DISCOUNTED CONTINUOUS TIME MARKOV DECISION PROCESSES: THE CONVEX ANALYTIC APPROACH

A.B. Piunovskiy

The University of Liverpool, England.
piunov@liverpool.ac.uk

Abstract: The convex analytic approach which is dual, in some sense, to dynamic programming, is useful for the investigation of multicriteria control problems. It is well known for discrete time models, and the current paper presents similar results for the continuous time case. Namely, we define and study the space of occupation measures, and apply the abstract convex analysis to the study of constrained problems. Finally, we briefly consider a meaningful example on a controlled bicriteria Markovian queue.

Keywords: Stochastic systems, Optimal control, Jump process, Markov Decision Processes (MDP), Constraints, Queuing theory

1. INTRODUCTION

It is well known that dynamic programming approach, so effective in the standard optimal control theory, is inconvenient for the study of multiobjective problems. In such situations, the so called "convex analytic approach" proved to be useful for discrete-time models (Altman, 1999; Hernandez-Lerma and Laserre, 1999; Piunovskiy, 1997). The author does not know any systematic application of this method to the case of continuous-time jump models, and the present report has to fill partially this gap.

The convex analytic approach mainly deals with occupation measures (see Section 3). We present Theorem 1 where the key properties of those measures are formulated. Perhaps, the main result is Lemma 2, because all the subsequent statements can be proved based on the well known facts about the discrete-time occupation measures. Section 4 contains the results about constrained problems, and in Section 5, a new meaningful example is presented. Let us underline Remark 1 clarifying the duality between the dynamic programming and convex analytic approaches.

Note that some constrained discounted jump models were considered earlier (Feinberg, 2004; Guo and Hernandez-Lerma, 2003; Lai and Tanaka, 1991; Piunovskiy, 1998). Several other constrained versions were studied in the papers by Piunovskiy (2004a, 2004b). First linear programs for jump Markov models were suggested by Mine and Tabata (1970). At the same time, an exhaustive mathematically coherent convex-analytical study was absent so far.

Lai and Tanaka (1991) and Guo and Hernandez-Lerma (2003) considered only Markov strategies and developed the Lagrange multipliers technique (cf. Theorem 3), without analysing the occupation measures. Feinberg (2004) started with general strategies and showed how to reduce the model to discrete-time MDP. After that, one can obviously apply the corresponding convex analytic approach. The article (Piunovskiy, 1998) contains some preliminaries, mainly for a non-Markov model.

The theoretical results of the present report can be deduced, in principle, from (Feinberg, 2004) and old facts about discrete-time MDP (Altman, 1999; Hernandez-Lerma and Laserre, 1999; Piunovskiy, 1997). But the reasoning is absolutely different. First of all, we characterize occupation measures in Lemma 2 and prove the sufficiency of stationary strategies in Corollary 1. After that, we study linear programs (10), (11). All the statements are formulated in terms of the original continuous time model; results on discrete-time MDP are invoked only to speed up the proof of Theorem 1 which states the main properties of the occupation measures space. This approach seems to be more natural and understandable for the readers, since we avoid jumping from one model to another.

2. MATHEMATICAL MODEL

The model below is constructed by analogy with (Kitaev and Rykov, 1995).

Let X be a countable states space and $(A, \mathcal{B}(A))$ be a Borel actions space. Letting $\Omega^0 \triangleq (X \times \mathbb{R}_+)^\infty$ join to Ω^0 all the sequences of the form $(x_0, \theta_1, x_1, \ldots, \theta_{n-1}, x_{n-1}, \infty, x_\infty, \infty, x_\infty, \ldots)$, $n \geq 1$, where $\forall k = 1, 2, \ldots, n-1$, $\theta_k \neq \infty$ and $x_k \neq x_\infty$; x_∞ is an isolated point added to X. As a consequence we obtain the basic measurable space $(\Omega, \mathcal{F})$, where $\mathcal{F}$ is the natural σ-algebra. Set $T_0 \triangleq 0$, $T_n \triangleq \theta_1 + \theta_2 + \ldots + \theta_n$, $T_\infty \triangleq \lim_{n\to\infty} T_n$, and

$$\xi_t \triangleq \sum_{n\geq 0} I\{T_n \leq t < T_{n+1}\} x_n + I\{T_\infty \leq t\} x_\infty. \tag{1}$$

Suppose a measurable function $\lambda(j|i, a) \geq 0$ is defined for $j \neq i$ and put $\lambda(i, a) \triangleq \sum_{j \in X \setminus i} \lambda(j|i, a)$.

