Regularity of the value function of optimal stopping problems

Adriana Ocejo Monge

Department of Statistics
University of Warwick

March 6, 2012

What is an optimal stopping problem? In general...

Given a stochastic process $X = (X_t)_{t\geq 0}$, an optimal stopping problem is to find two things: (1) The value of

$$v = \sup_{\tau} \mathbb{E}\left[\int_0^{\tau} f(X_t)dt + g(X_\tau)\right]. \tag{1}$$

where the supremum is taken over a specified class of stopping times, and

(2) A stopping rule τ^* such that

$$v = \mathbb{E}\left[\int_0^{\tau^*} f(X_t)dt + g(X_{\tau^*})\right]. \tag{2}$$

We will concentrate in the special case where $f \equiv 0$, i.e., in the problem

$$\sup_{\tau} \mathbb{E}\, g(X_\tau).$$

What is an optimal stopping problem? In general...

Given a stochastic process $X = (X_t)_{t \geq 0}$, an optimal stopping problem is to find two things: (1) The value of

$$v = \sup_{\tau} \mathbb{E}\left[\int_0^{\tau} f(X_t)dt + g(X_\tau)\right]. \tag{1}$$

where the supremum is taken over a specified class of stopping times, and

(2) A stopping rule τ^* such that

$$v = \mathbb{E}\left[\int_0^{\tau^*} f(X_t)dt + g(X_{\tau^*})\right]. \tag{2}$$

We will concentrate in the special case where $f \equiv 0$, i.e., in the problem

$$\sup_{\tau} \mathbb{E}\, g(X_\tau).$$

... in particular

If X is a Markov process, one can formulate the problem as

$$v(x) = \sup_{\tau} \mathbb{E}_x\, g(X_\tau). \tag{3}$$

where the expectation is with respect to the measure $P_x(X_0 = x) = 1$.

This approach is convenient:

$\surd$ When the initial state of the process is relevant.

$\surd$ At time t, the decision to stop or to continue *only depends* on the present state of X_t.

$\surd$ The problem becomes a problem of optimal stopping for a random path in the state space E (instead of the probability space Ω), which in general is $E = \mathbb{R}^n$.

Stochastic Analysis: Doob inequalities

It is well-known that if B is a standard Brownian motion and τ is any stopping time for B with $\mathbb{E}\tau < \infty$ then Doob's L^2- maximal inequality holds:

$$\mathbb{E}\left(\max_{0\le t\le\tau}|B_t|^2\right) \le 4\mathbb{E}|B_\tau|^2.$$

(Actually it holds for any càd martingale or positive submartingale X)

If the Brownian motion starts at any given point $x \ge 0$ under P_x, i.e., $P_x(B_0 = x) = 1$, one can show that

$$\mathbb{E}_x\left(\max_{0\le t\le\tau}|B_t|^2\right) \le 4\mathbb{E}_x|B_\tau|^2 - 2x^2,$$

by solving the optimal stopping problem

$$V(y,s) = \sup_\tau \mathbb{E}_{y,s}(S_\tau - c\tau)$$

with the Markov process (Y,S), where, for $0 \le y \le s$,

$$Y_t = |B_t|^2, \qquad S_t = \left(\max_{0\le r\le t}|B_r|^2\right) \vee s, \qquad Y_0 = y = x^2, \quad S_0 = s.$$

The idea is to realize that, setting $v(c) = \sup_\tau \mathbb{E}_{y,y}(S_\tau - c\tau)$ for $c > 0$,

$$\mathbb{E}_x\left(\max_{0\le t\le\tau}|B_t|^2\right) \le \inf_{c>0}\{v(c) + c\mathbb{E}_x\tau\}.$$

Stochastic Analysis: Doob inequalities

It is well-known that if B is a standard Brownian motion and τ is any stopping time for B with $\mathbb{E}\tau < \infty$ then Doob's L^2- maximal inequality holds:

$$\mathbb{E}\left(\max_{0\le t\le\tau}|B_t|^2\right) \le 4\mathbb{E}\,|B_\tau|^2.$$

(Actually it holds for any càd martingale or positive submartingale X)

