Optimal stopping for partially observed piecewise deterministic Markov processes

Adrien Brandejsky, Benoîte de Saporta, François Dufour
INRIA and Univ. Bordeaux

ANR-09-SEGI-004 Fautocoes

XIème congrès franco-roumain de mathématiques appliquées

Outline

Definition of piecewise deterministic Markov processes

Davis (80's)

General class of non-diffusion dynamic stochastic hybrid models: deterministic motion punctuated by random jumps.

Applications

Engineering systems, operations research, management science, economics, dependability and safety, maintenance,. . .

Dynamics

Hybrid process $X_t = (m_t, y_t)$

- discrete mode $m_t \in \{1, 2, \dots, p\}$
- Euclidean state variable $y_t \in \mathbb{R}^n$

Local characteristics for each mode m

- E_m open subset of $\mathbb{R}^d$, ∂E_m its boundary and $\overline{E}_m$ its closure
- Flow ϕ_m: $\mathbb{R}^d \times \mathbb{R} \to \mathbb{R}^d$ deterministic motion between jumps, one-parameter group of homeomorphisms
- Intensity λ_m: $\overline{E}_m \to \mathbb{R}_+$ intensity of random jumps
- Markov kernel Q_m on $(\overline{E}_m, \mathcal{B}(\overline{E}_m))$ selects the post-jump location

Two types of jumps

- $t^*(m,y)$ deterministic exit time

$$t^*(m,y) = \inf\{t > 0 : \phi_m(y,t) \in \partial E_m\}$$

- law of the first jump time T_1

$$\mathbb{P}_{(m,y)}(T_1 > t) = \begin{cases} \mathrm{e}^{-\int_0^t \lambda_m\left(\phi_m(y,s)\right)ds} & \text{if} \quad t < t^*(m,y) \\ 0 & \text{if} \quad t \geq t^*(m,y) \end{cases}$$

- T_1 has a density on $[0, t^*(m,y)[$ but has an atom at $t^*(m,y)$:

