Estimation in Reversible Markov Chains

David H. ANNIS, Peter C. KIESSLER, Robert LUND, and Tara L. STEUBER

This article examines estimation of the one-step-ahead transition probabilities in a reversible Markov chain on a countable state space. A symmetrized moment estimator is proposed that exploits the reversible structure. Examples are given where the symmetrized estimator has superior asymptotic properties to those of a naive estimator, implying that knowledge of reversibility can sometimes improve estimation. The asymptotic mean and variance of the estimators are quantified. The results are proven using only elementary results such as the law of large numbers and the central limit theorem.

KEY WORDS: Asymptotic variance; Reversibility; Transition probability estimation.

1. INTRODUCTION

This article studies estimation of the transition probabilities in a time-reversible Markov chain $\{X_t\}_{t=0}^{\infty}$. The chain's state space S is taken as a countable subset of $\{0, 1, \ldots\}$. The chain is assumed to be irreducible, aperiodic, and positive recurrent. Such chains have a unique limiting distribution with $\lim_{t\to\infty} \Pr[X_t = j | X_0 = i] = \pi_j$ for every $i \in S$, where $\pi_j > 0$ for $j \in S$. The one-step-ahead transition matrix $\mathbf{P} = (p_{i,j})_{i,j\in S}$ has (i, j)th entry $p_{i,j} = \Pr[X_{t+1} = j | X_t = i]$. The chain is assumed to be time-homogeneous in that $p_{i,j}$ does not depend on t. The data are assumed sampled from a stationary chain; sufficient for this is that $\Pr[X_0 = k] = \pi_k$ for all states $k \in S$.

The chain is said to be reversible if

$$\pi_i p_{i,j} = \pi_j p_{j,i}$$

for each pair of states i and j. Reversibility implies that the long-term flow rate from state i to j equals that from state j to i. Kolmogorov's criterion allows one to assess reversibility directly from the $p_{i,j}$'s; specifically, the chain is reversible if and only if

$$p_{i,i_1} p_{i_1,i_2} \cdots p_{i_k,i} = p_{i,i_k} p_{i_k,i_{k-1}} \cdots p_{i_1,i} \qquad (1)$$

for each $k \geq 2$ and all states $i, i_1, \ldots, i_k$ (Kijima 1997; Ross 2007). It is not clear whether one can statistically assess reversibility from a realization of a chain; however, the chain cannot be reversible if there exist i and j with $p_{i,j} > 0$ and $p_{j,i} = 0$. The works by Diaconis and Stroock (1991), Kijima (1997), Chen (2005), Stroock (2005), and Ross (2007) are good references for general properties of reversible chains.

David H. Annis is ???, Wachovia Bank, Charlotte, NC 28288 (E-mail: *dave@sportsquant.com*). Peter C. Kiessler is ???, Robert Lund is ???, Tara L. Steuber is ???, Department of Mathematical Sciences, Clemson University, Clemson, SC 29634-0975. The authors thank the referee and editor for helpful suggestions. Robert Lund's research was supported by National Science Foundation grant DMS 0905570.

Several broad classes of Markov chains, including random walks on graphs, birth and death chains, and many Markov chain Monte Carlo generated chains, are known to be reversible. For one example, a discrete-time birth and death chain on $S = \{0, 1, \ldots\}$ is a chain that can only move one unit from its current position, either up or down, in any nonboundary transition. Specifically, the nonzero entries in the transition matrix have the form $p_{i,i+1} = \alpha_i$ and $p_{i,i-1} = 1 - \alpha_i$ when $i \geq 1$ (we take $p_{0,1} = \alpha_0$ and $p_{0,0} = 1 - \alpha_0$ where $\alpha_0 > 0$ so that the chain will be aperiodic). A second example of a reversible chain is a random walk on a graph. Here, S is a finite set and there is a collection of bivariate pairs of states called edges. The walk can transition from i to j only when the state pair (i, j) is an edge. It may be helpful to think of various U.S. cities as the states in the chain, with an edge existing between cities i and j when it is possible to fly directly from city i to j. The cost of traveling directly from city i to j is $w_{i,j}$. Symmetry is assumed in that one can fly directly from j to i if it is possible to fly directly from i to j; we also take $w_{i,j} = w_{j,i}$. The probability of undergoing a transition from i to j is proportional to its cost in that

$$p_{i,j} = \frac{w_{ij}}{\sum_{j\in S} w_{i,j}}.$$

See the books by Kijima (1997), Stroock (2005), and Ross (2007) for further examples of reversible chains.

