


Martin Jacobsen James W. Pitman

Birth, Death, and Conditioning of Markov Chains

Preprint
November
1975

17

Institute of Mathematical Statistics
University of Copenhagen



M. Jacobsen and J.W. Pitman^{*}

BIRTH, DEATH, AND CONDITIONING OF MARKOV CHAINS

Preprint 1975 No. 17

INSTITUTE OF MATHEMATICAL STATISTICS
UNIVERSITY OF COPENHAGEN

November 1975

^{*} Department of Statistics, University of California, Berkeley.

BIRTH, DEATH, AND CONDITIONING OF MARKOV CHAINS[†]

by

M. Jacobsen
University of Copenhagen

and

J.W. Pitman
University of California
Berkeley, California

1. INTRODUCTION

Given a Markov chain X_0, X_1, \dots with stationary transition probabilities, we investigate random times τ with the property that the joint distribution of the pre- τ fragment $(X_0, \dots, X_{\tau-1})$ and the post- τ fragment $(X_\tau, X_{\tau+1}, \dots)$ can be described by saying that one or other of these fragments is Markovian with stationary transition probabilities, and that the two fragments are conditionally independent given the position of the inner endpoint of the Markovian fragment at $\tau-1$ or τ . Such a description of the joint law of the pre- τ and post- τ processes for a random time τ will be called a path decomposition. For some examples of more sophisticated path decompositions which provided motivation for the present study see Williams [8], [9], Jacobsen [2], Pitman [4], [5], Pittenger and Shih [6]. Following Meyer, Smythe and Walsh [3] we refer to those random times τ for which the post- τ fragment is Markov as birth times, and to those for which the pre- τ fragment is Markov as death times. We show that for discrete time Markov chains with countable state space the analogues of the types of birth times and death times considered by Meyer, Smythe and Walsh for continuous time processes, namely optional, cooptional, terminal and coterminial times, all admit the additional conditional independence property described above, and that from these special types of random times it

[†]This research was prepared with the support of National Science Foundation Grant MPS 74-18967 and the Danish Natural Science Research Council.

is possible to construct the most general random times determined by the evolution of the Markov chain which allow this kind of path decomposition.

To make things precise we assume that $(X_n, n \in \mathbb{N})$ is the coordinate process defined on the space Ω of all sequences in a countable set J indexed by the non-negative integers \mathbb{N} , equipped with the usual product σ -field F . We say that a probability P on (Ω, F) is Markov, or Markov(p), if P is the distribution of a Markov chain with stationary transition probabilities p , i.e. if under P the sequence (X_n) is itself such a Markov chain. Background on this framework may be found in Freedman [1], and fuller definitions follow at the end of this introduction.

A random time $\tau = \tau(\omega)$ is now an F -measurable function of sequences $\omega \in \Omega$ with values in the extended time set $\mathbb{N} \cup \{\infty\}$. Given a Markov(p) probability P and a random time τ , we say that τ is a birth time for P if the P -distribution of the post- τ process is Markov(q) for some transition matrix q , and say that τ is a regular birth time for P if in addition the pre- τ and post- τ processes are conditionally independent given X_τ on $(\tau < \infty)$. Put another way, τ is a regular birth time for P if there is a transition matrix q such that conditional on the pre- τ process, $\tau < \infty$, and $X_\tau = x$ (for each state x) the P distribution of the post- τ process is Markov(q) with starting state x . According to the strong Markov property each optional (stopping) time τ is a regular birth time for every Markov probability P , and in this case $q = p$. If all states are recurrent it will be seen that every regular birth time is a.s. equal to an optional time, but if there are transient states there will usually be many regular birth times τ for which q differs from p , e.g. the last time τ that a certain set of states H is visited, when the post- τ process is like the original process conditioned

never to hit H .

It turns out quite generally that the Markov chain which emerges at a regular birth time τ can be described by conditioning a Markov chain with the same transition probabilities as the original. With this in mind we determine in Section 2 the collection of all events $C \in \mathcal{F}$ with the property that when a Markov probability P on (Ω, \mathcal{F}) is conditioned on C , another Markov probability results. The result of Section 2 is then applied in Section 3 to give a complete description of all regular birth times for a Markov probability P . It is shown that there is a class of random times B with the property that for each Markov probability P

- (i) every $\tau \in B$ is a regular birth time for P
- (ii) every regular birth time for P is P a.s. equal to a random time in B .

This canonical collection of regular birth times may be roughly described as comprising 'optional times after coterminal times'. It is interesting that the conditional independence hypothesis involved in regularity is quite essential for this type of result. We show that there exists no such canonical collection of plain birth times by exhibiting two Markov probabilities P and Q with the same null sets together with a random time τ which is a birth time for P but not for Q .

In Sections 4 and 5 we consider death times. We say that τ is a death time for P if the P distribution of the pre- τ process $(X_0, \dots, X_{\tau-1})$ is Markov(r) for some sub-stochastic transition matrix r , and say that τ is a regular death time for P if in addition the pre- τ and post- τ processes are conditionally independent given $X_{\tau-1}$ on $(0 < \tau < \infty)$. Since the time reversal of a Markov fragment with finite lifetime and stationary transition probabilities is again Markov with stationary transition probabilities, the regularity

condition for a death time is equivalent to demanding that there is some sub-stochastic transition matrix \hat{r} such that conditional on the post- τ process, $0 < \tau < \infty$ and $X_{\tau-1} = x$ (for each state x), the reversed pre- τ fragment $(X_{\tau-1}, X_{\tau-2}, \dots, X_0)$ is Markov(\hat{r}) with starting state x . Thus the notion of a regular death time may be viewed as the dual under time reversal to the notion of a regular birth time.

In Section 5 we prove the existence of a canonical class \mathcal{D} of regular death times which roughly speaking comprises 'co-optional times prior to terminal times'. This result is like a dual to the existence of the class \mathcal{B} of regular birth times, but owing to the impossibility of reversing on $(\tau = \infty)$ we are unable to bridge between the two results by any direct use of time reversal. For the death time theorem we instead make use of a new method developed in Section 4, exploiting a functional equation satisfied by certain conditional probabilities associated with any death time, regular or not. Once again we show that there is no canonical collection of plain death times.

Section 6 is devoted to random times which are both regular birth times and regular death times. We show that for nice transition matrices p these times are essentially either terminal times or coterminial times, and give a detailed description of the associated path decompositions. Finally, in Section 7 we discuss possible extensions of our results to Markov processes with more general time set or state space.

We set out now the basic notation and conventions which will be used throughout. We take it right from the start that our countable state space J contains a conventional coffin state Δ , but Δ is for use only after killing operations. Except where otherwise specified we assume that we are given a fixed Markov probability P on the sequence space (Ω, \mathcal{F}) , with initial

distribution $\lambda = (\lambda(x), x \in J)$ and transition probabilities $p = (p(x,y), x,y \in J)$ which are arbitrary subject to the coffin state conventions that $\lambda(\Delta) = 0$, $p(\Delta, \Delta) = 1$, $p(x, \Delta) = 0$, $x \neq \Delta$. Thus P can be any Markov probability concentrating on the space $\Omega_0 \subset \Omega$ of all sequences in the state space $J \setminus \{\Delta\}$. For $x \in J$ we denote by P^x the probability on (Ω, F) which is Markov(p) with starting state x : thus $P = \sum_{x \in J} \lambda(x) P^x$.

For $n \in \bar{N} = N \cup \{\infty\}$ we define the coordinate maps $X_n: \Omega \rightarrow J$, killing operators $K_n: \Omega \rightarrow \Omega$ and shift operators $\theta_n: \Omega \rightarrow \Omega$ as follows: for $\omega = (\omega_0, \omega_1, \dots) \in \Omega$, $n \in N$,

$$\begin{aligned} X_n(\omega) &= \omega_n, \\ K_n(\omega) &= (\omega_0, \dots, \omega_{n-1}, \Delta, \Delta, \dots), \\ \theta_n(\omega) &= (\omega_n, \omega_{n+1}, \dots), \end{aligned}$$

while for $n = \infty$

$$X_\infty(\omega) = \Delta, \quad K_\infty(\omega) = \omega, \quad \theta_\infty(\omega) = \omega_\Delta,$$