If the Brownian motion starts at any given point $x \ge 0$ under P_x, i.e., $P_x(B_0 = x) = 1$, one can show that

$$\mathbb{E}_x\left(\max_{0\le t\le\tau}|B_t|^2\right) \le 4\mathbb{E}_x\,|B_\tau|^2 - 2x^2,$$

by solving the optimal stopping problem

$$V(y,s) = \sup_\tau \mathbb{E}_{y,s}(S_\tau - c\tau)$$

with the Markov process (Y,S), where, for $0 \le y \le s$,

$$Y_t = |B_t|^2, \qquad S_t = \left(\max_{0\le r\le t}|B_r|^2\right) \vee s, \qquad Y_0 = y = x^2, \quad S_0 = s.$$

The idea is to realize that, setting $v(c) = \sup_\tau \mathbb{E}_{y,y}(S_\tau - c\tau)$ for $c > 0$,

$$\mathbb{E}_x\left(\max_{0\le t\le\tau}|B_t|^2\right) \le \inf_{c>0}\{v(c) + c\mathbb{E}_x\tau\}.$$

Stochastic Analysis: Doob inequalities

It is well-known that if B is a standard Brownian motion and τ is any stopping time for B with $\mathbb{E}\tau < \infty$ then Doob's L^2- maximal inequality holds:

$$\mathbb{E}\left(\max_{0\leq t\leq\tau}|B_t|^2\right) \leq 4\mathbb{E}\,|B_\tau|^2.$$

(Actually it holds for any càd martingale or positive submartingale X)

If the Brownian motion starts at any given point $x \geq 0$ under P_x, i.e., $P_x(B_0 = x) = 1$, one can show that

$$\mathbb{E}_x\left(\max_{0\leq t\leq\tau}|B_t|^2\right) \leq 4\mathbb{E}_x\,|B_\tau|^2 - 2x^2,$$

by solving the optimal stopping problem

$$V(y,s) = \sup_\tau \mathbb{E}_{y,s}(S_\tau - c\tau)$$

with the Markov process (Y,S), where, for $0 \leq y \leq s$,

$$Y_t = |B_t|^2, \qquad S_t = \left(\max_{0\leq r\leq t}|B_r|^2\right)\vee s, \qquad Y_0 = y = x^2, \quad S_0 = s.$$

The idea is to realize that, setting $v(c) = \sup_\tau \mathbb{E}_{y,y}(S_\tau - c\tau)$ for $c > 0$,

$$\mathbb{E}_x\left(\max_{0\leq t\leq\tau}|B_t|^2\right) \leq \inf_{c>0}\{v(c) + c\mathbb{E}_x\tau\}.$$

Statistics: sequential testing

At time $t = 0$ we begin to observe a Poisson process $X = (X_t)_{t \geq 0}$ with intensity $\lambda > 0$, which is either λ_0 or λ_1. The true value of λ is not known to us.

Problem: to decide ASAP and with a minimal error probability the true value of λ. We now formalize this mathematically:

- On a probability space $(\Omega, \mathcal{F}; P_\pi, \pi \in [0,1])$ where

$$P_\pi = \pi P_1 + (1 - \pi) P_0$$

  we assume that the r.v. $\lambda = \lambda(\omega)$ takes two values λ_1 and λ_0 according to the a priori distribution P_π, i.e.

$$P_\pi(\lambda = \lambda_1) = \pi \qquad \text{and} \qquad P_\pi(\lambda = \lambda_0) = 1 - \pi.$$

- Concerning the observable process X,

$$P_\pi(X \in \cdot \mid \lambda = \lambda_i) = P_i(X \in \cdot),$$

  where $P_i(X \in \cdot)$ coincides with the distribution of a Poisson process with intensity λ_i.

Statistics: sequential testing

At time $t = 0$ we begin to observe a Poisson process $X = (X_t)_{t \geq 0}$ with intensity $\lambda > 0$, which is either λ_0 or λ_1. The true value of λ is not known to us.

