Optimal stopping of partially observed piecewise deterministic Markov processes

Adrien Brandejsky, Benoîte de Saporta, François Dufour

Institut de Mathématiques de Bordeaux
INRIA Bordeaux Sud-Ouest

Outline

Definition of piecewise deterministic Markov processes

Davis (80's)

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

Applications

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

Dynamics

Hybrid process $X_t = (m_t, y_t)$

- discrete mode $m_t \in \{1, 2, \ldots, p\}$
- 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,
- 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

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

Iterative construction

Starting point

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

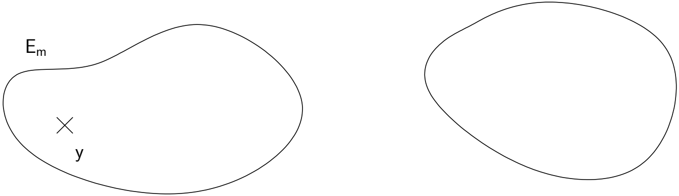

Iterative construction

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

$$\mathbb{P}_{(m,y)}(S_1 > t) = \mathbf{1}_{\{t > t^*(m,y)\}} \exp\Big(- \int_0^t \lambda_m(\phi_m(y,s)ds\Big)$$
$$X_t = (m, \phi_m(y,t)), \quad t < T_1$$

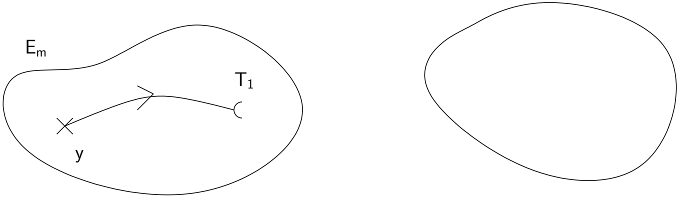

Iterative construction

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

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

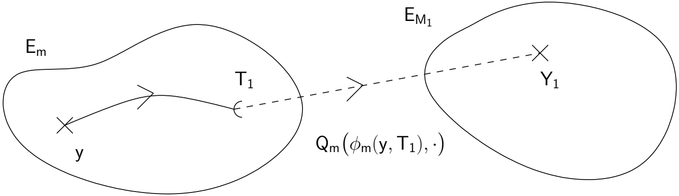

Iterative construction

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

$$\mathbb{P}_{(M_1,Y_1)}(S_2 > t) = \mathbf{1}_{\{t > t^*(m,y)\}} \exp\Big(- \int_0^t \lambda_m(\phi_m(y,s)ds\Big)$$
$$X_{T_1+t} = (M_1, \phi_{M_1}(Y_1, t)), \quad t < S_2$$

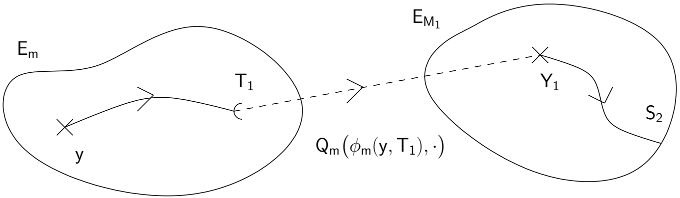

Iterative construction

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

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

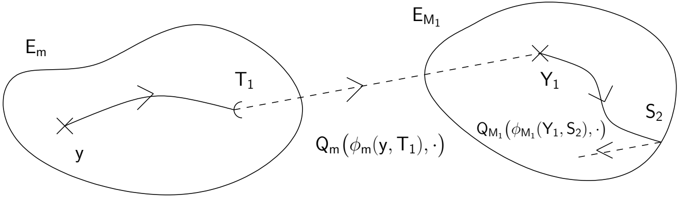

Embedded Markov chain

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

Natural embedded Markov chain

- Z_n new mode and location after n-th jump,
  $S_n = T_n - T_{n-1}$, time between two jumps
- Z_0 starting point, $S_0 = T_0 = 0$

Important property

(Z_n, S_n) is a discrete-time Markov chain
$(Z_n, S_n) \rightsquigarrow$ the distribution of the PDMP.

