Numerical method for optimal stopping of piecewise deterministic Markov processes

Benoîte de Saporta François Dufour

INRIA Bordeaux Sud-Ouest
University of Bordeaux
funded by ANR-09-SEGI-004

July 15 2010 – Liverpool

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,. . .

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, 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 when the process starts in mode m at y:

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

- law of the first jump time T_1 starting from (m, y)

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

Remark

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

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

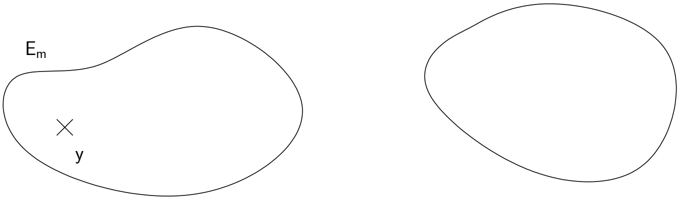

Iterative construction

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

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

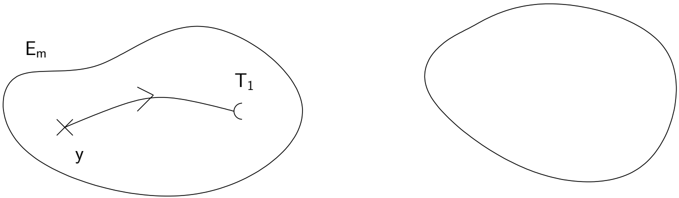

Iterative construction

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

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

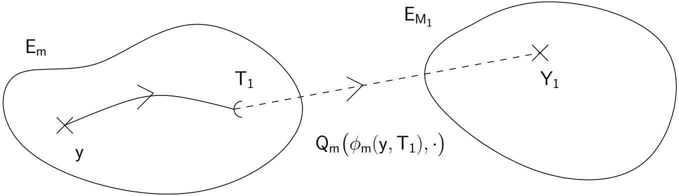

Iterative construction

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

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

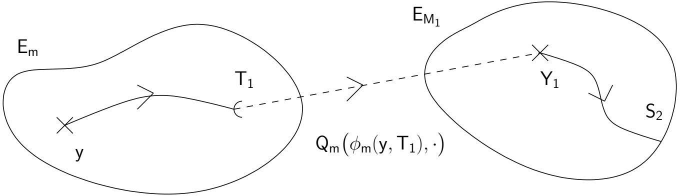

Iterative construction

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

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

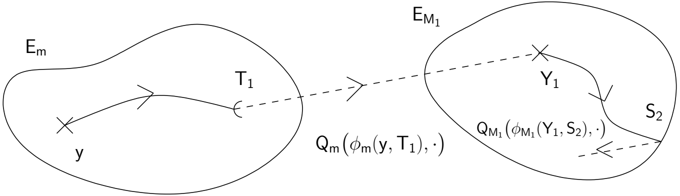

Embedded Markov chain

$\{X_t\}$ 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

Proposition

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

Industrial example

Industrial collaboration: EADS Astrium

- crack propagation,
- material subject to corrosion and randomly exposed to different stressing ambiences.

Material subject to corrosion

Model

- $m_t \in \{1, 2, 3\}$ describes the ambience
- d_t loss of thickness
- γ_t duration of the initial protection against corrosion
- ρ_t rate of corrosion

The process starts in ambience 1: $m_0 = 1$, $d_0 = 0mm$,

$$\gamma_0 \sim 1 - \exp(-[t/\alpha]^\gamma), \qquad \rho_0 \sim \mathcal{U}[\rho_1^1, \rho_1^2]$$

Dynamics

- Deterministic succession of ambiences: $1 \to 2 \to 3 \to 1 \cdots$
- Time spent in ambience $i \sim Exp(\lambda_i)$
- ρ_t is constant on $[T_n, T_{n+1}[$ and in ambience i, $\rho \sim \mathcal{U}[\rho_i^1, \rho_i^2]$
- $\gamma_{T_{n+1}} = 0$ if $\gamma_{T_n} < S_{n+1}$, $\quad \gamma_{T_{n+1}} = S_{n+1} - \gamma_{T_n}$ otherwise
- On $[T_n, T_{n+1}[$, in ambience i

$$d_t = \begin{cases} 0, \text{ if } t \leq \gamma_{T_n}, \\ \rho_{T_n}\Big(t - (\gamma_{T_n} + c_i) + c_i exp(-[t - \gamma_{T_i}]/c_i)\Big), \text{ otherwise.} \end{cases}$$

