

TOTAL VARIATION APPROXIMATION FOR QUASI-STATIONARY DISTRIBUTIONS

A. D. BARBOUR,* *Universität Zürich*

P. K. POLLETT,** *University of Queensland*

Abstract

Quasi-stationary distributions, as discussed in Darroch and Seneta (1965), have been used in biology to describe the steady state behaviour of population models which, while eventually certain to become extinct, nevertheless maintain an apparent stochastic equilibrium for long periods. These distributions have some drawbacks: they need not exist, nor be unique, and their calculation can present problems. In this paper, we give biologically plausible conditions under which the quasi-stationary distribution is unique, and can be closely approximated by distributions that are simple to compute.

Keywords: Quasi-stationary distribution; total variation distance; stochastic logistic model

2010 Mathematics Subject Classification: Primary 60J28; 92D25; 92D30

1. Introduction

The logistic growth model of Verhulst (1838) was the first to describe mathematically the evolution of a population to a nonzero equilibrium, contrasting with the Malthusian law of exponential growth. Its stochastic version, a Markov chain X in continuous time in which $X(t)$ represents the number of individuals at time t in a population in a prescribed area A , has transition rates

$$q_{i,i+1} = bi \quad \text{and} \quad q_{i,i-1} = di + \frac{ei^2}{A} \quad \text{for } i \geq 1, \quad q_{ij} = 0 \quad \text{otherwise}, \quad (1.1)$$

where b and d are the *per capita* rates of birth and natural mortality, and there is an additional *per capita* death rate ex , due to crowding, at population density $x = i/A$. The stochastic model has the drawback that its equilibrium distribution assigns probability 1 to the state 0, population extinction, irrespective of the initial state. This apparently negates the most valuable property of Verhulst's model, its ability to allow an equilibrium other than extinction. However, if $b > d$ and A is large, the population density $X(t)/A$ can be expected to remain near the 'carrying capacity' $\kappa := (b - d)/e$ for a very long time, in an apparent (and often biologically relevant) nonextinct stochastic equilibrium.

Darroch and Seneta (1965), building on the work of Yaglom (1947) in the context of branching processes, introduced the concept of a quasi-stationary distribution, in an attempt to reconcile

Received 11 May 2010; revision received 26 July 2010.

* Postal address: Angewandte Mathematik, Universität Zürich, Winterthurerstrasse 190, CH-8057 Zürich, Switzerland.
Email address: a.d.barbour@math.uzh.ch

Research supported in part by Schweizerischer Nationalfonds Projekt Nr. 20–107935/1.

** Postal address: Department of Mathematics, University of Queensland, Brisbane, QLD 4072, Australia.

Research supported in part by the Australian Research Council Centre of Excellence for Mathematics and Statistics of Complex Systems.

these at first sight inconsistent properties of the model. In a discrete-time Markov chain X consisting of an absorbing state 0 together with a single finite transient aperiodic class C , the limiting conditional probabilities

$$m_j := \lim_{t \rightarrow \infty} P_i[X(t) = j \mid X(t) \in C], \quad i, j \in C, \tag{1.2}$$

exist, and are the same for each $i \in C$. The $m_j, j \in C$, also determine a quasi-stationary distribution, in the sense that

$$m_k = \sum_{j \in C} m_j p_{jk} / \sum_{j \in C} m_j \sum_{k \in C} p_{jk}, \tag{1.3}$$

where $P := (p_{jk})$ denotes the one-step probability transition matrix. If, however, C is countably infinite, the situation is much less satisfactory; there may be no quasi-stationary distribution, or exactly one, or infinitely many, and determining which of these is the case may be a difficult problem. Even when there is a unique quasi-stationary distribution, its calculation can pose substantial problems, unless equations (1.3) happen to have an obvious solution, because the probabilistic definition (1.2) involves conditioning on an event which, in the limit as $t \rightarrow \infty$, has probability 0. This appears to make the quasi-stationary distribution unsatisfactory for typical biological applications.

In this paper we give conditions, simply expressed in terms of the properties of the process X , under which things are in fact much simpler. Under the conditions of Theorem 2.1, there is exactly one quasi-stationary distribution, and it can be approximated to a specified accuracy by the equilibrium distribution π^μ of a ‘returned process’ X^μ . What is more, under slightly more stringent conditions, the distribution of $X(t)$ is shown in Theorem 2.2 to be close to the quasi-stationary distribution for long periods of time.

The returned process, introduced in Bartlett (1960, pp. 24–25) and used in Ewens (1963), (1964) in a population genetics setting, is a Markov process that evolves exactly like X , up to the time at which 0 is reached, but is then instantly returned to a random state in C , chosen according to the probability measure μ . The mapping $\mu \mapsto \pi^\mu$, studied in the paper of Ferrari *et al.* (1995), is contractive under our conditions, and iterating the mapping leads to the unique quasi-stationary distribution m on C , which satisfies $m = \pi^m$. In many practical applications, including the stochastic logistic model of (1.1) when A is large, iteration is unnecessary, inasmuch as *any* distribution π^μ is extremely close to m . Furthermore, since π^μ is a genuine equilibrium distribution, its computation does not involve conditioning on sets of vanishing probability, and is hence typically much simpler.