Optimal stopping problem

Stop the process in order to maximize a reward g

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

- compute the value function $V \rightsquigarrow$ best possible performance
- compute an optimal stopping time $\tau \rightsquigarrow$ control strategy

Classical problem: completely observed

Σ set of stopping times w.r.t. 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 noisy observations of (X_t) are available, Σ set of stopping times w.r.t. the natural filtration $(\mathcal{F}_t^Y)$ of the observation process.

Methodology

- derive the filtering process: $\Pi_t = \mathbb{E}[X_t \mid \mathcal{F}_t^Y]$
- transform the initial problem into a completely observed one: the new state variables $\rightsquigarrow$ the filtering process (Π_t).

Main drawback

- New problem is completely observed but of infinite dimension.

State of the art (numerical approximation)

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

- 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
  - the dynamic of a PDMP can be described the discrete-time Markov chain (Z_n, S_n),
  - stopping times w.r.t. continuous filtration.
- reformulation $\rightsquigarrow$ non standard problem
  - derive new dynamic programming equations
  - operators not Lipschitz

Outline

Observation process

- Marked point process

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

  - S_n is perfectly observed
  - Z_n is observed through a noise:

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

- The filtration of reference is given by $\mathcal{F}_t^Y = \sigma(Y_s, s \leq t)$.

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
   - Approximation
   - Convergence

Filtering process

Hypothesis

Finite number of possible values for Z_n:
$Z_n \in E_0 = \{x_1, \ldots, x_q\} \quad \longleftrightarrow \quad Q(x, E_0) = 1$ for any $x \in E$
with $t_i^* = t^*(x_i)$, $t_1^* \leq t_2^* \cdots \leq t_q^*$

Definition

Filtering process $\Pi_n = (\Pi_n^1, \ldots, \Pi_n^q)$ where

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

We have a recursive construction

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

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
   - Approximation
   - Convergence

Reformulated problem

Reformulation of the partially observed optimal stopping problem

- Antagonistic features : the underlying process is a discrete-time process (Z_n, S_n) and the stopping strategy is continuous.
- The problem is similar to optimal stopping problem for PDMP but different because (Π_n, S_n) is not the underlying Markov chain of some PDMP

Use

- Markov property for (Π_n, S_n)
- Special properties of the stopping times w.r.t. $\{\mathcal{F}_t^Y\}$

Dynamic programming equation

Some operators

$$
\begin{aligned}
Gv(\pi, u) &= \mathbf{E}[v(\Pi_1)\mathbb{1}_{\{S_1 \leq u\}} | \Pi_0 = \pi], \\
Hh(\pi, u) &= \sum_{i=1}^{q} h \circ \Phi(x_i, u)\pi^i \mathbb{1}_{\{u < t_i^*\}} e^{-\Lambda(x_i, u)}, \\
J(v, h)(\pi, u) &= Hh(\pi, u) + Gv(\pi, u), \\
L(v, h)(\pi) &= \sup_{u \geq 0} J(v, h)(\pi, u).
\end{aligned}
$$

where $v \in B(\mathcal{M}_1(E_0))$ and $h \in B(E)$.

Dynamic programming equation

The Bellman equation

The sequence $(v_n)_{0 \leq n \leq N}$ of real-valued functions is defined on $\mathcal{M}_1(E_0)$ by

$$\begin{cases} v_N(\pi) &= \sum_{i=1}^{q} g(x_i)\pi^i, \\ v_{n-1}(\pi) &= L(v_n, g)(\pi), \quad 1 \leq n \leq N. \end{cases}$$

Theorem

For all $1 \leq n \leq N$ and $\pi \in \mathcal{M}_1(E_0)$, one has

$$\sup_{\sigma \in \Sigma_n^Y} \mathbf{E}[g(X_\sigma)|\Pi_0 = \pi] = v_{N-n}(\pi),$$

with $\Sigma_n^Y = \left\{\sigma \in \Sigma^Y \text{ such that } \sigma \leq T_n \text{ a.s.}\right\}$.