The material is inefficient when the thickness loss is greater than $0.2mm$

Trajectories

Examples of trajectories for the loss of thickness d_t

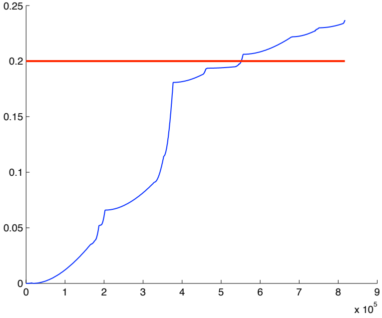

Trajectories

Examples of trajectories for the loss of thickness d_t

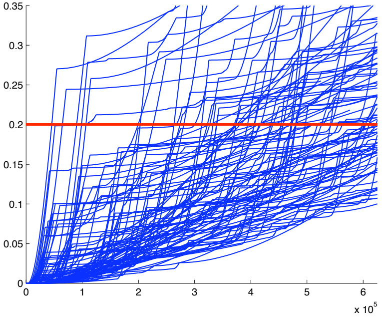

Definition

- Reward function g
- Time horizon N-th jump T_N
- $\mathcal{M}_N$ set of all stopping times $\tau \leq T_N$

Optimal stopping problem

- compute the value function

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

- find an (ε-)optimal stopping time τ^* that reaches $V(x)(-\varepsilon)$

References for PDMP's

- O. Costa & M. Davis (88)
- O. Costa & F. Dufour (00)
- U. Gugerli (86)
- D. Gatarek (91)
- S. Lenhart & Y. Liao (85)

Application to maintenance optimization

- $X_t = (m_t, y_t)$ state of a machine/structure at time t
- T_n failure of some components/changes of ambience

Optimal stopping

Find an optimal balance between

- changing the components too early/often
- no maintenance leading to a total breakdown

Our aim

Objective

Propose a numerical method

- to evaluate the value function
- to compute an ε-optimal stopping rule

with error bounds

Strategy

Adapt numerical procedures for optimal stopping of SDE's

Numerical method for diffusion processes

Bally, Pagès, Pham, Printems 98–05

Y_t continuous-time Markov diffusion process

1. time discretization (Euler scheme) : $Y_k = Y_{k\Delta t}$ discrete-time Markov chain with continuous state space
2. quantization : replace Y_k by a random variable $\widehat{Y}_k$ taking values in a finite state space
3. replace the conditional expectations in the dynamic programming equation by finite sums

Assumptions + Lipschitz-continuous reward function $\Longrightarrow$ convergence rate of the approximated value function to the original one

Numerical method for diffusion processes

Bally, Pagès, Pham, Printems 98–05

Y_t continuous-time Markov diffusion process

1. time discretization (Euler scheme) : $Y_k = Y_{k\Delta t}$ discrete-time Markov chain with continuous state space
2. quantization : replace Y_k by a random variable $\widehat{Y}_k$ taking values in a finite state space
3. replace the conditional expectations in the dynamic programming equation by finite sums

Assumptions + Lipschitz-continuous reward function $\Longrightarrow$ convergence rate of the approximated value function to the original one

Numerical method for diffusion processes

Bally, Pagès, Pham, Printems 98–05

Y_t continuous-time Markov diffusion process

1. time discretization (Euler scheme) : $Y_k = Y_{k\Delta t}$ discrete-time Markov chain with continuous state space
2. quantization : replace Y_k by a random variable $\widehat{Y}_k$ taking values in a finite state space
3. replace the conditional expectations in the dynamic programming equation by finite sums

Assumptions + Lipschitz-continuous reward function $\Longrightarrow$ convergence rate of the approximated value function to the original one

Numerical method for diffusion processes

Bally, Pagès, Pham, Printems 98–05

Y_t continuous-time Markov diffusion process

1. time discretization (Euler scheme) : $Y_k = Y_{k\Delta t}$ discrete-time Markov chain with continuous state space
2. quantization : replace Y_k by a random variable $\widehat{Y}_k$ taking values in a finite state space
3. replace the conditional expectations in the dynamic programming equation by finite sums

