \pi(u)p(v|u) = k_u \frac{k_{uv}}{k_u} = k_{uv},$$

which implies that $P_\pi(u, v) = k_{uv} = k_{v^*u^*} = P_\pi(v^*, u^*)$. Therefore, the \mathcal{X} -valued, order- r Markov chain represented by the random walk satisfies the reversibility condition in Proposition 1.2. Proving the second statement is now straightforward. Let $V_n, n \in \mathbb{N}$, be the first-order representation of an irreducible, order- r Markov chain, with transition probabilities $p'(v|u)$. By the Perron–Frobenius theorem, V has a unique stationary distribution π' . Assign edge weights to the de Bruijn graph on \mathcal{X}^r , setting $k_{uv} = \pi'(u)p'(v|u)$. Since the order- r Markov chain is reversible, it follows directly from Proposition 1.2 that the edge weights satisfy conditions (1)–(3). \square

3. From a reinforced random walk to the conjugate prior. An edge-reinforced random walk (ERRW) is a random walk on an finite, undirected graph, where every edge-weight is increased by 1 each time it is crossed. Since Diaconis and Coppersmith defined this process [9], we have learned that it is partially exchangeable and, by de Finetti’s theorem for Markov chains, a mixture of Markov chains [8]. The mixing measure, which lives on the space of reversible Markov chains, was more recently characterized in the literature [12]. Diaconis and Rolles showed that this distribution is a conjugate prior for the reversible Markov chain, much as the Beta distribution, arising from a Polyá urn scheme, is a conjugate prior for sequences of i.i.d. binary random variables [10].

Here, we construct a conjugate prior for higher-order reversible Markov chains via a reinforced random walk in \mathcal{X}^r , making use of de Finetti’s theorem for Markov chains. This process is markedly different from an ERRW in \mathcal{X}^r due to the structure of a reversible Markov chain with memory, although it is designed to be partially exchangeable.

Let α be any sequence on \mathcal{X} and v a sequence shorter than α . Define the function $J'_\alpha(v)$, which counts the number of times that v appears in α , and $J''_\alpha(v)$, which counts the number of times that v appears in α followed by at least one state. Fix w , a stationary measure for an irreducible, reversible, order- r Markov chain. Also

fix $v_0 \in \mathcal{X}^r$. Let β be a palindromic sequence that starts with v_0 and ends with v_0^* . Choose a positive constant c , such that for all $u \in \mathcal{X}^{r+1}$, $w(u) - cJ'_\beta(u) > 0$. Now, given a sequence η , starting with v_0 , and any sequence v , define the functions

$$(4) \quad w'(\eta, v) = w(v) + c(J'_\eta(v) + J'_{\eta^*}(v) - J'_\beta(v)) \quad \text{and}$$

$$(5) \quad w''(\eta, v) = w(v) + c(J''_\eta(v) + J''_{\eta^*}(v) - J''_\beta(v)).$$

When η represents the path of a stochastic process in \mathcal{X}^r up to time n (formally, $\eta = \overline{v_0 \cdots v_n}$), we will use the notation $w'_n(v) \equiv w'(\eta, v)$ and $w''_n(v) \equiv w''(\eta, v)$.

DEFINITION 3.1. The *reinforced random walk of order r* is a stochastic process $Y_n, n \in \mathbb{N}$, on \mathcal{X}^r with distribution Q_{w, v_0} . The initial state is v_0 with probability 1. For any admissible path v_0, \dots, v_n , the conditional transition probability

$$Q_{w, v_0}(Y_{n+1} = u | Y_0 = v_0, \dots, Y_n = v_n) = \frac{w'_n(\overline{v_n u})}{w''_n(v_n)}$$

whenever v_n, u is admissible and zero otherwise.

REMARK 3.2. The law Q_{w, v_0} also depends on β and c . These parameters are constant in the following discussion, so they are omitted from the notation for conciseness. When $r = 1$, this process is equivalent to an ERRW. In this case, the palindrome is unnecessary because the terms involving β in w' and w'' can be modeled with a different w . For $r \geq 2$, this is not the case, and β is essential for partial exchangeability (see Proposition 3.5).

REMARK 3.3. This process admits an interpretation as a reinforcement scheme of the circuit process. Consider a circuit process of order r with stationary probability $w(u) = \sum_{\gamma \in \mathcal{C}} w_\gamma J_\gamma(u)$ for all $u \in \mathcal{X}^{r+1}$. In addition, consider three weighted sequences: the palindrome β , a sequence η that represents the path of the reinforced process from the initial state v_0 up to the current state, and the reversed path η^* . These are depicted in Figure 3 along with their weights $-c, c$

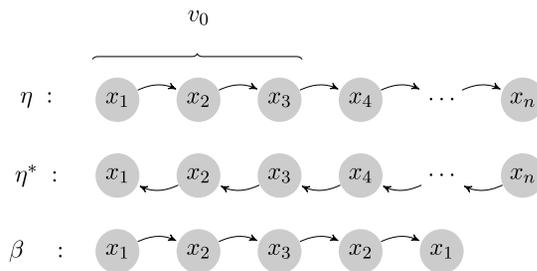


FIG. 3. Auxiliary sequences in the order- r reinforced random walk.

and c , respectively. As in the circuit process, we move along any circuit or sequence that contains the current state with probability proportional to its weight. The reinforcement is accomplished by elongating the paths η and η^* .

REMARK 3.4. The process is also a reinforcement scheme of a modified reversible random walk of order r . Consider a weighted de Bruijn graph, where for every admissible u, v , $k_{uv} = w(\overline{uv})$. Then, for every \overline{uv} in the palindrome β , subtract c from k_{uv} . The reinforcement scheme will consist of a random walk on the resulting graph, where after every transition $v_i \rightarrow v_{i+1}$ we increase both $k_{v_i v_{i+1}}$ and $k_{v_{i+1}^* v_i^*}$ by c . Accordingly, if $\overline{v_i v_{i+1}}$ is a palindrome, the weight $k_{v_i v_{i+1}}$ is increased by $2c$.

PROPOSITION 3.5. *The reinforced random walk of order r is partially exchangeable in the sense of Diaconis and Freedman [8].*

PROOF. We must show that the probability $Q_{w, v_0}(v_0, \dots, v_n)$ of any admissible path v_0, \dots, v_n is a function of the initial state v_0 and the transition counts between every pair of states. For any pair u, v in \mathcal{X}^r with $A(v) = \Omega(u)$, let $C(u, v)$ be the total number of transitions $u \rightarrow v$, and $v^* \rightarrow u^*$. We will show the stronger statement that v_0 and C are sufficient statistics for the reinforced random walk.

Let us first establish some properties that are conserved in the process. For every $u \in \mathcal{X}^{r+1}$, the initial weights $w'_0(u)$ and $w'_0(u^*)$ are equal. This is direct from the definition in equation (4) because: w defines a reversible Markov chain of order r ; the functions J'_{v_0} and $J'_{v_0^*}$ are zero for both u and u^* ; and β is a palindrome, so if it contains u , it also contains u^* , and $J'_\beta(u) = J'_\beta(u^*)$. This property is maintained after every transition $v_n \rightarrow v_{n+1}$, because the weights may both be increased by c if $\overline{v_n v_{n+1}}$ is u or u^* , or both remain constant otherwise.

For every $v \neq v_0$ in \mathcal{X}^r , the initial weights $w''_0(v) = w''_0(v^*)$. This is direct from equation (5) because: w is reversible, both J''_{v_0} and $J''_{v_0^*}$ are zero for v and v^* , and the sequence β is a palindrome, so for every transition starting at v there will be another starting from v^* . The last fact is not necessarily true for v_0 , because unless v_0 itself is a palindrome, β will contain a transition starting from it, but no transition starting from v_0^* . So, in the beginning, $w''_0(v_0) = w''_0(v_0^*) - c$. When a transition occurs from v_0 to v_1 , the weights $w''_1(v_0)$ and $w''_1(v_0^*)$ become equal, while $w''_1(v_1) = w''_1(v_1^*) - c$, provided v_1 is not a palindrome. Hence, this singularity is preserved by the last state visited by the process.

The probability $Q_{w, v_0}(v_0, \dots, v_n)$ is a ratio of two products. In the numerator, we find a factor of the form $w'_i(\overline{uv})$ for every admissible transition $u \rightarrow v$, while in the denominator, we find a corresponding weight $w''_i(u)$. It is easy to check that the numerator is only a function of C . Every transition $u \rightarrow v$ or $v^* \rightarrow u^*$ adds a new factor of $w'_i(\overline{uv})$, which is always greater than the previous one by c . If \overline{uv} is a

palindrome, then every new factor of $w'_t(\overline{uv})$ is increased by $2c$. So, the numerator can be computed from the initial weights and C .