$$\mathbb{P}_{(m,y)}(T_1 = t^*(m,y)) > 0$$

Iterative construction

Starting point

$$X_0 = Z_0 = (m, y)$$

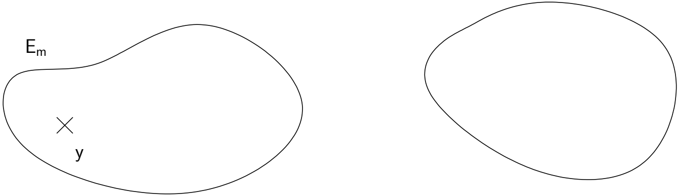

Iterative construction

X_t follows the deterministic flow until the first jump time $T_1 = S_1$

$$X_t = \big(m, \phi_m(y, t)\big), \quad t < T_1$$

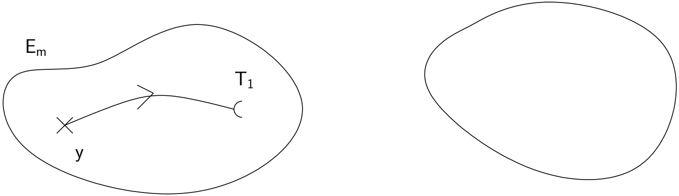

Iterative construction

Post-jump location $Z_1 = (M_1, Y_1)$ selected by

$$Q_m(\phi_m(y, T_1), \cdot)$$

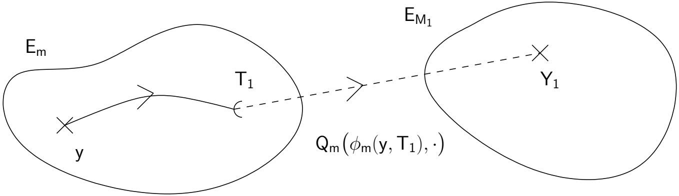

Iterative construction

X_t follows the flow until the next jump time $T_2 = T_1 + S_2$

$$X_{T_1+t} = \big(M_1, \phi_{M_1}(Y_1, t)\big), \quad t < S_2$$

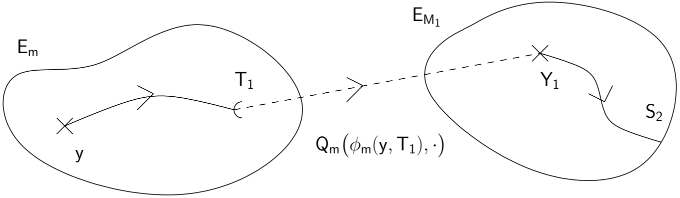

Iterative construction

Post-jump location $Z_2 = (M_2, Y_2)$ selected by

$$Q_{M_1}\left(\phi_{M_1}(Y_1, S_2), \cdot\right) \ldots$$

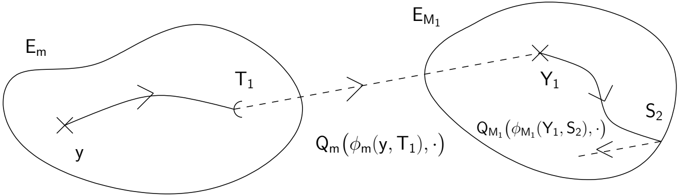

Embedded Markov chain

$\{X_t\}$ strong Markov process (M.H.A. Davis)

Natural embedded Markov chain

- Z_0 starting point, $S_0 = 0$, $S_1 = T_1$
- Z_n new mode and location after n-th jump,
  $S_n = T_n - T_{n-1}$, time between two jumps

Important property

(Z_n, S_n) is a discrete-time Markov chain
Only source of randomness of the PDMP

Optimal stopping problem

Stop the process in order to maximize a reward g

$$V(x) = \sup_{\tau \in \mathcal{M}} \mathbb{E}_x[g(X_\tau)]$$

- compute the value function V best possible performance
- compute an optimal stopping time τ

Classical problem

$\mathcal{M}$ set of stopping times for the natural filtration of (X_t)

- dynamic programming equation [Gugerli 86]
- numerical approximation based on a discretization of (Z_n, S_n) [de Saporta, Dufour, Gonzalez 10]

Partial observations

Only a noisy observation of (X_t) is available, $\mathcal{M}$ set of stopping times for the natural filtration $(\mathcal{F}_t^{\mathcal{O}})$ of the observation process

Methodology

- introduce the filter process $\Pi_t = \mathbb{E}[X_t \mid \mathcal{F}_t^{\mathcal{O}}]$
- transform the initial problem into a completely observed one for the filter process

Main drawback

- infinite dimension of the filter

State of the art

[Pham, Runggaldier, Sellami 05]
Optimal stopping under partial observation for discrete time Markov chains with finite state space

- absolute continuity assumption for the observation process
- reformulation as a standard optimal stopping problem for a continuous state space Markov chain
- numerical approximation of the value function based on discretization of the filter process

Specificities of PDMP's

- continuous time process
  - work with (Z_n, S_n),
  - $Z_n \in \{x_1, \dots, x_q\}$
  - stopping times remain continuous
- distribution of PDMP has singular components
  - study of the filter not straightforward
- non standard reformulated problem
  - derive new dynamic programming equations
  - operators not Lipschitz

Outline

Observation process

- S_n perfectly observed
- Z_n observed through a noise

$$Y_n = \phi(Z_n) + W_n$$

- continuous time observation process

$$Y_t = \sum_{n=0}^{\infty} \mathbb{1}_{[T_n, T_{n+1}[}(t) Y_n$$

- filtration $\mathcal{F}_t^{\mathcal{O}} = \sigma(Y_s, s \leq t)$

Simple example

PDMP: $E = [0,1[$,

$\Phi(x,t) = x + t$, $\lambda(x) = 3x$, $Q(x,\cdot) = x\delta_{1/4}(\cdot) + (1-x)\delta_{1/2}(\cdot)$

Observation process: $Y_n = Z_n + W_n$, $W_n \sim \mathcal{N}(0, 0.1^2)$

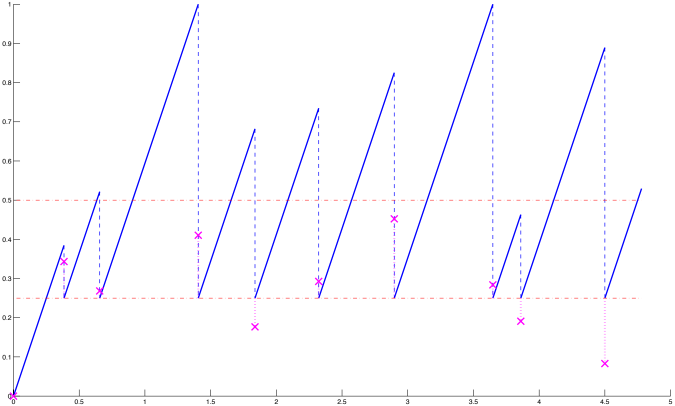

Outline

Filter process

Assumption

Finite number of possible values for Z_n,
$x_1, \dots, x_q$ with $t_i^* = t^*(x_i)$, $t_1^* \leq t_2^* \cdots \leq t_q^*$