Assumptions + Lipschitz-continuous reward function $\Longrightarrow$ convergence rate of the approximated value function to the original one

Numerical method for diffusion processes

Bally, Pagès, Pham, Printems 98–05

Y_t continuous-time Markov diffusion process

1. time discretization (Euler scheme) : $Y_k = Y_{k\Delta t}$ discrete-time Markov chain with continuous state space
2. quantization : replace Y_k by a random variable $\widehat{Y}_k$ taking values in a finite state space
3. replace the conditional expectations in the dynamic programming equation by finite sums

Assumptions + Lipschitz-continuous reward function $\Longrightarrow$ convergence rate of the approximated value function to the original one

Quantization

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

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

Asymptotic properties

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

$$\lim_{K\to\infty} K^{p/d} \min_{|\Gamma|\leq K} \|X - \widehat{X}^\Gamma\|_p^p = J_{p,d}\Big(\int |h|^{d/(d+p)}(u)du\Big),$$

where $P_x(du) = h(u)\lambda_d(du) + \nu$ with $\nu \perp \lambda_d$ and $J_{p,d}$ a constant.

Algorithms

There exist algorithms providing

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

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

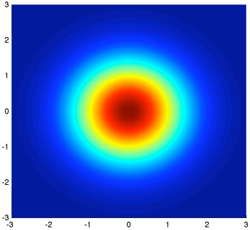

Algorithms

There exist algorithms providing

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

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

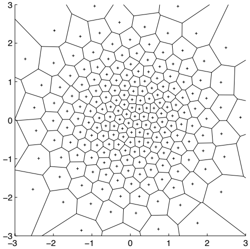

Algorithms

There exist algorithms providing

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

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

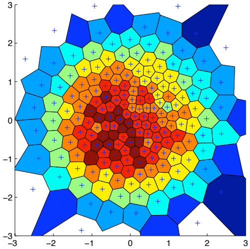

Some references

Quantization technique have been developed recently in

- numerical probability Pages (98)
- nonlinear filtering Pages & Pham (05)
- optimal stochastic control in finance Bally& Pages (03); Pages & Pham &Printemps (05); Bally & Pages & Printemps (05)

Specificities of PDMP's

- jumps at random times
- indicator functions in the dynamic programming equation

Solution

- use the embedded Markov chain (Z_n, S_n)
- be careful with the time grids

Iterative resolution

Backward dynamic programming equation (U. Gugerli, 1986):

- $v_N = g$
- $v_n = L(v_{n+1}, g)$ for $n \leq N-1$

$$v_0(x) = \sup_{\tau \in \mathcal{M}_N} \mathbb{E}_x[g(X_\tau)] = V(x)$$

$$\begin{aligned} L(w,g)(x) \\ = \quad & \sup_{t \geq 0} \Big[\int_0^{t \wedge t^*(x)} \lambda Q w(\phi(x,s)) \mathrm{e}^{-\Lambda(x,s)} ds + g\big(\phi(x, t \wedge t^*(x))\big) \mathrm{e}^{-\Lambda(x, t \wedge t^*(x))} \Big] \\ & \vee \int_0^{t^*(x)} \lambda Q w(\phi(x,s)) \mathrm{e}^{-\Lambda(x,s)} ds + Q w(\phi(x, t^*(x))) \mathrm{e}^{-\Lambda(x, t^*(x))}. \end{aligned}$$

Probabilistic interpretation

Backward dynamic programming equation

- $v_N(Z_N) = g(Z_N)$
- $v_n(Z_n) = L(v_{n+1}, g)(Z_n)$ for $n \leq N-1$

$$v_0(Z_0) = \sup_{\tau \in \mathcal{M}_N} \mathbb{E}_x[g(X_\tau)]$$

$$\begin{aligned} v_n(Z_n) &= L(v_{n+1}, g)(Z_n) \\ &= \sup_{u \leq t^*(Z_n)} \left\{ \mathbb{E}\left[v_{n+1}(Z_{n+1}) \mathbf{1}_{\{S_{n+1} < u\}} + g\big(\phi(Z_n, u)\big) \mathbf{1}_{\{S_{n+1} \geq u\}} \;\middle|\; Z_n \right] \right\} \\ &\quad \vee \mathbb{E}\left[v_{n+1}(Z_{n+1}) \;\middle|\; Z_n \right] \end{aligned}$$