Problem: to decide ASAP and with a minimal error probability the true value of λ. We now formalize this mathematically:

- On a probability space $(\Omega, \mathcal{F}; P_\pi, \pi \in [0,1])$ where

  $$P_\pi = \pi P_1 + (1-\pi) P_0$$

  we assume that the r.v. $\lambda = \lambda(\omega)$ takes two values λ_1 and λ_0 according to the a priori distribution P_π, i.e.

  $$P_\pi(\lambda = \lambda_1) = \pi \qquad \text{and} \qquad P_\pi(\lambda = \lambda_0) = 1 - \pi.$$

- Concerning the observable process X,

  $$P_\pi(X \in \cdot \mid \lambda = \lambda_i) = P_i(X \in \cdot),$$

  where $P_i(X \in \cdot)$ coincides with the distribution of a Poisson process with intensity λ_i.

To test sequentially the hypothesis $H_1 : \lambda = \lambda_1$ and $H_0 : \lambda = \lambda_0$, assume we can decide

- A stopping time $\tau = \tau(\omega)$ to stop the observation of X, and
- A terminal decision function $d = d(\omega)$ which indicates that either H_0 or H_1 should be accepted.

Each *decision rule* (τ, d) implies looses:

- $c\mathbb{E}_\pi\tau$, $c > 0$- due to a cost of observation.
- $aP_\pi(d = 0, \lambda = \lambda_1) + bP_\pi(d = 1, \lambda = \lambda_0)$, $a, b > 0$- due to a wrong terminal decision.

The *total average loss* of the decision rule (τ, d) is

$$L_\pi(\tau, d) = \mathbb{E}_\pi\left(c\,\tau + a\,I_{(d=0,\lambda=\lambda_1)} + b\,I_{(d=1,\lambda=\lambda_0)}\right)$$

and the problem is then to compute

$$V(\pi) = \inf_{(\tau,d)} L_\pi(\tau, d).$$

The optimal (τ^*, d^*) is called π*-Bayes decision rule*.

To test sequentially the hypothesis $H_1 : \lambda = \lambda_1$ and $H_0 : \lambda = \lambda_0$, assume we can decide

- A stopping time $\tau = \tau(\omega)$ to stop the observation of X, and
- A terminal decision function $d = d(\omega)$ which indicates that either H_0 or H_1 should be accepted.

Each *decision rule* (τ, d) implies looses:

- $c\mathbb{E}_\pi \tau$, $c > 0$- due to a cost of observation.
- $aP_\pi(d = 0, \lambda = \lambda_1) + bP_\pi(d = 1, \lambda = \lambda_0)$, $a, b > 0$- due to a wrong terminal decision.

The *total average loss* of the decision rule (τ, d) is

$$L_\pi(\tau, d) = \mathbb{E}_\pi \left(c\,\tau + a\, I_{(d=0,\lambda=\lambda_1)} + b\, I_{(d=1,\lambda=\lambda_0)} \right)$$

and the problem is then to compute

$$V(\pi) = \inf_{(\tau,d)} L_\pi(\tau, d).$$

The optimal (τ^*, d^*) is called π*-Bayes decision rule*.

To test sequentially the hypothesis $H_1 : \lambda = \lambda_1$ and $H_0 : \lambda = \lambda_0$, assume we can decide

- A stopping time $\tau = \tau(\omega)$ to stop the observation of X, and
- A terminal decision function $d = d(\omega)$ which indicates that either H_0 or H_1 should be accepted.

Each *decision rule* (τ, d) implies looses:

- $c\mathbb{E}_\pi \tau$, $c > 0$- due to a cost of observation.
- $aP_\pi(d = 0, \lambda = \lambda_1) + bP_\pi(d = 1, \lambda = \lambda_0)$, $a, b > 0$- due to a wrong terminal decision.

The *total average loss* of the decision rule (τ, d) is

$$L_\pi(\tau, d) = \mathbb{E}_\pi \left(c\,\tau + a\, I_{(d=0,\lambda=\lambda_1)} + b\, I_{(d=1,\lambda=\lambda_0)} \right)$$

and the problem is then to compute

$$V(\pi) = \inf_{(\tau,d)} L_\pi(\tau, d).$$

The optimal (τ^*, d^*) is called π*-Bayes decision rule*.

Mathematical Finance: American options

Let X be an asset price and assume that it follows the SDE

$$dX_t = \sigma(X_t)dB_t + b(X)dt, \qquad X_0 = x > 0.$$

An American-type option is a contract which gives the holder (or writer) the right to sell (or buy) the asset at any time before some fixed time T, and if she decides to exercise the option at time t then she receives a payoff $g(X_t)$.