Dynamic programming equation

Reformulation of the cost function in terms of (Π_n, S_n)

Proposition:

Let $\sigma \in \Sigma^Y$ and $n \geq 1$. For all $\pi \in \mathcal{M}_1(E_0)$ one has

$$\begin{aligned} &\mathbf{E}[g(X_{\sigma \wedge T_n}) | \Pi_0 = \pi] \\ &= \sum_{k=0}^{n-1} \sum_{i=1}^{q} \mathbf{E}[\mathbb{1}_{\{T_k \leq \sigma\}} \mathbb{1}_{\{R_k < t_i^*\}} g \circ \Phi(x_i, R_k) e^{-\Lambda(x_i, R_k)} \Pi_k^i | \Pi_0 = \pi] \\ &\quad + \sum_{i=1}^{q} \mathbf{E}[\mathbb{1}_{\{T_n \leq \sigma\}} g(x_i) \Pi_n^i | \Pi_0 = \pi], \end{aligned}$$

where $(R_k)_{k \in \mathbb{N}}$ is the sequence of non negative random variables associated to σ satisfying $R_k \in \mathcal{F}_{T_k}^Y$.

Dynamic programming equation

Proof:

We split $\mathbf{E}[g(X_{\sigma \wedge T_n})|\Pi_0 = \pi]$ into several terms depending on the position of σ w.r.t. the jump times T_k

$$\begin{aligned}
&\mathbf{E}[g(X_{\sigma \wedge T_n})\big|\Pi_0 = \pi] \\
&= \sum_{k=0}^{n-1} \sum_{i=1}^{q} \mathbf{E}[\mathbb{1}_{\{T_k \leq \sigma < T_{k+1}\}} \mathbb{1}_{\{Z_k = x_i\}} g \circ \Phi(x_i, R_k)|\Pi_0 = \pi] \\
&+ \sum_{i=1}^{q} \mathbf{E}[\mathbb{1}_{\{T_n \leq \sigma\}} \mathbb{1}_{\{Z_n = x_i\}} g(x_i)|\Pi_0 = \pi].
\end{aligned}$$

Dynamic programming equation

Proof:

We split $\mathbf{E}[g(X_{\sigma\wedge T_n})|\Pi_0 = \pi]$ into several terms depending on the position of σ w.r.t. the jump times T_k

$$\mathbf{E}[g(X_{\sigma\wedge T_n})\big|\Pi_0 = \pi]$$
$$= \sum_{k=0}^{n-1}\sum_{i=1}^{q} \mathbf{E}[\mathbb{1}_{\{T_k\leq\sigma<T_{k+1}\}}\mathbb{1}_{\{Z_k=x_i\}}g\circ\Phi(x_i, R_k)|\Pi_0 = \pi]$$
$$+\sum_{i=1}^{q}\mathbf{E}[\mathbb{1}_{\{T_n\leq\sigma\}}\Pi_n^i g(x_i)|\Pi_0 = \pi].$$

Dynamic programming equation

Proof:

We split $\mathbf{E}[g(X_{\sigma\wedge T_n})|\Pi_0=\pi]$ into several terms depending on the position of σ w.r.t. the jump times T_k

$$
\begin{aligned}
&\mathbf{E}[g(X_{\sigma\wedge T_n})\big|\Pi_0=\pi] \\
&=\sum_{k=0}^{n-1}\sum_{i=1}^{q}\mathbf{E}[\mathbb{1}_{\{T_k\le\sigma\}}\mathbb{1}_{\{Z_k=x_i\}}g\circ\Phi(x_i,R_k)\mathbf{E}[\mathbb{1}_{\{S_{k+1}>R_k\}}|\mathfrak{F}_{T_k}]|\Pi_0=\pi] \\
&+\sum_{i=1}^{q}\mathbf{E}[g(x_i)\mathbb{1}_{\{T_n\le\sigma\}}\Pi_n^i|\Pi_0=\pi].
\end{aligned}
$$

Dynamic programming equation

Proof:

We split $\mathbf{E}[g(X_{\sigma\wedge T_n})|\Pi_0=\pi]$ into several terms depending on the position of σ w.r.t. the jump times T_k

$$\begin{aligned}
&\mathbf{E}[g(X_{\sigma\wedge T_n})\big|\Pi_0=\pi]\\
&=\sum_{k=0}^{n-1}\sum_{i=1}^{q}\mathbf{E}[\mathbb{1}_{\{T_k\leq\sigma\}}\mathbb{1}_{\{Z_k=x_i\}}g\circ\Phi(x_i,R_k)\mathbb{1}_{\{R_k<t^*(Z_k)\}}e^{-\Lambda(Z_k,R_k)}|\Pi_0=\pi]\\
&+\sum_{i=1}^{q}\mathbf{E}[g(x_i)\mathbb{1}_{\{T_n\leq\sigma\}}\Pi_n^i|\Pi_0=\pi].
\end{aligned}$$