Definition

Filter process $\Pi_n = (\Pi_n^1, \dots, \Pi_n^q)$

$$\Pi_n^i = \mathbb{P}[Z_n = x_i \mid \mathcal{F}_{T_n}^{\mathcal{O}}]$$

Recursive construction

$$\Pi_n = \Psi(\Pi_{n-1}, Y_n, S_n)$$

with

$$\begin{aligned}
\Psi^i(\pi, y, s) &= \sum_{m=0}^{q-1} \mathbb{1}_{\{s \in]t^*_m; t^*_{m+1}[\}} \frac{\psi^i_m(\pi, y, s)}{\overline{\psi}_m(\pi, y, s)} + \sum_{m=1}^{q} \mathbb{1}_{\{s = t^*_m\}} \frac{\psi^{*i}_m(y)}{\overline{\psi^*}_m(y)} \\
\psi^i_m(\pi, y, s) &= \sum_{j=m+1}^{q} \pi^j \lambda(\Phi(x_j, s)) e^{-\Lambda(x_j, s)} Q(\Phi(x_j, s), x_i) f_W(y - \varphi(x_i)), \\
\psi^{*i}_m(y) &= Q(\Phi(x_m, t^*_m), x_i) f_W(y - \varphi(x_i)), \\
\overline{\psi}_m(\pi, y, s) &= \sum_{i=1}^{q} \psi^i_m(\pi, y, s), \quad \overline{\psi^*}_m(y) = \sum_{i=1}^{q} \psi^{*i}_m(y).
\end{aligned}$$

Recursive construction

$$\Pi_n = \Psi(\Pi_{n-1}, Y_n, S_n)$$

with

$$\Psi^i(\pi, y, s) = \sum_{m=0}^{q-1} \mathbb{1}_{\{s \in]t^*_m; t^*_{m+1}[\}} \frac{\psi^i_m(\pi, y, s)}{\overline{\psi}_m(\pi, y, s)} + \sum_{m=1}^{q} \mathbb{1}_{\{s = t^*_m\}} \frac{\psi^{*i}_m(y)}{\overline{\psi^*}_m(y)}$$

Proof Use

- $\mathcal{F}^{\mathcal{O}}_{T_n} = \sigma(Y_0, S_0, \ldots, Y_n, S_n)$
- independence between (Z_n, S_n) and W_n
- expression of conditional law of (Z_n, S_n) w.r.t. Z_{n-1}

Example

PDMP: $E = [0, 1[$, $x_1 = 0$, $x_2 = 1/4$, $x_3 = 1/2$,
$\Phi(x, t) = x + t$, $\lambda(x) = 3x$, $Q(x, \cdot) = x\delta_{1/4}(\cdot) + (1 - x)\delta_{1/2}(\cdot)$
Observation process: $Y_n = Z_n + W_n$, $W_n \sim \mathcal{N}(0, 0.1^2)$

	0	1	2	3	4	5	6	7	8	9	10
Z_n	0	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25
Y_n	0	0.34	0.27	0.41	0.18	0.29	0.45	0.28	0.19	0.08	0.26
Π_n	1	0.00	0.03	0.00	0.21	0.01	0.00	0.02	0.16	0.74	0.03
	0	0.67	0.91	0.29	0.78	0.87	0.13	0.89	0.83	0.26	0.92
	0	0.31	0.06	0.71	0.01	0.11	0.87	0.09	0.01	0.00	0.05

Aims

Optimal stopping stopping problem

$$V(\pi) = \sup_{\tau \in \mathcal{M}} \mathbb{E}[g(X_{\tau \wedge T_N}) \mid \Pi_0 = \pi]$$

$\mathcal{M}$ set of $(\mathcal{F}_t^{\mathcal{O}})$ stopping times

- reformulate the problem
- derive dynamic programming equations
- propose a numerical approximation of the value function
- propose a numerical construction for an ϵ-optimal stopping time

Outline

1. Introduction
   - Piecewise deterministic Markov processes
   - Optimal stopping
   - State of the art
2. Observation process
3. Filtering
4. Dynamic programming
5. Numerical method
   - Quantization
   - Convergence
   - Example