Suppose we observe the data $X_0, \ldots, X_t$ and wish to estimate the one-step-ahead transition probabilities $p_{i,j}$ for all states $i \neq j \in S$. The classical (naive) estimator of $p_{i,j}$ is

$$\hat{p}_{i,j}^{(N)}(t) = \frac{N_{i,j}(t)}{N_i(t)} 1_{[N_i(t)>0]}, \qquad (2)$$

where $1_{[A]}$ is an indicator that is 1 when the event A occurs and zero otherwise, $N_{i,j}(t)$ is the number of one-step-ahead transitions from i to j, and $N_i(t)$ is the number of times state i is visited up to time t. The indicator $1_{[N_i(t)>0]}$ in (2) is introduced to avoid division by zero. The counts $N_{i,j}(t)$ and $N_i(t)$ are

$$N_{i,j}(t) = \sum_{\ell=0}^{t-1} 1_{[X_\ell = i \cap X_{\ell+1} = j]} \quad \text{and} \quad N_i(t) = \sum_{\ell=0}^{t} 1_{[X_\ell = i]}. \qquad (3)$$

One may ask if a priori knowledge of a chain's reversibility aids transition probability estimation. In particular, is $\hat{p}_{i,j}^{(N)}(t)$ in (2) the best asymptotic estimator? This question was beautifully answered by Greenwood and Wefelmeyer (1999) and Greenwood, Schick, and Wefelmeyer (2001) who showed that the symmetrized (reversible) estimator

$$\hat{p}_{i,j}^{(R)}(t) = \frac{N_{i,j}(t) + N_{j,i}(t)}{2N_i(t)} 1_{[N_i(t)>0]} \qquad (4)$$

is not only preferable, but also asymptotically most efficient. Since the joint distributions of $(X_0, \ldots, X_t)$ and $(X_t, \ldots, X_0)$ are identical in reversible chains, the estimator in (4) can be viewed as merely averaging forward and backward versions of (2).

The goal of this article is to further understand estimation for reversible chains. In Section 2, the reversible and naive estimators are reformulated from a renewal-based perspective. In Section 3, we show that both estimators are asymptotically unbiased and calculate their asymptotic variances in a straightforward manner, using only the classic limit theorems from probability. Our work will show that the asymptotic variance of the reversible estimator is never larger than that of the naive estimator, that

$$\lim_{t\to\infty} \frac{\operatorname{Var}(\hat{p}_{i,j}^{(R)}(t))}{\operatorname{Var}(\hat{p}_{i,j}^{(N)}(t))} \in \left[\frac{1}{2}, 1\right],$$

and that both bounds are tight (i.e., there are examples where the reversible estimator is, asymptotically, twice as efficient). Implications of our results are that the naive and reversible estimators have the same asymptotic performance for a birth and death chain, but that the reversible estimator is more efficient in the case of a random walk on a graph.

2. REFORMULATION OF THE ESTIMATORS

This section uses renewal theory to express $\hat{p}_{i,j}^{(N)}(t)$ and $\hat{p}_{i,j}^{(R)}(t)$ in a form which facilitates their asymptotic analysis. Observe that the two estimators are identical when $i = j$; hence, we assume that $i \neq j$. The times at which the chain visits state i form a renewal sequence. Let $N_i(t)$ be the number of visits (renewals) to state i which have occurred up to time t. The renewal times partition the observed states into cycles, the ℓth cycle consisting of the succession of states visited between the ℓth and $(\ell+1)$st visits to state i. An initial sojourn of states prior to the beginning of the first cycle exists unless $X_0 = i$. Likewise, time t typically occurs during the interior times of a cycle; hence, the last cycle may be incomplete.