where $\omega_\Delta = (\Delta, \Delta, \dots)$ is the dead sequence. For a random time $\tau: \Omega \rightarrow \bar{N}$ we define F -measurable mappings X_τ , K_τ and θ_τ in the obvious way: e.g. $X_\tau(\omega) = X_{\tau(\omega)}(\omega)$, $\omega \in \Omega$. Thus X_τ gives the position of the process (X_n) at time τ , K_τ describes the strictly pre- τ fragment $(X_0, \dots, X_{\tau-1})$ by identifying it with the more manageable process $(X_0, \dots, X_{\tau-1}, \Delta, \Delta, \dots)$, while θ_τ describes the post- τ fragment $(X_\tau, X_{\tau+1}, \dots)$. For $n \in N$ we define F_n to be the sub- σ -field of F generated by X_0, \dots, X_n , and denote by A_n the countable collection of all atoms of F_n , i.e. all events A of the form $A = (X_k = x_k, 0 \leq k \leq n)$ for some $x_0, \dots, x_n \in J$. For a random time τ define F_τ , the σ -field of events up to and including time τ , to be the σ -field

generated by $K_{\tau+1}$. This agrees with the usual definition for an optional time τ , and especially $F_\tau = F_n$ for the constant time $\tau = n$. We are only ever interested in the trace of F_τ on the space Ω_0 of sequences avoiding the coffin state Δ , and we find that for each $n \in \bar{\mathbb{N}}$ the trace of F_τ on the event $\Omega_0(\tau=n)$ is identical to the trace of F_n on $\Omega_0(\tau=n)$, and that the event $\Omega_0(\tau<\infty)$ is the union of the countable collection $\{\Omega_0 A(\tau=n), A \in \mathcal{A}_n, n \in \mathbb{N}\}$ of atoms of F_τ , where here and throughout, AB stands for the intersection of two events A and B in F .

2. CONDITIONED MARKOV CHAINS

In this section we solve the following problem: given a Markov chain with stationary transition probabilities, on what events determined by the evolution of the Markov chain can one condition to obtain a new Markov chain with stationary transition probabilities? For any probability P on (Ω, F) and $C \in F$ with $P(C) > 0$ let P_C or $P(\cdot|C)$ denote the probability on (Ω, F) obtained by conditioning on C :

$$P_C(F) = P(F|C) = P(FC)/P(C), \quad F \in F.$$

Thus the problem becomes: given that P is Markov, for which $C \in F$ is P_C again Markov? We start by defining various collections of events contained in F :

(2.1) Definition. Let

$$C_0 = \{C: C \in F, C = (X_0 \in H) \text{ for some } H \subset J\},$$

$$C_* = \{C: C \in F, C = [(X_n, X_{n+1}) \in V, n \in \mathbb{N}] \text{ for some } V \subset J \times J\},$$

$$C_\infty = \{C: C \in F, C = (\theta_1 \in C)\}.$$

Thus $C_0 = F_0$ is the σ -field of initial events generated by X_0 , C_∞ is the σ -field of invariant events, but the collection C_* of events which constrain all the transitions to be of a certain type is not a σ -field at all.

(2.2) Definition. Let

$$C_+ = \{C: C \in F, C = C_* C_\infty \text{ for some } C_* \in C_*, C_\infty \in C_\infty\},$$

$$C = \{C: C \in F, C = C_0 C_* C_\infty \text{ for some } C_0 \in C_0, C_* \in C_*, C_\infty \in C_\infty\}.$$

Events in C_+ will be called coterminal events, anticipating the connection between these events and coterminal times which is described in the next section. Events in C are intersections of initial events and coterminal events.

Now each of the collections C_β , $\beta = 0, *, \infty$, is readily seen to have the property that if P is Markov then so is P_C whenever $C \in C_\beta$ and $P(C) > 0$, and it follows by repeated conditioning that the class C of all intersections of events from these collections must again have this property. The central result of this section is that no matter what Markov probability P we consider, the events in C are up to P -equivalence the only events for which P_C is Markov:

(2.3) Theorem. Suppose P is Markov and C is an event with $P(C) > 0$. Then P_C is Markov if and only if C is P -equivalent to an event in C .

The theorem is an immediate consequence of Lemma (2.5) and Proposition (2.10) below. Proofs of these results take up the remainder of the section, but we mention first a simple corollary:

(2.4) Corollary. Suppose the Markov probability P makes all states recurrent. Then P_C is Markov if and only if C is P -equivalent to an initial event.

Proof. If P makes all states recurrent then it is found that every coterminial event is P -equivalent to an initial event (see Freedman [1], 1.120).

(2.5) Lemma. If C is a coterminial event, then there are events $F_n \in \mathcal{F}_n$ such that

$$(i) \quad C = F_n(\theta_n \in C), \quad n \in \mathbb{N}.$$

Conversely, if C is an event such that

$$(ii) \quad C = F_1(\theta_1 \in C)$$

for some $F_1 \in \mathcal{F}_1$, then C is a coterminial event. Furthermore, if (ii) holds only P^x a.s. for all $x \in J$, where P^x is Markov(p) starting at x , then there is a coterminial event which is P^x -equivalent to C for all x .

Proof. The first assertion is obvious. For the converse suppose that $C \in \mathcal{F}$ satisfies (ii). Since $F_1 = ((X_0, X_1) \in V)$ for some $V \subset J \times J$ we get

$$C = ((X_0, X_1) \in V)(\theta_1 \in C),$$

whence

$$C = ((X_{k-1}, X_k) \in V, 1 \leq k \leq n)(\theta_n \in C), \quad n \geq 1,$$

by iteration. But intersecting this identity over all $n \geq m$ gives us

$$C = C_V(\theta_n \in C, n \geq m)$$

where $C_V = ((X_{k-1}, X_k) \in V, 1 \leq k < \infty) \in \mathcal{C}_*$, and taking the union of this identity over all m gives $C = C_V C_\infty$ where $C_\infty = \liminf_{n \rightarrow \infty} (\theta_n \in C)$ is invariant, so that C is indeed a coterminial event. For the final assertion the same

sequence of identities is justified P^x a.s. for all x by using the fact that if two events F_1 and F_2 agree P^x a.s. for all x then so too do the events $(\theta_n \in F_1)$ and $(\theta_n \in F_2)$ for each $n \in \mathbb{N}$.

Our efforts now are directed toward establishing the final result (2.10) required for the proof of Theorem (2.3).

(2.6) Notation. Recall that A_n denotes the countable collection of all atoms of F_n . Now for $y \in J$ let A_{ny} denote the subcollection of A_n comprising those atoms contained in the event $(X_n = y)$.

We observe that a probability P on (Ω, F) is Markov if and only if for each $y \in J$ the P_A distribution of the post- n process θ_n remains constant as A varies over all events in A_{ny} with $P(A) > 0$ and n varies over \mathbb{N} . When P is Markov(p) this constant distribution is of course the probability p^y .

(2.7) Definition. Let $n \in \mathbb{N}$. For each event A in A_n , and each event $F \in F$, define a set $F_A \subset \Omega$, the section of F beyond A as follows: For $A = (X_k = x_k, 0 \leq k \leq n) \in A_n$, F_A comprises those sequences $\omega = (\omega_0, \omega_1, \dots)$ such that $\omega_0 = x_n$ and the sequence $(x_0, \dots, x_n, \omega_1, \omega_2, \dots)$ is in F .

Then F_A is an event in F and we shall make repeated use of the identity

$$(2.8) \quad AF = A(\theta_n \in F_A), \quad A \in A_n, \quad F \in F.$$

Notice that if $A \in A_{ny}$ then $F_A \subset (X_0 = y)$, $F \in F$.

(2.9) Lemma. Suppose P is Markov(p), $C \in F$ with $P(C) > 0$. Then for $A \in A_{ny}$ with $P(AC) > 0$, the P_{AC} distribution of θ_n is $P_{CA}^y = P^y(\cdot | C_A)$.

Proof. For $B \in F$, $A \in A_{ny}$ with $P(AC) > 0$, we have

$$\begin{aligned} P_{AC}(\theta_n \in B) &= \frac{P[AC(\theta_n \in B)]}{P(AC)} = \frac{P[A(\theta_n \in C_A B)]}{P[A(\theta_n \in C_A)]} \\ &= \frac{P_A(\theta_n \in C_A B)}{P_A(\theta_n \in C_A)} = \frac{P^Y(C_A B)}{P^Y(C_A)} = P^Y(B|C_A) . \end{aligned}$$

(2.10) Proposition. Suppose P is Markov(p), $C \in F$ with $P(C) > 0$. Then P_C is Markov if and only if there exists an event $D \in F$ such that

$$(i) \quad C = C_0 D \quad P \text{ a.s.}$$

for some initial event $C_0 \in C_0$, and

(ii) for each $n \in N$ there is an event $F_n \in F_n$ with

$$D = F_n(\theta_n \in D) \quad P^X \text{ a.s.}$$

for all $x \in J$.