The main results, Theorems 2.1 and 2.2, are proved in Section 2. In Section 3, as an illustration, we discuss the application of the theorems to birth-and-death processes, of which the stochastic logistic model (1.1) is an example. Because of their relatively simple structure, birth-and-death processes have already been widely studied; in a biological context, Cattiaux *et al.* (2009) discussed their quasi-stationary distributions, as well as those of analogous diffusion models, which were also examined in detail in Steinsaltz and Evans (2004). In this context, the key quantities appearing in our theorems can be relatively easily estimated. However, our theorems are equally applicable to processes with more complicated structure.

2. A general approximation

Let X be a stable, conservative, and nonexplosive pure-jump Markov process on a countable state space, consisting of a single transient class C together with a cemetery state 0. For any

probability distribution μ on C , define the modified process X^μ with state space C to have exactly the same behaviour as X while in C , but, on reaching 0, to be instantly returned to C according to the distribution μ . Thus, if Q denotes the infinitesimal matrix associated with X , and Q^μ denotes the infinitesimal matrix belonging to X^μ , we have

$$q_{ij}^\mu = q_{ij} + q_{i0}\mu_j \quad \text{for } i, j \in C. \tag{2.1}$$

In this section, under a rather simple set of conditions, we show that the quasi-stationary distribution m of X is unique, and can be approximated in total variation to a prescribed accuracy by the *stationary* distribution of X^μ , for an arbitrary choice of μ . We give a bound on the total variation distance between m and π^μ that is expressed solely in terms of hitting probabilities and mean hitting times for the process X , and which is the same for all μ . The bound is such that it can be expected to be small in circumstances in which the process X typically spends a long time in C in apparent equilibrium, before being absorbed in 0 as a result of an ‘exceptional’ event. If the bound is not, as it stands, small enough for practical use, it can be improved geometrically fast by iteration of the return mapping $\mu \mapsto \pi^\mu$.

Our basic conditions are as follows.

Condition A. *There exist $s \in C$, $p > 0$, and $T < \infty$ such that, uniformly for all $k \in C$,*

- (i) $p_k := P_k[X \text{ hits } s \text{ before } 0] \geq p$;
- (ii) $E_k[\tau_{\{s,0\}}] \leq T < \infty$.

Here, P_k and E_k refer to the distribution of X conditional on $X(0) = k$, and

$$\tau_A := \inf\{t > 0: X(t) \in A, X(s) \notin A \text{ for some } s < t\},$$

the infimum over the empty set being taken to be ∞ . Condition A(i) can be expected to be satisfied in reasonable generality; Condition A(ii), although satisfied by the stochastic logistic model, is not so immediately natural.

We now introduce the quantity

$$U := \sum_{k \in C} \frac{q_{k0}}{q_k E_k(\tau_{\{k,0\}})}, \tag{2.2}$$

where, as usual, $q_k := -q_{kk} = \sum_{j \in C \cup \{0\} \setminus \{k\}} q_{kj}$, and $q_k < \infty$ because X is conservative. To interpret the meaning of U , observe that a renewal argument for X^μ , with renewal epochs the visits to any specific $j \in C$, shows that

$$\pi^\mu(j) = \frac{q_j^{-1}}{E_j(\tau_{\{j\}}^\mu)} \leq \frac{1}{q_j E_j(\tau_{\{j,0\}})}. \tag{2.3}$$

In particular, if X has a quasi-stationary distribution m , it follows from (2.2) that

$$U \geq \sum_{i \in C} \pi^m(i)q_{i0} = \sum_{i \in C} m(i)q_{i0} = \lambda_m,$$

where λ_m is the rate at which the X -process, starting in the quasi-stationary distribution m , leaves C : $P_m[X(t) \in C] = e^{-\lambda_m t}$. Thus, U acts as a computable upper bound for any λ_m . Note that p , T , and U are all quantities that can reasonably be bounded using a knowledge of the process X .

In the remainder of this section, we show that the quasi-stationary distribution m exists, is close to any π^μ , and well describes the long-time behaviour of X prior to absorption in 0 , as long as UT/p is small enough. Our first main result is the following.

Theorem 2.1. *Suppose that Condition A is satisfied, and that $2UT/p < 1$. Then X has a unique quasi-stationary distribution m , and, for any probability measure μ on C , we have*

$$d_{TV}(m, \pi^\mu) \leq \frac{2UT}{p}.$$

Remark. Of course, for the theorem to imply that π^μ is a sharp approximation to m , we need U to be small enough (and, therefore, certainly finite). In many applications, X can only jump to 0 from a small number of states in C , and, if the quasi-equilibrium really behaves like a genuine equilibrium for long periods of time, the quantity $E_k(\tau_{\{k,0\}})$, for each such k , can be expected to contain a large contribution from paths that, after leaving k , spend a very long time ‘in equilibrium’ in other states of C , before either returning to k or being absorbed in 0 . In such applications, as in the next section, these two features combine to make U small.

To prove the theorem, we first need some preparatory results. We first show that, under Condition A, the mean time to hitting the state s is uniformly bounded for all return processes X^μ and all initial states.

Lemma 2.1. *Under Condition A, for all probability measures μ on C and all $r \in C$, we have*

$$E_r \tau_{\{s\}}^\mu \leq \frac{T}{p} < \infty,$$

where τ_A^μ is defined similarly to τ_A , but with the process X^μ in place of X .

