Statistics and Probability Letters 91 (2014) 41–46

Contents lists available at ScienceDirect

Statistics and Probability Letters

journal homepage: www.elsevier.com/locate/stapro

On finite long run costs and rewards in infinite Markov chains

Hendrik Baumann [a,*], Werner Sandmann [b]

[a] Clausthal University of Technology, Department of Applied Stochastics and Operations Research, 38678 Clausthal-Zellerfeld, Germany

[b] University of Derby, School of Computing and Mathematics, Derby, DE22 1GB, UK

ARTICLE INFO

Article history:
Received 13 June 2013
Received in revised form 1 April 2014
Accepted 4 April 2014
Available online 12 April 2014

Keywords:
Infinite Markov chain
Finite long run cost
Finite long run reward
Foster–Lyapunov-type criterion
Drift condition

ABSTRACT

Conditions for the finiteness of long run costs and rewards associated with infinite recurrent Markov chains that may be discrete or continuous in time are considered. Without resorting to results from the theory of Markov processes on general state spaces we provide instructive proofs in the course of which we derive auxiliary results that are of interest in themselves. Potential applications of the finiteness conditions are outlined in order to elucidate their high practical relevance.

1. Introduction

It is well known that for irreducible recurrent Markov chains, discrete or continuous in time, on a countable, possibly infinite state space $\mathcal{S}$ there exists an invariant measure $\psi = (\psi_i)_{i\in\mathcal{S}}$, which is unique up to a multiplicative constant. Furthermore, denoting irreducible discrete-time and continuous-time Markov chains on $\mathcal{S}$ by $(Y_n)_{n\in\mathbb{N}}$ and $(X_t)_{t\geq 0}$, respectively, for functions $f^{(1)}, f^{(2)} : \mathcal{S} \to \mathbb{R}$ with $\psi\left|f^{(1)}\right| < \infty$, $\psi\left|f^{(2)}\right| < \infty$,

$$\lim_{N\to\infty} \frac{\sum_{n=0}^{N} f^{(1)}(Y_n)}{\sum_{n=0}^{N} f^{(2)}(Y_n)} = \frac{\psi f^{(1)}}{\psi f^{(2)}}, \quad \text{resp.,} \quad \lim_{t\to\infty} \frac{\int_0^t f^{(1)}(X_s)\, ds}{\int_0^t f^{(2)}(X_s)\, ds} = \frac{\psi f^{(1)}}{\psi f^{(2)}}, \tag{1}$$

with probability 1, see Chung (1960, pp. 85–86, 203–209), where $f^{(1)}, f^{(2)}$, written as functions in the usual way or represented as column vectors, may be interpreted as cost or reward functions. This motivates to study conditions for the finiteness of $\psi f = \sum_{i\in\mathcal{S}} \psi_i f(i)$ with $f : \mathcal{S} \to \mathbb{R}$ in the case of irreducible recurrent Markov chains. Both in discrete time and in continuous time we consider 'Foster–Lyapunov-type criteria' involving generalized 'drift conditions'. More precisely, for discrete-time Markov chains $(Y_n)_{n\in\mathbb{N}}$ with transition probability matrix $P = (p_{ij})_{i,j\in\mathcal{S}}$ and a function $g : \mathcal{S} \to \mathbb{R}$ the drift function $d_g : \mathcal{S} \to \mathbb{R}$ is defined by

$$d_g(i) = \mathrm{E}[g(Y_n) - g(Y_{n-1}) | Y_{n-1} = i] = \sum_{j\in\mathcal{S}} p_{ij} g(j) - g(i), \tag{2}$$

* Corresponding author. Tel.: +49 5323722417.
E-mail addresses: hendrik.baumann@tu-clausthal.de, hendrik-baumann@t-online.de (H. Baumann), w.sandmann@derby.ac.uk (W. Sandmann).

http://dx.doi.org/10.1016/j.spl.2014.04.005
0167-7152/

that is, when writing g and d_g in the form of column vectors, $d_g = Pg - g$. For continuous-time Markov chains $(X_t)_{t\geq 0}$ with generator matrix $Q = (q_{ij})_{i,j\in\mathcal{S}}$ and a function $g : \mathcal{S} \to \mathbb{R}$ the drift function $d_g : \mathcal{S} \to \mathbb{R}$ is defined by

$$d_g(i) = \frac{d}{dt}\mathrm{E}[g(X_t)|X_t = i] = \sum_{j\in\mathcal{S}} q_{ij} g(j) \tag{3}$$

that is, when writing g and d_g in the form of column vectors, $d_g = Qg$. In both cases, $d_g(i)$ is the (generalized) drift in state i with respect to g.

Theorem 1. *Let $(Y_n)_{n\in\mathbb{N}}$ be an irreducible discrete-time Markov chain with transition probability matrix $P = (p_{ij})_{i,j\in\mathcal{S}}$, $\mathcal{C} \subset \mathcal{S}$ finite, and let $f, g : \mathcal{S} \to \mathbb{R}_{\geq 0}$ meet the conditions*

(C1) $\forall i \in \mathcal{S} \setminus \mathcal{C} : \ d_g(i) \leq -f(i)$,
(C2) $\forall i \in \mathcal{C} : \ d_g(i) < \infty$,
(C3) $\forall r < \infty : \ |\{i \in \mathcal{S} : \ g(i) \leq r\}| < \infty$.