Proof. Fix P and $C \in F$ with $P(C) > 0$, and define $I \subset J$ to be the essential range of (X_n) under P_C :

$$I = \{y \in J: P_C(X_n=y) > 0 \text{ for some } n\} .$$

For $y \in J$ define A_{ny}^+ to be the collection of all atoms A in A_{ny} with $P_C(A) > 0$, and set $A_y^+ = \bigcup_n A_{ny}^+$. Thus A_y^+ is non-empty if and only if $y \in I$. Now P_C is Markov if and only if for each $y \in I$ the P_{AC} distribution of θ_n is constant as A varies over A_{ny}^+ and n varies over N . Thus by (2.9) P_C is Markov if and only if for each $y \in I$ the probabilities $P^Y(\cdot|C_A)$ are identical, $A \in A_y^+$, i.e. if and only if the events C_A are P^Y a.s. identical, $A \in A_y^+$. But if there is a D satisfying (i) and (ii), then we clearly have

for all $A \in A_y^+$

$$C_A = (X_0=y)D \quad P^y \text{ a.s.},$$

hence P_C is Markov. Conversely, if P_C is Markov, say Markov(q), then for each $y \in I$ we can select a representative event C_A with $A \in A_y^+$, call it D_y , set $D = \bigcup_{y \in I} D_y$, and then have for each $y \in I$ the identity

$$(2.11) \quad P^y(\cdot|D) = P^y(\cdot|D_y) = P^y(\cdot|C_A) = Q^y(\cdot), \quad A \in A_y^+,$$

where Q^y is Markov(q) starting at y . Obviously this D satisfies (i) with $C_0 = (X_0 \in H)$ for $H = \{y: P_C(X_0=y) > 0\}$, and this D also satisfies (ii), as can be seen by the following argument. For any $y \in I$, P_D^y is Markov(q) so that $D = D_A$ P^z a.s. for any $A \in A_{nz}$ with $P^y(AD) > 0$. Consequently $AD = A(\theta_n \in D)$ P^y a.s. for any $A \in A_n$ with $P^y(AD) > 0$ and taking the union over all such atoms A we arrive at a representation $D = F_{ny}(\theta_n \in D)$ P^y a.s. with $F_{ny} \in F_n$. Defining $F_n = \bigcup_{y \in I} (X_0=y)F_{ny}$ we now see that $D = F_n(\theta_n \in D)$ P^x a.s. for arbitrary $x \in I$, and since $P^x(D) = P^x(F_n) = 0$ for $x \notin I$, D satisfies (ii), and the proof is complete.

3. REGULAR BIRTH TIMES

The main result of this section is Theorem (3.9) which describes all regular birth times for a Markov probability P in terms of certain fundamental birth times associated with the coterminal events of the previous section, i.e. the coterminal times of Meyer, Smythe and Walsh [3].

Using the notation defined at the end of the introduction, a random time τ is a regular birth time for P if and only if a P conditional distribution of θ_τ given F_τ is equal to Q^x on $(\tau < \infty, X_\tau = x)$, where Q^x is Markov(q)

with starting state x for some transition matrix q . Put another way, τ is a regular birth time for P if and only if under P the post- τ sequence $(X_{\tau+n}, n \in \mathbb{N})$ is Markov(q) with respect to the increasing sequence of σ -fields $(F_{\tau+n}, n \in \mathbb{N})$.

Suppose now that C is a coterminal event as defined in (2.2), i.e.

$$(3.1) \quad C = C_V C_\infty$$

where

$$(3.2) \quad C_V = [(X_{n-1}, X_n) \in V, 1 \leq n < \infty]$$

for some $V \subset J \times J$, and C_∞ is invariant.

(3.3) Definition. The coterminal time associated with C is the random time τ_C defined by

$$\tau_C = \inf\{n \in \mathbb{N}: \theta_n \in C\}.$$

(Here and elsewhere we use the convention $\inf \emptyset = \infty$, $\sup \emptyset = 0$).

Since for coterminal events C

$$(3.4) \quad (\theta_k \in C) \subset (\theta_m \in C), \quad 0 \leq k \leq m < \infty,$$

we have the identity

$$(3.5) \quad (\tau_C \leq n) = (\theta_n \in C), \quad n \in \mathbb{N}.$$

In particular, if $C = C_V$, then τ_C is the time that the last transition in V^C is completed:

$$\tau_{C_V} = \sup\{n \geq 1: (X_{n-1}, X_n) \in V^C\},$$

while if $C = C_\infty$ is invariant, then

$$\tau_{C_\infty} = 0 \text{ on } C_\infty, \quad \infty \text{ on } C_\infty^c,$$

and in general for $C = C_V C_\infty$ we have that τ_C is simply the maximum of τ_{C_V} and τ_{C_∞} .

It is easy to check that τ_C is indeed a coterminal time as defined by Meyer, Smythe and Walsh in [3], i.e. that the random time $\tau = \tau_C$ has the properties

$$(3.6) \quad \begin{aligned} (i) \quad & \tau \circ \theta_n = (\tau - n)^+, \quad n \in \mathbb{N}, \\ (ii) \quad & \tau \circ K_n = \tau \text{ on } (\tau < n), \quad n \in \mathbb{N}. \end{aligned}$$

Conversely, if τ is a coterminal time, then $C = (\tau=0)$ is a coterminal event and $\tau = \tau_C$. To see this observe that (3.6)(i) implies

$$(3.7) \quad (\tau \leq n) = (\tau \circ \theta_n = 0) = (\theta_n \in C), \quad n \in \mathbb{N},$$

so that by (3.5) it suffices to show that C is a coterminal event. But by (3.6)(ii) for $n = 2$ and (3.7) for $n = 1$ we have

$$C = (\tau = 0, \tau < 2) = (\tau \circ K_2 = 0, \tau < 2) = (\tau \circ K_2 = 0, \theta_1 \in C),$$

and since $(\tau \circ K_2 = 0) \in F_1$, we conclude from (2.5)(ii) that C is a coterminal event.

We shall see shortly that each coterminal time $\tau = \tau_C$ is a birth time for each Markov probability P : indeed a P conditional distribution for θ_τ given F_τ equals P_C^X on $(X_\tau = x)$, where $P_C^X = Q^X$ is Markov(q) for some q by (2.3). This is just the analogue in the present context of Theorem 5.1 of [3]. For a more detailed description of the path decomposition at τ_C giving

the transition probabilities q of the post- τ_C process, see Section 6.

(3.8) Definition. Let B denote the class of all random times τ of the form

$$\tau = \tau_C + \rho$$

where τ_C is the coterminal time associated with a coterminal event C , and ρ is an optional time for the increasing sequence of σ -fields $(F_{\tau_C+n}, n \in \mathbb{N})$, i.e. $(\rho=n) \in F_{\tau_C+n}, n \in \mathbb{N}$.

Once we know that each τ_C is a regular birth time for P , it follows at once from the strong Markov property of the sequence $(X_{\tau_C+n}, n \in \mathbb{N})$ adapted to $(F_{\tau_C+n}, n \in \mathbb{N})$ that each $\tau \in B$ is again a regular birth time for P . Our principal result is that the random times in B form a complete, canonical collection of birth times in that no matter what Markov probability P we start off with, every regular birth time for P is P -equivalent to a random time in B .

(3.9) Theorem. A random time τ is a regular birth time for a Markov probability P if and only if τ is P -equivalent to a random time in B .

The proof of this theorem takes up the rest of the section, but we mention first the following corollary:

(3.10) Corollary. If P makes all states recurrent then every regular birth time for P is P -equivalent to a stopping time.

Proof. Just as for (2.4).

(3.11) Definition. A random time τ is a conditional independence time for a Markov probability P if under P the pre- τ and post- τ processes are conditionally independent given X_τ , i.e. if there is a conditional distribution of θ_τ given F_τ within $(\tau < \infty)$ which is a function of X_τ alone.

(3.12) Lemma. A random time τ is a conditional independence time for P if and only if there are events $F_n \in \mathcal{F}_n$, $n \in \mathbb{N}$, $G \in \mathcal{F}$ with

$$(3.13) \quad (\tau = n) = F_n(\theta_n \in G) \quad P \text{ a.s.}, \quad n \in \mathbb{N},$$

and there is then a conditional distribution of θ_τ given F_τ which equals P_G^X on $(\tau < \infty, X_\tau = x)$.

Remark. The proof can easily be sharpened to show that τ is a conditional independence time for P if and only if τ is P a.s. equal to a τ^* with $(\tau^* = n) = F_n(\theta_n \in G)$ exactly for some $F_n \in \mathcal{F}_n$, $G \in \mathcal{F}$. Every such time is thus a.s. equal to a splitting time, defined in Jacobsen [2] as a random time τ for which $(\tau = n) = F_n(\theta_n \in G_n)$ for some $F_n \in \mathcal{F}_n$, $G_n \in \mathcal{F}$. The present argument will also show that splitting times are characterized by conditional independence of the pre- τ and post- τ processes given both X_τ and τ .