Proof. Recursively define

$$\begin{aligned} \tau_{\{s,0\}}^{\mu,1} &:= \tau_{\{s,0\}}^\mu, \\ \tau_{\{s,0\}}^{\mu,j} &:= \inf\{t > \tau_{\{s,0\}}^{\mu,j-1} : X^\mu(t) \in \{s, 0\}, X^\mu(u) \notin \{s, 0\} \text{ for some } \tau_{\{s,0\}}^{\mu,j-1} < u < t\} \end{aligned}$$

for $j \geq 2$, and, for $j \geq 1$, let $Z_j := \mathbf{1}[X^\mu(\tau_{\{s,0\}}^{\mu,l}) = 0, 1 \leq l \leq j]$, taking $Z_0 = 1$. Then it follows that

$$\tau_{\{s\}}^\mu = \sum_{j \geq 1} (\tau_{\{s,0\}}^{\mu,j} - \tau_{\{s,0\}}^{\mu,j-1}) Z_{j-1}.$$

Now $E_r \tau_{\{s,0\}}^{\mu,1} \leq T$ by Condition A(ii), and, for $j \geq 2$,

$$E[(\tau_{\{s,0\}}^{\mu,j} - \tau_{\{s,0\}}^{\mu,j-1}) Z_{j-1} \mid \mathcal{F}_{\tau_{\{s,0\}}^{\mu,j-1}}] = Z_{j-1} \sum_{k \in C} \mu_k E_k \tau_{\{s,0\}} \leq T Z_{j-1}$$

by Condition A(ii), where $\mathcal{F}_{\tau_{\{s,0\}}^{\mu,j-1}}$ denotes the σ -field of events up to the stopping time $\tau_{\{s,0\}}^{\mu,j-1}$. Then, for $j \geq 1$,

$$E[Z_j \mid \mathcal{F}_{\tau_{\{s,0\}}^{\mu,j-1}}] \leq (1 - p) Z_{j-1}$$

by Condition A(i). Hence, for $j \geq 1$ and any $r \in C$, it follows that

$$E_r[(\tau_{\{s,0\}}^{\mu,j} - \tau_{\{s,0\}}^{\mu,j-1}) Z_{j-1}] \leq T(1 - p)^{j-1},$$

and so $E_r \tau_{\{s\}}^\mu \leq T/p$, as required.

It follows in particular from Lemma 2.1 that $E_s \tau_{\{s\}}^\mu \leq T/p < \infty$, so that X^μ is positive recurrent on C ; denote its stationary distribution by π^μ . Then, for any $f: C \rightarrow \mathbb{R}$ bounded, set

$$h_f^\mu(j) := - \int_0^\infty \{E_j f(X^\mu(t)) - \pi^\mu(f)\} dt, \quad j \in C.$$

To show that the integral is well defined, note that

$$|h_f^\mu(j)| \leq \int_0^\infty 2\|f\| d_{TV}(\mathcal{L}(X^\mu(t) \mid X^\mu(0) = j), \pi^\mu) dt,$$

where $\|\cdot\|$ denotes the supremum norm. The latter integral is finite provided that $E_s[(\tau_{\{s\}}^\mu)^2] < \infty$, by the coupling inequality (see Lindvall (2002, Equation (2.8))) and from Pitman (1974, Corollary 1, Equation (1.23) with $r = 2$). That this is the case follows from the next lemma.

Lemma 2.2. *Under Condition A, for all probability measures μ on C , we have*

$$E_s[(\tau_{\{s\}}^\mu)^2] < \infty.$$

Proof. Writing $\tau := \tau_{\{s\}}^\mu$, note that

$$\tau^2 = \left(\int_0^\infty \mathbf{1}[\tau > t] dt \right)^2 = 2 \int_0^\infty \mathbf{1}[\tau > t] \left(\int_t^\infty \mathbf{1}[\tau > u] du \right) dt. \tag{2.4}$$

Now, from Lemma 2.1 and by the Markov property, we have

$$E \left[\int_t^\infty \mathbf{1}[\tau > u] du \mid \mathcal{F}_t^\mu \right] \leq \left(\frac{T}{p} \right) \mathbf{1}[\tau > t],$$

where \mathcal{F}_t^μ denotes the history of X^μ up to time t . Hence, taking expectations in (2.4), it follows that

$$E_s[\tau^2] \leq 2 \left(\frac{T}{p} \right) E_s \tau \leq 2 \left(\frac{T}{p} \right)^2,$$

again from Lemma 2.1, completing the proof.

It is shown in the proof of Theorem 2.2 below that the distribution of $\tau_{\{s\}}^\mu$ actually has an exponential tail.

The functions h_f^μ are central to the argument to come. First, we show that they are bounded and Lipschitz, with appropriate constants.