If the holder chooses the stopping time $\tau \leq T$, the payoff has present value (after lots of considerations)

$$\mathbb{E}_x\, e^{-r\tau} g(X_\tau), \qquad r > 0 \text{ is the interest rate.}$$

Since the holder wants the highest reward but she doesn't know what the price will be in the future, she should take the stopping time τ^* such that

$$\mathbb{E}_x\, e^{-r\tau^*} g(X_\tau) = \sup_{\tau \leq T} \mathbb{E}_x\, e^{-r\tau} g(X_\tau),$$

and this again is an optimal stopping problem.

Mathematical Finance: American options

Let X be an asset price and assume that it follows the SDE

$$dX_t = \sigma(X_t)dB_t + b(X)dt, \qquad X_0 = x > 0.$$

An American-type option is a contract which gives the holder (or writer) the right to sell (or buy) the asset at any time before some fixed time T, and if she decides to exercise the option at time t then she receives a payoff $g(X_t)$.

If the holder chooses the stopping time $\tau \leq T$, the payoff has present value (after lots of considerations)

$$\mathbb{E}_x\, e^{-r\tau} g(X_\tau), \qquad r > 0 \text{ is the interest rate.}$$

Since the holder wants the highest reward but she doesn't know what the price will be in the future, she should take the stopping time τ^* such that

$$\mathbb{E}_x\, e^{-r\tau^*} g(X_\tau) = \sup_{\tau \leq T} \mathbb{E}_x\, e^{-r\tau} g(X_\tau),$$

and this again is an optimal stopping problem.

What do I study? The setting is:

Consider a two-dimensional strong Markov process $(X, Y) = (X_t, Y_t, t \geq 0)$ with state space $\mathbb{R} \times S$, $S \subseteq \mathbb{R}$, where X solves the SDE

$$dX = a(X)YdB \tag{4}$$

and Y is any of two classes:

1) *Regime-switching*: Y is an irreducible continuous-time MC independent of B.
2) *Diffusion*: Y solves an SDE of the type

$$dY = \eta(Y)dB^Y + \theta(Y)dt \tag{5}$$

where B and B^Y might be correlated and a, η, θ are measurable functions.

Note that Y does NOT depend on X.

The problem is...

...the regularity of the value function $v(x, y)$. Specifically, the monotonicity and continuity of $v(x, y)$ with respect to $y \in S$, where

$$v(x, y) = \sup_{0 \leq \tau \leq T} \mathbb{E}_{x,y}[e^{-q\tau} g(X_\tau)], \quad (x, y) \in \mathbb{R} \times S, \qquad (6)$$

with

- $q > 0$, $T \in [0, \infty]$,
- the gain function $g : \mathbb{R} \to \mathbb{R}$ is a measurable function,
- and the supremum is taken over all stopping times with respect to the filtration generated by (X, Y).

For ease of presentation we will assume $T = \infty$.

Regime-switching case

Assume that Y is a continuous-time MC with Q-matrix (q_{ij}) taking values on $S = \{y_i : i = 1, 2, \dots\} \subset (0, \infty)$ and,

$$X_t = x + \int_0^t a(X_s) Y_s \, dB_s,$$

The goal is to show that for fixed $x \in \mathbb{R}$ and $y, y' \in S$, it holds that

$$\textit{if} \quad y \leq y' \quad \textit{then} \quad v(x, y) \leq v(x, y'). \tag{7}$$

Recall that

$$v(x, y) = \sup_{\tau} \mathbb{E}_{x,y} [e^{-q\tau} g(X_\tau)]. \quad (x, y) \in \mathbb{R} \times S,$$

Immediate difficulties:

- $\mathbb{E}_{x,y}$ and $\mathbb{E}_{x,y'}$ are defined for different measures.
- X_t depends on the value of Y_t, so path comparison is not effective.