We have left to show that the denominator is only dependent on v_0 and C . Note that the transition counts from v or v^* are a function of C and v_0 , because every event $v \rightarrow u$ is a transition from v , while every event $u^* \rightarrow v^*$ is followed by a transition from v^* , unless this is the final state, which is determined by v_0 . After every transition from v or v^* , we add a factor of $w''_t(v)$ or $w''_t(v^*)$ to the denominator. At any time t , these weights differ by c (if v is not a palindrome), but the factor added is always the smaller of the two. Between two transitions, each of these weights is reinforced by c , so consecutive factors differ by that amount. If v is a palindrome, there is no distinction between $w''_t(v)$ and $w''_t(v^*)$, and consecutive factors differ by $2c$. \square

LEMMA 3.6. *Suppose that in the reinforced random walk, we visit v and v^* in \mathcal{X}^r infinitely often a.s., and let τ_n be the n th time we visit either state. The process Y_{τ_n} is a mixture of Markov chains. Furthermore, if D_n is the ratio of the number of visits to v^* and v by τ_n , D_n converges a.s. to a finite limit D_∞ .*

PROOF. We claim that if Y_n is partially exchangeable, so is Y_{τ_n} . It is sufficient to show that the probability of a sequence Y_{τ_n} is invariant upon block transpositions, which generate the group of permutations that preserve transition counts ([8], Proposition 27). The probability of a path $v_{\tau_1}, \dots, v_{\tau_n}$ in Y_{τ_n} is the sum of the probabilities of all paths $v_0, v_1, \dots, v_{\tau_n}$ in Y_n that map to it. Denote this set of paths Θ . After a transposition of v -blocks or v^* -blocks, the probability of the path in Y_{τ_n} is equal to the sum of the probabilities of a different set of paths Θ' in Y_n . However, it is easy to see that this transposition of v -blocks or v^* -blocks defines a bijection from Θ to Θ' , and the probability of each path and its transposition is the same, because Y_n is partially exchangeable. Therefore, Y_{τ_n} is partially exchangeable. Furthermore, we assume that v and v^* are recurrent, so by de Finetti's theorem for Markov chains Y_{τ_n} is a mixture of Markov chains with a unique measure μ on the space of 2 by 2 transition matrices [8]. Note that both states are recurrent with probability 1, so the subset of transition matrices where one of the states is transient has μ -measure zero. This implies that μ -a.s. the transition matrix is irreducible, and since the state space is finite, both states are positive-recurrent. Therefore, D_n converges a.s. to a finite limit. \square

PROPOSITION 3.7. *The reinforced random walk of order r traverses every edge $v \rightarrow u$ with $w(\overline{uv}) > 0$ infinitely often, almost surely.*

PROOF. As \mathcal{X} is finite, we must visit at least one state in \mathcal{X}^r infinitely often, so without loss of generality, let this state be v . Let τ_n be the n th time we visit v , and \mathcal{F}_n be $\sigma(Y_1, \dots, Y_{\tau_n})$. For u with v, u admissible and $w(\overline{vu}) > 0$, let A_n be

the event that $Y_{\tau_{n+1}} = u$. Also, let $p_n = Q_{w,v_0}(A_n|\mathcal{F}_n)$. By Lévy’s extension of the Borel–Cantelli lemma (Lemma A.2),

$$\lim_{n \rightarrow \infty} \frac{\sum_{m=1}^n 1_{A_m}}{\sum_{m=1}^n p_m} = 1 \quad \text{on} \quad \left\{ \sum_{m=1}^{\infty} p_m = \infty \right\}.$$

Therefore, to show that the transition $v \rightarrow u$ is observed infinitely often with probability 1, it is sufficient to show that $\sum_m p_m = \infty$ a.s. The conditional probability p_m is just $w'_{\tau_m}(\overline{vu})/w''_{\tau_m}(v)$. Let $B_{m,k}$ be the event that we observe v^* fewer than km times between τ_1 and τ_m . On $B_{m,k}$, we can lower-bound p_m using the minimum possible value of $w'_{\tau_m}(\overline{vu})$, which is its initial value, and the maximum possible value of $w''_{\tau_m}(v)$, which is $(k + 1)mc$. Thus,

$$\begin{aligned} p_m &= Q_{w,v_0}(A_m \cap B_{m,k}|\mathcal{F}_n) + Q_{w,v_0}(A_m \cap B_{m,k}^C|\mathcal{F}_n) \\ &\geq \mathbf{1}_{B_{m,k}} \frac{w'_{\tau_1}(\overline{vu})}{(k + 1)mc}. \end{aligned}$$

Now, consider the event $\{D_\infty < N\}$. On this set, for any $k > N$, we will be in $B_{m,k}$ for all but finitely many m , which implies $\sum_m p_m = \infty$, by the previous inequality. But, by Lemma 3.6 we have $Q_{w,v_0}\{D_\infty < \infty\} = 1$, so noting $\{D_\infty < \infty\} = \bigcup_{N \in \mathbb{N}} \{D_\infty < N\}$ we conclude that $\sum_m p_m = \infty$ Q_{w,v_0} -a.s., and A_m happens infinitely often. Since w defines an irreducible Markov chain, the proposition follows by induction. \square

Propositions 3.7 and 3.5 are sufficient to show by de Finetti’s theorem for Markov chains [8] that the reinforced random walk of order r is a mixture of Markov chains on \mathcal{X}^r , or

$$(6) \quad Q_{w,v_0}(v_0, \dots, v_n) = \int_{\mathcal{T}} P_{v_0}^T(v_0, \dots, v_n) d\phi_{w,v_0}(T),$$

where $P_{v_0}^T$ is the distribution of a Markov chain started at v_0 and parametrized by the matrix T , \mathcal{T} is the space of $\mathcal{X}^r \times \mathcal{X}^r$ stochastic matrices and ϕ_{w,v_0} is a unique measure on the Borel subsets of this space. Let $\mathcal{T}' \subseteq \mathcal{T}$ be the set of matrices that represent irreducible, reversible Markov chains of order r .

PROPOSITION 3.8. *The reinforced random walk of order r is a mixture of reversible Markov chains of the same order, or $\phi_{w,v_0}(\mathcal{T}') = 1$.*

PROOF. This is a special case of Proposition 4.6. \square

4. Variable-order, reversible Markov chains. The number of parameters of a Markov chain grows as $|\mathcal{X}|^r$ with the order, r , which renders higher-order models impractical in many statistical applications. In this section, we investigate a family of models with finite memory length which do not suffer from this curse of dimensionality.

DEFINITION 4.1. A *variable-order Markov chain* is a Markov chain of order r with the constraint that for every *history* h in the set $\mathcal{H} \subseteq \{v \in \mathcal{X}^q : q < r\}$, if two states $u, u' \in \mathcal{X}^r$ both end in h , the transition probabilities $p(v|u)$ and $p(v|u')$ are equal for every $v \in \mathcal{X}^r$.

In essence, this is a discrete process which upon reaching a sequence $h \in \mathcal{H}$ loses memory of what preceded it. When \mathcal{H} is empty, we recover a general Markov chain of order r . Variable-order Markov chains have proven useful in applications where there is long memory only in certain directions. The literature on the subject can be traced to Rissanen [15] and Weinberger [17], who developed tree-based algorithms for estimating the set of histories efficiently in the context of compression. Bühlmann and Wyner proved several consistency results on these algorithms [7], and the former later addressed the problem of model selection [6]. For an evaluation of different algorithms in applications, see [4].

It is worth noting that MacQueen mentioned variable-order Markov chains in an unpublished abstract. However, there is a marked difference between his definition and Bühlmann and Wyner's, which relates to the closure properties of \mathcal{H} . MacQueen requires that if h is in \mathcal{H} , then so are all the sequences that begin with h . Intuitively, this means that the process cannot recover memory once it is lost. Bühlmann and Wyner do not impose this constraint. However, this is guaranteed when the process is reversible.

PROPOSITION 4.2. Let $X_n, n \in \mathbb{N}$, be an irreducible, reversible, variable-order Markov chain with histories \mathcal{H} . If $h \in \mathcal{H}$, then h^* is also a history; additionally, any sequence that has h as a prefix is also in \mathcal{H} .