Backward resolution

Backward dynamic programming equation

- $v_N(Z_N) = g(Z_N)$
- $v_n(Z_n) = L(v_{n+1}, g)(Z_n)$ for $n \leq N-1$

$$v_0(Z_0) = \sup_{\tau \in \mathcal{M}_N} \mathbb{E}_x[g(X_\tau)]$$

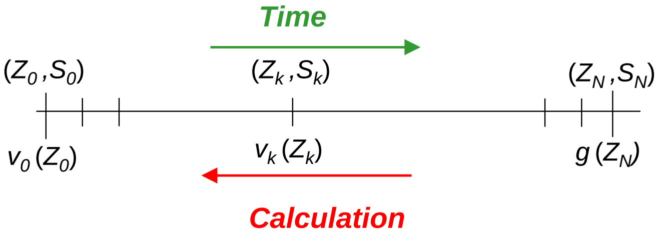

Discretization

Approximation of the value function

- $v_N(Z_N) = g(Z_N)$
- $v_n(Z_n) = L(v_{n+1}, g)(Z_n)$ for $n \leq N - 1$

$$L(v_{n+1}, g)(Z_n) = \sup_{u \leq t^*(Z_n)} \Big\{ \mathbb{E}\Big[v(Z_{n+1}) \mathbf{1}_{\{S_{n+1} < u\}} + g\big(\phi(Z_n, u)\big) \mathbf{1}_{\{S_{n+1} \geq u\}} \;\Big|\; Z_n \Big] \Big\} \vee \mathbb{E}\big[v(Z_{n+1}) \;\big|\; Z_n \big]$$

Discretization

Time discretization

- $v_N(Z_N) = g(Z_N)$
- $v_n(Z_n) = L(v_{n+1}, g)(Z_n)$ for $n \leq N-1$

$$
\begin{aligned}
& L_d(v_{n+1}, g)(Z_n) \\
&= \max_{u \in G(Z_n)} \left\{ \mathbb{E}\left[v(Z_{n+1}) \mathbf{1}_{\{S_{n+1} < u\}} + g\big(\phi(Z_n, u)\big) \mathbf{1}_{\{S_{n+1} \geq u\}} \;\middle|\; Z_n \right] \right\} \\
&\quad \vee \mathbb{E}\left[v(Z_{n+1}) \;\middle|\; Z_n \right]
\end{aligned}
$$

Discretization

Quantization

- $v_N(Z_N) = g(Z_N)$
- $v_n(Z_n) = L(v_{n+1}, g)(Z_n)$ for $n \leq N-1$

$$
\begin{aligned}
&\widehat{L}_d(v_{n+1}, g)(Z_n) \\
&= \max_{u \in G(Z_n)} \left\{ \mathbb{E}\left[v(\widehat{Z}_{n+1}) \mathbf{1}_{\{\widehat{S}_{n+1} < u\}} + g\big(\phi(Z_n, u)\big) \mathbf{1}_{\{\widehat{S}_{n+1} \geq u\}} \;\Big|\; \widehat{Z}_n \right] \right\} \\
&\quad \vee \mathbb{E}\left[v(Z_{n+1}) \;\Big|\; \widehat{Z}_n \right]
\end{aligned}
$$

Discretization

Approximation of the value function

- $\widehat{v}_N(\widehat{Z}_N) = g(\widehat{Z}_N)$
- $\widehat{v}_n(\widehat{Z}_n) = \widehat{L}_d(\widehat{v}_{n+1}, g)(\widehat{Z}_n)$ for $n \leq N-1$

$$\widehat{L}_d(v_{n+1}, g)(Z_n) = \max_{u \in G(Z_n)} \left\{ \mathbb{E}\left[v(\widehat{Z}_{n+1}) \mathbf{1}_{\{\widehat{S}_{n+1} < u\}} + g\big(\phi(Z_n, u)\big) \mathbf{1}_{\{\widehat{S}_{n+1} \geq u\}} \;\Big|\; \widehat{Z}_n \right] \right\} \vee \mathbb{E}\left[v(Z_{n+1}) \;\Big|\; \widehat{Z}_n \right]$$