θ_n play the role of inter-jump intervals or sojourn times; T_n are the jump moments, x_n is the state of the controlled process on $[T_n, T_{n+1})$, and $\lambda(\cdot)$ is the jumps intensity of ξ_t if action a is chosen.

Introduce the integer-valued random measure

$$\mu(\omega, dt, i) = \sum_{n\geq 1} I\{T_n < \infty\} I\{x_n = i\} \delta_{T_n}(dt)$$

(where $\delta_y(\cdot)$ is the Dirac measure concentrated at the point y) and also the right-continuous family of σ-algebras $\{\mathcal{F}_t\}_{t\geq 0}$: $\mathcal{F}_t \triangleq \sigma\{x_0, \mu([0, \tau] \times i), \tau \in [0, t], i \in X\}$; $\mathcal{F}_\infty \triangleq \bigvee_{t\geq 0} \mathcal{F}_t$. Let $\mathcal{P}$ be the σ-algebra of predictable sets on $\Omega \times \mathbb{R}^0_+$ related to $\{\mathcal{F}_t\}_{t\geq 0}$.

Definition. A $\mathcal{P}$-measurable transition probability $\pi(\cdot|\omega, t)$ on $(A_\infty, \mathcal{B}(A_\infty))$ is called a strategy π; also $\pi(a_\infty|\omega, t) = I\{T_\infty \leq t\}$ is a standing assumption, a_∞ being an isolated point added to A; $A_\infty \triangleq A \cup \{a_\infty\}$. A strategy is called a selector and denoted by φ whenever there exists a predictable A_∞-valued process $\varphi(\omega, t)$ such that $\pi(\Gamma^A|\omega, t) = I\{\Gamma^A \ni \varphi(\omega, t)\}$ $\forall \Gamma^A \in \mathcal{B}(A)$. A strategy is called stationary if it has the form $\pi(\cdot|\xi_{t-}(\omega))$. A selector $\varphi(\xi_{t-}(\omega))$ is called stationary, too.

Evidently, for any control strategy π, the random measure

$$\nu^\pi(\omega, dt, i) \triangleq \left[\int_A \pi(da|\omega, t)\lambda(i|\xi_{t-}(\omega), a)\right] dt \tag{2}$$

is predictable and $\nu^\pi(\omega, \{t\} \times X) = \nu^\pi(\omega, [T_\infty, \infty) \times X) = 0$. So, for any strategy π and any initial distribution P_0 on X (which is supposed to be given and fixed), there exists a unique probability measure $P^\pi_{P_0}$ on $(\Omega, \mathcal{F})$ such that $P^\pi_{P_0}\{x_0 = i\} = P_0(i)$ and ν^π is a predictable projection of the measure μ (Kitaev and Rykov, 1995).

If a strategy π is chosen then we have a stochastic basis $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, P^\pi_{P_0})$ which is always assumed to be completed, and the main controlled random process ξ_t thereon. In what follows, $E^\pi_{P_0}$ denotes the mathematical expectation with respect to $P^\pi_{P_0}$. Since the initial distribution is assumed to be fixed, we usually omit that index: P^π, E^π.

Main Optimization Problem. Suppose the measurable loss rates

$$r_0(i, a), \quad r_1(i, a), \quad \ldots, \quad r_N(i, a)$$

are defined on $X \times A$. For a fixed strategy π and for any $n = 0, 1, \ldots, N$ we define $R_n(\pi)$ as follows

$$R_n(\pi) = E^\pi\left[\int_0^\infty e^{-\alpha t}\left[\int_A \pi(da|\omega, t) r_n(\xi_{t-}, a)\right] dt\right], \tag{3}$$

where $\alpha > 0$ is a fixed discount factor. Here and below, we assume that all the integrals are well defined, e.g. each function $r(\cdot), s_n(\cdot), n = 1, 2, \ldots, N$ is bounded (below or above).