Proof. Working on atoms as in the proof of (2.9), let A_{nx}^+ be the collection of all atoms A of F_n contained in $(X_n = x)$ with $P(AG_n) > 0$, where $G_n = (\tau = n)$, so that $(X_n = x)G_n$ is P a.s. equal to the union of the sets AG_n over all A in A_{nx}^+ . Defining G_{nA} as in (2.7) to be the section of G_n beyond A , we have from (2.9) that the P conditional distribution of θ_n given AG_n is $P^X(\cdot | G_{nA})$, $A \in A_{nx}^+$. But τ is a conditional independence time if and only if this conditional distribution is a function of x alone for all $A \in A_{nx}^+$, $n \in \mathbb{N}$, i.e. if and only if for some G_x in \mathcal{F} with

$G_x \subset (X_0 = x)$ we have for each $x \in J$

$$(3.14) \quad G_{nA} = G_x \quad P^x \text{ a.s.}, \quad A \in A_{nx}^+.$$

But if (3.13) holds we have (3.14) with $G_x = G(X_0 = x)$, while if (3.14) holds we get (3.13) with $G = \bigcup_{x \in J} G_x$, and F_n the union of all $A \in A_n$ with $P(AG_n) > 0$.

Proof of Theorem (3.9). We have that τ is a regular birth time for P if and only if τ is a conditional independence time for P and the P distribution of θ_τ is Markov. Thus the 'if' part follows at once from (3.12) and (2.3). For the converse, suppose τ is a regular birth time for P . Then (3.12) shows that there are events $F_n \in F_n$, $G \in F$ with

$$(3.15) \quad (\tau = n) = F_n(\theta_n \in G) \quad P \text{ a.s.},$$

and also that the conditional distribution of θ_τ given F_τ is P_G^x on $(\tau < \infty, X_\tau = x)$. Since the regular birth time property tells us that this probability is Markov(q) for some q not depending on x , we deduce from Theorem (2.3) that for every x with $P(\tau < \infty, X_\tau = x) > 0$, G is P^x a.s. equal to a coterminal event C which because of (2.10)(ii) and (2.5)(ii) may be chosen so as not to depend on x . It now follows that we can replace the set $G_x = G(X_0 = x)$ which appears in (3.14) by the set $C_x = C(X_0 = x)$ to deduce that in fact

$$(3.16) \quad (\tau = n) = F_n(\theta_n \in C) \quad P \text{ a.s.}, \quad n \in \mathbb{N}.$$

But now define

$$(3.17) \quad C_n = F_n(\theta_n \in C)$$

and define a random time τ' by setting

$$\begin{aligned}\tau' &= n \text{ on } C_n \setminus \bigcup_{k=0}^{n-1} C_k, \quad n \in \mathbb{N} \\ &= \infty \text{ on } \left(\bigcup_{k=0}^{\infty} C_k \right)^c.\end{aligned}$$

(Note that we could not define τ' to be n on C_n because the sets C_n might not be disjoint). We clearly have $\tau' = \tau$ P a.s., and we now conclude by showing that $\tau' \geq \tau_C$ and that $\rho = \tau' - \tau_C$ is an optional time for $(F_{\tau_C+m}, m \in \mathbb{N})$. But from (3.17), (3.4) and (3.5) we have

$$(\tau' \leq n) = \bigcup_{k=0}^n C_k \subset \bigcup_{k=0}^n (\theta_k \in C) = (\tau_C \leq n), \quad n \in \mathbb{N},$$

which implies $\tau' \geq \tau_C$. Furthermore, setting $\rho = \tau' - \tau_C$ we have for $m \in \mathbb{N}$ the identity

$$(3.18) \quad (\rho \leq m) = \bigcup_{n=m}^{\infty} (\tau_C = n-m, \tau' \leq n).$$

But

$$(3.19) \quad (\tau_C = n-m, \tau' \leq n) = (\tau_C = n-m) \left[\bigcup_{k=0}^n C_k \right] = \bigcup_{k=0}^n B_k$$

where $B_k = (\tau_C = n-m)C_k$. Now if $k < n-m$ we have by (3.17) and (3.5) that $C_k \subset (\theta_k \in C) = (\tau_C \leq k)$ so that $B_k = \emptyset$, while if $n-m \leq k \leq n$ we have from (3.5) and (3.4) that $(\tau_C = n-m) \subset (\theta_{n-m} \in C) \subset (\theta_k \in C)$ so that using (3.17) we get $B_k = (\tau_C + m = n)F_k$ where $F_k \in F_k \subset F_n$. In either case we see that $B_k \in F_{\tau_C+m}$, and thus working back through (3.19) to (3.18) we get also that $(\rho \leq m) \in F_{\tau_C+m}$, which is to say that ρ is an optional time of $(F_{\tau_C+m}, m \in \mathbb{N})$. The proof is complete.

We conclude this section with some remarks about plain birth times. For an interesting example fix $m \geq 1$ and define

$$\tau = \inf\{n \geq 1: (X_n, \dots, X_{n+m}) = (X_0, \dots, X_m)\}.$$

If P makes all states recurrent then τ is P a.s. finite and since $\tau+m$ is a stopping time and $(X_\tau, \dots, X_{\tau+m}) = (X_0, \dots, X_m)$ on $(\tau < \infty)$ it is easy to see that the P distribution of θ_τ is P . Thus τ is a birth time for P , but certainly not a regular birth time, since knowledge of the pre- τ process completely determines the first m -moves of the post- τ process. Examples show that τ can fail to be a birth time if there are transient states. But it is of greater interest to use the idea behind this example to construct two Markov probabilities P and R with the same null sets and a τ which is a birth time for P but not for R , since this shows that there exists no canonical class B^* of plain birth times with the property that τ is a birth time for a Markov probability P if and only if τ is P -equivalent to a time in B^* .

(3.20) Example. Ignore the coffin state Δ and let $J = \{1, 2, 3\}$. Define transition matrices p and r on J by

$$p = \begin{bmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix}, \quad r = \begin{bmatrix} 0 & \frac{1}{4} & \frac{3}{4} \\ 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix}.$$

Let P and R be Markov(p) and Markov(r) respectively, both starting at 1. Then P and R have the same null sets but the random time τ defined by

$$\tau = \begin{cases} \inf\{n: X_n = 2, X_{n+1} = 2\} & \text{if } X_1 = 2 \\ \inf\{n: X_n = 2, X_{n+1} = 3\} & \text{if } X_1 = 3 \end{cases}$$

is a birth time for P only.

4. DEATH TIMES

Given a Markov(p) probability P on (Ω, F) we now investigate death times for P , i.e. random times τ such that under P the distribution of the pre- τ fragment $(X_0, \dots, X_{\tau-1})$ is Markov with stationary sub-stochastic transition probabilities, or, what is the same, that the killed process $K_\tau = (X_0, \dots, X_{\tau-1}, \Delta, \Delta, \dots)$ is Markov with stationary transition probabilities.

Let τ be a death time for P . If we let J_+ denote the essential range of the pre- τ path,

$$J_+ = \{x \in J \setminus \{\Delta\} : P(\tau > n, X_n = x) > 0 \text{ for some } n \in \mathbb{N}\},$$

then the transition probabilities $q(x, y)$ of the killed chain are well defined for all $x, y \in J_+ \cup \{\Delta\}$ and induce a family of probabilities $\{Q^x, x \in J_+ \cup \{\Delta\}\}$ on Ω which concentrate on paths ω which remain forever within this restricted state space $J_+ \cup \{\Delta\}$. Clearly if $P(\tau > 0, X_0 = x) > 0$ then Q^x is identical to the P^x distribution of K_τ given $(\tau > 0)$. This conclusion may be either false or meaningless if $P(\tau > 0, X_0 = x) = 0$, but we shall see that matters can be rectified by redefining $\tau(\omega)$ properly for paths ω starting at such points x .

We consider now the whole family of probabilities $\{P^x\}_{x \in J}$ where P^x on (Ω, F) is Markov(p) with starting state x . For each random time τ we define $J_\tau \subset J$ by

$$J_\tau = \{x \in J \setminus \{\Delta\} : P^x(\tau > 0) > 0\},$$

and say that τ is a death time for the family $\{P^x\}_{x \in J}$ if there is a

transition matrix q on $J_\tau \cup \{\Delta\}$ such that for each $x \in J_\tau$ the P^x distribution of K_τ given $(\tau > 0)$ is Q^x , where Q^x on (Ω, F) is Markov(q) with starting state x . In particular Q^x concentrates on the set of sequences in $J_\tau \cup \{\Delta\}$.

