Stationary flows and uniqueness of invariant measures

FRANÇOIS BACCELLI
École Normale Supérieure
Paris, France

TAKIS KONSTANTOPOULOS
Heriot-Watt University
Edinburgh, UK

Abstract

In this short paper, we consider a quadruple $(\Omega, \mathscr{A}, \vartheta, \mu)$, where $\mathscr{A}$ is a σ-algebra of subsets of Ω, and ϑ is a measurable bijection from Ω into itself that preserves the measure μ. For each $B \in \mathscr{A}$, we consider the measure μ_B obtained by taking cycles (excursions) of iterates of ϑ from B. We then derive a relation for μ_B that involves the forward and backward hitting times of B by the trajectory $(\vartheta^n \omega, n \in \mathbb{Z})$ at a point $\omega \in \Omega$. Although classical in appearance, its use in obtaining uniqueness of invariant measures of various stochastic models seems to be new. We apply the concept to countable Markov chains and Harris processes.

Keywords and phrases. STATIONARY FLOWS, INVARIANT MEASURES, UNIQUNESS, HARRIS CHAINS

AMS 2000 *subject classifications.* Primary 28D05,60G10; secondary 60J10,60J05

1   Introduction

This paper was initiated from the following question. It is classical that, for a Markov chain $(X_n, n \geq 0)$ with a countable state space S that possesses a positive recurrent state $b \in S$, there is at least one invariant probability measure $\pi^{(b)}$ on S which is defined by the usual "cycle formula":

$$\pi^{(b)}(A) = \frac{1}{E_b \mathsf{t}_b} E_b \sum_{n=0}^{\mathsf{t}_b - 1} \mathbf{1}(X_n \in A),$$

where t_b is the first return time to b. To ensure that $\pi^{(b)}$ is the only invariant probability measure we need, in addition, to ensure that the only positive recurrent states are those that communicate with b (this holds, for instance, if the chain is irreducible). There are several proofs of uniqueness, ranging from analytic (e.g. by means of the Perron-Frobenius theorem which itself can be proved in a number of ways–see, e.g. Lind and Marcus (1995) for a geometric proof) to probabilistic (e.g. by means of applying the Doeblin coupling construction: this requires, in addition, aperiodicity–see, e.g. Thorisson (2000)). The question we posed is whether there is a way to prove uniqueness directly from the way that $\pi^{(b)}$ is constructed by the cycle formula. If so, can we do this for Markov chains in a general state space? And finally, how "Markovian" is the proof of uniqueness (can the "local" character of definition of π_b be extended to other processes)?

In answering the question, we abstracted the problem and lifted it to a general measurable space $(\Omega, \mathscr{A})$ endowed with a measurable bijective transformation ϑ that preserves some measure μ. The point of view appears to be new, although the tools used below are quite natural in Ergodic Theory and in the construction of Palm Probabilities. The origin of these tools can be traced, as far as we can tell, to a paper by Kac (1947). In Section 2 we define, for each $B \in \mathscr{A}$, the forwards and backwards hitting times of B by the iterates of ϑ (called $T_B, \widetilde{T}_B$, respectively) and the measure

$$\mu_B(A) = \int_B d\mu \sum_{n=0}^{T_B-1} \mathbf{1}_{\vartheta^{-n}A}.$$

Theorem 1 states the basic formula of interest:

$$\mu_B(A) = \mu(A, \widetilde{T}_B < \infty).$$

It can be read as: on the event that B has been visited in the past at least once, the measures μ_B and μ coincide. Thus, if $\mu(B) > 0$, Poincaré's recurrence lemma (recalled as Lemma 1), $\mu_B = \mu$ for all B. In Section 3, we consider a Markov chain on a countable set S. Assuming irreducibility and positive recurrence, the previous observation immediately yields a unique probability measure π on S such that $\pi\mathsf{P} = \pi$, which answers the original question. Finally, in Section 4, we consider a Harris chain and show uniqueness of the invariant probability measure constructed by means of cycles away from a recurrent regeneration set R.