Lemma 2.3. *For all $j \in C$,*

$$|h_f^\mu(j) - h_f^\mu(s)| \leq \frac{2\|f\|T}{p}.$$

Proof. For any $j \in C \setminus \{s\}$, we can write

$$\begin{aligned} -h_f^\mu(j) &= \int_0^\infty E_j[(f(X^\mu(t)) - \pi^\mu(f)) \mathbf{1}[\tau_{\{s\}}^\mu \leq t]] dt \\ &\quad + \int_0^\infty E_j[(f(X^\mu(t)) - \pi^\mu(f)) \mathbf{1}[\tau_{\{s\}}^\mu > t]] dt. \end{aligned} \tag{2.5}$$

Then, by the strong Markov property, we have

$$\begin{aligned} E_j[(f(X^\mu(t)) - \pi^\mu(f)) \mathbf{1}[\tau_{\{s\}}^\mu \leq t]] &= E_j[E_s[(f(X^\mu(t - \tau_{\{s\}}^\mu)) - \pi^\mu(f)) \mathbf{1}[\tau_{\{s\}}^\mu \leq t]]] \\ &= \int_0^t g_{js}^\mu(v) E_s[f(X^\mu(t - v)) - \pi^\mu(f)] dv, \end{aligned}$$

where g_{js}^μ denotes the probability density of the random variable $\tau_{\{s\}}^\mu$ for the process started at $j \neq s$. Hence, it follows that

$$\begin{aligned} &\int_0^\infty E_j[(f(X^\mu(t)) - \pi^\mu(f)) \mathbf{1}[\tau_{\{s\}}^\mu \leq t]] dt \\ &= \int_0^\infty dt \int_0^\infty dv g_{js}^\mu(v) E_s[f(X^\mu(t - v)) - \pi^\mu(f)] \mathbf{1}\{v \leq t\}. \end{aligned}$$

Now, since

$$\begin{aligned} &\int_0^\infty \mathbf{1}\{v \leq t\} |E_s[f(X^\mu(t - v)) - \pi^\mu(f)]| dt \\ &\leq 2\|f\| \int_0^\infty d_{TV}(\mathcal{L}(X^\mu(t) \mid X^\mu(0) = s), \pi^\mu) dt \\ &< \infty, \end{aligned}$$

we can use Fubini's theorem to conclude that

$$\begin{aligned} &\int_0^\infty E_j[(f(X^\mu(t)) - \pi^\mu(f)) \mathbf{1}[\tau_{\{s\}}^\mu \leq t]] dt \\ &= \int_0^\infty g_{js}^\mu(v) \left\{ \int_v^\infty E_s[f(X^\mu(t - v)) - \pi^\mu(f)] dt \right\} dv \\ &= \int_0^\infty g_{js}^\mu(v) h_f^\mu(s) dv \\ &= h_f^\mu(s). \end{aligned}$$

Hence, from (2.5) and Lemma 2.1, it follows that

$$|h_f^\mu(j) - h_f^\mu(s)| \leq 2\|f\| E_j[\tau_{\{s\}}^\mu] \leq \frac{2\|f\|T}{p},$$

as required.

In particular, the function h_f^μ is itself bounded.

A similar argument, by conditioning on the time of the first jump, shows that

$$h_f^\mu(j) = -q_j^{-1}\{f(j) - \pi^\mu(f)\} + \sum_{k \in C, k \neq j} q_j^{-1} q_{jk}^\mu h_f^\mu(k), \tag{2.6}$$

and the sum in (2.6) is absolutely convergent because h_f^μ is bounded. This can be rewritten in the form

$$(Q^\mu h_f^\mu)(j) = f(j) - \pi^\mu(f), \quad j \in C, \tag{2.7}$$

so that, for any bounded f and any probability measures μ and ν on C , we have

$$\pi^\mu(Q^\nu h_f^\nu) = \pi^\mu(f) - \pi^\nu(f). \tag{2.8}$$

In the terminology of Stein’s method, (2.7) determines h_f^μ to be the solution h of the Stein equation $(Q^\mu h)(j) = f(j) - \pi^\mu(f)$ for the distribution π^μ , corresponding to the given function f . Also, by Dynkin’s formula we have

$$\pi^\mu(Q^\mu h) = 0 \tag{2.9}$$

for any bounded function h (for the special case $h = h_f^\mu$, this follows from (2.8)). We are now in a position to prove Theorem 2.1.

Proof of Theorem 2.1. Take any probability measures μ and ν on C . Then (2.8) gives

$$\pi^\mu(Q^\nu h_f^\nu) = \pi^\mu(f) - \pi^\nu(f),$$

whereas (2.9) gives $\pi^\mu(Q^\mu h_f^\nu) = 0$. Taking the difference, we obtain

$$\pi^\mu(f) - \pi^\nu(f) = \pi^\mu(Q^\nu h_f^\nu - Q^\mu h_f^\nu). \tag{2.10}$$

Now, for bounded h and any $i \in C$,

$$(Q^\nu h - Q^\mu h)(i) = \sum_{k \in C} q_{ik}^\nu (h(k) - h(i)) - \sum_{k \in C} q_{ik}^\mu (h(k) - h(i)),$$

with both sums absolutely convergent, and, from (2.1), it then follows that

$$\begin{aligned} (Q^\nu h - Q^\mu h)(i) &= q_{i0} \sum_{k \in C} (\nu(k) - \mu(k))(h(k) - h(i)) \\ &= q_{i0} \sum_{k \in C} (\nu(k) - \mu(k))(h(k) - h(s)), \end{aligned}$$

since $\sum_{k \in C} \nu(k) = \sum_{k \in C} \mu(k) = 1$. Hence, from (2.10), we have

$$\pi^\mu(f) - \pi^\nu(f) = \sum_{i \in C} \pi^\mu(i) q_{i0} \sum_{k \in C} (\nu(k) - \mu(k))(h_f^\nu(k) - h_f^\nu(s)),$$

and, from Lemma 2.3, this gives

$$|\pi^\mu(f) - \pi^\nu(f)| \leq \sum_{i \in C} \pi^\mu(i) q_{i0} 2 \|f\| \left(\frac{T}{p}\right) \|\nu - \mu\|_{TV}.$$