Let $C_\ell = 1$ if the ℓth cycle begins with a transition from state i to state j; otherwise, set $C_\ell = 0$. It follows that

$$N_{i,j}(t) = \sum_{\ell=1}^{N_i(t-1)} C_\ell$$

and

$$\hat{p}_{i,j}^{(N)}(t) = \frac{\sum_{\ell=1}^{N_i(t-1)} C_\ell}{N_i(t)} 1_{[N_i(t)>0]}. \tag{5}$$

Set $D_\ell = 1$ if the ℓth cycle ends in state j; otherwise, set $D_\ell = 0$. For edge effects induced by the initial and possibly incomplete last cycle, set $E_1(t) = 1$ if the trajectory of states before the first cycle (before visiting state i for the first time) ends in state j; otherwise, take $E_1(t) = 0$. Take $E_2(t)$ as unity only when the observed data end with a transition from j to i: $E_2(t) = 1_{[X_{t-1}=j, X_t=i]}$. Then

$$N_{j,i}(t) = E_1(t) + \sum_{\ell=1}^{N_i(t-1)-1} D_\ell + E_2(t).$$

It now follows that

$$\hat{p}_{i,j}^{(R)}(t) = \frac{\sum_{\ell=1}^{N_i(t-1)-1} \left(\frac{C_\ell + D_\ell}{2}\right) + E_1(t) + E_2(t) + E_3(t)}{N_i(t)} \times 1_{[N_i(t)>0]}, \tag{6}$$

where $E_3(t) = C_{N_i(t-1)}$ is a third edge effect. Other renewal representations are possible, but we have taken care to write all statistics as functions of $X_0, \ldots, X_t$ only.

We now collect a few limiting results needed to calculate the asymptotic bias and variance of the estimators. All convergences are as $t \to \infty$. Since the chain is aperiodic and positive recurrent, $N_i(t) \to \infty$ and $N_i(t)/t \to \pi_i$ with probability 1. The random vectors (C_ℓ, D_ℓ) are independent and identically distributed (iid). By the strong Markov property, the probability that a cycle begins with a transition from i to j is $p_{i,j}$; hence, $E[C_\ell] = p_{i,j}$. Since the chain is reversible, the probability that a cycle ends with a transition from j to i is the same as that a cycle begins with a transition from i to j: $E[D_\ell] = p_{i,j}$. Using $C_\ell = C_\ell^2$ and $D_\ell = D_\ell^2$, we have

$$\operatorname{Var}(C_\ell) = \operatorname{Var}(D_\ell) = p_{i,j} - p_{i,j}^2.$$

We next compute $\mathrm{E}(C_\ell D_\ell)$. Observe that $C_\ell D_\ell$ is either zero or unity, with unity occurring if and only if $C_\ell = 1$ and $D_\ell = 1$. But $C_\ell = 1$ and $D_\ell = 1$ when the ℓth cycle begins with a transition from i to j and ends in state j. Since state i cannot be visited during the interior times of this cycle, $C_\ell D_\ell = 1$ with probability $p_{i,j} \sum_{k=0}^{\infty} {}_i p_{i,j}^{(k)} p_{j,i}$, where ${}_i p_{i,j}^{(k)}$ is the "taboo probability" that starting from state i, the chain is in state j at time t and the first return time to state i is greater than k. Here, the adjective "taboo" indicates that state i must be avoided during the interior times in the cycle. It follows that $\mathrm{E}(C_\ell D_\ell) = p_{i,j} \sum_{k=0}^{\infty} {}_i p_{i,j}^{(k)} p_{j,i}$ and the variance of $(C_\ell + D_\ell)/2$ is

$$\begin{aligned}
&\operatorname{Var}\left(\frac{C_\ell + D_\ell}{2}\right) \\
&\quad = \frac{1}{4}\left[2p_{i,j} + 2p_{i,j} \sum_{k=0}^{\infty} {}_i p_{i,j}^{(k)} p_{j,i} - 4p_{i,j}^2\right] \\
&\quad = \frac{1}{2}\left[(p_{i,j} - p_{i,j}^2) + \left(p_{i,j} \sum_{k=0}^{\infty} {}_i p_{i,j}^{(k)} p_{j,i} - p_{i,j}^2\right)\right].
\end{aligned}$$

Finally, note that $E_k(t)/N_i(t)^p \to 0$ with probability 1 for $k = 1, 2, 3$ and any $p > 0$.

3. EXPECTATION AND VARIANCE

The three theorems to follow show that both estimators are consistent and asymptotically unbiased and determine their asymptotic variances. All convergences are as $t \to \infty$ unless otherwise noted.

Theorem 1. The asymptotic mean of $\hat{p}_{i,j}^{(N)}(t)$ and $\hat{p}_{i,j}^{(R)}(t)$ is $p_{i,j}$.