Convergence rate

Theorem

Lipschitz assumptions on ϕ, λ, Q, t^* and g

$$|v_0(x) - \widehat{v}_0(x)| \leq C\sqrt{EQ}$$

C explicit constant,
EQ quantization error

$\sqrt{\cdot}$ due to the indicator functions

Indicator functions

Dealing with indicator functions

Set η s.t. $\forall s \in G(\widehat{Z}_0)$, $s + \eta < t^*(\widehat{Z}_0)$

$$\left\| \max_{u \in G(x)} \mathbf{E}_{Z_0}\left[\left| \mathbf{1}_{\{S_1 < u\}} - \mathbf{1}_{\{\widehat{S}_1 < u\}} \right| \right] \right\|_2 \leq \frac{1}{\eta} \| S_1 - \widehat{S}_1 \|_2 + C\eta$$

Easy cases:

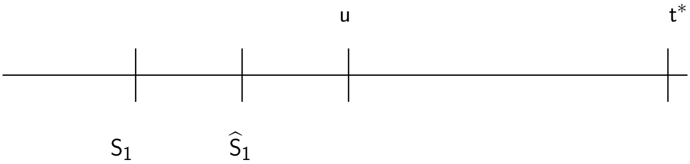

Indicator functions

Dealing with indicator functions

Set η s.t. $\forall s \in G(\widehat{Z}_0)$, $s + \eta < t^*(\widehat{Z}_0)$

$$\left\| \max_{u \in G(x)} \mathbf{E}_{Z_0}\left[\left|\mathbf{1}_{\{S_1 < u\}} - \mathbf{1}_{\{\widehat{S}_1 < u\}}\right|\right] \right\|_2 \leq \frac{1}{\eta} \|S_1 - \widehat{S}_1\|_2 + C\eta$$

Easy cases:

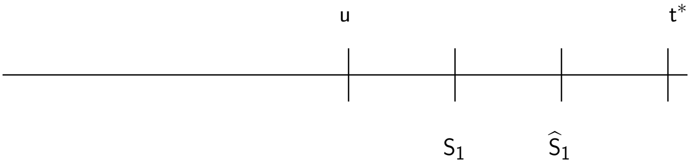

S_1 and $\widehat{S}_1$ are either side of u

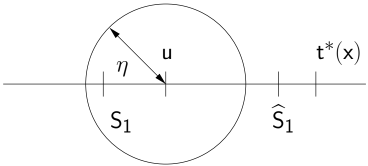

$$\left|\mathbf{1}_{\{S_1<u\}}-\mathbf{1}_{\{\widehat{S}_1<u\}}\right|\leq \mathbf{1}_{\{|S_1-u|\leq\eta\}}+\mathbf{1}_{\{|S_1-\widehat{S}_1|>\eta\}}$$

$$\begin{aligned}\mathbf{E}_{Z_0}\left[\mathbf{1}_{\{u-\eta\leq S_1\leq u+\eta\}}\right] &= \mathbf{E}_{Z_0}\left[\mathbf{1}_{\{u-\eta\leq S_1\leq u+\eta\}}\right]\\ &= \int_{u-\eta}^{u+\eta}\lambda(\phi(z_0,u))du \;\leq\; C\eta\end{aligned}$$

S_1 and $\widehat{S}_1$ are either side of u

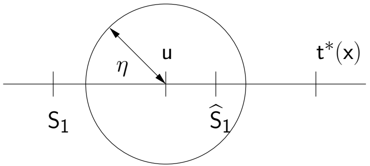

$$\left|\mathbf{1}_{\{S_1<u\}}-\mathbf{1}_{\{\widehat{S}_1<u\}}\right| \leq \mathbf{1}_{\{|S_1-u|\leq\eta\}}+\mathbf{1}_{\{|S_1-\widehat{S}_1|>\eta\}}$$

$$\begin{aligned}\mathbf{E}_{Z_0}\left[\mathbf{1}_{\{u-\eta\leq S_1\leq u+\eta\}}\right] &= \mathbf{E}_{Z_0}\left[\mathbf{1}_{\{u-\eta\leq S_1\leq u+\eta\}}\right] \\ &= \int_{u-\eta}^{u+\eta}\lambda(\phi(z_0,u))du \;\leq\; C\eta\end{aligned}$$

Optimal stopping time : τ^*

$$\mathbb{E}_x[g(X_{\tau^*})] = v_0(x) = \sup_{\tau \in \mathcal{M}_N} \mathbb{E}_x[g(X_\tau)]$$

Existence?