Dynamic programming equation

Proof:

We split $\mathbf{E}[g(X_{\sigma\wedge T_n})|\Pi_0=\pi]$ into several terms depending on the position of σ w.r.t. the jump times T_k

$$\begin{aligned}
&\mathbf{E}[g(X_{\sigma\wedge T_n})\big|\Pi_0=\pi]\\
&=\sum_{k=0}^{n-1}\sum_{i=1}^{q}\mathbf{E}[\mathbb{1}_{\{T_k\leq\sigma\}}\Pi_k^i g\circ\Phi(x_i,R_k)\mathbb{1}_{\{R_k<t_i^*\}}e^{-\Lambda(x_i,R_k)}|\Pi_0=\pi]\\
&+\sum_{i=1}^{q}\mathbf{E}[g(x_i)\mathbb{1}_{\{T_n\leq\sigma\}}\Pi_n^i|\Pi_0=\pi].
\end{aligned}$$

Dynamic programming equation

The following Proposition proves that v_{N-n} is an upper bound for the value function of the problem with horizon T_n.

Proposition

For all $1 \leq n \leq N$ and $\pi \in \mathcal{M}_1(E_0)$, one has

$$\sup_{\sigma \in \Sigma_n^Y} \mathbf{E}[g(X_\sigma)|\Pi_0 = \pi] \leq v_{N-n}(\pi).$$

Proof:

We prove this result by induction on n.

Dynamic programming equation

n=1 The previous result yields

$$
\begin{aligned}
&\mathbf{E}[g(X_{\sigma \wedge T_1})|\Pi_0 = \pi] \\
&= \sum_{i=1}^{q} \mathbf{E}[\mathbb{1}_{\{R_0 < t_i^*\}} g \circ \Phi(x_i, R_0) e^{-\Lambda(x_i, R_0)} \Pi_0^i | \Pi_0 = \pi] \\
&+ \sum_{i=1}^{q} \mathbf{E}[\mathbb{1}_{\{T_1 \leq \sigma\}} g(x_i) \Pi_1^i | \Pi_0 = \pi].
\end{aligned}
$$

we recognize that the first term of the right hand side is $Hg(\pi, R_0)$ and the second term is $Gv_N(\pi, R_0)$.

Dynamic programming equation

and so,

$$\begin{aligned}\mathbf{E}[g(X_{\sigma\wedge T_1})|\Pi_0=\pi] &= J(v_N,g)(\pi,R_0) &\le& \sup_{u\ge 0} J(v_N,g)(\pi,u)\\ && =& L(v_N,g)(\pi)=v_{N-1}(\pi).\end{aligned}$$

$\boxed{n\ge 2}$ Assume $\mathbf{E}[g(X_\tau)|\Pi_0=\pi]\le v_{N-(n-1)}(\pi)$, for all $\tau\in\Sigma^Y_{n-1}$.
As in the case $n=1$,

$$\mathbf{E}[g(X_{\sigma\wedge T_n})|\Pi_0=\pi]=Hg(\pi,R_0)+\mathbf{E}[\Xi\mathbb{1}_{\{S_1\le R_0\}}|\Pi_0=\pi].$$

where

$$\begin{aligned}\Xi &= \mathbf{E}\Big[\sum_{k=1}^{n-1}\sum_{i=1}^{q}\mathbb{1}_{\{T_k\le\sigma\}}\mathbb{1}_{\{R_k<t_i^*\}}g\circ\Phi(x_i,R_k)e^{-\Lambda(x_i,R_k)}\Pi_k^i\\ &\quad+\sum_{i=1}^{q}\mathbb{1}_{\{T_n\le\sigma\}}g(x_i)\Pi_n^i\Big|\mathfrak{F}^Y_{T_1}\Big].\end{aligned}$$