2   The master formula

Let $(\Omega, \mathscr{A})$ be a measurable space and $\vartheta : \Omega \to \Omega$ a measurable bijection. For $A, B \in \mathscr{A}$ define the following functions:

$$T_B \equiv T_B(\omega; \vartheta) := \inf\{n \in \mathbb{N} : \ \vartheta^n \omega \in B\} \tag{1a}$$

$$M_B(A) \equiv M_B(A, \omega; \vartheta) := \sum_{0 \le n < T_B(\omega; \vartheta)} \mathbf{1}(\vartheta^n \omega \in A), \tag{1b}$$

stressing that both take values in $\mathbb{N} \cup \{+\infty\} := \{1, 2, \dots\} \cup \{+\infty\}$, where $\inf \varnothing = +\infty$. The definition of $M_B(A)$ requires giving a meaning to the quantity ϑ^{T_B}. We let $\Omega_B = \{T_B < \infty\}$, and define $\vartheta^{T_B} : \Omega_B \to \Omega_B$ by

$$\left(\vartheta^{T_B}\right)(\omega) := \vartheta^{T_B(\omega)}(\omega), \quad \omega \in \Omega_B.$$

On $\Omega - \Omega_B$, we define ϑ^{T_B} rather arbitrarily, e.g. by letting it to be the identity on it. We can easily see that ϑ^{T_B} is invertible with $\left(\vartheta^{T_B}\right)^{-1} = \vartheta^{-T_B}$, and where ϑ^{-T_B} is defined in a similar way. We shall also need (1a)-(1b) when using ϑ^{-1} in place of ϑ:

$$\widetilde{T}_B \equiv T_B(\omega; \vartheta^{-1}) := \inf\{n \in \mathbb{N} : \ \vartheta^{-n} \omega \in B\}$$

$$\widetilde{M_B}(A) \equiv M_B(A, \omega; \vartheta^{-1}) := \sum_{0 \le n < T_B(\omega; \vartheta^{-1})} \mathbf{1}(\vartheta^{-n} \omega \in A).$$

The interpretation is that $M_B(A)$ evaluated at ω is the number of times the forward trajectory $(\omega, \vartheta\omega, \vartheta^2\omega, \dots)$ visits the set A up to (and not including) the time it visits the set B. Similarly, $\widetilde{M}_B(A)$ refers to the backward trajectory $(\omega, \vartheta^{-1}\omega, \vartheta^{-2}\omega, \dots)$. There is a certain "duality" between $M_B(A)$ and $\widetilde{T}_B$ on one hand and $\widetilde{M}_B(A)$ and T_B on the other, once we integrate against an invariant measure. We discuss this next. Recall first the following standard lemma:

Lemma 1 (Poincaré recurrence). *If the measure μ on $(\Omega, \mathscr{A})$ is preserved by ϑ then, for all $B \in \mathscr{A}$,*

$$\mu(B) = \mu(B, \widetilde{T}_B < \infty) = \mu(B, T_B < \infty). \tag{2}$$

Proof. This follows from

$$\mu(B^c, \widetilde{T}_B = \infty) = \lim_{n \to \infty} \mu(B^c \cap \vartheta B^c \cap \cdots \cap \vartheta^{n-1}B^c)$$
$$= \lim_{n \to \infty} \mu(\vartheta B^c \cap \vartheta^2 B^c \cap \cdots \cap \vartheta^n B^c) = \mu(\widetilde{T}_B = \infty),$$

and similarly for T_B. $\qquad\square$

In other words, $\widetilde{T}_B < \infty$ and $T_B < \infty$, μ-a.e. on B. This is used in proving:

Theorem 1. *If the measure μ on $(\Omega, \mathscr{A})$ is preserved by ϑ, then, for all $A, B \in \mathscr{A}$,*

$$\mu_B(A) := \int_B M_B(A)d\mu = \int_A \mathbf{1}(\widetilde{T}_B < \infty)d\mu, \tag{3a}$$

$$\widetilde{\mu}_B(A) := \int_B \widetilde{M}_B(A)d\mu = \int_A \mathbf{1}(T_B < \infty)d\mu. \tag{3b}$$