PROOF. Let P_π be the stationary law of the chain. If $h \in \mathcal{H}$, then for any pair $a, b \in \mathcal{X}^q$, where q and the length of h sum to r , $P_\pi(X_1, \dots, X_{r+q} = ahb | X_1, \dots, X_r = ah)$ is independent of a , or

$$\frac{P_\pi(ahb)}{P_\pi(ah)} = C \quad \forall a \in \mathcal{X}^q.$$

This implies

$$\frac{P_\pi(hb)}{P_\pi(h)} = \frac{\sum_{a \in \mathcal{X}^q} P_\pi(ahb)}{\sum_{a \in \mathcal{X}^q} P_\pi(ah)} = \frac{\sum_{a \in \mathcal{X}^q} P_\pi(ah)C}{\sum_{a \in \mathcal{X}^q} P_\pi(ah)} = \frac{P_\pi(ahb)}{P_\pi(ah)}.$$

Using the fact that P_π is invariant upon time reversal and rearranging factors, we obtain

$$\frac{P_\pi(b^*h^*a^*)}{P_\pi(b^*h^*)} = \frac{P_\pi(h^*a^*)}{P_\pi(h^*)}.$$

The left-hand side is equal to $P_\pi(X_1, \dots, X_{r+q} = b^*h^*a^* | X_1, \dots, X_r = b^*h^*)$, which by the previous identity is independent of b^* . As this is true for any $a \in \mathcal{X}^q$,

h^* must be a history in \mathcal{H} . To prove the second part of the statement, suppose h is a prefix of g . Since h^* is in \mathcal{H} , and g^* ends in h^* , then by definition $g^* \in \mathcal{H}$. Using the first result, we conclude that $g \in \mathcal{H}$. \square

We will define a reinforcement scheme, which like the one in the previous section is recurrent, partially exchangeable and, by de Finetti’s theorem, a mixture of Markov chains. But, in this case, the mixing measure is restricted to the variable-order, reversible Markov chains with a fixed set of histories \mathcal{H} . As before, we begin with a stationary, reversible function w , an initial state $v_0 \in \mathcal{X}^r$, and a palindromic sequence β that starts with v_0 . Let the function $f : \mathcal{X}^r \mapsto \mathcal{H}$ map any sequence to its shortest ending in \mathcal{H} .

DEFINITION 4.3. The *variable-order, reinforced random walk* is a stochastic process $Z_n, n \in \mathbb{N}$, on \mathcal{X}^r with measure H_{w,v_0} . The initial state is v_0 with probability 1. For any admissible path v_0, \dots, v_n , the conditional transition probability

$$H_{w,v_0}(Z_{n+1} = u | Z_0 = v_0, \dots, Z_n = v_n) = \frac{w'_n(\overline{f(v_n)u})}{w''_n(f(v_n))}$$

whenever v_n, u is admissible and zero otherwise.

REMARK 4.4. This process is a reinforced circuit process, just like the one defined in Remark 3.3, with the difference that in computing the transition probabilities, instead of taking the current state to be the sequence $v_n \in \mathcal{X}^r$, we let it be the shortest ending of v_n in \mathcal{H} , or $f(v_n)$.

PROPOSITION 4.5. *The variable-order, reinforced random walk is partially exchangeable in the sense of Diaconis and Freedman.*

This proof is deferred to the [Appendix](#). One can show that this process is recurrent following the same argument of Proposition 3.7. In the proof of Proposition 3.7, we use a shortest history h in place of v , and Lemma 3.6 still holds for h and h^* . Recurrence and partial exchangeability imply

$$(7) \quad H_{w,v_0}(v_0, \dots, v_n) = \int_{\mathcal{T}} P_{v_0}^T(v_0, \dots, v_n) d\psi_{w,v_0}(T)$$

for a unique measure ψ_{w,v_0} characterized by the function w , and the initial state, in addition to the parameters β, c and \mathcal{H} , which we keep fixed. In the [Appendix](#), we show that ψ_{w,v_0} is restricted to the reversible, variable-order Markov chains with histories \mathcal{H} .

PROPOSITION 4.6. *Let $T'' \subseteq \mathcal{T}$ be the set of transition matrices representing an irreducible, reversible, variable-order Markov chain where every $h \in \mathcal{H}$ is a history. Then, $\psi_{w,v_0}(T'') = 1$.*

5. Bayesian analysis. In Section 3, we defined a family of measures in the space of order- r , reversible Markov chains, and in Section 4 we extended it to variable-order, reversible Markov chains. In the following, we will show that these distributions are conjugate priors for a Markov chain of order r . We discuss properties of the prior relevant to Bayesian analysis, such as a natural sampling algorithm and closed-form expressions for some important moments.

DEFINITION 5.1. Consider a variable-order, reinforced random walk $Z_n, n \in \mathbb{N}$, with distribution H_{w, v_0} and take any admissible path $e = v_0, \dots, v_n$. We define $Z_n^{(e)}, n \in \mathbb{N}$, to be the process with law

$$\begin{aligned} H_{w, v_0, e}(v_n, u_1, \dots, u_m) \\ = H_{w, v_0}(Z_{n+1} = u_1, \dots, Z_{n+m+1} = u_m | Z_1 = v_1, \dots, Z_m = v_m). \end{aligned}$$

In words, $Z^{(e)}$ is the continuation of a variable-order reinforced random walk after traversing some fixed path e . We can rewrite the law

$$(8) \quad H_{w, v_0, e}(v_n, u_1, u_2, \dots, u_m) = \frac{H_{w, v_0}(v_1, \dots, v_n, u_1, \dots, u_m)}{H_{w, v_0}(v_1, \dots, v_n)},$$

which makes it evident that $Z^{(e)}$ is partially exchangeable, because for a fixed e , the numerator only depends on the transition counts in v_n, u_1, \dots, u_m , while the denominator is constant. It is also not hard to see that the process visits every state infinitely often with probability 1. Therefore, by de Finetti’s theorem for Markov chains, it is a mixture of Markov chains with a mixing measure that will be denoted $\psi_{w, v_0, e}$.

PROPOSITION 5.2. Suppose we model a process $W_n, n \in \mathbb{N}$, as a reversible, variable-order Markov chain with histories $\mathcal{H} \subseteq \{v \in \mathcal{X}^q : q < r\}$, and we assign a prior ψ_{w, v_0} to the transition probabilities, T . Given an observed path, $e = v_0, \dots, v_n$, the posterior probability of T is $\psi_{w, v_0, e}$. In consequence, the family of measures

$$\mathcal{D} = \{\psi_{w, v_0, e} : e \text{ an admissible path starting in } v_0\}$$

is closed under sampling.

PROOF. Consider the event $W_n = v_n, W_{n+1} = u_1, \dots, W_{n+1+m} = u_m$. By Bayes rule, the posterior probability of this event given the observation is the prior probability of $W_1 = v_1, \dots, W_n = v_n, W_{n+1} = u_1, \dots, W_{n+1+m} = u_m$ divided by the prior probability of $W_1 = v_1, \dots, W_n = v_n$. By equation (8), this posterior is equal to $H_{w, v_0, e}$. Let $\rho(T)$ be the posterior distribution of T given the observation, then for any u_1, \dots, u_m and any $m > 0$,

$$H_{w, v_0, e}(v_n, u_1, \dots, u_m) = \int_{\mathcal{T}} P_{v_n}^T(v_n, u_1, \dots, u_m) d\rho(T).$$

By de Finetti’s theorem for Markov chains, the mixing measure $\psi_{w,v_0,e}$ is unique; therefore, we must have $\rho = \psi_{w,v_0,e}$. \square

In the next proposition, we show that the variable-order, reinforced random walk may be used to simulate from the conjugate prior ψ_{w,v_0} (or using a similar argument, a posterior of the form $\psi_{w,v_0,e}$). Let $\{V^{(i)} = v_1^{(i)}, v_2^{(i)}, \dots, v_n^{(i)}\}_{i \in \{1, \dots, k\}}$ be independent samples of the reinforced random walk with initial parameters w and v_0 . For any sequence $u \in \mathcal{X}^{r+1}$, consider the random variable $n^{-1}w'_n(u)$, the weight defined in equation (4) for a sample path with distribution H_{w,v_0} , normalized by the path’s length. Define the empirical estimate, $n^{-1}w'_{n,k}(u)$, to be the mean of this random variable evaluated at the paths $\{V^{(i)}\}_{i \in \{1, \dots, k\}}$. Also, let P_π^T be the stationary law of an order- r Markov chain with transition probabilities T . We have seen that $\{P_\pi^T(u) : u \in \mathcal{X}^{r+1}\}$ has a one-to-one correspondence with T .