Dynamic programming equation

$$\Xi = \mathbf{E}\Big[\sum_{k=1}^{n-1} \sum_{i=1}^{q} \mathbb{1}_{\{T_k \leq \sigma\}} \mathbb{1}_{\{R_k < t_i^*\}} g \circ \Phi(x_i, R_k) e^{-\Lambda(x_i, R_k)} \Pi_k^i + \sum_{i=1}^{q} \mathbb{1}_{\{T_n \leq \sigma\}} g(x_i) \Pi_n^i \Big| \mathfrak{F}_{T_1}^Y \Big].$$

Markov property: $\Pi_k = \Pi_{k-1} \circ \theta$, $T_k = T_1 + T_{k-1} \circ \theta$ where θ translation operator of the $(\mathfrak{F}_{T_n}^Y)_{n\in\mathbb{N}}$-Markov chain $(\Pi_n, Y_n, S_n)_{n\in\mathbb{N}}$.
Moreover, on $\{T_1 \leq \sigma\}$, $\sigma = T_1 + \widetilde{\sigma} \circ \theta$ and $R_k = \widetilde{R}_{k-1} \circ \theta$.

$$\rightsquigarrow \Xi = w(\Pi_1) \text{ with } w(\pi) = \mathbf{E}[g(X_{\widetilde{\sigma} \wedge T_{n-1}}) | \Pi_0 = \pi].$$

Dynamic programming equation

The induction hypothesis gives

$$\Xi \leq v_{N-(n-1)}(\Pi_1),$$

and so

$$\mathbf{E}[g(X_{\sigma \wedge T_n})|\Pi_0 = \pi] = Hg(\pi, R_0) + \mathbf{E}[\Xi \mathbb{1}_{\{S_1 \leq R_0\}}|\Pi_0 = \pi].$$

Dynamic programming equation

The induction hypothesis gives

$$\Xi \leq v_{N-(n-1)}(\Pi_1),$$

and so

$$\mathbf{E}[g(X_{\sigma \wedge T_n})|\Pi_0 = \pi] \leq Hg(\pi, R_0) + \mathbf{E}[v_{N-(n-1)}(\Pi_1)\mathbb{1}_{\{S_1 \leq R_0\}}|\Pi_0 = \pi].$$

Dynamic programming equation

The induction hypothesis gives

$$\Xi \le v_{N-(n-1)}(\Pi_1),$$

and so

$$\mathbf{E}[g(X_{\sigma\wedge T_n})|\Pi_0=\pi] \le Hg(\pi,R_0)+\mathbf{E}[v_{N-(n-1)}(\Pi_1)\mathbb{1}_{\{S_1\le R_0\}}|\Pi_0=\pi].$$

In the second term, we recognize the operator G and one has

$$\begin{aligned}\mathbf{E}[g(X_{\sigma\wedge T_n})|\Pi_0=\pi] &\le Hg(\pi,R_0)+Gv_{N-(n-1)}(\pi,R_0)\\ &= J(v_{N-(n-1)},g)(\pi,R_0)\\ &\le \sup_{u\ge 0} J(v_{N-(n-1)},g)(\pi,u)\\ &= L(v_{N-(n-1)},g)(\pi) = v_{N-n}(\pi),\end{aligned}$$

that proves the induction. □

Dynamic programming equation

We now prove the reverse inequality by constructing a sequence of ϵ-optimal stopping times.

Definition

For $\epsilon > 0$, $1 \leq n \leq N$ and for $\pi \in \mathcal{M}_1(E_0)$, we define

$$r_n^\epsilon(\pi) = \inf \{ u > 0 : \ J(v_{N-n}, g)(\pi, u) > v_{N-n-1}(\pi) - \epsilon \} .$$

Consider $R_{1,0}^\epsilon = r_0^\epsilon(\Pi_0)$ and for $2 \leq n \leq N$,

$$\begin{cases} R_{n,0}^\epsilon & = r_{n-1}^{\epsilon/2}(\Pi_0), \\ R_{n,k}^\epsilon & = r_{n-1-k}^{\epsilon/(2^{k+1})}(\Pi_k) \mathbb{1}_{\{R_{n,k-1}^\epsilon \geq S_k\}} \text{ for } 1 \leq k \leq n-1, \end{cases}$$

and finally set

$$U_n^\epsilon = \sum_{k=1}^{n} R_{n,k-1}^\epsilon \wedge S_k.$$

Dynamic programming equation

The following lemma describes the effect of the translation operator θ on the sequence $(R^{\epsilon}_{n,k})_{1\leq n\leq N, 0\leq k\leq n-1}$.