Proof. We only need to show the first identity.

$$\int_B M_B(A)d\mu = \mu(A \cap B) + \sum_{n \geq 1} \mu(B \cap \vartheta^{-1}B^c \cap \cdots \cap \vartheta^{-n}B^c \cap \vartheta^{-n}A)$$
$$= \mu(A \cap B) + \sum_{n \geq 1} \mu(\vartheta^n B \cap \vartheta^{n-1}B^c \cap \cdots \cap B^c \cap A)$$
$$= \mu(A \cap B) + \sum_{n \geq 1} \mu(\widetilde{T}_B = n, A \cap B^c)$$
$$= \mu(A \cap B) + \mu(A \cap B^c, \widetilde{T}_B < \infty) = \mu(A, \widetilde{T}_B < \infty), \tag{4}$$

where the Poincaré recurrence formula (and more precisely its consequence that $\mu(A \cap B, \widetilde{T}_B = \infty) = 0$) was used to obtain the last equality. $\qquad\square$

Proposition 1 (strong invariance). *If the measure μ on $(\Omega, \mathscr{A})$ is preserved by ϑ, then, its restriction $\mu(\cdot \cap B)$ on some $B \in \mathscr{A}$ is preserved by ϑ^{T_B} and by $\vartheta^{\widetilde{T}_B}$, i.e., for all $A, B \in \mathscr{A}$,*

$$\mu(B \cap \vartheta^{-T_B}A) = \mu(B \cap \vartheta^{-\widetilde{T}_B}A) = \mu(A \cap B).$$

Note: The terminology *strong invariance* is by analogy to the strong Markov property.

Proof. Since, due to the Poincaré recurrence, $\mu(B \cap \vartheta^{-T_B} A \cap (\Omega - \Omega_B)) = 0$, we have

$$
\begin{aligned}
\mu(B \cap \vartheta^{-T_B} A) &= \sum_{n=1}^{\infty} \mu(B \cap \vartheta^{-n} A, T_B = n) \\
&= \sum_{n=1}^{\infty} \mu(B \cap \vartheta^{-n} A \cap \vartheta^{-1} B^c \cap \cdots \cap \vartheta^{-(n-1)} B^c \cap \vartheta^{-n} B) \\
&= \sum_{n=1}^{\infty} \mu(B \cap \vartheta^{-1} B^c \cap \cdots \cap \vartheta^{-(n-1)} B^c \cap \vartheta^{-n}(A \cap B)) \\
&= \sum_{n=1}^{\infty} \mu(\vartheta^n B \cap \vartheta^{n-1} B^c \cap \cdots \cap \vartheta B^c \cap A \cap B) \\
&= \sum_{n=1}^{\infty} \mu(\widetilde{T}_B = n, A \cap B) = \mu(\widetilde{T}_B < \infty, A \cap B) = \mu(A \cap B),
\end{aligned}
$$

where the latter equality again follows from the Poincaré recurrence (2). The second assertion is proved in the same manner. □

Proposition 2. *If the measure μ on $(\Omega, \mathscr{A})$ is preserved ϑ, then, for all $B \in \mathscr{A}$, the measures $\mu_B(\cdot)$, $\widetilde{\mu}_B(\cdot)$, defined by (3a), (3b), respectively, are also preserved by ϑ.*

Proof. Note that
$$
M_B(\vartheta^{-1} A) - M_B(A) = \mathbf{1}_{\vartheta^{-T_B} A} - \mathbf{1}_A.
$$
Using this and Proposition 1 we obtain
$$
\mu_B(\vartheta^{-1} A) - \mu_B(A) = \int_B M_B(\vartheta^{-1} A) d\mu - \int_B M_B(A) d\mu = \mu(B \cap \vartheta^{-T_B} A) - \mu(B \cap A) = 0.
$$
□

Some remarks:

(i) Since $M_B(\Omega) = T_B$, $\widetilde{M}_B(\Omega) = \widetilde{T}_B$, we have, from Theorem 1,
$$
\int_B T_B d\mu = \mu(\widetilde{T}_B < \infty), \qquad \int_B \widetilde{T}_B d\mu = \mu(T_B < \infty). \tag{5}
$$
Thus, if $\mu = P$ is a probability measure and if E denotes integration with respect to P, then
$$
E T_B \mathbf{1}_B = P(\widetilde{T}_B < \infty) \le 1.
$$
If, in addition, $P(B) > 0$ then $P(\widetilde{T}_B < \infty) \ge P(\widetilde{T}_B < \infty, B) = P(B) > 0$ and so
$$
E(T_B \mid B) = \frac{1}{P(B \mid \widetilde{T}_B < \infty)},
$$
where, as usual, $E(T_B \mid B) = \frac{E T_B \mathbf{1}_B}{P(B)}$. This is slightly more general than Kac' formula (Kac (1947)). Similar formula holds, of course, for $E(\widetilde{T}_B \mid B)$:
$$
E(\widetilde{T}_B \mid B) = \frac{1}{P(B \mid T_B < \infty)}.
$$

4

(ii) Since $T_B \geq 1$, $\widetilde{T}_B \geq 1$, (5) implies that:

$$\mu(B) \leq \mu(\widetilde{T}_B < \infty) \wedge \mu(T_B < \infty).$$

This leads to the following equivalences:

$$\mu(B) > 0 \iff \mu(T_B < \infty) > 0 \iff \mu(\widetilde{T}_B < \infty) > 0.$$

Indeed, if $\mu(B) > 0$ then $\mu(\widetilde{T}_B < \infty) > 0$ from the last inequality. Conversely, if $\mu(B) = 0$ then (4) shows that $\mu(\widetilde{T}_B < \infty) = 0$.

(iii) The function $B \mapsto \mu_B(A)$ can be thought of as a pre-capacity. Indeed, let

$$\Psi(\omega) := \{\vartheta^{-1}\omega, \vartheta^{-2}\omega, \ldots\}$$

and consider it as a random set. Then

$$\mu_B(A) = P(A, \Psi \cap B \neq \varnothing)$$

is the pre-capacity functional of the random set Ψ (see Molchanov (2005).) We avoid using the terminology capacity because there no topological properties of Ψ are introduced. An interesting problem would be to investigate properties of the function $\mu_B(A)$ jointly in A, B.

(iv) Theorem 1 and Proposition 2 should of course be linked to the cycle formula of Palm calculus, and Proposition 1 to the invariance of the Palm measure. The main point here is that within this discrete time setting, there is no need to invoke the general theory (Baccelli and Brémaud (2003)).

(v) Some results do not require the invertibility of ϑ. For instance, $\mu(B) = \mu(B, T_B < \infty) = \mu(B, \cup_{n=1}^{\infty} \vartheta^{-n}B)$ holds for any μ-preserving measurable map ϑ (Lemma 1). However, the main formulae (3a)-(3b) that exhibit the "duality" between forward and backward iterates of ϑ, do require invertibility. On the other hand, even without using Theorem 1 and Propositions 1-2, we can show that the measure $\nu_B(A) := \mu(A, T_B < \infty)$ satisfies $\nu_B(\vartheta A) = \nu_B(A)$ directly. To do this, note that $\{T_B \circ \vartheta < \infty\} = B \cup \{T_B < \infty\}$ and write

$$\begin{aligned}
\mu(\vartheta A, T_B < \infty) &= \mu(A, T_B \circ \vartheta < \infty) \\
&= \mu((A \setminus B) \cup (A \cap B), \ B \cup \{T_B < \infty\}) \\
&= \mu(A \setminus B, T_B < \infty) + \mu(A \cap B, T_B < \infty) \\
&= \mu(A, T_B < \infty).
\end{aligned}$$

This, incidentally, gives a second proof of Proposition 2.

3 Uniqueness in Markov chains

Suppose that $\mathsf{P} = [p_{i,j}]$ is a stochastic matrix on a countable state space $S = \{a, b, c, \ldots, i, j, \ldots\}$, i.e.

$$p_{i,j} \geq 0, \quad \sum_{k \in S} p_{i,k} = 1, \quad i, j \in S.$$

Assume that it is

(i) irreducible (each i *communicates* with each j in S),

(ii) positive recurrent (starting from some i the expected return time to i has finite expectation).