Obviously every death time for the family $\{P^x\}_{x \in J}$ is a death time for each Markov(p) probability P . Conversely we have

(4.1) Proposition. Let τ be a death time for a Markov(p) probability P . Then there is a death time τ^* for the family $\{P^x\}_{x \in J}$ such that $\tau^*(\omega) = \tau(\omega)$ for all paths ω starting at points $x \in J$ with $P(\tau > 0, X_0 = x) > 0$.

Remark. If $P = P^y$ for a fixed state y then τ^* and τ agree P a.s.

Proof. As before let J_+ be the essential range of the pre- τ path under P , and suppose K_τ is Markov(q) under P . For those paths ω starting at an $x \in J \setminus J_+$ we set $\tau^*(\omega) = 0$, while for those starting at an $x \in J_+$ with $P(\tau > 0, X_0 = x) > 0$ we put $\tau^*(\omega) = \tau(\omega)$. Finally, if $x \in J_+$ and $P(\tau > 0, X_0 = x) = 0$ we find an $m \geq 1$ and $x_0, \dots, x_m \in J_+$ with $x_m = x$ such that $P[A(\tau > m)] > 0$ where $A = (X_0 = x_0, \dots, X_m = x_m)$, and then for $\omega = (\omega_0, \omega_1, \dots)$ with $\omega_0 = x$ define

$$\tau^*(\omega) = (\tau(x_0, \dots, x_{m-1}, \omega_0, \omega_1, \dots) - m)^+.$$

Obviously $J_{\tau^*} = J_+$ and it remains only to check that the P^x distribution of K_{τ^*} given $(\tau^* > 0)$ is Q^x for $x \in J_+$ with $P(\tau > 0, X_0 = x) = 0$, where Q^x is Markov(q) starting at x . For this it suffices to show that for $n \in \mathbb{N}$, $B = (X_0 = y_0, \dots, X_n = y_n)$, where $y_0 = x$, $y_1, \dots, y_n \in J_+$, we have

$$P^x[B(\tau^* > n) | \tau^* > 0] = Q^x(B),$$

where $Q^X(B) = q(y_0, y_1) \cdots q(y_{n-1}, y_n)$ by definition.

But let m and n be as above. Then by first using the Markov property of P and then the fact that $\tau^* \circ \theta_m = (\tau - m)^+$ on A we have

$$\begin{aligned} P^X[B(\tau^* > n)] &= P[A(\theta_m \in B, \tau^* \circ \theta_m > n)]/P(A) \\ &= P[A(\theta_m \in B, \tau > m+n)]/P(A) . \end{aligned}$$

Similarly $P^X(\tau^* > 0) = P[A(\tau > n)]/P(A)$, whence

$$P^X[B(\tau^* > n) | \tau^* > 0] = P[A(\theta_m \in B, \tau > m+n)]/P[A(\tau > n)]$$

which equals $Q^X(B)$ because the P distribution of K_τ is Markov(q).

This result reduces the problem of describing the death times for a given Markov probability P to that of characterizing the death times for a family $\{P^x\}_{x \in J}$. Such a characterization is provided by the following proposition. Write $P(\cdot | F_n)$ for any of the identical conditional probabilities $P^X(\cdot | F_n)$ and given a random time τ introduce

$$f(x) = P^X(\tau > 0) , \quad Z_n = P(\tau > n | F_n) , \quad n \in \mathbb{N} .$$

Note that $Z_0 = f(X_0)$ and recall that $J_\tau = \{x \in J \setminus \{\Delta\} : f(x) > 0\}$.

(4.2) Proposition. A random time τ is a death time for the family $\{P^x\}_{x \in J}$ if and only if for each $x \in J \setminus \{\Delta\}$, $m, n \in \mathbb{N}$ the identity

$$Z_{m+n} = \begin{cases} Z_m Z_n \circ \theta_m / f(X_m) & \text{on } (X_m \in J_\tau) \\ 0 & \text{on } (X_m \notin J_\tau) \end{cases}$$

holds P^X a.s.

Proof. With Q^x the P^x distribution of K_τ given $(\tau > 0)$, the condition that τ be a death time for $\{P^x\}$ is equivalent to requiring that for all $m, n, k \in \mathbb{N}$, $x_0, \dots, x_{m+n}, x \in J_\tau$,

$$(4.3) \quad Q^{x_0}(x_0=x_0, \dots, x_{m+n}=x_{m+n}) = Q^{x_0}(x_0=x_0, \dots, x_m=x_m) Q^{x_m}(x_0=x_m, \dots, x_n=x_{m+n}) ,$$

$$(4.4) \quad Q^{x_0}(x_k \notin J_\tau \cup \{\Delta\}) = 0 .$$

Introducing the atoms $A = (x_0=x_0, \dots, x_m=x_m)$, $B = (x_0=x_m, \dots, x_n=x_{m+n})$, it is seen that (4.3) and (4.4) are equivalent to

$$(4.5) \quad P^{x_0}[A(\theta_m \in B, \tau > m+n)] = P^{x_0}[A(\tau > m)] P^{x_m}[B(\tau > n)] / f(x_m)$$

$$(4.6) \quad P^{x_0}(x_k \notin J_\tau, \tau > k) = 0 .$$

But the left side of (4.5) equals

$$P^{x_0}[Z_{m+n}; A(\theta_m \in B)]$$

while the right side becomes

$$P^{x_0}(Z_m; A) P^{x_m}(Z_n; B) / f(x_m) = P^{x_0}[Z_m Z_n \circ \theta_m / f(x_m); A(\theta_n \in B)] .$$

Since A and B are arbitrary atoms for paths within J_τ and since (4.6) is equivalent to demanding that Z_k, Z_{k+1}, \dots vanish P^x a.s. on $(x_k \notin J_\tau)$ the result follows.

We finish this section with an example of two Markov probabilities P and R , with the same null sets, and a random time τ which is a death time for P but not for R , thus proving that there exists no collection of random times \mathcal{D}^* such that τ is a death time for a Markov probability P if and only if

τ is P a.s. equal to a random time in \mathcal{D}^* (cf. Example (3.20)).

(4.7) Example. Ignoring the coffin state Δ let $J = \{1,2,3,4\}$. Define transition matrices p and r on J by

$$p = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad r = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{4} & \frac{3}{4} \\ 0 & \frac{3}{4} & 0 & \frac{1}{4} \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Let P and R be $\text{Markov}(p)$ and $\text{Markov}(r)$ respectively, both starting at 1. Clearly P and R have the same null sets. Define ρ as the time of entry into $\{2,3\}$, σ as the time to absorption in state 4, put $\delta = \sigma - \rho$ and finally let τ be the minimum of ρ and δ .

Under both P and R the random times ρ and δ are independent; under P they both have a geometric distribution, whence so too does τ , but under R the random time ρ is geometric while δ is not, and thus τ is not geometric either. Since the pre- τ process never leaves state 1 it is therefore Markov under P but not under R .

5. REGULAR DEATH TIMES

In this section we consider death times τ for a Markov probability P on (Ω, \mathcal{F}) with the additional property, discussed in the introduction, that conditional on the final position $X_{\tau-1}$ of the pre- τ process on $(0 < \tau < \infty)$ the pre- τ and post- τ processes are independent. Such a death time for P will be referred to as regular.

(5.1) Definition. Let \mathcal{D} denote the class of all random times τ of the form

$$\tau = \tau_{V,F} = \sup\{n: 1 \leq n \leq \tau_V, \theta_{n-1} \in F\}$$

for some $V \subset J \times J$, $F \in \mathcal{F}$, where τ_V is the terminal time associated with V :

$$\tau_V = \inf\{n: n \geq 1, (X_{n-1}, X_n) \in V\}.$$

Our main result is that the random times in \mathcal{D} form a complete canonical collection of regular death times for each Markov probability P^x on (Ω, \mathcal{F}) with a fixed starting state x . For random starts \mathcal{D} must be enlarged to include all times of the form $\tau 1(X_0 \in H)$ for $H \subset J$, but we shall ignore this trivial complication.

(5.2) Theorem. A random time τ is a regular death time for a Markov probability P^x with starting state x in $J \setminus \{\Delta\}$ if and only if τ is P^x -equivalent to a random time in \mathcal{D} .

The proof of this theorem will be taken up later in the section. We mention first some facts about the collection \mathcal{D} which derive from the following result.

(5.3) Proposition. The random times in \mathcal{D} are characterized by the following two structural properties: for every $n \in \mathbb{N}$

$$(i) \quad (\tau > n) = F_n(\tau \circ \theta_n > 0) \text{ for some } F_n \in \mathcal{F}_n,$$

$$(ii) \quad \tau \circ \theta_n = \tau - n \text{ on } (\tau > n).$$

Indeed, if τ satisfies (i) and (ii), then $\tau = \tau_{V,F}$ with $F = (\tau > 0)$ and V such that $F_1 = ((X_0, X_1) \in V^c)$, and then

$$F_n = (\tau_V > n) = ((X_k, X_{k+1}) \in V^c, 0 \leq k \leq n-1) .$$

Proof. See Lemma (5.4) below. The present criterion is just a restatement of the criterion (a) of that lemma.