Regime-switching case

Assume that Y is a continuous-time MC with Q-matrix (q_{ij}) taking values on $S = \{y_i : i = 1, 2, \dots\} \subset (0, \infty)$ and,

$$X_t = x + \int_0^t a(X_s) Y_s \, dB_s,$$

The goal is to show that for fixed $x \in \mathbb{R}$ and $y, y' \in S$, it holds that

$$\textit{if} \quad y \leq y' \quad \textit{then} \quad v(x, y) \leq v(x, y'). \tag{7}$$

Recall that

$$v(x, y) = \sup_{\tau} \mathbb{E}_{x,y} [e^{-q\tau} g(X_\tau)]. \quad (x, y) \in \mathbb{R} \times S,$$

Immediate difficulties:

- $\mathbb{E}_{x,y}$ and $\mathbb{E}_{x,y'}$ are defined for different measures.
- X_t depends on the value of Y_t, so path comparison is not effective.

Working on only ONE probability space

One difficulty is that $\mathbb{E}_{x,y}$ *and* $\mathbb{E}_{x,y'}$ *are defined for different measures.*

To overcome this we choose an arbitrary $(\Omega, \mathcal{F}, P)$ BIG enough to carry three processes Y, Y', W such that

- Y has the same law as the original MC under $P_{x,y}$,
- Y' has the same law as the original MC under $P_{x,y'}$, and
- W is a Brownian motion independent of (Y, Y').

Technical consideration: $\mathcal{F}_t = \mathcal{F}_t^W \vee \mathcal{F}_\infty^Y \vee \mathcal{F}_\infty^{Y'}$.

We will construct two processes X and X' on $(\Omega, \mathcal{F}, P)$ such that they solve (weakly) the original SDE.

Then we can write

$$v(x, y) = \sup_\tau \mathbb{E}\,[e^{-q\tau} g(X_\tau)]$$

$$v(x, y') = \sup_{\tau'} \mathbb{E}\,[e^{-q\tau'} g(X'_{\tau'})]$$

where τ and τ' are stopping times with respect to the filtration generated by (X, Y) and (X', Y'), respectively.

Working on only ONE probability space

One difficulty is that $\mathbb{E}_{x,y}$ *and* $\mathbb{E}_{x,y'}$ *are defined for different measures.*

To overcome this we choose an arbitrary $(\Omega, \mathcal{F}, P)$ BIG enough to carry three processes Y, Y', W such that

- Y has the same law as the original MC under $P_{x,y}$,
- Y' has the same law as the original MC under $P_{x,y'}$, and
- W is a Brownian motion independent of (Y, Y').

Technical consideration: $\mathcal{F}_t = \mathcal{F}_t^W \vee \mathcal{F}_\infty^Y \vee \mathcal{F}_\infty^{Y'}$.

We will construct two processes X and X' on $(\Omega, \mathcal{F}, P)$ such that they solve (weakly) the original SDE.

Then we can write

$$v(x,y) = \sup_\tau \mathbb{E}\left[e^{-q\tau} g(X_\tau)\right]$$

$$v(x,y') = \sup_{\tau'} \mathbb{E}\left[e^{-q\tau'} g(X'_{\tau'})\right]$$

where τ and τ' are stopping times with respect to the filtration generated by (X, Y) and (X', Y'), respectively.

Time-change method

Another difficulty is that X_t *depends on the value of* Y_t.

To overcome this, we use a time-change method: given the MC Y, define

$$A_t = \int_0^t Y_u^2 du, \quad \textit{and} \quad \Gamma_t = \inf\{s \geq 0 : A_s > t\},$$

so that $A = \Gamma^{-1}$. Consider $G = (G_t)_{t \geq 0}$ as a unique strong solution to the SDE

$$G_t = G_0 + \int_0^t a(G_s)\, dW_s \qquad G_0 = x.$$

Note that G does not depend on Y.

Also define the local martingale $M_t = \int_0^t dW_u / Y_{\Gamma_u}$. The process $B = M \circ A = (M_{A_t})_{t \geq 0}$ is a Brownian motion.

For $X = G \circ A$ and $B = M \circ A$ we have

$$X_t = x + \int_0^t a(X_s) Y_s\, dB_s, \; t \geq 0, \; a.s.,$$

Similarly with Y' to get $X' = G \circ A'$ and $B' = M' \circ A'$.

Time-change method

Another difficulty is that X_t *depends on the value of* Y_t.