Proof. By the strong law of large numbers, as $m \to \infty$,

$$\frac{1}{m}\sum_{\ell=1}^{m} C_\ell \to p_{i,j} \qquad \text{and} \qquad \frac{1}{m}\sum_{\ell=1}^{m}\left(\frac{C_\ell + D_\ell}{2}\right) \to p_{i,j}$$

with probability 1. But since $N_i(t)$ is integer-valued and converges to infinity and $N_i(t-1)/N_i(t) \to 1$ with probability 1,

$$\frac{1}{N_i(t)} \sum_{\ell=1}^{N_i(t-1)} C_\ell \to p_{i,j}$$

and

$$\frac{1}{N_i(t)} \sum_{\ell=1}^{N_i(t-1)-1} \left(\frac{C_\ell + D_\ell}{2}\right) \to p_{i,j}$$

with probability 1. Also, $E_k(t)/N_i(t) \to 0$ for $k = 1, 2, 3$ and $1_{[N_i(t)>0]} \to 1$ with probability 1. Using these results and (5) and (6), we infer that $\hat{p}_{i,j}^{(N)}(t) \to p_{i,j}$ and $\hat{p}_{i,j}^{(R)} \to p_{i,j}$ with probability 1. Since both $\hat{p}_{i,j}^{(N)}(t)$ and $\hat{p}_{i,j}^{(R)}(t)$ are nonnegative and bounded above by unity, the convergence of $E[\hat{p}_{i,j}^{(N)}(t)]$ and $E[\hat{p}_{i,j}^{(R)}(t)]$ to $p_{i,j}$ follows from the dominated convergence theorem.

Theorem 2. As $t \to \infty$, we have the following distributional convergence:

$$\sqrt{t}\big(\hat{p}_{i,j}^{(N)}(t) - p_{i,j}\big) \xrightarrow{\mathcal{D}} N\left(0, \frac{p_{i,j} - p_{i,j}^2}{\pi_i}\right) \stackrel{\mathcal{D}}{=} N\left(0, \frac{\operatorname{Var}(C_1)}{\pi_i}\right). \tag{7}$$

Proof. A careful analysis based on (5) and cases provides

$$\big(\hat{p}_{i,j}^{(N)}(t) - p_{i,j}\big) = \left[\frac{\sum_{\ell=1}^{N_i(t-1)} (C_\ell - p_{i,j})}{N_i(t-1)}\right] \frac{N_i(t-1)}{N_i(t)} \times 1_{[N_i(t-1)>0]} - p_{i,j} 1_{[N_i(t-1)=0]}. \tag{8}$$

To handle the edge-effect term in (8), note that

$$\sqrt{t}\, p_{i,j} 1_{[N_i(t-1)=0]} \xrightarrow{\mathcal{P}} 0$$

due to $\Pr[N_i(t-1) = 0] = \Pr(\tau_1 > t-1) \le E[\tau_1]/(t-1)$, which is justified by Markov's inequality. Here, τ_1 is the first time the chain visits state i; $E[\tau_1]$ is finite by the assumed positive recurrence. Observe that $N_i(t-1)/N_i(t) \to 1$ and $1_{[N_i(t-1)>0]} \to 1$ (all with probability 1). An application of Slutzky's theorem now shows that our work is done if we simply prove that

$$\frac{\sqrt{t}}{N_i(t-1)} \sum_{\ell=1}^{N_i(t-1)} (C_\ell - p_{i,j}) \xrightarrow{\mathcal{D}} N\left(0, \frac{\operatorname{Var}(C_1)}{\pi_i}\right). \tag{9}$$

To verify (9), apply the central limit theorem to the iid sequence $\{C_\ell\}$ to infer that as $m \to \infty$,

$$\frac{1}{\sqrt{m}} \sum_{\ell=1}^{m} (C_\ell - p_{i,j}) \xrightarrow{D} N(0, \operatorname{Var}(C_1)).$$

Since $N_i(t) \to \infty$, theorem 17.1 in the book by Billingsley (1968) gives

$$\frac{1}{\sqrt{N_i(t-1)}} \sum_{\ell=1}^{N_i(t-1)} (C_\ell - p_{i,j}) \xrightarrow{D} N(0, \operatorname{Var}(C_1)),$$

which implies (7) and (9) when combined with $\sqrt{t/N_i(t-1)} \to \sqrt{1/\pi_i}$ and $\operatorname{Var}(C_1) = p_{i,j} - p_{i,j}^2$.