The multiple-objective optimization problem

$$R_n(\pi) \longrightarrow \min_\pi, \qquad n = 0, 1, \ldots, N \tag{4}$$

is usually inconsistent. In what follows, we intend to investigate the constrained version

$$R_0(\pi) \longrightarrow \min_\pi, \quad R_n(\pi) \leq d_n, \quad n = 1, 2, \ldots, N, \tag{5}$$

under appropriate fixed constants $d_1, d_2, \ldots, d_N$.

3. OCCUPATION MEASURES AND THEIR PROPERTIES

Definition. Two strategies π^1 and π^2 are called equivalent if the equality

$$S(\pi^1) = S(\pi^2) \tag{6}$$

holds for any bounded non-negative measurable loss function $s(x,a)$. (Here and below $S(\pi)$ is defined similarly to (3).)

Definition. The occupation measure for the strategy π is the probability measure $\eta^\pi(i, da)$ on the space $X \times A$ which is defined in the following way:

$$\eta^\pi(i, \Gamma^A) = \alpha E^\pi \left[\int_0^\infty e^{-\alpha t} I\{\xi_{t-} = i\} \pi(\Gamma^A|\omega, t) dt \right]. \tag{7}$$

Lemma 1. Two strategies π^1 and π^2 are equivalent if and only if the corresponding occupation measures coincide: $\eta^{\pi^1} = \eta^{\pi^2}$.

Our goal is to prove that any occupation measure is generated by some *stationary* strategy meaning that many different strategies are indeed equivalent.

Condition 1. $\lambda(i,a) \le K < \infty$.

Lemma 2. Let Condition 1 be satisfied. Then a measure $\eta(i, \Gamma^A)$ on $X \times A$ is an occupation measure if and only if equality

$$\alpha\eta(i,A) - \sum_{k \in X \setminus i} \int_A \lambda(i|k,a)\eta(k,da) + \int_A \lambda(i,a)\eta(i,da) = \alpha P_0(i) \tag{8}$$

holds for each $i \in X$.

Corollary 1. Under Condition 1, for any strategy π, there exists a stationary strategy π^s such that the corresponding occupation measures coincide: $\eta^\pi = \eta^{\pi^s}$. Moreover, π^s can be found from the representation

$$\eta^\pi(i, \Gamma^A) = \eta^\pi(i, A)\pi^s(\Gamma^A|i). \tag{9}$$

Let $\mathcal{D}$ be the space of all occupation measures in the model, endowed with the weak topology.

Theorem 1. Let Condition 1 be satisfied.

(a) The space $\mathcal{D}$ is convex.

(b) A point η is extreme in $\mathcal{D}$ if and only if there exists a stationary selector φ such that $\eta = \eta^\varphi$.

(c) If the jumps intensity $\lambda(i|j,a)$ and the total rate out of j, $\lambda(j,a)$, are continuous (in a) then $\mathcal{D}$ is closed in $\mathcal{P}(X \times A)$.

(d) If, additionally, A is compact then $\mathcal{D}$ is a metrizable compact and coincides with the closed convex hull of the set $\mathcal{D}^\varphi$ of the occupation measures generated by stationary selectors.

The (indirect) proof is based on the following reasonings. Let introduce denotation

$$p_{ij}(a) \triangleq \begin{cases} \lambda(j|i,a)/K, & \text{if } j \ne i; \\ \frac{K-\lambda(i,a)}{K}, & \text{if } j = i. \end{cases}$$

Obviously, p defines the transition probability for a (discrete time) MDP with the same state space X, initial distribution P_0, and action space A. Suppose $\beta = \frac{K}{\alpha+K}$ is the discount factor. Then the space of occupation measures $\tilde{\mathcal{D}}$ in this model coincides with the space of probability measures η satisfying equation

$$\eta(i,A) - (1-\beta)P_0(i) - \beta \sum_{k \in X} \int_A p_{ki}(a)\eta(k,da) = 0$$

(see e.g. (Piunovskiy, 1997)), which is identical with (8). Hence $\tilde{\mathcal{D}} = \mathcal{D}$. Now properties (a)–(d) follow from the corresponding theorems proved for discounted MDP. (See the books by Altman (1999), Hernandez-Lerma and Lasserre (1999), and Piunovskiy (1997), and references therein.)