Lemma

For $n \geq 2$ and $1 \leq k \leq n-1$, on the set $\{T_1 \leq U^{2\epsilon}_n\}$, one has

$$R^{\epsilon}_{n-1,k-1} \circ \theta = R^{2\epsilon}_{n,k}.$$

Equipped with this preliminary result, we may now prove that $(U^{\epsilon}_n)_{1\leq n\leq N}$ is a sequence of ϵ-optimal stopping times

Theorem

For all $1 \leq n \leq N$ and $\epsilon > 0$, one has $U^{\epsilon}_n \in \Sigma^Y_n$ and

$$\mathbf{E}[g(X_{U^{\epsilon}_n})|\Pi_0 = \pi] \geq v_{N-n}(\pi) - \epsilon.$$

Dynamic programming equation

Proof:

We prove this result by induction on n.

$\boxed{\text{n=1}}$ By using the same arguments as previously

$$\mathbf{E}[g(X_{R^{\epsilon}_{1,0}\wedge S_1})|\Pi_0 = \pi] = Hg(\pi, r^{\epsilon}_0) + Gv_N(\pi, r^{\epsilon}_0) = J(v_N, g)(\pi, r^{\epsilon}_0).$$

The definition of r^{ϵ}_0 yields $J(v_N, g)(\pi, r^{\epsilon}_0) \geq v_{N-1}(\pi) - \epsilon$ thus one has

$$\mathbf{E}[g(X_{R^{\epsilon}_{1,0}\wedge S_1})|\Pi_0 = \pi] \geq v_{N-1}(\pi) - \epsilon.$$

Dynamic programming equation

$\boxed{2 \leq n \leq N}$ Assume $\mathbf{E}[g(X_{U^{\epsilon}_{n-1}})|\Pi_0 = \pi] \geq v_{N-(n-1)}(\pi) - \epsilon$, for $\epsilon > 0$.
The first proposition yields

$$\mathbf{E}[g(X_{U_n^{2\epsilon}})|\Pi_0 = \pi]$$
$$= \sum_{k=0}^{n-1} \sum_{i=1}^{q} \mathbf{E}\left[\mathbb{1}_{\{T_k \leq U_n^{2\epsilon}\}} \mathbb{1}_{\{R_{n,k}^{2\epsilon} < t_i^*\}} g \circ \Phi(x_i, R_{n,k}^{2\epsilon}) e^{-\Lambda(x_i, R_{n,k}^{2\epsilon})} \Pi_k^i \middle| \Pi_0 = \pi\right]$$
$$+ \sum_{i=1}^{q} \mathbf{E}[\mathbb{1}_{\{T_n \leq U_n^{2\epsilon}\}} g(x_i) \Pi_n^i | \Pi_0 = \pi].$$

The term for $k = 0$ equals $Hg(\pi, r_{n-1}^{\epsilon})$ since $R_{n,0}^{2\epsilon} = r_{n-1}^{\epsilon}(\Pi_0)$.
Taking the conditional expectation w.r.t. $\mathfrak{F}_{T_1}^{Y}$ in the other terms gives

Dynamic programming equation

$$\mathbf{E}[g(X_{U_n^{2\epsilon}})|\Pi_0=\pi] = Hg(\pi, r_{n-1}^{\epsilon}) + \mathbf{E}[\Xi' \mathbb{1}_{\{T_1 \le U_n^{2\epsilon}\}}|\Pi_0=\pi],$$

with

$$\begin{aligned}\Xi' &= \mathbf{E}\Big[\sum_{k=1}^{n-1}\sum_{i=1}^{q} \mathbb{1}_{\{T_k \le U_n^{2\epsilon}\}} \mathbb{1}_{\{R_{n,k}^{2\epsilon} < t_i^*\}} g \circ \Phi(x_i, R_{n,k}^{2\epsilon}) e^{-\Lambda(x_i, R_{n,k}^{2\epsilon})} \Pi_k^i \\ &\quad + \sum_{i=1}^{q} \mathbb{1}_{\{T_n \le U_n^{2\epsilon}\}} g(x_i) \Pi_n^i \big| \mathfrak{F}_{T_1}^{Y}\Big].\end{aligned}$$