A similar argument proves the following result, the essential change being that (6) is used in place of (5), and $\operatorname{Var}((C_1 + D_1)/2)$ replaces $\operatorname{Var}(C_1)$.

Theorem 3. As $t \to \infty$,

$$\sqrt{t}\big(\hat{p}_{i,j}^{(R)}(t) - p_{i,j}\big) \xrightarrow{\mathcal{D}} N\left(0, \frac{(p_{i,j} - p_{i,j}^2) + \big(p_{i,j} \sum_{k=0}^{\infty} {}_i p_{i,j}^{(k)} p_{j,i} - p_{i,j}^2\big)}{2\pi_i}\right) \stackrel{\mathcal{D}}{=} N\left(0, \frac{\operatorname{Var}((C_1 + D_1)/2)}{\pi_i}\right). \tag{10}$$

In terms of asymptotic efficiencies, we have now shown that

$$\lim_{t\to\infty} \frac{\operatorname{Var}(\hat{p}_{i,j}^{(R)}(t))}{\operatorname{Var}(\hat{p}_{i,j}^{(N)}(t))} = \frac{\operatorname{Var}\big(\frac{C_1+D_1}{2}\big)}{\operatorname{Var}(C_1)} = \frac{\sigma_R^2}{\sigma_N^2}, \tag{11}$$

where

$$\sigma_N^2 = \frac{p_{i,j} - p_{i,j}^2}{\pi_i} \qquad \text{and} \qquad \sigma_R^2 = \frac{(p_{i,j} - p_{i,j}^2) + \big(p_{i,j} \sum_{k=0}^{\infty} {}_i p_{i,j}^{(k)} p_{j,i} - p_{i,j}^2\big)}{2\pi_i}. \tag{12}$$

Observe that $\sum_{k=0}^{\infty} {}_i p_{i,j}^{(k)} p_{j,i} \le \sum_{k=0}^{\infty} \Pr_i[\eta_i = k+1] \le 1$, η_i denoting the time of first return to state i and $\Pr_i$ indicating the initial condition $X_0 = i$. Using this in (12) shows that $\sigma_R^2 \le \sigma_N^2$. In the next section, we will show that $\sigma_R^2/\sigma_N^2 \ge 1/2$.

4. LOWER BOUNDS FOR σ_R^2/σ_N^2

We start with two examples. In the first, C_ℓ and D_ℓ are perfectly correlated and the asymptotic efficiency of the naive and reversible estimators is unity. In the second example, C_ℓ and D_ℓ are uncorrelated and the reversible estimator is twice as efficient as the naive estimator.

Consider a birth and death chain. This chain is skip-free in that from state $i \ge 1$, the only possible transitions are to states $i-1$ and $i+1$. The transition probabilities are $p_{i,i+1} = \alpha_i$ and $p_{i,i-1} = 1 - \alpha_i$, where $\alpha_i \in [0, 1]$ (at state 0, we take $p_{0,1} = \alpha_0$ and $p_{0,0} = 1 - \alpha_0$). Assuming $\alpha_i > 0$ for all $i \ge 0$ and $\alpha_i < 1/2$ for all large i, the chain is irreducible, aperiodic, positive recurrent, and reversible and has a limiting distribution with form

$$\pi_j = \begin{cases} K, & j = 0 \\ \dfrac{\alpha_1 \cdots \alpha_{j-1}}{(1-\alpha_1)\cdots(1-\alpha_j)} K, & j > 0. \end{cases}$$

Here, the constant K is such that the limiting distribution has unit mass.

The only nonzero $p_{i,j}$'s occur when $j = i-1$ or $j = i+1$. When $j = i+1$, then if $C_\ell = 1$, the ℓth cycle starts with a transition from i to $i+1$ and, by the skip-free property, must end with a transition from $i+1$ to i. Hence, $D_\ell = 1$ for this cycle. If $C_\ell = 0$, then the ℓth cycle starts with a transition from i to $i-1$ and, by the skip-free property, must end with a transition from $i-1$ to i. Hence, $D_\ell = 0$ for this cycle. It now follows that $\operatorname{Var}((C_\ell + D_\ell)/2) = \operatorname{Var}(C_\ell)$. Thus, for skip-free chains, the reversible and naive estimators have the same asymptotic efficiency.