PROPOSITION 5.3. For any bounded, real-valued function $g(P_\pi^T(\cdot))$,

$$(9) \quad \lim_{n \rightarrow \infty} \lim_{k \rightarrow \infty} g(\{n^{-1}w'_{n,k}(u) : u \in \mathcal{X}^{r+1}\}) \stackrel{a.s.}{=} \int_T g(P_\pi^T) d\psi_{w,v_0}(T).$$

PROOF. The empirical estimate $g(\{n^{-1}w'_{n,k}(u) : u \in \mathcal{X}^{r+1}\})$ is the average of i.i.d. observations, so by the strong law of large numbers, w.p.1,

$$\lim_{k \rightarrow \infty} g(\{n^{-1}w'_{n,k}(u) : u \in \mathcal{X}^{r+1}\}) = H_{w,v_0}[g(\{n^{-1}w'_n(u) : u \in \mathcal{X}^{r+1}\})],$$

where the right-hand side is the expectation in a reinforced random walk with parameters w, v_0 . In the proof of Proposition 4.6, we showed that $w'_n(u)$ converges H_{w,v_0} -a.s. Taking the limit as $n \rightarrow \infty$, by dominated convergence,

$$\begin{aligned} & \lim_{n \rightarrow \infty} \lim_{k \rightarrow \infty} g(\{n^{-1}w'_{n,k}(u) : u \in \mathcal{X}^{r+1}\}) \\ &= H_{w,v_0} \left[\lim_{n \rightarrow \infty} g(\{n^{-1}w'_n(u) : u \in \mathcal{X}^{r+1}\}) \right]. \end{aligned}$$

Conditional on a variable T measurable on its tail σ -field with distribution ψ_{w,v_0} , the reinforced random walk is a Markov chain with law $P_{v_0}^T$. We know $w'_n(u)$ converges $P_{v_0}^T$ -a.s. to $P_\pi^T(u)$, so equation (9) follows. \square

Several moments of H_{w,v_0} have closed-form expressions. In particular, the mean likelihood $P_{v_0}^T$ of any path beginning in v_0 is just the probability of the path in the reinforced random walk by equation (7). From the proof of Proposition 4.5, one can deduce a closed-form expression for the law of the variable-order reinforced random walk as a function of the transition counts in a path (see Supplement [2]). From a realization of the transition counts as a path, one can also compute the law H_{w,v_0} by modeling a random walk with reinforcement.

The expectation of cycle probabilities with a prior ψ_{w,v_0} on T may also be computed exactly.

PROPOSITION 5.4. For any cyclic path v, v_1, \dots, v_n, v , not necessarily including v_0 , the expectation of $P_v^T(v, v_1, \dots, v_n, v)$ with prior ψ_{w, v_0} on T has a closed-form expression, provided $w'_0(u)$ is greater than $3c$ for all $u \in \mathcal{X}^{r+1}$.

PROOF. Find the shortest cycle v, \dots, v_0, \dots, v with positive weight w . Then, for any transition matrix T in the support of ψ_{w, v_0} , we have

$$(10) \quad P_v^T(v, v_1, \dots, v_n, v) = \frac{P_{v_0}^T(v, v_1, \dots, v_n, v, \dots, v_0, \dots, v)}{P_{v_0}^T(v, \dots, v_0, \dots, v)}.$$

Taking the expectation with a measure ψ_{w, v_0} on T , we obtain

$$\begin{aligned} & \int_T P_v^T(v, v_1, \dots, v_n, v) d\psi_{w, v_0}(T) \\ &= \int_T \frac{P_{v_0}^T(v, v_1, \dots, v_n, v, \dots, v_0, \dots, v)}{P_{v_0}^T(v, \dots, v_0, \dots, v)} d\psi_{w, v_0}(T). \end{aligned}$$

By Bayes theorem, the product of the likelihood $P_{v_0}^T(v, v_1, \dots, v_n, v, \dots, v_0, \dots, v)$ and the prior $d\psi_{w, v_0}(T)$ is equal to the marginal prior probability of the path $v, v_1, \dots, v_n, v, \dots, v_0, \dots, v$ times the posterior of T :

$$\begin{aligned} & \int_T P_v^T(v, v_1, \dots, v) d\psi_{w, v_0}(T) \\ &= H_{w, v_0}(v, v_1, \dots, v, \dots, v_0, \dots, v) \int_T \frac{1}{P_{v_0}^T(v, \dots, v_0, \dots, v)} d\psi_{w_p, v_0}(T), \end{aligned}$$

where w_p are the weights parametrizing the posterior of T given the path $v, v_1, \dots, v, \dots, v_0, \dots, v$. To solve the integral on the right-hand side, let us rewrite it using Bayes theorem and equation (7),

$$\begin{aligned} & H_{w_{pp}, v_0}^{-1}(v, \dots, v_0, \dots, v, \dots, v_0, \dots, v) \\ & \times \int_T \frac{P_{v_0}^T(v, \dots, v_0, \dots, v, \dots, v_0, \dots, v)}{P_{v_0}^T(v, \dots, v_0, \dots, v)} d\psi_{w_{pp}, v_0}, \end{aligned}$$

where w_{pp} are the weights w_p reduced by the cycle $v, \dots, v_0, \dots, v, \dots, v_0, \dots, v$. These weights are positive because of the assumption $w'_0(u) > 3c$ for all u , which could certainly be relaxed in some cases. Applying equations (7) and (10) once more, the last expression becomes

$$H_{w_{pp}, v_0}^{-1}(v, \dots, v_0, \dots, v, \dots, v_0, \dots, v) H_{w_{pp}, v_0}(v, \dots, v_0, \dots, v),$$

which completes our derivation. \square

The ability to compute these expectations exactly makes it possible to use Bayes factors for model comparison [11]. Given some data \mathbf{X} and two probabilistic models, where each model i has a prior measure $P^{(i)}$ and parameters θ_i , a Bayes factor

quantifies the relative odds between them. It is formally defined as,

$$(11) \quad \frac{P^{(1)}(\mathbf{X})}{P^{(2)}(\mathbf{X})} = \frac{\int P^{(1)}(\mathbf{X}|\theta_1) dP^{(1)}(\theta_1)}{\int P^{(2)}(\mathbf{X}|\theta_2) dP^{(2)}(\theta_2)},$$

the ratio between the marginal probabilities of the data under each model. Each marginal probability is sometimes referred to as the *evidence* for the corresponding model. Diaconis and Rolles apply Bayes factors to compare a number of models on different data sets. They consider reversible Markov chains, general Markov chains, and i.i.d. models [10], assigning conjugate priors which facilitate computing the marginal probabilities in equation (11).

The conjugate priors introduced here facilitate similar comparisons, where the family of models under consideration is expanded to include reversible Markov chains that differ in their length of memory. For some data \mathbf{X} , one can define two variable-order reversible Markov models, with different histories, $\mathcal{H}^{(1)}$ and $\mathcal{H}^{(2)}$. In each case, we assign a conjugate prior, $\psi_{w,v_0}^{(1)}$ and $\psi_{w,v_0}^{(2)}$, respectively, to the transition probability matrix. To make the prior uninformative in some sense we could set w to be uniform for all $u \in \mathcal{X}^{r+1}$ and let β be the shortest palindrome starting with v_0 , for example. The constant c is set to 1. The Bayes factor is then

$$\frac{P^{(1)}(\mathbf{X})}{P^{(2)}(\mathbf{X})} = \frac{\int_{\mathcal{T}} P_{v_0}^T(\mathbf{X}) d\psi_{w,v_0}^{(1)}(T)}{\int_{\mathcal{T}} P_{v_0}^T(\mathbf{X}) d\psi_{w,v_0}^{(2)}(T)}.$$

We have seen that the expectations on the right-hand side can be computed exactly when \mathbf{X} is a path starting at v_0 or any cyclic path. In the following example, we apply this test to finite data sets simulated from a lumped Markov chain.

EXAMPLE 5.5 (Order estimation for a lumped reversible Markov chain). A random walk was simulated on the 9-state graph shown in Figure 4, from which we omitted self-edges on every state, all weighted by 1. The observation was lumped into the 3 *macrostates* separated by the dashed lines. This is meant

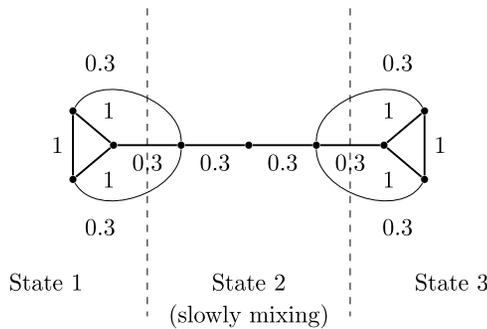


FIG. 4. A lumped reversible Markov chain.

to illustrate a natural experiment, where the difference between the states within each macrostate is obscured by the measurement. From the resulting sequence, we take the initial macrostate and every 7th macrostate thereafter to form a path \mathbf{X} of length 1000 in $\mathcal{X} = \{1, 2, 3\}$.

We test 4 reversible Markov models, that differ in the length of memory:

1. A first-order, reversible Markov chain.
2. A second-order, reversible Markov chain.
3. A variable-order model with maximum order 2, where states 1 and 3 are histories. Intuitively, only state 2 has “memory.”
4. A variable-order model with maximum order 2, where states 2 and 3 are histories. Intuitively, only state 1 has “memory.”