ϵ-optimal stopping time : $\hat{\tau}$

$$v_0(x) - \epsilon \leq \mathbb{E}_x[g(X_{\hat{\tau}})] \leq v_0(x)$$

Optimal stopping time : τ^*

$$\mathbb{E}_x[g(X_{\tau^*})] = v_0(x) = \sup_{\tau \in \mathcal{M}_N} \mathbb{E}_x[g(X_\tau)]$$

Existence?

ϵ-optimal stopping time : $\hat{\tau}$

$$v_0(x) - \epsilon \leq \mathbb{E}_x[g(X_{\hat{\tau}})] \leq v_0(x)$$

Proposition of a computable stopping rule $\hat{\tau}$

- explicit iterative construction
- no extra computation
- true stopping time in $\mathcal{M}_N$

Proposition of a computable stopping rule $\hat{\tau}$

- explicit iterative construction
- no extra computation
- true stopping time in $\mathcal{M}_N$

Optimality

Theorem

Same assumptions

$$|v_0(x) - \mathbb{E}_x[g(X_{\hat{\tau}})]| \leq C_1 EV + C_2 \sqrt{EQ}$$

C_1, C_2 explicit constants
EV value function error
EQ quantization error

Provides another approximation of the value function via Monte Carlo simulations

Material subject to corrosion

Model

- $m_t \in \{1, 2, 3\}$ describes the ambience
- d_t loss of thickness
- γ_t duration of the initial protection against corrosion
- ρ_t rate of corrosion

The process starts in ambience 1: $m_0 = 1$, $d_0 = 0mm$,

$$\gamma_0 \sim 1 - \exp(-[t/\alpha]^\gamma), \qquad \rho_0 \sim \mathcal{U}[\rho_1^1, \rho_1^2]$$

Dynamics

- Deterministic succession of ambiences: $1 \to 2 \to 3 \to 1 \cdots$
- Time spent in ambience $i \sim Exp(\lambda_i)$
- ρ_t is constant on $[T_n, T_{n+1}[$ and in ambience i, $\rho \sim \mathcal{U}[\rho_i^1, \rho_i^2]$
- $\gamma_{T_{n+1}} = 0$ if $\gamma_{T_n} < S_{n+1}$, $\quad \gamma_{T_{n+1}} = S_{n+1} - \gamma_{T_n}$ otherwise
- On $[T_n, T_{n+1}[$, in ambience i

$$d_t = \begin{cases} 0, \text{ if } t \leq \gamma_{T_n}, \\ \rho_{T_n}\Big(t - (\gamma_{T_n} + c_i) + c_i exp(-[t - \gamma_{T_i}]/c_i)\Big), \text{ otherwise.} \end{cases}$$

The material is inefficient when the thickness loss is greater than $0.2mm$

Reward function

Reward function g depends only on the loss of thickness

- Early maintenances are penalized
- The material is inefficient when the loss is greater than $0.2mm$