Properties of the observation and filter processes

Property of the filter process

(Π_n), (Π_n, S_n) and (Π_n, S_n, Y_n) are Markov chains

$\mathcal{F}^{\mathcal{O}}$ stopping times

$$\sigma \wedge T_{n+1} = (T_n + R_n) \wedge T_{n+1} \qquad \text{on } \{\sigma \geq T_n\}$$

with R_n $\mathcal{F}^{\mathcal{O}}_{T_n}$-measurable

Reformulated problem

Reformulated optimal stopping problem

- similar to optimal stopping for PDMP
- different because (Π_n, S_n) not underlying Markov chain of some PDMP

Use

- Markov property for (Π_n, S_n)
- Fine structure of $\mathcal{F}^{\mathcal{O}}$ stopping times

Dynamic programming equations

- Initialization $v_N(\pi) = \sum_{i=1}^{q} g(x_i)\pi^i$
- Iteration $v_n(\pi) = L(v_{n+1}, g)(\pi)$, for $n < N$
- $v_0(\pi) = V(\pi)$
- recursive construction of ϵ-optimal stopping time

$$
\begin{aligned}
&L(v, g)(\pi) \\
&= \max_{0\leq m\leq q-1}\left\{\sup_{t_m^*\leq u<t_{m+1}^*}\sum_{i=1}^{q}\mathbb{E}\left[\Pi_n^i h\big(\Phi(x_i, u)\big)\mathbb{1}_{\{S_{n+1}>u\}} + v(\Pi_{n+1})\mathbb{1}_{\{S_{n+1}\leq u\}}|\Pi_n = \pi\right]\right\} \\
&\quad \vee \mathbb{E}[v(\Pi_{n+1})|\Pi_n = \pi]
\end{aligned}
$$

Outline

Quantization

[Pagès 98], [Bally, Pagès 03, 05], [Pagès, Pham, Printems 04]...

Quantization of a random variable $X \in L^p(\mathbb{R}^d)$

Approximate X by $\widehat{X}$ taking finitely many values such that $\|X - \widehat{X}\|_p$ is minimum

- finite weighted grid Γ with $|\Gamma| = K$
- $\widehat{X} = p_\Gamma(X)$ closest neighbour projection

Asymptotic properties

If $E[|X|^{p+\eta}] < +\infty$ for some $\eta > 0$ then

$$\lim_{K\to\infty} \min_{|\Gamma|\le K} \|X - \widehat{X}^\Gamma\|_p \simeq K^d$$

Algorithms

There exist algorithms providing

- grids Γ
- law of $\widehat{X}$
- transition probabilities for quantization of Markov chains

Example: $\mathcal{N}(0, I_2)$

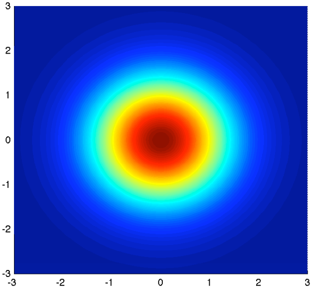

Algorithms

There exist algorithms providing

- grids Γ
- law of $\widehat{X}$
- transition probabilities for quantization of Markov chains

Example: $\mathcal{N}(0, I_2)$

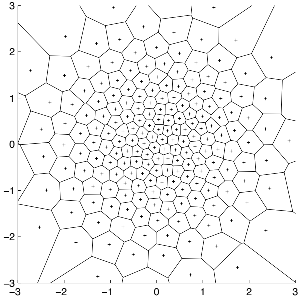

Discretization of the dynamic programming operator

- recursion on functions v_n turns into recursion on random variables $v_n(\Pi_n)$

$$
\begin{aligned}
V_n(\Pi_n) &= L(v_{n+1}, g)(\Pi_n) \\
&= \max_{0\le m\le q-1} \left\{ \sup_{t^*_m \le u < t^*_{m+1}} \sum_{i=1}^{q} \mathbb{E}\Big[\Pi^i_n h\big(\Phi(x_i, u)\big) \mathbb{1}_{\{S_{n+1}>u\}} + v_{n+1}(\Pi_{n+1}) \mathbb{1}_{\{S_{n+1}\le u\}} | \Pi_n\Big] \right\} \\
&\qquad \vee \ \mathbb{E}[v(\Pi_{n+1})|\Pi_n]
\end{aligned}
$$