Thus, it follows that

$$\|\pi^\nu - \pi^\mu\|_{TV} \leq \frac{2T}{p} \sum_{i \in C} \pi^\mu(i) q_{i0} \|\nu - \mu\|_{TV},$$

and (2.3) then implies that

$$\|\pi^\nu - \pi^\mu\|_{TV} \leq \frac{2UT}{p} \|\nu - \mu\|_{TV}.$$

This, by the Banach fixed point theorem, establishes the first part of the theorem, and the second part follows by taking $\nu = m$, and using the fact that, for probability measures F and G , $d_{TV}(F, G) = \frac{1}{2} \|F - G\|_{TV}$.

We now turn our attention to the distribution of $X(t)$ for fixed values of t , starting from any initial distribution, and compare it to m , the distribution at any time of the return process X^m started in the quasi-stationary distribution m . We begin by taking the initial state of X to be s , and remark later that this restriction makes little difference, provided that s is hit at least once.

Theorem 2.2. *Let $B := Tq_s/p \geq 1$. Then, under Condition A and if also $2UT/p < 1$, there is a constant K such that*

$$d_{TV}(\mathcal{L}_s(X(t)), m) \leq Ut + KB\sqrt{\frac{T}{pt}} + \left(\frac{2}{e}\right)^{pt/16T} =: \eta(t).$$

Remark. Hence, if $UB^2T/p \ll 1$, the distribution $\mathcal{L}_s(X(t))$ is close to m for all times t such that

$$\frac{B^2T}{p} \ll t \ll U^{-1}.$$

Proof of Theorem 2.2. The argument is based on coupling two copies $X^{(1)}$ and $X^{(2)}$ of the return process X^m , with $X^{(1)}$ in equilibrium and with $X^{(2)}$ starting in s . The coupling is achieved by forcing $X^{(1)}$ to follow the same sequence of states as $X^{(2)}$ after the first time that it hits s , and to have identical residence times in all states other than s . Define

$$\tau_{\{s\}}^0(1) := \inf\{t \geq 0: X^{(1)}(t) = s\}, \quad \tau_{\{s\}}^0(2) := 0,$$

and let

$$\tau_{\{s\}}^n(l) := \inf\{t \geq \tau_{\{s\}}^{n-1}(l): X^{(l)}(t-) \neq X^{(l)}(t) = s\}, \quad l = 1, 2,$$

denote the n th return time of $X^{(l)}$ to s . Then, if $\tau_{\{s\}}^0(1) = v$, we have $\tau_{\{s\}}^n(1) = v + W + T_{n,s}(1)$ and $\tau_{\{s\}}^n(2) = W + T_{n,s}(2)$, where W denotes the common time spent in states other than s between the first and n th visits to s , and $T_{n,s}(l)$ denotes the total time spent in s by the process $X^{(l)}$ on its first n visits there. Note that $T_{n,s}(l) \sim q_s^{-1}G(n, 1)$, $l = 1, 2$, where $G(n, 1)$ is the gamma distribution with shape parameter n and unit scale parameter, and that the $T_{n,s}(l)$ are independent of W and $\tau_{\{s\}}^0(1)$. Hence,

$$\begin{aligned} d_{TV}(\mathcal{L}(\tau_{\{s\}}^n(1) \mid \tau_{\{s\}}^0(1) = v), \mathcal{L}(\tau_{\{s\}}^n(2))) \\ = d_{TV}(\delta_v * q_s^{-1}G(n, 1), q_s^{-1}G(n, 1)) \\ \leq c_G q_s v n^{-1/2} \end{aligned}$$

for a suitable constant c_G , where δ_v denotes the point mass at v . Hence, for any $n \geq 1$, we can couple $X^{(1)}$ and $X^{(2)}$ by arranging that $\tau_{\{s\}}^n(1) = \tau_{\{s\}}^n(2)$, with the two processes to be run identically thereafter, and the probability of this coupling failing, conditional on $\tau_{\{s\}}^0(1) = v$, is at most $c_G q_s v n^{-1/2}$. Thus, in particular,

$$d_{TV}(\mathcal{L}(X^{(1)}(t)), \mathcal{L}(X^{(2)}(t))) \leq P[\tau_{\{s\}}^n(2) > t] + \frac{c_G q_s n^{-1/2} T}{p}, \tag{2.11}$$

using Lemma 2.1. It now remains to show that we can reach the bound given in the theorem by choosing n almost as a multiple of t .