4. MAIN OPTIMIZATION PROBLEM: SOLVABILITY AND SUFFICIENT CLASSES OF STRATEGIES

In view of Corollary 1, when analysing problem (4), it is sufficient to consider only (randomized) stationary strategies. It is convenient to rewrite (4) and (5) in the form

$$\sum_{i \in X} \int_A r_n(i,a)\eta(i,da) \longrightarrow \inf_{\eta \in \mathcal{D}}, \tag{10}$$

$$n = 0, 1, \dots, N;$$

$$\begin{cases} \sum_{i \in X} \int_A r_0(i,a)\eta(i,da) \longrightarrow \inf_{\eta \in \mathcal{D}}, \\ \sum_{i \in X} \int_A r_n(i,a)\eta(i,da) \le d_n, \; n = 1, \dots, N. \end{cases} \tag{11}$$

In acordance with Lemma 2, $\eta \in \mathcal{D}$ iff equality (8) is satisfied. Hence (11) is a linear program (LP) on the space of finite measures which will be called *primal*.

Remark 1. Suppose $N = 0$, i.e. there are no constraints, and assume for simplicity that function r_0 is bounded. If we consider η as a (non-negative) element of the linear space of finite signed measures then the dual space consists of real bounded functions on $X \times A$. In this framework, the dual program to (11) looks as follows:

find a bounded function $w(i)$ on X such that

$$\begin{aligned} &(1-\beta)\sum_{i\in X} w(i)P_0(i) \longrightarrow \sup \\ &r_0(i,a) + \beta \sum_{j\in X} w(j)p_{ij}(a) - w(i) \geq 0. \end{aligned} \tag{12}$$

One can easily check that LP (12) is equivalent to solving the following problem

$$\alpha w(i) = \inf_{a\in A}\{(\alpha+K)r_0(i,a)$$

$$+ \sum_{j\in X\setminus i} w(j)\lambda(j|i,a) - w(i)\lambda(i,a)\}, \quad i \in X,$$

which, up to the constant $(\alpha + K)$, gives the Bellman function for the (scalar) problem (4) (Guo and Hernandez-Lerma, 2003; Kitaev and Rykov, 1995). The value $\frac{\alpha}{\alpha+K}\sum_{i\in X} w(i)P_0(i)$ coincides with the value of LP (12). Therefore, the convex analytic approach and dynamic programming are represented by the couple of *dual* linear programs.

Definition. A model is called *semicontinuous* if Condition 1 is satisfied, all the assumptions of Item (d) of Theorem 1 hold, and $\forall n = 0, 1, \ldots, N$ function $r_n(\cdot)$ is lower semicontinuous and bounded below. To put it different,
$\lambda(i, a) \leq K < \infty$;
$\lambda(j|i, a)$ and $\lambda(i, a)$ are continuous functions;
A is compact;
$r_n(i, a)$ are lower semicontinuous and bounded below functions.

Theorem 2. In a semicontinuous model, if there exists $\hat{\eta} \in \mathcal{D}$ satisfying all the inequalities in (11), then constrained problem (11) is solvable, and the solution can be found in the class of occupation measures generated by stationary strategies.

Condition 2. (Slater) There exists a point $\hat{\eta} \in \mathcal{D}$ such that all the inequalities in (11) are strict.

Theorem 3. (Kuhn-Tucker) Suppose a model satisfies Condition 2. Then a point $\eta^* \in \mathcal{D}$ is a solution to (11) if and only if there exists a vector $Y_* \in (\mathbb{R}^0_+)^N$ for which one of the following two equivalent assertions holds:

(a) the pair (η^*, Y_*) is a saddle point of the Lagrange function L defined on $\mathcal{D} \times (\mathbb{R}^0_+)^N$:

$$L(\eta, Y) \triangleq \sum_{i\in X}\int_A r_0(i,a)\eta(i,da)$$

$$+ \sum_{n=1}^{N} Y_i \left[\sum_{i\in X}\int_A r_n(i,a)\eta(i,da) - d_n\right].$$

To put it differently,

$$L(\eta^*, Y) \leq L(\eta^*, Y_*) \leq L(\eta, Y_*)$$

at any $(\eta, Y) \in \mathcal{D} \times (\mathbb{R}^0_+)^N$.

(b) All the inequalities in (11) are fulfilled, $L(\eta^*, Y_*) = \min_{\eta\in\mathcal{D}} L(\eta, Y_*)$, and the condition of *complementary slackness*

$$\sum_{n=1}^{N} Y_{*i} \left[\sum_{i\in X}\int_A r_n(i,a)\eta^*(i,da) - d_n\right] = 0$$

is satisfied.