For each model i , we assign a prior $\psi_{w,v_0}^{(i)}$ to the transition matrix, where v_0 is the initial state in \mathbf{X} , $w(u) = 2$ for all $u \in \mathcal{X}^3$ and β is the shortest palindrome starting with v_0 . We compared the 4 models using 50 independent realizations of the lumped Markov chain and found that model 3 had the highest evidence in 72% of the cases, while model 2 was selected in all the remaining cases. In Figure 5, we report a boxplot of the logarithm of the Bayes factors comparing models 1, 2, and 4 against model 3.

This represents compelling evidence for model 3. The result is not entirely surprising given that this model gives memory to state 2, which is slowly mixing, as indicated in Figure 4. The fact that the most complex model (model 2) is not nec-

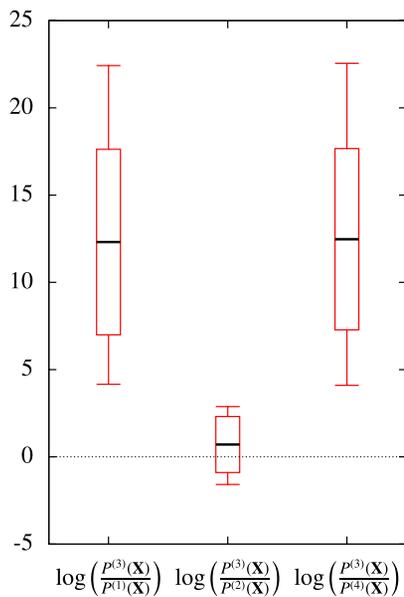


FIG. 5. Boxplot of logarithmic Bayes factors computed from 50 independent datasets.

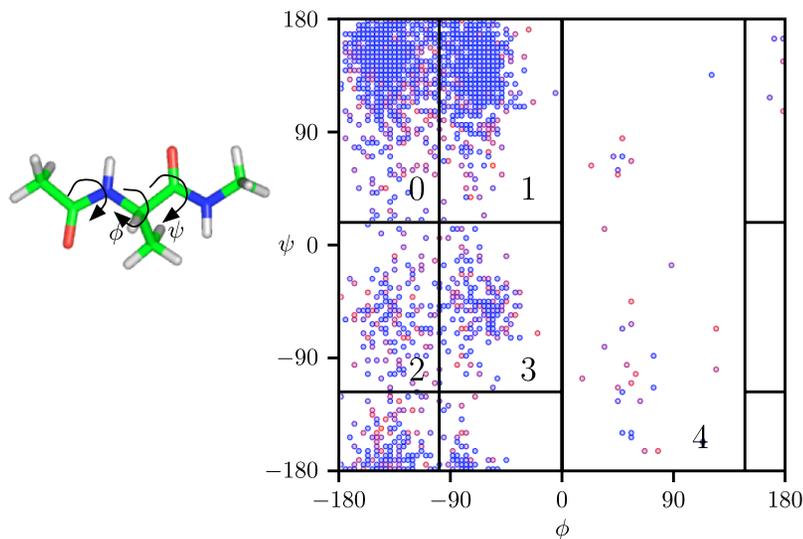


FIG. 6. The structure of *Ace-Ala-Nme* is described by two dihedral angles, ϕ and ψ . The periodic map on the right shows a partition of conformational space into 5 states. The colored markers indicate the free energy of bins centered at each point, which reveals the metastable nature of this molecule's dynamics.

essarily selected showcases the automatic penalty for model complexity in Bayes factors.

We conclude this section with two applications of Bayesian analysis of reversible Markov chains to molecular dynamics (MD). An MD simulation approximates the time-reversible dynamics of a molecule in solvent. The trajectories produced by a simulation are discretized in space and time.

EXAMPLE 5.6. The terminally blocked alanine dipeptide, shown in Figure 6, is a common test system for Markov models of MD. The conformational space of the molecule, which is represented in the figure in a two-dimensional projection, is partitioned into 5 states. The states are believed to be metastable due to the basins that characterize the free-energy function, also plotted in the figure. This metastability allows one to approximate the dynamics of the molecule, projected onto the partition, as a reversible Markov chain. The approximation will be good when the discrete time interval at which a trajectory is sampled is larger than the timescale for equilibration within every state, but smaller than the timescale of transitions.

Few statistical validation methods are available for Markov models of MD. Bacallado, Chodera and Pande used a Bayesian hypothesis test to compare different partitions of conformational space [3]. Here, we apply Bayes factors to test a first-order Markov model on a fixed partition, by comparing it to second-order and

TABLE 1

Molecular dynamics simulation of the alanine dipeptide. The entries in the table are the transition counts $(x_1, x_2) \rightarrow (x_2, x_3)$ in the trajectory \mathbf{X} , which has initial state $(0, 4)$

x_1	x_2	x_3					x_1	x_2	x_3				
		0	1	2	3	4			0	1	2	3	4
0	0	261	187	13	2	0	3	0	5	13	2	0	0
	1	188	144	13	11	0		1	5	4	2	1	0
	2	12	4	9	15	0		2	4	3	16	5	0
	3	5	1	0	1	0		3	2	5	3	3	0
	4	1	0	0	0	0		4	0	0	0	0	0
1	0	180	143	22	5	0	4	0	1	0	0	0	0
	1	141	125	5	5	0		1	0	0	0	0	0
	2	4	3	10	4	0		2	0	0	0	0	0
	3	4	1	10	3	0		3	0	0	0	0	0
	4	0	0	0	0	0		4	0	0	0	0	0
2	0	16	13	3	0	0							
	1	16	4	1	1	0							
	2	12	12	37	11	0							
	3	9	5	15	6	0							
	4	0	0	0	0	0							

variable-order models on the same partition. The data \mathbf{X} are the transition counts in a single MD trajectory of 1767 steps sampled at an interval of 6 picoseconds, as recorded in Table 1. The prior parameters w and β are the same as in the previous example. The results of the model comparison are summarized in the following table.

Model (i)	$\log P^{(i)}(\mathbf{X})$
First order	-1846
Variable order 0	-1824
Variable order 1	-1825
Variable order 2	-1844
Variable order 3	-1846
Variable order 4	-1847
Second order	-1800

The state describing each variable order model is the only state in the model that has a memory of length 2 (the only state that is not a history). There seems to be substantial evidence in favor of a second-order model. Adding memory to states seen in a large number of transition makes a bigger difference, as expected. This result is in accordance with certain exploratory observations which indicate that at

the timescale of 6 picoseconds, the effect of water around the molecule, neglected in our state definitions, persists.

EXAMPLE 5.7. The alanine pentapeptide is a longer polymer that exhibits a higher degree of structural and dynamical complexity. Buchete and Hummer partition the conformational space of the molecule into 32 states by chemical conventions [5]. An MD trajectory¹ in conformational space was projected onto this partition, and an exploratory analysis suggested that the effects of memory decay after 500 picoseconds. Accordingly, we take a conformation from the trajectory every 500 picoseconds to form a sequence \mathbf{X} of 1885 steps in $\mathcal{X} = \{0, \dots, 31\}$.

As in previous examples, we tested models with varying lengths of memory. Each model was assigned a conjugate prior, this time setting $w(u) = 1/32$ for all $u \in \mathcal{X}^3$. Of all the variable-order models where a single state has a memory of length 2 and all others are histories, we found that only 4 models were strongly selected over a first-order model. In the following table, we show the logarithm of the evidence for each of these models, a first-order model and a variable-order model that gives a memory of length 2 to all 4 states.

Model (i)	$\log P^{(i)}(\mathbf{X})$
First order	-4090.0
Variable order 14	-4015.5
Variable order 15	-3814.5
Variable order 30	-3860.3
Variable order 31	-3301.6
Variable order 14, 15, 30, 31	-2964.3

This represents compelling evidence for a model that gives memory to states 14, 15, 30 and 31. It is interesting to contrast inferences based on this model to those based on a first-order Markov model. To do this, we computed 1000 approximate posterior samples of the transition matrix in each case. This was done by simulating a reinforced random walk, which is a mixture of variable-order Markov chains with the posterior distribution of T as a mixing measure (see Proposition 5.3). The reinforced random walk was simulated 10^7 steps to obtain each sample.

In Figure 7, we histogram stationary probabilities of the transition matrices sampled from the posterior. In particular, we show plots for the stationary probabilities of states 14, 15, 30 and 31. In the variable-order model, we define $\pi(x) = \sum_{y \in \mathcal{X}} \pi(xy)$. The inferences of each model in this case are very similar.

The largest eigenvalues of the transition matrix are also of interest because they are related to different modes of relaxation. Each eigenvalue λ is associated with a timescale $-\tau_{\text{lag}}/\log \lambda$, which is useful in exploratory analysis. Here, τ_{lag} is the

¹Simulated with the Amber-GSs forcefield at 300K in explicit solvent.