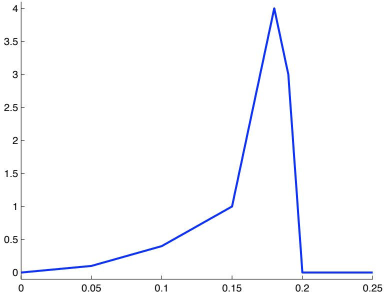

Iterative stopping rule

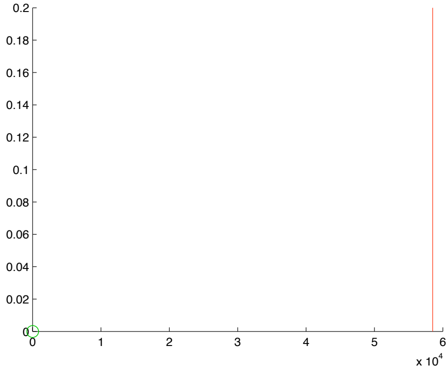

Iterative stopping rule

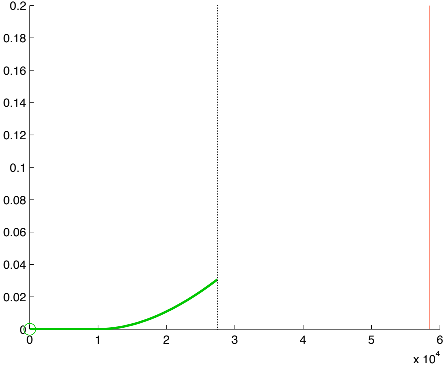

Iterative stopping rule

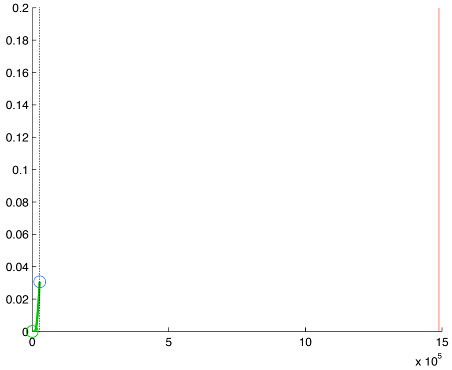

Iterative stopping rule

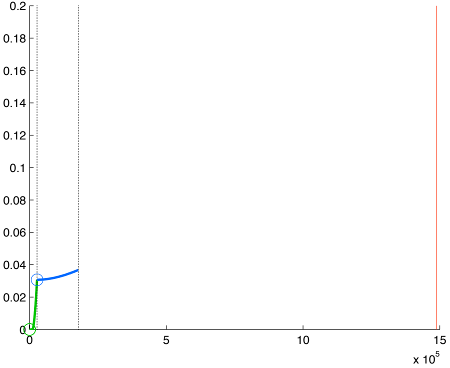

Iterative stopping rule

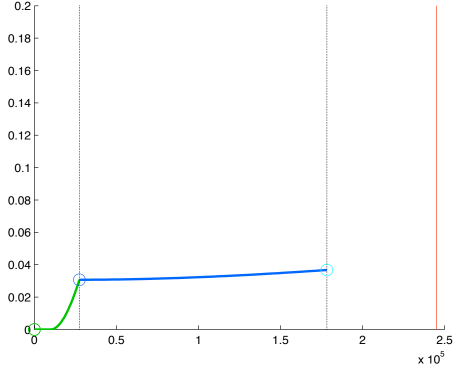

Iterative stopping rule

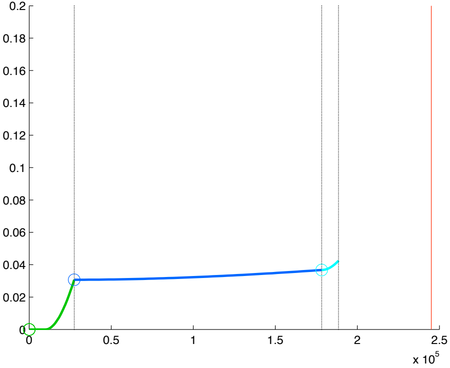

Iterative stopping rule

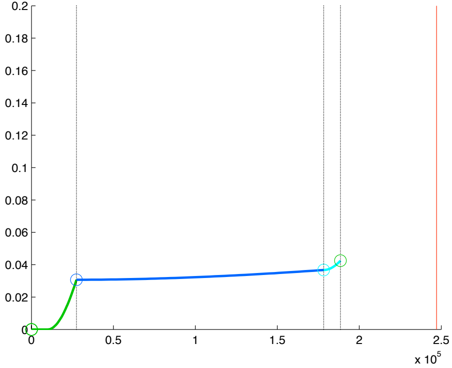

Iterative stopping rule

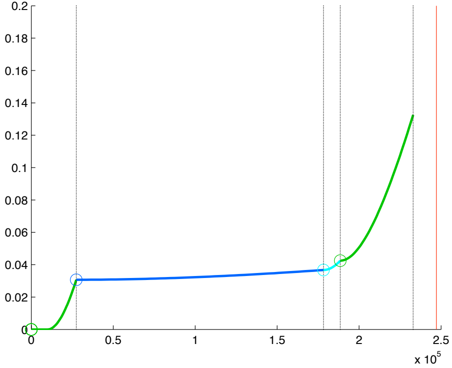

Iterative stopping rule