A vector $Y_* \in (\mathbb{R}^0_+)^N$ exhibits the above properties if and only if Y_* provides the maximum

$$g(Y) \triangleq \inf_{\eta\in\mathcal{D}} L(\eta, Y) \longrightarrow \max_{Y\in(\mathbb{R}^0_+)^N}.$$

Theorem 4. Suppose a semicontinuous model satisfies Condition 2 and the functions $r_n(\cdot)$, $n = 1, 2, \ldots, N$ are continuous and bounded. Then there exists a solution to problem (11) which has the form of a mixture of $(N + 1)$ stationary selectors:

$$\eta^* = \sum_{n=1}^{N+1} \gamma_n \eta^{\varphi_n}, \qquad \gamma_n \in [0,1], \qquad \sum_{n=1}^{N+1} \gamma_n = 1.$$

(Here η^φ is the occupation measure generated by a stationary selector φ.)

All these theorems (and many other) follow from the corresponding statements for discounted (discrete time) MDP presented e.g. in the books by Altman (1999), Hernandez-Lerma and Lasserre (1999), and Piunovskiy (1997): see the end of Section 3.

5. EXAMPLE

Let us consider a switch (router) transforming information. Packets (file segments sent from remote terminals) arrive to the router one after another and, following rapid consideration, are stored in the router buffer. After that, they are transmitted further using the First In – First Out basis, transmission time depending on the length of the packet. In the absence of congestion, the intensity of the input stream increases until $\overline{\lambda}$ which is connected with the so called 'window size'. In the case of congestion, the switch can send feedback signals to the terminals asking them to decrease the window size; which results in the decrease of the input intensity, up to the low boundary $\underline{\lambda}$. The service rate μ remains constant. Therefore, we intend to consider the $M/M/1/\infty$ queueing system with a *controlled input stream*.

In the terms of Section 2, $X = \{0, 1, 2, \ldots\}$, $A = [\underline{\lambda}, \overline{\lambda}]$; for $j \neq i \in X$, $\lambda(j|i,a) = \begin{cases} a, & \text{if } j = i+1, \\ \mu, & \text{if } j = i-1, \\ 0, & \text{otherwise}. \end{cases}$ The random process ξ_t describes the number of packets in the system. Put $r_0(x,a) = x$ and $r_1(x,a) = -a$. Now problem (5), for a negative value $-d$ on the right side, coinsides with the following:

minimize the discounted queue length

$$R_0(\pi) = E^\pi \left[\int_0^T e^{-\alpha t} \xi_{t-} dt \right] \to \min$$

under the constraint on the discounted throughput $R_1(\pi) \leq -d$, i.e.

$$S(\pi) = E^\pi \left[\int_0^T e^{-\alpha t} \lambda(t) dt \right] \geq d.$$

(Here $\lambda(t) = \int_A a\pi(da|\omega, t)$ is the intensity of the input stream at time moment t.)

Application of the theory developed above leads to the following qualitative result. Solution to the constrained problem stated is given by the threshold stationary selector

$$\varphi(\xi_{t-}) = \begin{cases} \overline{\lambda}, & \text{if } \xi_{t-} < i^*, \\ \gamma\overline{\lambda} + (1-\gamma)\underline{\lambda}, & \text{if } \xi_{t-} = i^*, \\ \underline{\lambda}, & \text{if } \xi_{t-} > i^*. \end{cases}$$

The values of the parameters $i^* \geq -1$ and $\gamma \in (0, 1]$ can be effectively calculated, but the expressions are cumbersome.

6. CONCLUSION

The convex analytic approach is particularly useful for the study of multiple-criteria control problems. We managed to obtain all the main theorems because, after Lemma 2 is established, one can easily use many known results from the discrete-time theory. It should be emphasized that the jumps intensity was assumed to be uniformly bounded. If this condition is not satisfied, the theory becomes much more complicated.

If there are no constraints ($N = 0$) then Theorem 2 gives the well known result on the suffieciency of stationary strategies. Moreover, Theorem 4 says that the solution can be found in the class of stationary selectors. These facts are known long ago (Yushkevich, 1980).