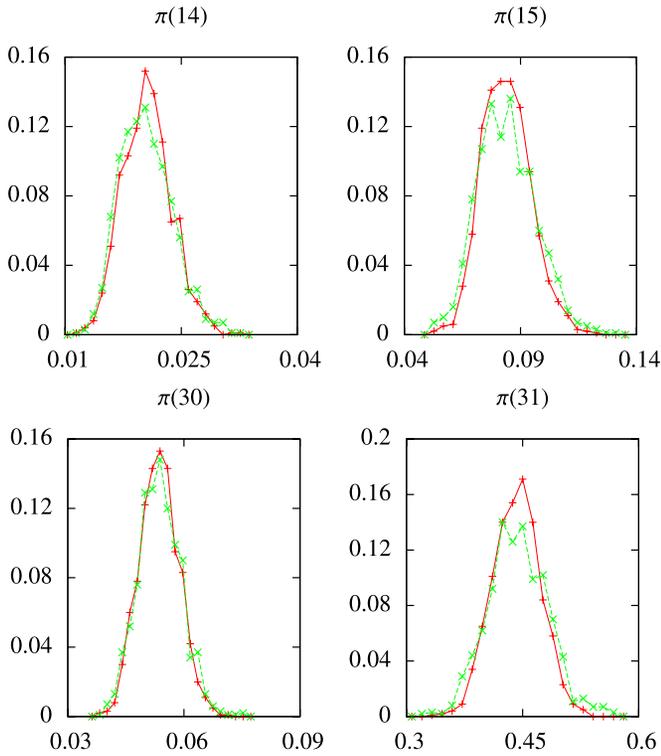


FIG. 7. Histograms of 1000 posterior samples of the stationary probabilities of states 14, 15, 30, 31. The red solid lines correspond to the first-order Markov model, and the green dashed lines to the variable-order Markov model that gives a memory of length 2 to states 14, 15, 30 and 31.

length in time of one step of the Markov chain, or 500 picoseconds. In Figure 8, we histogram posterior samples of the three largest nonunit eigenvalues and their associated timescales. In this case, the inferences of each model are quite different, with the variable-order model predicting larger eigenvalues and timescales.

6. Conclusions. We define a reinforcement scheme for the higher-order, reversible Markov chain that extends the ERRW on an undirected graph. Several properties of the ERRW, like recurrence and partial exchangeability, were shown to generalize to this process. Other properties may also generalize but were not pursued here. In particular, we can mention the uniqueness results of Johnson [18] and Rolles [16], and the fact that mixtures of measures in \mathcal{D} are weak-star dense in the space of all priors [10].

The reinforced random walk leads to a conjugate prior that facilitates estimation and hypothesis testing of reversible processes in which the effects of memory decay after some time. Certain statistical problems remain a challenge, such as inferring the transition matrix with a *fixed* stationary distribution. In applications, it

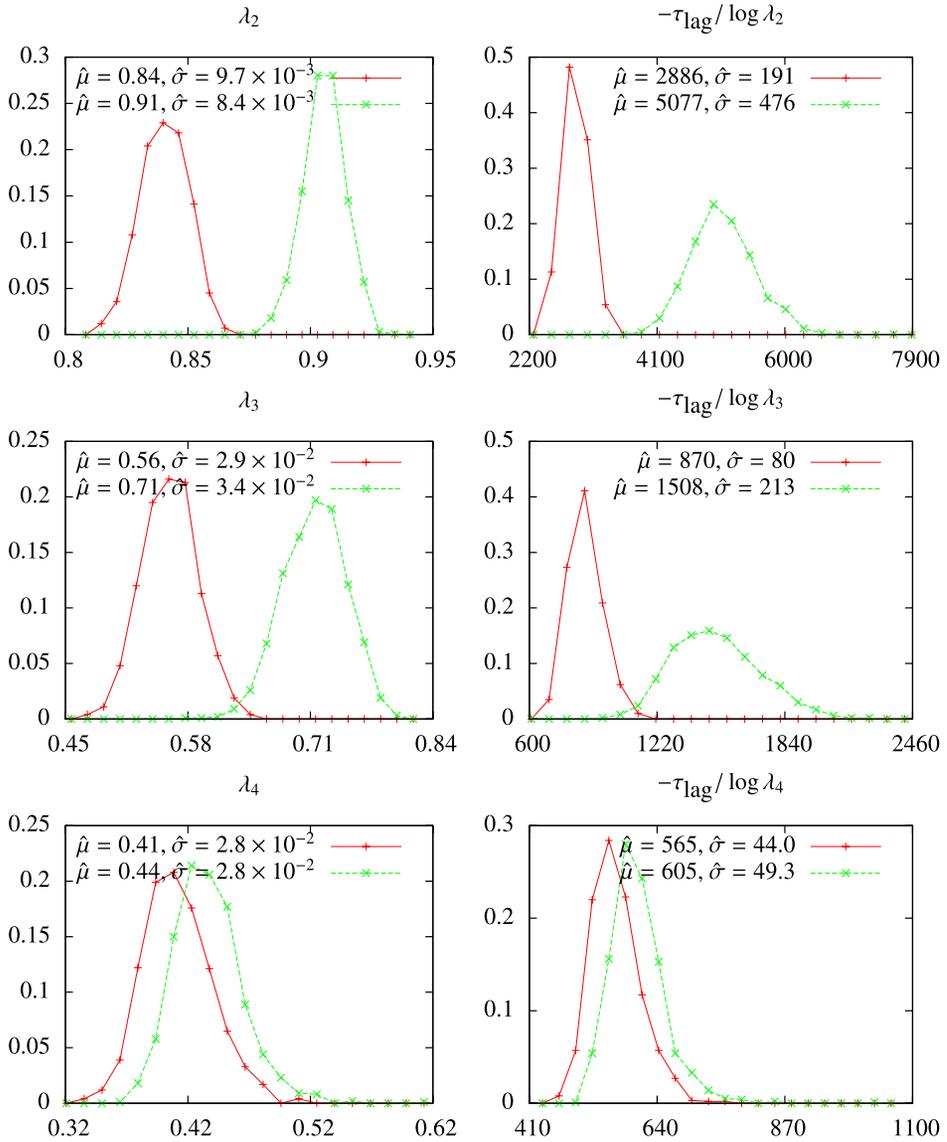


FIG. 8. Histograms of 1000 posterior samples of the second, third and fourth largest eigenvalues of the transition matrix, as well as the timescales associated with these eigenvalues. The red solid lines correspond to the first-order Markov model, and the green dashed lines to the variable-order Markov model that gives memory to states 14, 15, 30 and 31. In both cases, we compute the eigenvalues of the transition matrix for the process $V_n, n \in \mathbb{N}$, in \mathcal{X}^2 . All sample means $\hat{\mu}$ and standard deviations $\hat{\sigma}$ are shown.

will become important to evaluate the objectivity of the prior and to determine the optimal value of its parameters in this sense.

From a practical point of view, we only discussed Bayesian updating for data sets composed of a single Markov chain starting with probability 1 from the initial state v_0 used in the prior. Numerical algorithms are needed to perform inference with data sets composed of multiple chains. A starting point could be the method developed by Bacallado, Chodera and Pande to apply the prior of Diaconis and Rolles to first-order, reversible Markov chains [3].

APPENDIX

In the following, we use the notation defined in the first paragraph of Section 2.

PROPOSITION A.1 (Kolmogorov’s criterion). *Let $X_n, n \in \mathbb{N}$, be an irreducible order- r Markov chain with transition probabilities p . Then X_n is reversible if and only if for any cyclic admissible path $v_0, v_1, \dots, v_n, v_0$,*

$$(12) \quad p(v_1|v_0)p(v_2|v_1) \cdots p(v_0|v_n) = p(v_0^*|v_1^*)p(v_1^*|v_2^*) \cdots p(v_n^*|v_0^*).$$

PROOF. The “only if” statement is straightforward. By the definition of the stationary distribution and reversibility

$$\begin{aligned} p(v_1|v_0)p(v_2|v_1) \cdots p(v_0|v_n) &= \frac{P_\pi(\overline{v_0v_1})}{\pi(v_0)} \frac{P_\pi(\overline{v_1v_2})}{\pi(v_1)} \cdots \frac{P_\pi(\overline{v_nv_0})}{\pi(v_n)} \\ &= \frac{P_\pi(\overline{v_1^*v_0^*})}{\pi(v_0^*)} \frac{P_\pi(\overline{v_2^*v_1^*})}{\pi(v_1^*)} \cdots \frac{P_\pi(\overline{v_0^*v_n^*})}{\pi(v_n^*)} \\ &= p(v_0^*|v_1^*)p(v_1^*|v_2^*) \cdots p(v_n^*|v_0^*). \end{aligned}$$