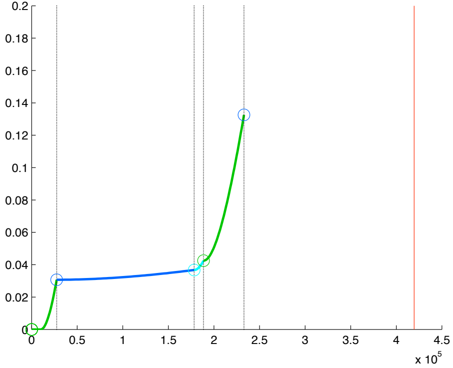

Iterative stopping rule

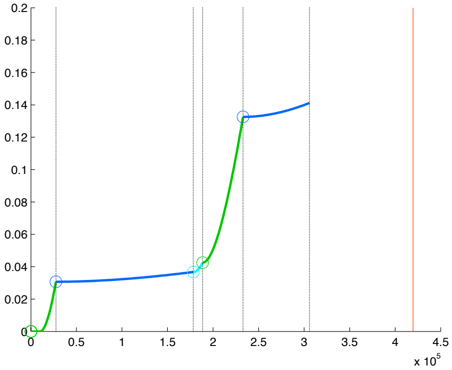

Iterative stopping rule

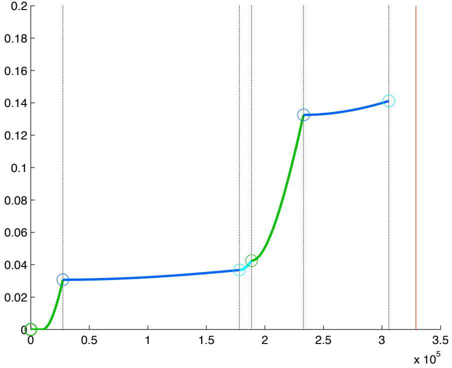

Iterative stopping rule

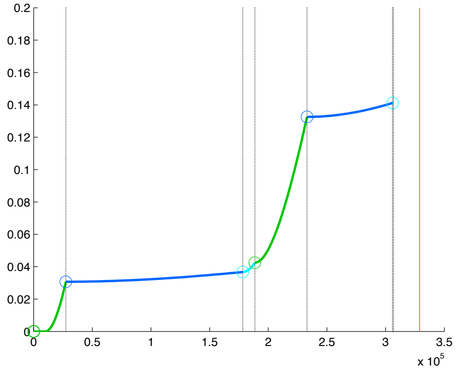

Iterative stopping rule

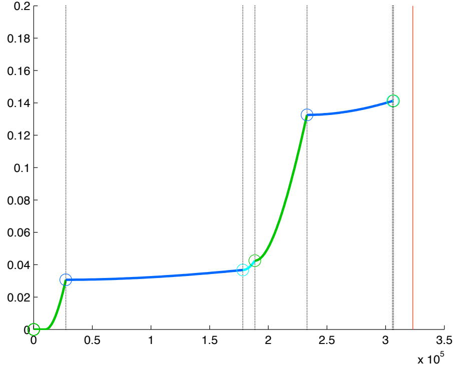

Iterative stopping rule

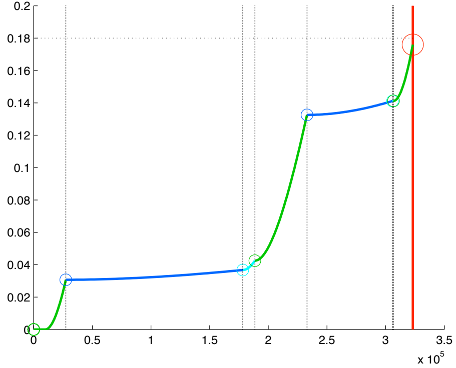

Results

Number of points in the quantization grids	Approximated value function	Monte Carlo approximated value function
10	2.48	0.94
50	2.70	1.84
100	2.94	2.10
200	3.09	2.63
500	3.39	3.15
1000	3.56	3.43
2000	3.70	3.60
5000	3.82	3.73