Dynamic programming equation

Markov property in the term Ξ'. One has for $n \geq 2$ and $1 \leq k \leq n-1$ on the event $\{T_1 \leq U_n^{2\epsilon}\} = \{S_1 \leq R_{n,0}^{2\epsilon}\}$:

$$R_{n-1,k-1}^{\epsilon} \circ \theta = R_{n,k}^{2\epsilon}.$$

Thus, on $\{T_1 \leq U_n^{2\epsilon}\}$

$$\begin{aligned} U_n^{2\epsilon} &= S_1 + \sum_{k=2}^{n} R_{n,k-1}^{2\epsilon} \wedge S_k = T_1 + \sum_{k=2}^{n} (R_{n-1,k-2}^{\epsilon} \circ \theta) \wedge (S_{k-1} \circ \theta) \\ &= T_1 + U_{n-1}^{\epsilon} \circ \theta. \end{aligned}$$

Dynamic programming equation

Consequently, $\mathbb{1}_{\{T_k \le U_n^{2\epsilon}\}} = \mathbb{1}_{\{T_{k-1} \le U_{n-1}^{\epsilon}\}} \circ \theta$ and thus, for

$$\begin{aligned} \Xi' &= \mathbf{E}\Big[\sum_{k=1}^{n-1} \sum_{i=1}^{q} \mathbb{1}_{\{T_k \le U_n^{2\epsilon}\}} \mathbb{1}_{\{R_{n,k}^{2\epsilon} < t_i^*\}} g \circ \Phi(x_i, R_{n,k}^{2\epsilon}) e^{-\Lambda(x_i, R_{n,k}^{2\epsilon})} \Pi_k^i \\ &\qquad + \sum_{i=1}^{q} \mathbb{1}_{\{T_n \le U_n^{2\epsilon}\}} g(x_i) \Pi_n^i \Big| \mathfrak{F}_{T_1}^Y \Big]. \end{aligned}$$

we have

$$\Xi'(\Pi_1) = w'(\Pi_1) \text{ with } w'(\pi) = \mathbf{E}\left[g(X_{U_{n-1}^{\epsilon}}) \middle| \Pi_0 = \pi \right].$$

Dynamic programming equation

Moreover, thanks to the induction assumption, one has $w'(\pi) \geq v_{N-(n-1)}(\pi) - \epsilon$ so that one obtains

$$\Xi' \geq v_{N-(n-1)}(\Pi_1) - \epsilon,$$

and so, one obtains

$$\begin{aligned}
&\mathbf{E}[g(X_{U_n^{2\epsilon}})|\Pi_0 = \pi] \\
&\qquad \geq Hg(\pi, r_{n-1}^{\epsilon}) + \mathbf{E}[v_{N-(n-1)}(\Pi_1)\mathbb{1}_{\{S_1 \leq r_{n-1}^{\epsilon}\}}|\Pi_0 = \pi] - \epsilon \\
&\qquad = J(v_{N-(n-1)}, g)(\pi, r_{n-1}^{\epsilon}) - \epsilon \\
&\qquad \geq v_{N-n}(\pi) - 2\epsilon,
\end{aligned}$$

from the definition of r_{n-1}^{ϵ}, showing the result. $\square$

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
   - Approximation
   - Convergence

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

$$\min_{|\Gamma| \leq K} \|X - \widehat{X}^\Gamma\|_p \underset{K\to\infty}{=} O(K^{1/d})$$

Quantization

There exist algorithms providing

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

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

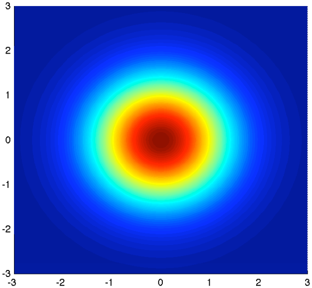

Quantization

There exist algorithms providing

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

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

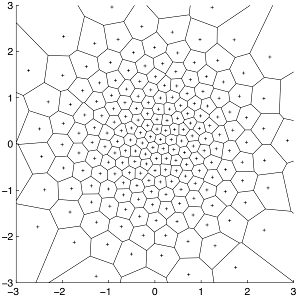

Dynamic programming equation

The Bellman equation

The sequence $(v_n)_{0 \leq n \leq N}$ of real-valued functions is defined on $\mathcal{M}_1(E_0)$ by

$$\begin{cases} v_N(\pi) & = \sum_{i=1}^{q} g(x_i)\pi^i, \\ v_{n-1}(\pi) & = L(v_n, g)(\pi), \quad 1 \leq n \leq N. \end{cases}$$