These properties depend entirely on the matrix P. It is classical that:

Theorem 2. *If (i) and (ii) hold then there is a unique probability π on S such that $\pi \mathsf{P} = \pi$.*

We wish to show this by using the idea developed in the previous section.

Proof of existence It is uniqueness that is novel here. Existence of such a π is immediately answered by the "cycle formula": Let $(X_0, X_1, \ldots)$ be a realisation of the Markov chain with transition probability matrix P. Fix some state b, let

$$\mathsf{t}_b := \inf\{n \geq 1 : X_n = b\},$$

and define the probability $\pi^{(b)}$ on S by

$$\pi^{(b)}(a) = \frac{E_b \sum_{n=0}^{\mathsf{t}_b - 1} \mathbf{1}(X_n = a)}{E_b \mathsf{t}_b}, \quad a \in S,$$

where E_b is expectation conditional on $X_0 = b$. That this $\pi^{(b)}$ is an invariant probability measure (satisfies $\pi^{(b)} \mathsf{P} = \pi^{(b)}$ is standard (see, e.g. Brémaud (1999)). It is important to note that $\pi^{(b)}$ depends entirely on the stochastic matrix P only. $\qquad\square$

Proof of uniqueness To show uniqueness, we work at the level of sequences, i.e. with the space $\Omega = S^{\mathbb{Z}}$, whose elements are denoted by $\omega = (\omega_n, n \in \mathbb{Z})$, equipped with the cylinder σ-algebra $\mathscr{A}$. We consider the natural shift

$$\vartheta : (n \mapsto \omega_n) \mapsto (n \mapsto \omega_{n+1}),$$

which is obviously $\mathscr{A}$-measurable and invertible. We are thus in the setup of the earlier section. Consider a probability π on S satisfying $\pi \mathsf{P} = \pi$, and let P be the probability measure on $(S^{\mathbb{Z}}, \mathscr{A})$ defined by

$$P(\{\omega \in \Omega : \omega_m = i_m, \ldots, \omega_n = i_n\}) = \pi(i_m) p_{i_m, i_{m+1}} \cdots p_{i_{n-1}, i_n},$$
$$i_m, \ldots, i_n \in S, \quad m, n \in \mathbb{Z}, \quad m \leq n. \quad (6)$$

6

Consider also the random variables

$$X_n(\omega) := \omega_n, \quad \omega \in \Omega, \quad n \in \mathbb{Z}.$$

Under P, the sequence (X_n) is a Markov chain with transition probability matrix P. Clearly, the measure P is preserved by ϑ and, by Proposition 2, so are the measures

$$P_B(\cdot) = E\mathbf{1}_B M_B(\cdot) = \int_B M_B(\cdot) dP,$$

where T_B, M_B are given by (1a)-(1b), for any $B \in \mathscr{A}$. Fix some $b \in S$, and consider the set

$$B = \{\omega \in \Omega : \ \omega_0 = b\}.$$

Observe that

$$\mathsf{t}_b(\omega) := \inf\{n \geq 1 : \omega_n = b\} = T_B(\omega), \quad \widetilde{\mathsf{t}}_b(\omega) := \inf\{n \geq 1 : \omega_{-n} = b\} = \widetilde{T}_B(\omega).$$

By Theorem 1,

$$P_B(A) = E\mathbf{1}_B M_B(A) = P(A, \widetilde{T}_B < \infty) = P(A, \widetilde{\mathsf{t}}_b < \infty), \quad A \in \mathscr{A}. \tag{7}$$

By (i) and (ii) we have $P(\mathsf{t}_b < \infty) = 1$, $P(\widetilde{\mathsf{t}}_b < \infty) = 1$, and so (7) yields

$$P_B(A) = P(A), \quad A \in \mathscr{A},$$

and $E\mathbf{1}_B T_B = 1$. Therefore,

$$P(A) = P_B(A) = E\mathbf{1}_B M_B(A) = \frac{E\mathbf{1}_B M_B(A)}{E\mathbf{1}_B T_B} = \frac{E_b \sum_{n=0}^{\mathsf{t}_b-1} \mathbf{1}(\vartheta^n \omega \in A)}{E_b \mathsf{t}_b}.$$