To overcome this, we use a time-change method: given the MC Y, define

$$A_t = \int_0^t Y_u^2 du, \quad \textit{and} \quad \Gamma_t = \inf\{s \geq 0 : A_s > t\},$$

so that $A = \Gamma^{-1}$. Consider $G = (G_t)_{t \geq 0}$ as a unique strong solution to the SDE

$$G_t = G_0 + \int_0^t a(G_s)\, dW_s \qquad G_0 = x.$$

Note that G does not depend on Y.

Also define the local martingale $M_t = \int_0^t dW_u / Y_{\Gamma_u}$. The process $B = M \circ A = (M_{A_t})_{t \geq 0}$ is a Brownian motion.

For $X = G \circ A$ and $B = M \circ A$ we have

$$X_t \;=\; x + \int_0^t a(X_s) Y_s \, dB_s,\ t \geq 0,\ a.s.,$$

Similarly with Y' to get $X' = G \circ A'$ and $B' = M' \circ A'$.

Putting all together

Observe that for $X = G \circ A$ and τ stopping time w.r.t. $(\mathcal{F}_{A_t})_{t\geq 0}$,

$$v(x,y) = \sup_\tau \mathbb{E} e^{-q\tau} g(X_\tau) = \sup_\tau \mathbb{E} e^{-q\Gamma_{A_\tau}} g(G_{A_\tau})$$
$$= \sup_\rho \mathbb{E} e^{-q\Gamma_\rho} g(G_\rho).$$

where we set $\rho = A_\tau$.

Similarly for $X' = G \circ A'$ and τ' stopping time w.r.t. $(\mathcal{F}_{A'_t})_{t\geq 0}$,

$$v(x,y') = \sup_{\tau'} \mathbb{E} e^{-q\tau} g(X'_{\tau'}) = \sup_\rho \mathbb{E} e^{-q\Gamma'_\rho} g(G_\rho)$$

where we set $\rho = A'_{\tau'}$.

Remark two important facts:

- In both cases, ρ is a stopping time w.r.t. SAME $(\mathcal{F}_t)_{t\geq 0}$
- The dependence on the Markov chain is ONLY on the discount factors Γ and Γ' which are the inverses of

$$A_t = \int_0^t Y_u^2 du, \qquad \text{and} \qquad A'_t = \int_0^t (Y'_u)^2 du$$

Therefore, we WANT to obtain $A_t \leq A'_t$ for all $t \geq 0$ a.s.

Putting all together

Observe that for $X = G \circ A$ and τ stopping time w.r.t. $(\mathcal{F}_{A_t})_{t\geq 0}$,

$$v(x,y) = \sup_\tau \mathbb{E} e^{-q\tau} g(X_\tau) = \sup_\tau \mathbb{E} e^{-q\Gamma_{A_\tau}} g(G_{A_\tau}) = \sup_\rho \mathbb{E} e^{-q\Gamma_\rho} g(G_\rho).$$

where we set $\rho = A_\tau$.

Similarly for $X' = G \circ A'$ and τ' stopping time w.r.t. $(\mathcal{F}_{A'_t})_{t\geq 0}$,

$$v(x,y') = \sup_{\tau'} \mathbb{E} e^{-q\tau} g(X'_{\tau'}) = \sup_\rho \mathbb{E} e^{-q\Gamma'_\rho} g(G_\rho)$$

where we set $\rho = A'_{\tau'}$.

Remark two important facts:

- In both cases, ρ is a stopping time w.r.t. SAME $(\mathcal{F}_t)_{t\geq 0}$
- The dependence on the Markov chain is ONLY on the discount factors Γ and Γ' which are the inverses of

$$A_t = \int_0^t Y_u^2 du, \qquad \text{and} \qquad A'_t = \int_0^t (Y'_u)^2 du$$

Therefore, we WANT to obtain $A_t \leq A'_t$ for all $t \geq 0$ a.s.

Coupling of Markov chains

Suppose that $\tilde{Y}$ and Y' are independent MC's with the same intensity matrix Q but $\tilde{Y}_0 = y$ and $Y'_0 = y'$, with $y \leq y'$. Assume that Q is tridiagonal.