To prove the “if” statement, choose an arbitrary state u ; then, for any v , since the chain is irreducible, there is an admissible path $u, v_1, v_2, \dots, v_n, v$ with positive probability. Define

$$(13) \quad \pi'(v) = B \frac{p(v_1|u)p(v_2|v_1) \cdots p(v|v_n)}{p(v_n^*|v^*)p(v_{n-1}^*|v_n^*) \cdots p(u^*|v_1^*)},$$

where B is a positive constant. Note that this expression does not depend on the sequence v_1, \dots, v_n chosen. Take a different sequence z_1, \dots, z_m . Let $t \in \mathcal{X}^r$ be a palindrome, then because the chain is irreducible, we can find a sequence v, t_1, t_2, \dots, t with positive probability, and it is easy to see from equation (12) that the palindrome $v, t_1, t_2, \dots, t, \dots, t_2^*, t_1^*, v^*$ has positive probability. We can construct another palindrome $u^*, s_1, s_2, \dots, s_2^*, s_1^*, u$ in the same way. Multiplying equation (13) by factors of 1,

$$\begin{aligned} B \frac{p(v_1, v_2, \dots, v|u)}{p(v_n^*, v_{n-1}^*, \dots, u^*|v^*)} &= B \frac{p(v_1, v_2, \dots, v|u)}{p(v_n^*, v_{n-1}^*, \dots, u^*|v^*)} \frac{p(t_1, t_2, \dots, v^*|v)}{p(t_1, t_2, \dots, v^*|v)} \\ &\quad \times \frac{p(z_m^*, z_{m-1}^*, \dots, u^*|v^*)}{p(z_1, z_2, \dots, v|u)} \frac{p(s_1, s_2, \dots, u|u^*)}{p(s_1, s_2, \dots, u|u^*)} \end{aligned}$$

$$\begin{aligned} &\times \frac{p(z_1, z_2, \dots, v|u)}{p(z_m^*, z_{m-1}^*, \dots, u^*|v^*)} \\ &= B \frac{p(z_1, z_2, \dots, v|u)}{p(z_m^*, z_{m-1}^*, \dots, u^*|v^*)}. \end{aligned}$$

The first four terms equal 1 because the numerator and denominator are the probabilities of the same cycle forward and backward, which are equal by equation (12). Now, we check that $\pi'(v)$ satisfies the reversibility conditions specified in the **Introduction**. First, we show that $\pi'(v) = \pi'(v^*)$. Take a path $u, z_1, \dots, z_\ell, v^*$ with positive probability, and the previously found palindrome $u^*, s_1, s_2, \dots, s_2^*, s_1^*, u$, then applying the same method,

$$\begin{aligned} \pi'(v) &= B \frac{p(v_1, v_2, \dots, v|u)}{p(v_n^*, v_{n-1}^*, \dots, u^*|v^*)} \\ &= B \frac{p(v_1, v_2, \dots, v|u)}{p(v_n^*, v_{n-1}^*, \dots, u^*|v^*)} \frac{p(s_1, s_2, \dots, u|u^*)}{p(s_1, s_2, \dots, u|u^*)} \\ &\quad \times \frac{p(z_\ell^*, z_{\ell-1}^*, \dots, u^*|v)}{p(z_1, z_2, \dots, v^*|u)} \frac{p(z_1, z_2, \dots, v^*|u)}{p(z_\ell^*, z_{\ell-1}^*, \dots, u^*|v)} \\ &= B \frac{p(z_1, z_2, \dots, v^*|u)}{p(z_\ell^*, z_{\ell-1}^*, \dots, u^*|v)} = \pi'(v^*). \end{aligned}$$

From this, and equation (13) we deduce that for any admissible $v, z, \pi'(v)p(z|v) = \pi'(z^*)p(v^*|z^*)$. Since the state space is finite, we can choose B such that π' sums to 1. We have shown that the weights $k_{v,z} \equiv \pi'(v)p(z|v)$ satisfy the conditions of a reversible random walk with memory, so by Proposition 2.3 the process with transition probabilities p represents a reversible, order- r Markov chain. \square

PROOF OF PROPOSITION 4.5. The probability $H_{w,v_0}(v_0, \dots, v_m)$ is a product of transition probabilities, to which the n th transition contributes a factor of

$$(14) \quad p_n = \frac{w'_{n-1}(\overline{f(v_{n-1})v_n})}{w''_{n-1}(f(v_{n-1}))}.$$

We know that $f(v_n)$ cannot be longer than $\overline{f(v_{n-1})v_n}$ by Proposition 4.2; let $L(v_{n-1}, v_n)$ be the set of histories of v_n that are shorter than $f(v_{n-1})$. If this set is nonempty, let us multiply equation (14) by factors of 1, to obtain the following factor for the n th transition:

$$(15) \quad p_n = \frac{w'_{n-1}(\overline{f(v_{n-1})v_n})}{w''_{n-1}(f(v_{n-1}))} \prod_{z \in L(v_{n-1}, v_n)} \frac{w'_{n-1}(z^+)}{w''_{n-1}(z^+)},$$

where z^+ is the ending of v_n that is longer than z by 1. The added factor equals 1 because, if $w'_{n-1}(z^+) \neq w''_{n-1}(z^+)$, then $f(v_{n-1})$ must end on z , which by definition is a history shorter than $f(v_{n-1})$, a contradiction.

Consider all the possible factors in the numerator of $H_{w,v_0}(v_0, \dots, v_m)$. Take any $h \in \mathcal{H}$ that is *minimal*, meaning that it does not end in another history. For any $a \in \mathcal{X}$, we will see a factor $w'(ha)$ after every transition through ha . The conjugate factor $w'(ah^*)$ will appear every time we go through ah^* , because:

- If $A(ah^*) \in \mathcal{H}$, it is minimal by the closure properties of \mathcal{H} , so $w'(ah^*)$ will be the numerator of the first factor in equation (15).
- Otherwise, the minimal history in the transition ending in ah^* will be longer than h^* , and there will be an added factor in equation (15) with $w'(ah^*)$ in the numerator. Conversely, note that the factor $w'(ah^*)$ is only added to the numerator of equation (15) when we go through ah^* for some minimal h , because we required that $h^* \in L(v_{n-1}, v_n)$, so $h \in \mathcal{H}$ and does not end in another history.

As in the proof of Proposition 3.5, we argue that every new factor $w'(ha)$ or $w'(ah^*)$ is increased by c with respect to the previous one (or by $2c$ if ha is a palindrome). Therefore, the numerator of $H_{w,v_0}(v_0, \dots, v_m)$ is only a function of the transition counts and the initial state.

Finally, consider all the factors in the denominator of $H_{w,v_0}(v_0, \dots, v_m)$. Take any minimal history h . We will see a factor $w''(h)$, for every transition through h . The conjugate factor $w''(h^*)$ will appear every time we go through h^* , because:

- If h^* is also minimal, then $w''(h^*)$ will be in the denominator of the first factor in equation (15).
- Otherwise, we know that $A(h^*)$ is not a history, so the transition ending in h^* must have a history at least as long as h^* , which is longer than the history $\Omega(h^*)$. So, $w''(h^*)$ will appear in the denominator of a factor added in equation (15). Conversely, we only add factors of $w''(h^*)$ to the denominator of equation (15) when we go through h^* for a minimal h , because we required $\Omega(h^*) \in L(v_{n-1}, v_n)$ which implies h minimal.

As before, every new factor $w''(h)$ or $w''(h^*)$ will be increased by c with respect to the previous one (or by $2c$ if h is a palindrome). Therefore, the denominator is a function of the transition counts and the initial state, and the process is partially exchangeable. \square

PROOF OF PROPOSITION 4.6. Let $\vec{C}_n(u, v)$ be the transition counts from u to v in the first n steps of a stochastic process on \mathcal{X}^r . Also, define $\vec{C}_n(u) \equiv \sum_{v \in \mathcal{X}^r} \vec{C}_n(u, v)$, which counts the visits to u . Remember \mathcal{T}'' is the set of irreducible transition matrices for variable-order, reversible Markov chains where all $h \in \mathcal{H}$ are histories. Define the event D , that the set $\{\vec{C}_n(u, v)/\vec{C}_n(u) : \forall u, v \text{ admissible}\}$ converges to a transition probability matrix in \mathcal{T}'' .