The example considered in Section 5 is of its own interest. Similar model with long-run average losses was studied in depth in (Piunovskiy, 2004b).

7. ACKNOWLEDGEMENTS

The author is grateful to Prof. O.Hernandez-Lerma (Cinvestav-IPN, Mexico) for the fruitful discussion during his visit in Liverpool in May 2004, supported by the London Mathematical Society, Scheme 2 grant. Besides, the example from Section 5 was discussed with Dr. E.Altman, Dr. K.Avrachenkov, and Mr. U.Ayesta (INRIA, project 'MISTRAL', France) during the workshops in Sophia-Antipolis and in Liverpool in 2002-03 supported by The British Council (Franco-British Partnership program. Project No. PN02.045).

8. APPENDIX

Sketch of the proof of Lemma 2. Let us fix a strategy π and the corresponding occupation measure η^π. As is known (Kitaev and Rykov, 1995), the measures

$$K^\pi(t, i, \Gamma^A) \stackrel{\triangle}{=} E^\pi[\pi(\Gamma^A|\omega, t) I\{\xi_{t-} = i\}] \text{ and}$$

$$p_i^\pi(t) \stackrel{\triangle}{=} P^\pi\{\xi_t = i\} = P^\pi\{\xi_{t-} = i\} = K^\pi(t, i, A)$$

on $X \times A$ and on X correspondingly, satisfy the forward Kolmogorov equation

$$p_i^\pi(t) = P_0(i) + \int_0^t \sum_{k \in X \setminus i} \int_A K^\pi(s, k, da)\lambda(i|k, a) ds$$

$$- \int_0^t \int_A K^\pi(s, i, da)\lambda(i, a) ds.$$

Hence, using the Fubini theorem and omitting the algebraic calculations, we have

$$\eta^\pi(i, A) = \alpha \int_0^\infty e^{-\alpha t} p_i^\pi(t) dt = P_0(i)$$

$$+ \sum_{k \in X \setminus i} \int_A \lambda(i|k, a) \times \int_0^\infty e^{-\alpha t} K^\pi(t, k, da) dt$$

$$- \int_A \lambda(i, a) \int_0^\infty e^{-\alpha t} K^\pi(t, i, da) dt$$

and equality (8) follows.

Suppose now that a measure η satisfies (8) and disintegrate it:

$$\eta(i, \Gamma^A) = \eta(i, A)\pi^s(\Gamma^A|i),$$

where π^s is a certain stochastic kernel. After that, one can show that $\eta = \eta^{\pi^s}$.

REFERENCES

Altman, E. (1999). *Constrained Markov Decision Processes*, Chapman and Hall/CRC Press, London.

Feinberg, E.A. (2004) Continuous discounted jump Markov decision processes: a discrete-event approach, *Mathem. Oper. Research*, **29**, 492-524.

Guo, X. and O.Hernandez-Lerma (2003). Constrained continuous-time Markov control processes with discounted criteria, *Stoch. Anal. Appl.*, **21**, 379-400.

Hernandez-Lerma, O. and J.-B.Lasserre (1999). *Further Topics on Discrete-Time Markov Control Processes*, Springer, New York.

Kitaev, M.Y. and V.V.Rykov (1995). *Controlled Queueing Systems*, CRC Press, New York.

Lai, H.C. and K.Tanaka (1991). On continuous-time discounted stochastic dynamic programming, *Appl. Math. Optim.*, **23**, 155-169.

Mine, H. and Y.Tabata (1970). Linear programming and continuous Markovian decision problems, *J. Appl. Prob.*, **7**, 657-666.

Piunovskiy, A.B. (1997). *Optimal Control of Random Sequences in Problems with Constraints*, Kluwer Academic Publishers, Dordrecht-Boston-London.

Piunovskiy, A.B. (1998). A controlled jump discounted model with constraints, *Theory Probab. Appl.*, **42**, 51-72.

Piunovskiy, A.B. (2004a). Optimal interventions in countable jump Markov processes, *Mathem. Oper. Research*, **29**, 289-308.

Piunovskiy, A.B. (2004b). Bicriteria optimization of a queue with a controlled input stream, *Queueing Systems*, **48**, 159-184.

Yushkevich, A.A. (1980). Controlled jump Markov models, *Theory Probab. Appl.*, **25**, 244-266.