Define the coupling time

$$C = \inf\{t \geq 0 : \tilde{Y}_t = Y'_t\}$$

and the Markov chain

$$Y_t = \begin{cases} \tilde{Y}_t & \text{if } t < C \\ Y'_t & \text{if } t \geq C. \end{cases}$$

We have that Y has the same law as $\tilde{Y}$ and $Y_0 = \tilde{Y}_0 = y$.

Moreover, the MC $\tilde{Y}$ does NOT overtake Y' before C, then

$$Y_t \leq Y'_t \qquad \forall t \geq 0.$$

Therefore, we can ALWAYS obtain

$$A_t \leq A'_t \qquad \forall t \geq 0 \quad \text{a.s.}$$

as we required for showing $v(x, y) \leq v(x, y')$.

Diffusion case

Now the two-dimensional strong Markov process (X, Y) is given by

$$dX = a(X)YdB$$
$$dY = \eta(Y)dB^Y + \theta(Y)dt.$$

Again, it holds true that for fixed $x, y, y' \in \mathbb{R}$,

$$\textit{if} \quad y \leq y' \quad \textit{then} \quad v(x, y) \leq v(x, y').$$

The technique is adapted from the MC case, i.e., it is based on time-change and coupling but with *more technicalities*.

From the monotonicity result we can also show continuity of $v(x, y)$ in the parameter y, at least in the finite-horizon ($T < \infty$) case.

Diffusion case

Now the two-dimensional strong Markov process (X, Y) is given by

$$dX = a(X)YdB$$
$$dY = \eta(Y)dB^Y + \theta(Y)dt.$$

Again, it holds true that for fixed $x, y, y' \in \mathbb{R}$,

$$\textit{if} \quad y \leq y' \quad \textit{then} \quad v(x, y) \leq v(x, y').$$

The technique is adapted from the MC case, i.e., it is based on time-change and coupling but with *more technicalities*.

From the monotonicity result we can also show continuity of $v(x, y)$ in the parameter y, at least in the finite-horizon ($T < \infty$) case.

Game against nature

The time-change/coupling method proved to be very powerful in comparison of processes.

Consider the following *game against nature*:

The payoff of the game is given by

$$f_{x,y}(\tau, Q) = \mathbb{E}_{x,y} e^{-q\tau} g(X^Q_\tau)$$

where X^Q stands for the solution to $dX = a(X)YdB$ if Y has intensity matrix Q.

Suppose that

- *nature* chooses Q so as to minimize f, while
- *the player* chooses the stopping rule τ so as to maximize f.

Then we could say that

- Nature plays the worst strategy against the player –> If Woody Allen is the player he would say "I am at two with nature."
- or, the player plays the best strategy assuming the worst situation –> pessimistic view.

Game against nature

The time-change/coupling method proved to be very powerful in comparison of processes.

Consider the following *game against nature*:

The payoff of the game is given by

$$f_{x,y}(\tau, Q) = \mathbb{E}_{x,y} e^{-q\tau} g(X_\tau^Q)$$

where X^Q stands for the solution to $dX = a(X)YdB$ if Y has intensity matrix Q.

Suppose that

- *nature* chooses Q so as to minimize f, while
- *the player* chooses the stopping rule τ so as to maximize f.

Then we could say that

- Nature plays the worst strategy against the player –> If Woody Allen is the player he would say "I am at two with nature."
- or, the player plays the best strategy assuming the worst situation –> pessimistic view.

The value functions of the player and nature are given by

$$\begin{aligned} v^P(x,y) &= \sup_\tau \inf_Q f_{x,y}(\tau, Q), \\ v^N(x,y) &= \inf_Q \sup_\tau f_{x,y}(\tau, Q). \end{aligned} \tag{8}$$

One can check that $v^P(x,y) \leq v^N(x,y)$ and so the player wants to achieve the equality so as to maximize its value function.

Suppose there exist strategies τ^* and Q^* such that

$$\begin{aligned} \tau^* &= \arg\max f_{x,y}(\cdot, Q^*), \\ Q^* &= \arg\min f_{x,y}(\tau^*, \cdot). \end{aligned} \tag{9}$$

It follows that

$$f_{x,y}(\tau, Q^*) \leq f_{x,y}(\tau^*, Q^*) \leq f_{x,y}(\tau^*, Q) \tag{10}$$

for all strategies τ and Q. If the above inequalities hold then $v^P(x,y) = v^N(x,y)$.