So, if we pick

$$A := \{\omega \in \Omega : \ \omega_0 = a\},$$

we conclude that $\pi(a) = \pi^{(b)}(a)$ for all $a \in S$. Thus, an arbitrary invariant probability measure π must be equal to the specific measure $\pi^{(b)}$; whence the uniqueness. $\square$

Remarks:

(i) The last argument directly proves that

$$\frac{E_b \sum_{n=0}^{\mathsf{t}_b-1} \mathbf{1}(X_n = a)}{E_b \mathsf{t}_b} = \frac{E_c \sum_{n=0}^{\mathsf{t}_c-1} \mathbf{1}(X_n = a)}{E_c \mathsf{t}_c},$$

the so-called exchange formula of (discrete-index) Palm theory (see also Konstantopoulos and Zazanis (1995)).

(ii) Only the existence proof used the Markov property. The uniqueness proof was at the level of stationary processes.

(iii) In essence, uniqueness follows from the following two facts:

- Unique determination of the Palm measure: thanks to the Markov setting considered here, the Palm law of a cycle starting from a given state until the chain returns to this state is uniquely determined by the transition matrix;

- Slivnyak's inverse construction: this construction shows that the stationary law of a point process is fully determined by its Palm measure. (See Slivnyak (1962).)

Again, the main point here is that there is no need to invoke the general theory.

(iv) The same argument can be used to show the weaker result:

Theorem 3. *Suppose that (ii) holds (every state is positive recurrent) Let $S = \cup_{i \geq 1} S_i$, be the decomposition of S into its irreducible components. Let $b_i \in S_i$, for all $i \geq 1$. Then every probability π on S such that $\pi \mathsf{P} = \pi$ is a convex combination of the measures $\pi^{(b_i)}$.*

4   Uniqueness in Harris chains

The method explained above can also be applied to yield a proof of uniqueness for the invariant probability measure of a positive Harris recurrent chain.

A Markov process (X_n) with values in a Polish space $(S, \mathscr{S})$ and transition kernel

$$K(x, \cdot) = P_x(X_1 \in \cdot)$$

is called *Harris recurrent* or, simply, Harris chain (Asmussen (2003)) if it possesses a recurrent regeneration set $R \in S$. This means that

(i)
$$P_x(\mathsf{t}_R < \infty) = 1, \quad x \in S,$$

where
$$\mathsf{t}_R := \inf\{n \in \mathbb{N} : X_n \in R\};$$

(ii) there is a probability measure λ on $(S, \mathscr{S})$, an $\varepsilon > 0$, and $\ell \in \mathbb{N}$, such that

$$K^\ell(x, \cdot) \geq \varepsilon \lambda(\cdot), \quad x \in \mathbb{R},$$

where
$$K^\ell(x, \cdot) = P_x(X_\ell \in \cdot).$$

The chain is called *positive Harris recurrent* if, in addition to (i) and (ii) we also have

(iii)
$$E_\lambda \mathsf{t}_R < \infty,$$

where, as usual, E_λ denotes expectation with respect to $P_\lambda(\cdot) := \int_S \lambda(dx) P_x(\cdot)$.

We here give a proof of the following:

Theorem 4. *A positive Harris recurrent chain possesses a unique invariant probability measure.*

Note that this theorem is proved in the paper of Athreya and Ney (1978) by different methods and only in the case $\ell = 1$. There is a substantial difference between the $\ell = 1$ and $\ell > 1$ cases in that the cycles defined by the iterates of the stopping time $\mathfrak{t}$ (see (11) below) are not independent.