Now $\tau_{\{s\}}^n(2)$ is a sum of n independent random variables, each with distribution $\mathcal{L}_s(\tau_{\{s\}}^m)$, where $\tau_{\{s\}}^m$ is defined as in Lemma 2.1. By that lemma and Markov's inequality, it follows that

$$P_r \left[\tau_{\{s\}}^m \geq \frac{2T}{p} \right] \leq \frac{1}{2}, \quad r \in C,$$

and, hence, that

$$P_s \left[\tau_{\{s\}}^m \geq \frac{2kT}{p} \right] \leq 2^{-k} = \exp \left\{ -\frac{p \log 2}{2T} \frac{2kT}{p} \right\} \quad \text{for all } k \geq 1.$$

Thus, the distribution $\mathcal{L}_s(\tau_{\{s\}}^m)$ is stochastically bounded above by that of

$$\frac{2T}{p} \left(1 + \frac{1}{\log 2} E \right),$$

where E has a standard exponential distribution. Hence, the distribution of $\tau_{\{s\}}^n(2)$ is stochastically bounded above by that of

$$\frac{2T}{p} \left(n + \frac{1}{\log 2} G_n \right),$$

where $G_n \sim G(n, 1)$. The inequality $P[G_n \geq 2n] \leq (2/e)^n$ thus implies that

$$P \left[\tau_{\{s\}}^n(2) > \frac{2Tn}{p} \left(1 + \frac{2}{\log 2} \right) \right] \leq \left(\frac{2}{e} \right)^n.$$

So, for any fixed t , using $1 + 2/\log 2 \leq 4$, we take $n = n_t := \lfloor tp/8T \rfloor$ in (2.11), giving

$$d_{TV}(\mathcal{L}(X^{(1)}(t)), \mathcal{L}(X^{(2)}(t))) \leq \left(\frac{2}{e} \right)^{n_t} + \frac{c_G q_s n_t^{-1/2} T}{p},$$

from which it follows that, for $t \geq 16T/p$,

$$d_{TV}(\mathcal{L}(X^{(1)}(t)), \mathcal{L}(X^{(2)}(t))) \leq \left(\frac{2}{e} \right)^{pt/16T} + 4c_G \frac{(Tq_s/p)^{3/2}}{\sqrt{q_s t}}. \tag{2.12}$$

We first observe that $\mathcal{L}(X^{(1)}(t)) = m$ for all t . Then we have

$$P[\tau_{\{0\}}(1) \leq t] = 1 - e^{-\lambda_m t} \leq Ut,$$

where $\tau_{\{0\}}(1) := \inf\{t \geq 0 : X^{(1)}(t) = 0\}$. On the event that $X^{(1)}$ and $X^{(2)}$ are successfully coupled at $\tau_{\{s\}}^{n_t} \leq t$, it thus follows that the event that neither hits 0 before t has probability at least $1 - Ut$, and, on this event, $X^{(2)}(t)$ is also the value of an X -process starting in s , since $X^{(2)}$ has had no visits to 0 before t . This, together with (2.12), completes the proof.

Remark. Denoting by $A(\{s\}, \{0\})$ the event that X hits s before 0, the same argument can be used to show that $d_{TV}(\mathcal{L}_k(X(t) \mid A(\{s\}, \{0\})), \mathcal{L}_s(X^m(t)))$ is at most $\eta(t)$ for any $k \in C$, under the conditions of Theorem 2.2. Hence, conditional on the event that X hits s before reaching 0, the distribution of $X(t)$ starting from any $k \in C$ is also close to m for all times t such that

$$\frac{B^2 T}{p} \ll t \ll U^{-1},$$

provided that $UB^2T/p \ll 1$. Thus, the quasi-stationary distribution m is then indeed the appropriate long-time approximation to the distribution of X in C , for times $t \ll U^{-1}$.

Note also that the coupling used in Theorem 2.2 may be very pessimistic, only making use of the residence times in s . For most processes, the variability in the remaining residence times and in the possible sequences of states can be exploited to get sharper bounds. However, in the examples for which we make computations below, the quantity B^2T/p is of only polynomial order in the size of the system, whereas U^{-1} is exponentially large; hence, even this crude estimate is more than adequate.

3. Birth-and-death processes

Consider now a birth-and-death process with $C = \{1, 2, \dots, N\}$ (for $C = \mathbb{N}$, replace N by ∞ in what follows) having birth rates $b_j > 0$, $1 \leq j < N$, with $b_0 = 0$ and $b_N = 0$ if $N < \infty$, and with strictly positive death rates d_j , $j \in C$. It is convenient to introduce the quantities $(\alpha_j, j \in C)$, where $\alpha_1 = 1$ and, for $j > 1$,

$$\alpha_j = \frac{b_1 \cdots b_{j-1}}{d_2 \cdots d_j}.$$

The return process with $\mu = \delta_{\{1\}}$, equivalent to redefining d_1 to be 0, is then recurrent if $\alpha_+ := \sum_{j \geq 1} \alpha_j < \infty$, in which case $\pi^\mu(j) = \alpha_j / \alpha_+$, so that its computation is very easy. We now wish to investigate when this distribution can be used as a reasonable approximation to the effective steady state behaviour of the process.