We note that for an optional time τ we have $(\tau > n) \in F_n$ so that the two conditions (i) and (ii) above collapse to (ii) alone. Since an optional time satisfying (ii) is by definition a terminal time, we find that an optional time is in \mathcal{D} if and only if it is a terminal time. Furthermore it is plain from the original definition of \mathcal{D} , that τ is a terminal time if and only if τ is the first time that the path either enters H or completes a jump in V for some $H \subset J, V \subset J \times J$ (cf. Walsh and Weil [7]).

For a cooptional time we have by definition that $\tau \circ \theta_n = (\tau - n)^+$, and since this obviously implies both (i) and (ii) we see that \mathcal{D} includes all cooptional times. We find also that every cooptional time τ can be represented as $\tau = \sup\{n: n \geq 1, \theta_{n-1} \in F\}$ where $F = (\tau > 0)$, and since for cooptional τ we must have $(\tau > 0) = (\theta \in G)$ for some $G \in \mathcal{F}$ we deduce that τ is a cooptional time if and only if for some $G \in \mathcal{F}$

$$\tau = \sup\{n: n \in \mathbb{N}, \theta_n \in G\} .$$

Finally we mention two important closure properties of \mathcal{D} which are readily checked using Proposition (5.3): if σ and τ are two random times in \mathcal{D} , then so are (i) their minimum $\sigma \wedge \tau$, and (ii) the random time $\sigma \circ K_\tau$ obtained by applying σ after first killing at τ , provided $\sigma \leq \inf\{n \in \mathbb{N}: X_n = \Delta\}$.

(5.4) Lemma. Each of the following conditions (a) and (b) is necessary and sufficient for a random time τ to belong to \mathcal{D} :

(a) For each $m \in \mathbb{N}$ there is an $F_m \in \mathcal{F}_m$ such that

$$(\tau > m+n) = F_m(\tau \circ \theta_m > n) , \quad n \in \mathbb{N} .$$

(b) There is an $F_1 \in \mathcal{F}_1$ such that for $\bar{F}_n = F_1(\theta_1 \in F_1) \cdots (\theta_{n-1} \in F_1)$ we have

$$\begin{aligned} (i) \quad & (\tau = n+1) = \bar{F}_n(\tau \circ \theta_n = 1) , \\ (ii) \quad & (\tau = \infty) = \bar{F}_n(\tau \circ \theta_n = \infty) , \end{aligned} \quad n \geq 1 .$$

Proof. If $\tau = \tau_{V,F} \in \mathcal{D}$ then (a) is immediate with $F_m = (\tau_V > m)$. Conversely, if τ satisfies (a) we will show that $\tau = \tau_{V,F}$ with $F = (\tau > 0)$ and the V defined by $F_1 = ((X_0, X_1) \in V^C)$. As a preliminary we claim that for $m \geq 1$ the F_m appearing in (a) can be replaced by the \bar{F}_m determined from F_1 as in (b). This is trivial for $m = 1$, while if true for m and all n we find that

$$\begin{aligned} (\tau > (m+1) + n) &= (\tau > m + (n+1)) = \bar{F}_m(\tau \circ \theta_m > n+1) \\ &= \bar{F}_m[\theta_m \in F_1(\tau \circ \theta > n)] = \bar{F}_{m+1}(\tau \circ \theta_{m+1} > n) , \end{aligned}$$

establishing the fact for $m+1$, and the assertion is verified. To show now that $\tau = \tau_{V,F}$ for the given V and F it suffices to check that $(\tau_{V,F} > m) = (\tau > m)$ for $m \in \mathbb{N}$, i.e. that

$$\bigcup_{n \geq m} (\tau_V > n, \theta_n \in F) = \bar{F}_m(\tau \circ \theta_m > 0) .$$

But since $\bar{F}_m = (\tau_V > m)$ and $(\tau \circ \theta_m > 0) = (\theta_m \in F)$, the right hand event is identical to $(\tau_V > m, \theta_m \in F)$, and we must therefore show that for each $n \geq m$ this event contains the event

$$(5.5) \quad (\tau_V > n, \theta_n \in F) = (\tau_V > m, \tau_V \circ \theta_m > n-m, \theta_n \in F) .$$

But because $(\tau > 0) \supset (\tau > 1) = F_1(\tau \circ \theta > 0)$, we find by induction that $(\tau > 0) \supset \bar{F}_k(\tau \circ \theta_k > 0) = (\tau_V > k, \theta_k \in F)$ and using this with $k = n-m$ in (5.5) yields the desired conclusion.

For the second characterization (b) it is obvious that if $\tau = \tau_{V,F}$ is in \mathcal{D} , then both (i) and (ii) hold with $F_1 = ((X_0, X_1) \in V^C)$, while if these identities obtain it is easy to check that (a) holds with $F_m = \bar{F}_m$.

The same proofs show that if $\Omega^* \subset \Omega$ is invariant under θ , i.e. if $\Omega^* \subset (\theta \in \Omega^*)$, then knowing that either (5.4)(a) or (5.4)(b) holds just on Ω^* , e.g. that

$$(\tau = n+1)\Omega^* = \bar{F}_n(\tau \circ \theta_n = 1)\Omega^* , \quad (\tau = \infty)\Omega^* = \bar{F}_n(\tau \circ \theta_n = \infty)\Omega^* ,$$

we can deduce that $\tau = \tau_{V,F}$ on Ω^* for some $V \subset J \times J$, $F \in \mathcal{F}$.

We also note that an obvious modification of Lemma (3.12) tells us that for a random time τ and a Markov probability P the pre- τ and post- τ processes are conditionally independent given $X_{\tau-1}$ and the event $(0 < \tau < \infty)$ if and only if for every $n \in \mathbb{N}$ there exists $F_n \in \mathcal{F}_n$, $G \in \mathcal{F}$ such that

$$(5.6) \quad (\tau = n+1) = F_n(\theta_n \in G) \quad P \text{ a.s. .}$$

Proof of Theorem (5.2). Fix $x \in J \setminus \{\Delta\}$. Assuming first that $\tau = \tau_{V,F}$ P^x a.s., it follows from (5.4)(a) that for $Z_n = P(\tau_{V,F} > n | \mathcal{F}_n)$ we have $Z_{m+n} = 1_{F_m} Z_n \circ \theta_m$ a.s. where here and throughout this proof 'a.s.' on its own means ' P^y a.s. for all $y \in J \setminus \{\Delta\}$ '. But Proposition (4.2) is then easily applied to show that $\tau_{V,F}$ is a death time for P^x , and hence so too is τ . Since (5.4)(b) applied to $\tau_{V,F}$ shows that (5.6) holds with $G = (\tau_{V,F} = 1)$ it follows that τ is a regular death time for P^x .

Now assume conversely that τ is a regular death time for P^X . In particular by (5.6), $(\tau = n+1) = F_n(\theta_n \in G)$ P^X a.s. for some $F_n \in \mathcal{F}_n$, $G \in \mathcal{F}$. Replacing τ by the τ^* constructed in Proposition (4.1) we get a random time which is P^X a.s. equal to τ and which is a death time for the family $\{P^y\}_{y \in J}$. Furthermore, it is not difficult to verify that $(\tau^* = n+1) = F_n^*(\theta_n \in G)$ a.s., $n \in \mathbb{N}$, where F_n^* is the union of the events $F_{n,y}$ over y in J_{τ^*} , with $F_{n,y}$ the section of F_{n+n} beyond the atom $A_y = (X_0 = x_0, \dots, X_n = y)$ used to define τ^* on paths starting at $y \in J_{\tau^*}$ (see (2.7) and the proof of (4.1)). Thus we may as well drop the stars and take it from the start that τ is such that there are events $F_n \in \mathcal{F}_n$, $G \in \mathcal{F}$ with

$$(5.7) \quad (\tau = n+1) = F_n(\theta_n \in G) \text{ a.s.}$$

We now aim to show that such a τ satisfies (i) and (ii) of (5.4)(b) a.s. Both in this and later arguments there will be much tacit use of the fact that if two events F_1 and F_2 are a.s. equal then so too are $(\theta_n \in F_1)$ and $(\theta_n \in F_2)$ for each $n \in \mathbb{N}$. Writing

$$M = \{y \in J \setminus \{\Delta\} : P^y(\tau = 1) > 0\}$$

we first claim that

$$(5.8) \quad (\tau = n+1, X_n \in M) = F_n(X_n \in M, \tau \circ \theta_n = 1) \text{ a.s.}$$