Proof of Theorem 4. Existence is standard (see Asmussen (2003)) and requires construction of the chain on a suitable probability space. We repeat the construction here. In addition to the chain, we consider a sequence (ζ_n) of i.i.d. Bernoulli random variables taking values 1 or 0 with probability ε or $1 - \varepsilon$ respectively. Informally, whenever $X_n \in R$ distribute $X_{n+\ell}$ according to λ if $\zeta_n = 1$ or according to $\dfrac{K^\ell(x, \cdot) - \varepsilon\lambda(\cdot)}{1 - \varepsilon}$ if $\zeta_n = 0$, and, conditional on $(X_n, X_{n+\ell})$, distribute $(X_{n+1}, \ldots, X_{n+\ell-1})$ by respecting the given Markov kernel. Otherwise, if $X_n \notin R$, then ignore ζ_n and continue the chain as usual. Formally, we define an ℓ-th order Markov chain (X_n, ζ_n) with values in $S \times \{0, 1\}$ via the following: Let $G(dx_1, \ldots, dx_{\ell-1} | x, y)$ be the conditional distribution of $(X_1, \ldots, X_{\ell-1})$ given that $X_0 = x, X_\ell = y$, i.e.

$$G(dx_1, \ldots, dx_{\ell-1} | x, y) := \frac{K(x, dx_1) \cdots K(x_{\ell-2}, dx_{\ell-1}) K(x_{\ell-1}, dy)}{\int_{S^{\ell-1}} K(x, dx_1') \cdots K(x_{\ell-2}', dx_{\ell-1}') K(x_{\ell-1}', dy)}, \qquad (8)$$

where the integration in the denominator is with respect to the variables $x_1', \ldots, x_{\ell-1}'$ and the ratio is to be understood as a Radon-Nikodým derivative with respect to y. Then let

$$P(X_{n+i} \in dx_i, 1 \le i \le \ell \mid X_n = x, \zeta_n = \alpha)$$

$$= \begin{cases} \lambda(dx_\ell) G(dx_1, \ldots, dx_{\ell-1} | x, x_\ell), & \text{if } x \in R, \ \alpha = 1 \\ \dfrac{K^\ell(x, dx_\ell) - \varepsilon\lambda(dx_\ell)}{1 - \varepsilon} G(dx_1, \ldots, dx_{\ell-1} | x, x_\ell), & \text{if } x \in R, \ \alpha = 0 \\ K(x, dx_1) \cdots K(x_{\ell-1}, dx_\ell), & \text{otherwise} \end{cases} \qquad (9)$$

and finally require that, for all n,

$$P(X_{n+i} \in dx_i, \zeta_i = \alpha_i, 1 \le i \le \ell \mid X_m, \zeta_m, m \le n)$$
$$= p(\alpha_1) \cdots p(\alpha_\ell) \ P(X_{n+i} \in dx_i, 1 \le i \le \ell \mid X_n, \zeta_n), \quad (10)$$

where $\alpha_i \in \{0, 1\}$ and $p(0) := \varepsilon$, $p(1) := 1 - \varepsilon$.

It is easy to see that (X_n) is a realisation of the Harris chain with the given transition kernel K, and that (ζ_n) is an i.i.d. sequence; the two sequences are dependent.

Consider
$$\mathfrak{t} := \inf\{n : X_{n-\ell} \in R, \ \zeta_{n-\ell} = 1\}, \qquad (11)$$

9

(so that X_t has distribution λ) and define

$$\pi(\cdot) := \frac{E_\lambda \sum_{n=0}^{t-1} \mathbf{1}(X_n \in \cdot)}{E_\lambda t}, \tag{12}$$

It is now standard to check that $\pi(\cdot)$ is an invariant probability measure for the chain (X_n).