Hence we look for the *saddle point* (τ^*, Q^*)!

The value functions of the player and nature are given by

$$\begin{aligned} v^P(x,y) &= \sup_{\tau} \inf_{Q} f_{x,y}(\tau, Q), \\ v^N(x,y) &= \inf_{Q} \sup_{\tau} f_{x,y}(\tau, Q). \end{aligned} \tag{8}$$

One can check that $v^P(x,y) \leq v^N(x,y)$ and so the player wants to achieve the equality so as to maximize its value function.

Suppose there exist strategies τ^* and Q^* such that

$$\begin{aligned} \tau^* &= \arg\max f_{x,y}(\cdot, Q^*), \\ Q^* &= \arg\min f_{x,y}(\tau^*, \cdot). \end{aligned} \tag{9}$$

It follows that

$$f_{x,y}(\tau, Q^*) \leq f_{x,y}(\tau^*, Q^*) \leq f_{x,y}(\tau^*, Q) \tag{10}$$

for all strategies τ and Q. If the above inequalities hold then $v^P(x,y) = v^N(x,y)$.

Hence we look for the *saddle point* (τ^*, Q^*)!

The value functions of the player and nature are given by

$$\begin{aligned} v^P(x,y) &= \sup_{\tau} \inf_{Q} \; f_{x,y}(\tau, Q), \\ v^N(x,y) &= \inf_{Q} \sup_{\tau} \; f_{x,y}(\tau, Q). \end{aligned} \tag{8}$$

One can check that $v^P(x,y) \le v^N(x,y)$ and so the player wants to achieve the equality so as to maximize its value function.

Suppose there exist strategies τ^* and Q^* such that

$$\begin{aligned} \tau^* &= \arg\max \; f_{x,y}(\cdot, Q^*), \\ Q^* &= \arg\min \; f_{x,y}(\tau^*, \cdot). \end{aligned} \tag{9}$$

It follows that

$$f_{x,y}(\tau, Q^*) \le f_{x,y}(\tau^*, Q^*) \le f_{x,y}(\tau^*, Q) \tag{10}$$

for all strategies τ and Q. If the above inequalities hold then $v^P(x,y) = v^N(x,y)$.

Hence we look for the *saddle point* (τ^*, Q^*)!

Solution

One can find a saddle point (τ^*, Q^*), i.e,

$$f_{x,y}(\tau, Q^*) \leq f_{x,y}(\tau^*, Q^*) \leq f_{x,y}(\tau^*, Q).$$

The left-hand side inequality is simply the solution to the optimal stopping problem

$$f(\tau^*, Q^*) = \sup_{\tau} f_{x,y}(\cdot, Q^*),$$

which is known to exist.

The right-hand side inequality can be solved using time-change and coupling!

Work in progress is to study the problem in the case where Y is a controlled diffusion process.

Solution

One can find a saddle point (τ^*, Q^*), i.e,

$$f_{x,y}(\tau, Q^*) \leq f_{x,y}(\tau^*, Q^*) \leq f_{x,y}(\tau^*, Q).$$

The left-hand side inequality is simply the solution to the optimal stopping problem

$$f(\tau^*, Q^*) = \sup_{\tau} f_{x,y}(\cdot, Q^*),$$

which is known to exist.

The right-hand side inequality can be solved using time-change and coupling!

Work in progress is to study the problem in the case where Y is a controlled diffusion process.

Solution

One can find a saddle point (τ^*, Q^*), i.e,

$$f_{x,y}(\tau, Q^*) \leq f_{x,y}(\tau^*, Q^*) \leq f_{x,y}(\tau^*, Q).$$

The left-hand side inequality is simply the solution to the optimal stopping problem

$$f(\tau^*, Q^*) = \sup_{\tau} f_{x,y}(\cdot, Q^*),$$

which is known to exist.

The right-hand side inequality can be solved using time-change and coupling!

Work in progress is to study the problem in the case where Y is a controlled diffusion process.

Thank you!