From the recurrence of the variable-order, reinforced random walk and equation (7), it is evident that the set of irreducible Markov chains has measure 1

under ψ_{w,v_0} . In this set, the variables $\{\bar{C}_n(u, v)/\bar{C}_n(u) : \forall u, v \text{ admissible}\}$ converge almost surely to the transition probabilities, so for any $T \notin \mathcal{T}''$ irreducible, $P_{v_0}^T(D) = 0$. Furthermore, by Lemma A.3, D happens almost surely in the variable-order, reinforced random walk. Putting this into equation (7), we have

$$H_{w,v_0}(D) = 1 = \int_{\mathcal{T}} P_{v_0}^T(D) d\psi_{w,v_0}(T) \leq \int_{\mathcal{T}''} d\psi_{w,v_0}(T),$$

which implies the proposition. \square

LEMMA A.2 (Lévy). *Consider a sequence of events $B_k \in \mathcal{F}_k, k \in \mathbb{N}$, in some filtration $\{\mathcal{F}_k\}$. Let $b_n = \sum_{k=1}^n \mathbf{1}_{B_k}$ be the total number of events occurring among the first n , and let $s_n = \sum_{k=1}^n P(B_k | \mathcal{F}_{k-1})$ be the sum of the first n conditional probabilities. Then, for almost every ω :*

- If $s_n(\omega)$ converges as $n \rightarrow \infty$, then $b_n(\omega)$ has a finite limit.
- If $s_n(\omega)$ diverges, then $b_n(\omega)/s_n(\omega) \rightarrow 1$.

LEMMA A.3. $H_{w,v_0}(D) = 1$.

PROOF. For any u in $\{\mathcal{X}^q : q \leq r + 1\}$, the variables $n^{-1}w'_n(u)$ and $n^{-1}w''_n(u)$ are functions of $\{n^{-1}\bar{C}_n(u, v) : \forall u, v \text{ admissible}\}$, therefore they converge almost surely, because the reinforced random walk is a mixture of irreducible Markov chains for which the latter converge. The reinforcement scheme defined in Definition 4.3 imposes some constraints on the limits of $n^{-1}w'_n(u)$ and $n^{-1}w''_n(u)$. Note that $w''_n(u), w''_n(u^*), w'_n(u)$ and $w'_n(u^*)$ never differ by more than c ; we also know that the reinforced random walk is positive recurrent (it is a mixture of irreducible, finitely-valued Markov chains), so almost surely

$$\begin{aligned} \lim_{n \rightarrow \infty} n^{-1}w''_n(u) &= \lim_{n \rightarrow \infty} n^{-1}w''_n(u^*) = \lim_{n \rightarrow \infty} n^{-1}w'_n(u) \\ (16) \qquad \qquad \qquad &= \lim_{n \rightarrow \infty} n^{-1}w'_n(u^*) > 0. \end{aligned}$$

Denote this limit $w_\infty(u) = w_\infty(u^*)$. It is also easy to see that if $u \in \mathcal{X}^q$, then for all $s > q$,

$$(17) \qquad \sum_{\{v \in \mathcal{X}^s : v \text{ ends in } u\}} w_\infty(v) = w_\infty(u).$$

Now, let τ_n be the n th visit to $u \in \mathcal{X}^r$ and let B_n be the event that we make a transition to v at τ_n . Define

$$p_n(f(u), v) \equiv H_{w,v_0}(B_n | \sigma(Y_1, \dots, Y_{\tau_n})) = \frac{w'_{\tau_n}(\overline{f(u)v})}{w''_{\tau_n}(f(u))}.$$

We know $p_n(f(u), v)$ converges a.s. to $w_\infty(\overline{f(u)v})/w_\infty(f(u)) > 0$. Therefore, $\sum_n p_n(f(u), v) = \infty$ a.s., and by Lévy’s extension of the Borel–Cantelli lemma (Lemma A.2),

$$\begin{aligned} \frac{\sum_{m=1}^n \mathbf{1}_{B_m}}{\sum_{m=1}^n p_m(f(u), v)} &\rightarrow 1 \quad \text{a.s.} \\ \implies \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{m=1}^n \mathbf{1}_{B_m} &= \lim_{k \rightarrow \infty} \frac{\vec{C}_k(u, v)}{\vec{C}_k(u)} = \frac{w_\infty(\overline{f(u)v})}{w_\infty(f(u))}. \end{aligned}$$

This means that $\{\vec{C}_n(u, v)/\vec{C}_n(u) : \forall u, v \text{ admissible}\}$ converges H_{w, v_0} -a.s. to a set of transition probabilities, $w_\infty(\overline{f(u)v})/w_\infty(f(u))$, for a variable-order Markov chain with histories \mathcal{H} . To show that this Markov chain is reversible, note that w_∞ is the stationary distribution, because

$$\begin{aligned} \sum_{\left\{ \begin{array}{l} u \in \mathcal{X}^r: \\ u, v \text{ admissible} \end{array} \right\}} w_\infty(u) \frac{w_\infty(\overline{f(u)v})}{w_\infty(f(u))} &= \sum_{\left\{ \begin{array}{l} h \in \mathcal{H} \text{ minimal:} \\ h, v \text{ admissible} \end{array} \right\}} \sum_{\{u: f(u)=h\}} w_\infty(u) \frac{w_\infty(\overline{hv})}{w_\infty(h)} \\ &= \sum_{\left\{ \begin{array}{l} h \in \mathcal{H} \text{ minimal:} \\ h, v \text{ admissible} \end{array} \right\}} w_\infty(h) \frac{w_\infty(\overline{hv})}{w_\infty(h)} \\ &= w_\infty(v), \end{aligned}$$

where we used equation (17) in the last two identities. By equation (16), w_∞ satisfies the conditions for reversibility. Therefore, $H_{w, v_0}(D) = 1$. \square

Acknowledgments. The author would like to thank Persi Diaconis and Vijay Pande for valuable suggestions, and Lutz Maibaum for providing molecular dynamics datasets.

SUPPLEMENTARY MATERIAL

Law of a variable-order, reinforced random walk (DOI: 10.1214/10-AOS857 SUPP; .pdf). We provide a closed form expression for this law as a function of transition counts and suggest how it could be useful.

REFERENCES

[1] ALDOUS, D. and FILL, J. Reversible Markov chains and random walks on graphs. Unpublished manuscript.
 [2] BACALLADO, S. (2010). Supplement to “Bayesian analysis of variable-order, reversible Markov chains.” DOI:10.1214/10-AOS857SUPP.
 [3] BACALLADO, S., CHODERA, J. D. and PANDE, V. (2009). Bayesian comparison of Markov models of molecular dynamics with detailed balance constraint. *J. Chem. Phys.* **131** 045106.

- [4] BEGLEITER, R., EL-YANIV, R. and YONA, G. (2004). On prediction using variable order Markov models. *J. Artificial Intelligence Res.* **22** 385–421. [MR2129473](#)
- [5] BUCHETE, N. and HUMMER, G. (2008). Coarse master equations for peptide folding dynamics. *J. Phys. Chem. B* **112** 6057–6069.
- [6] BÜHLMANN, P. (2000). Model selection for variable length Markov chains and tuning the context algorithm. *Ann. Inst. Statist. Math.* **52** 287–315. [MR1763564](#)
- [7] BÜHLMANN, P. and WYNER, A. J. (1999). Variable length Markov chains. *Ann. Statist.* **27** 480–513. [MR1714720](#)
- [8] DIACONIS, P. and FREEDMAN, D. (1980). de Finetti theorem for Markov chains. *Ann. Probab.* **8** 115–130. [MR0556418](#)
- [9] 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 Univ. Press, New York. [MR1008047](#)
- [10] DIACONIS, P. and ROLLES, S. (2006). Bayesian analysis for reversible Markov chains. *Ann. Statist.* **34** 1270–1292. [MR2278358](#)
- [11] KASS, R. and RAFTERY, A. (1995). Bayes factors. *J. Amer. Statist. Assoc.* **90** 773–795.
- [12] KEANE, M. and ROLLES, S. (2000). Edge-reinforced random walk on finite graphs. In *Infinite Dimensional Stochastic Analysis* 217–234. R. Neth. Acad. Arts Sci., Amsterdam. [MR1832379](#)
- [13] KELLY, F. P. (1979). *Reversibility and Stochastic Networks*. Wiley, Chichester. [MR0554920](#)
- [14] MACQUEEN, J. (1981). Circuit processes. *Ann. Probab.* **9** 604–610. [MR0624686](#)
- [15] RISSANEN, J. (1983). A universal data compression system. *IEEE Trans. Inform. Theory* **29** 656–664. [MR0730903](#)
- [16] ROLLES, S. (2003). How edge-reinforced random walk arises naturally. *Probab. Theory Related Fields* **126** 243–260. [MR1990056](#)
- [17] WEINBERGER, M., RISSANEN, J. and FEDER, M. (1995). A universal finite memory source. *IEEE Trans. Inform. Theory* **41** 643–652.
- [18] ZABELL, S. (1982). W. E. Johnson’s sufficientness postulate. *Ann. Statist.* **10** 1091–1099. [MR0673645](#)

DEPARTMENT OF STRUCTURAL BIOLOGY
STANFORD UNIVERSITY
CLARK CENTER, S296
STANFORD, CALIFORNIA 94305
USA
E-MAIL: sergiob@stanford.edu