Theorem

For any $\pi \in \mathcal{M}_1(E_0)$, one has

$$\sup_{\sigma \in \Sigma_N^Y} \mathbf{E}[g(X_\sigma) | \Pi_0 = \pi] = v_0(\pi).$$

Approximation of the Bellman equation

The Bellman equation: $\sup_{\sigma\in\Sigma_N^Y} \mathbf{E}[g(X_\sigma)|\Pi_0=\pi] = v_0(\Pi_0)$.

$$\begin{cases} v_N(\Pi_N) &= g(\Pi_N), \\ v_{n-1}(\Pi_{n-1}) &= L(v_n, g)(\Pi_n), \quad 1 \le n \le N. \end{cases}$$

- recursion for the random variables $v_n(\Pi_n)$

$$\begin{aligned} v_n(\Pi_n) &= L(v_{n+1}, g)(\Pi_n) \\ &= \sup_u \left\{ \sum_{i=1}^q \Pi_n^i h\big(\Phi(x_i,u)\big) \mathbb{1}_{\{u<t_i^*\}} e^{-\Lambda(x_i,u)} + \mathbb{E}\Big[v_{n+1}(\Pi_{n+1}) \mathbb{1}_{\{S_{n+1}\le u\}} |\Pi_n\Big] \right\} \\ &\quad \vee \mathbb{E}[v_{n+1}(\Pi_{n+1})|\Pi_n] \end{aligned}$$

Approximation of the Bellman equation

The Bellman equation: $\sup_{\sigma \in \Sigma_N^Y} \mathbf{E}[g(X_\sigma)|\Pi_0 = \pi] = v_0(\Pi_0)$.

$$\left\{ \begin{array}{rl} v_N(\Pi_N) & = g(\Pi_N), \\ v_{n-1}(\Pi_{n-1}) & = L(v_n, g)(\Pi_n), \quad 1 \leq n \leq N. \end{array} \right.$$

- recursion for the random variables $v_n(\Pi_n)$
- discretize the intervals $]t_m^*; t_{m+1}^*[$ with regular grids G_m

$$L^d(v, g)(\Pi_n) = \max_{u \in G} \left\{ \sum_{i=1}^{q} \Pi_n^i h\big(\Phi(x_i, u)\big) \mathbb{1}_{\{u < t_i^*\}} e^{-\Lambda(x_i, u)} + \mathbb{E}\left[v(\Pi_{n+1}) \mathbb{1}_{\{S_{n+1} \leq u\}} | \Pi_n \right] \right\} \vee \mathbb{E}[v(\Pi_{n+1})|\Pi_n]$$

Approximation of the Bellman equation

The approximation of $v_0(\Pi_0)$: $\widehat{v}_0(\widehat{\Pi}_0)$.

$$\begin{cases} \widehat{v}_N(\widehat{\Pi}_N) &= g(\widehat{\Pi}_N), \\ \widehat{v}_{n-1}(\widehat{\Pi}_{n-1}) &= L(\widehat{v}_n, g)(\widehat{\Pi}_{n-1}), \quad 1 \leq n \leq N. \end{cases}$$

- recursion for the 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 its quantized approximation $(\widehat{\Pi}_n, \widehat{S}_n)$

$$\widehat{L}^d(v, g)(\widehat{\Pi}_n)$$
$$= \max_{u \in G} \left\{ \sum_{i=1}^{q} \widehat{\Pi}^i_n h\big(\Phi(x_i, u)\big) \mathbb{1}_{\{u < t^*_i\}} e^{-\Lambda(x_i, u)} + \mathbb{E}\left[v(\widehat{\Pi}_{n+1}) \mathbb{1}_{\{\widehat{S}_{n+1} \leq u\}} | \widehat{\Pi}_n \right] \right\}$$
$$\vee \, \mathbb{E}[v(\widehat{\Pi}_{n+1}) | \widehat{\Pi}_n]$$

Convergence

Theorem

Lipschitz conditions

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

Construction of a computable ϵ stopping time