In order to apply Theorems 2.1 and 2.2, we need to choose a state $s \geq 1$, and find values for p, T, B , and U . For p , let $r_k, k \geq 1$, be the probability that the process starting in k hits s before it hits 0, where $s \geq 1$. If $k > s$ then $r_k = 1$. Otherwise, $r_0 = 0, r_s = 1$, and

$$(b_k + d_k)r_k = b_k r_{k+1} + d_k r_{k-1}, \quad k = 1, 2, \dots, s - 1,$$

leading to $r_k = \sigma_k / \sigma_s$, where

$$\sigma_0 = 0 \quad \text{and} \quad \sigma_k = \sum_{j=1}^k \frac{1}{d_j \alpha_j} \quad \text{for } k = 1, \dots, s.$$

Since σ_k is nondecreasing in k , we can take

$$p = r_1 = \frac{1}{d_1 \sigma_s} \tag{3.1}$$

for any state $s \in C$.

For T , we first note that, for $1 \leq k \leq s, E_k[\tau_{\{s,0\}}]$ is bounded above by the expected time it takes the process, modified so that $d_1 = 0$, to reach s starting from k , and (see Anderson (1991, Chapter 8))

$$E_k[\tau_{\{s\}}] = \sum_{j=k}^{s-1} E_j[\tau_{\{j+1\}}] = \sum_{j=k}^{s-1} \frac{1}{b_j \alpha_j} \sum_{i=1}^j \alpha_i$$

for the modified process. As this quantity is decreasing in k ,

$$T_1 := \max_{1 \leq k \leq s} E_k[\tau_{\{s,0\}}] \leq \sum_{j=1}^{s-1} \frac{1}{b_j \alpha_j} \sum_{i=1}^j \alpha_i < \infty. \tag{3.2}$$

For $k > s$,

$$E_k[\tau_{\{s,0\}}] = E_k[\tau_{\{s\}}] = \sum_{j=s+1}^k E_j[\tau_{\{j-1\}}] = \sum_{j=s+1}^k \frac{1}{d_j \alpha_j} \sum_{i=j}^{\infty} \alpha_i$$

(again see Anderson (1991, Chapter 8)). Since the latter quantity is increasing in k , we may take

$$T := \max(T_1, T_2), \quad \text{where} \quad T_2 := \sum_{j=s+1}^{\infty} \frac{1}{d_j \alpha_j} \sum_{i=j}^{\infty} \alpha_i. \tag{3.3}$$

Note that then Condition A(ii) holds if T_2 , the so-called ‘ D series’, converges, and that $T_2 < \infty$ is a necessary and sufficient condition for a birth-and-death process to have a unique quasi-stationary distribution (see van Doorn (1991, Part 2 of Theorem 3.2)). Note also that $B := Tq_s/p$ can be bounded using (3.1)–(3.3), together with the fact that $q_s = b_s + d_s$.

Finally, the quantity U can be evaluated as

$$U = \frac{d_1}{(b_1 + d_1) E_1[\tau_{\{1,0\}}]} = \frac{d_1}{1 + b_1 E_2[\tau_{\{1\}}]} = \frac{d_1}{\sum_{j=1}^N \alpha_j}, \tag{3.4}$$

because, also from Anderson (1991, Chapter 8),

$$E_i[\tau_{\{i-1\}}] = \frac{1}{d_i \alpha_i} \sum_{j=i}^N \alpha_j = \frac{1}{b_{i-1} \alpha_{i-1}} \sum_{j=i}^N \alpha_j,$$

and, in particular, since $\alpha_1 = 1$,

$$1 + b_1 E_2[\tau_{\{1\}}] = 1 + \sum_{j=2}^N \alpha_j = \sum_{j=1}^N \alpha_j.$$

In order to apply Theorems 2.1 and 2.2 in practice, we need to be able to bound the quantities p, T, B , and U by assigning concrete expressions in terms of the b_j and d_j to replace (3.1)–(3.4). Simple estimates can be derived under the assumptions that the death rates d_j are increasing in j , and that the ratios b_j/d_j are decreasing, with $b_1/d_1 > 1$. If this is the case, define $s \geq 1$ in such a way that $b_s/d_s \geq 1 > b_{s+1}/d_{s+1}$, and let $1 \leq s_1 \leq s < s_2$ be such that

$$\frac{b_{s_1}}{d_{s_1}} =: \rho_1 > 1 > \rho_2 := \frac{b_{s_2}}{d_{s_2}}.$$

Then

$$x_j := \frac{d_{j+1} \alpha_{j+1}}{d_1} = \prod_{l=1}^j \frac{b_l}{d_l}$$

is maximal at $j = s$, and

$$x_j \geq \rho_1^{s_1 \wedge j}, \quad 0 \leq j \leq s, \quad \frac{x_j}{x_l} \leq \rho_2^{j-l}, \quad j \geq l \geq s_2. \tag{3.5}$$

Hence, from (3.1), we have the bound

$$p = 1 / \sum_{j=0}^{s-1} x_j^{-1} = \left\{ \sum_{j=0}^{s_1-1} x_j^{-1} + \sum_{j=s_1}^{s-1} x_j^{-1} \right\}^{-1} \geq \left\{ \frac{\rho_1}{\rho_1 - 1} + (s - s_1) \rho_1^{-s_1} \right\}^{-1},$$

the final inequality following from (3.5). Then, by (3.5) and because d_j is increasing in j , $d_{s_1} \alpha_{j+1} \geq d_{j+1} \alpha_{j+1} \geq d_1 \rho_1^j$ for $0 \leq j < s_1$, and so (3.4) implies that

$$U \leq \left\{ \sum_{j=1}^{s_1} \frac{\alpha_j}{d_1} \right\}^{-1} \leq d_{s_1} (\rho_1 - 1) \rho_1^{-s_1}.$$