To see this observe that by (5.7), $(X_0 \in M) = F_0(P^{X(0)}G > 0)$ so that also $(X_n \in M) = (\theta_n \in F_0, P^{X(n)}G > 0)$. But then, again using (5.7)

$$(\tau = n+1, X_n \in M) = F_n(\theta_n \in F_0G, P^{X(n)}G > 0) \text{ a.s.}$$

and (5.8) follows. Introducing now

$$f(y) = P^y(\tau > 0), \quad Z_n = P(\tau > n | F_n), \quad W_n = P(\tau = n+1 | F_n),$$

it follows from Proposition (4.2) that

$$Z_n = W_n + Z_n[f(X_n)]^{-1} 1(X_n \in J_\tau) P^{X(n)}(\tau > 1) \quad \text{a.s.}$$

so that because $(Z_n > 0) \subset (X_n \in J_\tau)$,

$$(5.9) \quad (W_n > 0) = (Z_n > 0, X_n \in M) \quad \text{a.s.}$$

Because $(\tau = n+1) \subset (W_n > 0)$ we find in particular that $(\tau = n+1) \subset (X_n \in M)$ so that (5.8) may be written

$$(\tau = n+1) = F_n(X_n \in M, \tau \circ \theta_n = 1) \quad \text{a.s.}$$

whence $(W_n > 0) = F_n(X_n \in M)$, and comparing with (5.9) it develops that

$$(\tau = n+1) = (W_n > 0, \tau \circ \theta_n = 1) = (Z_n > 0, X_n \in M, \tau \circ \theta_n = 1) \quad \text{a.s.}$$

Since $(\tau \circ \theta_n = 1) \subset (X_n \in M)$ and τ is a death time for $\{P^x\}_{x \in J}$, Proposition (4.2) implies that (i) of (5.4)(b) holds a.s. with $F_1 = (Z_1 > 0)$. To establish (ii) of (5.4)(b) observe that by the martingale convergence theorem

$1(\tau = \infty) = \lim_{m \uparrow \infty} Z_m$ a.s. But because of Proposition (4.2) that limit equals

$$Z_n[f(X_n)]^{-1} 1(X_n \in J_\tau) \lim_m Z_m \circ \theta_m = Z_n[f(X_n)]^{-1} 1(X_n \in J_\tau, \tau \circ \theta_n = \infty) \quad \text{a.s.}$$

for every $n \in \mathbb{N}$, and the desired conclusion that $(\tau = \infty) = (Z_n > 0, \tau \circ \theta_n = \infty)$ a.s. is now immediate.

Thus, for each $y \in J \setminus \{\Delta\}$, ignoring a P^y -null set L_y , we have that τ satisfies the identities of (5.4)(b) for all n with $F_1 = (Z_1 > 0)$. But the set $L = \bigcup_{y \in J \setminus \{\Delta\}} (X_0 = y) L_y$ is a null set for all P^y , $y \in J \setminus \{\Delta\}$ simultaneously

and consequently, defining $\Omega^* = (\bigcup_{n \in \mathbb{N}} (\theta_n \in L))^c$ we see that for $n \in \mathbb{N}$ the identities

$$\begin{aligned} (\tau = n+1)\Omega^* &= (Z_1 > 0, \dots, Z_1 \circ \theta_{n-1} > 0, \tau \circ \theta_n = 1)\Omega^*, \\ (\tau = \infty)\Omega^* &= (Z_1 > 0, \dots, Z_1 \circ \theta_{n-1} > 0, \tau \circ \theta_n = \infty)\Omega^*, \end{aligned}$$

hold exactly, and since Ω^* is invariant under θ , the remark following Lemma (5.4) shows that on Ω^* , τ is of the form $\tau_{V,F}$ for some $V \subset J \times J$, $F \in \mathcal{F}$. Since also $P^y(\Omega^*) = 1$ for every y , in particular for $y = x$, we conclude that τ is P^x a.s. equal to some $\tau_{V,F}$, and the proof is complete.

As a final comment on death times, it may be observed that there exist random times τ which are death times for all Markov probabilities simultaneously without being regular. A simple example is obtained by taking an integer $a \geq 2$ and defining

$$(\tau > 0) = \Omega, \quad (\tau > n) = (X_0 = \dots = X_{an}), \quad n \geq 1.$$

6. BIRTH AND DEATH TIMES

We now consider random times which are both regular death times and regular birth times for each Markov(p) probability P . For such a random time τ it is seen that a path decomposition specifying the joint distribution of the pre- τ and post- τ processes can be given in terms of just four quantities determined by τ and p , namely the function $f: J \rightarrow [0,1]$ and the three transition matrices q, r and s on J such that under the probability P^x on (Ω, \mathcal{F}) which makes (X_n) Markov(p) starting at x we have

- (6.1) (i) $P^x(\tau > 0) = f(x)$, $x \in J$.
- (ii) Conditional on $\tau > 0$ the P^x distribution of the pre- τ process is Markov(q) starting at x ,
- (iii) Conditional on $0 < \tau < \infty$ and a pre- τ path with $X_{\tau-1} = y$ the P^x distribution of X_τ is $r(y, \cdot)$,
- (iv) Conditional on $\tau < \infty$, the pre- τ path and $X_\tau = z$, the P^x distribution of the post- τ process is Markov(s) starting at z .

With f , q , r and s specified by (6.1) the path decomposition involved in (6.1) can be expressed more intuitively by saying that the following probabilistic motion describes a Markov chain with stationary transition probabilities p : Start at x , and then with probability $1 - f(x)$ move off according to a Markov chain with transition probabilities q ; when (if ever) this chain dies, look back at the position y where the chain was at the instant before it died and instead of dying make a single transition according to $r(y, \cdot)$; if this gets you to state z (where $z = x$ if there was no motion according to q) complete the motion by moving forevermore according to a Markov chain with transition probabilities s starting at z .

There are two basic kinds of random times which induce a path decomposition of this kind: terminal times and coterminial times. We first indicate how the parameters f , q , r , and s are obtained for these times, and then show that for nice transition matrices p these are essentially the only random times inducing such a path decomposition.

Suppose first that τ is a terminal time. Then as mentioned below (5.3) there is a subset H of J and a subset V of $H \times H$ such that

$$(\tau > 0) = (X_0 \in H), \quad (\tau > n) = ((X_k, X_{k+1}) \in V, 0 \leq k \leq n-1), \quad n \geq 1.$$

With H and V so defined it easily is checked that the parameters f , q , r and s are given by

$$\begin{aligned} f(x) &= 1_H(x), \quad x \in J, \\ q(x,y) &= \begin{cases} p(x,y)1_V(x,y) & \text{if } y \neq \Delta \\ 1 - \sum_{z \neq \Delta} q(x,z) & \text{if } y = \Delta \end{cases} \\ r(x,y) &= \begin{cases} p(x,y)1_{V^C}(x,y) / \sum_z p(x,z)1_{V^C}(x,z) & \text{if } x \in H \\ \text{arbitrary} & \text{if } x \in H^C, \end{cases} \\ s(x,y) &= p(x,y), \end{aligned}$$

where 1_B stands for the indicator function of a subset B of either J or $J \times J$.

For τ a coterminal time there is a subset V of $J \times J$ and an invariant event C_∞ such that

$$(\tau \leq n) = (\theta_n \in C), \quad n \in \mathbb{N},$$

where C is the coterminal event $((X_k, X_{k+1}) \in V, k \in \mathbb{N})C_\infty$. Define functions f , \bar{f} , g and h from J to $[0,1]$ by setting for $x \in J$

$$\begin{aligned} f(x) &= P^x(\tau > 0) = P^x(C^C) \\ \bar{f}(x) &= P^x(\tau = 0) = P^x(C) \\ g(x) &= P^x(\tau = 1) = P^x[(X_0, X_1) \in V^C, \theta_1 \in C] \\ h(x) &= P^x(\tau = \infty) = P^x[((X_n, X_{n+1}) \in V^C \text{ infinitely often}) \cup C_\infty^C]. \end{aligned}$$

Then f , \bar{f} and g are related by the identities $f + \bar{f} = 1$ and

$$g(x) = \sum_y p(x,y)1_{V^C}(x,y)f(y), \quad x \in J,$$

and it may also be observed that f is p -excessive, that h is p -harmonic, and

that the Riesz decomposition of f is $f = Ug + h$ where $U = \sum_{n=0}^{\infty} p^n$ is the potential operator associated with p . The parameters f , q , r and s for the path decomposition are now readily seen to be specified as follows: f has already been defined,

$$q(x,y) = \begin{cases} p(x,y)f(y)/f(x) & \text{if } f(y) > 0, y \neq \Delta, \\ 0 & \text{if } f(y) = 0, y \neq \Delta, \\ 1 - \sum_{z \neq \Delta} q(x,z) & \text{if } y = \Delta, \end{cases}$$

$$r(x,y) = \begin{cases} p(x,y)1_{V^c}(x,y)f(y) & \text{if } g(y) > 0, \\ \text{arbitrary} & \text{if } g(y) = 0, \end{cases}$$

$$s(x,y) = \begin{cases} p(x,y)1_V(x,y)\bar{f}(y)/\bar{f}(x) & \text{if } \bar{f}(x) > 0, \\ \text{arbitrary} & \text{if } \bar{f}(x) = 0. \end{cases}$$

In this case parts (i), (ii) and (iv) of the path decomposition statement (6.1) are the discrete analogues of Theorems (2.1) and (5.1) of Meyer, Smythe and Walsh [3]. Part (iii) provides the inner link between the pre- τ and post- τ processes which is required for the full statement of the path decomposition.