Then $(Y_n)_{n\in\mathbb{N}}$ is recurrent and for any invariant measure ψ, $\psi f < \infty$.

Theorem 2. *Let $(X_t)_{t\geq 0}$ be an irreducible continuous-time Markov chain with generator matrix $Q = (q_{ij})_{i,j\in\mathcal{S}}$, $\mathcal{C} \subset \mathcal{S}$ finite, and let $f, g : \mathcal{S} \to \mathbb{R}_{\geq 0}$ meet the conditions* (C1)–(C3). *Then Q is regular (it uniquely defines $(X_t)_{t\geq 0}$, the Feller process of Q), and $(X_t)_{t\geq 0}$ is recurrent and for any invariant measure ψ, $\psi f < \infty$.*

Note that the special case of Theorem 1 for $f(i) = 1$ is a famous criterion for positive recurrence of discrete-time Markov chains. With $|\mathcal{C}| = 1$ it is due to Foster (1953), in the slightly more general case of arbitrary finite $\mathcal{C}$ it was proven by Pakes (1969). A similar result is given by Theorem 14.0.1 in Meyn and Tweedie (1993), where (C1), (C2) and $f \geq 1$ are proved to be sufficient for f-ergodicity, that is positive recurrence and $\pi f < \infty$ for the invariant distribution π. Comparing Theorem 1 with this result, we have to emphasize that Theorem 1 only requires nonnegativity of f, that is $f \geq 0$, whereas Theorem 14.0.1 in Meyn and Tweedie (1993) requires $f \geq 1$. Hence, Theorem 1 poses a weaker condition on f. This has two important consequences:

- Theorem 1 does not guarantee *positive* recurrence. For instance, let $\mathcal{S} = \mathbb{N}$, $p_{i,i+1} = p_{i,i-1} = \frac{1}{2}$ for $i \geq 1$, $f = 0$ and $g(i) = i$. Then $d_g(i) = 0$ for $i \geq 1$, and hence, conditions (C1)–(C3) are met, but obviously, we have null recurrence.
- (C3) cannot be omitted: Let P be an arbitrary irreducible transition probability matrix, and choose $g = 1$. Then $d_g(i)$ converges for all $i \in \mathcal{S}$ with $d_g(i) = 0$. Hence, (C1) and (C2) are met with $f = 0$. Since P can be transient, Theorem 1 would not be true without condition (C3).

Likewise, the special case of Theorem 2 for $f(i) = 1$ is a famous criterion for regularity and positive recurrence of continuous-time Markov chains, which is due to Theorem 2.3 in Tweedie (1975). Appropriate functions g as in the theorems are often called Lyapunov functions, the conditions in the theorems are also referred to as Foster–Lyapunov-type criteria and those on d_g as (generalized) drift conditions.

The recurrence of $(Y_n)_{n\in\mathbb{N}}$ follows from Theorem 3.3 in Tweedie (1975), so that an invariant measure ψ exists. Hence, it remains to prove the finiteness of ψf. In principle, one may deduce it from Tweedie (1983, 1988), where discrete-time Markov processes on general state spaces are considered, but we shall provide an alternative instructive proof specifically for countable state spaces in the course of which we derive auxiliary results that are of interest in themselves. Then we prove Theorem 2 via the embedded jump chain.

2. Proof of Theorem 1

First note that when proving positive recurrence, one can exploit the fact that $\lim_{n\to\infty} p_{ij}^{(n)} > 0$ for some $j \in \mathcal{S}$ is sufficient for positive recurrence, where $p_{ij}^{(n)}$ denotes the n-step transition probability from state i to state j, but for proving Theorem 1 with general nonnegative f we cannot use this approach, because even in the case of positive recurrence, where $\lim_{n\to\infty} p_{ij}^{(n)}$ exists and is positive for all $i, j \in \mathcal{S}$, it is not guaranteed that ψf is finite. Alternatively, from Tweedie (1983, Theorem 1), where conditions for the existence of moments with respect to the unique stationary distribution of a discrete-time Markov chain with general state space are given, one may deduce the special case of countable state spaces and then extend it to our case of arbitrary invariant measures.

Instead, we will exploit an interpretation of g as follows. For positive recurrence, $f(i) = 1$, an appropriate function g is given by $g(j) = h_{j,\mathcal{C}} = \mathrm{E}[\tau_{\mathcal{C}}|Y_0 = j]$ where $\tau_{\mathcal{C}}$ is the first passage time to the set $\mathcal{C}$ and thus $h_{j,\mathcal{C}}$ is the mean first passage time from j to $\mathcal{C}$, see Theorem 3.2 in Tweedie (1975). Furthermore, any function g satisfying the conditions of Theorem 1 for $f(i) = 1$ yields upper bounds for the mean first passage times $h_{j,\mathcal{C}} \leq g(j)$. Thus, $h_{j,\mathcal{C}}$ is finite for every $j \in \mathcal{S}$ and $(Y_n)_{n\in\mathbb{N}}$ is positive recurrent. Such an approach for proving Foster–Lyapunov-criteria can for example be found in Breuer and Baum (2005, pp. 30–31). We will appropriately generalize this proof technique. The main idea consists in replacing $\tau_{\mathcal{C}}$ by

$$\tau_{\mathcal{C}}^{(f)} = \sum_{k=0}^{\tau_{\mathcal{C}}-1} f(Y_k), \tag{4}$$