As a second example, consider an iid chain. Specifically, $X_0, X_1, \ldots$ are independent and have the common probability mass function $\Pr[X_i = j] = \pi_j$ with $\pi_j > 0$ for all j. Such a sequence can be regarded as a Markov chain with the transition probabilities $p_{i,j} = \pi_j$. The stationary distribution is $\{\pi_i\}_{i=0}^{\infty}$ and the chain is easily shown to be irreducible, aperiodic, positive recurrent, and reversible.

To calculate σ_R^2, note that the taboo probability is

$$\sum_{k=0}^{\infty} {}_i p_{i,j}^{(k)} = \sum_{k=0}^{\infty} (1-\pi_i)^k \pi_j = \pi_i^{-1} \pi_j.$$

It follows from (12) that

$$\begin{aligned}\sigma_R^2 &= \frac{1}{2}(\pi_j - \pi_j^2 + \pi_j \pi_i^{-1} \pi_j \pi_i - \pi_j^2) \\ &= \frac{\pi_j - \pi_j^2}{2} = \frac{\sigma_N^2}{2}.\end{aligned}$$

Hence, $\hat{p}_{i,j}^{(R)}(t)$ is asymptotically twice as efficient as $\hat{p}_{i,j}^{(N)}(t)$.

We close by showing that $\operatorname{Cov}(C_\ell, D_\ell) \ge 0$. With this and (11), we have $1/2 \le \sigma_R^2/\sigma_N^2 \le 1$ and the two examples above provide cases where the relative efficiencies of $1/2$ and 1 are achieved.

Theorem 4. C_ℓ and D_ℓ are nonnegatively correlated; that is, $\operatorname{Cov}(C_\ell, D_\ell) \ge 0$.

Proof. Because of the binary structure of C_ℓ and D_ℓ, it suffices to show that $\Pr(C_\ell = 1, D_\ell = 1) \ge \Pr(C_\ell = 1)\Pr(D_\ell = 1)$. To this end, we note that since

$$\begin{aligned}&\Pr(C_\ell = 1)\Pr(D_\ell = 1) \\ &\quad = [\Pr(C_\ell = 1, D_\ell = 1) + \Pr(C_\ell = 1, D_\ell = 0)] \\ &\qquad \times [\Pr(C_\ell = 1, D_\ell = 1) + \Pr(C_\ell = 0, D_\ell = 1)] \\ &\quad = \Pr(C_\ell = 1, D_\ell = 1)[1 - \Pr(C_\ell = 0, D_\ell = 0)] \\ &\qquad + \Pr(C_\ell = 1, D_\ell = 0)\Pr(C_\ell = 0, D_\ell = 1),\end{aligned}$$

it suffices to show that

$$\begin{aligned}&\Pr(C_\ell = 1, D_\ell = 1)\Pr(C_\ell = 0, D_\ell = 0) \\ &\quad \ge \Pr(C_\ell = 1, D_\ell = 0)\Pr(C_\ell = 0, D_\ell = 1). \qquad (13)\end{aligned}$$

Since $\Pr(C_\ell = 1, D_\ell = 0)$ is the probability that a cycle begins with a transition from i to j and ends with a transition from some state other than j to i, we have

$$\Pr(C_\ell = 1, D_\ell = 0) = \sum_A p_{i,j} p_{j,k_1} \cdots p_{k_n,i},$$

where $A = \bigcup_{n=1}^{\infty} \{(k_1, \ldots, k_n); k_h \ne i \text{ for } h = 1, \ldots, n \text{ and } k_n \ne j\}$. Similarly, since $\Pr(C_\ell = 0, D_\ell = 1)$ is the probability a cycle begins with a transition from i to some state other than j and ends with a transition from j to i,

$$\Pr(C_\ell = 0, D_\ell = 1) = \sum_B p_{i,l_1} \cdots p_{l_m,j} p_{j,i},$$

where $B = \bigcup_{m=1}^{\infty} \{(l_1, \ldots, l_m); l_h \ne i \text{ for } h = 1, \ldots, m \text{ and } l_1 \ne j\}$.