We now establish a characterization of terminal and coterminial times by the path decomposition (6.1).

(6.2) Theorem. Let the Markov probabilities $\{P^x, x \in J\}$ be induced by a transition matrix p on J which makes all states in $J \setminus \{\Delta\}$ form a single closed communicating class. Then a random time τ is both a regular birth time and a regular death time for P^x if and only if τ is P^x -equivalent to a random time which is either terminal or coterminial.

Remark. The characterization fails without some hypothesis on p . For instance if p induces two closed communicating classes A and B and a transient state x from which absorption into either A or B is possible, then τ could be equal to a terminal time on paths entering A and to a coterminial time on paths entering B .

Proof. The 'if' part is contained in Theorems (3.9) and (5.2). For the 'only if' part observe first that if τ is a regular birth time and a regular death time for P^X , then by the same results

$$(6.3) \quad \tau = \tau_C + \rho \quad P^X \text{ a.s.},$$

$$(6.4) \quad \tau = \tau_{V,F} \quad P^X \text{ a.s.}$$

where τ_C is a coterminial time associated with some coterminial event C , ρ is an optional time for $\{F_{\tau_C+n}, n \geq 0\}$ and $\tau_{V,F}$ is given as in (5.1) for some $V \subset J \times J$, $F \in \mathcal{F}$. In particular, if $\tau_V \geq 1$ is the terminal time associated with V , we have

$$(6.5) \quad \tau_C \leq \tau \leq \tau_V \quad P^X \text{ a.s.}$$

Now C is the intersection of an invariant event C_∞ with an event requiring that all transitions belong to some subset, V'^C say, of $J \times J$. By (6.5) therefore $\tau_C = \tau = \tau_V = \infty$, P^X a.s. on C_∞^C so that $C_\infty^C \subset ((X_n, X_{n+1}) \in V'^C, n \in \mathbb{N})$, P^X a.s., which because the states form one communicating class is possible only if either $P^X C_\infty^C = 0$ or transitions in V are impossible, i.e. we must have that either τ_C equals P^X a.s. the last time σ_V , a transition in V' is completed or else $\tau_V = \infty$, P^X a.s. But in the first case the inequality $\sigma_V \leq \tau_V$ shows that it is impossible to perform a transition in V' after a

transition in V , and again because we have one communicating class, this means that either V' or V consists of transitions which are impossible under p . Thus either $\tau_C = 0$, P^X a.s. or $\tau_V = \infty$, P^X a.s. and by (6.3), (6.4) this means that either there is an optional time ρ which equals τ , P^X a.s., or there is an event $F \in \mathcal{F}$ such that $\tau = \tau_{\emptyset, F} = \sup\{n \geq 1: \theta_{n-1} \in F\}$, P^X a.s.

We have already seen that if, e.g. τ is itself an optional time satisfying (6.4) exactly, then τ is a terminal time. If we only know that $\tau = \rho$, P^X a.s. with ρ optional it follows that $P^X(\tau > n | \mathcal{F}_n) = 1_{(\tau > n)}$, P^X a.s. while because of (6.4) the same conditional probability equals $1_{(\tau_V > n)} P^{X(n)}(\tau_{V, F} > 0)$. But then

$$(\tau > n) = (\tau_V > n, P^{X(m)}(\tau_{V, F} > 0) > 0, 0 \leq m \leq n) \quad P^X \text{ a.s.}$$

showing that τ is P^X -equivalent to a terminal time.

Similarly, if $\tau = \tau_{\emptyset, F}$ P^X a.s. we have $P^X(\tau \leq n | \mathcal{F}_n) = P^{X(n)}(\tau_{\emptyset, F} = 0)$ while because of (6.3), $(\tau \leq n) = F_n(\theta_n \in C)$ for some $F_n \in \mathcal{F}_n$ so that the conditional probability equals $1_{F_n} P^{X(n)}_C$ and consequently $F_n(P^{X(n)}_C > 0) = (P^{X(n)}(\tau_{\emptyset, F} = 0) > 0) \quad P^X$ a.s. whence

$$(\tau \leq n) = F_n(\theta_n \in C, P^{X(n)}_C > 0) = (\theta_n \in C, P^{X(m)}(\tau_{\emptyset, F} = 0) > 0, m \geq n) \quad P^X \text{ a.s.,}$$

showing that τ is P^X -equivalent to a coterminal time. The theorem is proved.

7. POSSIBLE EXTENSIONS

Though we have restricted ourselves in this paper to Markov processes in discrete time with countable state space, the concepts of birth times, death times, and conditional independence times can all be formulated for Markov

processes with more general time set or state space. We conclude in this section with some comments on the difficulties involved in extending our results to apply to these situations.

A few of the results do carry over to apply to Markov processes with abstract measurable state space and time set T either \mathbb{N} or $[0, \infty)$. Copying the definition from Section 4 of a death time for a family of Markov probabilities, a generalization of Proposition (4.2) remains valid, and by adopting an analogous definition of a conditioning event for a family rather than a single Markov probability, it can be shown that F is such an event if and only if there are events $C \in \mathcal{F}$, $F_t \in \mathcal{F}_t$, $t \in T$ such that $F = F_t(\theta_t \in C)$, P^x a.s. for every x and every $t \geq 0$ (cf. (2.10) and Jacobsen [2], Lemma 1).

As for our other results, the assumption of a countable state space is used chiefly to avoid measure theoretical problems in the proofs of the harder 'only if' assertions of Theorems (2.3), (3.9) and (5.2), while the restriction to discrete time is essential for our treatment of conditional independence across a random time. The basic criterion for deciding whether a random time possesses the conditional independence property is Lemma (3.12), the proof of which relies on the fact that within the set $A(\tau=n)$, where A is an arbitrary atom in \mathcal{F}_n , the conditional probability law of θ_τ given the pre- τ field may be determined as a conditional P^y -probability given the event G_{nA} which is the section of $(\tau=n)$ beyond A . A generalization of this to processes in continuous time fails, partly because the conditioning event G_{nA} may now have measure 0 for more than a negligible collection of atoms A , and partly because, even when this is not the case, it is no longer obvious that the desired conditional probability given \mathcal{F}_τ results. For some criteria for conditional independence and some of the subtleties involved see Jacobsen [2], Pittenger and Shih [6].

REFERENCES

- [1] Freedman, D. (1971). Markov Chains, Holden-Day, San Francisco.
- [2] Jacobsen, M. (1974). Splitting times for Markov processes and a generalized Markov property for diffusions. Z. Wahrscheinlichkeitstheorie und Verw. Gebiete 30, 27-43.
- [3] Meyer, P.A., Smythe, R.T., and Walsh, J.B. (1972). Birth and death of Markov processes. Proc. 6th Berkeley Symposium Math. Statist. Prob. Univ. Calif. III 295-305.
- [4] Pitman, J.W. (1974). Path decomposition for conditional Brownian motion. Inst. of Math. Statist., Univ. of Copenhagen, Preprint No. 11.
- [5] Pitman, J.W. (1975). 1-dimensional Brownian motion and the 3-dimensional Bessel process. To appear in Adv. Appl. Prob.
- [6] Pittenger, A.O., Shih, C.T. (1973). Coterminal families and the strong Markov property. Trans. Amer. Math. Soc. 182, 1-42.
- [7] Walsh, J., and Weil, M. (1972). Representation de temps terminaux et applications aux fonctionelles additives et aux systèmes de Lévy. Ann. Sci. École Norm. Sup., 4^e serie, V5, 121-155.
- [8] Williams, D. (1970). Decomposing the Brownian path. Bull. Amer. Math. Soc. 76, 871-873.
- [9] Williams, D. (1974). Path decomposition and continuity of local time for one-dimensional diffusions. Proc. London Math. Soc. Third Series, 28, 738-768.