For T_2 , we note from (3.5) that

$$\frac{d_i \alpha_i}{d_j \alpha_j} = \frac{x_{i-1}}{x_{j-1}} \leq \rho_2^{(i \vee s_2) - (j \vee s_2)}, \quad i \geq j \geq s + 1,$$

so that, from (3.3),

$$T_2 = \sum_{j=s+1}^{\infty} \frac{1}{d_j \alpha_j} \sum_{i=j}^{\infty} \alpha_i = \sum_{i=s+1}^{\infty} \frac{1}{d_i} \sum_{j=s+1}^i \frac{x_{i-1}}{x_{j-1}} \leq \sum_{i=s+1}^{\infty} \frac{1}{d_i} \left\{ s_2 - s + \frac{1}{1 - \rho_2} \right\}.$$

A similar argument based on (3.2) then gives

$$T_1 \leq \sum_{j=1}^{s-1} \frac{1}{b_j \alpha_j} \sum_{i=0}^j \alpha_i \leq \left\{ s - s_1 + \frac{\rho_1}{\rho_1 - 1} \right\} \sum_{i=1}^{s-1} \frac{1}{b_i}.$$

Thus, if, for instance, d_j grows at most polynomially fast in j , with the sum $\sum_{j \geq 1} d_j^{-1} < \infty$, and if s_1 and $s - s_1$ are large and of comparable size, then T/p is roughly of order s^2 and $B = Tq_s/p$ of polynomial order in s , whereas U is geometrically small with s , making UT/p very small indeed.

More precise calculations for the stochastic logistic model of (1.1) as $A \rightarrow \infty$, with $s = \lfloor \kappa A \rfloor$, give

$$T = O(\log A), \quad p \geq 1 - \frac{d}{b}, \quad U \leq \left\{ 1 + \frac{b-d}{b+d} \right\}^{-\kappa A/2}, \quad B = O(A \log A),$$

so that UT/p is geometrically small in A as $A \rightarrow \infty$. Thus, for the stochastic logistic model, the unique quasi-stationary distribution can be very closely approximated by any return distribution, as long as A is large. Entirely similar estimates are true for the SIS epidemic model, which models the number of susceptibles in a closed population of size N , to be thought of as large but finite. The process is a birth-and-death process on $\{0, 1, \dots, N\}$ having rates

$$b_i := \lambda i \left(1 - \frac{i}{N} \right) \quad \text{and} \quad d_i := \mu i, \quad 0 \leq i \leq N;$$

in this case, UT/p is geometrically small in N if $\mu < \lambda$, and there is a quasi-stationary distribution close to $s := \lfloor N(1 - \mu/\lambda) \rfloor$.

Acknowledgements

We thank the anonymous referee for a number of helpful suggestions. ADB wishes to thank the School of Mathematical Sciences at Monash University, the Australian Research Council Centre of Excellence for Mathematics and Statistics of Complex Systems, and the Institute Mittag-Leffler for their warm hospitality, while part of this work was accomplished.

References

ANDERSON, W. J. (1991). *Continuous-Time Markov Chains*. Springer, New York.
 BARTLETT, M. S. (1960). *Stochastic Population Models in Ecology and Epidemiology*. Methuen, London.
 CATTIAUX, P. *et al.* (2009). Quasi-stationary distributions and diffusion models in population dynamics. *Ann. Prob.* **5**, 1926–1969.

- DARROCH, J. N. AND SENETA, E. (1965). On quasi-stationary distributions in absorbing discrete-time Markov chains. *J. Appl. Prob.* **2**, 88–100.
- EWENS, W. J. (1963). The diffusion equation and a pseudo-distribution in genetics. *J. R. Statist. Soc. B* **25**, 405–412.
- EWENS, W. J. (1964). The pseudo-transient distribution and its uses in genetics. *J. Appl. Prob.* **1**, 141–156.
- FERRARI, P. A., KESTEN, H., MARTÍNEZ, S. AND PICCO, P. (1995). Existence of quasi-stationary distributions. A renewal dynamic approach. *Ann. Prob.* **23**, 501–521.
- LINDVALL, T. (2002). *Lectures on the Coupling Method*, 2nd edn. Dover, Mineola, NY.
- PITMAN, J. W. (1974). Uniform rates of convergence for Markov chain transition probabilities. *Z. Wahrscheinlichkeitsth.* **29**, 193–227.
- STEINSALTZ, D. AND EVANS, S. N. (2004). Markov mortality models: implications of quasistationarity and varying initial distributions. *Theoret. Pop. Biol.* **65**, 319–337.
- VAN DOORN, E. A. (1991). Quasi-stationary distributions and convergence to quasi-stationarity of birth-death processes. *Adv. Appl. Prob.* **23**, 683–700.
- VERHULST, P.-F. (1838). Notice sur la loi que la population poursuit dans son accroissement. *Corresp. Math. Phys.* **10**, 113–121.
- YAGLOM, A. M. (1947). Certain limit theorems of the theory of branching processes. *Dok. Akad. Nauk SSSR* **56**, 795–798.