Thus,

$$\begin{aligned}&\Pr(C_\ell = 1, D_\ell = 0)\Pr(C_\ell = 0, D_\ell = 1) \\ &\quad = \sum_A \sum_B p_{i,j} p_{j,k_1} \cdots p_{k_n,i} p_{i,l_1} \cdots p_{l_m,j} p_{j,i}.\end{aligned}$$

An application of Kolmogorov's criteria for reversibility in (1) gives

$$\begin{aligned}&\Pr(C_\ell = 1, D_\ell = 0)\Pr(C_\ell = 0, D_\ell = 1) \\ &\quad = p_{i,j} p_{j,i} \Bigg(\sum_A \sum_B p_{i,k_n} \cdots p_{k_1,j} p_{j,l_m} \cdots p_{l_1,i} \Bigg).\end{aligned}$$

Since n and m are both at least 1 and l_1 and k_n do not equal j, each term in the double summation is the probability of some cycle that begins with a transition from i to some state other than j and ends with a transition from some state other than j to i. Thus, the term inside the parentheses is less than or equal to $\Pr(C_\ell = 0, D_\ell = 0)$ and

$$\begin{aligned}&\Pr(C_\ell = 1, D_\ell = 0)\Pr(C_\ell = 0, D_\ell = 1) \\ &\quad \le p_{i,j} p_{j,i} \Pr(C_\ell = 0, D_\ell = 0). \qquad (14)\end{aligned}$$

Because one way for a cycle to have $C_\ell = 1$ and $D_\ell = 1$ is to make a transition from i to j and then immediately back to i, we have

$$p_{i,j} p_{j,i} \le \Pr(C_\ell = 1, D_\ell = 1). \qquad (15)$$

Combining (14) and (15) gives (13) and completes the proof.

5. CONCLUSION AND COMMENTS

Reversibility is a structural property inherited by many Markov chains. Reversibility can be exploited in some cases to obtain transition probability estimates that have smaller asymptotic variances than naive estimators based on ratios of counts. The improvement in the asymptotic efficiency of a reversible estimate, relative to a naive estimate, is quantified in (11). In cases where the chain possesses the so-called skip-free property, such as the birth and death chain in Section 1, there is no improvement; in other cases, such as the random walk on a graph, some improvement may be possible. In any case, the reversible estimator's asymptotic variance can be no lower than half the naive estimator's asymptotic variance.

[Received April 2009. Revised December 2009.]

REFERENCES

Billingsley, P. (1968), *Convergence of Probability Measures*, New York: Wiley. [3]

Chen, M. F. (2005), *Eigenvalues, Inequalities and Ergodic Theory*, London: Springer. [1]

Diaconis, P., and Stroock, D. (1991), "Geometric Bounds for Eigenvalues of Markov Chains," *The Annals of Applied Probability*, 1, 36–61. [1]

Greenwood, P. E., and Wefelmeyer, W. (1999), "Reversible Markov Chains and Optimality of Symmetrized Empirical Estimators," *Bernoulli*, 5, 109–123. [1]

Greenwood, P. E., Schick, A., and Wefelmeyer, W. (2001), Comment on "Inference for Semiparametric Models: Some Questions and an Answer," by P. J. Bickel and J. Kwon, *Statistica Sinica*, 11, 892–905. [1]

Kijima, M. (1997), *Markov Processes for Stochastic Modeling*, London: Chapman & Hall. [1]

Ross, S. M. (2007), *Introduction to Probability Models* (9th ed.), Burlington, MA: Academic Press. [1]

Stroock, D. (2005), *An Introduction to Markov Processes*, New York: Springer. [1]

META DATA IN THE PDF FILE

Following information will be included as pdf file Document Properties:

Title : Estimation in Reversible Markov Chains
Author : David H. Annis, Peter C. Kiessler, Robert Lund, Tara L. Steuber
Subject : The American Statistician, Vol.0, No.0, 2010, 1-6
Keywords: Asymptotic variance, Reversibility, Transition probability estimation

THE LIST OF URI ADRESSES

Listed below are all uri addresses found in your paper. The non-active uri addresses, if any, are indicated as ERROR. Please check and update the list where necessary. The e-mail addresses are not checked – they are listed just for your information. More information can be found in the support page: *http://www.e-publications.org/ims/support/urihelp.html*.

```
--- mailto:dave@sportsquant.com [2:pp.1,1] Check skip
200 http://www.amstat.org [2:pp.1,1] OK
200 http://dx.doi.org/10.1198/tast.2010.09083 [2:pp.1,1] OK
302 http://pubs.amstat.org/loi/tas [2:pp.1,1] Found
```