To prove uniqueness, we shall again consider the same construction defined by (8), (9) and (10), and, in addition, we shall assume that the chain is stationary and therefore defined over the index set $\mathbb{Z}$. Specifically, our probability space is $\Omega = (S \times \{0,1\})^{\mathbb{Z}}$, equipped with the natural cylinder σ-algebra $\mathscr{A}$. A typical element of Ω is denoted by $\omega = ((x_n, \zeta_n), n \in \mathbb{Z})$. The shift is again the natural one:

$$\vartheta : (n \mapsto (x_n, \zeta_n)) \mapsto (n \mapsto (x_{n+1}, \zeta_{n+1})).$$

The probability measure P on $(\Omega, \mathscr{A})$ is such that it makes the coordinate process an ℓ-th order Markov chain with transition kernel defined through (8), (9) and (10), and is invariant under ϑ. (Thus, we have created a setup $(\Omega, \mathscr{A}, \vartheta, P)$, as in Section 2, where P plays the rôle of μ and, here, $P(\Omega) = 1$.) We now prove that there can be only one such P. To this end, let

$$B := \{\omega = (x, \zeta) \in \Omega : \ x_{-\ell} \in R, \ \zeta_{-\ell} = 1\}.$$

By our assumptions, $P(T_B < \infty) = 1$, $P(\widetilde{T}_B < \infty) = 1$. By Theorem 1,

$$P_B(A) = E\mathbf{1}_B M_B(A) = P(A, \widetilde{T}_B < \infty) = P(A), \quad A \in \mathscr{A},$$

and $E\mathbf{1}_B T_B = 1$. But

$$P_B(A) = \frac{E\mathbf{1}_B M_B(A)}{E\mathbf{1}_B T_B} = \frac{E\left[\sum_{n=0}^{t-1} \mathbf{1}(\vartheta^n \omega \in A) \mid x_{-\ell} \in R, \zeta_{-\ell} = 1\right]}{E[t \mid x_{-\ell} \in R, \zeta_{-\ell} = 1]} = \frac{E_\lambda \sum_{n=0}^{t-1} \mathbf{1}(\vartheta^n \omega \in A)}{E_\lambda t},$$

since, by construction, $P(C|x_{-\ell} \in R, \zeta_{-\ell} = 1) = P_\lambda(C)$ for any C in the σ-algebra generated by $(\omega_n, n \geq 0)$. Taking $A := \{\omega = (x, \zeta) \in \Omega : \ x_0 \in \cdot\}$ we conclude that any ϑ-invariant probability measure P that preserves the given Markovian structure must have a marginal given by (12). This proves uniqueness. $\qquad \square$

Final note: The proof of uniqueness, again, uses arguments that do not rely on the Markov property. As such, it would be worth exploiting it further in stochastic scenarios with absence of Markovian property.

References

ASMUSSEN, S. (2003). *Applied Probability and Queues*, 2nd ed. Springer-Verlag.

ATHREYA, K.B. AND NEY, P. (1978). A new approach to the limit theory of recurrent Markov chains. *Trans. Amer. Math. Soc.* **245**, 493-501.

BACCELLI, F. AND BRÉMAUD, P. (2003). *Elements of Queueing Theory.* Springer-Verlag.

BRÉMAUD, P. (1999). *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues.* Springer Verlag.

KAC, M. (1947). On the notion of recurrence in discrete stochastic processes. *Bull. AMS* **53**, 1002-1010.

KONSTANTOPOULOS, T. AND ZAZANIS, M. (1995). A discrete time proof of Neveu's exchange formula. *J. Appl. Probability* **32**, 917-921.

LIND, D. AND MARCUS, B. (2000). *An Introduction to Symbolic Dynamics and Coding.* Cambridge University Press.

MOLCHANOV, I. (2005). *Theory of Random Sets.* Springer-Verlag.

SLIVNYAK, I.M. (1962). Some properties of stationary flows of homogeneous random events. *Th. Prob. Appl.* **7**, 336-341.

THORISSON, H. (2000). *Coupling, Stationarity, and Regeneration.* Springer.

AUTHORS' ADDRESSES:

FRANÇOIS BACCELLI
DÉPARTEMENT D'INFORMATIQUE
ÉCOLE NORMALE SUPÉRIEURE
45 RUE D'ULM
F-75230 PARIS CEDEX 05, FRANCE
E-MAIL: `Francois.Baccelli@ens.fr`

TAKIS KONSTANTOPOULOS
SCHOOL OF MATHEMATICAL SCIENCES
HERIOT-WATT UNIVERSITY
EDINBURGH EH14 4AS, UK
E-MAIL: `takis@ma.hw.ac.uk`