Discretization of the dynamic programming operator

- recursion on functions v_n turns into recursion on random variables $v_n(\Pi_n)$
- discretize the intervals $[t^*_m; t^*_{m+1}[$ with regular grids G_m

$$
\begin{aligned}
&L^d(v,g)(\Pi_n) \\
&= \max_{0\le m\le q-1}\left\{\max_{u\in G_m}\sum_{i=1}^{q}\mathbb{E}\left[\Pi^i_n h\big(\Phi(x_i,u)\big)\mathbb{1}_{\{S_{n+1}>u\}}+v(\Pi_{n+1})\mathbb{1}_{\{S_{n+1}\le u\}}|\Pi_n\right]\right\} \\
&\qquad \vee\ \mathbb{E}[v(\Pi_{n+1})|\Pi_n]
\end{aligned}
$$

Discretization of the dynamic programming operator

- recursion on functions v_n turns into recursion on random variables $v_n(\Pi_n)$
- discretize the intervals $[t^*_m; t^*_{m+1}[$ with regular grids G_m
- replace (Π_n, S_n) by some quantized approximation

$$\widehat{L}^d(v,g)(\widehat{\Pi}_n)$$
$$= \max_{0\leq m\leq q-1}\left\{\max_{u\in G_m}\sum_{i=1}^{q}\mathbb{E}\left[\widehat{\Pi}^i_n h\big(\Phi(x_i,u)\big)\mathbb{1}_{\{\widehat{S}_{n+1}>u\}} + v(\Pi_{n+1})\mathbb{1}_{\{\widehat{S}_{n+1}\leq u\}}|\widehat{\Pi}_n\right]\right\}$$
$$\vee\ \mathbb{E}[v(\widehat{\Pi}_{n+1})|\widehat{\Pi}_n]$$

Convergence

Theorem

Lipschitz conditions

$$\|\widehat{v}_0(\widehat{\Pi}_0) - V(\Pi_0)\|_p \leq cQE^{1/2}$$

Construction of a computable ϵ stopping time

Numerical results

No theoretical value

Comparison with $\mathbb{E}[\sup X_t] = 0.997$

Number of quantized points	Approx. value function	ϵ-optimal stopping time
50	0.924	0.932
100	0.926	0.938
200	0.931	0.940
500	0.934	0.942

Stopped trajectories

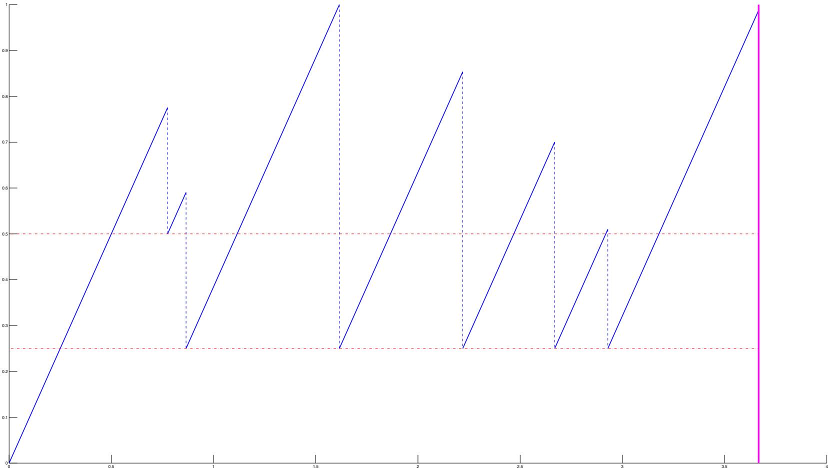

Stopped trajectories

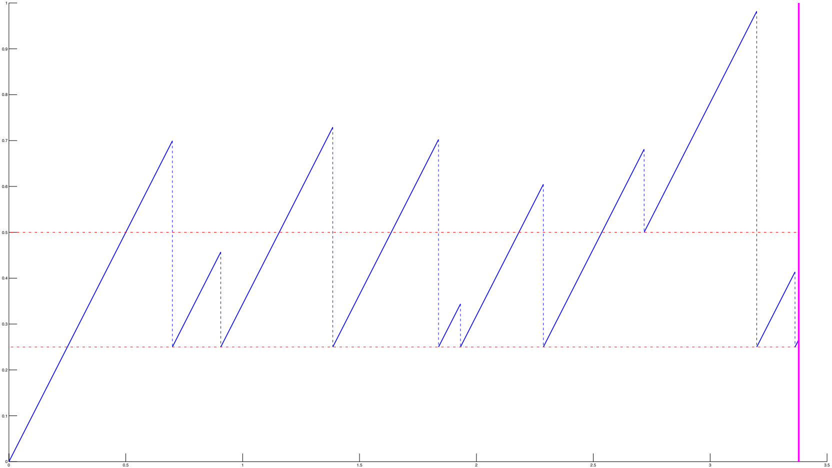

Merci